

ZBasic Language Reference Manual

Including Information on the
ZX Series Microcontrollers

Version 4.3.2

Publication History

November 2005	First publication
May 2006	Added information on structures and other updates
August 2006	Updated for new ZX models
October 2006	Added new information on tasks and memory allocation
February 2007	Added information on new ZX models
August 2007	Added information on a new ZX model
March 2008	Added information on new ZX models
July 2008	Updated native mode device information
January 2009	Added information on a new ZX model
January 2010	Added information on object-oriented extensions and new devices
June 2010	Added information on new ZX models
October 2010	Added information on new ZX models
March 2011	Updated for new compiler features
September 2011	Updated for generic target devices
March 2012	Updated for new compiler features
January 2013	Updated for new compiler features
January 2014	Updated for new devices
February 2015	Updated for new compiler features and devices
September 2015	Updated for new compiler features and devices

Disclaimer

Elba Corp. makes no warranty regarding the accuracy of or the fitness for any particular purpose of the information in this document or the techniques described herein. The reader assumes the entire responsibility for the evaluation of and use of the information presented. The Company reserves the right to change the information described herein at any time without notice and does not make any commitment to update the information contained herein. No license to use proprietary information belonging to the Company or other parties is expressed or implied.

Critical Applications Disclaimer

ELBA CORP. PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE OR TO BE USED IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS IN LIFE-SUPPORT OR SAFETY DEVICES OR SYSTEMS, CLASS III MEDICAL DEVICES, NUCLEAR FACILITIES, APPLICATIONS RELATED TO THE DEPLOYMENT OF AIRBAGS, OR ANY OTHER APPLICATIONS WHERE DEFECT OR FAILURE COULD LEAD TO DEATH, PERSONAL INJURY OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE (INDIVIDUALLY AND COLLECTIVELY, "CRITICAL APPLICATIONS"). FURTHERMORE, ELBA CORP. PRODUCTS ARE NOT DESIGNED OR INTENDED FOR USE IN ANY APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE OR AIRCRAFT. CUSTOMER AGREES, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE ELBA CORP. PRODUCTS, TO THOROUGHLY TEST THE SAME FOR SAFETY PURPOSES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF ELBA CORP. PRODUCTS IN CRITICAL APPLICATIONS.

The following notice applies solely to the Scintilla and SciTE software on which the ZBasic IDE is based:

License for Scintilla and SciTE, Copyright © 1998-2004 by Neil Hodgson

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

Trademarks

ZBasic, ZX-24, ZX-24a, ZX-24n, ZX-24p, ZX-24r, ZX-24s, ZX-24t, ZX-24x, ZX-24u, ZX-40, ZX-40a, ZX-40n, ZX-40p, ZX-40r, ZX-40s, ZX-40t, ZX-44, ZX-44a, ZX-44n, ZX-44p, ZX-44r, ZX-44s, ZX-44t, ZX-328n, ZX-328l, ZX-32n, ZX-32l, ZX-1280, ZX-1280n, ZX-1281, ZX-1281n, ZX-32a4, ZX-128a4u and ZX-128a1 are trademarks of Elba Corp.

ZX-24e, ZX-24ae, ZX-24ne, ZX-24pe, ZX-24nu, ZX-24pu, ZX-24ru, ZX-24su, ZX-24xu, ZX-328nu, ZX-128e, ZX-128ne, ZX-1281e and ZX-1281ne are trademarks of Oak Micros used under license from Elba Corp.

AVR is a registered trademark of Atmel Corp.

BasicX, BX-24 and BX-35 are trademarks of NetMedia, Inc.

PBasic is a trademark and Basic Stamp is a registered trademark of Parallax, Inc.

Visual Basic is a registered trademark of Microsoft Corp.

Arduino is a trademark of the Arduino Team.

Other brand and product names are trademarks or registered trademarks of their respective owners.

Table of Contents

Chapter 1 - Introduction.....	1
1.1 The ZBasic System Library	3
1.2 The ZX Microcontroller Family.....	3
1.3 Support for Generic Microcontrollers	4
1.4 Conventions	5
Chapter 2 - ZBasic Language Elements.....	7
2.1 Identifiers	7
2.2 Data Types	7
2.3 Modules	8
2.3.1 The Options Section.....	8
Option <pin>	9
Option <port>.....	9
Option Base	10
Option Explicit.....	10
Option Language	11
Option AllocStr	11
Option StringSize	11
Option Strict	12
Option TargetCPU	12
Option TargetDevice	12
Option DeviceParameter	12
Option PortPinEncoding	13
Option CodeLimit.....	13
Option SignOn.....	13
Option DefaultISR.....	14
Option CallFunctions	14
Option TaskStackMargin	14
Option ExtRamConfig	14
Option RamSize	15
Option HeapSize	15
Option HeapReserve.....	15
Option HeapLimit	16
Option MainTaskStackSize	16
Option TxQueueSize	16
Option RxQueueSize	16
Option ConsoleSpeed	16
Option Com1Speed.....	17
Option Console.....	17
Option ComChannels	17
Option RTC	18

Option AtnChar.....	18
Option Include	18
Option Objects.....	19
Option Overload	19
Option Namespaces.....	19
Option NameStyle.....	19
Option CodeType.....	20
Option Arduino	20
Option Notice.....	20
Option SerialReadStrobe.....	20
Option X10Interrupt	21
Option OCModulateEnable	21
Option Library.....	21
2.3.2 The Definitions Section.....	21
Defining Constants	22
Defining Variables.....	23
Defining Arrays of Variables.....	23
Defining Subroutines	24
Defining Functions	26
2.4 Expressions	28
2.4.1 Operator Precedence	28
2.4.2 Operator Associativity	29
2.4.3 Arithmetic Operators	30
2.4.4 Logical Operators.....	30
2.4.5 Comparison Operators.....	31
2.4.6 Miscellaneous Operators	31
2.4.7 No "Short Circuit" Evaluation	31
2.5 Statements	31
2.5.1 Assignment Statement.....	32
2.5.2 Call Statement.....	32
2.5.3 CallTask Statement	33
2.5.4 Console.WriteLine and Console.WriteLine Statements	33
2.5.5 Debug.Print Statement.....	34
2.5.6 Do-Loop Statement and Variants	35
2.5.7 Exit Statement	36
2.5.8 For-Next Statement	36
2.5.9 Goto Statement	38
2.5.10 If-Then-Else Statement.....	38
2.5.11 Single-line If-Then Statement	40
2.5.12 Select-Case Statement.....	40
2.5.13 Set Statement.....	42
2.5.14 While-Wend Statement.....	42

2.5.15 With Statement.....	42
2.6 Literals	43
2.6.1 Boolean Literals	43
2.6.2 Numeric Literals - Integral Values	43
2.6.3 Numeric Literals - Real Values	44
2.6.4 String Literals.....	44
2.6.5 Built-in Binary Constants	44
2.7 Comments.....	45
2.8 Line Continuation and Multiple Statements Per Line	45
2.9 Persistent Variables.....	46
2.10 Program Memory Data Items.....	47
2.11 String Types.....	51
2.11.1 Bounded Strings	51
2.11.2 Fixed-Length Strings	51
2.12 Variable Initialization	52
2.13 Type Conversions.....	52
2.14 Parameter Passing Conventions.....	53
2.15 Program and Data Item Properties	54
2.16 Default Visibility.....	55
Chapter 3 - Advanced Topics.....	57
3.1 Scope and Lifetime	57
3.2 Enumerations.....	60
3.3 Serial Channels	62
3.4 Queues	62
3.4.1 System Queues.....	64
3.5 Multitasking	64
3.5.1 Advanced Multi-tasking Options	66
3.6 Semaphores.....	67
3.7 Built-in Variables	67
3.7.1 Special Function Registers.....	68
3.7.2 System Variables and Constants	68
3.8 Built-in Constants	73
3.9 Conditional Compilation Directives	74
3.10 Error Directive	77
3.11 Notice Directive	77
3.12 Include Directive	78
3.13 Include_path Directive	78
3.14 Using Conditional Directives in Project and Argument Files	78
3.15 Preprocessor Symbols.....	79
3.16 Array Data Order.....	80
3.17 Recursion in Subroutines and Functions	80
3.18 Using Default Parameter Values in Subroutines and Functions	81

3.19 Subroutine and Function Overloads	81
3.20 Aliases	82
3.21 Based Variables.....	83
3.22 Based Procedures	85
3.23 Reference Variables.....	86
3.24 Sub-byte Types	87
3.24.1 Forcing Byte Alignment.....	88
3.25 Structures.....	88
3.25.1 Structures in Persistent Memory and Program Memory	91
3.26 Unions.....	91
3.27 Using Namespaces	92
3.28 Data Type Implementation Information.....	93
3.28.1 User-defined Type Details.....	94
3.28.2 String Data Type Details.....	94
3.28.3 String Address and String Type	95
3.29 Controlling the Heap Size and Main() Task Stack Size.....	96
3.30 Task Management.....	97
3.30.1 Task Control Block	97
3.31 Dynamic Memory Allocation	98
3.32 Exception Handling	99
3.33 Run Time Stack Checking.....	100
Chapter 4 - ZBasic Object-Oriented Extensions	103
4.1 Enabling Object-Oriented Extensions.....	103
4.2 Defining a Class	103
4.3 Defining Class Methods	105
4.4 Object Creation Issues.....	106
4.5 Object Destruction Issues.....	108
4.6 Object Assignment Issues	108
4.7 Object Self-reference and Parent Reference	109
4.8 Explicit Class References and Default Namespace References	109
4.9 Class Sections.....	109
4.10 Static Class Members	110
4.11 Inheritance	111
4.12 Abstract Classes, Abstract Methods.....	114
4.13 Final Classes	115
4.14 Using Mixins	115
4.15 Using the Const Attribute for Methods	117
4.16 Based Objects, Reference Objects.....	117
4.17 Miscellaneous Class Elements.....	118
Chapter 5 - Compiling for Generic Target Devices.....	120
5.1 Overview.....	120
5.2 Loading Application Code onto the Target Device	120

5.3 Target Device Parameters	120
ClockFrequency	121
ClockPrescaler	121
Package.....	122
RTCFrequency.....	122
RTCScale	122
RTCError	123
TimerSpeed1Divisor, TimerSpeed2Divisor	123
TimerOCPin	123
HWUartSpeed	124
SWUartDivisor, SWUartMinSpeed, SWUartMaxSpeed, SWUartBaseSpeed	124
ZBasicBootloader.....	125
BootloaderAddress	125
BootloaderSize	126
WriteWordAddress, WritePageAddress.....	126
clkCtrl, psCtrl, oscCtrl, xoscCtrl, pllCtrl.....	126
5.4 ZBasic Bootloader.....	127
Chapter 6 - Special Considerations for Native Mode Devices.....	129
6.1 Using Inline C and Assembly Code	129
6.2 Defining and Using External Subroutines, Functions and Variables	130
6.3 Defining Interrupt Service Routines	132
6.4 Executing Blocks of Code Atomically	133
6.5 Attributes for Procedures and Variables	134
6.6 Considerations for Task Stack Size	135
6.7 Creating and Using Object Libraries.....	136
6.8 Importing Identifiers from External Modules	136
Chapter 7 - Compatibility Issues	139
7.1 Known Differences and Compatibility Between ZBasic and BasicX	139
7.2 Known Differences and Compatibility Between ZBasic and Visual Basic	140
Chapter 8 - The ZBasic IDE	141
8.1 Using the Editor.....	141
8.1.1 Basic Editing	141
8.1.2 Special Code Editing Features	141
8.1.3 Expand/Collapse	142
8.1.4 Auto-Completion	142
8.1.5 Call Tips.....	142
8.2 Project Configuration, Compiling and Downloading.....	143
8.3 Compiling and Downloading Individual Files	144
8.4 Setting Serial Port Options	144
8.5 Setting Device Options	145
8.6 Setting Target Options	145
8.7 Setting the Downloader Command	145

8.8 Downloading Without Using DTR Signaling	146
8.9 Updating Device Firmware	147
8.10 Setting the Debug Output Limit	147
8.11 Other Configurable Items	148
Chapter 9 - Compiler Guide.....	149
9.1 Compiler Invocation.....	149
9.2 Compiler Options in Detail.....	151
9.3 Error and Warning Messages	160
9.3.1 Controlling Warnings	160
9.3.2 Internal Errors	161
Chapter 10 - Downloader Utility	163
10.1 Firmware Updates	164
10.1.1 Emergency Update Procedure	165
10.2 Device Configuration.....	165
10.3 Downloader API.....	167
Appendix A - Reserved Words	169
Appendix B - Supported Target Devices	170
Appendix C - Pre-Defined Structures.....	171
Appendix D - ZX-24 Series Hardware Reference	175
D.1 External Connections	175
D.2 Pin Configuration.....	177
D.2.1 Standard Pins	177
Detailed Pin Descriptions	178
D.2.2 Expansion Pins	180
Detailed Pin Descriptions	181
Appendix E - ZX-40 Series Hardware Reference	183
E.1 ZX-40 Series Specifications.....	183
E.2 ZX-40 Series Required External Components	183
Appendix F - ZX-44 Series Hardware Reference	189
F.1 ZX-44 Series Specifications.....	189
F.2 ZX-44 Series Required External Components.....	189
Appendix G - ZX-1281 Series Hardware Reference	195
G.1 ZX-1281 Series Specifications.....	195
G.2 ZX-1281 Series Required External Components	195
Appendix H - ZX-1280 Series Hardware Reference	199
H.1 ZX-1280 Series Specifications	199
H.2 ZX-1280 Series Required External Components	199
Appendix I - ZX-328n and ZX-328I Hardware Reference	203
I.1 ZX-328 Series Specifications	203
I.2 ZX-328 Series Required External Components	203
Appendix J - ZX-32n and ZX-32I Hardware Reference	207
J.1 ZX-32 Series Specifications	207

J.2 ZX-32 Series Required External Components	207
Appendix K - ZX-32a4 and ZX-128a4u Hardware Reference.....	211
K.1 ZX-32a4 and ZX-128a4u Specifications	211
K.2 ZX-32a4 and ZX-128a4u Required External Components	211
Appendix L - ZX-128a1 Hardware Reference	215
L.1 ZX-128a1 Specifications	215
L.2 ZX-128a1 Required External Components.....	215
Appendix M - External Components for Generic Target Devices.....	219
Appendix N - Processor Resource Description Files	224

The ZBasic Language and ZX Series Microcontrollers

Chapter 1 - Introduction

The purpose of this manual is to describe the elements of the ZBasic language and the ZX series microcontrollers for those who have, at a minimum, a rudimentary knowledge of programming concepts and an understanding of basic electronics. It is not intended to teach programming skills nor electronics - there are many texts available for that purpose.

The ZBasic programming language will be familiar to anyone who has used Microsoft's Visual Basic language or NetMedia's BasicX language. ZBasic was designed to achieve a high level of compatibility with these languages but it also incorporates some additional capabilities to improve its utility in the microcontroller environment. Even if you aren't familiar with either of these two particular dialects of the Basic language, if you have used another fairly modern dialect of you will find many familiar elements. In contrast, ZBasic is quite different from archaic dialects of Basic like the PBasic language developed by Parallax for use on the Basic Stamp series of microcontrollers. Even so, moving from PBasic to ZBasic should be a relatively easy and rewarding experience.

ZBasic incorporates the modern block structured programming elements that help make programs easier to write, easier to maintain and more reliable. Moreover, it provides a good range of data types allowing efficient representation of the data items you need in your application. Beyond that, it has facilities to implement multi-tasking systems, allowing complex problems to be solved more easily.

All ZBasic programs comprise, at a minimum, one subroutine. That one required subroutine must be named `Main()` and it must be defined as taking no parameters. Additionally, the `Main()` subroutine must have "public" visibility as opposed to being private to a particular module. The minimal ZBasic program is shown below, in its entirety.

```
' the minimal program
Sub Main()
End Sub
```

This do-nothing program illustrates several points. First is the structure of comments as illustrated on the first line of the program. A comment consists of an apostrophe (sometimes called a single quote mark) and all characters following it all the way to the end of the line. It is not necessary for a comment to be on a line by itself. Characters occurring before the apostrophe are processed as if the comment portion of the line weren't present.

The second point illustrated by the program above is the manner in which subroutines are defined. A subroutine definition begins with the keyword `Sub`. That is followed by at least one space or tab character and then the name of the subroutine. Following the subroutine name is the parameter list, details of which are discussed later in this document. The parameter list is always enclosed in a pair of parentheses which must be present even if there are no parameters as is the case above. All of the statements, if any, between the `Sub` line and the `End Sub` line constitute the body of the subroutine.

Here is another program, intended for a 24-pin ZX device, that is a bit more complicated.

```
Const redLED as Byte = 25      ' define the pin for the red LED
Const grnLED as Byte = 26      ' define the pin for the green LED

Sub Main()
    ' configure the pins to be outputs, LEDs off
    Call PutPin(redLed, zxOutputHigh)
    Call PutPin(grnLed, zxOutputHigh)
```

```

' alternately blink the LEDs
Do
    ' turn on the red LED for one half second
    Call PutPin(redLed, zxOutputLow)
    Call Delay(0.5)
    Call PutPin(redLed, zxOutputHigh)

    ' turn on the green LED for one half second
    Call PutPin(grnLed, zxOutputLow)
    Call Delay(0.5)
    Call PutPin(grnLed, zxOutputHigh)
Loop
End Sub

```

The first two lines above define some constant values. It is a good idea to use constant values like this instead of using the value explicitly in many different places. One benefit of doing so is that it makes the program somewhat easier to understand assuming, of course, that the name chosen for the constant is suggestive of its intended use or meaning. A second benefit is that if you later wish to change the value, you can do so easily by modifying it in only one place. This way you'll avoid the common error of changing a constant value in some places but not in others.

It is important to note that alphabetic case is not significant in ZBasic. You may type keywords, variable names, etc. in upper case, lower case or mixed case and it makes no difference. The subroutine `Main()` is the same as the subroutine `main()`. Many programmers use alphabetic case to improve readability. Some even use it to remind them of the visibility attribute of the subroutine, constant, etc. One convention for this is to begin all public names with an upper case letter and begin all private names with a lower case letter. You're free to adopt these conventions or not as you see fit.

The public `Main()` routine is where execution begins when you run a program. That's why every program must have a public subroutine `Main()`. In this program, the first two lines of `Main()` utilize a call to a System Library subroutine that sets the state of an I/O pin on the processor. The first parameter that is provided to the `PutPin()` subroutine is the pin number that we wish to configure and we've used the previously defined constants to do so. The second parameter to the `PutPin()` subroutine is a value that indicates whether the pin should be an output or an input and, additionally, its characteristics. In this case, we've used a built-in constant `zxOutputHigh` which both configures the pin to be an output and sets it to be at a logic high level. The LEDs that are present on the processor are illuminated when the pin to which they are attached is at logic low level. The effect of these first two lines, then, is to make the pins associated with both LEDs outputs and to make sure the LEDs are off.

The next part of the code is the block beginning with `Do` and ending with `Loop`. This construct is useful for continually repeating a sequence of statements. The logic flow is that each of the statements inside the `Do . . . Loop` construct is executed in turn. When the `Loop` statement is reached, control transfers back to the top of the Do loop and the process repeats. ZBasic has several different kinds of control structures for looping, each with different useful characteristics. There are also control structures to allow you to write statements that are executed only if certain conditions prevail. The set of allowable statements is described in detail later in this manual.

Inside of the Do loop, there are two sequences of statements that perform similar functions. The first three lines of code within the Do loop turn on the red LED for a brief period of time. This is done by first setting the corresponding pin to a logic low, then calling the System Library subroutine to delay for a half second and finally setting the LED pin to a logic high. The second sequence does the same for the green LED.

A subroutine is nothing more than a collection of statements that can be executed by invoking the subroutine name. Although it is possible to implement all of the logic of your application inside the subroutine `Main()`, this is usually ill advised except for fairly simple applications like the example above. It is usually better to decompose your application into logical elements and to implement the functionality of each element using a subroutine or a function. The difference between a subroutine and a function is that a function returns a value. Because of this, a function can be used wherever a value may be used.

A subroutine is invoked by using the `Call` statement that specifies the name of the subroutine to invoke and, optionally, the parameters to pass to the subroutine. A function is invoked by using its name in place of a value (directly or in an expression) and specifying the parameters, if any, that are to be passed to the function.

This concludes a rather brief introduction to the ZBasic language. The remaining sections of this manual describe in more detail the various elements of ZBasic including data types, variables, constants, statements, subroutines and functions. Also, more advanced topics like multi-tasking, queues, serial communications and others are presented.

1.1 The ZBasic System Library

The ZBasic System Library provides a rich collection of several hundred subroutines and functions that you can use to quickly add functionality to your application. The routines fall into several fundamental categories including mathematical functions, string-oriented routines, I/O-related routines, type conversion functions, etc. The routines are fully documented in the ZBasic System Library Reference Manual.

1.2 The ZX Microcontroller Family

The ZBasic language was designed to be well suited for programming microcontrollers. It is based on a subset of Microsoft's popular Visual Basic (VB6) with modifications and extensions to address the special needs of microcontroller programming. The ZX series microcontrollers were designed to run ZBasic programs efficiently. At the time of publication, the ZX series devices available from Elba Corp. include several models with different capabilities, based on several different CPU types as shown in the table below. The various ZX devices will be referred to collectively in this document as simply ZX unless the context requires a specific reference to a particular model.

Underlying CPU Type for ZX Devices

Device	CPU	Status
ZX-24, ZX-40, ZX-44	mega32	discontinued
ZX-24a, ZX-40a, ZX-44a	mega64	discontinued
ZX-24p, ZX-40p, ZX-44p, ZX-24n, ZX-40n, ZX-44n	mega64P	current
ZX-24r, ZX-40r, ZX-44r, ZX-24s, ZX-40s, ZX-44s, ZX-40t, ZX-44t	mega128P	current
ZX-1281, ZX-1281n	mega1281	current
ZX-1280, ZX-1280n	mega1280	current
ZX-328n, ZX-3281, ZX-32n, ZX-321	mega328P	current
ZX-24x, ZX-32a4	xmega32A4	current
ZX-24u, ZX-128a4u	xmega128A4U	current
ZX-128a1	xmega128A1	current

The ZX-24, ZX-24a, ZX-24p, ZX-24n, ZX-24r and ZX-24s are 24-pin DIP format modules that are pin-compatible with the Parallax Basic Stamp and the NetMedia BX-24 microcontrollers. However, many improvements have been made with respect to those pioneering products in order to provide a more powerful and flexible programming platform. Because of its ready-to-use configuration it is an ideal starting point for someone just beginning to use microcontrollers. Yet, it has enough powerful capabilities for experts and advanced users as well. See Appendix D for more detailed information on the ZX-24 series devices including several suggested connection diagrams and detailed descriptions of the function of each pin. The ZX-24x has the same basic pinout as other 24-pin ZX devices but it operates on a maximum of 3.6 volts. Moreover, some of the special functions (USART, PWM, I2C, etc.) are on different pins as compared to the other 24-pin ZX devices. Additional information on pin assignments can be found in the ZBasic System Library Reference Manual.

The ZX-40, ZX-40a, ZX-40p, ZX-40n, ZX-40r, ZX-40s and ZX-40t are 40-pin DIP format integrated circuits that are pre-programmed with the ZX control program. However, in order to use them you must add a few external components such as a regulated supply, a crystal, a memory chip and serial interface circuitry. The compensating advantages of the 40-pin parts are reduced cost, more design flexibility and the availability of more I/O pins. These attributes make the 40-pin ZX well-suited for advanced projects and commercial applications. See Appendix E for more detailed information on the 40-pin ZX devices including several suggested connection diagrams.

The ZX-44, ZX-44a, ZX-44p, ZX-44n, ZX-44r, ZX-44s and ZX-44t are 44-pin TQFP format integrated circuits that are pre-programmed with the ZX control program. Similarly, the ZX-1280 and ZX-1281 are 100-pin and 64-pin, respectively, TQFP format integrated circuits that are pre-programmed with the ZX control program. All of these devices are bare CPUs that require some additional support components. Because of their smaller size, they are well suited for larger volume applications or those that require compact size. See Appendix F for more detailed information on the 44-pin ZX devices including several suggested connection diagrams. See Appendix G for more detailed information on the suggested ZX-1281 circuitry and Appendix H for the ZX-1280.

The ZX-328n and ZX-328l are a 28-pin DIP format integrated circuit that is programmed with the ZX native mode control program. In order to use the ZX-328n or ZX-328l you must add a few external components such as a regulated supply, a crystal and serial interface circuitry. See Appendix I for more detailed information on the ZX-328n and ZX-328l including several suggested connection diagrams.

The ZX-32n and ZX-32l are a 32-pin TQFP format integrated circuit that is programmed with the ZX native mode control program. In order to use the ZX-32n or ZX-32l you must add a few external components such as a regulated supply, a crystal and serial interface circuitry. See Appendix J for more detailed information on the ZX-32n and ZX-32l including several suggested connection diagrams.

The ZX-32a4, ZX-128a4u and ZX-128a1 are 44-pin TQFP and 100-pin TQFP format ATxmega integrated circuits that are programmed with the ZX native mode control program. In order to use the ZX-32a4, ZX-128a4u or ZX-128a1 you must add a few external components such as a regulated supply, a crystal and serial interface circuitry. See Appendix K and Appendix L and for more detailed information on the ZX-32a4, ZX-128a4u and ZX-128a1, respectively, including several suggested connection diagrams. Note, particularly that the ATxmega-based ZX devices run at twice the speed of most other ZX devices and operate at a maximum of 3.6 volts.

It should be noted that ZBasic is a superset of NetMedia's BasicX language, incorporating many improvements and advanced features. Many, perhaps even most, BasicX programs can be recompiled using the ZBasic compiler and the resulting code should run exactly as it would on the BX-24 microcontroller except for differences due to the fact that the ZX is twice as fast as the BX-24. See Chapter 7 for more information on compatibility issues.

Some ZX devices run in "native mode" meaning that the ZBasic application is compiled to native object code for the processor. In contrast, other ZX devices contain a "virtual machine" (commonly referred to as a VM) and ZBasic applications are compiled to instruction codes that are executed by the VM. The primary advantage of the native mode devices is that the application executes faster but the disadvantage is that the resulting program is generally larger than the corresponding VM mode program.

Some ZX devices mentioned in this manual have an "e" or "u" suffix, e.g. ZX-24pe, ZX-24nu. These "extended" devices are produced by Oak Micros under license from Elba Corp. and are available in both VM and native mode versions.

1.3 Support for Generic Microcontrollers

In addition to the ZX microcontroller family, the ZBasic compiler can produce code for generic microcontrollers including more than 100 of the Atmel AVR line comprising members of the ATtiny, ATmega, ATxmega lines as well as some Atmel CAN and USB devices. See Appendix B for a complete list of supported generic microcontrollers. Also, see Chapter 5 for information about creating applications for generic target devices.

1.4 Conventions

In this manual, specific examples of things that can appear in your ZBasic program are shown in Courier typeface as are descriptions of the syntax of program elements. Otherwise, discussion and general references to program elements will be shown in Arial typeface.

When describing the syntax of a ZBasic language element, sometimes portions of it are optional. This fact is indicated by enclosing the optional portion with left and right square brackets. Literal elements, those that must appear as typed (but without regard to alphabetic case, of course) appear as themselves. Often, it is useful to describe an element generally rather than giving a specific instance. When this is done, the description will be italicized and enclosed with left and right angle brackets and subsequent text will describe the characteristics of the generalized item.

Here is an example that uses these conventions:

```
[Private] Const <name> As <type> = <value>
```

The idea being conveyed is that this program element may begin with the `Private` keyword or it may be omitted, i.e. it is optional. The keywords `Const` and `As` are specified literally as is the equal sign. The name of the constant to be defined, its type and value are all specified using a descriptive word or phrase enclosed in angle brackets. These placeholder components will be explicitly described in the discussion of the program element.

In some instances, there may be more than one optional element that may appear in a given position. In these cases, the vertical bar character is used to indicate that any one element may be used, but not more than one, of course.

```
{Public / Private} Const <name> As <type> = <value>
```

This conveys the idea that either the keyword `Public` or the keyword `Private` may be present but it is also permissible for neither to be present.

In other cases, there may be a set of elements from which exactly one must be chosen.

```
{Public / Private / Dim} <name> As <type>
```

The curly braces indicate that one of the keywords `Public`, `Private` or `Dim` must precede the `<name>` element.

In various places in this document you will find a text box with information about BasicX compatibility. Most often, the text will describe how some particular feature is different when the BasicX compatibility mode is enabled. In a few instances, the text will describe how ZBasic is fundamentally different from BasicX. This information is most useful to those who are upgrading to ZBasic from BasicX. If that is not the case for you, you may safely ignore such information.

The microcontroller for which you are compiling your application is referred to as the *target device*. A target device may be a ZX device or it may be a generic target device. When referring to target devices, the phrase *ZX device* is used when the discussion applies only to ZX devices. Otherwise, the target device is referred to generally as a *ZBasic device* or *ZBasic target device* when the discussion applies to both ZX devices and generic target devices.

Chapter 2 - ZBasic Language Elements

The logic of a program is described in the ZBasic language using identifiers, keywords, literal values, expression operators, statements, etc. These elements are all more fully described in this chapter.

2.1 Identifiers

In ZBasic, as in most computer languages, all constants, variables, subroutines and functions have a name. Such names are generally referred to as identifiers and each computer language has rules that specify how an identifier may be formed. In ZBasic an identifier must begin with an alphabetic character (A-Z, a-z) and may contain zero or more additional characters which may be alphabetic, numeric (0-9) or an underscore. There is virtually no maximum number of characters that an identifier may contain (in reality it is limited by the amount of memory available on your computer) but as a practical matter identifiers seldom exceed 15 to 20 characters. Beyond that they become somewhat cumbersome to type.

As mentioned previously, the alphabetic case of the letters of an identifier is insignificant. The variable `myVar` is considered the same as the variable `MyVar`.

There is a set of identifiers that are reserved for special purposes and cannot be otherwise used in your program. Some of the reserved words are data type names, some are used as keywords in ZBasic statements (`If`, `For`, `Sub`, etc.). Some reserved identifiers have no current use in ZBasic but they are reserved nonetheless because they are keywords in other Basic dialects and may be incorporated into ZBasic in the future. If you attempt to use a reserved word in a role other than its predetermined role the compiler will issue an error message pointing out the unacceptable use. See Appendix A for a complete list of reserved words.

2.2 Data Types

The table below describes the fundamental data types available in ZBasic. There are a few additional data types that are used for special purposes that will be described later.

Fundamental Data Types	
Data Type Name	Range of Values
Boolean	True, False
Byte	0 to 255
Integer	-32,768 to 32,767
UnsignedInteger	0 to 65,535
Long	-2,147,483,648 to 2,147,483,647
UnsignedLong	0 to 4,294,967,295
Single	approximately ±1.4e-45 to ±3.4e+38 and 0.0
String	0 to 255 characters

The special type `DataAddress` is a synonym for the type needed to hold a RAM address. On AVR-based devices it is the same as `UnsignedInteger` but on larger devices it is the same as `UnsignedLong`.

It is important to note that some mathematical operations on floating point values (type `Single`) result in certain special values that indicate an exceptional result. Among these special values are ones that represent positive infinity, negative infinity and a general category referred to as "not a number", called `Nan`. The System Library function `SngClass()` returns a value indicating the class to which a floating point value belongs. See the ZBasic System Library Reference Manual for more information on the `SngClass()` function.

In addition to the fundamental data types enumerated in the table above, ZBasic supports several additional special-purpose types. There are two additional integral-value types that are collectively referred to as sub-byte types – `Bit` and `Nibble`. These are useful for reducing the amount of space

used for small-valued data but they are not quite as efficient with respect to code size as using `Byte` types. See Section 3.24 for more information on these special types. Also, two special string types are available – Bounded Strings and Fixed Length strings. These are supported largely for compatibility with BasicX but may be useful in special situations. See Section 2.11.1 and Section 2.11.2 for more information on Bounded Strings and Fixed Length strings, respectively.

ZBasic is a strongly typed language that doesn't allow you to freely intermix different data types. This helps prevent programming errors or unexpected program behavior that results from unanticipated type conversions. However, a series of well-defined type conversion functions is provided that allow you to make explicit type conversions when necessary. See Section 2.13 for more information. Lastly, Section 3.28 contains technical details of the implementation of the various data types. This information is generally not necessary for most programming purposes but it is sometimes useful to know.

2.3 Modules

A ZBasic program may comprise one or more modules each contained within a separate file. When a module is compiled, a module name is derived from the name of the file. The module name is mostly for internal use by the compiler but you'll find references to the module name in the generated map file. Also, in some special situations you may need to use the module name as part of an identifier in order to refer to a variable or constant in your code. This is usually not necessary but it is good to remember that it is available. See Section 3.1 for more information on resolving identifier references using the module name.

The module name is derived by first removing the path prefix, if any, and the extension (typically `.bas`). The remaining characters of the filename are examined and any character that is not a letter, a digit or an underscore is replaced with an underscore. For the most part, this process converts the filename into a legal ZBasic identifier.

As an example, consider the filename “`C:\temp\my test code-1.bas`”. When you compile this module the derived module name is `my_test_code_1`. If you compile a file whose corresponding module name does not begin with a letter, the compilation will proceed as normal. However, you will not be able to refer to that module by its module name since it will not be a legal ZBasic identifier.

A module is divided conceptually into two sections: an options section which, if present, must be first, and a definitions section. The definitions section may also be omitted but with no definitions, the module serves no useful purpose.

2.3.1 The Options Section

The options section contains directives that tell the compiler how to process the definitions section that follows. In the absence of any option directives, the compiler uses certain default settings as described below. This means that it is perfectly reasonable, and quite common, not to have any option directives in a module. The options section may also contain comments. This is useful for documenting what you're trying to accomplish with each option directive.

Note that some of the settings affected by the option directives described in this section may also be affected by compiler command line options. Option directives have precedence over command line options. See Section 10.1 for more information on command line options.

Unless otherwise indicated, option directives have global scope and may only appear in the first module compiled.

Option <pin>

Option <pin> <pin-state>

This option directive provides a way to configure an I/O pin to be an input or an output. If it is being configured to be an input, you can further specify whether or not the pull-up resistor is enabled. If it is being configured to be an output, you can further specify the logic level that you want to be output. This option directive may only appear in the first module compiled - usually the one containing the `Main()` subroutine.

The `<pin>` element specifies the pin to be configured. The first way to specify the pin to configure is to use the word `Pin` followed immediately by a physical pin number, e.g. `Pin20`. Of course, the set of pin numbers that can be specified in this way depends on the ZX processor that you are using. For example, for the 24-pin ZX devices the allowable values are `Pin5` through `Pin20` as well as `Pin25`, `Pin26` and `Pin27`. The second way to specify the pin to be configured is to give the port letter and a bit number of the port with a period separating them, e.g. `C.0`. The advantage of the second method is that it is largely independent of the ZX processor that you're using.

The `<pin-state>` element may be one of the entries in the table below:

Pin Configuration Values	
Pin State	Meaning
<code>zxInputTriState</code>	The pin should be a tri-state input (pull-up disabled).
<code>zxInputPullUp</code>	The pin should be an input with the pull-up enabled.
<code>zxOutputLow</code>	The pin should be an output set to logic zero.
<code>zxOutputHigh</code>	The pin should be an output set to logic one.

Alternatively, for compatibility with the `Option Port` directive, the `<pin-state>` element may be a quoted string containing a single character `T`, `P`, `0` or `1` corresponding to the entries in the table above. For convenience when using the 24-pin ZX devices, `RedLED` and `GreenLED` are allowed as synonyms for `Pin25` and `Pin26` respectively. When using these synonyms the `<pin-state>` must be specified as `On` or `Off`.

By default, all pins are configured to be inputs with the pullup disabled.

Examples

```
Option Pin5 zxOutputLow      ' make pin 5 an output at logic zero
Option Pin20 "T"             ' make pin 20 a tri-state input
Option C.0 zxOutputHigh     ' make bit 0 of Port C an output at logic one
Option RedLED On             ' turn on the red LED
```

Option <port>

Option <port> <config-string>

This option directive provides a way to configure an entire I/O port at once instead of doing so pin by pin. For each bit of the port being configured to be an input, you can further specify whether or not the pull-up resistor is enabled. For each bit of the port being configured to be an output, you can further specify the logic level that you want to be output. This option directive may only appear in the first module compiled - usually the one containing the `Main()` subroutine.

The allowable values for the `<port>` element are `PortA`, `PortB`, `PortC`, etc. Note, however, that `PortB` and `PortD` will rarely be used with the 24 pin ZX devices since the many of the pins of those ports are not directly available. The `<config-string>` element must be a series of characters, enclosed in quote marks, which specify for each bit one of four configuration states. The allowable configuration characters are described in the table below. The `<config-string>` must contain exactly 8

configuration characters, the leftmost of which corresponds to the bit 7 of the port and the rightmost of which corresponds to the least bit 0 of the port. For devices that are missing particular bits on a port or have analog-only pins, you must use the x configuration character as a placeholder for the missing bit(s).

Certain bits of PortB and PortD have dedicated uses on some ZX devices that require a specific configuration. Due to this requirement, those bits are protected from being changed by user-specified configuration. By default, all bits are configured to be inputs with the pullup disabled except for the special purpose bits.

Port Configuration Designators

Configuration Character	Meaning
T	The corresponding bit should be a tri-state input (pull-up disabled).
P	The corresponding bit should be an input with the pull-up enabled.
D	The corresponding bit should be an input with the pull-down enabled.
0	The corresponding bit should be an output set to logic zero.
1	The corresponding bit should be an output set to logic one.
X	The corresponding bit is not a valid digital I/O pin for the device.

Note that it is not necessary to configure your ports using this option directive. The configuration may also be done using the System Library subroutine `PutPin()` or the built-in registers like `Register.DDRA`. Also, note that if you specify configuration directives for both the pins and the port containing them, the last occurring directive will prevail.

Example

```
Option PortA "TPTP0001"
```

Option Base

Option Base <value>	Default: 0
----------------------------------	------------

When an array is defined, the lower bound for each dimension may be specified explicitly (the recommended practice) or it may be omitted. In the latter case, the compiler will utilize the default lower bound which you can select using the `Option Base` directive. The `<value>` element of the directive must be either 0 or 1. It is important to note that this directive has local scope, i.e., it only affects arrays defined in the module in which the directive appears.

This option does not affect arrays defined without either the upper bound or the lower bound specified (e.g. a Program Memory data item, an array that is a parameter to a procedure, a based array, or a `ByRef` array). In these cases, the lower bound is always assumed to be 1.

Example

```
Option Base 1
```

Option Explicit

Option Explicit	Default: off
------------------------	--------------

Traditionally, the purpose of the `Option Explicit` directive has been to require that each variable in the program be explicitly defined. Historically, early Basic dialects allowed programs to be written so that when the compiler encountered a new variable name, it automatically created a definition for it. Although this may have been thought to be a nice feature, it turned out to be the source of many programming errors and bugs. If you accidentally misspelled a variable name, a new variable was created entirely unbeknownst to you. Because of this serious drawback implicit variable definition is not supported by ZBasic. The `Option Explicit` directive is accepted by the compiler for compatibility reasons but neither its presence nor its absence affects any aspect of the compiler's operation.

Option Language

Option Language { BasicX | ZBasic }

Default: ZBasic

By default the compiler processes modules using the rules and defaults for the ZBasic language. This option directive can be used to instruct the compiler to process the module using the syntax rules defined by a specific language variant. If Option Language BasicX is specified, the compiler will process the module in BasicX compatibility mode. This may be useful in certain peculiar situations if you have existing modules written for BasicX. In most cases, existing BasicX code will compile correctly in ZBasic mode with few, if any, changes.

If you need to utilize BasicX compatibility mode you should be aware that none of the enhanced language features nor any of the enhanced System Library routines of ZBasic will be available in that module. Also, using Option Language BasicX changes the defaults for some of the other options as listed in the table below. You may still change the prevailing setting of these other options by utilizing the related option directive either before or after this option directive.

Option Directive	Default for BasicX	Default for ZBasic
Strict	On	Off
AllocStr	Off	On

This option has local scope, i.e., it only affects compilation of the module in which the option appears.

Option AllocStr

Option AllocStr [On / Off | Default]

Default: On

By default, native compilation mode uses an allocation strategy for strings called dynamic string allocation. In contrast, BasicX uses an allocation strategy called static allocation. This option directive allows you to select the dynamic allocation strategy (Option AllocStr On) or the BasicX-compatible static allocation strategy (Option AllocStr Off). Specifying Option AllocStr Default sets the option to the default for the currently selected language. If neither On, Off nor Default is specified, the value On is assumed. See Section 3.28.2 for more information on the two allocation strategies.

This option has local scope, i.e., it only affects string variables defined in the same module in which the option appears.

Examples

```
Option AllocStr  
Option AllocStr Off
```

Option StringSize

Option StringSize {<value> | Default}

Default: 20

When string variables are defined, space to hold the characters of the string is allocated either statically or dynamically (see Option AllocStr, above). When the space is statically allocated a fixed number of bytes of space is set aside for the string at compile-time thus setting the maximum size of that string. This option specifies the number of characters that should be reserved for statically allocated string storage. See the description of the Option AllocStr directive for more information about dynamically allocated string storage. If Option StringSize Default is specified the default value for the string size is used. This is useful if the string size was set using a command line option and you want to restore it to the default for this particular module. The maximum allowable string size is 255.

This option has local scope, i.e., it only affects string variables defined in the same module in which the option appears.

Examples

```
Option StringSize 25  
Option StringSize Default
```

Option Strict

Option Strict [On Off Default]	Default: Off
---	--------------

This option directive, supported for BasicX compatibility, enables or disables so-called “strict syntax checking”. You can enable strict mode by using `Option Strict` by itself or by including the keyword `On`. You disable strict syntax checking by using `Option Strict Off`. Specifying `Option Strict Default` sets the option to the default for the currently selected language.

The implications of strict syntax checking are noted in the description of each affected element but a summary of the effects is given here.

- the loop index variable of a For loop has restrictions on lifetime, visibility and accessibility
- logical operators like Not, And and Or may not be used with signed integral data types

This option has local scope, i.e., it only affects compilation of the module in which the option appears.

Example

```
Option Strict Off
```

Option TargetCPU

Option TargetCPU <CPU-type>	Default: ZX24
--	---------------

This option has been deprecated. Use `Option TargetDevice` instead.

Option TargetDevice

Option TargetDevice <device-name>	Default: ZX24
--	---------------

Due to differences between the various target devices, the compiler needs to know for which device it should compile the code. A complete list of supported device names is given in Appendix B.

Example

```
Option TargetDevice ZX24a
```

Option DeviceParameter

Option DeviceParameter <parameter-name> <parameter-value>	Default: see text
--	-------------------

This directive allows specifying a device configuration parameter usually for a generic target device (however, some can be used for native mode ZX devices). The `<parameter-name>` element has the form of an identifier (beginning with an alphabetic character and containing only alphabetic and numeric characters). The `<parameter-value>` element may be in the form of an identifier, a quoted string, or a numeric value depending on the particular parameter. See Section 5.3 for more information on device configuration parameters.

Examples

```
Option DeviceParameter ClockFrequency 16000000
Option DeviceParameter Package "PDIP-40"
```

Option PortPinEncoding

Option PortPinEncoding {On / Off}

Default: see text

When the compiler encounters an I/O port pin designator like C.2 it can convert it to the corresponding physical pin number for the target device or it can convert it to a composite value whose fields specify the port and the bit number. This directive controls which of these conversions is performed. When the option is off, the result is a pin number and when it is on the result is the encoded port/pin value. The encoded result is a Byte value that has the binary form 1PPPnnn where PPPP represents a port index (PortA=0, PortB=1, etc.) and nnn represents the bit number (0-7). The option is on by default for all generic target devices and for ZX devices based on the mega328P, mega128, mega1280, mega1281, xmega32A4 and xmega128A1 and off by default for other devices.

Example

```
Option PortPinEncoding On
```

Option CodeLimit

Option CodeLimit <limit-value>

Default: none

This option can be used to have the compiler check the size of the generated code and issue an error message if it exceeds a specified size. The limit value is specified by a decimal number, optionally using the suffix K or k to denote a multiple of 1024.

Examples

```
Option CodeLimit 32K
Option CodeLimit 32768
```

Option SignOn

Option SignOn {On / Off}

Default: On

This option can be used to control the generation of a sign-on message when the ZX begins running after a reset. The flag to control the sign-on is stored in Persistent Memory of the processor. This is important to remember because if you download one program to the ZX that, say, turns it off and then you download another program that has no Option SignOn directive at all, the sign-on will still be in the off state. In other words, if you want to ensure that it is in particular state for your particular program, be sure to include this directive. Of course, if you've never turned it off, it will still be in the On state.

Example

```
Option SignOn Off
```

Option DefaultISR

Option DefaultISR {On / Off}

Default: see discussion

For native mode devices, this option can be used to control whether a "catch all" ISR is automatically included in your application. By default, it is included for ZX devices and not included with generic target devices. Note that this option has no effect if you explicitly define a default ISR in your application source code.

Example

```
Option DefaultISR Off
```

Option CallFunctions

Option CallFunctions {On / Off}

Default: Off

For native mode devices, this option can be used to allow your application code to invoke functions as if they were subroutines (the returned value is discarded). Note that even when CallFunctions is enabled, functions returning user-defined types like structures, unions and class objects are not allowed to be called as if they were subroutines.

The presence of this option directive in a module affects only that module; the command line option --call-functions can be used to affect multiple modules.

Example

```
Option CallFunctions On
```

Option TaskStackMargin

Option TaskStackMargin <margin-value>

Default: 10

When the compiler compares the size of the stack allocated to a task to the estimated task stack usage it pads the estimate with a safety margin. This option can be used to specify a value for that safety margin different from the default of 10 bytes. The safety margin may be any non-negative decimal integral value.

Example

```
Option TaskStackMargin 20
```

Option ExtRamConfig

Option ExtRamConfig [On | Off | <constant-expression>]

Default: (see text)

For ZBasic target devices that support external RAM, you may control whether external RAM is enabled using this directive. Using the keyword On enables external RAM support in the default configuration (16-bit address, no wait states). If you need a different configuration, you may specify a constant expression whose 16-bit value is written to the external RAM configuration registers of the CPU. For ATmega-based devices, the high byte is written to the register XMCRB register and the low byte is written to XMCRA register. (Note that for the mega128 CPU, the high bit of the XMCRA register is undefined but the configuration value must have that bit asserted in order for the external RAM interface to be enabled.) For ATxmega-based target devices, the low byte is written to the EBI_CTRL register while the high byte is written to the EBI_CS1_CTRLB register.

For ATmega-based device, the default external RAM configuration value is &H0080. For ATxmega-based devices, the default external RAM configuration value is &H014d.

Examples

```
Option ExtRamConfig On
Option ExtRamConfig &H0084  ' enable with 1 wait state (ATmega)
```

Option RamSize

Option RamSize <constant-expression>

Default: (see text)

The compiler compares the aggregate size of the statically allocated data items to the amount of RAM available in the target device. If the aggregate size is too large, the compiler issues a warning to that effect. For ZBasic devices that support external RAM (e.g. ZX-1281), the compiler must know the actual size of User RAM (including external RAM) in order to avoid generating meaningless warnings. That is the purpose of this directive, which may only be used on ZBasic devices that support external RAM. The example below shows the simplest way of specifying the augmented User RAM size. The example is for the case where the maximum amount of additional RAM was added.

Example

```
Option RamSize 65536 - Register.RamStart
```

Option HeapSize

Option HeapSize <constant-expression>

Default: (see text)

This directive controls the allocation of RAM between the string heap and the Main() task stack. The effect is to set a hard limit beyond which the heap will not grow, thereby preventing the heap from encroaching on the Main() task stack. If a size less than the minimum size (which varies by device) is specified, the minimum is used instead. Also, for devices with external RAM, the special value of 65535 (&xFFFF) can be used to specify that all external RAM should be used for the heap. See the section on setting heap and task stack sizes for more information including the default heap size.

For VM mode devices the directive has an additional effect. When the compiler performs various checks regarding the use of RAM, it takes into account an amount of RAM reserved for the heap (from which memory for strings, for example, is allocated). The size specified for the heap will effect whether the compiler will issue a warning about excessive use of RAM in a particular circumstance.

Example

```
Option HeapSize 500
```

Option HeapReserve

Option HeapReserve <constant-expression>

Default: (see text)

After compilation, the compiler analyzes the use of RAM and compares it to the amount of RAM available on the target device. If a specific task stack size is specified for the Main() task, the compiler adds the amount of statically allocated data, the Main() task stack size and the minimum heap size and compares this total to the amount of available RAM. If the aggregate total exceeds the available RAM an error message will be issued.

This directive allows you to specify the minimum space you want to reserve for the heap. See the section on setting heap and task stack sizes for more information including the default heap size.

Option HeapLimit

Option HeapLimit <constant-expression>	Default: n/a
---	--------------

This directive provides a different way of allocating RAM between the string heap and the task stacks. The value provided is interpreted as a RAM address and specifies the limit beyond which the heap will not grow. This directive is perhaps more useful with target devices that support external RAM. See the section on setting heap and task stack sizes for more information.

Option MainTaskStackSize

Option MainTaskStackSize <constant-expression>	Default: n/a
---	--------------

This directive provides a different way of allocating RAM between the string heap and the task stacks. The value provided is interpreted as the desired size of the task stack for the Main() task and the heap limit is set at the end of the task stack. See the section on setting heap and task stack sizes for more information.

Option TxQueueSize

Option RxQueueSize

Option TxQueueSize <constant-expression>	Default: 25
Option RxQueueSize <constant-expression>	Default: 50

These directives can be used to modify the sizes of the default transmission and reception queues for the console channel (typically Com1). They are effective only for native mode devices such as the ZX-24n and generic target devices. The value of zero for one or the other is useful for setting up the console for transmit-only or receive-only operation.

Example

```
Option TxQueueSize 100
```

Option ConsoleSpeed

Option ConsoleSpeed [<constant-expression> / Default]	Default: see text
--	-------------------

This directive can be used to set the default speed of the serial channel assigned to the console (typically Com1 but see Option Console). It is not supported on VM devices with a VM version older than v3.0.4 nor is it supported on any ZX device based on the mega32. The set of allowable baud rates is given in the table below; the shaded entries are supported only on native mode devices.

300	600	1200	2400	4800	9600	14400
19200	28800	38400	57600	76800	115200	230400

It is important to note that for VM devices, the speed selection is made part of the download image and is written to Persistent Memory when the code is downloaded. If, later, a program is downloaded that does not contain Option ConsoleSpeed directive, the Persistent Memory location containing the console speed indicator will be unchanged from the earlier setting.

The default console speed is 19200 for all ZX devices and it is zero for all generic target devices. For native mode devices, the console speed can be set to zero thereby requesting that the console serial channel not be open initially. In this case, if you want to use the console serial channel you must specifically open it. For ZX devices, setting the console speed to zero when Com1 is the console has the further effect of allowing Com1 to be fully closed when CloseCom() is called. This is in contrast to the normal ZX behavior of always keeping Com1 open. The primary advantage to specifying the speed as zero is that it results in a smaller executable size provided that your application does not explicitly open the console serial channel.

Example

```
Option ConsoleSpeed 9600
```

Option Com1Speed

<code>Option Com1Speed [<constant-expression> Default]</code>	Default: 19200
---	----------------

This option has been deprecated. Use Option ConsoleSpeed instead.

Option Console

<code>Option Console <serial-channel-designator></code>	Default: see discussion
<code>Option Console <pin-designator> [Inverted]</code>	
<code>Option Console None</code>	

This option, available only on native mode ZBasic target devices (both ZX and generic), allows you to specify the serial channel (using the first form) to which output will be routed when Debug.Print, Console.Write or Console.WriteLine are used and from which input will be obtained when Console.Read or Console.ReadLine is used. If used, the directive must appear in the first module compiled and it must appear after the target device is specified. Any serial channel designator that is valid for the target device (e.g., Com1 through Com12) may be used. Note that if the specified serial channel is not open with transmission enabled (i.e. a transmit queue has been specified) at the time output is generated, the resulting character stream will be silently discarded. Also, the console read functions will never return if the specified serial channel is either not open or not enabled for reception.

The default console channel is Com1 for devices with at least one hardware UART. The compile-time constant Option.ConsoleChannel can be used in your code when it is necessary to perform different actions depending on the console channel setting (1 implies Com1, 2 implies Com2, etc.). The value of zero for Option.ConsoleChannel means that a serial channel is not being used for console output.

The second form allows you to specify a pin to be used for the transmission of console data using the “bit-bang” method. This is most useful on devices that do not have a hardware UART or when you don’t want to dedicate a timer for the software UARTs. When this form is used, the pin specified is made an output and set to the “idle” state (high for non-inverted mode, low for inverted mode). It is important to note that when the bit-bang mode is employed the console is transmit-only; data reception is not supported. Note, also, that there is no output queue associated with the console in this mode. The compile-time constant Option.ConsolePin can be used in your code if you need to refer to the specified pin; zero means that no console pin as been specified.

The third form allows you to specify that your application won’t use the console at all. This can reduce the code size and RAM use of your application.

Examples

```
Option Console Com2      ' use Com2 for console I/O
Option Console A.2       ' use pin A.2 for console output
Option Console Pin 12 Inverted ' use pin 12 for console output (inverted)
```

Option ComChannels

<code>Option ComChannels [<constant-expression>]</code>	Default: 4
---	------------

This directive, useful only on native mode devices, can be used to reduce the amount of RAM needed by an application by specifying a lower number of software UART channels to be supported. For example, if your application only uses one software UART channel you can save approximately 75 bytes of RAM by using this directive to specify that only one channel is needed. The value given must be in the range 0-4.

The value zero can be used to tell the compiler that your application uses no software UART channels at all. Doing so can reduce code size in cases where the compiler cannot otherwise deduce which serial channels are used by the application.

Example

```
Option ComChannels 1           ' request support for one SW UART channel
```

Option RTC

Option RTC [On Off]	Default: (see text)
--------------------------------	---------------------

This option, intended primarily for use with ZX devices, can be used to disable the RTC. If an application does not need the RTC, turning it off with this option will reduce code size and make the RTC timer available for other purposes. The pre-defined constant Register.RTCEnabled can be used determine if the RTC is enabled for an application, particularly useful in the expression of a #if conditional.

Example

```
Option RTC Off
```

Option AtnChar

Option AtnChar [<constant-expression> Default Off]	Default: off
---	--------------

This directive can be used to specify a special character that can be used instead of DTR signaling to perform downloads to a ZBasic device. This is useful in situations where the communication channel does not support DTR or is not capable of toggling DTR fast enough to be recognized as an attention signal. The range of allowable values for the ATN character is 0 to 31 (&H00 to &H1F).

It is important to note that the ATN character is made part of the download image and is written to Persistent Memory when the code is downloaded. If, later, a program is downloaded that does not contain Option AtnChar directive, the Persistent memory location containing the ATN character value will be unchanged from the earlier setting.

The pre-defined variable Register.ATNChar can also be used to set or disable the ATN character. See section 3.7.2 for more details.

Example

```
Option AtnChar &H04
```

Option Include

Option Include [Public Private] <item-list>	Default: none
--	---------------

This directive can be used to add one or more pre-defined structures to the application. If the Private keyword is present, the structures will be private to the module in which the directive appears; otherwise the structures will be public.

The <item-list> element consists of one or more names of pre-defined structures separated by commas. A list of the available pre-defined structures and their definitions is available in Appendix C.

Example

```
Option Include Private Port_t, Timer16_t
```

Option Objects

Option Objects

Default: none

This option enables the ZBasic object-oriented extensions. If used, it must appear for the first time in the first module compiled. If it does appear in the first module compiled then its presence in other modules is silently ignored. This option implicitly enables subroutine/function overloads as well.

Example

Option Objects

Option Overload

Option Overload

Default: none

This option enables the definition of subroutine/function overloads. If used, it must appear for the first time in the first module compiled. If it does appear in the first module compiled then its presence in other modules is silently ignored. Note that this option is redundant if Option Objects is present.

Example

Option Overload

Option Namespaces

Option Namespaces

Default: none

This option enables the definition of namespaces. It may be used in any module and its occurrence enables the definition of namespaces in that module and all subsequently processed modules. Nevertheless, it is generally advisable to place it in the first module compiled. When this option is used, the identifier *Namespace* becomes a keyword and cannot be used as a procedure or variable name.

Example

Option Namespaces

Option NameStyle

Option NameStyle <name-style>

Default: C

This option can be used to change the allowable format of a ZBasic identifier. There are two supported name styles, Basic and C. Under the Basic name style, an identifier must begin with an alphabetic character and thereafter contain only alphabetic and numeric characters and underscores. With the C name style, the rules are expanded to allow an identifier begin with an underscore. The C name style is supported primarily to allow accessing procedures and variables in code written in C and/or assembly language in cases where the names begins with an underscore.

Example

Option NameStyle Basic

Option CodeType

Option CodeType <code-type>

Default: C

This option, available only when compiling for native mode devices, specifies the type of intermediate code to generate corresponding to the ZBasic application which is then converted to target processor instructions by the back-end compiler. The default is to generate C code but in some cases, it may be useful to direct the ZBasic compiler to generate C++ code by specifying the <code-type> as C++ or CPP.

Example

```
Option CodeType C++
```

Option Arduino

Option Arduino ["<arduino-directory>"]

Default: off

This directive, available only when compiling for native mode devices, is used to specify that an application is being built using code from the Arduino code base. The optional <arduino-directory> specifies a directory where the ZBasic-Arduino compatibility libraries and include files can be found.

The presence of this option has several effects. Firstly, it directs the compiler to generate C++ code and enables ZBasic object extensions. Secondly, if the optional <arduino-directory> is given, that directory and all of the sub-directories thereof will be searched for any needed header (.h) files and the ArdLib sub-directory thereof will be expected to contain the pre-built ZBasic-Arduino compatibility library for the target device, e.g. ArdLLib/libArdZB_mega328p.a.

See the ZBasic application note for more details on building applications incorporating Arduino code.

Example

```
Option Arduino "C:/projects/ZBasicArduino"
```

Option Notice

Option Notice {On | Off}

Default: on

This directive can be used to enable or disable output from the #notice directives.

Example

```
Option Notice Off
```

Option SerialReadStrobe

Option SerialReadStrobe "<pin-designator>"

Default: none

This directive, available only when compiling for native mode devices, specifies a pin to toggle just before reading the first of nine bit window samples when the function SerialGetByte() is called. When observed with a logic analyzer or similar instrument, this is useful for confirming that the sample is being taken near the middle of the start bit. The pin, which you must configure as an output prior to invoking SerialGetByte(), may be specified in the form <port>. <bit> or as a physical pin number as shown in the two examples below. This option is available on all ATxmega devices and those ATmega and ATTiny devices that can toggle an output pin by writing a bit to the PINx register.

Examples

```
Option SerialReadStrobe A.2           ' use pin A.2 for the strobe
```

```
Option SerialReadStrobe Pin 12      ' use pin 12 for the strobe
```

Option X10Interrupt

Option X10Interrupt <external interrupt designator>

Default: INT0

This option, available only on ATtiny and ATmega-based ZBasic devices, allows you to specify the external interrupt to be used for the zero-crossing input needed for the low-level X-10 routines. If used, it must appear for the first time in the first module compiled and it must appear after the target device is specified.

This option is not supported on VM devices with a VM version older than v3.0.6 and, moreover, only INT0 or INT2 may be specified for VM devices. For native mode devices, it is not supported on xmega-based devices and on mega-based devices the choice is limited to the external interrupts supported by the device in the range INT0 through INT7

Example

```
Option X10Interrupt INT2
```

Option OCModulateEnable

Option OCModulateEnable <value>

Default: see discussion

This option, available only on ZBasic target devices that support OutputCapture modulation, specifies the value to be used to enable output of the modulation carrier timer. The default value, which varies by target device, configures the modulation carrier timer to toggle the timer's compare match output on each compare match. In certain special situations, a different value may be needed to achieve the desired result. The value must be specified such that it is suitable for logically ORing it into the timer's control register after having cleared the set of bits that control output mode.

Option Library

Option Library "<object-library-filename>"

Default: none

This option, which can only be used for native mode devices, specifies the name of an object library that should be linked in when the executable is built. Typically, this directive will appear in a declarations file that identifies the public entities contained in the library. If the specified filename does not have a path prefix, it is assumed that the specified name is relative to the directory of the module containing the directive.

This option may be used in multiple modules as needed.

Example

```
Option Library "mylib.a"
```

2.3.2 The Definitions Section

The definitions section of a ZBasic program may contain constant definitions, variable definitions, subroutine definitions and function definitions. There may be any number of each of these types of definitions and the definitions may occur in any order. It is a common practice, however, to place constant and variable definitions at the top of the definitions section followed by subroutine and function definitions. On the other hand, some programmers prefer to define the constants and variables closer to the routine or routines that use them.

BasicX Compatibility Note

In BasicX mode, all variables and constants must be defined before the first subroutine or function in a module.

Each of these program items may be defined to be `Public` or `Private`. A public item is visible to other modules and may be referenced in the definitions contained in other modules. A private item is visible only within the module in which it is defined. Generally speaking, unless there is a specific need for an item to be public, it should be private. If you make something private and later decide that you need to reference it in another module, it is a simple matter to change the definition from private to public.

Defining Constants

It is often convenient to define constants that can be used in other parts of the program. Doing so generally helps clarify the purpose of the value, assuming a reasonably descriptive name is chosen, and also facilitates easier maintenance and modification of the program.

The syntax for defining a constant is as follows:

```
[Public / Private] Const <name> As <type> = <value>
```

If neither `Public` nor `Private` is specified, the constant will be private. The `<name>` must be a legal identifier as described in Section 2.1. The `<type>` must be one of the fundamental data type names described in Section 2.2 or an `Enum` type (described in Section 3.2). Lastly, the `<value>` element must be value or an expression that has a constant value and is the same type as (or compatible with) the specified `<type>`.

In many cases, the `<value>` will be a simple numeric literal like `-55` or `3.14159`. In the case of string constants, it may be a literal string like `"Hello, world!"`. However, it is sometimes convenient to define a constant in terms of another constant. Consider the example below.

```
Private Const Pi as Single = 3.14159
Const TwoPi as Single = Pi * 2.0
```

You may also use certain System Library functions in the constant's value expression. The restriction is that the expression must be able to be evaluated at compile time.

```
Private Const Pi as Single = ACos(-1.0)
Private Const InitialValue as Single = Sin(Pi / 2.0)
```

The definition of the value of `pi` in the manner shown in the first example is useful because it results in the maximum accuracy of the constant value.

String constants are sometimes useful as well.

```
Public Const VersionNum as String = "V1.0"
Public Const VersionDate as String = "Oct 2005"
Public Const VersionStr as String = VersionNum & " " & VersionDate
Public Const VersionDateUC as String = UCASE(VersionDate)
```

You may define multiple constants of the same or different types on the same line.

```
Const c1 as Integer = 7, c2 as Single = 3e10
```

BasicX Compatibility Note

In BasicX mode, UnsignedInteger, UnsignedLong and String constants are not supported. Also, constant expressions cannot utilize built-in functions.

Defining Variables

To define a variable at the module level (as opposed to within a subroutine or function, described later) the syntax is:

```
{Public | Private | Dim} <name> As <type>
```

Dim has exactly the same effect as Private, i.e., the variable will only be directly accessible to code within the module.

Examples

```
Dim ival as Integer, pulseCount as Byte  
Private busy as Boolean  
Dim msg as String
```

The first example above shows two different variables being defined on the same line.

The initial value of a variable depends on its type and how it is defined. See Section 2.12 for information on variable initialization.

Defining Arrays of Variables

Variables which hold a single value like those discussed above are called scalar variables. ZBasic also supports arrays of variables. An array of a fundamental type may be defined using the syntax:

```
{Public | Private | Dim} <name>(<dim-spec-list>) As <type>
```

The *<dim-spec-list>* is a list of up to 8 dimension specifications each separated from the next by a comma. A dimension specification consists of a constant expression giving the upper bound of elements along that dimension or two constant expressions separated by the keyword To specifying the lower bound and upper bound, respectively, of the elements along that dimension.

When only the upper bound is given, the lower bound defaults to either zero or one depending on whether or not an Option Base directive is in effect or not. If no Option Base directive has been specified, the lower bound is zero.

Note that in order to be passed as a parameter to a subroutine or function an array must have a lower bound of 1. For this reason it is probably more common to define arrays with a lower bound of 1. Many people find it easier to think about arrays this way as well. The default lower bound of zero is kept for compatibility reasons.

Examples

```
Option Base 0  
Dim ival(5) as Integer
```

This defines an array of Integer values with 6 elements. The lower bound is zero and the upper bound is 5.

```
Option Base 1
...
Dim ival(5) as Integer
```

This defines an array of `Integer` values with 5 elements. The lower bound is 1 and the upper bound is 5. The lower bound may be explicitly specified as well. The second example below illustrates the use of a negative value as one of the bounds. Note that the upper bound must be greater than or equal to the lower bound.

```
Public pulseCount(1 to 6) as Byte
Private itemData(-3 to 45, 1 to 4) as Boolean
```

You may also define an array of strings but only when `Option AllocStr` is in effect as it is by default.

```
Dim msg(1 To 4) as String
```

BasicX Compatibility Note

In BasicX mode, each array dimension must have at least two elements. Also, arrays of type `String` are not supported.

There is no checking, either at compile time or at run time, for array index underflow or overflow. If you write code that uses an index outside of the defined range of indices, the results are undefined.

Defining Subroutines

A subroutine is a collection of statements that can be executed by using the subroutine name in a `Call` statement. The advantage of creating subroutines is that we can think of them as logical blocks instead of thinking about all of the details that are dealt with by the statements within the subroutine.

A subroutine may be defined as taking zero or more parameters. If it has parameters defined, you must supply a value for each of the parameters when you invoke the subroutine using the `Call` statement. The syntax for defining a subroutine is shown below.

```
[Public | Private] Sub <name> ( [<parameter-list>] )
    [<statements>]
End Sub
```

If neither `Public` nor `Private` is specified, `Public` is assumed.

The `<parameter-list>` consists of zero or more parameter specifications, each separated from the next by a comma. Note that the parentheses are required even when there are no parameters. The parameters given in the subroutine's definition are called the "formal parameters". The parameters that appear in each invocation of the subroutine are referred to as the "actual parameters".

The syntax for a formal parameter specification is:

```
[ByVal | ByRef] <name> As <type>
```

The `<name>` element is the name by which the passed parameter is known within the subroutine. The `<type>` element is one of the fundamental types listed in the table in Section 2.2. You may also specify a default value for a parameter, a topic that is discussed in more detail in Section 3.18.

The keywords `ByVal` and `ByRef` refer to the method by which the parameter is passed to the subroutine. The keyword `ByVal` means that the parameter is passed to the subroutine "by value" while `ByRef` means

that the parameter is passed “by reference”. Each of these parameter passing conventions has its own advantages and disadvantages and the full explanation of the difference between these two parameter passing methods is given in Section 2.14. Suffice to say here that the primary difference is whether or not the called subroutine is able to change the value of the passed parameter from the perspective of the caller. Some variable types may be passed by one of the methods but not by the other. Again, see Section 2.14 for complete details. The default passing convention, if neither `ByVal` nor `ByRef` is specified, is `ByRef`.

To specify a formal parameter that is an array, simply add a set of parentheses following the parameter name. For example,

```
Sub mySub(ByRef data() as Byte)
End Sub
```

Only one-dimensioned arrays whose lower bound is 1 may be passed as parameters and they must be passed by reference. The upper bound is indeterminate – it is the responsibility of the programmer to ensure that the subroutine does not access elements beyond the upper bound of the passed array. Often, programmers will include a parameter that specifies the upper bound so that the code may safely operate with different sizes of arrays.

```
Sub mySub(ByRef data() as Byte, ByVal count as Byte)
End Sub
```

Lastly, the `<statements>` element in the subroutine definition represents zero or more valid ZBasic statements that are described in Section 2.5 of this manual.

Example

```
Private Const redLED as Byte = 25      ' define the pin for the red LED
Private Const grnLED as Byte = 26      ' define the pin for the green LED

Public Sub Main()
    ' configure the pins to be outputs, LEDs off
    Call PutPin(redLed, zxOutputHigh)
    Call PutPin(grnLed, zxOutputHigh)

    ' alternately blink the LEDs
    Do
        ' turn on the red LED
        Call Blink(redLed)

        ' turn on the green LED
        Call Blink(grnLed)
    Loop
End Sub

Private Sub Blink(ByVal pin as Byte)
    ' turn on the LED connected to the specified pin for one half second
    Call PutPin(pin, zxOutputLow)
    Call Delay(0.5)
    Call PutPin(pin, zxOutputHigh)
End Sub
```

In this program, we have factored out the code that turns an LED on and off into a subroutine named `Blink()`. In the definition of `Blink()`, `pin` is called the formal parameter. In `Main()` where `Blink()` is invoked, the parameters `redLed` and `grnLed` are the actual parameters.

By factoring out the code that was common to blinking the two LEDs we have simplified the program. The details of how an LED is blinked are encapsulated in the definition of `Blink()`. No longer does the `Main()` subroutine need to know how to blink an LED; it just calls the subroutine to handle all of the

details of blinking an LED connected to a specific pin and provides the necessary data for `Blink()` to do the work. In this case, all that is needed is a pin number.

If we wished to do so, we could add another parameter to the `blink` subroutine to specify how long we want the LED illuminated.

```
Private Sub Blink(ByVal pin as Byte, ByVal duration as Single)
    ' turn on the LED connected to the specified pin for the time given
    Call PutPin(pin, zxOutputLow)
    Call Delay(duration)
    Call PutPin(pin, zxOutputHigh)
End Sub
```

With this definition, we would need to add a second actual parameter in each call. For example,

```
Call Blink(redLed, 0.5)
```

It is recommended that you invoke subroutines as shown using the `Call` keyword as shown. However, for compatibility with other Basic dialects it is also possible, although not recommended, to invoke a subroutine by using its name as if it were a statement. If this is done, the actual parameters are not allowed to have enclosing parentheses as illustrated below.

```
Blink redLed, 0.5
```

One powerful aspect of subroutines is that variables and constants may be defined within a subroutine itself. When this is done, the variable or constant is private to the subroutine and cannot be directly accessed from any other routine. Although it is common to place such definitions near the beginning of the subroutine, the definitions may occur anywhere in the subroutine as long as they occur before their first use.

Note that the normal execution sequence of a subroutine may be altered by using the `Exit Sub` statement. When this statement is executed, it causes control to return to the caller immediately bypassing the remaining statements in the subroutine.

Defining Functions

A function is a collection of statements that can be executed by using the function name in place of a value or in an expression. The advantage of creating functions is that we can think of them as logical blocks instead of thinking about all of the details that are dealt with by the statements within the functions.

Like a subroutine, a function may have zero or more parameters defined. If it has parameters defined, you must supply values for each of the parameters when you invoke the function.

The syntax for defining a function is very similar to that for defining a subroutine and is shown below.

```
[Public | Private] Function <name> ( [<parameter-list>] ) As <type>
    [<statements>]
End Function
```

The primary difference is the use of the keyword `Function` in place of `Sub` and the specification of a type for the value to be returned by the function. Like a subroutine, if neither `Public` nor `Private` is specified, `Public` is assumed. The `<parameter-list>` syntax is identical to that for a subroutine.

Example

```
' compute the factorial of the value provided
Public Function Factorial(ByVal val as Long) As Long
    Dim factVal as Long
    factVal = 1
```

```

Do While (val > 1)
    factVal = factVal * val
    val = val - 1
Loop
Factorial = factVal
End Function

```

This function implements a mathematical operation called factorial. The value of N factorial (usually written $N!$) is the product of all the integers from 1 up to and including N. The value of 0! is, by definition, 1. The value of 3! is 6 and so on.

This example introduces a couple of statements that we haven't seen yet but they are fairly straightforward. Notice that a variable named `factVal` was defined within the function. This variable is called a "local variable" and it is not visible to any code outside of the function. The other aspect of this local variable that is different from variables defined at the module level is that this variable only takes up space while the function is executing. When the function returns, the space used by the local variable is reclaimed by the system and can be used for other purposes. See Section 3.1 for more information on the concepts of scope and lifetime of variables.

The second line of the function shows a value being assigned to the variable `factVal`. Instead of using the literal constant value as shown on the right hand side of the equal sign, we could just as well have written an expression that involved several values (perhaps including function invocations) and operators like addition, subtraction, multiplication, etc. This statement is known as an assignment statement and is described in more detail in Section 2.5.1.

The third line of code illustrates a useful variation on the `Do` loop that we have already seen. In this case, the `Do` has a condition associated with it. The condition is tested before every pass through the loop and as long as the conditional expression evaluates to the Boolean value `True`, the statements within the loop are executed again. In this case, the condition tests if the `val` parameter has a value greater than 1, the statements in the loop will be executed. Otherwise, control will transfer to the first statement following the `Loop` statement.

Within the loop, there are two more assignment statements, each of which has an expression on the right hand side, the first involving multiplication and the second involving subtraction. The meaning of the expressions should be self-evident.

The last statement of the function, just before the `End Function`, illustrates how the value to be returned by the function is set. It is useful to think of there being a variable automatically defined that has the same name as the function as well as the same type. We call this a return value variable. Within the function, you can refer to the return value variable, in most cases, just as you would any other variable or parameter. In the example above, instead of introducing a new variable `factVal`, we could have instead simply used the reference to the return value variable `Factorial`. The function's code is re-written below using this idea.

```

Function Factorial(ByVal val As Long) As Long
    Factorial = 1
    Do While (val > 1)
        Factorial = Factorial * val
        val = val - 1
    Loop
End Function

```

As with subroutines, variables and constants may be defined within a function. When this is done, the variable or constant is private to the function and cannot be directly accessed from any other routine. Although it is common to place such definitions near the beginning of the function, the definitions may occur anywhere in the subroutine as long as they occur before their first use.

Note that the normal execution sequence of a function may be altered by using the `Exit Function` statement. When executed, this statement causes control to return to the caller immediately. Also, it is important to know that the return value variable of a function is not automatically initialized. If your

function returns without having assigned a value to the return value variable, the return value will have an undefined value.

Once a function is defined, it may be used anywhere a variable may be used, e.g. in an expression. One exception is that a function name may not be used on the left hand side of an assignment.

```
Dim lval as Long  
lval = Factorial(4)
```

If a function is defined as taking zero parameters, it may be invoked by giving its name without the parentheses following it. This form is supported for compatibility reasons but its use is discouraged. If the parentheses are present a reader of the code knows immediately that it is a function invocation as opposed to the use of a variable or constant.

BasicX Compatibility Note

In BasicX compatibility mode, when a function is defined as returning an UnsignedInteger or UnsignedLong type, the very first line of the function must be a Set statement

2.4 Expressions

Expressions are an important part of most ZBasic programs. They provide the means by which your programs implement the mathematical, logical and comparison operations necessary for your application. Generally, anywhere a value may be used, an expression may be used as well. An expression consists of one or more values, called operands, and one or more operators that indicate the function to perform on the operands.

Example

```
b = b * 5 + 3
```

In this assignment statement, the value being assigned to the variable `b` is an expression comprising three operands and two operators. Some operators, like both of those in the expression above, are called “binary operators” because they require two operands. Other operators require only one operand and are called “unary operators”. The available operators and their characteristics are described in subsequent sections.

ZBasic is a strongly typed language. This means that operands supplied for binary operators must generally be of the same type. The only exception to this rule is the exponentiation operator which allows a restricted mixing of types as described in Section 2.4.3.

The order of evaluation of the components of an expression is governed by operator precedence and associativity as described in the next two sections. However, that order may be overridden by the use of parentheses.

2.4.1 Operator Precedence

Consider again the example expression `b = b * 5 + 3`. Each of the operators requires two operands but it may not be immediately clear what the operands are in each case. It could be that the expression above means to add 5 to 3 and then multiply the result by the value of `b`. On the other hand, it could mean to multiply the value of `b` by 5 and then add 3 to the result. The property that dictates the order in which operators are applied in this case is called “operator precedence”. An operator having higher precedence has priority over an operator with lower precedence and is therefore applied first. The table below depicts the precedence of ZBasic operators.

Operator Precedence

Precedence Level	Operators
11 (highest)	$^$
10	$+ \text{(unary)} - \text{(unary)}$
9	$* /$
8	\backslash
7	Mod
6	$+ -$
5	$\&$
4	$= < > <= >= <>$
3	Not
2	And
1	Or
0 (lowest)	Xor

The actual precedence level values are of no particular significance. All that matters is whether the precedence value of one operator is higher, equal to or lower than that of another operator. In the example used above, since the multiplication operator has higher precedence than the addition operator, the meaning of the expression is to multiply the value of the variable `b` by 5 and then add 3 to the result. It happens in this case that the order of application of the operators is left to right but that is not always the case. Consider this slightly different example:

```
b = b + 5 * 3
```

The meaning of this expression is to multiply 5 by 3 and then add the value of the variable `b` to the result. The relative precedence level of the operators requires that they be applied in an order that is not left to right.

BasicX Compatibility Note

In BasicX mode, a different operator precedence set is used, as shown in the table below, in order to be compatible with BasicX.

Operator Precedence	
Precedence Level	Operators
5 (highest)	$^$
4	$+ \text{(unary)} - \text{(unary)}$
3	Not
2	$* / \backslash \text{Mod} \text{And}$
1	$+ - \text{Or} \text{Xor} \&$
0 (lowest)	$= < > <= >= <>$

2.4.2 Operator Associativity

After studying the precedence table in the preceding section, the obvious question would be what happens when two operators have the same precedence?

```
b = b - 5 + 3
```

In case of equal precedence, the operators are applied in the order dictated by the associativity of the operator. Except for the exponentiation operator, all operators in ZBasic are left associative. This means that given equal precedence operators they are applied in left to right order. Exponentiation is right associative meaning that they will be applied in right to left order.

However, no matter what the precedence levels are you can force the operators to be applied in any order that you wish by utilizing parentheses. Consider the two examples below.

```
b = (b - 5) + 3
b = b - (5 + 3)
```

The parentheses indicate that the expression within should be evaluated first after which the result may be used as an operand in another operation.

2.4.3 Arithmetic Operators

The arithmetic operators are listed in the table below along with the permitted operand types. The result is the same type as the operands.

Arithmetic Operators			
Function	Type	Operator	Permitted Operand Types
Negation	Unary	-	any numeric
Addition	Binary	+	any numeric
Subtraction	Binary	-	any numeric
Multiplication	Binary	*	any numeric ¹
Division, Integer	Binary	\	any integral ¹
Division, Real	Binary	/	Single
Modulus	Binary	Mod	any numeric ¹
Exponentiation	Binary	^	any numeric ²

Notes:

¹ In BasicX mode, operating on UnsignedLong is not supported.

² The result will be the same type as the left operand. In BasicX mode, the left operand must be Single and the right operand must be either Single or Integer.

For Single operands, dividing by zero produces a special value indicating either positive or negative infinity depending on the sign of the dividend. Dividing 0.0 by 0.0 produces a special value called NaN, representing a value that is “Not a Number”. Similarly, performing the Mod operation with Single values using a divisor of zero produces a NaN. For all other operand types the result of using a zero divisor for either division or Mod is undefined. The System Library function SngClass() returns a value indicating the general classification of a value of type Single. See the description of SngClass() in the System Library Reference Manual for more information.

2.4.4 Logical Operators

The logical operators are listed in the table below along with the permitted operand types. The result is the same type as the operand(s).

Logical Operators			
Function	Type	Operator	Permitted Operand Types
Logical AND (conjunction)	Binary	And	Boolean
Bitwise AND	Binary	And	any integral ¹
Logical OR (disjunction)	Binary	Or	Boolean
Bitwise OR	Binary	Or	any integral ¹
Logical XOR (exclusive disjunction)	Binary	Xor	Boolean
Bitwise XOR	Binary	Xor	any integral ¹
Logical Complement	Unary	Not	Boolean
Bitwise Complement	Unary	Not	any integral ¹

Notes:

¹ If Option Strict is enabled, signed types are not allowed.

2.4.5 Comparison Operators

The comparison operators are listed in the table below along with the permitted operand types. The result type is Boolean.

Comparison Operators			
Function	Type	Operator	Permitted Operand Types
Equality	Binary	=	any
Inequality	Binary	<>	any
Greater Than	Binary	>	String or any numeric
Greater Than or Equal To	Binary	>=	String or any numeric
Less Than	Binary	<	String or any numeric
Less Than or Equal To	Binary	<=	String or any numeric

2.4.6 Miscellaneous Operators

The remaining operators to be described are the string concatenation operator (&) and the address operator (@). For the string contenation operator, both operands must be type `String` and the result will be type `String`. Note that the + operator may also be used for concatenating strings. The sole difference between using & and using + is that the former supports automatic value-to-string conversion while the latter does not.

The address operator is used to obtain the address of a variable, array element or procedure; it is an alternate means of referring to the `DataAddress` and `CodeAddress` properties. In use, the address operator is placed immediately before the variable or procedure name as illustrated below.

```
Dim myVar as Byte
Dim addr as UnsignedInteger
addr = @myVar ' same as myVar.DataAddress
```

2.4.7 No "Short Circuit" Evaluation

It is important to note that in ZBasic, as in most Basic dialects, every term in an expression is always evaluated irrespective of the intermediate results. This is a technical detail that is significant only when an expression contains function invocations and the act of invoking one or more of the functions involved has “side effects” like modifying a global variable, modifying a parameter passed by reference, or changing the state of the hardware. Consider the evaluation of the conditional expression in the `If` statement below when the value of the variable `a` is, say, 10.

```
If (a > 3) Or (foo() > 10) Then
    [other statements]
End If
```

When the expression on the left side of the `Or` operator is evaluated the result will be True. Because of this fact we know that the resulting value of the entire conditional expression will also be True – nothing on the right hand side can possibly affect the outcome. Nonetheless, the expression on the right hand side of the `Or` operation will still be evaluated and thus the function `foo()` will be invoked. Some other computer languages, notably C/C++ and Java, implement the concept of “short circuit evaluation”. In those languages, the evaluation of an expression stops as soon as the result is known. If that were the case here, the right hand side of the `Or` expression would not be evaluated and, hence, the function `foo()` would not be invoked. To reiterate, ZBasic does not implement short circuit evaluation.

2.5 Statements

Within a subroutine or function you can define variables and use statements to implement the logic required for the functionality of the routine. This section describes the types of statements available. ZBasic statements may be divided into two general categories: simple and compound. An example of a simple statement is the assignment statement where the entire statement is expressed on one line (ignoring possible line continuations). In contrast, a compound statement comprises two or more lines and may contain other statements within it. In many respects, it is convenient to think of a compound statement as if it were a single statement even though it may have many constituent statements.

2.5.1 Assignment Statement

The assignment statement is perhaps the most basic and most often used statement in a program. The syntax of an assignment statement is shown below in two forms, one for assigning a value to a scalar variable and one for assigning a value to an array element.

```
<var-name> = <value>
<var-name>( <index-list> ) = <value>
```

In both cases, the *<value>* element may be a simple value or a complex expression involving one or more other variables, constants, functions, etc. However, the type of *<value>* must match the type of the variable or array element to which it is being assigned. There is no automatic type conversion.

When assigning a value to an array element, you must specify the index or indices of the particular element of interest. For a single-dimensional array, you will specify a single value for the index of the desired element. For multi-dimensional arrays, you must specify a value for each of the indices separated from one another by a comma.

Example

```
Dim i as Integer
Dim ia(1 to 5) as Integer
Dim board(1 to 12, 1 to 12) as Byte

i = 2
ia(i) = (ia(3) + 2) / 4
board(i, 6) = CByte(ia(2))
```

2.5.2 Call Statement

The Call statement is used to invoke a subroutine. The syntax is:

```
Call <subroutine-name> ( [<parameter-list>] )
```

The optional *<parameter-list>* must contain the proper number of parameters each of the correct type for the subroutine being invoked. If more than one parameter is given each parameter must be separated from the next by a comma.

When this statement is executed the supplied parameters, if any, are pushed on the stack and control is transferred to the first statement of the subroutine. When the subroutine finishes executing control resumes with statement following the Call statement.

Although not recommended, for compatibility with other Basic dialects it is permissible to omit the Call keyword. If this is done the parentheses surrounding the parameter list must also be omitted. Note, particularly, the third example below that seems to violate this rule. However, it does not because a parenthesized expression is, in fact, an expression.

Examples

```
Call PutPin(12, 0)
PutPin 12, 0
Delay (1.0)
```

2.5.3 CallTask Statement

The CallTask statement is used to start a task. See Section 3.5 for more information on using tasks. The basic syntax to invoke a task is:

```
CallTask <task-name>, <task-stack>
```

In this case, the *<task-name>* element must be the name of a user-defined subroutine (usually one that takes no parameters). The *<task-stack>* must be the name of a Byte array that will serve as the stack for the task. For compatibility with BasicX, the *<task-name>* may be enclosed in quote marks.

Example

```
Dim ts1(1 to 40) as Byte
CallTask task1, ts1
```

A task may also be passed parameters when it is invoked. The syntax for doing so is similar to that for invoking a subroutine that requires parameters.

```
CallTask <task-name>(<parameter-list>), <task-stack>
```

See the discussion of the CallTask statement in the ZBasic System Library Reference manual for more details on the allowed parameter types. This syntax is not supported in BasicX compatibility mode.

Example

```
Dim ts1(1 to 40) as Byte
CallTask task1(&H100), ts1
```

For advanced users, the task stack may also be specified by giving its address explicitly. This topic is discussed in section 3.5.1, Advanced Multi-tasking Options.

2.5.4 Console.WriteLine Statements

These statements (with a syntax more akin to object-oriented methods) are similar to Debug.Print but they are limited to displaying one string at a time. They are supported for compatibility with Visual Basic. The syntax of the statements is:

```
Console.WriteLine( <string-expression> )
Console.WriteLine( <string-expression> )
```

The difference between these two statements is that the latter also outputs a carriage return/line feed following the string while the former does not.

Examples

```
Dim i as Integer
Dim s as String

Console.WriteLine(CStr(i))
Console.WriteLine("i = " & CStr(i))
Console.Write("s = ")
Console.WriteLine(s)
Console.WriteLine(" ")
```

This sequence of statements is written to produce exactly the same result as the sequence of Debug.Print statements above. Note how the string concatenation operation is used in the second Console.WriteLine() invocation to produce a single string. In the final example, an empty string is output followed by a carriage return/line feed.

Note that there are counterparts to these statements called Console.Read and Console.ReadLine that allow you to retrieve data from the Com1 serial port. However, these are technically functions (since they both return a value) and are therefore not described here. See the ZBasic System Library Reference manual for information on these functions.

2.5.5 Debug.Print Statement

One technique for debugging a program is called “print statement debugging”. The idea is that to determine how your program is executing you insert statements into the program that display the values of important variables or simply display a distinctive message so that you know what the program is doing. Debug.Print is a special statement intended just for this purpose. The syntax is:

```
Debug.Print [<string-list>][ ; ]
```

The *<string-list>* element represents zero or more string expressions separated from one another by a semicolon. Each of the string expressions is evaluated in turn, from left to right, and the string result of each is output to the console channel. If the optional trailing semicolon is omitted, a carriage return/line feed will also be output so that the next time something is output to the console channel it will appear on a new line. If you don't want the subsequent output to be on a new line, simply add the semicolon at the end of the list. This is often done when you need to compute several different values to output. You can use a separate Debug.Print statement for each value and keep them all on the same output line by ending all but the last with a semicolon. It is permitted, syntactically, to specify an empty string list but still include the trailing semicolon. However, this construction does nothing.

Note that each of the items to be displayed must be a string. You can use the CStr() function to produce a string from any value.

Examples

```
Dim i as Integer
Dim s as String

Debug.Print CStr(i)
Debug.Print "i = ";CStr(i)
Debug.Print "s = ";
Debug.Print s
Debug.Print
```

The last example, with an empty *<string-list>* and no trailing semicolon, will simply send a carriage-return/linefeed to the console channel.

The unusual form of `Debug.Print` is due to its heritage from Visual Basic. It should probably be called the `Print` method of the system `Debug` object but it even departs from the traditional syntax of the methods that are part of object-oriented languages. Nonetheless, it is included for compatibility as well as its utility.

2.5.6 Do-Loop Statement and Variants

This compound statement, briefly mentioned earlier in this document, is the basic repetition construct in ZBasic. The syntax is:

```
Do  
    [<statements>]  
Loop
```

This construct causes the sequence of zero or more statements to be repeatedly executed. However, execution of the loop may be terminated using an `Exit Do` statement at which point control will transfer to the first statement following the `Loop` statement. Note that the `Do-Loop` compound statement may be nested and the `Exit Do` only terminates the innermost `Do-Loop` that contains it.

Example

```
Do  
    <other-statements>  
    Do  
        <other-statements>  
        If (i > 5) Then  
            <other-statements>  
            Exit Do  
        End If  
    Loop  
    <other-statements>  
Loop
```

Here, the `Exit Do` only terminates the inner `Do-Loop`; the outer one continues to iterate.

There are four other variations on this basic looping concept, all of which involve a condition for continuing the iteration. The syntax of the four variations is as follows:

```
Do While <boolean-expression>  
    [<statements>]  
Loop  
  
Do Until <boolean-expression>  
    [<statements>]  
Loop  
  
Do  
    [<statements>]  
Loop While <boolean-expression>  
  
Do  
    [<statements>]  
Loop Until <boolean-expression>
```

As you can see, the difference between the four variations is whether the test is at the top of the loop or at the bottom of the loop and, secondly, the logic sense of the condition. Testing the condition at the top of the loop means that the statements within the loop may be executed zero or more times. Testing the condition at the bottom of the loop means that the statements will always be executed at least once.

The difference between using `While` and `Until` is nothing more than a logic inversion. The construct `Do While Not <boolean-expression>` is logically equivalent to `Do Until <boolean-expression>`.

There is no fixed limit on how deeply Do loops may be nested. The actual limit is governed by how much memory is available to the compiler. For all practical purposes, there is no limit.

BasicX Compatibility Note

In BasicX mode, the nesting of Do loops is limited to 10 for compatibility.

2.5.7 Exit Statement

The Exit statement allows you to terminate the execution of a loop, a function or a subroutine earlier than it otherwise would. This is most commonly used when a condition is detected by the code that prevents further normal processing. The syntax for the Exit statement is:

```
Exit <exit-type>
```

The *<exit-type>* element may be Do, For, Sub or Function but the use of each is restricted to use within a Do-Loop, For-Next, subroutine and function, respectively. Any other use will result in a compiler error.

Example

```
Do
    <other-statements>
    If (i > 5) Then
        Exit Do
    End If
    <other-statements>
Loop
```

When an `Exit Do` is executed within nested Do-Loop statements only the innermost Do-Loop that contains the exit statement will be terminated. Control will be transferred to the first statement following the terminated Do-Loop. The same idea applies to an `Exit For` within nested For-Next statements.

2.5.8 For-Next Statement

The For-Next compound statement is another form of looping that provides controlled iteration using a loop index variable. The syntax for a For-Next loop is:

```
For <var> = <start-expr> To <end-expr> [ Step <step-expr> ]
    [<statements>]
Next [<var>]
```

The *<var>* element, referred to as the loop index variable, must refer to a previously defined scalar variable (i.e. not an array element) that is either a numeric type or an enumeration type. The *<start-expr>*, *<end-expr>* and optional *<step-expr>* elements must all be the same type as the loop index variable.

Example

```
Dim i as Integer
For i = 1 To 10
    Debug.Print CStr(i)
Next i
```

When the For statement begins execution, the *<start-expr>* is evaluated and the resulting value is assigned to the loop index variable. Then, before executing any statements contained within the body of the For-Next statement, the *<end-expr>* is evaluated and the value of loop index variable is compared to that value. If the value of the loop index variable is less than or equal to the value of the *<end-expr>* the statements within the body of the For-Next are executed. This is called the “loop entry test” because it controls whether or not the loop statements are executed.

At the bottom of the loop, marked by the Next statement, the loop index variable is modified in preparation for the next iteration of the For-Next loop. If the optional Step *<step-expr>* is present its value is added to the loop index variable, otherwise the loop index variable is simply augmented by 1. Then control is transferred back to the top of the loop where the loop entry test is performed again.

The logic of the For loop in the example above may be exactly duplicated using other statements as illustrated below. You can see that the For loop allows you to express the same logic more concisely.

```
Dim i as Integer  
  
i = 1  
Do While (i <= 10)  
    Debug.Print CStr(i)  
    i = i + 1  
Loop
```

There are several things to note about the For-Next loop. Firstly, it is important to be aware that although the values of the *<end-expr>* and *<step-expr>* elements are used multiple times, the expressions are only evaluated once, when the execution of the For-Next begins. This distinction is only important, of course, if these expressions contain references to other variables that might be modified during loop execution or if they involve function calls. Secondly, because the loop entry test checks to see if the loop index variable is less than or equal to the *<end-expr>* you must choose a data type for the loop index variable that is capable of representing a value that is greater than the value of the *<end-expr>*.

Otherwise, the loop entry test will always be true. Consider this errant example:

```
Dim i as Byte  
  
For i = 0 to 255  
    <statements>  
Next i
```

This loop will never terminate because the value of *i* is always less than or equal to 255 since it is a Byte type. Where possible, the compiler will issue a warning if it detects these problematic conditions. Even so, you should develop the habit of considering this potential problem every time you code a For-Next loop.

The third important aspect of the For-Next statement is that if the *<step-expr>* evaluates to a negative value, the sense of the loop entry test changes. In this case, the loop index variable is tested to see if it is greater than or equal to the *<end-expr>* and, if so, the statements of the loop are executed. Also, in this case the caveat noted above about the range of the loop index variable changes. The loop index variable must be capable of representing a value that is less than the value of the *<end-expr>*. When a For loop is used with an unsigned data type, the step value is considered to be negative if the most significant bit of the value is a 1.

One other note: it is permissible for the *<step-expr>* to evaluate to zero. This will cause the For loop to execute indefinitely. The For-Next loop may be terminated at any time by using the Exit For statement.

The presence of the *<var>* on the Next statement is optional. However, if it is present, it must match the name of the loop index variable of the For loop with which it is associated. There is no fixed limit on how deeply For-Next loops may be nested. The actual limit is governed by how much memory is available to the compiler. For all practical purposes, there is no limit.

BasicX Compatibility Note

In BasicX mode, the For-Next statement is much more restrictive. The loop index variable must be a scalar integral type and must also be a local variable. Referring to the loop index variable in a Next statement is not supported. The *<step-expr>* is restricted to a constant expression that evaluates at compile time to either 1 or -1. Lastly, For-Next loops may be nested to a maximum depth of 10 for compatibility.

When Option Strict is enabled, there are additional restrictions that apply. Firstly, the loop index variable must be local to the routine; it cannot be defined at the module level. Secondly, the loop index variable is not allowed to be used or modified outside of the For-Next loop except that it can be used as the loop index variable in a subsequent For-Next loop. Thirdly, inside the For-Next loop the loop index variable is read-only. Any attempt to modify the loop index variable, or pass it by reference to another routine will result in an error message from the compiler.

One final note: although loop index variables of type Single are allowed, some experienced programmers advise against doing so. This is due to the fact that not all real numbers can be exactly represented as a Single value. Consequently, using a Single loop index variable may not produce the expected results. It is often better to use an integral loop index variable along with an auxiliary real variable to accomplish the desired objective.

2.5.9 Goto Statement

The Goto statement allows you to transfer control to a specific point in the sequence of statements that comprise a subroutine or function. The point to which control is transferred is marked by a label statement. The label statement is simply an identifier followed by a colon appearing on a line by itself (except that it may be followed by a comment).

Because it interferes with the normal program flow, the Goto statement can be overused resulting in a program that is difficult to understand and, therefore, difficult to maintain. Some programmers believe that a Goto statement should never be used. Others believe that it is acceptable to use a Goto but only if the alternative code structure is even less palatable. The latter strategy is probably the best to adopt.

Example

```
Goto doOtherStuff
    <other-statements>
doOtherStuff:
    <other-statements>
```

2.5.10 If-Then-Else Statement

The if-then-else compound statement is the basic decision making construct in ZBasic. In its simplest form, the syntax is:

```
If <boolean-expression> Then
    <statements>
End If
```

The *<boolean-expression>* element is an expression whose value is of type Boolean. It most often involves one of the conditional operators that allow you to compare the values of two expressions but it may also simply be the invocation of a function whose return type is Boolean.

The *<statements>* element represents zero or more ZBasic statements possibly including other If statements. This allows you to create nested decision-making statements of arbitrary complexity.

Although the construction described above is useful, it is often the case that you want your program to execute a certain set of statements if a condition is true but you want it to execute a different set of statements if the condition is false. The If statement allows this logic using the syntax shown below.

```
If <boolean-expression> Then  
    <statements>  
Else  
    <statements>  
End If
```

Sometimes you'll want to test several different conditions and execute a different set of statements in each case. The If statement allows this logic using the following syntax:

```
If <boolean-expression> Then  
    <statements>  
ElseIf <boolean-expression> Then  
    <statements>  
Else  
    <statements>  
End If
```

The ElseIf portion of the statement, including the associated statements, may occur zero or more times. The Else portion of the statement, along with its associated statements, may occur zero or one times. Note that the logic of the If-Then-Else statement is designed so that at most one set of statements gets executed. This construct represents a series of tests that are performed sequentially. The first such test that produces a Boolean True result will cause the statements associated with that test to be executed. If none of the tests produce a True result and an Else clause exists, the statements associated with the Else will be executed. In all cases, after the set of statements is executed, control transfers to the first statement following the End If.

If the same expression is being repeatedly tested against different values, it is more efficient to use the Select-Case statement described in Section 0.

Examples

```
If (i > 3) Then  
    Call PutPin(12, zxOutputLow)  
Else  
    j = 55  
    Call PutPin(12, zxOutputHigh)  
End If  
  
If i > 3 Then  
    Call PutPin(12, zxOutputLow)  
ElseIf (i > 0) Then  
    j = 0  
Else  
    j = 55  
    Call PutPin(12, zxOutputHigh)  
End If
```

Note that the conditional expression is not required to be enclosed in parentheses. Many programmers are accustomed to other languages where they are required and therefore do so out of habit. Others believe that the parentheses improve the readability and use them for that reason. You're free to adopt whichever practice suits you.

One other comment on style is in regard to indentation. The examples used in this document indent the statements within compound statements like If-Then-Else in order to improve readability. The compiler ignores spaces and tabs except to the extent that they separate identifiers, keywords, etc. You're free to adopt any indentation style that you deem appropriate.

There is no fixed limit on how deeply If-Then statements may be nested. The actual limit is governed by how much memory is available to the compiler. For all practical purposes, there is no limit.

2.5.11 Single-line If-Then Statement

Sometimes, it is convenient to express conditional logic concisely using a single-line If statement. The form is similar to the multi-line form except that the Else If clause is not supported and there is no End If. The syntax of a single-line If statement is:

```
If <boolean-expression> Then <statement> [ Else <statement> ]
```

The *<boolean-expression>* element is the same as described in the previous section. The *<statement>* element represents one ZBasic statement such as an assignment statement or a Call statement. Note, that it is permissible for *<statement>* to be multiple statements each separated from the next by a colon. Moreover, the line continuation character may be used to distribute the *<statement>* over multiple lines.

Examples

```
If (a > b) Then a = b  
If (a > b) Then a = b Else b = 10  
If flag Then flag = False : Call MySub(25)
```

BasicX Compatibility Note

The single-line If statement is not supported in BasicX compatibility mode.

2.5.12 Select-Case Statement

The Select-Case compound statement is a multi-way branch statement that can be used in place of an If-Then-ElseIf chain in certain situations. The syntax is shown below.

```
Select Case <test-expr>  
Case <case-expr-list>  
    [<statements>]  
    ...  
Case Else  
    [<statements>]  
End Select
```

The *<test-expr>* element, known as the selection expression, gives a value that will be tested against the value(s) given in zero or more standard case clauses. Each standard case clause begins with the word Case and is followed by a list of one or more expressions, each of which must evaluate to the same type as *<test-expr>*. If multiple expressions are given, they must be separated from one another by a comma. The remainder of the case clause consists of zero or more ZBasic statements. The type of the selection expression may be Boolean, an enumeration, any numeric type or String. The practical value of using a Boolean type is somewhat limited, however – it's simpler to just use an If-Then statement.

There may be at most one default case clause introduced by the keywords Case Else. The remainder of the default case clause consists of zero or more ZBasic statements. If the default case clause is present, it must be the final case clause.

The Select-Case statement executes by first evaluating the *<test-expr>*. Then the resulting value is compared with the value of each of the expressions in the *<case-expr-list>* of the first standard case clause, if present. The evaluation of the case expressions and the comparison with the test value is done in order, left to right. As soon a case expression is found whose value is equal to the test value, the statements associated with that case clause are executed and then control transfers to the first statement following the End Select. If none of the expressions in the first case clause match the *<test-expr>* value, the process is repeated with the second standard case clause and so on until all of the standard case clauses have been tested. When all of the standard case clauses have been tested without finding a matching expression value, if a default case clause exists the statements associated with it are executed.

There are two special forms of case expressions that may be used in the *<case-expr-list>* of a standard case clause. The first special form is the range expression. This takes the form of two expressions separated by the keyword To. Both expressions must evaluate to the same type as *<test-expr>*. The *<test-expr>* value will be deemed to select the case clause if the value is greater than or equal to the value of the expression to the left of the To keyword and less than or equal to the value of the expression to the right of the To keyword. Effectively, the range expression specifies an inclusive range.

The second special form may be used to implement special test conditions. It has the syntax:

Is <conditional-operator> <expression>

The *<conditional-operator>* element may be any one of the six conditional operators: =, <>, <, <=, >, and >=. The *<expression>* element must be an expression that evaluates to the same type as the *<test-expr>*. Note that the construction Is = <expr> yields the same result as simply specifying the expression value alone.

There is no fixed limit on how deeply Select-Case statements may be nested. The actual limit is governed by how much memory is available to the compiler. For all practical purposes, there is no limit.

Example

```
Select Case i * j
Case 3
    j = 5
    Call PutPin(12, zxOutputLow)

Case 4, 5 To 20, 27
    j = 1
    Call PutPin(13, zxOutputLow)

Case 3, 100, Is > 200, j
    j = 0

Case Else
    j = -1
End Select
```

In the example above, if the selection expression evaluates to 3 the statements of the first case clause will be executed. The fact that the third case clause also has a case value of 3 is of no consequence. Also note that the case expressions are evaluated every time they are tested. This fact must be kept in mind for two reasons. Firstly, if the case expression contains a variable whose value changes between successive executions of the Select-Case statement (a situation that is strongly discouraged), the case clause that is selected may change even if the selection expression value does not change. Secondly, if any of the case clause expressions involves a function call, the function may or may not be invoked depending on the value of the selection expression and the values of the various expressions in the case clauses preceding it.

BasicX Compatibility Note

In BasicX mode, the use of String and Single types is not supported nor is the construction Is <op> <expr>. Moreover, there must be at least one standard case clause.

2.5.13 Set Statement

This statement is only allowed as the first statement of a function that returns an UnsignedInteger or UnsignedLong type. It must precede all other statements and variable definitions. It is supported for compatibility with BasicX and is required in BasicX compatibility mode but it is otherwise ignored. The syntax is shown below.

```
Set <function-name> = New <type>
```

The <function-name> element must match the name of the function containing the Set statement and the <type> must match the function's type.

Example

```
Function myFunc() as UnsignedInteger
    Set myFunc = New UnsignedInteger

    Dim I as Integer

    <other-statements>
End Function
```

2.5.14 While-Wend Statement

For compatibility with other dialects of Basic, ZBasic includes support for an alternative to the Do While – Loop construct. The syntax is:

```
While <boolean-expression>
    [<statements>]
Wend
```

Note that this compound statement is logically equivalent to the Do While variation of the Do-Loop statement. The one difference is that Exit Do cannot be used to terminate a While-Wend statement.

BasicX Compatibility Note

In BasicX mode, the While-Wend statement is not supported.

2.5.15 With Statement

The With statement allows you to use a shorthand notation to refer to some objects. The syntax for the With statement is:

```
With <prefix>
    <other-statements>
End With
```

Between, the With and End With statements, any reference to an identifier that begins with a period will be treated as if it had the series of characters identified by <prefix> immediately preceding the period.

Example

```
tick = Register.RTCTick ' the long way
With Register
    <other-statements>
    tick = .RTCTick ' the short way, implies Register.RTCTick
    <other-statements>
End With
```

Note that the entire construct, from `With` to `End With`, is treated much like a compound statement in that it cannot be split across other statement boundaries. It is important to note, however, that this is not a true compound statement with block scoping. Variables and constants defined within a `With` block are visible to statements that follow it.

In addition to `Register`, other useful `<prefix>` designations are `Console`, `Debug`, `Option`, `Version` and `Module`. Also, the `<prefix>` may specify a portion of a structure member reference, allowing shorthand access to structure members. See Section 3.25 for more information on using structures.

BasicX Compatibility Note

In BasicX compatibility mode, only `With Register` is supported.

2.6 Literals

Boolean, numeric and string constant values are often used in programming. These are called literals because they represent the literal value of the number or string that you have in mind as opposed to a variable whose value may change over time.

2.6.1 Boolean Literals

Boolean literals are the keywords `true` and `false`, in upper, lower or mixed case. These literals are of type `Boolean`.

2.6.2 Numeric Literals - Integral Values

A decimal integral numeric literal consists of decimal digits optionally with a leading plus or minus sign to indicate a positive or negative literal value.

Integral literals may be specified in hexadecimal (base 16) by beginning the literal with an ampersand and the letter H (upper or lower case) followed by one or more hexadecimal digits (0-9, A-F, a-f). In BasicX compatibility mode, a trailing ampersand is either allowed, required or disallowed depending on the specific value and whether or not Strict mode is enabled. In native ZBasic mode the trailing ampersand is always allowed, is never required and has no effect whatsoever on the resulting value.

Integral literals may also be specified in binary (base 2) by beginning the literal with an ampersand and the letter B (upper or lower case) followed by one or more binary digits (0-1). To improve readability, you may also include one or more underscores within the digit string provided that each underscore occurs between two digits. For compatibility with other Basic dialects, an upper or lower case X may be used in place of the radix indicator B. Binary literals are not supported in BasicX compatibility mode.

Examples of decimal, hexadecimal and binary integral literals:

```
124
+16
-357
```

```
&HabCd  
&H8000&  
&B0010_0011
```

2.6.3 Numeric Literals - Real Values

A real numeric literal consists of one or more decimal digits, optionally with a leading plus or minus sign, followed by either a decimal point and one or more decimal digits or the letter E followed by one or more decimal digits also optionally prefixed by a plus or minus sign. If the fractional part is present, it may also be followed by an exponent specification. For compatibility with other Basic dialects, a literal that would otherwise be an integral literal will be interpreted as a real literal if it is immediately followed by an exclamation mark. Real literals have the type `Single`.

Examples of real literals:

```
3.14159  
+6.02e23  
3e10  
-300!
```

BasicX Compatibility Note

In BasicX compatibility mode, real literals must either contain a decimal point or have a type designation suffix ! or # to force them to be recognized as type `Single`. A literal like 12e2 represents the integer value 1200.

2.6.4 String Literals

A string literal consists of zero or more characters enclosed in quotation marks. Note that a string may not be continued on the next line by ending the first line with an underscore. However, you may use the concatenation operator in conjunction with the underscore continuation to span line boundaries. The compiler will combine the operands to the concatenation operator as long as they are both string constants. String literals have the type `String`.

Examples of String Literals:

```
" "  
"Hello, world!"  
"The quick brown fox" & _  
" jumped over the lazy dog."  
"Hello, ""Joe""!"
```

The third and fourth lines above show how to use the concatenation operator and line continuation to construct longer strings. Note that the underscore must be the last character on the line and that there must be a space or tab character preceding it. The last example shows how to include a quotation mark within a string literal. Two consecutive quote marks are reduced to one in the actual string. If you want two adjacent quote marks in the string, you'll have to double each of them.

2.6.5 Built-in Binary Constants

Although they are technically not numeric literals, ZBasic provides some built-in `Byte` constants that serve the same purpose. The constants begin with the letters `BX` and are followed by exactly 8 binary digits (0-1). There may be an underscore between any pair of binary digits to enhance readability. These constants are of type `Byte` and may only be used where a `Byte` type is allowed. These built-in constants are supported for compatibility with BasicX. It is recommended that new applications use binary literals (described in Section 2.6.2) since they are more generally useful.

Examples

```
Bx0100_1101  
BX01_00_11_01
```

2.7 Comments

Comments may be placed on a line by themselves or at the end of a line containing other program text. A comment begins with an apostrophe and continues to the end of the line. A comment may be continued on the next line in the manner described in Section 2.8.

Note: The BasicX compiler does not allow comments to be continued but Visual Basic does. ZBasic allows comment continuation both in native mode and in BasicX compatibility mode.

Examples

```
'This is a comment.  
a = 23      ' this is a comment, too  
b = 55      ' and because this comment ends with an underscore _  
it continues on the next line
```

2.8 Line Continuation and Multiple Statements Per Line

A statement may be continued across multiple lines by ending each line except the last with an underscore preceded by at least one space or tab character. Except for spaces, tabs and/or a comment, no characters other than end-of-line characters may follow the underscore for it to be considered a line continuation character. The maximum aggregate size of a line, whether continued across multiple lines or not, is 1000 characters.

In the example below the beginning of the If statement is continued to the following line. This is often useful to help make more complex expressions more readable.

Example

```
If (GetPin(20) = 1) And _  
    (GetPin(12) = 0) Then  
    Call PutPin(5, 0)  
End If
```

While the line continuation capability allows you to create statements that span multiple lines, it is sometimes convenient to place multiple statements on one line. In ZBasic, as in many other Basic dialects, you may accomplish this by using a colon to separate each pair of statements on the line.

Example

```
Dim i as Integer  
Dim j as Integer, k as Byte  
  
i = 0 : j = 1 : k = 2
```

2.9 Persistent Variables

You may define variables that are stored in the processor's internal EEPROM, referred to in this document as Persistent Memory. It is called persistent because the values that you store there are retained even if the system is powered down or reset. This characteristic makes persistent variables useful for storing configuration information for your application and other similar information that your application needs to be preserved.

A persistent variable is defined at the module level using the syntax:

```
{Public | Private | Dim} <name> as Persistent <type>
```

Using the keyword `Dim` has the same effect as using `Private`. Within a subroutine or function, a persistent variable is defined using the syntax.

```
Dim <name> as Persistent <type>
```

In both cases, the `<type>` element may be any numeric type (e.g. Byte, Integral, Single, etc.), Boolean or a user-defined type (structure or enumeration). A persistent string must be defined using the bounded string syntax (Section 2.11.1), i.e.

```
Dim <name> as Persistent BoundedString(<size-expr>)
```

The `<size-expr>` element must be a constant integral expression that specifies the number of bytes to reserve for the persistent string's characters.

Examples

```
Dim kbdAttached as Persistent Boolean  
Private signOnMsg as Persistent BoundedString(25)
```

It is important to note that the implementation of the `PersistentString` type is identical to that of the `BoundedString` type and is therefore not protected from overwriting the boundaries of the data item. To protect against overwriting, it is advisable to explicitly limit the size of the string to be written.

Arrays of persistent variables may be defined as well. To do so, simply add the array dimension list to the variable name in the same manner as for regular variables. If no lower bound is specified, the default array base applies.

Example

```
Dim freq(1 to 10) as PersistentInteger
```

It is important to note that persistent variables are not initialized by the system. They have values based on whatever data happens to be at the Persistent Memory address to which they are assigned. The compiler assigns Persistent Memory addresses in the order that modules are compiled and, within modules, in the order the variables are defined.

To avoid problems of unexpected address order changes, it is highly recommended that all persistent variables be defined in a single module. Also, it is recommended when you add more persistent variables to an existing application that you add them following the definitions of the previously existing persistent variables. Deleting persistent variables or inserting new persistent variables in the midst of existing ones may cause problems because it will change the address to which subsequent variables are assigned.

You'll probably want to build into your application a way to initialize all of your persistent variables to a known state. This initialization only needs to be done once, when the application is first installed (the `FirstTime()` function may be useful for this purpose). It may also be useful, however, to be able to do this at other times as well. Another useful technique is to include a persistent variable whose only purpose is to indicate that the persistent variables have been properly initialized. For this to work, you would need to choose a value that is unlikely to otherwise occur. It may even be advisable to place such

a “sentinel” variable at both the beginning and the end of the group of persistent variables, decreasing the likelihood of false positive or false negative indications.

Persistent variables have an associated property named `DataAddress`. The value of this property is the address of the data item in Persistent Memory. The type of the property is `UnsignedInteger` for compatibility with the `PersistentPeek()` subroutine.

Example

```
Dim freq(1 to 10) as PersistentInteger  
Dim addr as UnsignedInteger  
  
addr = freq.DataAddress
```

It is possible, also, to use the `DataAddress` property to get the address of an element of a persistent array. To accomplish this, simply add parentheses following the property name and specify the index or indices of the item of interest. The example below will result in `addr` having the Persistent Memory address of the fourth data value of the `freq()` array.

```
addr = freq(4).DataAddress
```

Persistent structures may also be defined, see Section 3.25 for more details. A Persistent variable may also be defined using `Based` keyword, see Section 3.21 for more details.

For compatibility with BasicX, an alternate syntax is also supported for defining persistent variables as shown below.

```
{Public | Private | Dim} <name> as [New] <persistent-type>
```

Caution: although Persistent Memory data items can be modified, the memory in which they are stored has a write cycle limit of approximately a million writes. Writing to a particular address in excess of this limit may cause the memory to become unreliable. Also, writing to Persistent Memory is much slower than writing to RAM-based variables.

The keyword `New` is optional except in BasicX compatibility mode when it is required. The `<persistent-type>` element specifies the type of persistent variable being defined and may be one of the following special types:

<code>PersistentBoolean</code>	<code>PersistentByte</code>
<code>PersistentInteger</code>	<code>PersistentLong</code>
<code>PersistentSingle</code>	

BasicX Compatibility Note

In BasicX mode, all persistent variables must be defined at the module level and neither arrays nor structures of persistent variables are supported. Also, the `DataAddress` property cannot be used to determine the address of a persistent variable.

2.10 Program Memory Data Items

It is often useful to have available initialized arrays of data that are stored in Program Memory. This is advantageous for two reasons. Firstly, you don't need run-time code to initialize the arrays and, secondly, you don't use up the more scarce RAM space for data that seldom, if ever, changes.

You define an initialized Program Memory data item in a manner similar to the way that you define RAM-based data items. There are two differences, however. Firstly, you need to specify a data source from which the array is initialized. Secondly, you don't explicitly specify the array dimensions. Instead, the

dimensions are deduced from the content of the initialization data and the lower bound of each index is always 1. The definition syntax for a Program Memory data item is:

```
{Public | Private | Dim} <name> as [New] <progmem-type>(<init-data>)
```

For Program Memory data items defined within a subroutine or function, the Public and Private keywords are disallowed because they would serve no useful purpose.

The supported *<progmem-type>* items are:

One-dimensional types (vector types):

ByteVectorData	ByteVectorDataRW
IntegerVectorData	IntegerVectorDataRW
LongVectorData	LongVectorDataRW
SingleVectorData	SingleVectorDataRW
StringVectorData	

Two-dimensional types (table types):

ByteTableData	ByteTableDataRW
IntegerTableData	IntegerTableDataRW
LongTableData	LongTableDataRW
SingleTableData	SingleTableDataRW
StringTableData	

The types ending with *RW* may be both read and written while the remaining types can read but they cannot have values assigned to them. (Of course, you can still modify the read-only types by using the System Library routine *PutProgMem()* but this is generally only used for special circumstances.)

The initialization data may be provided in two ways. The first way is to provide a file name, enclosed in quote marks, as the *<init-data>* element in the syntax description above. If the filename is not specified using an absolute path (i.e. beginning with the root directory and/or a drive letter), the path prefix (if any) of the current module is appended to the front of the filename. Note, however, that if an include path is specified on the command line, a filename that is specified with a relative path will, instead, be sought in among the directories specified in the include path list. See Section 10.2 for more information on the include path option.

The content of the file should be a textual representation of the initialization data. For integral types, the data values may be expressed in decimal, in hexadecimal (using a &H or &h prefix), or in binary (using a &B or &b prefix). In the latter case, an underscore may exist between any pair of digits. For real types, the data values may be expressed in integral, decimal or scientific notation format. See the *Values()* System Library routine for a description and examples of the acceptable formats of values. For string types, the data values should be zero or more characters enclosed in quotation marks. To include a quotation mark in the string it must appear twice in succession.

For the one-dimensional types, one or more values may be specified per line. When multiple values are given per line they must be separated by a comma and/or white space (space or tab characters). The first value on a line may be preceded by white space. Following the last value on a line, there may be a comma and/or white space and/or a comment (introduced by an apostrophe). The number of elements in the vector will be exactly the number of properly formatted data values in the file.

For the two-dimensional types, the values for each row of the table must be placed on a separate line. The column values on each line must be separated by a comma and/or white space. A comment may follow the last column value on a line. The number of valid column values must be the same for each row.

For either type, the initialization data file may contain blank lines and lines containing only a comment optionally preceded by white space. For the Byte types, data values may also be specified using a quoted string. In this case, the ASCII value of each character of the string is used as a data value. As usual, a quote may be included in the string by using two quotes in succession. For the string types, values may be specified by concatenating strings and/or byte values by separating each pair of components with a plus sign. The examples below include samples of each of these special cases.

The second method to provide initialization data is to use an in-line initializer that consists of a pair of curly braces bracketing the initialization data itself. The form of the in-line initializer data is essentially the same as the content of the initialization file described above but appearing between curly braces directly in your source file.

Here is an example of the content of an initialization file for a `ByteVectorData` type:

```
' this is a data file
&H55
2      ' comment
3

4, 5
' another comment
5,
&Haa
```

Below is an example of the initialization data for a `SingleTableData` type:

```
.30103,           3.14159 ' log of 2 and pi
-200.,          1e05
+6.02E+23     100
```

Here are examples of in-line initializers for one-dimensional and two-dimensional types.

```
Dim d1 as ByteVectorData({ 20, &Hff, &H20, "row" })

Dim strList as StringVectorData({
    "alpha", "bravo", "charlie", "delta", "echo", "fox" + &H5f + "trot"
})

Dim tbl as New SingleTableData({
    ' column 1           column 2
    .30103,             3.14159
    -200.,              1e05
    +6.02E+23,         100
})
```

The values specified in an in-line initializer may be literal constants as shown above or they may be named constants that are visible within the module. For example,

```
Const cval as Byte = &H20
Dim d1 as ByteVectorData({ 20, &Hff, cval, "row" })
```

Program Memory data items have an associated property named `DataAddress`. The value of this property is the address of the data item in Program Memory. The type of the property is `Long` for compatibility with BasicX and the `GetProgMem()` subroutine.

Example

```
Dim addr as Long

addr = tbl.DataAddress
```

It is possible, also, to use the `DataAddress` property to get the address of a particular Program Memory data item. To accomplish this, simply add parentheses following the property name and specify the index or indices of the item of interest. The example below will result in `addr` having the Program Memory address of the second data value of the first row of the table.

```
addr = tbl.DataAddress(2, 1)
```

Caution: Program Memory data tables are arranged in memory in row-major order, i.e. the column values for the first row, followed by the column values of the second row, etc. This is a direct result of scanning the initialization data row by row. When you index a data table, you must specify the column index first and the row index second. This is backward with the respect to the way matrices are often visualized, i.e. (row, column). This strategy was adopted to maintain compatibility with BasicX. See Section 3.16 for more information on array data order.

Note that the `UBound()` function is useful with Program Memory data items to determine the dimensions of the vectors and tables. `LBound()` will always return 1 since initialized Program Memory data items are always 1-based.

For special situations, the data initialization file may be supplied in “raw” form. This means that the initialization file contains actual binary data as opposed to containing formatted data items. You instruct the compiler to interpret the initialization file in raw mode by appending `Attribute(raw)` to the end of the definition as in the example below. Note, particularly, that raw mode can only be used for numeric types (i.e. not String types) and can only be used for vector types (i.e. not table types). Also, the data file must contain an integral number of data elements. For example, for `IntegerVectorData` the data file must contain an even number of bytes.

```
Dim d1 as ByteVectorData("mydata.dat") Attribute(raw)
```

Program Memory variables may also be defined using a syntax similar to that used for defining RAM-based variables, using the keyword attribute `ProgMem` preceding the type name. For example,

```
Dim d1(1 to 20) as ProgMem Byte
```

This defines and reserves space for an array of bytes in Program Memory. Variables defined in this way will be zero-filled. Strings in Program Memory may be defined as well using the bounded string syntax. In this case, the string will have an initial value representing an empty (zero length) string.

```
Dim ps as ProgMem BoundedString(15)
```

Program memory structures may also be defined, see Section 3.25 for more details. A Program Memory variable may also be defined using `Based` keyword, see Section 3.21 for more details.

Caution: although Program Memory data items can be modified, the memory in which they are stored has a write cycle limit. For ZX devices with external Program Memory (e.g. the ZX-24a), the limit is approximately a million writes. For ZX devices with internal Program Memory (e.g. the ZX-24n) and for all generic target devices, the limit is approximately 100,000 writes. Writing to a particular address more than this may cause the memory to become unreliable. Also, writing to Program Memory is much slower than writing to RAM-based variables.

BasicX Compatibility Note

In BasicX mode, Program Memory string types are not supported nor are any vector types other than the `Byte` types. Also, in-line initializers are not supported, the `DataAddress` property cannot be used to determine the address of an individual Program Memory data element, and quoted strings cannot be used to specify data values. Finally, Program Memory data items may only be defined at the module level.

You may completely omit the initialization data from the definition of a Program Memory data item, including the parentheses that normally enclose it. If you do this, you must use the `Source` method to specify the initialization data as shown below. This alternate initialization mechanism is supported for backward compatibility with BasicX and is not recommended for new applications. Note, particularly, that this somewhat odd construction involving a `Call` does not produce any run-time executable code. It is merely a signal to the compiler to read the initialization data from the specified file.

```
Dim d1 as New ByteVectorData
```

```

Sub Main()
    Dim b as Byte

        ' specify the initialization data
    Call d1.Source("mydata.txt")

    b = d1(2)
End Sub

```

Only one method of specifying the initialization data can be used for any particular Program Memory data item. Attempting to specify the initialization data multiple times will result in a compiler error even if the data supplied in the multiple cases is identical.

2.11 String Types

ZBasic supports several variations of the fundamental type **String**. You may define a string variable thus:

```
Dim msg as String
```

The amount of space required for this variable definition and the maximum size of the string that it can represent varies depending compiler command line options and Option Directives. By default, the maximum string size is 255 characters. See Section 3.27 for more information on the implementation details of the various string data types.

2.11.1 Bounded Strings

A bounded string is nothing more than a way to specify a string having a maximum length that may be different than the default string length. A bounded string is defined using the syntax:

```
{Public | Private | Dim} <name> as [New] BoundedString(<size-expr>)
```

When defining a bounded string, you replace the *<size-expr>* with a constant integral expression specifying the number of bytes to allocate for the string's characters. It is important to be aware that assigning a string that has more characters than a bounded string variable's capacity will cause locations following that variable to be overwritten. Therefore, it is imperative that you write your code to explicitly limit the length of a string assigned to a bounded string variable.

Examples

```

Const slen as Integer = 7
Dim msg as BoundedString(15)
Dim msg as BoundedString(slen + 2)

```

The first definition will create a string variable that can hold up to 15 characters; the second will hold 9 characters.

For compatibility with BasicX, the alternate syntax shown below is also supported. New applications should use the definition syntax given above since it allows the use of an expression to specify the length.

```
{Public | Private | Dim} <name> as New BoundedString_<length>
```

2.11.2 Fixed-Length Strings

For compatibility with BasicX, ZBasic supports fixed-length strings. These are similar to bounded strings but with two important distinctions. Firstly, the string size is constant and equal to the specified fixed size.

If a string value is assigned that has fewer characters than a fixed-length string variable's specified size, the remaining characters will be filled with spaces. Secondly, if a string value is assigned having more characters than a fixed-length string's specified size, the excess characters will be discarded.

A fixed-length string is defined using the syntax:

```
{Public | Private | Dim} <name> as [New] String * <size-expr>
```

The keyword `New` is optional except in BasicX mode where it is required for compatibility reasons. When defining a fixed-length string the `<size-expr>` should be a constant integral expression specifying the number of bytes to allocate for the string's characters.

Example

```
Dim msg as String * 15
```

This definition will create a string variable that always contains exactly 15 characters.

2.12 Variable Initialization

All statically allocated variables are initialized by the system immediately prior to `Main()` beginning execution. For `String` types, this means that the bytes comprising the variable are set to represent an empty string. For all other types, the constituent bytes are set to zero. Variables defined at the module level and those defined using `Static` within a subroutine or a function are statically allocated and are, therefore, initialized.

Dynamically allocated variables are not initialized by the system except for `String` types which are initialized to represent an empty string. Variables defined using `Dim` within a subroutine or a function are dynamically allocated.

When you define a variable you may provide an initial value by adding an equal sign and the desired value (which must be a constant valued expression) following the variable's type. Initialization is not supported for arrays, structures, Based or Alias variables nor for Program Memory or Persistent Memory data items.

Examples

```
Dim count as Integer = 5  
Dim str as String = "column"
```

2.13 Type Conversions

The ZBasic language is strongly typed meaning, for example, that it is not allowed to assign the value of a constant, variable or parameter of one type to a variable or parameter of a different type. There are two apparent exceptions to the strong-type regimen. The first exception is with respect to integral numeric literals. An integral literal is considered to have a universal integral type (32-bit internally) so it can be assigned to a parameter or variable of any integral type (`Byte`, `Integer`, `UnsignedInteger`, etc.). Note that the presence of a plus sign or minus sign on a numeric literal does not change this interpretation so it is allowable to assign the value `-1` to an unsigned variable type.

The second apparent exception to the strong typing rules occurs with the System Library routines. Many of these routines will accept two or more data types for some of their parameters. It is as though several different versions of the library routines exist, differing only in the types of the parameters that they accept. This computer science concept is known as polymorphism.

The System Library routines include a set of functions for performing type conversions. The first set, `CBool()`, `CByte()`, `CInt()`, `CUInt()`, `CLng()`, `CULng()`, `CSng()` and `CStr()` allows, with some exceptions, conversion of a value of an arbitrary type to the target type. The second set, `FixB()`,

`FixI()`, `FixUI()`, `FixL()`, and `FixUL()` are specifically for converting `Single` values to the target type. The difference between using `FixI()` and `CInt()`, for example, to convert a `Single` value is the rounding method used. The final set of conversion functions, `CType()` and `To<enum>()` is for converting an integral value to an enumeration member. See the ZBasic System Library Reference Manual for more details on these conversion functions.

2.14 Parameter Passing Conventions

When a subroutine or function is defined, part of the definition specifies the parameters and their types that are expected by the routine. These parameters are referred to as the “formal parameters”. When a subroutine or function is invoked, parameters must be provided that match the formal parameters in order, number and type. These parameters are referred to as the “actual parameters” for each invocation.

As discussed earlier, parameters may be passed to subroutines and functions either “by value” or “by reference”. The differences between these two methods are subtle but important. When passed by value, the actual parameter may be an individual value like a constant or variable or it may be an expression. In either case the value of the actual parameter is calculated and the resulting value is passed to the routine. Within the called routine, the passed value may be utilized in any manner; it may even be modified and this modification will have no effect on any of the constituent elements of the passed parameter value. Effectively, the called routine gets its own private copy of the passed value.

When passed by reference, the address of the actual parameter is passed to the called routine. This implies that the actual parameter must be an individual variable or another parameter and not an expression nor a constant. Since the called routine has the address of the actual parameter, it is able to both read from and write to the actual parameter unless otherwise restricted. The ability of a routine to modify a variable passed by reference is often useful, especially in cases where a routine needs to produce multiple values for use by the caller. On the other hand, if used carelessly, it can be the source of errant program operation that is difficult to diagnose.

There are restrictions on whether a particular variable type may, may not, or must be passed by value or by reference. For example, a persistent variable cannot be passed by reference because the called routine is expecting the address of a RAM-based variable. Also, for efficiency reasons some variables are always passed to routines by providing the variable address, even if the definition of the routine specifies the parameter is to be passed by value. In such cases, the compiler treats the parameter as being read-only, effectively enforcing the semantics of pass-by-value. Any attempt to modify a read-only variable in the called routine or to pass it by reference to another routine will be detected and reported as an error by the compiler. Arrays may only be passed by reference and then only if they are RAM-based, single-dimension and have a lower bound of 1.

Allowed Parameter Passing Methods

Actual Parameter Type	Pass By Value	Pass By Reference
Constant or expression, any type	Yes	No
RAM-based variable or array element	Yes ¹	Yes ²
RAM-based single-dimension array, 1-based	No	Yes
RAM-based single-dimension array, not 1-based	No	No
RAM-based multi-dimensional array	No	No
Persistent variable or array element	Yes	No
Persistent Memory array	No	No
Program Memory array element	Yes	No
Program Memory array	No	No

Notes:

¹ String types and structures are read-only within the called routine when passed by value.

² The sub-byte types, Bit and Nibble, cannot be passed by reference. Also, in BasicX mode, `UnsignedInteger` and `UnsignedLong` types are read-only within the called routine when passed by reference.

The table below gives the number of bytes of stack space required to pass different variable types using the two passing methods. This table applies only to VM mode devices. For native mode devices, parameters are passed in registers in most cases and, therefore, require no stack space.

Actual Parameter Type	Pass By Value	Pass By Reference
Boolean, Byte	1	2
Bit, Nibble	1	2^2
Integer, UnsignedInteger, Enum	2	2
Long, UnsignedLong, Single	4	2
String ¹ , structure	2	2
Array, any type	n/a	2

Notes:

- ¹ Persistent strings, Program Memory strings and strings returned by functions all require 4 bytes of temporary data space (local to the caller) plus the 2-byte reference when passed to a routine other than a System Library routine.
- ² Sub-byte types like Bit and Nibble may only be passed by reference if they are byte aligned. See Section 3.24.1 for details.

2.15 Program and Data Item Properties

Most data items, whether located in RAM, Persistent Memory or Program Memory have an associated property called `DataAddress` that evaluates to the address of the data item. The `DataAddress` property is applied to a data item by appending it to the data item's name with a period separating them as illustrated by the example below.

Example

```
Dim b as Byte
Dim addr as UnsignedInteger

Sub Main()
    addr = b.DataAddress
End Sub
```

The `DataAddress` property can be applied to arrays and structures as well. When used with arrays it is best to append it after the array indices, if any. For most data items, the type of the `DataAddress` property is `UnsignedInteger`. However, for compatibility with `GetProgMem()` and other routines related to Program Memory, the type of the `DataAddress` property for Program Memory data items is `Long`.

Along similar lines, subroutines and functions have an associated property called `CodeAddress` whose type is `Long`. The `CodeAddress` property is employed in a similar manner as the `DataAddress` property is as shown by the example below. Of course, use of the `CodeAddress` property of a subroutine is not limited to the code in the subroutine itself. It can be applied to any subroutine or function that is visible to the code. In short, if you can invoke the subroutine or function, you can also get its address via the `CodeAddress` property.

Example

```
Dim addr as Long

Sub Main()
    addr = Main.CodeAddress
End Sub
```

It should be noted that for native mode devices, the `CodeAddress` property returns a word address instead of the usual byte address. In most cases, this fact is transparent but it should be kept in mind for special circumstances such as passing the address of a procedure to `GetProgMem()` for example. One advantage of this anomaly is that procedure addresses for ZBasic devices containing up to 128K of Program Memory may be represented in 16-bit data items on native mode devices.

2.16 Default Visibility

For subroutines, functions, along with variables and other data items defined at the module level, the default visibility (i.e. whether the entity is Public or Private by default) varies depending on the type of entity as shown in the table below. Some of the entities, e.g. Structure and Class types, are described later in this manual.

Default Visibility by Entity Type

Entity Type	Default Visibility
Constant	Private
Variable	Private
Subroutine	Public
Function	Public
Structure	Public
Enumeration	Public
Class	Public

Chapter 3 - Advanced Topics

This chapter provides additional technical information on topics that were introduced earlier in this document. Also, some more advanced concepts are introduced.

3.1 Scope and Lifetime

There are two important attributes of variables that have been alluded to in earlier discussion – scope and lifetime. The scope of a variable reflects its visibility in the sense of where it can be accessed by name. A variable defined within a subroutine or function has local scope meaning that it is only directly accessible to code within that routine. A variable defined outside of any routine has module scope if it is declared `Private` and global scope if it is declared `Public`. Module scope means that only routines within that module can access it directly. Global scope means that any routine in the application can access it directly.

These three scoping levels, global, module and local, form a hierarchy that controls the visibility of the variables. Global scope is at the outermost level of the hierarchy, module scope is at the next inner level and local scope is at the next inner level to that. There are additional inner levels of scope created by compound statements, which topic is discussed further below.

At any particular level of the scoping hierarchy, variables that are in the same scope level or farther outward are visible. It is possible to define variables with the same name at different scoping levels. This does not cause a conflict because the compiler resolves a reference to a particular variable name by searching the current scoping level first and then proceeding outward in the hierarchy until the variable name is found or not as the case may be. Variables are said to “hide” same-named variables that exist at outer scoping levels. In most cases, the hidden variables can still be accessed but more information has to be added to the variable name to clarify to the compiler which variable is being referenced. This concept may be clarified by an example.

Example

Module T1:

```
Public i as Integer
```

Module T2:

```
Private i as Integer
```

```
Sub foo()
    Dim i as Integer

    i = 5      ' this refers to the locally defined variable
    t2.i = 5   ' this refers to the private variable at the module level
    t1.i = 5   ' this refers to the public variable in module T1
End Sub
```

The second and third references to the variable `i` in the example above are qualified by the addition of the module name containing the definition of the desired variable. You may add module qualification to any module level and global level variable reference if you wish but it is generally only done when required to resolve the reference to the intended variable.

Although the preceding discussion focused on variables, the same scope concept applies to all identifiers – variables, constants, subroutines and functions. A local constant named `count` will hide a module level variable of the same name. Because the inadvertent hiding of identifiers is a common cause of programming errors the compiler, by default, issues a warning about the hiding. The warning can be disabled if desired. See Section 10.2 for specific information on various compiler options.

The second important attribute is lifetime. This concept refers to how long storage space is reserved for a variable. For variables defined at the module level, the lifetime is indefinite. They exist as long as the program is running. For variables defined within a routine, the lifetime normally begins when the routine begins execution and it ends when the routine finishes execution. Because these variables are dynamically created and destroyed, they are referred to as dynamic variables. This is in contrast to the module level variables which are static variables – they exist for the duration of the program's execution.

Sometimes, it is convenient to have a variable that is visible only to the routine in which it is defined but which is also static. ZBasic supports this concept by allowing the use of the `Static` attribute following the `Dim` keyword in a variable definition within a routine. The `Static` attribute tells the compiler to allocate space for the variable alongside the module level variables but since it is defined within a routine, i.e., it has local scope, only the code in that routine can directly access the variable.

Example

```
Private Sub mySub()
    Dim var1 as Integer
    Dim Static var2 as Integer

    [other code here]
End Sub
```

The difference between `var1` and `var2` is that space is allocated on the stack for `var1` when the routine begins executing while the space for `var2` exists as long the program is executing. Another difference is that dynamically allocated variables like `var1` have an undefined value immediately after they are created, you have to add code to initialize them. By default, the compiler will issue a warning if you write code that uses the value of a dynamically allocated variable before it is initialized. Note that `String` variables are a special case in that they are automatically initialized to a zero length.

In contrast, statically allocated variables like `var2` and all module level variables are initialized to zero just before the `Main()` begins running. Each time `mySub()` is invoked, the value of `var1` is undefined but the value of `var2` is whatever was assigned to it last. A consequence of this difference comes into play if `mySub()` is recursively invoked. Each invocation of `mySub()` will have its own private version of `var1` but they will all share the same `var2`.

BasicX Compatibility Note

In BasicX mode, variables cannot be defined as `Static`.

For all compound statements (If-Then, Do-Loop, For-Next, Select-Case and While-Wend), you may define additional variables within the body of the compound statement. When this is done, those variables will only be directly accessible to statements within the compound statement, including any nested compound statements. This is another example of local scope described above. If a variable so defined has the same name as a variable defined in an enclosing compound statement, in the routine itself, or at the module level, the newly defined variable obscures the same-named variable defined at the outer level rendering it inaccessible by normal means.

Example

Module Test:

```
Dim i as Integer
```

```

Sub Main()
    Dim i as Byte, j as Byte

    i = 44
    Test.i = 55
    For j = 0 to 1
        Dim i as Byte

        For i = 1 to 3
            Dim i as String

            i = "Hello"
            Debug.Print i
            Debug.Print CStr(Main.i)
            Debug.Print CStr(Test.i)
        Next i
    Next j
End Sub

```

The first Debug.Print will display the string "Hello" because it refers to the variable *i* that is defined in the body of the innermost For loop. The second Debug.Print will display the value 44 because it refers to the variable *i* defined at the outermost level of the subroutine using the subroutine name qualifier. The third Debug.Print will display the value 55 because it refers to the module-level variable *i* using the module qualifier. Within the body of the innermost For loop, there is no way to access the loop index variable because it is hidden by the local definition. There is no qualifier that can be added to a variable reference to resolve to outer block scope levels except the block formed by the routine definition itself.

Since ZBasic implements true block scoping, one advantage to using variables defined within compound statements is that, in addition to the restricted visibility, the stack space used by the variables can be reused by local variables defined in subsequent compound statements.

Example

```

If (j > 5) Then
    Dim i as Byte
    i = 12
    Call PutPin(i, zxOutputLow)
Else
    Dim s as Single
    s = 3.14159
    debug.print CStr(s)
End If

For k = 1 to 10
    Dim b as Byte
    b = GetPin(13)
    Debug.print CStr(b)
Next k

```

In this example, the three variables *i*, *s* and *b* share the same stack space. That works because none of them are “active” at the same time.

BasicX Compatibility Note

In BasicX mode, block scoping of variables is not supported. It is permitted to define variables within a compound statement but the effect is the same as if they were defined at the beginning of the routine.

3.2 Enumerations

In some situations it is convenient to be able to refer to the values of a variable by a name rather than by a numeric value. An enumeration type essentially allows you to define a new data type and name the set of values for that type. The syntax for defining an enumeration is:

```
[Public | Private] Enum <name>
    <member-name> [= <constant-expr>]
    ...
End Enum
```

In this syntax, *<member-name>* is an identifier that names a member of the enumeration. The optional *<constant-expr>* represents a value that you want to be associated with that member name. For any member that does not have an explicit member value specified, a member value will be automatically assigned that is one larger than the preceding member or zero for the first member.

The ellipsis in the syntax above indicates that there may be zero or more additional member definitions. Member names must be unique within the set of members for each enumeration. A particular member name may, however, be used in multiple enumerations. See the discussion below for information about how ambiguity is resolved.

If neither `Private` nor `Public` is specified on an enumeration definition, the enumeration is public. An enumeration may be defined at the module level or it may be defined within a subroutine or function, either at the outer level or within any inner block. In the latter case, the `Public` and `Private` keywords have no useful purpose and are therefore not allowed.

Examples

```
Enum Pet
    Dog
    Bird
    Snake
End Enum
```

In this case, the members will be assigned values of 0, 1 and 2 respectively.

```
Enum Mammal
    Cat
    Dog
    Elephant = 5
    Horse
End Enum
```

Here the members will have the values 0, 1, 5 and 6 respectively. Note that if explicit values are specified, they must be larger than the value assigned, explicitly or implicitly, to the preceding member. The value associated with enumeration members is unsigned.

After an enumeration has been defined, the enumeration name may be used as a *<type>* in a variable, constant or structure definition. Enumerations may also be used as the *<type>* in the formal parameter list of a subroutine or function definition. Note, however, that a `Public` subroutine or function cannot be defined with a parameter that is a `Private` `Enum`.

Example

```
Public animal as Mammal

animal = Elephant
animal = Mammal.Dog
```

In the second assignment, the enumeration name is used as a qualifier on the enumeration member name. This is always allowed but it is only required when ambiguity exists. See the discussion below for additional information on resolving ambiguity.

The numeric value of a member may be obtained by using System Library type conversion functions. Continuing the example from above:

```
Dim i as Integer, j as Integer  
...  
i = CInt(animal)  
i = CInt(Mammal.Dog)
```

It is also possible to convert an integral value to an enumeration member. There are two ways to accomplish this. The first, and recommended, way is to use the System Library function `CType()`. This function takes two parameters, the first being the integral value to convert and the second being the name of the enumeration.

```
animal = CType(3, Mammal)
```

The conversion will be performed even if the value specified does not actually correspond to any member of the enumeration so this type of conversion must be used carefully.

Wherever it is used, an enumeration name may be qualified with the module name containing the enumeration. Assume that the enumeration defined above exists in a module named `Test`.

```
animal = Test.Mammal.Dog  
animal = CType(3, Test.Mammal)
```

Qualification using the module name is only necessary in unusual cases but it is well to remember that it is allowed for the situations where you need it.

For an enumeration defined within a subroutine or function, the enumeration may be qualified with the subroutine/function name, optionally preceded by a module name qualification. For example, if the enumeration `Mammal` is defined in the subroutine `Main()` contained in the module `Test`, the following constructions are permitted within the `Main()` subroutine.

```
animal = Main.Mammal.Dog  
animal = Test.Main.Mammal.Dog
```

When resolving a reference to a member name, the following order is used. If the member reference occurs within a subroutine or function, the name is first checked against the set of members of all enumerations defined within that subroutine or function. If only one enumeration has a matching member name, no ambiguity exists. If no match was found, the name is next checked against the set of members of all enumerations defined to be private to the module. If only one enumeration has a matching member name, no ambiguity exists. If no private enumerations have a matching member, the public enumerations of all modules are checked next. Again, if only one enumeration has a matching member name, no ambiguity exists. This search order may be overridden by qualifying the member name with the name of the enumeration to which it belongs. The reference may be further qualified by adding the module name. This allows access to a public enumeration that is being hidden by an enumeration that is private to the module.

The only operations that can be performed on enumeration variables are comparison using relational operators, assignment and type conversion.

BasicX Compatibility Note

In BasicX compatibility mode, enumerations may only be defined at the module level and qualification of an enumeration is not supported. Also, the values of enumeration members are only 8-bits wide in BasicX while they are 16-bit values in ZBasic.

An alternate type conversion method is supported for compatibility reasons but is a bit awkward to use. For each defined enumeration, there is a special System Library conversion function just for that enumeration whose name is `To<enum>` where `<enum>` is replaced with the enumeration name. This special conversion function takes only one parameter, that being the value to convert.

```
animal = ToMammal( 3 )
```

One drawback to this conversion method is that there is no way to qualify the enumeration with the module name in which it is defined. It is recommended that all new applications use the `CType()` conversion function.

3.3 Serial Channels

All of the ZBasic devices support multiple serial channels. Channel 1 (Com1) is implemented using a USART hardware on the microcontroller chip (if available). On ZX devices, when the program begins running, Com1 is configured to run at 19.2K baud, 8 data bits, 1 stop bit. Your program can send and receive data using the default configuration by utilizing the console I/O routines like `Console.Write()`, `Console.Read()` and the related `Debug.Print` command. For generic target devices, Com1 is not open by default but “Option ConsoleSpeed” can be used to specify that it should be initially open.

Some ZBasic devices have additional hardware-based serial channels but the remaining serial channels are implemented in software. This strategy allows a lot of flexibility in choosing which pins are used for transmission and reception but it also imposes a processing overhead on the system, even when characters are not being actively sent and received. Because of this, only one additional software serial channel, COM3, is enabled for use by default. If you wish to utilize additional software serial channels (COM4 to COM6) you must call the System Library routine `ComChannels()` to specify both the number of channels desired and the maximum baud rate that may be used. See the description of the `ComChannels()` routine in the ZBasic System Library Reference manual for more details.

Since serial channels 3-6 rely on interrupts to achieve the necessary timing for serial I/O, if interrupts are disabled for a substantial fraction of the bit time (the inverse of the baud rate) the integrity of the transmitted or received characters may suffer. Typically, the maximum acceptable interrupt disable time is about 25% of the bit time of the fastest channel. If the fastest channel is running at 9600 baud, the interrupt disable time should be kept below 25 μ S or so. Many of the I/O routines that utilize Timer1 (e.g. `PulseOut`) disable interrupts in order to achieve precise timing. Those that do disable interrupts have a caveat to that effect in their respective descriptions in the ZBasic System Library Manual.

See the descriptions of `DefineCom()`, `OpenCom()` and `CloseCom()` in the ZBasic System Library Manual for more details on setting up and using a serial channel.

3.4 Queues

A queue is a fundamental data structure that is widely used in computer programs. Its primary distinguishing feature is that data items are extracted from the queue in the same order in which they were inserted. This is called first-in, first-out or FIFO order.

One of the uses for a queue in ZBasic is as a temporary buffer for data going to and coming from one of the serial channels. Another common use is as a medium through which to pass data between tasks. The producer of the data puts data in the queue and the consumer of the data takes it from the queue.

ZBasic queues are of a fixed length, that length being determined at the time that the queue is prepared for use using `OpenQueue()`. If the data producer inserts data faster than the consumer removes it, the queue will eventually become full. Further attempts to add data to the queue will stall until enough data is removed from the queue to make space for the new data.

Although the data in queues is nominally byte oriented, you can put data of any type into a queue. All that is necessary is for the producer and consumer of the data to agree on the nature and meaning of the

data. For example, the producer could copy several Single data values to a queue and as long as the consumer copies them out to Single variables all will be well.

Prior to using a queue it must be initialized by using the System Library routine OpenQueue(). An example of this is shown below.

```
Dim myQueue(1 to 40)

Call OpenQueue(myQueue, SizeOf(myQueue))
```

This call prepares the queue for use by initializing the first 9 bytes of the specified array with queue management data. The remainder of the array is used for the data to be held by the queue, in this case the queue can hold up to 31 bytes of data since it was defined as being 40 bytes long. This implies that the smallest array usable as a queue is 10 bytes. A pre-defined constant, System.MinQueueSize, can be used in your programs to represent this minimum size.

Note the use of the `SizeOf()` function. Although the second parameter could just as well have been the literal value 40, if you later changed the size of the `myQueue` array you'd have to also remember to change the second parameter to match. If you made the queue larger, opening the queue with a smaller value would have no ill effect other than wasting RAM. On the other hand, if you reduced the size of the queue array but left the larger value as the second parameter, the data area inside the queue would overlap adjacent variables leading to puzzling results. For this reason it is highly recommend to use `SizeOf()` in this and other similar situations rather than hard-coded constants. An alternative method is to use a defined constant as shown below.

```
Const myQueueSize as Integer = System.MinQueueSize + 30
Dim myQueue(1 to myQueueSize)

Call OpenQueue(myQueue, myQueueSize)
```

Either of these methods of improving the maintainability of your program is acceptable. Which you choose is more of a stylistic issue than a technical one. The code generated by the compiler in the two cases is identical.

There are several System Library routines available for adding data to and extracting data from a queue as well as some querying functions for determining the status of the queue. Once initialized as shown above, the code fragment below will place some data in the queue.

```
Dim punct(1 to 2) as Byte

punct(1) = &H2C
punct(2) = &H20
Call PutQueueStr(myQueue, "Hello")    ' add a string
Call PutQueue(myQueue, punct, 2)       ' add some individual bytes
Call PutQueueStr(myQueue, "world")    ' add another string
Call PutQueueByte(myQueue, &H21)       ' add an exclamation point
```

After the data is in the queue, the following code fragment will extract it and display it on the console.

```
Do While (StatusQueue(myQueue)    ' add a string
          Dim b as Byte
          Call GetQueue(myQueue, b, 1)      ' retrieve a byte from the queue
          Debug.Print Chr(b);
Loop
Debug.Print
```

It is important to note that the queue insertion routines all wait until there is sufficient space available in the queue for the data being inserted before inserting any data. This may lead to deadlock situations, particularly if the size of the data being inserted is larger than the queue's data area.

3.4.1 System Queues

The console I/O routines like Console.Read() and Console.Write() use the input and output queues associated with Com1. Because of their special use in this manner, the queues associated with Com1 are called system queues. The values Register.RxQueue and Register.TxQueue give the address of the system input queue and system output queue respectively. The function CByteArray() can be used to convert these values to a reference to a Byte array allowing them to be passed to the queue-related routines. The example code below shows how to determine if there is any data available in the system input queue using this technique.

```
If StatusQueue(CByteArray(Register.RxQueue)) Then  
    b = Console.Read()  
End If
```

The system queues are initially set to be small pre-defined queues. If your code opens Com1, the queues that you provide with the OpenCom() invocation become the system queues until you subsequently close Com1. At that time, the system queues will revert to the pre-defined queues.

Of the two pre-defined queues, the output queue is the smallest, having space for only a few data bytes. Because of this, if you send a fairly long string to the output queue using Console.Write() your application will experience a delay until the string's characters can all be transferred to the output queue. If this causes a problem in your application you can define a larger queue to be used as the output queue. The example code below does so while retaining the pre-defined input queue.

```
Private sysOutQueue(1 to 40) as Byte  
  
Call OpenQueue(sysOutQueue, SizeOf(sysOutQueue))  
Call CloseCom(1, 0, 0)  
Call OpenCom(1, Option.Com1Speed, _  
            CByteArray(Register.RxQueue), sysOutQueue)
```

For native mode devices (e.g. the ZX-24n) you can specify the default sizes for the system queues using the directives Option TxQueueSize and Option RxQueueSize.

3.5 Multitasking

The concept of multitasking is a very powerful one although it can be, at first, somewhat confusing. In a single task system, there is only one program and it runs continuously until it is terminated. During the course of execution of the program it may perform several different functions, e.g. checking the keyboard to see if a character is available, updating the display unit, etc. The fundamental idea behind multitasking is that these activities are divided into related groups, each of which is called a task. Each task executes for a period of time and then the next task executes. There are two basic types of multitasking: cooperative and preemptive. In cooperative multitasking it is left up to each task to determine when to suspend itself and allow the next task to run. In preemptive multitasking a part of the operating system called a task scheduler determines, generally, how long a task is allowed to run before switching to the next task.

The multitasking in ZBasic is, fundamentally, preemptive multitasking. The task scheduler switches tasks on each RTC tick (approximately 1.95mS). A task can, however, lock itself to avoid (with some exceptions) giving up execution. This capability should be used sparingly because it defeats the equitable sharing of processing resources.

It is not necessary, of course, to structure your program to utilize the multitasking capabilities. If there is only one task defined (the subroutine Main()) the task scheduler will simply allow it to run continuously. If your application might benefit from being divided into tasks, it is simple to define them. Each task has two core elements: a task main routine and a task stack. The task main routine is the (usually parameterless) subroutine with which the task begins execution. The task stack is a portion of RAM set aside exclusively for the task. The example below shows how to define a task stack and how to activate a second task.

```

Dim taskStack(1 to 80) as Byte

Sub Main()
    CallTask "MyTask", taskStack
    Do
        Debug.Print "Hello from Main"
        Call Delay(1.0)
    Loop
End Sub

Sub MyTask()
    Do
        Debug.Print "Hello from MyTask"
        Call Delay(2.0)
    Loop
End Sub

```

This simple program has two tasks: Main and MyTask. The Main task is created automatically for you and its task stack is automatically allocated all of the remaining User RAM after explicitly defined variables are allocated. In contrast, additional tasks such as MyTask have to be explicitly invoked using the CallTask statement and each task's stack must also be explicitly allocated, for example by defining a Byte array as shown above. A task is said to be “active” if it has been invoked by CallTask and has not yet terminated. Both of the tasks above never terminate so they are always active. The example below shows a situation where a task terminates.

```

Dim taskStack(1 to 50) as Byte

Sub Main()
    CallTask "MyTask", taskStack
    Do
        Debug.Print "Hello from Main"
        Call Delay(1.0)
    Loop
End Sub

Sub MyTask()
    Dim i as Integer
    For i = 1 to 5
        Debug.Print "Hello from MyTask"
        Call Delay(2.0)
    Next i
End Sub

```

In this example, the task MyTask will output the message 5 times and then terminate. After the task terminates, its task stack could be used for another task. It is important to note that tasks cannot use the same task stack if they will be active at the same time.

The examples above were chosen because they clearly illustrate, when executed, that the separate tasks experience interleaved execution. A more realistic example is shown (necessarily incompletely) below.

```

Dim taskStack(1 to 50) as Byte

Sub Main()
    CallTask "MyTask", taskStack
    Do
        <other-important-stuff>
    Loop
End Sub

Sub MyTask()
    ' set the interval timer to 5 seconds

```

```

Call SetInterval(5.0)
Do
    WaitForInterval()
    <stuff-to-do-periodically>
Loop
End Sub

```

In this partial example, a task is set up to perform some important activity periodically. If this were part of a system to control the chlorine level in a pool, `MyTask` might read a transducer that indicates chlorine level and, if it's below a certain level, cause a known amount of chlorine solution to be injected into the circulation stream.

One important aspect of creating and using tasks is the allocation of a task stack. The task stack has two components: a task control block and the portion used as an execution stack. The task control block is a fixed size (see Section 3.30 for details) and resides either at the beginning (for VM mode devices) or at end (for native mode devices) of the task stack. The remainder of the task stack is dedicated to the execution stack for the task. The execution stack is used for local variables in the task routine, parameters passed to other subroutines and functions that might be invoked as well as local variables used in those routines, and space required for expression evaluation.

If a task stack is allocated that is much larger than is actually necessary, User RAM is wasted since no other task can use the excess space. On the other hand, if a task stack is smaller than required, data items preceding or following the task stack in memory will be overwritten. This will generally cause your application to misbehave and may result in the processor resetting.

This begs the question, of course, of what is the optimum task stack size. Unfortunately, this is not as easy to answer as it might seem. The difficulty in answering this question lies in the fact that stack use is dynamic and may depend on external factors such as a signal applied by an external device or a user pressing a key. Not only that, the stack use may depend on the order or the timing of such external events. Further, recursive invocation of routines adds to the dynamic nature of stack use and is impossible to account for (in most cases) using any kind of static analysis.

The ZBasic compiler automatically estimates the minimum task stack size for each task in your application. Information about the estimate, including the stack use of individual routines, is displayed in the .map file (assuming that one is generated as is the default). If the size of the stack that you have allocated is smaller than the estimated size plus a safety margin, the compiler will generate a warning indicating the deficiency and indicating the minimum allocation required. By default, the safety margin is 10 bytes but you may specify a larger or smaller safety margin (including zero) using `Register.StackMargin`.

To assist you in finely tuning the task stack size a special System Library function is provided that determines, at run time, how much unused space exists in the task stack. By exercising your application and periodically checking the amount of unused task stack space exists, you may be able to determine empirically a more suitable minimum task stack size. See the description of `System.TaskHeadRoom()`

3.5.1 Advanced Multi-tasking Options

For advanced users, a task's main routine may also be defined with parameters. This may be useful to provide information for the task to perform its function rather than doing so using global variables. Parameters for the task are specified in a comma-separated list enclosed in parentheses just as they are for a subroutine invocation. The number and types of the parameters specified must match those of the task being invoked.

For special circumstances, it is possible to specify the task's stack by giving its address as an integral value. Generally, it will also be advisable to specify the size of the stack since the compiler will be unable to deduce the size. Unless the task stack size is specified or can be deduced, no stack overrun checking can be performed. For native mode ZBasic devices the task size must be explicitly specified or it must be determinable by the compiler at compile time, otherwise, the compiler will issue an error message.

The syntax for these advanced options is described in the ZBasic System Library Reference manual. Also, see Section 3.30 for additional information on task management.

3.6 Semaphores

In computer science a semaphore is a mechanism used to control access to a shared resource, the idea being to prevent more than one actor from attempting to use, modify or update a resource simultaneously. As a physical analogy, consider a long single-lane tunnel on a roadway. A driver waiting to enter at one end would like to be certain that no car enters the tunnel from the opposite end while he is enroute. A semaphore indicating that a vehicle is in the tunnel would, if properly observed, help prevent simultaneous use of the shared resource. Clearly, however, if not all drivers understand the meaning of the semaphore or if a driver ignores the semaphore an accident is likely to occur. So it is with semaphores in a computer program.

In a single task system there is no need for semaphores because there are not multiple actors to coordinate. However, in a multitasking system (or a multiprocessor system) because there are multiple actors there is a need to coordinate access to shared resources such as a serial channel, a timer a data structure, etc. Whether or not you need to use a semaphore in your program depends on how it is structured and what resources are shared between tasks.

The essential element of a semaphore is known in computer science as an *atomic test and set* operation. The basic idea is that a method is needed to test a Boolean variable to see if it is already set true and if it is not set then set it to true. The adjective *atomic* in this case refers to the fact that the testing phase and the setting phase are indivisible. That is, once a task begins the process of testing, no other task is able to begin testing until the first task has completed the test and set. This prevents what is called a *race condition*.

The atomic test and set operation in ZBasic is provided by the System Library routine `Semaphore()`. To implement a semaphore you must define a Boolean variable and ensure that it is set to `False` initially. Then, whenever a task wants to use the shared resource it must first call `Semaphore()` passing the previously defined variable as a parameter. If the semaphore is already set, the function will return `False` indicating that the resource is busy. If the semaphore is not already set, the function will set it to `True` and return `True` indicating that the semaphore has been successfully obtained. Once the task is finished with the resource, it must set the semaphore variable back to `False` again so that the next user may successfully acquire a semaphore on the resource. The example code below illustrates the sequence.

```
Dim serSem as Boolean  
  
serSem = False  
  
' wait until we get the semaphore  
Do While (Not Semaphore(serSem))  
Loop  
  
' now we can use the controlled resources  
[add code here]  
  
' finished with the resources, release the semaphore  
serSem = False
```

3.7 Built-in Variables

The set of pre-defined registers comprises two sub-groups: actual CPU registers and control program variables. All of these built-in variables must be referenced using the `Register` prefix or within a `With Register` compound statement.

Example

```
Dim tick as Long  
tick = Register.RTCTick
```

In addition to the particular registers available for each underlying processor, three special register variables are provided that correspond to the PORT, PIN and DDR registers associated with a particular I/O pin. The form of these three special register functions is: Register.Port(<pin>), Register.Pin(<pin>) and Register.DDR(<pin>), respectively. In all three cases, the property .DataAddress may be appended to get the address of the register associated with the specified pin.

Example

```
val = Register.Port(C.2)                      ' reads Register.PortC  
Register.Port(C.2) = val                      ' writes to Register.PortC  
addr = Register.Pin(A.4).DataAddress          ' address of Register.PinA
```

These three special register variables will often be used in conjunction with the ZBasic System Library function PortMask() which returns the bit mask for an I/O port corresponding to a specified pin.

Example

```
Const pin as Byte = C.7  
Call SetBits(Register.DDR(pin), PortMask(pin), &Hff)    ' make a pin an output
```

3.7.1 Special Function Registers

Each supported target processor has a number of special function registers (SFRs), some of which are 8 bits wide and some of which are 16 bits wide. The SFRs can be accessed in application code by adding the prefix Register. to the SFR name, e.g. Register.PortC for the register PORTC. For the details on the purpose and use of each SFR, see the manufacturer's datasheet for the particular target processor.

The SFRs for the supported targets are described in a series of XML resource description files located in a subdirectory of the ZBasic installation directory. See Appendix N for more information on the resource description files.

3.7.2 System Variables and Constants

Some of the ZBasic register values allow access to certain variables used by the system control program as described below. Unless otherwise indicated, the registers are readable and writable. However, for some of the variables, modifying their values may have undesirable or unpredictable effects on your program. See the descriptions of the individual items for more information.

BasicX Compatibility Note

In BasicX mode, only a subset of the variables described below is available, being limited to those that are also present in BasicX. Those that are available in BasicX mode are specifically identified in their descriptions.

Register.BootVersion

This UnsignedInteger value gives the version number of the bootloader installed on the ZX device. The high byte gives the major value while the low byte gives the minor value. For example, if the bootloader version number is v1.4, the value of Register.BootVersion will be &H0104.

Register.ATNChar

This read/write Byte value gives represents a special trigger character that, if received on the Com1 serial channel, will put the ZBasic device into “command mode”. The purpose of the ATN character is to facilitate downloading code to the device without using the normal DTR toggling trigger. The value of Register.ATNChar will act as a trigger only if its value is in the range &H00 to &H1f (0 to 31 decimal). The default value of Register.ATNChar is 0Hff, effectively disabling the feature.

Register.ResetFlags

Whenever the processor is reset, the cause of the reset is noted and stored in an internal variable that is available to user programs as Register.ResetFlags. The bits of the Byte value have the meaning shown in the table below.

ResetFlags Bit Semantics		
Reset Source	ATmega/ATtiny	ATxmega
Power On Reset	&H01	&H01
External Reset	&H02	&H02
Brown Out Reset	&H04	&H04
WatchDog Reset	&H08	&H08
Debug Interface Reset		&H10
Software Reset		&H20
Spike Detection Reset		&H40

The value of this register is set just before the ZBasic Device begins executing your program. The value is not used by the system in any manner so you may modify it to suit the needs of your application. Note that more than one bit may be set.

Register.RTCDay **Register.RTCTick**

These two register values (both available in BasicX compatibility mode) represent the current state of the real time clock (RTC). Register.RTCTick is a Long value that is incremented on each RTC tick (see Register.RTCTickFrequency). After 24 hours of continuous execution, it will reach its maximum value and will then roll over to zero. At the same time, the value of Register.RTCDay, type UnsignedInteger, will be incremented. Day number zero represents January 1, 1999 (for compatibility with BasicX). When the system is reset or powered up both of these values are initialized to zero.

Register.RTCFastTick

This Byte value changes at a rate equal to that of Register.RTCTick times the value of RTCScale. For ZX devices, RTCScale is 2 except for devices running at 7.37MHz where it is 1. For generic target devices, RTCScale is a configurable parameter. For most purposes, the value of Register.RTCFastTick should be considered to be read-only. Changing it will affect the accuracy of the RTC and may interfere with normal task switching.

Register.RTCStopWatch

This UnsignedInteger value is incremented on each RTC tick (see Register.RTCTickFrequency), the same as Register.RTCTick. However, you may reset this value to zero at any time to facilitate simpler elapsed time calculations without affecting the RTC's timekeeping. This register variable is available in BasicX compatibility mode.

Register.RTCTickFrequency

This read-only UnsignedInteger value indicates the number of RTC ticks that will occur per second. For most ZX devices currently available, the RTC tick frequency is 512Hz. For generic target devices RTCTickFrequency is a configurable device parameter.

Register.RTCTimerFrequency

This read-only UnsignedLong value indicates the frequency at which the TCNT register of the RTC timer changes value. For most ZX devices currently available, the RTC timer frequency is 230,400Hz. For generic target devices RTCTimerFrequency is dependent on the configurable parameters CPUFrequency, RTCScale and RTCTickFrequency. This value is useful, for example, to convert the value returned by GetElapsedMicroTime() to seconds.

Register.CPUFrequency

This read-only UnsignedLong value indicates the frequency of the CPU clock, in Hertz. For most ZX devices currently available, the CPU frequency is 14,745,600Hz. For generic target devices CPUFrequency is a configurable device parameter.

Register.SeedPRNG

This register, having type Long, represents the “seed” value used by the built-in pseudo-random number generator. With a given seed value, the random number generator will always return the same sequence of values. Usually, you wouldn’t want this type of repeatability but for some purposes it is useful. See the descriptions for the System Library routines Rnd() and Randomize() for more details.

```
Register.Timer0Busy  
Register.Timer1Busy  
Register.Timer2Busy  
Register.Timer3Busy  
Register.Timer4Busy  
Register.Timer5Busy  
Register.Timer6Busy  
Register.Timer7Busy  
Register.TimerIOBusy  
Register.TimerRTCBusy  
Register.TimerUARTBusy
```

These Boolean values indicate when the processor’s built-in timers are being used. Your code can pass one of these register values as the parameter to the Semaphore() function in order to get exclusive access to the corresponding timer. The last two entries are synonyms for one of the earlier entries. This is useful because the actual timer used, for example, for timing I/O functions varies amongst ZBasic devices. See the Resource Usage subsection entitled Timers in the ZBasic System Library manual for more information on Timer use.

Register.TimerSpeed1 **Register.TimerSpeed2**

These two registers, both Byte values, represent the timer pre-scaler value used by several System Library routines for the I/O Timer. Setting the value of these registers other than by direct assignment will produce undefined results. See the Resource Usage subsection entitled Timers in the ZBasic System Library manual for more information on Timer use.

Register.IOScaling

This read/write Boolean value controls (for native mode ZBasic devices only) whether the results of certain ZBasic System Library routines are scaled to match the units of NetMedia's BasicX devices (which run at 7.3728 MHz) for code compatibility. The initial value is True for all mega-based ZX running at 14.7456 MHz and false for all other ZX devices and all generic target devices. Generally, it is inadvisable to change the value for devices having a default initial value of False.

Register.TaskStackMain **Register.TaskStackCurrent**

These UnsignedInteger values represent the address of the task stack for the Main() task and the current task, respectively. The value, which is read-only, can be passed to the various task management functions by using the System Library function CByteArray(). See Section 3.30 for more details on Task Management.

Register.TaskMain **Register.TaskCurrent**

These UnsignedInteger values represent the address of the task control block for the Main() task and the current task, respectively. The values, which are read-only, are the same as the respective task stack addresses (see above) for VM mode devices but not so for native mode devices. See Section 3.30 for more details on the task control block structure.

Register.StackMargin

This Byte value specifies how close to the end of the stack the stack pointer for a task may approach before triggering a stack fault. The default value is 6. See Section 3.33 for more information on Run Time Stack Checking. This built-in is useful only for VM code devices such as the ZX-24.

Register.FaultType **Register.FaultData** **Register.FaultData2**

These values give information about the last detected system fault. Register.FaultType is a Byte value that indicates the fault type. Register.FaultData and Register.FaultData2 are UnsignedInteger values that provide additional data about the fault. See Section 3.33 for more information on Run Time Stack Checking.

Register.RxQueue **Register.TxQueue**

These UnsignedInteger registers contain the address of the queues currently associated with Com1. These queues used for Console.Read(), Console.Write() and related routines. The values, which are read-only, can be passed to the various queue functions by using the System Library function CByteArray(). See Section 3.4.1 for more information on the system queues.

Register.Console.EOL

This Byte value represents the character that the system will recognize as the end-of-line character. It is initially set to the value of a line feed character (&H0a). See the discussion of Console.ReadLine() in the ZBasic System Library Reference for more information on how it is used.

Register.Console.Echo

This Boolean value, initially set to True, controls whether characters received by `Console.Read()` and `Console.ReadLine()` are echoed back to the sending device. See the descriptions of these two functions in the ZBasic System Library Reference for more information.

Register.Console.Speed

Deprecated – use `Option.ConsoleSpeed` instead. This read-only `UnsignedInteger` value gives the default console speed, the baud rate for which Com1 is configured when the system begins running.

Register.SignOn

This Boolean register contains the flag that controls whether or not the ZX issues a sign-on message when it starts up after reset. It is a Persistent Memory value that may also be set or cleared by using the `Option.SignOn` directive. See Section 2.3.1 for more information.

```
Register.CodeSize  
Register.RamSize  
Register.RamStart  
Register.RamUsed  
Register.PersistentSize  
Register.PersistentStart  
Register.PersistentUsed
```

These system values may be useful for diagnostic and other purposes. `Register.CodeSize` is a Long constant that indicates the number of bytes of Program Memory consumed by your program together with any Program Memory data items that it defines.

`Register.RamSize` is an `UnsignedInteger` constant that indicates the total number of bytes of User RAM that is directly available to your program. `Register.RamStart` is an `UnsignedInteger` constant that indicates address at which User RAM begins. `Register.RamUsed` is an `UnsignedInteger` constant that indicates the total number of bytes of User RAM that your program statically allocates. The difference between `Register.RamSize` and `Register.RamUsed` represents the number of bytes that will be allocated automatically for the task stack for the `Main()` task.

Similarly, `Register.PersistentSize` is an `UnsignedInteger` value that indicates the size of Persistent Memory, in bytes, that is available to your program. `Register.PersistentStart` is an `UnsignedInteger` value that indicates address at which User Persistent Memory begins and `Register.PersistentUsed` indicates the number of byte of Persistent Memory actually used by your program. These values are all read-only.

Register.HeapEnd

This read-only `UnsignedInteger` value indicates the lower bound of the memory allocation heap. The heap grows from the high end of RAM toward the beginning of RAM. As blocks of memory are allocated from the heap (for String variables or to satisfy `System.Alloc()` requests), the value of `Register.HeapEnd` will decrease. As blocks of memory are returned to the heap (as String variables change or go out of scope, or due to calls to `System.Free()`), the value of `Register.HeapEnd` may or may not increase. The value of `Register.HeapEnd` will not increase until the allocated block of memory closest to its current value is freed. At that time, `Register.HeapEnd` will increase to be near the lowest still-allocated block.

Register.ExtRamConfig

This UnsignedInteger value indicates the current external RAM configuration. Of course, it is only meaningful on ZBasic devices that support external RAM, e.g. the ZX-1281. Although you may change this value during execution, the configuration of the external RAM interface will not be affected until the next time the device resets (i.e. following power up, WatchDog reset, download, etc.).

Register.FirstTime

This read-only Boolean value is identical to that which would be returned by the `FirstTime()` function. Note, however, that reading the value via `Register.FirstTime` does not reset the flag like invoking the `FirstTime()` function does.

Register.UserSP

This UnsignedInteger register contains the value of the stack pointer for the current task at the moment it is referenced. This value may be useful in estimating the stack usage for a task. For native mode devices (e.g. the ZX-24n) the value of `Register.UserSP` is identical to `Register.SP`.

Caution: modifying this value will probably cause your program to malfunction.

Register.SP

This UnsignedInteger register contains the value of the hardware stack pointer of the underlying processor at the moment it is referenced. For VM code devices, there are few, if any, practical uses for this value. See, instead, `Register.UserSP`. For native mode devices, the value of `Register.SP` is identical to `Register.UserSP` and indicates the value of the current task's stack pointer at moment it is referenced.

Caution: modifying this value will probably cause your program to malfunction.

3.8 Built-in Constants

The compiler provides several built-in constants. These constants may be used in conditional directives and they may also be used as if they were constants defined by the normal means. The entries in the table below that begin with `Option` typically derive their value from the corresponding Option directives, the corresponding compiler command line options or their respective default values.

Name	Type	Description
<code>Module.Name</code>	String	The name of the module being compiled.
<code>Module.Number</code>	UnsignedInteger	The ordinal number of the module being compiled.
<code>Option.AllocStr</code>	Boolean	Indicates if dynamic string allocation is in effect.
<code>Option.AtnChar</code>	Byte	Gives the value of the ATN character (if non-zero).
<code>Option.Base</code>	Integer	Gives the array base value for the current module.
<code>Option.CodeLimit</code>	UnsignedLong	Specifies the code size limit (zero if none is specified).
<code>Option.Com1Speed</code>	Long	Deprecated. Use <code>Option.ConsoleSpeed</code> .
<code>Option.ComChannels</code>	Byte	Indicates the maximum number of SW UART channels.
<code>Option.ConsoleChannel</code>	Byte	Indicates the console serial channel (e.g. 1 means Com 1).
<code>Option.ConsolePin</code>	Byte	Gives the pin to be used for console output (zero if none).
<code>Option.ConsoleSpeed</code>	Long	Gives the current default console channel baud rate.
<code>Option.CPUFamily</code>	String	Deprecated. Use <code>Option.TargetFamily</code> .
<code>Option.CPUType</code>	String	Deprecated. Use <code>Option.TargetMCU</code> .
<code>Option.ExtRamAble</code>	Boolean	Indicates if the device supports external RAM use.
<code>Option.ExtRamEnabled</code>	Boolean	Indicates if external RAM use is enabled.
<code>Option.HeapLimit</code>	UnsignedInteger	Indicates the specified heap limit.
<code>Option.HeapReserve</code>	UnsignedInteger	Indicates the default or specified minimum heap size.

Option.HeapSize	UnsignedInteger	Indicates the specified or default heap size.
Option.Language	String	Specifies the language mode that is in effect.
Option.MainTaskStackSize	UnsignedInteger	Indicates the specified size for the main task stack.
Option.Objects	Boolean	Indicates if object-oriented extensions are enabled.
Option.Overload	Boolean	Indicates if procedure overloading is enabled.
Option.PortPinEncoding	Boolean	Indicates port pin encoding is enabled.
Option.RTCEnabled	Boolean	Indicates if the RTC is enabled for the target device.
Option.Strict	Boolean	Indicates if Strict mode is in effect.
Option.StringSize	Integer	Gives the default string size for static string allocation.
Option.TargetCode	String	The type of code being generated ("ZVM" or "Native").
Option.TargetDevice	String	The device for which code is being compiled, e.g. "ZX24a".
Option.TargetFamily	String	The family to which the target device belongs, e.g. "ATtiny", "ATmega" or "ATxmega".
Option.TargetIsZX	Boolean	Indicates if the target device is a ZX device.
Option.TargetMCU	String	Specifies the MCU used on the target device.
Option.TargetPackage	String	Specifies the package selected for a generic target device, e.g. "DIP-40". For ZX devices it is an empty string.
Option.TimerIO	String	Specifies the timer used for IO on the target device.
Option.TimerRTC	String	Specifies the timer used for the RTC on the target device.
Option.TimerIO	String	Specifies the timer used for SW UARTs on the target device.
Pin.GreenLED	Byte	Gives the pin designator for the green LED (if available).
Pin.RedLED	Byte	Gives the pin designator for the red LED (if available).
Pin.YellowLED	Byte	Gives the pin designator for the yellow LED (if available).
System.JumpBufSize	Integer	Gives the size required for a buffer used for SetJmp().
System.MinQueueSize	Integer	Gives the minimum size for an array to be used for a queue.
Version.Major	Integer	Gives the major portion of the compiler version number.
Version.Minor	Integer	Gives the minor portion of the compiler version number.
Version.Variant	Integer	Gives the variant portion of the compiler version number.
Version.Value	UnsignedLong	Gives the composite value of the compiler version number. The value corresponding to v1.2.3 is &H010203.
Version.String	String	Gives the compiler version number as a string, e.g. "1.2.3".

With regard to the built-in constants referring to the version number, consider the version number v1.2.3. In this version number, the major value is 1, the minor value is 2 and the variant is 3.

Additionally, for each target device there exists a set of built-in pin name constants of the form <port>. <bit> where <port> is a letter (case insignificant) referring to an I/O port and <bit> is a digit in the range 0-7 referring to a bit of the port. The value of the pin name is of type Byte and it may be the pin number for the target device corresponding to that port bit or it may be an encoded port/pin designator (as controlled by Option PortPinEncoding). For example, when the target device is ZX24, the pin name C.7 has the value 5 if Option PortPinEncoding is off.

3.9 Conditional Compilation Directives

The ZBasic compiler supports conditional compilation. This means that you can add conditional constructs to your code to specify that a portion of the code should be or should not be processed by the compiler. This is useful in order to create source code that can be compiled in different ways. Such flexibility can be used to create special versions of your application for different markets or for different customers, etc. It also provides a fast way to logically remove blocks of code from your program while leaving the source code intact so that it can be easily restored.

A key element of conditional compilation is the ability to define special identifiers and to give them values. These identifiers can then be used in conditional expressions that control whether or not a block of code will be processed normally by the compiler or ignored completely. The compiler supports the definition of conditional identifiers on the command line but you can also define them in your source code as well using the following syntax:

```
#define <identifier> [ [=] <expression> ]
```

Here, *<identifier>* is the name of the conditional identifier that you want to define. You may also give it a value, represented by *<expression>*, which may be an integral or string type. If you do not specify a value, a default value of 1 is used. The expression may include literals or identifiers previously defined using #define and some built-in constants (see Section 3.8).

Examples

```
#define EXPERIMENTAL
#define Version 23
#define Greeting = "Hello"
```

If you attempt to define a conditional identifier that is already defined, you will get an error message to that effect. If you want to redefine a conditional identifier you must first “undefine” the existing one using the directive:

```
#undef <identifier>
```

If the specified identifier is not actually defined, no error message will be issued so you may freely use this directive to ensure that no definition exists prior to defining a conditional identifier. Note that undefining an identifier that was defined on the command line only has effect in the current module. All other modules will see the original value.

Once you have defined your conditional identifiers, you may use them in conditional directives that are similar to If statements. The first two forms presented below are complementary.

```
#ifdef <identifier>
<other-text>
#endif

#ifndef <identifier>
<other-text>
#endif
```

The first form specifies that if the given *<identifier>* is defined, the compiler should process the text up to the matching #endif but if the *<identifier>* is not defined, the compiler should ignore the text up to the matching #endif. The second form has the opposite effect.

Example

```
#ifdef EXPERIMENTAL
    Call TestSetup(i)
#endif
```

This allows the subroutine call to be compiled into the application if EXPERIMENTAL is defined, otherwise it is left out.

As you might have already guessed, the conditional syntax also allows an #else clause.

```
#ifdef EXPERIMENTAL
    ver = "X006"
#else
    ver = "V1.4"
#endif
```

An alternate form of conditional directive allows you to specify an expression involving conditional identifiers and integer or string literals, the Boolean value of which determines whether the code within the conditional block is processed normally or not.

```
#if <expression>
<other-text>
#elseif <expression>
<other-text>
#else
<other-text>
#endif
```

The `<expression>` element may be any legal ZBasic expression involving constants, conditional identifiers, ZBasic operators and ZBasic System Library functions that can be evaluated at compile time. The usual type-compatibility rules apply, e.g., you cannot add an integral value and a string value. Additionally, the special operator `defined()` may be used in a conditional expressions. The parameter to the `defined()` operator should be an identifier and the result will be non-zero or zero depending on whether the identifier is defined or not.

Examples

```
#if defined(EXPERIMENTAL)
    Debug.Print "Experimental version"
#endif

#if defined(Pin.RedLED)
    Call PutPin(Pin.REDLED, zxOutputLow)
#endif
```

The `#elseif` clause in the conditional construction may appear zero or more times. The `#else` clause may appear at most once. The `<other-text>` element represents arbitrary text and may contain other conditional constructs. There is no practical limit on the nesting of conditionals.

Examples

```
#if Version >= 23
    #ifdef EXPERIMENTAL

        ' prepare the external circuitry and activate it
        Call TestSetup(i)
        Call PutPin(12, j)
    #endif
#endif

#if (Version >= 23) And (EXPERIMENTAL <> 0)
    ' prepare the external circuitry and activate it
    Call TestSetup(i)
    Call PutPin(12, j)
#endif
```

The two examples above have the same effect.

Conditional identifiers may be used in definitions, statements, and expressions as if they were constants defined using the `Const` definition (but the converse is not true).

```
#define Version "V1.0"
#define arraySize 26

Dim myArray(1 To arraySize) as Byte

debug.print Version
```

One implication of this is that adding a definition of a conditional identifier may result in a compiler message related to a `Const` definition warning about duplicate definitions. Conditional identifiers defined

using a compiler option are visible in all modules while those defined in a particular module are only visible in that module. Also note that conditional identifiers are essentially module-level constants. This is true even if they are defined in a procedure. A consequence of this is that in spite of whether a conditional identifier is defined inside a procedure or not, its value is visible to all subsequent conditional expressions in the module.

A useful element of a conditional expression is the function `TargetHas()`. This built-in function, what takes a String parameter and available at compile-time only, returns True or False depending of the availability for the current target of the ZBasic System Library routine named in the string.

Example

```
#if TargetHas( "OpenPWM8" )
    Debug.Print "OpwnPWM8() is available."
#endif
```

BasicX Compatibility Note

Conditional constructs are not supported in BasicX compatibility mode.

3.10 Error Directive

Sometimes it is useful to be able to purposely generate an error message in order to remind yourself of some detail or condition that requires attention. Often, this is used in conjunction with conditional directives to point out that an incompatible set of conditions exists.

The form of the error directive is shown below.

```
#error <message>
```

All of the characters beginning with the first non-white space character following `#error` up to the end-of-line will appear as the error message. If the last character on the line is an underscore and is preceded by a space or tab character, it is treated as a continued line and all of the characters on the next line up to the end-of-line will also be part of the error message.

Example

```
#if PLATFORM = "alpha"
Const Frequency As Single = 42.347
#elseif PLATFORM = "beta"
Const Frequency As Single = 45.347
#else
#error No platform specified.
#endif
```

3.11 Notice Directive

This directive is similar to the `#error` directive discussed in the previous section in that it adds a string to the error output (unless disabled by the option `--notice=off` or the directive `Option Notice Off`). However, its use does not increase the error count like `#error` does.

The form of the error directive is shown below.

```
#notice <message>
```

3.12 Include Directive

You may use an include directive in a source file to cause another source file to be compiled as well. The form of the include directive is:

```
#include "<filename>"
```

The `<filename>` element is the name of the file that you want to have compiled. If the filename is not specified using an absolute path (i.e. beginning with the root directory and/or a drive letter), the path prefix (if any) of the current module will be appended to the front of the filename. Note, however, that if an `--include-path` option is specified on the command line or a `#include_path` directive precedes the `#include` in the source code, a filename specified with a relative path will, instead, be sought among the directories specified in the include path list. See Section 10.2 for more information on the `--include-path` option. See the next section for information about the `#include_path` directive.

Note that the effect of the include directive is no different than if you had instructed the compiler directly to compile the file. It is compiled as a separate module.

BasicX Compatibility Note

Include directives are not supported in BasicX compatibility mode.

3.13 Include_path Directive

You may specify a list of directories in which to search for files to be included via `#include` in the cases where the filename is not absolute. The include path is also used when importing definitions from external C/C++ header files using the `#import` directive (see section 6.8). The include path directive has the following form.

```
#include_path "<directory_list>"
```

Note that this directive may be used as many times as needed but the path list specified most recently preceding a particular `#include` or `#import` directive will be the one in effect for that operation. Whitespace preceding and following directory names in the path list is ignored and the entire semicolon-separated list must be enclosed in double quotes. Further, either forward slash or backslash may be used as directory separators within the pathnames. If a component of the directory list ends with a slash (or backslash) followed by an at sign (@), that directory and all subordinate directories will be added to the include path.

Two special "macros" may be used in the directory list. Firstly, the character sequence `$(ZBASIC)` will be replaced by the full pathname to the subdirectory of the ZBasic installation directory containing the file `zbasic.h`. Secondly, the character sequence `$(INCPATH)` will be replaced with the include path that existed before the directive was processed. The latter macro allows directories to be added to the existing path.

Examples

```
#include_path "C:\arduino-0022\hardware\arduino\cores\arduino"  
#include_path "C:\My Projects\subdir; C:/mydir2;D:/subdir/dir3"
```

3.14 Using Conditional Directives in Project and Argument Files

The ZBasic compiler optionally supports the use of the conditional directives, as described in Section 3.9, in project and argument files. Because conditionals are introduced with a # character and that same character is recognized as a comment introduction character in project and argument files, support for conditionals in these files is not enabled by default.

You may enable support for conditionals in project and argument files using the compiler option --allow-conditionals. If this option appears as the first line in a project or argument file, conditionals will be enabled for that file and for all subsequently processed project and argument files. If the --allow-conditionals option appears on a line other than the first line, conditionals will be enabled for subsequently processed files but not for the file in which the option appeared. Alternately, conditionals may be enabled only for a specific project or argument file by placing the special comment #!allow-conditionals as the first line of the file.

When creating conditionals in project and argument files, you may utilize symbols previously defined via command line options as well as the built-in constant symbols described in Section 3.8.

3.15 Preprocessor Symbols

The ZBasic compiler supports several “preprocessor symbols” that can be used in your program code. These are special sequences of characters that the compiler recognizes very early in the compilation process. When they are recognized, they are replaced with a specific series of characters, different for each symbol. Conceptually, the replacement process is very much like a “global search and replace” operation in a text editor where context is not taken into account. The supported preprocessor symbols are shown in the table below.

Symbol	Description
<code>__DATE__</code>	The month, day and year when compiled, e.g. Oct 21, 2005.
<code>__TIME__</code>	The hour, minute and second when compiled, e.g. 14:16:04.
<code>__FILE__</code>	The name of the file being compiled.
<code>__LINE__</code>	The line number of the file on which the symbol occurs.

Although these are primarily intended as an aid to the testing procedure for the compiler, they may be useful in other circumstances. For example, you can cause the compilation date and time to be put in Program Memory using the following instructions:

```
Private Const compDate as String = "__DATE__"  
Private Const compTime as String = "__TIME__"  
Private compData as StringVectorData({ compDate, compTime })
```

During the preprocessing, the symbols above will be replaced by their respective character values. Of course, there will need to be some code that refers to `compData`, otherwise the compiler will optimize it out.

One caveat is that, as mentioned earlier, the substitution is done without regard to context. This will be a problem if you attempt to define a variable thus:

```
Dim a__LINE__b as Integer
```

If this happens to occur on line 23 of the file the compiler will see this as

```
Dim a23b as Integer
```

However, on a different line a reference to `a__LINE__b` will yield a different variable name, probably causing the compiler to complain about use of an undefined variable. Since the preprocessor symbols are case sensitive but ZBasic identifiers are not, the simple workaround is to spell the variable name differently, e.g.

```
Dim a_line_b as Integer
```

3.16 Array Data Order

RAM-based arrays are stored sequentially in memory with the first index varying the fastest and the last index varying the slowest. Consider a two-dimensional array defined as:

```
Dim ba(1 To 3, 1 To 2) as Byte
```

The bytes of this array are assigned addresses in memory sequentially as:

```
ba(1,1) ba(2,1) ba(3,1) ba(1,2) ba(2,2) ba(3,2)
```

Whether this constitutes row-major order or column-major depends on whether you consider the first index of a two-dimensional array to be the column and the second index the row or vice versa. To a large extent, this is a non-issue as long as you remember that the first index varies the fastest.

When thinking about Program Memory data tables, on the other hand, the perception does matter because the initializer data is arranged in rows and columns and you need to know how to get the array element that you want. To do so, always specify the column index first and the row index second. This order was adopted to maintain compatibility with the BasicX compiler.

3.17 Recursion in Subroutines and Functions

A subroutine or function may be invoked recursively. This means that among the statements in the routine there is one or more that invokes the same subroutine or function again, either directly or indirectly. Clearly, the recursive invocation must be conditional so that at some point the recursion ceases.

In Section 2.3.2, Defining Functions, an example function was given for computing factorial. Here is the same function written using recursion.

```
Function Factorial(ByVal val as Integer) As Integer
    If (val > 1) Then
        Factorial = Factorial(val - 1) * val
    Else
        Factorial = 1
    End If
End Function
```

Although it may look confusing at first, the logic is fairly simple. The idea is based on the fact that the value of N factorial is equal to N times the factorial of (N-1). That logic is directly expressed in the second line of code. Note that the identifier Factorial is used in two distinctly different ways in the second line of the function. On the left side of the equal sign, the identifier Factorial refers to the return value variable while that on the right side of the equal sign is a recursive invocation of the Factorial function itself. The compiler is able to distinguish the two uses by the presence of parentheses following the name. Since the return value variable can never be an array, the two types of references are easily distinguishable.

The negative aspect of recursion is that it can consume a large amount of stack space. Each time a function or subroutine invocation is performed, the processor allocates additional stack space. The extra space is used to hold some tracking information to allow the processor to return to executing the code that immediately follows the invocation. Also, stack space is required for any parameters that are passed to the subroutine/function and for any variables that are defined inside the subroutine/function. Lastly, in the case of a function only, stack space is required to hold the return value from the function. Clearly, Factorial(10000) would require more stack space than is available even if all the RAM were dedicated to the stack.

BasicX Compatibility Note

The BasicX compiler does not allow direct recursive invocation of a function but the Basic X mode of the ZBasic compiler does.

3.18 Using Default Parameter Values in Subroutines and Functions

ZBasic supports the designation of default values for parameters to subroutines and functions. This saves time when typing statements and makes the code easier to read when a particular routine is usually invoked with one or more parameters having the same constant value. Specification of the default value is done by adding an equal sign and a constant expression following the type specification of the formal parameter in the routine definition as illustrated in the example below. In the first call to the subroutine `foo`, the second parameter is omitted so the compiler automatically adds the specified default value for the second parameter.

```
Sub Main()
    Call foo(3)
    Call foo(3, 5)
End Sub

Sub foo(ByVal size as Byte, ByVal cnt as Byte = 1)
End Sub
```

If any parameter has a default value specified, all parameters following that parameter must also have a default value. Also, only parameters that are passed ByVal may have a default value. Some parameter types, like arrays and structures for example, are always passed ByRef even if they are defined as ByVal and therefore cannot have a default value specification.

3.19 Subroutine and Function Overloads

Overloading is the computer science term for defining two or more procedures having the same name but having different formal parameter lists. The compiler selects which of the overloads to invoke based on the actual parameters specified in a particular call. Several of the ZBasic System Library procedures have multiple overloads (e.g. `GetADC()`) and you may optionally define overloads in your own application code.

Normally, the compiler will disallow the definition of multiple subroutines/functions with the same name. If you wish to use this capability in your application, you must specifically enable it. This can be achieved by using the directive `Option Overload` in the first module compiled. Alternately, the compiler option `-overload` has the same effect; it must occur on the command line or in the .pjT file before the first .bas file processed. It is important to note that subroutine/function overloading is automatically enabled if ZBasic object-oriented extensions are enabled. See Chapter 4 for more information on the object-oriented extensions.

As mentioned earlier, the compiler determines which of several identically named procedures to invoke depending on the number and types of the parameters given in each particular instance. The procedure name along with the types and passing method of the parameters (if any) is referred to as the procedure's "signature". If the compiler cannot find a procedure with a signature compatible with a particular invocation, it will generate an error message to that effect. Likewise, if there is more than one procedure whose signature is compatible with a particular invocation, the compiler will issue an error message indicating an ambiguous reference and enumerating the possible matches.

It is important to note that the fact that a procedure is a subroutine or a function is not part of the signature. Moreover, a function's return type is not part of the signature, either. This means that it is not permissible to create both a function and a subroutine with the same name and parameter list nor is it permissible to create two functions that differ only in return type.

Example

```
Option Overload      ' enable overloading

Sub foo(ByVal x as Integer)
    Debug.Print "foo ByVal"
End Sub

Sub foo(ByRef x as Integer)
    Debug.Print "foo ByRef"
End Sub

Sub foo(ByVal s as String)
    Debug.Print "foo string"
End Sub

Dim i as Integer
Sub Main()
    Call foo(5)          ' this call is unambiguous
    Call foo(i)          ' this call is ambiguous
    Call foo("abc")     ' this call is unambiguous
End Sub
```

In the example above, for the first invocation of `foo()` in `Main()` the compiler determines that the first overload of `foo()` having the `ByVal` parameter is the only one whose signature is compatible with the invocation. This is not ambiguous because the constant value cannot be passed by reference so the overload of `foo()` with the `ByRef` parameter cannot possibly be used.

In contrast, for the second invocation of `foo()`, the first two overloads are both compatible resulting in an ambiguous reference.

3.20 Aliases

Occasionally, it is useful to be able to access a variable or parts of a variable as different types at different times. Although this can be accomplished by using the System Library routines `BlockMove()` or `RamPeek()`/`RamPoke()` it is simpler and more efficient to use the concept of an alias. Simply stated, defining an alias tells the compiler to generate code to access a variable or part of a variable as if it were a different type. To be clear, no new data space is allocated by defining an alias. It simply provides a different way of accessing previously defined space.

The syntax for defining an alias is similar to that for defining a variable. For example, the syntax for defining an alias at the module level is shown below.

```
{Public / Private / Dim} <name>[(<dim-list>)] As <type> Alias <var-ref>
```

As with normal variables, `Dim` has exactly the same effect as `Private`. Within a subroutine, a function or any block structure, a local alias may be defined using the syntax shown below.

```
Dim <name>[(<dim-list>)] As <type> Alias <var-ref>
```

In both cases the `<var-ref>` element is the name of a RAM variable or the name of another alias optionally including a parenthesized set of one or more constant index expressions. The parenthesized index list is only allowed, of course, if the referent item is an array.

Examples

```
Dim ival as Integer, fval as Single
Dim buf(1 to 20) as Byte

Dim b As Byte Alias fval
Dim c As Byte Alias buf(2)
Dim c2(1 to 3) As Byte Alias buf(3)
Dim bval(1 to 5) As Byte Alias ival
```

The first alias definition allows you to read/write the least significant byte of the `Single` value `fval`. The second definition allows direct access to the second byte of the `buf` variable. The third example shows how to define a sub-array within an array. The fourth example shows an alias being defined that spans more than one variable. Although the compiler allows this form its use is discouraged because the effect depends on the order in which the compiler chooses to allocate data items.

Recursive alias definitions are not allowed; an error message will be issued by the compiler when a recursive definition is detected. Note that is not allowed to define an alias that is a `String` type. You may, however, define an alias that overlays a `String` variable although this is not often useful.

One interesting use for an alias is when your application requires that a series of data items be arranged in a particular order in memory. Consider a situation where, for whatever reason, it would be convenient to have an `Integer` value, a `Byte` value and a `Single` value that are guaranteed to be arranged in sequence in memory. This can be accomplished with the definitions shown below.

```
Dim host(1 to 7) As Byte
Dim ival As Integer Alias host(1)
Dim bval As Byte Alias host(3)
Dim fval As Single Alias host(4)
```

This technique may be useful, for example, for reading and writing data packets to an external device.

One aspect of using aliases that requires careful thought and possibly some experimentation is that an alias of a fundamental type (e.g. `Byte`, `Integer`, etc.) must be defined so that it aligns on a byte boundary. If the target variable for the alias is also a fundamental type this will not be an issue because the fundamental types are always byte-aligned. On the other hand, sub-byte types may or may not be byte-aligned so defining an alias to a sub-byte type may result in a compiler error message indicating that it is not byte-aligned. See Section 3.24 for more information on this topic.

BasicX Compatibility Note

Aliases are not available in BasicX compatibility mode.

3.21 Based Variables

Based variables are a very powerful tool and their use is recommended for advanced programmers only. If used carelessly or without a complete understanding of their characteristics they may cause your program to malfunction in ways that are quite difficult to diagnose.

A based variable is similar to a procedure parameter that is passed `ByRef` in the respect that no space is allocated for the data item. Rather, the location (i.e. the address) of the based variable is specified using an integral expression that can be constant or computed at run-time. The effect that can be achieved using a based variable is similar to using an alias but a based variable is even more powerful because of the ability of the address to change at run time.

The syntax for defining a based variable at the module level is shown below.

```
{Public / Private / Dim} <name>[(<dim-list>)] As <type> Based <base-expr>
```

As with normal variables, Dim has exactly the same effect as Private. Within a subroutine, a function or any block structure, a local based variable may be defined using the syntax shown below.

```
Dim <name>[(<dim-list>)] As <type> Based <base-expr>
```

In both cases the <base-expr> element is an integral expression that gives the base address of the variable. Some examples will help clarify the concept.

```
Dim bv as Byte Based &H100
```

This defines a Byte variable whose address is a constant value.

```
Dim addr as Integer  
Dim bv as Byte Based addr
```

This defines a Byte variable whose address is given by the value of the Integer variable addr.

```
Dim addr as Integer  
Dim sel as Byte  
Dim fv as Single Based addr + CInt(sel * 3)
```

This defines a Single variable whose address is given by the value of an expression.

```
Dim addr as Integer  
Dim bv as Byte Based addr.DataAddress
```

This defines a Byte variable whose address is constant - the same as the address of the variable addr. Except for one important aspect, this has exactly the same effect as the following code.

```
Dim addr as Integer  
Dim bv as Byte Alias addr
```

The difference between an alias and a based variable is how they are handled by the compiler's optimizer. Normally, when a variable's value is referenced, the compiler's optimizer attempts to deduce the variable's value at that point and, if it is more efficient to do so, the compiler will generate code using the deduced value instead of the variable's value in memory. However, if a variable is an Alias or has an Alias that refers to it, the compiler may not attempt to make this optimization.

In contrast, the compiler does not attempt to determine if a based variable might be occupying the same space as a normal variable. Consequently, if a variable's value is changed by assigning to a based variable that occupies the same space, the compiler may generate code that is incorrect. To circumvent this potential problem, you may instruct the compiler not to make any assumptions about the value of a variable by using the Volatile attribute as in the following example.

```
Dim Volatile addr as Integer  
Dim bv as Byte Based addr.DataAddress
```

When defining an array variable that is Based, you may omit the <dim-list> element. In this case, the compiler will assume that the array is one-dimensional, that its lower bound is 1 and that its upper bound is a large value.

```
Dim addr as UnsignedInteger  
Dim ba() as Byte Based addr
```

It is also permissible to define a based Persistent Memory or Program Memory variable. This is accomplished by using the normal syntax for defining a persistent/program memory variable and appending the Based keyword together with an address expression.

```
Dim persVar as Persistent Integer Based &H300  
Dim progVar as ProgMem Single Based &H1000
```

BasicX Compatibility Note

Based variables are not available in BasicX compatibility mode.

3.22 Based Procedures

Like based variables, based procedures are a powerful tool intended to be used by advanced programmers who fully understand their nuances. No actual code space is consumed by a based procedure. Rather, declaring a based procedure simply tells the compiler how to generate an invocation of the procedure once its address is known.

The syntax for declaring a based subroutine or based function is shown below.

```
Declare Sub <name> ( [<parameter-list>] ) Based <addr-expr>  
Declare Function <name> ( [<parameter-list>] ) As <type> Based <addr-expr>
```

As with based variables, the *<addr-expr>* giving the address of the procedure to be invoked must have an integral type and can be either constant or non-constant (i.e., computed at run time). The based procedure declaration may be placed inside a subroutine or function in which case the declaration is private to that routine. At the module level, the `Declare` keyword may be proceeded by `Public` or `Private` and if neither is specified, the declaration will be public by default.

When using a based procedure, you must be very careful to be certain that the declaration matches the actual procedure that exists at the address that is specified. If the address given is not the beginning of a procedure that is the same type and has the same number and type of parameters, the result is unpredictable but will likely cause your program to malfunction.

BasicX Compatibility Note

Based procedures are not available in BasicX compatibility mode.

One issue that arises in conjunction with based procedures is that the compiler's algorithm for estimating the stack use of each procedure cannot determine how much stack space is used by the set of actual procedures that might be invoked by way of a based procedure. Consequently, the minimum task stack size will be listed as "indeterminate" for any task that invokes a procedure that, directly or indirectly, uses a based procedure.

As a solution to this problem, a special mechanism is provided for you to provide the compiler the set of procedures that might be called via a based procedure. Assuming that a complete set of called procedures is provided for each based procedure, the compiler will provide an accurate estimate of minimum task stack size. The syntax for providing call target information is shown below.

```
#pragma CallTargets ( <based-procedure> : <procedure-list> )
```

The *<procedure-list>* element is a comma-separated list of procedure names and/or names of initialized ProgMem data items. In the latter case, it is assumed that the ProgMem data item is an array of procedure addresses.

Example

```
Dim procTbl as IntegerVectorData({ @proc1, @proc2 })
Dim procAddr as UnsignedInteger
Declare Sub myProc() Based procAddr
#pragma CallTargets(myProc : procTbl, proc3)
```

If the `CallTargets` pragma occurs at the module level, it will be visible throughout the module. If it occurs within a procedure, it applies beginning at the line on which it occurs (overriding any `CallTargets` pragma at the module level for the same based procedure) through the end of the procedure or the next occurrence of a `CallTargets` pragma for the same based procedure. A special form of the `CallTargets` pragma, with no `<procedure-list>`, may be used to terminate the validity of an earlier occurring `CallTargets` pragma. This form is shown by example below.

```
#pragma CallTargets(myProc)
```

There is a special case where the compiler composes an implicit `CallTargets` list that may obviate the need for explicitly specifying the `CallTargets`. If a based procedure is defined where the base expression is a simple expression comprising an initialized `ProgMem` data item with an index, the compiler assumes that all of the elements of the `ProgMem` data item are call targets.

You may also specify a global call target list that will be visible from any module using the syntax below. The global call target list will be overridden by a non-global call target list present in a module or in a procedure.

```
#pragma GlobalCallTargets ( <based-procedure> : <procedure-list> )
```

The effect of multiple call-targets pragmas for a given procedure is cumulative.

3.23 Reference Variables

Similar to the concept of a `ByRef` parameter for a subroutine or function, you may define a variable as a reference. A reference variable is interpreted as a “pointer to” the actual variable rather than the variable itself. Consequently, a reference variable is always the same size (typically 2 bytes) irrespective of the size of the variable to which it refers.

A reference variable is defined in the same way as an ordinary variable but including the `ByRef` attribute as illustrated in the examples below. The second example illustrates a reference to an array of bytes. As with a procedure parameter that is an array, an array defined by reference is also single dimensioned with a lower bound of 1. Note that a reference variable, unlike other variables, cannot have an initialization value specified in the definition.

```
Dim fval as Single ByRef
Dim buf() as Byte ByRef
```

The compiler dereferences a reference variable when it is used just as it does with `ByRef` parameters so there are no extra syntactic elements required to perform dereferencing. This begs the question, however, of how the value of the reference gets set. When a reference variable is defined, the content of the space reserved for the reference is either uninitialized or set to zero depending on whether the definition is local to a procedure or at the module level. Since neither of these possible states is likely to refer to any useful memory, you must set the value of the reference before any use of the reference variable. Initialization of the reference is accomplished using the `DataAddress` property as illustrated below.

```
Dim f as Single
Dim fval as Single ByRef

fval.DataAddress = f.DataAddress
```

After the reference is initialized, all uses of the name of the reference variable (without the `.DataAddress` qualifier) access the memory to which the reference refers. Note that the address reference operator (@) can be used with the variable name instead of the `DataAddress` property both for setting the reference address and for obtaining the address of an ordinary variable. The last line in the preceding example could thus be rewritten with equivalent effect in any of the three following ways.

```
@fval = @f
fval.DataAddress = @f
@fval = f.DataAddress
```

The similarities of the effects of a reference variable and a based variable are likely obvious. There are, however, some important differences that may recommend the use of one over the other depending on the circumstances. In particular, it is important to note that a reference variable may be used as a member of a structure or class whereas a based variable cannot.

3.24 Sub-byte Types

In addition to the fundamental data types described in Section 2.2, ZBasic also supports `Bit` and `Nibble` data types. These are referred to as sub-byte types because they occupy less than a whole byte – 1 bit and 4 bits, respectively. In certain circumstances, these types may help reduce the total RAM usage of an application. As compared to packing and unpacking bits in your source code, using these types will be more efficient in both execution time and code space.

`Bit` and `Nibble` types may be used in most places where one of the fundamental data types may be used. You can define arrays of them, you can define `Bit` and `Nibble` constants, you can pass them as parameters and you can define functions returning these types. One limitation is that you cannot pass a `Bit` or `Nibble` variable to a subroutine/function by reference unless the variable is byte-aligned (see further discussion below). This implies, of course, that you may not be able to pass an array of sub-byte types as a parameter since arrays are always passed by reference. The reason for this limitation is that sub-byte variables do not necessarily begin on a byte boundary and there is no way for the called routine to know what the alignment might be. One way to work around this limitation is to define an integral-byte alias to the sub-byte type, pass the alias by reference and then define a new sub-byte alias in the called routine. This works because integral-byte aliases are required to be byte aligned. Note that when passed by value as a parameter, a sub-byte type parameter occupies an entire byte on the stack.

Special type conversion functions, `CBit()` and `CNibble()` are provided to facilitate the use of these types in combination with other types. See the ZBasic System Library Reference Manual for more information on the conversion functions.

In Section 3.20 Aliases, it was mentioned that there may be an issue with aliases of a fundamental type overlaying sub-byte types. The issue arises because fundamental data types must be byte-aligned and a particular `Bit` or `Nibble` variable may or may not be byte-aligned. In most cases it will be simpler to define a sub-byte alias to overlay variables of fundamental type. Doing the converse may require some experimentation to achieve the required byte-alignment. The alias defined below may or may not be accepted by the compiler depending on what other sub-byte types are defined at the same scope level and the order in which they are defined.

```
Dim ba(1 to 20) as Bit
Dim bval as Byte Alias ba(3)
```

Internally, the compiler collects together all of the sub-byte types at a given scope level (module, procedure, block) and allocates space for them collectively. The `Nibble` variables are allocated first in the order that they are defined followed by the `Bit` variables in the order that they are defined. The map file will show a special variable with a name like `@BitNib00`. This is the host variable that contains the individual `Bit` and `Nibble` values for a particular scope level.

BasicX Compatibility Note

Bit and Nibble data types are not available in BasicX compatibility mode.

3.24.1 Forcing Byte Alignment

As described earlier, sub-byte type variables are normally sub-allocated within a host variable. One consequence of this space-saving strategy is that a sub-byte type variable is often not aligned on a byte boundary, a circumstance that imposes limitations on the possible uses of the variable. One solution to this problem is to instruct the compiler to align a particular variable on a byte boundary. This is done by adding the ByteAlign attribute to the variable definition as shown below.

```
Dim ByteAlign ba(1 to 20) as Bit
```

Each variable that is so defined will occupy an integral number of bytes of memory. For the example above, three bytes will be allocated even though only 20 of the 24 available bits will actually be used.

Note that the ByteAlign attribute may be used with any variable type but it has no effect for those types that are inherently byte-aligned. Also, the ByteAlign attribute may be used in a Structure to force sub-byte types to be byte-aligned.

3.25 Structures

It is often useful to define a data type that is a collection of data elements. For example, if you write a program to manipulate dates it may be useful to define a data type that contains “year”, “month” and “day” elements. This is convenient because it allows you to think about or refer to the group of data elements as a single entity rather than as the individual constituent elements.

In ZBasic, as in many other programming languages, such a collection of data items is referred to as a “structure”. A structure is a user-defined data type similar in some respects to an enumeration. As with an enumeration, you define a structure by specifying the constituent elements. The syntax for defining a structure at the module level is:

```
[Public | Private] Structure <name>
  <member-definition>
  ...
End Structure
```

If neither `Private` nor `Public` is specified, the structure definition is public. The ellipsis (...) in the syntax above connotes that there may be zero or more additional member definitions. A structure definition must have least one member and may have a practically unlimited number of members.

A structure may be defined within a subroutine or function, either at the outer level or within any inner block. In this case, the `Public` and `Private` keywords have no useful purpose and are therefore not allowed.

After a structure has been defined, the structure name may be used as a `<type>` in a variable or structure definition. A structure may also be used as the `<type>` in the formal parameter list of a subroutine or function definition. Note, however, that a Public subroutine or function cannot be defined with a parameter that is a Private Structure.

A `<member-definition>` has the same syntax as that used to define a variable. As with an ordinary variable, a member may be a single data element or it may be an array. The syntax for a member definition is given by the two descriptions below – the first being for a non-array member and the second being for a member that is an array.

```
{Public / Private / Dim} <name> As <type>
or
{Public / Private / Dim} <name>(<dim-spec-list>) As <type>
```

As with ordinary variables, Dim has exactly the same effect as Private, i.e., the member will only be directly accessible to code within the module. In contrast, a Public member may be accessed by code outside of the module in which the structure is defined. The names of the members of each structure defined may be any legal identifier however a particular name may be used only once in each structure. The use of a name as a member in one structure does not preclude it from also being used as a member name in a different structure or as a variable, constant, parameter, etc. This circumstance is a result of the ZBasic scoping rules - a structure definition is an independent name scope.

The <type> specified for a member may be any of the pre-defined types like Integer, Byte, String, etc. or it may be a user-defined type like an enumeration or another structure. The amount of space required for each member in a structure is the same as for a variable of the same type (but see the discussion below relating to sub-byte types and alignment). It is important to note that recursive structure definitions, with members that are or contain (directly or indirectly) the structure being defined, are not allowed.

Examples

```
Structure MyDate
    Public year as UnsignedInteger
    Public month as Byte
    Public day as Byte
End Structure

Structure MyTime
    Public hour as Byte
    Public minute as Byte
    Public seconds as Single
End Structure

Structure MyTimeStamp
    Public tdate as MyDate
    Public ttime as MyTime
    Private isCurrent as Boolean
End Structure

Sub Main()
    Dim ts as MyTimeStamp

    Call GetTimeStamp(ts.tdate.year, ts.tdate.month, ts.tdate.day, _
                      ts.ttime.hour, ts.ttime.minute, ts.ttime.seconds)
End Sub
```

The example above illustrates how members of an instance of a structure are referenced. The variable ts is an instance of the MyTimeStamp structure. A member is referenced by appending the member name to the variable name, separating them with a period. For members that are structures, members of the subordinate structure are referred to similarly. No spaces are allowed in this construction. In cases where a variable or a member is an array, the index list directly follows the variable/member name, preceding the period.

```
Structure foo
    Dim b as Byte
    Dim ai(1 to 10) as Integer
    Dim ts(1 to 4) as MyTimeStamp
End Structure
```

```

Dim yr as UnsignedInteger
Dim i as Integer
Dim f as foo
Dim b as Byte

i = f.ai(b)
yr = f.ts(3).tdate.year

```

The address of a member of a variable that is a structure may be obtained by appending the `.DataAddress` property identifier to the reference or by using the `MemAddress()` function.

```

Dim addr as UnsignedInteger

addr = ts.tdate.DataAddress
addr = MemAddressU(ts.tdate)

```

Structures may be used in Alias and Based variable definitions. However, in these cases the structures may not contain members that are any of the string types. Structures may be passed to subroutines and functions either by reference or by value. If a structure is passed by value, the structure will be read-only within the procedure.

A variable that is a structure may be assigned to another variable that is the same type of structure using the standard assignment operator. In most cases, the compiler produces code to implement the copy operation using a simple memory-to-memory copy operation. However, in cases where the structure contains one or more members that are of the allocated string type (or has members that are structures which have such members), the compiler uses a special copy procedure that ensures the correct result.

Moreover, two variables that are the same type of structure may be compared for equality or inequality using the standard comparison operators, `=` and `<>`. The equality/inequality test is implemented using a byte-by-byte comparison of the content of two structures. If one or more members of the structure are the `BoundedString` type, the byte-by-byte comparison may result in a `False` value even though the strings are identical. This is because the currently-unused portion of the string store may contain byte values that are different between the two instances being compared. Similarly, comparison of structures containing allocated strings, while allowed, is not recommended because of the likelihood of resulting in false negatives.

Structures may contain members that are sub-byte types, `Bit` and `Nibble`. Members that are `Bit` type will be aligned on the next available bit boundary and those that are `Nibble` type will be aligned on the next available nibble boundary. If a member that is not a sub-byte type follows a sub-byte type member, that non-sub-byte member will be aligned on the next available byte boundary. Depending on how you define your structure, this may result in "holes" in the structure layout representing unused bits. The unused bits are generally of no consequence but it is important to note that these "holes" may interfere with the equality/inequality test for structures that are not initialized because the unused bits will have an indeterminate value. You can avoid this potential problem by performing a block initialization on locally defined structures.

Example

```

' this structure contains a "hole", unused bits following the ab member
Structure foo
    Dim b as Byte
    Dim ab(1 to 4) as Bit
    Dim b2 as Byte
End Structure

```

```

Sub Main()
    ' this structure is not automatically initialized
    Dim bar as foo

    ' this call zeroes out the entire structure
    Call MemSet(bar.DataAddress, SizeOf(bar), 0)

    ' other code follows
End Sub

```

Another method of creating a user-defined type, compatible with VB6, is also supported. The syntax for a Type definition is:

```

[Public | Private] Type <name>
    <member-definition>
    ...
End Type

```

In this case, the *<member-definition>* is the same as for defining a member of a Structure except that the visibility attribute (Public, Private, or Dim) is not allowed. The visibility of each member will be the same as the visibility of the Type itself.

BasicX Compatibility Note

Structure and Type definitions are not allowed in BasicX compatibility mode.

3.25.1 Structures in Persistent Memory and Program Memory

Structures in Persistent Memory and Program Memory may be defined using the syntax:

```
{Public / Private / Dim} <name> As Persistent <type>
```

or

```
{Public / Private / Dim} <name> As ProgMem <type>
```

where *<type>* is the name of a previously defined structure. The allowable members of a Persistent Memory or Program Memory structure are the intrinsic types Byte, Integer, UnsignedInteger, Long, UnsignedLong, Single, bounded strings, arrays of those types and other structures containing only those types. Arrays of structures in Persistent Memory or Program Memory may be defined by specifying the dimensions in the usual way.

It is permissible to directly assign between any combination RAM-based, Persistent Memory and Program Memory variables defined using the same structure. Similarly, direct comparison between like structures (equality and inequality only) is supported.

3.26 Unions

It is occasionally useful to define a data type that is a collection of data elements superimposed on one another, i.e. where all members occupy the same space. This yields an effect similar to defining a variable to be an alias for another. In ZBasic, as in many other programming languages, such a collection of data items is referred to as a “union”. A union is a user-defined data type similar in some respects to a structure. As with a structure, you define a union by specifying the constituent elements. The syntax for defining a union at the module level is:

```
[Public | Private] Union <name>
    <member-definition>
    ...
End Union
```

If neither `Private` nor `Public` is specified, the union definition is public. The ellipsis (...) in the syntax above connotes that there may be zero or more additional member definitions. A union definition must have least one member and may have a practically unlimited number of members.

A union may be defined within a subroutine or function, either at the outer level or within any inner block. In this case, the `Public` and `Private` keywords have no useful purpose and are therefore not allowed.

Examples

```
Union FloatBytes
    Public byteData(1 to 4) as Byte
    Public floatData as Single
End Union
```

```
Union MyData
    Public b as Byte
    Public u as UnsignedInteger
End Union
```

A variable defined as a `FloatBytes` type would occupy four bytes of memory since the largest member is four bytes. Similarly, a variable defined as a `MyData` type would occupy two bytes of memory.

3.27 Using Namespaces

In computer science, a namespace is a context within which all identifiers must be unique. Earlier in this document it was mentioned that a variable with a particular name, e.g. `myVar`, may be defined at the module level and also defined inside a procedure. This is possible because the module level represents one context and each procedure represents a separate context. Sometimes it is useful to define one or more additional contexts at the module level in which entities such as constants, variables, procedures, etc. may be defined without needing to worry about duplication of names defined elsewhere. That is the purpose of a ZBasic namespace.

By default, namespace support is not enabled. The compiler directive `Option Namespaces` or the command line option `--namespaces` must be used to enable namespace support (note, however, that namespace support is enabled automatically when ZBasic object support is enabled). A namespace is defined using the syntax below, essentially being a “wrapper” around the definitions of the entities defined within it.

```
[Public / Private] Namespace <name>
<definitions of variables, procedures, etc.>
End Namespace
```

If neither `Public` nor `Private` is specified, the namespace is private by default. It is important to note that all entities defined within a private namespace are forced to be private as well; warnings are emitted for entities defined with a `Public` modifier inside a private namespace. In contrast, private entities within a public namespace are allowed.

In order to refer to entities defined within a namespace, it is necessary to specify both the namespace and the entity name. Consider the example below.

```
Namespace foo
    Dim a as Integer
    Sub SetVal(ByVal v as Integer)
        a = v
    End Sub
End Namespace
```

```

    End Sub
End Namespace

Sub Main()
    Debug.Print foo::a
    Call foo::SetVal(10)
    Debug.Print foo::a
End Sub

```

Outside the namespace, references to entities defined in the namespace must be preceded by the namespace prefix (the namespace name and a double colon) as seen in the `Main()` subroutine. However, within the namespace itself the namespace prefix may be omitted.

It is important to note that the `foo` namespace could also have been written as shown below.

```

Namespace foo
    Dim a as Integer
End Namespace

Namespace foo
    Sub SetVal(ByVal v as Integer)
        a = v
    End Sub
End Namespace

```

Multiple occurrences of the same namespace are effectively aggregated into a single namespace definition. Note, however, that the compiler will issue an error if the visibility of the namespace (i.e. Public vs. Private) is not the same in all cases.

Although rarely necessary, it is permissible to place a namespace definition within a namespace definition, such nesting allowed to an arbitrary depth. In such a case, namespace prefixes must be added to identifiers beginning at the outer level and moving inward until the identifier itself is reached. For example, `foo::bar::myVar` would be used (external to namespace `foo`) to refer to the variable `myVar` defined within the namespace `bar` itself defined within the namespace `foo`.

One convenient use of a namespace is to include code in your application that was written by another party where some of the identifiers used conflict with existing identifiers in your application. You can wrap the other party's code in a namespace and eliminate identifier conflicts without otherwise changing the code at the expense of having to use the namespace prefix to access the identifiers contained therein.

3.28 Data Type Implementation Information

This section provides more information on the technical details of the fundamental data types and variants. Although this information is generally not needed to write properly functioning programs it is provided for those who are interested and for the special cases where knowledge of such implementation details may help you implement your application.

To review, the table of fundamental data types introduced earlier in this manual is reproduced here with an additional column indicating the amount of space required for each type.

Fundamental Data Types		
Data Type Name	Range of Values	Size in Bytes
Boolean	True, False	1
Byte	0 to 255	1
Integer	-32,768 to 32,767	2
UnsignedInteger	0 to 65,535	2
Long	-2,147,483,648 to 2,147,483,647	4
UnsignedLong	0 to 4,294,967,295	4
Single	approx. ±1.4e-45 to ±3.4e+38 and 0.0	4
String	0 to 255 characters	See 3.28.2

The Boolean type, while occupying an entire byte, will always contain one of two values under normal circumstances. The value True is represented by the value 255 and the value False by 0. If your application has a need for a lot of Boolean variables it may be more efficient to use the Bit data type described in Section 3.21. The primary disadvantage to using the Bit type is that it cannot be passed by reference.

The Single type is implemented using the data format specified in the IEEE 754 standard for single precision floating point numbers.

3.28.1 User-defined Type Details

The user-defined types in ZBasic are enumerations and structures. Enumerations are implemented using a two-byte value to represent the enumeration member value. Structures are laid out in memory exactly corresponding to the order in which the members are defined in the structure definition. Bit and Nibble type members are aligned on bit and nibble boundaries, respectively. All other member types are aligned on a byte boundary. Due to these alignment rules and depending on the specific structure definition, there may be unused bits and/or nibbles within the structure. The number of bytes consumed by a structure variable is the sum of the sizes of the members including the unused bits.

3.28.2 String Data Type Details

For the string data types, the storage requirements and implementation details vary depending on compiler command line options, Option Directives and the type of string. For the String data type, if Option AllocStr is enabled (as it is by default for ZBasic modules) each string variable requires four bytes of storage in User RAM space plus additional space allocated from the dynamic memory allocation heap to hold the characters of the string. This string storage strategy is called dynamic string allocation because the space to hold the string's characters is dynamically allocated and will grow and shrink as needed to accommodate the string value assigned to it. One additional advantage to using dynamically allocated strings is that the string storage will never be overrun because it is changed dynamically and limited to the maximum size of 255 characters automatically.

The 4 bytes of User RAM for a dynamically allocated string variable are used in the following manner:

Dynamically Allocated String Storage Layout

Offset	Length	Description
0	1	The current string length, in bytes.
1	1	A marker identifying the string location: &H0 – RAM (allocated from the heap) &H2 – Program Memory &H3 – Persistent Memory &H4 – RAM (not allocated) &H5 – RAM (not allocated, limited to 2 chars max.) &H6 – RAM (pointer refers to a statically allocated string)
2	2	For type &H5, the one or two characters of the string. For type &H6, the address of a statically allocated string variable (see below). For all other types, the address of the first character of the string.

For statically allocated strings, including Bounded Strings and Fixed Length strings, the space allocated for the variable is used in the following manner:

Statically Allocated String Storage Layout

Offset	Length	Description
0	1	The current string length, in bytes.
1	1	A marker identifying the string characteristics: &H00 – fixed allocation, variable length &HFF – fixed allocation, fixed length
2	N	The characters of the string where N is the defined string length.

For Fixed Length strings the byte at offset zero will be constant and the byte at offset 1 will have the value &HFF to indicate that it is a Fixed Length string.

Fixed-Length String Storage Layout

Offset	Length	Description
0	1	The string length, in bytes.
1	1	A marker identifying the string characteristics: &HFF – fixed allocation, fixed length
2	N	The characters of the string where N is the fixed string length.

3.28.3 String Address and String Type

Generally, you needn't be concerned about the technical details regarding strings described in the preceding section. For certain special situations, however, it may be useful directly access a string variable's inner data components. The `StrAddress()` function works with all of the string types described above and will return the address of the first character of the string storage. Note, however, that for dynamically allocated strings if the string's length is zero the returned address will be of no use (it will generally be zero). Moreover, the returned address may be an address in RAM, in Program Memory or in Persistent Memory. Depending on where the string's characters are stored you must use different System Library Functions to retrieve the string's characters. You can use the System Library function `StrType()` to determine the nature of the string's storage. It will return the second byte of the string variable's storage space. See the tables above for the values and their meaning.

BasicX Compatibility Note

In BasicX compatibility mode, neither `StrAddress()` nor `StrType()` is available.

3.29 Controlling the Heap Size and Main() Task Stack Size

The diagram below illustrates how RAM is allocated between defined variables, task stacks and the heap. The static allocation area, located at the beginning of RAM, includes all module-level variables defined by your application and all statically allocated system variables. This includes the task stacks that are statically allocated by your application.

For native mode devices, the task stack for Main() occupies all of the RAM between the end of the static allocation area and the heap limit. The task stack grows downward from the heap limit while the heap grows downward from the end of RAM. For native mode ZX devices, the heap limit is set at 512 bytes from the end of RAM by default. For generic ZBasic devices, the default heap size depends on a number of factors including RAM size and Program Memory page size.

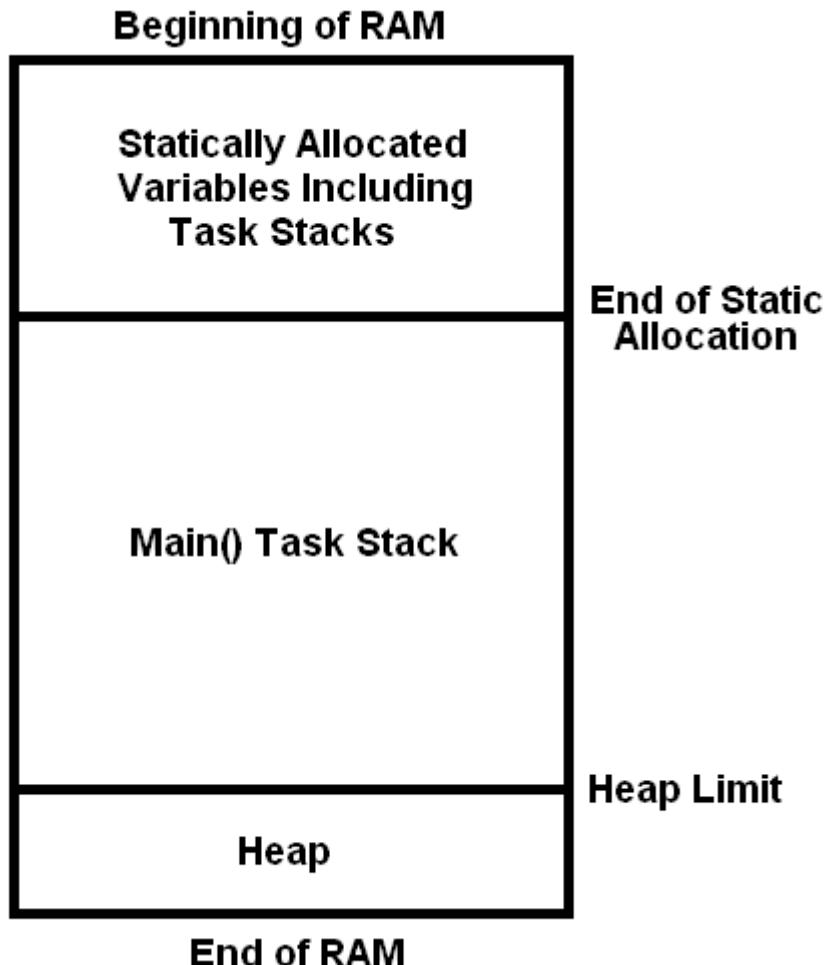
For VM mode devices, the space between the end of the static allocation area and the end of RAM is shared by default between the Main() task stack and the heap. The task stack grows upward from the end of the static allocation area while the heap grows downward from the end of RAM. There is no specific heap limit in the default case but you may specify one as described below.

If a heap limit exists, the heap will not grow beyond that point. When the heap is exhausted, further allocation attempts will fail, possibly resulting in malfunctioning of your application. The heap is used for storing most RAM-based strings (but not BoundedString or fixed-length string types) including strings that are returned by functions you define in your application and System Library functions that return strings. Additionally, for devices that use internal Flash memory for Program Memory (including all generic devices), when a Program Memory write operation is performed a temporary buffer is allocated from the heap. The size of this buffer is equal to the Flash page size for the microcontroller, typically 256 bytes. This additional temporary use of heap space must be considered if your application performs write operations on Program Memory. Lastly, the System Library function `System.Alloc()` may be used by your application to allocate memory from the heap. All of these uses must be considered when determining how much heap space is needed by your application. The function `System.HeapHeadRoom()` may be helpful for determining the minimum heap size for your application.

You can specify the heap limit setting in several different ways, each of which may be useful in different circumstances. Perhaps the simplest method of setting the heap limit is to specify the desired size for the Main() task stack. This can be done using the directive `Option MainTaskStackSize` in your Main module or by using the compiler option `--main-task-stack-size`. In both cases, the heap limit is set by adding the specified size to the address of the end of the static allocation area.

The second method of setting the heap limit is to specify the desired size of the heap using the directive `Option HeapSize` in your Main module or by using the compiler option `--heap-size`. The heap limit is set by subtracting the specified value from the address one past the end of RAM. For devices that have external RAM, the end of RAM is determined at startup. If you have external RAM, you may specify that all of the external RAM should be used for the heap by specifying the heap size using the special values 65535 or &HFFFF.

The final method for setting the heap limit is to specify an absolute address using the directive `Option HeapLimit` in your Main module or by using the compiler option `--heap-limit`. This method is recommended for use only by advanced programmers who understand completely how memory is allocated in a ZBasic application.



3.30 Task Management

The ZBasic System Library has several routines that are helpful for managing tasks including LockTask(), UnlockTask(), StatusTask(), ResumeTask(), RunTask() and ExitTask(). The latter four routines permit some advanced task management for special situations. One use of these routines is to implement a timeout on a task that is awaiting the completion of an event, e.g. an external interrupt or an input capture. Normally, when a task is set to wait for an event like these it will wait indefinitely.

3.30.1 Task Control Block

Each task has an associated Task Control Block (TCB) - a data area that occupies the first few bytes of the task stack in VM mode (the last few bytes of the task stack in native mode). In most cases, a program will not need to access the contents of a TCB. The information is included here for those rare circumstances when is needed. Use of this information is recommended for advanced programmers only.

When a task is activated its task control block is initialized and then inserted into a circular linked list immediately following the task control block of the then-current task. The table below gives some information on the structure of the task control block. It is important to note that this information is considered implementation detail subject to change as necessary.

Task Control Block Elements (VM Mode Devices)

Offset	Length	Description
0	1	Task status. See StatusTask() for details.
1	2	Remaining time to sleep (in RTC ticks).
3	2	Address of next task control block.
5	6	Task context: IP, BP, SP (valid only when not the current task).
11	1	Task control flags (used internally).
12	2	Address of the byte following the end of the task's stack.

Task Control Block Elements (Native mode Devices)

Offset	Length	Description
0	1	Task status. See StatusTask() for details.
1	2	Remaining time to sleep (in RTC ticks).
3	2	Address of next task control block.
5	2	Task context: SP (valid only when not the current task).
7	2	Task stack starting address.
9	2	Unused.
11	1	Task control flags (used internally).

Caution: directly modifying the task control block (other than the “time to sleep” value) is strongly discouraged as doing so will probably cause your program to malfunction.

3.31 Dynamic Memory Allocation

The ZBasic system maintains a dynamic memory allocation heap that is primarily used internally. For example, space is allocated from the heap automatically by the system to hold the characters of String variables. That allocated space is automatically returned to the heap when String variables change or go out of scope.

In some applications, it is useful to be able to allocate a block of memory to use for some period of time (perhaps for a data buffer, for example) and then to deallocate the block when it is no longer needed. The ZBasic System Library routines System.Alloc() and System.Free() provide access to the dynamic memory allocation heap for such purposes. While these routines provide a very useful functionality, they must be used with great care because careless or improper use can result in malfunctioning of the heap. There are two fundamental problems that can arise: heap exhaustion and heap corruption.

If blocks of memory are allocated from the heap and never freed, the heap will eventually become exhausted and subsequent allocation requests will fail. Since String variable types rely on proper heap function, assigning values to String variables will not have the desired effect when the heap is exhausted.

Heap corruption occurs when the heap management data structures are inadvertently modified. Such corruption can result from writing to a previously allocated block after it has been freed or writing to memory outside the bounds of a properly allocated but not yet freed block. The heap may also be corrupted by inadvertent writes to RAM within the bounds of the entire heap such as might result from the careless use of the RamPoke() or BlockMove() routines, by writing beyond the bounds of an array, or overflowing the stack assigned to a task. Heap corruption will also be the likely result of passing an invalid value to System.Free(). The only values that may be safely passed to System.Free() are those that have been returned by System.Alloc() but have not yet been passed to SystemFree(). As a special case, passing the value zero to SystemFree() is benign.

The value Register.HeapEnd, described in this document, may be useful for monitoring the state of the heap. For additional information, see the descriptions of System.Alloc() and System.Free() in the ZBasic System Library Reference manual. Also, for native mode devices (e.g. the ZX-24n) the function System.HeapHeadRoom() may be used to determine, at any time, the amount of unused space that remains in the heap. Moreover, for native mode devices the directive Option HeapSize can be used to specify an upper limit on the size of the heap thereby preventing it from encroaching on the task stacks.

3.32 Exception Handling

ZBasic implements a simple but effective form of exception handling using a concept borrowed from the C language. In the normal execution of a program, the call-return process is orderly and rigid. A routine can only return directly to the routine that called it. This forms a natural hierarchy that works well in most cases. However, it is sometimes the case that your program will detect a set of circumstances that preclude normal operation. In such cases, it is useful to be able to discard the normal call-return hierarchy and return directly to some distant caller, sending back a value to indicate the nature of the conditions that required the exceptional action.

The simple example below shows how to utilize the exception handling mechanism.

```
Dim jmpBuf1(1 to System.JumpBufSize) as Byte

Sub Main()
    debug.print "start test"

    ' initialize the jump buffer
    Select Case SetJmp(jmpBuf1)
    Case 0
        ' control came back from SetJmp()
        debug.print "calling foo()"
        Call foo()
        debug.print "normal return from foo()"
    Case 1
        ' control came back from LongJmp()
        debug.print "SetJmp() returned 1 via LongJmp()"
    End Select

    debug.print "end test"
End Sub

Sub foo()
    debug.print "in foo()"
    Call bar()
    debug.print "returning from foo()"
End Sub

Sub bar()
    debug.print "in bar()"
    Call LongJmp(jmpBuf1, 1)
    debug.print "returning from bar()"
End Sub
```

Running this program will result in the following debug output:

```
start test
calling foo()
in foo()
in bar()
SetJmp() returned 1 via LongJmp()
end test
```

Note that the last lines of neither `foo()` nor `bar()` were executed nor was the line after the call to `foo()` in `Main()`. If you comment out the call to `LongJmp()` in `bar()`, the program will produce the following debug output:

```
start test
calling foo()
in foo()
```

```

in bar()
returning from bar()
returning from foo()
normal return from foo()
end test

```

Note that a call to `LongJmp()` generally should not utilize a value of zero as the second parameter. Doing so will make it appear as though the original `SetJmp()` call is returning.

The jump buffer may also be a local variable if desired. The example code below is a modified version of the previous example showing how the jump buffer is passed down the hierarchy as a parameter. This technique may be used to create generalized subroutines that might return to one of several places depending on how it was called.

```

Sub Main()
    Dim jmpBuf1(1 To System.JumpBufSize) as Byte

    debug.print "start test"

    ' initialize the jump buffer
    Select Case SetJmp(jmpBuf1)
    Case 0
        ' control came back from SetJmp()
        debug.print "calling foo()"
        Call foo(jmpBuf1)
        debug.print "normal return from foo()"
    Case 1
        ' control came back from LongJmp()
        debug.print "SetJmp() returned 1 via LongJmp()"
    End Select

    debug.print "end test"
End Sub

Sub foo(ByRef jb() as Byte)
    debug.print "in foo()"
    Call bar(jb)
    debug.print "returning from foo()"
End Sub

Sub bar(ByRef jb() as Byte)
    debug.print "in bar()"
    Call LongJmp(jb, 1)
    debug.print "returning from bar()"
End Sub

```

3.33 Run Time Stack Checking

On VM mode devices (e.g. the ZX-24a), the control program on the ZX implements optional run-time stack overflow detection. You can enable and disable the checking at any time using the System Library subroutine `StackCheck()`. There is a small performance penalty for the stack check although the checking is only done after operations that add data to the stack.

When stack checking is enabled the stack pointer is compared against the “end of stack” value in the current task’s Task Control Block less the current setting of `Register.StackMargin`. If the stack pointer exceeds the limit a stack fault is generated and the processor is reset. When the processor begins running again `Register.ResetFlags` will indicate that a WatchDog reset occurred and `Register.FaultType` will have the value 1 indicating a stack fault condition occurred. Also, `Register.FaultData` will contain the address of the Task Control Block of the task that was running when the stack fault was detected. `Register.FaultData2` will contain the address of the instruction that was executing at the time of the fault. Note that `Register.FaultType` is a persistent value. If you

add code to your application to respond to the fault condition you'll want to reset Register.FaultType to zero after responding to it. Register.FaultData and Register.FaultData2 are also persistent but it will usually not be necessary to reset its value.

Chapter 4 - ZBasic Object-Oriented Extensions

The purpose of the object-oriented extensions for ZBasic is to bring the benefits of object-oriented programming - inheritance, polymorphism, and encapsulation - to ZBasic users. Although ZBasic historically has been derived from Microsoft Visual Basic (VB6), the decision was made not to implement the VB object-oriented features in ZBasic largely due to its weak implementation of the inheritance model. Since inheritance is one of the primary benefits of object oriented programming, implementing a weak version of it would do a great disservice to ZBasic users, especially those who are getting their first exposure to object-oriented programming. Instead, the realization of ZBasic object-oriented features takes ideas from other, more popular, object-oriented languages (e.g. C++ and Java) and implements them in ways that fit well in the ZBasic language structure.

The purpose of this chapter is to briefly describe the important aspects of the ZBasic object-oriented extensions. The intended audience is those who have a basic understanding of object-oriented programming principles. Those who do not meet this qualification should utilize the many print and Internet-based resources on the topic before continuing.

4.1 Enabling Object-Oriented Extensions

In order to use the ZBasic object-oriented extensions, you must enable them. This strategy was chosen so that the additional keywords needed by the extensions do not cause existing programs to fail to compile. Enabling the extensions can be achieved by using the directive `Option Objects` in the first module compiled. Alternately, the compiler option `--objects` has the same effect; it must occur on the command line or in the .pjT file before the first .bas file processed. One side effect of enabling object extensions is that it also allows overloaded subroutines and functions to be defined outside of a class. See section 3.19 for more information on procedure overloading.

4.2 Defining a Class

The characteristics of an object are described by defining a `Class`. In many ways, defining a class is very similar to defining a structure in ZBasic. However, instead of having members that are data items, a class generally will also have members that are subroutines or functions. As with a structure, the members of a class may have either `Public` or `Private` visibility attributes. A third visibility attribute, `Protected`, affords an additional dimension of access control (described later) that is useful when one class inherits characteristics from another class.

Similar to defining a structure, a `Class` is defined using the `Class...End Class` sequence. You may define as many classes in a module as you wish. A private class cannot be used by code in other modules while a public class can be.

```
[Public / Private] Class <name>
    <member-definition>
    ...
End Class
```

The ellipsis in the syntax above connotes that there may be zero or more additional member definitions. As with structures, if neither `Public` nor `Private` is specified, the class will be public, meaning that it will be visible outside of the module in which it is defined. After a class is defined, its name may be used as the type in a variable definition. Consider this simple example.

```
Public Class myClass
    Dim i as Integer
    Dim s as String
End Class
```

A `<member-definition>` that describes a data member has the same syntax as that used to define a variable. As with an ordinary variable, a member may be a single data element or it may be an array.

The syntax for a member definition is given by the two descriptions below – the first being for a non-array member and the second being for a member that is an array.

```
{Public / Private / Dim} <name> As <type>
```

or

```
{Public / Private / Dim} <name>(<dim-spec-list>) As <type>
```

As with ordinary variables, Dim has exactly the same effect as Private. The names of the members of each class may be any legal identifier but a particular name may be used only once in each class. The use of a name as a member in one class does not preclude it from also being used as a member name in a different class or structure or as a variable, constant, parameter, etc.

The <type> specified for a member may be any of the pre-defined types like Integer, Byte, String, etc. or it may be a user-defined type like an enumeration, a structure or another class. It is important to note that recursive class definitions, with members that are or contain (directly or indirectly) an object of the class being defined, are not allowed. Defining a class that contains one or more members that are of different class types, known as *composition*, is an important aspect of object-oriented programming.

Examples

```
Class MyDate
    Public year as UnsignedInteger
    Public month as Byte
    Public day as Byte
End Class

Class MyTime
    Public hour as Byte
    Public minute as Byte
    Public seconds as Single
End Class

Class MyTimeStamp
    Public tdate as MyDate
    Public ttime as MyTime
    Private isCurrent as Boolean
End Class

Sub Main()
    Dim ts as MyTimeStamp

    Call GetTimeStamp(ts.tdate.year, ts.tdate.month, ts.tdate.day, _
                     ts.ttime.hour, ts.ttime.minute, ts.ttime.seconds)
End Sub
```

The example above emphasizes the similarities between structures and classes; except for the different keyword used, the syntax is identical. As with structures, instances of classes (called objects) can be passed to subroutines and functions either by reference or by value. If an object is passed by value, it will be read-only within the receiving procedure.

A variable that is an instance of a class may be assigned to another variable that is an instance of the same class using the standard assignment operator. If the class has members that are an allocated string type or other objects, special, automatically generated “copy procedures” are invoked to perform the copying. This achieves the desired result in most cases, but there are certain special cases where the compiler cannot produce a logically correct copy using this strategy. For such cases, there is a way to specify how the object should be copied. See the discussion about assignment constructors in Section 4.6.

Two variables that are the same type of object may be compared for equality or inequality using the standard comparison operators, = and <>. As with structures, the equality/inequality test is implemented using a byte-by-byte comparison of the content of the two objects. If one or more members of the object are the BoundedString type, the byte-by-byte comparison may result in a False value even though the strings are identical. This is because the currently-unused portion of the string store may contain byte values that are different between the two instances being compared. Similarly, comparison of objects containing allocated strings, while allowed, is not recommended because of the likelihood of producing false negative results.

4.3 Defining Class Methods

So far, the discussion of classes has been limited to the similarities between structures and classes. However, classes, and instances of classes (objects), are much more powerful than structures. Most of the additional capability results from the ability to associate subroutines and functions with the class. In the parlance of object-oriented programming, a subroutine or function that is part of a class definition is known as a method. We will generally refer to the class' subroutines and functions as methods unless the context requires a clear distinction between a method that is a subroutine and one that happens to be a function.

A method is defined as being part of a class simply by including the desired subroutine/function definition within the class definition, right along with all of the class data members. Consider this simple example.

```
Class MyTime
    Private hour as Byte
    Private minute as Byte
    Private seconds as Single

    Public Sub SetTime(ByVal h as Byte, ByVal m as Byte, ByVal s as Single)
        hour = h
        minute = m
        Seconds = s
    End Sub
End Class
```

This simple class has three data members and one one method, whose ostensible purpose is to set the data members of the object to specific values. A class method is invoked the same way as any "normal" subroutine or function except that the name of the object to be acted upon must also be specified. The example below, which utilizes the class defined above, illustrates how this is done.

```
Dim t as MyTime
Sub Main()
    Call t.SetTime(12, 0, 0.0)
End Sub
```

In this example, the variable t is an instance of the class MyTime. In the parlance of object-oriented programming, t is said to be an object of type MyTime.

So far, the example class isn't particularly useful. It has a method for setting the data members of the object but, because the data members are defined to be `Private`, they can't be accessed by procedures other than class methods. Making data members private is an important part of the concepts of information hiding and encapsulation that are central to object-oriented programming. The general idea is that you make most (preferably, all) data members private and then provide accessor methods to return data values from the object to the caller and also provide *mutator* methods to allow a caller to modify an object's data members indirectly.

The `SetTime()` method above is a mutator method that happens to provide the ability to modify all of the data members in a single call. Depending on the nature and complexity of the classes that you define, it may be preferable to define mutator methods that modify individual data members, subsets of the class' data members, or all data members. One of the advantages of providing mutator methods for setting data member values is that the mutator method can enforce limits on values of the data members or

relationships between values of multiple data members. Such control would be difficult or impossible to enforce if outside entities were allowed to directly modify the data members. Moreover, requiring the use of mutator methods to change data members also allows the freedom to later change the way that the data members are stored without affecting any of the object's users. An example of enforcing data value ranges is shown below in the re-written `SetTime()` method. The class definition has also been augmented with accessor methods for the three data members.

```
Class MyTime
    Private hour as Byte
    Private minute as Byte
    Private seconds as Single

    Public Sub SetTime(ByVal h as Byte, ByVal m as Byte, ByVal s as Single)
        hour = Min(h, 23)
        minute = Min(m, 59)
        Seconds = Min(s, 59.999)
    End Sub

    Public Function GetHour() as Byte
        GetHour = hour
    End Function

    Public Function GetMinute() as Byte
        GetMinute = minute
    End Function

    Public Function GetSeconds() as Single
        GetSeconds = seconds
    End Function
End Class
```

4.4 Object Creation Issues

When an object is created, the data members have initial values that depend on where the object is defined. This is exactly the same situation as with non-object variables. For example, if an `Integer` variable is defined at the module level, its value is guaranteed to be initially zero. In contrast, if an `Integer` variable is defined within a subroutine or function (without the `Static` attribute), its initial value is indeterminate. `String` variables (both the allocated type and `BoundedString` type) have slightly different rules in that they are guaranteed initially to have zero-length string values no matter where the variable is defined. Similarly, fixed-length strings are guaranteed to have an initial value of a space-filled string with the specified length. These rules apply to the data members of objects (just as they do with structures), again depending on where the object is defined.

Beyond these fundamental initialization guarantees, when you define a class you can also define optional `constructor` methods that initialize the values of some or all of the data members. A constructor method is a subroutine that has the special name `_Create`. You may define several different constructor methods, all having the name `_Create`, but each of the constructor overloads must have different formal parameter lists so that the compiler can distinguish between them.

The expanded `MyTime` class definition below includes several constructor methods and the additional code shows how a particular constructor can be specified when an object is instantiated. It is important to note that when a particular constructor executes, the fundamental initialization (described above) will already have been performed.

```
Class MyTime
    Private hour as Byte
    Private minute as Byte
    Private seconds as Single
```

```

Public Sub _Create()
    Call SetTime()
End Sub

Public Sub _Create(ByVal h as Byte)
    Call SetTime(h)
End Sub

Public Sub _Create(ByVal h as Byte, ByVal m as Byte)
    Call SetTime(h, m)
End Sub

Public Sub _Create(ByVal h as Byte, ByVal m as Byte, ByVal s as Single)
    Call SetTime(h, m, s)
End Sub

Public Sub SetTime(ByVal h as Byte = 0, ByVal m as Byte = 0, _
                  ByVal s as Single = 0.0)
    hour = Min(h, 23)
    minute = Min(m, 59)
    Seconds = Min(s, 59.999)
End Sub

Public Function GetHour() as Byte
    GetHour = hour
End Function

Public Function GetMinute() as Byte
    GetMinute = minute
End Function

Public Function GetSeconds() as Single
    GetSeconds = seconds
End Function
End Class

Dim t as MyTime = _Create(12)

Sub Main()
    Debug.Print t.GetHour()
End Sub

```

Of course, in this particular case the four separate constructors could be replaced by a single constructor that uses default parameter values. (Just as with regular ZBasic procedures, the compiler automatically inserts the specified default value if you omit one or more of the right-most parameters on a particular invocation.) If that were done, the constructor would have the form shown below.

```

Public Sub _Create(ByVal h as Byte = 0, ByVal m as Byte = 0, _
                  ByVal s as Single = 0.0)
    Call SetTime(h, m, s)
End Sub

```

When you define an array of objects, you may specify different constructors for each element of the array as shown in the example below. In this particular case, explicit constructor invocations were specified only for three of the ten elements of the array of objects. The remaining elements are initialized using the *default constructor* – one that can be invoked with no parameters. You may also specify a constructor list for multi-dimension arrays. In that case, the constructors are applied sequentially to the elements of the array with the leftmost index varying the fastest. This corresponds with the way that arrays are laid out in memory.

```
Dim t2(1 to 10) as MyTime = { _Create(12), _Create(), _Create(1, 2) }
```

It should be noted that, although frequently done, it is not necessary to define a default constructor. If you don't define a default constructor, the compiler will automatically supply one for you that performs only the fundamental initialization. A constructor that has one or more parameters, all of which have default values specified, is effectively a default constructor.

One last issue that must be kept in mind when writing the code for a constructor is that, for objects defined at the module level, the constructor code will be executed before any tasks have been created. Because of this, you must avoid using any ZBasic System Library routines that rely on the multi-tasking system being initialized. In particular, the subroutines Sleep(), Delay() and related subroutines must not be used in a constructor for an object that might be defined at the object level. One way to work around this limitation is to add a special initialization method to the class that can be called from Main() to perform initialization that requires the multi-tasking system to be up and running. Another alternative for situations where a delay is needed, is to add a simple delay method to the class that effects a delay based on the changing value of Register.RTCTick.

4.5 Object Destruction Issues

When an object reaches the end of its lifetime, i.e. it goes out of scope, it may need to be "cleaned up". For example, if one or more data members are allocated string types, the memory allocated for the string store (if any) needs to be freed. The ZBasic compiler generates code that correctly deallocates string members of objects just as it does for individual string variables and structure members that are allocated string types.

Beyond that automatic cleanup activity, you may need to perform additional "cleanup" operations at the end of an object's life. For example, if your object acquires other resources that need to be released, e.g. a semaphore or memory allocated using System.Alloc(), you need to handle these operations yourself by writing an optional class method called a *destructor*. A class may have at most one destructor which is a parameterless subroutine with the special name _Destroy. In the code for your destructor you only need to handle the data members that aren't automatically handled by the compiler. For example, you needn't be concerned with string deallocation because that is handled automatically. The automatically generated cleanup code is executed after the code in your destructor executes.

4.6 Object Assignment Issues

A third optional method, known as an *assignment constructor*, may be defined to ensure that the assignment of one object to another (necessarily of the same type) is carried out properly. An assignment constructor is a subroutine with the special name _Assign() that takes a single ByVal parameter with the same type as the class. For most classes, it isn't necessary to define an assignment constructor because the ZBasic compiler correctly handles object assignment. For data members that are allocated strings, the compiler generates code that frees of the destination object member's existing string and then creates a copy of the source object member's string. For data members that are objects, the compiler generates code that invokes the object's assignment constructor or, if it has none, the default assignment process for that object. The same strategy is employed for data members that are structures having allocated string members. However, for all other data members the compiler generates code that performs a simple byte-wise copy.

If your class has data members that have special requirements, e.g. a data member that is an address of a block of memory allocated with System.Alloc(), you'll need to write your own specialized assignment constructor. It is important to note that if you do provide an assignment constructor, you are responsible for handling every data member of the class, not just those that require special treatment. An example of an assignment constructor for the MyTime class is shown below. In this case, the explicit assignment constructor is superfluous since it does no more than the default assignment process does. If an assignment constructor is not needed, it is advisable not to create one because it introduces a possible source of problems if new data members are added to the class and the assignment constructor is not updated to handle the new members.

```

Public Sub _Assign(ByVal src as MyTime)
    hour = src.hour
    minute = src.minute
    seconds = src.seconds
End Sub

```

4.7 Object Self-reference and Parent Reference

Occasionally, it is necessary or expedient to be able to refer to the particular object instance in a method. For such purposes, the special identifier `_this` is recognized. This identifier can be used as a qualifier on a member reference (e.g. `_this.hour`) or as a prefix to the `DataAddress` qualifier: `_this.DataAddress`.

Similarly, it is sometimes useful to be able to refer to a method of the parent class of an object. For this purpose, the special identifier `_parent` can be used as a prefix to a class method, e.g. `Call _parent.DoSomething()`. This is only necessary, of course, if the class in which the reference occurs has a method with the same signature (name plus parameter types), e.g. it is often useful for invoking a parent's constructor. If needed, the `_parent` prefix can be used multiple times to refer to a particular ancestor class, e.g. `_parent._parent._parent.DoSomething()`.

4.8 Explicit Class References and Default Namespace References

In some cases, it may be necessary to refer to a method or data member of a specific class, overriding the method or member that would otherwise be matched by a particular name. You can refer to the namespace of a specific class by adding a class name prefix to the identifier. The class name prefix has the form of the class' name followed by two colons, e.g. `MyTime::hour`. The class name prefix is useful inside of class methods as well as in non-class subroutines and functions and for referring to static methods of a class.

Occasionally, it may be necessary to refer to an identifier that is outside of a class when the class contains a method or data member having the same name. Adding the default namespace prefix (two colons) to the identifier allows you to refer to the namespace outside of the class hierarchy, e.g. `::foo`.

4.9 Class Sections

When defining a class, the various members have default visibility attributes identical to those that would apply if the same type of member were being defined outside of a class. For example, if a data member is defined using `Dim`, it will be private. Similarly the default visibility for a method is public just as it is for subroutines and functions outside of a class.

The differing default visibility attributes for different types of member can lead to confusion about the visibility of a particular element of a class. To mitigate this issue, one could adopt the practice of specifying the `Public` or `Private` attribute for every member. An alternative strategy is to use section labels in the class to override the default visibility of the members in the section. The `MyTime` class definition is rewritten below using section labels.

```

Class MyTime
Public:
    Sub _Create(ByVal h as Byte = 0, ByVal m as Byte = 0, _
                ByVal s as Single = 0.0)
        Call SetTime(h, m, s)
    End Sub

    Sub SetTime(ByVal h as Byte = 0, ByVal m as Byte = 0, _
                ByVal s as Single = 0.0)
        hour = Min(h, 23)
        minute = Min(m, 59)
    End Sub

```

```

        Seconds = Min(s, 59.999)
End Sub

Function GetHour() as Byte
    GetHour = hour
End Function

Function GetMinute() as Byte
    GetMinute = minute
End Function

Function GetSeconds() as Single
    GetSeconds = seconds
End Function

Private:
    Dim hour as Byte
    Dim minute as Byte
    Dim seconds as Single
End Class

```

Section labels may be used in any order and may appear multiple times. The order used in the example above is suggested but any suitable arrangement may be chosen. It is important to remember that private data members and methods are only visible to methods of the class. This is in contrast to the situation with structures where private members are visible to all procedures within the same module but not visible outside of the module in which the structure is defined.

4.10 Static Class Members

So far, all of the methods and data members defined in the examples have been object-specific. That means that invoking a method or accessing a data member had to be realized via an object reference, e.g. `t.SetTime()`. Within the methods themselves, the methods and data members of the class can be accessed without an explicit object reference because the object reference is implicit.

It is possible to define methods and data members that can be accessed without any object reference, explicit or implicit. To do so, simply add the `Static` keyword to the definitions as illustrated in the example below.

```

Class MyTime
Public:
    Sub _Create(ByVal h as Byte = 0, ByVal m as Byte = 0, _
                ByVal s as Single = 0.0)
        Call SetTime(h, m, s)
    End Sub

    Sub SetTime(ByVal h as Byte = 0, ByVal m as Byte = 0, _
                ByVal s as Single = 0.0)
        hour = Min(h, 23)
        minute = Min(m, 59)
        Seconds = Min(s, 59.999)
    End Sub

    Function GetHour() as Byte
        GetHour = hour
    End Function

    Function GetMinute() as Byte
        GetMinute = minute
    End Function

```

```

Function GetSeconds() as Single
    GetSeconds = seconds
End Function

Static Function Fmt(ByVal h as Byte, ByVal m as Byte, _
                    ByVal s as Single) as String
    Fmt = CStr(h) & ":" & CStr(m) & ":" & ::Fmt(s, 2)
End Function

Private:
    Dim hour as Byte
    Dim minute as Byte
    Dim seconds as Single
End Class

Sub Main()
    Debug.Print MyTime::Fmt(3, 10, 15.7)
End Sub

```

In this example, a static method named `Fmt()` is defined that produces a formatted time string. There are several important aspects of the example to note. Firstly, you may recall that there is a ZBasic System Library function called `Fmt()`. This example illustrates how you can create an overload for a System Library routine. Secondly, within the new `Fmt()` method, it was convenient to be able to use the System Library function of the same name. The example code shows how you can use the default namespace prefix to specify that you want to invoke the System Library routine rather than the `Fmt()` method of the `MyTime` class. Lastly, the invocation of `Fmt()` in the `Main()` subroutine shows how the classname prefix is used to specify that you want to invoke the `Fmt()` method of the `MyTime` class.

A static method may invoke only other static methods and procedures defined outside of the class. Also, it may access only static data members and other data items defined outside of the class. It cannot access non-static data members because those are specific to an instance.

A static data member can be similarly defined for a class. Access to the data member is achieved using the classname prefix just as with the static method. It is important to note that there is only one instance of each static data member in the application no matter how many objects of the class containing the static data member have been instantiated. The effect is the same as if a module-level variable were defined outside of the class; the advantage of making it part of the class is that it associates the static data item with the class to a greater degree than would be the case if a module-level variable were defined.

It is important to realize that a static method generally must be public for it to be of any value. If it were private (or protected), it would only be accessible to class (and subclass) methods. For the most part, the same is true for static data members although there may be special cases where private (or protected) static data items would be useful.

4.11 Inheritance

The example classes described thus far have derived all of their characteristics and capabilities from the methods and data members explicitly defined as part of the class. One of the great benefits of object-oriented programming comes from the ability to define a class that is based on the functionality of an existing class. Essentially, the idea is to take an existing class and extend it by adding some new capability without affecting the users of the existing class. Consider the example below.

```

Class A
Public:
    Enum Color
        Red
        Green
        Blue
    End Enum

```

```

Sub _Create(ByVal s as Byte = 0, ByVal c as Color = Red)
    m_size = s
    m_color = c
End Sub

Sub Identify()
    Debug.Print "I'm an A, size="; m_size
End Sub
Private:
    Dim m_size as Byte
    Dim m_color as Color
End Class

' Define a new class derived from an existing class.
Class B Extends A
Public:
    ' Define a constructor invoking a base class constructor.
    Sub _Create(ByVal s as Byte = 0, ByVal c as Color = Red, _
               ByVal w as Byte = 1)
        Call _parent._Create(s, c)
        m_weight = w
    End Sub

    Sub ShowWeight()
        Debug.Print "weight = "; m_weight
    End Sub

Private:
    Dim m_weight as Byte
End Class

Dim a1 as A = _Create(5, A::Green)
Dim b1 as B = _Create(10, B::Red, 100)

Sub Main()
    Call a1.Identify()
    Call b1.Identify()
    Call b1.ShowWeight()
End Sub

```

This example illustrates how the derived class B inherits some of its functionality from its base class A. The object b1 has a method named `Identify()` even though it is not explicitly present in the definition of class B; it inherits that method from its base class. This example also shows how to invoke a particular base class constructor in the body of the constructor of a derived class - the `_parent` prefix is used to refer to a constructor in the parent class. The base class constructor invocation, if present, must be the first statement in the constructor. If you don't explicitly invoke a base class constructor, the compiler will automatically include code that invokes the default constructor of the base class. You may also invoke a constructor of particular ancestor class instead of a constructor of the base class.

It is important to be aware that the private methods and data members of the base class are not accessible to the derived class. Because it is often useful for derived classes and base classes to share some method and data members that are not publically available, ZBasic supports a third visibility attribute called `Protected`. Base class methods and data members that have the `Protected` visibility attribute are not accessible to code outside of the class methods but they are accessible to code in the methods of a derived class. You can use the `Protected` keyword within a class definition anywhere that you can use the `Public` and `Private` keywords, including in section labels.

If you compile and run the example code above, you'll note that the invocation of `b1.Identify()` displays the same output as when `a1.Identify()` is invoked. It is likely that this behavior is not really what is wanted. Rather, it is more likely that the objects would identify themselves differently. This can

be corrected in several ways. Perhaps the most obvious way is to add a protected data member that contains all of or part of the identity string or some value to otherwise identify the object. While this will work, it is wasteful of space because every instance of the object will contain that data member. A better way to solve this problem is to add an `Identify()` method to the `B` class. A re-written version of the previous example appears below which has an `Identify()` method defined in both classes. The `m_size` data member was also made `Protected` so that it can be accessed by subclasses such as class `B`.

```

Class A
Public:
    Enum Color
        Red
        Green
        Blue
    End Enum

    Sub _Create(ByVal s as Byte = 0, ByVal c as Color = Red)
        m_size = s
        m_color = c
    End Sub

    Sub Identify()
        Debug.Print "I'm an A, size="; m_size
    End Sub
Protected:
    Dim m_size as Byte
Private:
    Dim m_color as Color
End Class

' Define a new class derived from an existing class.
Class B : A
Public:
    ' Define a constructor using a base class constructor.
    Sub _Create(ByVal s as Byte = 0, ByVal c as Color = Red, _
               ByVal w as Byte = 1) : _Create(s, c)
        m_weight = w
    End Sub

    Sub Identify()
        Debug.Print "I'm a B, size="; m_size; ", ";
        Call ShowWeight()
    End Sub

    Sub ShowWeight()
        Debug.Print "weight = "; m_weight
    End Sub

Private:
    Dim m_weight as Byte
End Class

Dim a1 as A = _Create(5, A::Green)
Dim b1 as B = _Create(10, B::Red, 100)

Sub Main()
    Call a1.Identify()
    Call b1.Identify()
End Sub

```

It is important to understand that when a class is derived from another class, the compiler generates code to ensure that the base class object is fully initialized before the constructor for the derived class executes. Similarly, for assignment constructors the base class assignment constructor is executed before that of the derived class. This order guarantees that the derived class object may safely access the methods and members of the base class in its constructors.

The situation with destructors is similar but the order is reversed. In the case of destructors the derived class destructor executes before the base class destructor. This order guarantees that the derived class may rely on the integrity of the base class data members if necessary.

One aspect of inheritance that must be carefully considered is that since the base class constructor executes before a derived class constructor, base class constructors must avoid invoking methods that would result in the derived class method executing before derived class' constructor has executed.

4.12 Abstract Classes, Abstract Methods

When designing an object-oriented solution, it is often useful to define a class that embodies some essential characteristics and functions that can then be inherited by several other classes. In some situations, the base class is "incomplete" in the sense that its only purpose is to serve as the base class for other classes. You can prevent such an incomplete class from being instantiated (i.e. defining an object of that type) by defining the base class to be *abstract*. This is done by adding the `Abstract` attribute to the class definition.

```
Abstract Class MyObject
Public:
Protected:
    Dim m_size as Integer
    Dim m_weight as Integer
Private:
End Class
```

If you try to create an object of the class `MyObject`, the compiler will produce an error message indicating that you cannot instantiate an abstract class. This is to remind you that the purpose of the class is as a building block for other classes.

You can also indicate that a class is abstract by defining the class with at least one abstract method as in the class definition below.

```
Class MyObject
Public:
    Abstract Sub Identify()
    End Sub
Protected:
    Dim m_size as Integer
    Dim m_weight as Integer
Private:
End Class
```

The result of both of the preceding class definitions is an abstract class. The difference is that when you define abstract methods, in addition to making the class abstract it also establishes a requirement that all derived classes must define a method that has the same signature, i.e. the same name and same number and types of parameters.

If you wish, you may include code in the body of the abstract method. One reason for doing so would be to define some essential functionality that many or most derived classes could use. The derived class method can invoke the abstract base class method using the class namespace prefix or the parent class prefix. For example, class A could be derived from `MyObject` and the `Identify()` methods of the `MyObject` and `A` classes could be rewritten as shown below, respectively. You should compile and run such an example to observe the result.

```

Abstract Sub Identify()
    Debug.Print "Howdy, ";
End Sub

Sub Identify()
    Call MyObject::Identify()
    Debug.Print "I'm an A, size="; m_size
End Sub

```

4.13 Final Classes

The classes described thus far have the characteristic that they may be (or, in the case of `Abstract` classes, must be) extended. It is also possible to define a class that may not be extended. This is done by using the `Final` attribute preceding the `Class` keyword.

```

Class A
Public:
    Sub Identify()
    End Sub
Protected:
    Dim m_size as Integer
    Dim m_weight as Integer
Private:
End Class

Final Class B Extends A
Public:
    Sub Identify()
    End Sub
Protected:
Private:
End Class

```

4.14 Using Mixins

Earlier in this chapter, the concept of inheritance (defining a class that inherits functionality from another class) was discussed. Also discussed was the concept of composition (defining a class containing one or more data members that are other objects).

To review, if class B is defined as inheriting from class A, then class B is said to satisfy the “*is a*” relationship (sometimes written as ISA) with respect to class A. This means that an object of type class B can be treated as if it is, in fact, an object of type class A. The public and protected methods and data members of class A are available to all methods of class B as if they were directly part of the definition of class B, etc.

In contrast, if class B is defined with a data member of type class A (i.e. using composition), then class B would be said to satisfy the “*has a*” relationship (sometimes written as HASA) with respect to class A. Only the public methods and data members of class A would be available to the methods of class B and the access would need to be qualified using the name of the data member.

The concept of a *mixin* combines aspects of both inheritance and composition. If class B includes class A as a mixin, then the public methods and data members of class A are accessible to the methods of class B as if they were directly part of the definition of class B (as with inheritance) but class B satisfies the HASA relationship with respect to class A (as with composition). Because you can specify multiple mixin classes, doing so provides some of the benefits of multiple inheritance but without having the difficult issues that arise with traditional multiple inheritance.

Including one or more other classes as mixins of a given class is accomplished using the `Includes` keyword as shown in the example below. If more than one mixin class is specified, the mixin class names are comma-separated.

```
Class A
Public:
    Sub DoSomething()
    End Sub
Protected:
Private:
End Class

Class B Includes A
Public:
    Sub Identify()
        Call DoSomething()
    End Sub
End Class

Dim MyB as B

Sub Main()
    Call MyB.DoSomething()
End Sub
```

One restriction on using mixins is that the public and protected methods and data members of each mixin must be unique among all of the mixins included and must not duplicate any method or data member names of the including class. This restriction is necessary because all public and protected methods and data members need to reside in the namespace of the including class and therefore must be unique. It is possible, however, to have methods of the same name in multiple mixins and/or the including class as long as they each have unique signatures, i.e. they must have different number and/or types of parameters. Note, however, that the including class may contain methods that overload methods of mixin classes. It is important to be aware that a mixin class method that is marked as `Abstract` *must* be overloaded by the including class.

Finally, if a class is defined using both inheritance and mixins, the `Extends <class name>` specification must appear before the `Includes <class name list>` specification, e.g.

```
Class D Extends B Includes A, C
Public:
    Sub DoSomething()
    End Sub
End Class
```

In rare cases, it is useful or necessary to be able to refer to the mixin element or its data members or methods as if they were made part of the containing class by composition. For this purpose, a member name for the mixin class element is automatically provided by the compiler consisting of the mixin class name with `_Mixin_` preceding it. For example, for a mixin class A the automatically created member name is `_Mixin_A`. Consequently, a public method of mixin A can be referred to using either of the following:

```
Call DoSomething()
Call _Mixin_A.DoSomething()
```

The two methods above are equivalent so there is no advantage to using the second, more verbose, option. However, if you want the address of the mixin element, the only way to achieve that goal is to use the prefixed name:

```
@_Mixin_A
```

4.15 Using the Const Attribute for Methods

Earlier in this chapter, it was stated that an object may be passed to a subroutine or function `ByRef` or `ByVal`. If it is passed `ByVal`, the object instance is considered to be “read-only” within the receiving procedure. As such, the procedure has read access to the public data members of the object but it is not allowed to modify those data members. Also, the receiving procedure may not pass the object `ByRef` to any other procedure or method. Additionally, in order to invoke a method of the object, that method must be declared to be “constant”, meaning that it isn’t allowed to modify the object or pass the object by reference to another method or procedure.

To define a class method as constant, add the keyword `Const` following the closing parenthesis of the parameter list. Optionally, for methods that are functions, you may instead place the `Const` keyword after the function’s return type. An example is shown below using a previously defined method.

```
Sub Identify() Const
    Call MyObject::Identify()
    Debug.Print "I'm an A, size="; m_size
End Sub
```

Of course, since the method above invokes the base class method of the same name, the latter method must also have the `Const` attribute. In summary, a `Const` method may invoke only other `Const` methods and may not modify any data members or any data members of the base class.

4.16 Based Objects, Reference Objects

It may be useful in some situations to define an object at a specific memory address. For example, you may want to allocate some memory using `System.Alloc()` to hold an object. Or, you may wish to define an object that occupies previously allocated space, such as a pre-defined buffer. In either event, the desired effect can be achieved by defining an object based at a given address. Consider the example below that uses the class `B` defined in Section 4.11.

```
Dim buf(1 to 20) as Byte

Sub Main()
    Dim b1 as B Based buf.DataAddress
    Call b1._Create(5, B::Blue, 25)
    Call b1.ShowWeight()
    Call b1._Destroy()
End Sub
```

There are several important aspects of this example that should be clearly understood. Firstly, when you create a based object you are responsible for all aspects of its management. You must ensure that sufficient space for the object exists at the address at which you base the object. The example above is poorly coded because the buffer may, in fact, be too small for the object. The example could be improved by replacing the upper bound of the buffer with `SizeOf(B)`. Secondly, you must explicitly invoke the constructor for the object before using any of its methods or data members because, prior to the constructor invocation, the object content is completely undefined. Note that explicit constructor invocation is disallowed for non-based objects but is (usually) required for based objects. Also, you are also responsible for invoking the object destructor (if necessary). Failure to do so may lead to memory “leaks”, i.e. allocated memory blocks that are not properly freed.

A second example, below, illustrates using a based object with allocated memory. Other than the means by which the space for the object is acquired, the issues are identical to those of the previous example.

```
Sub Main()
    Dim addr as UnsignedInteger
    addr = System.Alloc(SizeOf(B))
    Dim b1 as B Based addr
    Call b1._Create(5, B::Blue, 25)
```

```

    Call b1.ShowWeight()
    Call b1._Destroy()
    Call System.Free(addr)
End Sub

```

Another use for a based object is to be able to treat a base class object as if it were a derived object. Here again, it is your responsibility to ensure that such treatment is appropriate. The compiler will not issue any error messages or warnings about incompatible classes. As an example, consider the subroutine below.

```

Sub IdentifyObject(ByVal obj as A)
    Call obj.Identify()

    ' Define an object based on the object passed
    Dim objB as B Based obj.DataAddress
    Call objB.ShowWeight()
End Sub

```

If we knew, in certain cases, that the passed object were actually of class B (derived from class A), we could use a based object definition to be able to access the unique methods and data members of the B class. This capability is rarely needed but it is good to keep it in mind for those cases where it is. It is important to distinguish this particular use of a based object from the use case given in the first example in this section. In this use case the object is assumed to be validly constructed and, therefore, the constructor/destructor should not be invoked on the object as it is in the first example where the based object was defined as overlaying memory of undefined content.

One final note, you may optionally include the `Const` keyword between the class name and the `Based` keyword to define a read-only based object. With such a read-only based object definition, access is limited to `Const` methods and, moreover, data members may not be modified nor passed by reference to other procedures.

The ideas described above apply equally well to objects defined `ByRef`. Consider the example below that is a variation of the first example given above.

```

Dim buf(1 to 20) as Byte

Sub Main()
    Dim b1 as B ByRef
    b1.DataAddress = buf.DataAddress
    Call b1._Create(5, B::Blue, 25)
    Call b1.ShowWeight()
    Call b1._Destroy()
End Sub

```

4.17 Miscellaneous Class Elements

Class data members may be defined that have the `ProgMem` or `Persistent` attributes but in either case, the data member must be defined as `Static` so that there is only one instance of the data item. Such members may be either `Private` or `Public` as needed.

A `Declare` statement may be placed in a class definition. It may be either `Public` or `Private` but it must have the `Static` attribute.

You may define constants, enumerations and structures within a class definition. Such definitions may be private if it is desirable that they are only available to class methods or they may be public if that is useful.

BasicX Compatibility Note

The object-oriented features are not supported in BasicX compatibility mode.

Chapter 5 - Compiling for Generic Target Devices

5.1 Overview

Later versions of the ZBasic compiler support compiling ZBasic code for generic target devices, i.e. microcontrollers that are not ZX devices. This capability requires a special license (available from Elba Corp.). A complete list of the supported generic target devices is given in Appendix B.

A generic target device requires some external circuitry similar to that required for the 40-pin and 44-pin ZX devices. See Appendix M for information about the external circuitry that may be required.

If you are familiar with writing applications for ZX devices, there are some minor differences that need to be considered when making the change to using generic target devices. The biggest difference is that the serial console (typically serial channel 1) is not open by default on a generic target device as it is with a ZX device. This change was made for two reasons: not all generic target devices have a hardware UART (e.g. most ATtiny devices) and, secondly, many target devices have a rather small complement of Flash memory and omitting the code to open the serial channel by default reduces the size of the executable image. If you want to use the console serial channel in an application for a generic target device you can either add code to explicitly open it or, alternately, use the Option ConsoleSpeed directive to specify the default console baud rate.

A second difference between ZX and generic target applications relates to the RTC. If your generic target application includes the RTC, the several ZBasic System Library routines that rely on the presence of the RTC will behave in accordance with the RTC parameters you've specified. One example of this is the ZBasic System Library routine Sleep(). With an RTC present, the units of sleep interval are related to the period of the RTC that you specify. If no RTC is included, the units of sleep interval are 1 millisecond. Some ZBasic System Library routines, and multi-tasking generally, require the presence of the RTC; if you refer to them in an application that does not include the RTC the compiler will issue error messages.

5.2 Loading Application Code onto the Target Device

When using a generic target device, there are three possible scenarios for getting the application code programmed into the device. The most straightforward scenario is to use a specialized device programmer (e.g. the Atmel AVRISP) to write the application code to the target. When compiling code for this scenario the ZBasic compiler will produce an Intel Hex Format file containing the program image and (possibly) a second Intel Hex format file containing the data to be written to the device's EEPROM. Once you have these files you can use the command line or GUI utility to use the programming device to write the program and/or EEPROM data to the device. Note, also, that the ZBasic IDE provides a means to specify a command line to execute to facilitate downloading code to a generic device.

The two remaining scenarios are useful only for target devices that support having a bootloader. In this case, you would initially program the bootloader into the device using the same type of specialized device programmer referred to in the previous paragraph. After the bootloader is installed, the code and EEPROM images can be programmed into the device by interacting with the bootloader, typically over a serial or USB link. If the installed bootloader is a standard ZBasic bootloader (source code for which is provided with the compiler), the ZBasic IDE can download the application code to the target device in much the same manner as it does for ZX devices. Otherwise, if the bootloader is not compatible with the ZBasic bootloader, the IDE can download the code to the device only if you provide the IDE with a command line to do so. If you do not have access to a command line utility for downloading via the bootloader you'll have to user whatever application is provided for downloading via that bootloader outside of the IDE.

5.3 Target Device Parameters

In contrast to ZX devices, where operating parameters such as main clock frequency, RTC frequency, etc. are all predefined, with a generic target device you must specify several operating parameters to allow the ZBasic compiler to generate correct code. The parameters for generic target devices are specified using the --device-parameter command line option or the Option DeviceParameter

directive. In both instances the specification of a device parameter comprises a parameter name and a parameter value. The next section describes the available device parameters.

The following list enumerates the available device parameters, describes the parameter values, indicates which are required, and, when applicable, when they cannot be used. Note that neither the names nor the values are case sensitive. The procedure for determining appropriate values for some of the device parameters is fairly complicated (e.g. the SW UART parameters). The IDE provides a special dialog that performs the computations given some basic operating parameters. See Section 9.6 for more details.

The target device parameters may be specified in the project file using the command line option form (–device-parameter=*name*,*value*) or in the source code itself using the Option form (Option DeviceParameter *name* *value*). In the latter case, the value must be specified either as a decimal number, as an identifier (begins with a alphabetic character followed by zero or more alphanumeric characters), or as a quoted string. Examples of both forms are given with each description below. See the section of the ZBasic System Library manual entitled "Processor Speed and Device Configuration Issues" for information on the effect of generic target configuration values for various ZBasic System Library routines.

A few target device parameters can also be used with native mode ZX devices; the descriptions below explicitly indicate if they can be.

ClockFrequency

ClockFrequency < <i>value</i> >	Default: <i>none</i>
--	----------------------

This required parameter specifies the operating frequency of the target device (specified in Hertz). The value of this parameter influences all of the time-dependent aspects of ZBasic routines such as serial communication, I/O timing, delays, etc. as well as the operation of the real time clock (RTC) if it is enabled.

Examples

```
--device-parameter=ClockFrequency,16000000
Option DeviceParameter ClockFrequency 16000000
```

ClockPrescaler

ClockPrescaler < <i>value</i> >	Default: 1
--	------------

This optional parameter, valid only for ATtiny and ATmega generic devices and native mode ZX devices, specifies a division factor to be applied to the main clock source to achieve the desired clock frequency. The set of valid prescaler values is 1, 2, 4, 8, 16, 32, 64, 128 and 256. A device with a 16MHz crystal and a prescaler value of 2 will operate at 8MHz.

It is important to note that the maximum baud rate for the software serial channels is reduced in direct proportion to the clock prescaler. For example, on a ZX-24n using a prescaler value of 2 results in a maximum baud rate of 9600 as compared to the standard 19.2K baud. For ZX devices, the largest prescaler value allowed is the one that reduces the operating frequency to no lower than 230.4KHz.

Examples

```
--device-parameter=ClockPrescaler,2
Option DeviceParameter ClockPrescaler 2
```

Package

Package <package-name>

Default: see discussion

This parameter specifies the physical packaging of the target device. Many of the supported generic target devices are available in more than one package type and often the pin assignments as well as the available functionality varies based on the package. The set of recognized package identifiers is given below but only a few of the package types apply to any specific target device. Note, too, that some devices are available in multiple package types that have the same pin assignments, e.g. QFN-44 and MLF-44. In such cases it does not matter which package type is specified as long as the pin assignments match those of the actual package that you are using. Consequently, if the compiler complains that the package type you've specified is invalid for the target device, try specifying a different package type with the same pin assignments. Note that although not shown in the table below, it is permitted to specify a package name omitting the dash preceding the pin count, e.g. PDIP40 is equivalent to PDIP-40.

Recognized Package Types

PDIP-8	PDIP-14	PDIP-20	PDIP-28	PDIP-40
TQFP-28	TQFP-32	TQFP-44	TQFP-64	TQFP-100
SOIC-14	SOIC-20	TSSOP-20	VQFN-20	PLC-44
MLF-28	MLF-32	MLF-44	MLF-64	MLF-100
QFN-28	QFN-32	QFN-44	QFN-64	QFN-100

The default package type varies by target device and is selected from entries in the table above in left to right and top to bottom order.

Examples

```
--device-parameter=Package,TQFP-44
--device-parameter=Package,PDIP40
Option DeviceParameter Package "TQFP-44"
Option DeviceParameter Package PDIP40
```

RTCFrequency

RTCFrequency <value>

Default: 0

If your application uses multiple tasks or utilizes any of the RTC-related routines, you must specify the operating frequency of the real time clock (in Hertz). It is important to note that the specified RTC frequency may not be attainable with 100% accuracy given the main clock frequency, the available timer divisors, the RTC scale factor (see below) and the resolution of the timer used for the RTC. If the realizable RTC frequency varies from the specified frequency by more than the specified or default error threshold (see RTCError below) the compiler will issue a warning.

Examples

```
--device-parameter=RTCFrequency,500
Option DeviceParameter RTCFrequency 500
```

RTCScale

RTCScale <value>

Default: see discussion

At some operating frequencies, given the available divisors and resolution of the RTC timer, a scale factor is sometimes necessary to achieve a convenient RTC frequency. The RTCScale device parameter specifies the number of RTC timer interrupts that will occur for each RTC tick. Acceptable values for RTCScale are 1 and 2. Note that the rate of change of the ZBasic value Register.RTCFastTick is equal to the rate of change of Register.RTCTick times the RTCScale value. If you do not specify a value for this parameter, a value will be chosen by the compiler that results in an RTC frequency closest

to the requested value.

Examples

```
--device-parameter=RTCScale,2  
Option DeviceParameter RTCScale 2
```

RTCError

RTCError <value>	Default: 10
-------------------------	-------------

This parameter gives the desired error threshold, expressed in parts-per-million, for the realizable RTC frequency. If the realizable RTC frequency deviates by more than the threshold value the compiler will issue a warning.

Examples

```
--device-parameter=RTCError,100  
Option DeviceParameter RTCError 100
```

TimerSpeed1Divisor, TimerSpeed2Divisor

TimerSpeed1Divisor <value>	Default: 1
TimerSpeed2Divisor <value>	Default: 8

These parameters give the initial values for Register.TimerSpeed1 and Register.TimerSpeed2, respectively. These prescaler values are used to select the operating frequency of the I/O Timer for certain timer-base I/O routines as described in the ZBasic System Library manual.

Examples

```
--device-parameter=TimerSpeed2Divisor,64  
Option DeviceParameter TimerSpeed2Divisor 64
```

TimerOCPin

TimerOCPin <timer-pin-specification>	Default: see discussion
---	-------------------------

This parameter specifies a mapping of a timer compare output to a physical pin. It is only useful for certain devices, e.g. the ATtiny828, that support output compare pin mapping. The <timer-pin-specification> is a quoted string of the form "OC<timer><output>:<port>.<pin>" where

<timer> is a timer designator from the set {0..9}
<output> is an output compare designator from the set {A..Z}
<port> is a port designator from the set {A..Z}
<pin> is a port bit designator from the set {0..7}

The default mapping of timer compare outputs to pins is described in the ZBasic System Library Manual.

Examples

```
--device-parameter=TimerOCPin,"OC1A:B.5"  
Option DeviceParameter TimerOCPin "OC1A:B.5"
```

HWUartSpeed

HWUartSpeed <value>[,<value>]*

Default: none

This parameter is only useful with devices that support fractional baud rate generation (e.g. the ATxmega devices), allowing you to specify one or more baud rates that your application might use. Certain “baud rate friendly” operating frequencies (e.g. 14.7456MHz) are integral multiples of the common UART baud rate and thus allow standard baud rate generation with zero error. With other operating frequencies (e.g. 1.0MHz) only a few baud rates can be achieved that fall under the generally accepted 2% error margin unless the fractional baud rate generator is used. The set of baud rates specified by the (possibly multiple) instances of this parameter is used to generate a table of initialization values for your application thus avoiding the need to perform the computationally intensive procedure for determining the fractional baud rate setup values at run time. If you do not specify a set of baud rates using this parameter, a default list of standard baud rates ranging from 300 to 460800 is employed. Use of this device parameter is only needed if your target device is not running at a “baud rate friendly” frequency and you need to use a non-standard baud rate.

Note that only one baud rate may be specified when using the Option DeviceParameter form; use multiple instances to specify the set of baud rates desired.

Examples

```
--device-parameter=HWUartSpeed,9600,19200
Option DeviceParameter HWUartSpeed 9600
Option DeviceParameter HWUartSpeed 19200
```

SWUartDivisor, SWUartMinSpeed, SWUartMaxSpeed, SWUartBaseSpeed

SWUartDivisor <value>	Default: 0
SWUartMinSpeed <value>	Default: 0
SWUartMaxSpeed <value>	Default: 0
SWUartBaseSpeed <value>	Default: 0

These parameters are used to configure the software UART timer to provide the bit timing for the software UART channels (3-6). It is only necessary to specify values for these parameters if the software UART channels are used.

The parameter **SWUartDivisor** specifies the prescaler divisor value for the software UART timer causing the timer to run at that fraction of the main operating frequency. This value must be chosen to allow accurate bit timing of the software UART over the range of speeds desired and, further, to permit suitable timer limit values over the range of baud rates intended to be used. The timer limit value is computed as ($\text{ClockFrequency} / \text{SWUartDivisor} / (4 * \text{base baud rate})$). This value must fall in the range given in the table below for all possible *base baud rates* (as defined below).

Allowable Range of Timer Limit Values

8-bit timer (e.g. mega, tiny)	16-bit timer (e.g. xmega)
1 < timerLimit <= 256	1 < timerLimit <= 65,536

The values of the **SWUartMinSpeed** and **SWUartMaxSpeed** parameters must be in the range 300-19200 inclusively and the minimum speed must be less than or equal to the maximum speed. For devices with 8-bit timers, the value of the **SWUartBaseSpeed** may be specified as greater than the minimum speed and less than or equal to the maximum speed.

The set of *base baud rates* for which the timer limit value must be checked ranges from **SWUartMaxSpeed** on the high end to the greater of **SWUartMinSpeed** and **SWUartBaseSpeed** (if specified) on the low end. Essentially, the **SWUartBaseSpeed** parameter provides a way to avoid exceeding the timer limit range while still supporting the lower baud rates. The example set of parameter values shown below would be suitable for a target device running at 16MHz and having an 8-bit software

UART timer. For a device running at the same speed and having a 16-bit software UART timer, the SWUartBaseSpeed parameter is not needed.

Examples

```
--device-parameter=SWUartDivisor,8
--device-parameter=SWUartMinSpeed,300
--device-parameter=SWUartMaxSpeed,19200
--device-parameter=SWUartBaseSpeed,2400
Option DeviceParameter SWUartDivisor 8
Option DeviceParameter SWUartMinSpeed 300
Option DeviceParameter SWUartMaxSpeed 19200
Option DeviceParameter SWUartBaseSpeed 2400
```

ZBasicBootloader

ZBasicBootloader { true/false }	Default: false
--	----------------

This parameter specifies whether or not the target device contains a standard ZBasic bootloader. The source code for the standard ZBasic bootloader is supplied with the ZBasic compiler and the bootloader code is pre-built for all supported devices (not all target devices are capable of having a bootloader). Installing the ZBasic bootloader requires a special device programmer (e.g. Atmel AVRISP MkII) but once the bootloader is installed your application can be downloaded to the target device over the serial port.

The ZBasic bootloader also provides the necessary support for writing to Program Memory. For cases where the ZBasic bootloader is not present, other device parameters are provided to support writing to Program Memory for those targets and applications that need that capability.

Examples

```
--device-parameter=ZBasicBootloader,true
Option DeviceParameter ZBasicBootloader true
```

BootloaderAddress

BootloaderAddress <value>	Default: varies by device
----------------------------------	---------------------------

If the ZBasicBootloader parameter is true, the ZBasic compiler assumes that the bootloader resides at a specific address (which varies by target device). If you rebuild the ZBasic bootloader to reside at a non-standard address you must use this parameter (or BootloaderSize, below) to tell the compiler what that address is. The value, which may be given in decimal or hexadecimal (with either a 0x or &H prefix), should specify the byte address of the first location of the bootloader. The value may optionally include a k or K suffix to indicate a multiple of 1024. The examples given below all specify the same bootloader address. This parameter is not supported on target devices that reserve a fixed-size portion of Flash memory exclusively for a bootloader (e.g. xmega devices).

Examples

```
--device-parameter=BootloaderAddress,30720
--device-parameter=BootloaderAddress,30K
--device-parameter=BootloaderAddress,0x7800
--device-parameter=BootloaderAddress,&H7800
Option DeviceParameter BootloaderAddress 30720
Option DeviceParameter BootloaderAddress "30K"
Option DeviceParameter BootloaderAddress "0x7800"
Option DeviceParameter BootloaderAddress "&H7800"
```

BootloaderSize

BootloaderSize <value>	Default: varies by device
-------------------------------------	---------------------------

This parameter provides an alternate means of specifying the bootloader starting address. The bootloader starting address is computed by subtracting the stated size from the total Flash memory size of the target device. The examples given below all specify the same bootloader starting address. This parameter is not supported on target devices that reserve a fixed-size portion of Flash memory exclusively for a bootloader (e.g. xmega devices).

Examples

```
--device-parameter=BootloaderSize,2048
--device-parameter=BootloaderSize,2K
--device-parameter=BootloaderSize,0x800
--device-parameter=BootloaderSize,&H800
Option DeviceParameter BootloaderSize 2048
Option DeviceParameter BootloaderSize "2K"
Option DeviceParameter BootloaderSize "0x800"
Option DeviceParameter BootloaderSize "&H800"
```

WriteWordAddress, WritePageAddress

WriteWordAddress <value>	Default: 0
WritePageAddress <value>	Default: 0

If your target device contains a bootloader other than the ZBasic bootloader and your application writes to Program Memory, your bootloader must provide entry points for *write word* and *write page* functions and you must specify the addresses of those entry points using these device parameters. The C function prototypes of the service routines must match those given below.

```
void writeWord(uint16_t addrLo, uint16_t dataWord);
void writePage(uint16_t addrLo, uint16_t addrHi);
```

On devices with less than 64K bytes of Flash memory, the second parameter of the writePage function is omitted. Otherwise, the second parameter is the 64K page (zero-based) to which the previously constructed page should be written). The source code for the ZBasic bootloader (for which source code is provided) contains entry points which meet these requirements; they are named `writeFlashWord` and `writeFlashPage`.

Examples

```
--device-parameter=WriteWordAddress,0x812
--device-parameter=WritePageAddress,&H814
Option DeviceParameter WriteWordAddress "0x812"
Option DeviceParameter WritePageAddress "&H814"
```

clkCtrl, psCtrl, oscCtrl, xoscCtrl, plICtrl

clkCtrl <value>	Default: 0
psCtrl <value>	Default: 0
oscCtrl <value>	Default: 1
xoscCtrl <value>	Default: 0
plICtrl <value>	Default: 0

These parameters, usually required for and applicable only to xmega devices, are used to specify the clock system initialization values written to the corresponding xmega registers as shown in the table below.

ATxmega Register Mapping	
Parameter	xmega Register
clkCtrl	CLK_CCTRL
psCtrl	CLK_PSCTRL
oscCtrl	OSC_CTRL
xoscCtrl	OSC_XOSCCTRL
pllCtrl	OSC_PLLCTRL

Details regarding the appropriate values for these parameters can be found in the Atmel Xmega documentation. The example values shown below are those used by the xmega-based ZX devices that use a 14.7MHz crystal and run at 29.5MHz. They would work equally well with any 12-16MHz crystal frequency if an internal 2X operating speed is desired. The default values configure the xmega to run on the internal 2MHz RC oscillator.

Examples

```
--device-parameter=clkCtrl,0x04
--device-parameter=psCtrl,0x03
--device-parameter=oscCtrl,0x08
--device-parameter=xoscCtrl,0xcb
--device-parameter=pllCtrl,0xc8
Option DeviceParameter clkCtrl "0x04"
Option DeviceParameter psCtrl "0x03"
Option DeviceParameter oscCtrl "0x08"
Option DeviceParameter xoscCtrl "0xcb"
Option DeviceParameter pllCtrl "0xc8"
```

5.4 ZBasic Bootloader

The ZBasic Bootloader facilitates downloading of code to a generic target device without needing special programming hardware. Such hardware is needed, of course, for the initial installation of the bootloader but after bootloader installation application code can be downloaded without the programming hardware.

The ZBasic installer automatically installs the source code, build scripts and related files for the ZBasic Bootloader. These can be found in the zboot subdirectory of the ZBasic installation directory. Pre-built bootloaders for the supported devices are also installed (see zboot/boot). The pre-built bootloaders are configured for devices running at 16MHz (important for ATtiny and ATmega devices only) and a download speed of 115.2K baud. If you want a different download speed or your ATtiny/ATmega runs at a different speed, you'll have to rebuild to get a compatible bootloader image.

The easiest way to build a bootloader image is to use the batchfile makeboot.bat. Near the beginning of that file you'll find several configuration items including the operating frequency (F_CPU), the download baud rate (BAUD) and the target device (DEVICE). For xmega devices, you'll need to set BAUD_SEL to select the baud rate; BAUD and F_CPU aren't used in that case. After making the desired changes, run the batchfile to effect the build; the resulting .hex file containing the bootloader will then be found in the boot subdirectory in a file named for the device, e.g. zboot\boot\atmega644p_boot.hex for the ATmega644P.

If you are using Windows Vista or later, you'll need to copy the bootloader source directory from the installation directory to some other directory outside of the "Program Files" system directory because those later Windows versions do not allow applications to change files in the "Program Files" directory tree.

Chapter 6 - Special Considerations for Native Mode Devices

With the native mode devices like the ZX-24n (and all generic devices such as the Atmel ATmega1284p), the ZBasic compiler produces native object code that is executed by the microcontroller. This is in contrast to the VM mode devices like the ZX-24p where the ZBasic compiler produces instructions that are executed by the ZX virtual machine running on the underlying microcontroller. Because the native mode code runs directly on the underlying microcontroller instead of the more controlled virtual machine environment, there are several additional features that you can employ in your application. Moreover, there are several additional issues of which you must be aware as you write the code for your application or when you port your code from a VM mode device to a native mode device.

6.1 Using Inline C and Assembly Code

When compiling ZBasic code for a native mode device, the compiler first produces equivalent C code corresponding to the subroutines, functions and data definitions of each module. Then, the resulting code is compiled into native object code for the underlying processor and linked with the ZX Runtime Library to produce an executable file that can be downloaded into the target device. In some cases it may be useful or desirable to inject sequences of C or native assembly language code directly into the code stream that is produced by the ZBasic compiler.

The mechanism provided for doing this is similar for the two cases – the code lines to be injected are bracketed by lines containing the special directives #c and #endc or #asm and #endasm. All of the text lines between the beginning marker and the end marker, but excluding those marker lines, are copied verbatim to the output file corresponding to the module being compiled. It is important to note that no syntax checking or other analysis is performed on the code that is copied so you must be sure that the code is syntactically and semantically correct.

Inline code sequences that occur outside of any subroutine or function are included in the module's output file after the module's data definitions but before any of the generated subroutine/function definitions. Inline code sequences that occur inside of a subroutine or function are placed in the generated code at the corresponding position for the routines to which they belong.

Consider this simple example using inline C code:

```
Public b as Byte
#c
char ch;
#endif

Sub Main()
    b = 5
#c
    ch = zv_b;
#endif
End Sub
```

Note, particularly, that ZBasic public variable names are prefixed by `zv_` in the generated C code. This is done to avoid namespace conflicts that might occur if the ZBasic variable names were used in the generated C code without a prefix. Other elements of the ZBasic program have distinguishing prefixes as well, as shown in the table below.

Accessing ZBasic program elements using the prefixed names in inline C code is fairly straightforward. However, you must keep in mind that the ZBasic compiler may eliminate some variables that you define in your program if they are not used or if the ZBasic compiler's optimizer determines that they aren't needed. If you want to ensure that the ZBasic compiler will not eliminate a variable you've defined you can define it with the Volatile attribute. That attribute is a signal to the ZBasic compiler to not make any assumptions about the value of the variable at any given point in time.

Program Element Prefixes

Element	Prefix (if Public)	Prefix (if Private)
subroutine or function	zf_	mzf_
variable	zv_	mzv_
structure	zs_	mzs_
structure member	zm_	zm_
parameter		zp_
function return value		zr_

Accessing ZBasic program elements in inline assembly code is possible but is more complicated partly due to issues regarding the register allocation, code generation and optimization strategies employed by the compiler that translates the generated C code to native object code, discussion of which is beyond the scope of this document. That said, here is a simple example using inline assembly code.

```

Public b as Byte
#c
char ch;
#endc

Sub Main()
    b = 5
#asm
    ; save the registers used
    push    r30
    push    r31
    push    r24

    ; load the address of the variable
    ldi     r30, lo8(zv_b)
    ldi     r31, hi8(zv_b)

    ; load the variable's value and save it
    ld      r24, z
    sts    ch, r24

    ; restore the registers used
    pop    r24
    pop    r31
    pop    r30
#endasm
End Sub

```

It is important to note that inline assembly code must avoid altering certain registers that the compiler expects to remain unchanged. (The same is true, of course, of any assembly language code.) The code in the example above saves registers that it uses on the stack before using them and then restores their values afterward. There is a mechanism to inform the compiler as to which registers are altered in the inline assembly code thus allowing the compiler to automatically perform the save/restore if necessary. Discussion of that mechanism is an advanced topic that is beyond the scope of this document.

6.2 Defining and Using External Subroutines, Functions and Variables

For the native mode devices you may include .c files (C language), .S files (AVR assembly language), .o files (AVR object code) and .a files (AVR object code archives) in your project. When processing the component files of a project, the ZBasic compiler will note the presence of these special files but will otherwise ignore them except for including them in the final build phase. In order to invoke subroutines and functions contained in these special files, you must declare them so the ZBasic compiler knows how to check for invocation syntax errors and how to generate code to invoke them.

The syntax for declaring external routines is very similar to defining a ZBasic subroutine or function as shown below.

```
Declare Sub <name> ( [<parameter-list>] )  
Declare Function <name> ( [<parameter-list>] ) as <type>
```

Parameters of the external routines may be given a default value as described in Section 3.18. In cases where the external routine name is a ZBasic keyword or is not a valid ZBasic identifier, you can use the Alias clause to specify the actual name of the external routine as a quoted string while using the specified <name> to refer to it in ZBasic code.

```
Declare Sub <name> ( [<parameter-list>] ) Alias "<ext-name>"  
Declare Function <name> ( [<parameter-list>] ) as <type> Alias "<ext-name>"
```

Note that, with the exception of the String type, ZBasic data types are generally compatible with their similar counterparts in C and assembly language. To pass a string to an external routine, it is recommended that the string's characters be copied to a Byte array and then pass the array by reference to the external routine. Depending on the requirements of the external routine, the length of the string must also be passed or the string's characters must be null-terminated in the Byte array.

Example

```
Declare Sub extSub(ByVal ival as Integer)  
Call extSub(5)
```

The Declare keyword may be preceded by Public or Private. If neither is present, the declaration defaults to being public meaning that it will be visible in other modules. The primary use for declaring a private external routine is when you define a routine using inline C or assembly code in the same module.

Similarly, external variables may be declared so that they may be accessed in ZBasic code.

```
Declare <name> As <type> [ Alias "<ext-name>" ]
```

Likewise, for an external array:

```
Declare <name>(<dim-spec-list>) As <type> [ Alias "<ext-name>" ]
```

Making these declarations does not result in any variable space being allocated. It is up to you to provide additional code to make the external variables available. As with subroutine and function declarations, variable declarations may be Public or Private. In the absence of either attribute, a data declaration defaults to being private.

For external variables that are to be treated as constant, you append the Constant attribute to the declaration. With this attribute, the compiler will generate an error if you attempt to modify the variable or to pass it to a subroutine or function by reference. Note, also, that the external variable should actually be defined with the const attribute if it is defined in C code.

Example

```
Declare index as Integer Attribute(Constant)
```

If the external variable declaration has both an Alias and an Attribute specified, they may occur in either order. For compatibility with the Declare statement in Visual Basic, the Alias specification for an external subroutine or function declaration may be placed between the routine name and the opening parenthesis of the parameter list. The form shown above, however, is the recommended syntax because it is more consistent with the other uses of the Alias keyword.

For passing values to and receiving values from external functions, the CType() function may be useful for casting to/from ZBasic values. See the description of CType() in the ZBasic System Library Manual. Also, note that if you have header files for the external code modules you may find it more convenient to import the identifiers directly from the header file. See section 6.8 for more information.

It is also useful to note that if you include C, C++ or assembly language files in your project, when the compiler is invoked to process them the special identifier ZBASIC_APP is defined. You may use this identifier in conditionals to control the features or configuration of the included code.

6.3 Defining Interrupt Service Routines

For the native mode devices you may write special-purpose code to service hardware interrupts. This may be useful, for example, to add some interrupt-driven capability to your program that is not directly supported by ZBasic. The syntax for defining an interrupt service routine (ISR) is similar to that for defining a subroutine, illustrated here by example. Note that the definition of an ISR does not allow parameters but the parentheses are, nonetheless, required.

```
ISR Timer1_CompB()
  ' place the ISR code here
End ISR
```

The ISR name following the ISR keyword must be valid for the underlying processor. The set of valid ISR names may be found in one of the XML resource description files located in a subdirectory of the ZBasic installation directory.. See Appendix N for more information on the resource description files.

It is important to note that some ISRs are added to your application automatically depending on the set of ZBasic System Library routines that it uses and whether certain resources like the RTC and serial channels are used. See the ZBasic System Library Reference Manual for more information on the various ISRs that can be provided automatically. If you attempt to define an ISR having the same name as one that is automatically added the compiler will issue an error message indicating the attempt to define a duplicate ISR.

For the form of ISR definition shown in the example above, the compiler takes care of saving the necessary registers upon entry, establishing the standard register state, restoring registers upon exit and executing a “return from interrupt” instruction. For special cases, you may define a “naked” ISR – one in which none of this is done. Using this special form is advised only for advanced programmers, particularly those that understand the nuances of the underlying code.

An example of a naked ISR containing only a “return from interrupt” is shown below. This is commonly called a “stub” ISR. It is important to note that defining a naked ISR with no code statements in it at all results in undefined behavior. At a minimum, you should always include an assembly language “reti” as shown in the example below.

```
ISR Timer1_CompB() Attribute(Naked)
#asm
  reti
#endasm
End ISR
```

Sometimes, it is desirable to have the same ISR service multiple interrupts. This can be accomplished by defining an ISR that is aliased to another ISR as shown below.

```
ISR Timer1_CompA()
  ' place the common ISR code here
End ISR

ISR Timer1_CompB() Alias Timer1_CompA
End ISR
```

You may also define a special ISR that serves all interrupts that otherwise have no ISR assigned. The default ISR is defined by using the interrupt name “default” as shown in the example below.

```
ISR default()
End ISR
```

If you do not define a default ISR, the compiler supplies one automatically. The function of the automatically-supplied default ISR is to store in Persistent memory a fault code that indicates that an unhandled interrupt has occurred and then execute a watchdog reset. The Persistent system variable Register.FaultType will have the value 2 to indicate that an unhandled interrupt occurred.

6.4 Executing Blocks of Code Atomically

Compared to writing code for the VM mode devices, when writing code for the native mode devices (e.g. ZX-24n) you must be more aware of the possibility that a sequence of statements may be interrupted either by an Interrupt Service Routine or by a task switch. Unless special techniques are used, another task or ISR may get control at any time in sequence of ZBasic statements – even in the middle of reading or writing the bytes of a multi-byte data item.

If a sequence of statements must be executed without interruption (commonly referred to as a “critical section” or “atomic execution”), the only way to guarantee such atomicity is to disable interrupts before the statement(s) and re-enable interrupts after the statement(s). Clearly, one must do this carefully so that interrupts are disabled for the least time possible in order to avoid missing a frequently occurring interrupt and to ensure timely servicing of important interrupts.

ZBasic supports several different techniques for disabling and re-enabling interrupts for one or more statements. Perhaps the simplest technique, and the one that is recommended for most purposes, is to use the atomic block construction. An atomic block is introduced by the keyword `Atomic` and is terminated by the keyword sequence `End Atomic` as illustrated in the example below.

```
Dim i as Integer

Sub Main()
    Atomic
        i = 200
    End Atomic
End Sub
```

Of course, if this were the entire application there would be no need to be concerned about atomicity of the assignment to the variable `i`. However, if the variable `i` is also read or written by an ISR or by another task, the access to the variable should be protected as shown above.

The second method for guaranteeing atomicity of a sequence of one or more statements is to use the System Library routines `DisableInt()` and `EnableInt()` in matched pairs as shown in the example below.

```
Dim i as Integer

Sub Main()
    Dim stat as Byte
    stat = DisableInt()
    i = 200
    Call EnableInt(stat)
End Sub
```

The disadvantage of using this method is that the compiler does not attempt to verify that the disabling and enabling are properly paired. Consequently, you might inadvertently omit the call to `EnableInt()` which would prevent any further interrupts, effectively disabling serial I/O, the RTC, etc.

A third method of guaranteeing atomicity is to explicitly disable and re-enable interrupts as depicted below.

```
Dim i as Integer

Sub Main()
    Register.SREG = &H00
    i = 200
    Register.SREG = &H80
End Sub
```

Although this technique may be useful in rare cases, its use is strongly discouraged. The primary problem with this method is that it unconditionally re-enables interrupts without regard to whether or not interrupts were enabled beforehand. A related technique using inline assembly language code, having the same disadvantage, is shown below.

```
Dim i as Integer

Sub Main()
#asm
    cli
#endif
    i = 200
#endif
    sei
#endif
End Sub
```

6.5 Attributes for Procedures and Variables

This section describes several special attributes that you can apply to variables or procedures with native mode devices to obtain special effects. To apply special attributes, you list the desired attributes (see the table below) in a comma-separated list enclosed in parentheses following the keyword `Attribute`. This entire construction is placed at the end of a normal variable, subroutine or function definition as illustrated in the example below.

Special Attributes

Attribute	Valid For	Description
Inline	subroutine function	Instructs the compiler to "inline" the code for the subroutine/function instead of generating a call. Generally, this yields faster execution, sometimes at the expense of larger program size.
NoInline	subroutine function	Instructs the compiler not to "inline" the code for the subroutine/function and, instead, generate a call to the subroutine or function.
Used	subroutine function variable	Instructs the compiler to include the variable or procedure in the executable even if it appears not to be used. This is useful, for example if external C or assembly code needs to use a ZBasic data item or procedure but it is not otherwise used in the program.
"Alias:<name>"	subroutine function variable	Instructs the compiler to create an alias for the procedure or variable having the specified <code><name></code> . This is useful on native mode devices to allow external code to refer to a ZBasic procedure or variable by a pre-determined name.
Static	variable	Instructs the compiler allocate space for the variable statically. This attribute is redundant for variables defined at the module level since that is the default for those. In contrast, variables defined within a procedure

		are normally dynamically allocated, i.e. they exist only while the procedure is executing.
Volatile	variable	Instructs the compiler to not make any assumptions about the content of the variable. This should be used when another task or an ISR may be modifying the variable.
ByteAlign	variable	Instructs the compiler to allocate the variable so that it begins on a byte boundary. This is useful primarily for the sub-byte types Bit and Nibble.
Raw	variable	Applicable only to initialized Program Memory data items, this attribute tells the compiler to interpret the data in the initialization file as raw (binary) data rather than as textual data.
"PortPinEncoding:<val>"	variable	Applicable only to initialized Program Memory data items, this attribute tells the compiler how to evaluate pin constants (e.g. A.0). The <val> may be 0, Off, 1 or On; the first two result in a physical pin number while the latter two result in an encoded port/pin designator.

Example

```
Dim index as Integer Attribute(Volatile)

Sub foo() Attribute(Used,Inline)
End Sub
```

6.6 Considerations for Task Stack Size

As of v2.5.6, the ZBasic compiler produces an estimate of the minimum size for a native mode application's tasks. The estimate is produced by performing a static analysis of the generated native code beginning with the entry point of each task (including the default Main() task). The static analysis determines, if possible, the maximum stack depth assuming that all possible execution paths through the code are followed. Additionally, each ISR (both user-provided and system-provided) is analyzed using the same technique to determine the maximum stack depth it will use.

For each task, the required minimum task stack size is the sum of a) the maximum stack depth of the task itself, b) the size of the Task Control Block, c) the stack margin (default or user specified) and d) the largest of 1) the multi-tasking context size and 2) the stack depths of the individual ISRs. At the time of this writing, the Task Control block size is 13 bytes, the multi-tasking context size is 37 bytes (38 on ZBasic devices with more than 64K of Flash memory), the maximum stack depth of the system supplied ISRs is less than 20 bytes and the default stack margin is 10 bytes. This leads to a typical task stack size minimum that is the stack depth of the task itself plus 60 bytes (37 + 13 + 10) or 61 bytes (38 + 13 + 10).

The compiler outputs a list of the known tasks and the minimum task stack size for each to the .map file that is produced each time a project is compiled (unless it is explicitly requested otherwise). In most cases, the compiler stack depth estimation algorithm can successfully compute the maximum stack depth but there are certain situations where the stack effect cannot be deduced by examining the code. When this situation is encountered, the map file will list "???" as the minimum stack size in the task list contained in the .map file for the related task. In this situation, there are two possible alternative strategies. The first is to allocate a large task stack for the task in question, run the application (ensuring that as many of the execution paths as possible are taken) and then using the System Library function System.TaskHeadRoom() to determine how much unused space remains in that task's stack. Based on the results of this observation, you may choose to increase or decrease the size of that task's stack.

The second option is to explicitly inform the compiler of the maximum stack depth for the routine or ISR that contains the problematic code. This is done using the #pragma StackUse directive within the particular procedure or ISR, the syntax of which is shown below.

```
#pragma StackUse( <stack-depth> )
```

An alternate syntax that can be used outside of all procedures, but not for ISRs, is shown below.

```
#pragma StackUse( <procedure-name>, <stack-depth> )
```

The recommended method used to generate an estimate of the stack use is to examine the assembly listing file generated by the compiler in response to the `--list` option (which also requires the `--keep-files` option) and manually compute the maximum stack depth. Use of this method requires an understanding of AVR assembly code and is therefore recommended for advanced users only.

6.7 Creating and Using Object Libraries

It is often useful to create a library of code elements from which the essential elements are drawn when building an application. Such libraries can then be easily reused in several applications or distributed to other ZBasic users for use in building their applications. The ZBasic compiler supports the creation of object libraries only for native mode devices.

To instruct the ZBasic compiler to build a library instead of building a downloadable file as it normally does, add the option `--create-library` to your project file. If no filename is specified with that option, the library name will be same as the project name but with the extension `.a`. The compiler will create an additional file containing ZBasic declarations for all of the public elements in the library; this file will have the same name as the library but with the extension `.inc`. It is important to note that none of the ZBasic files compiled into the library should have a public subroutine named `Main()`.

To use the library in a ZBasic application, it is only necessary to add the name of the declarations file to that application's project file or, alternately, to include the declarations file in one of your modules (e.g. the `Main()` module) using the `#include` directive. Generally, the declarations file and the library file should be kept in the same directory. The declarations file contains a directive to tell the compiler the name of the corresponding library file – that's why it doesn't need to be explicitly mentioned.

If you create library files to distribute to others, you only need to supply the declarations file and the library file. The source code for the library needn't be distributed if you'd like to keep it proprietary.

6.8 Importing Identifiers from External Modules

One of the benefits of using a native mode device is that your application can incorporate code written by others even when code is not written in ZBasic. This can be accomplished by incorporating pre-compiled libraries or object files. Alternately, it can be done by including C, C++, or assembly language source code modules in your project.

The most direct way of providing access from ZBasic to variables, procedures, etc. in external modules is to write Declare statements (as described in Section 6.2) for each of the items to which access is desired from ZBasic. While this method works in most cases (but not, for example, for external C++ classes) it can be cumbersome and prone to error. A more convenient method, one that also supports accessing external C++ classes, is to import the definitions directly from an external header file (typically, a `.h` file). The syntax importing a header file is given below.

```
#import [<lang-id>] [Public / Private] "<import-file>" [<id-list>]
```

The optional `<lang-id>` element may be either `C` or `CPP` (not case sensitive) to indicate the C and C++ languages, respectively. If not specified, C++ is assumed. Ordinarily, it will not be necessary to specify the language. However, if you have header files that are designed with one set of features to be available when included in C code and a different set when included in C++ code you may wish to specify importing the header file specifically as a C language header.

If neither `Public` nor `Private` is specified the importing is done privately meaning that all imported identifiers will be private to the module. The optional `<id-list>` is a comma-separated list of identifiers

to be imported from the header file. Identifiers to exclude may be specified by adding a dash prefix. If the identifiers are not legal ZBasic identifiers (under the prevailing "name style") they must be enclosed in double quotes. If the optional *<id-list>* is present, only identifiers present in the list will be imported along with any other identifiers (e.g. structures, classes, etc.) necessary for the specified identifiers. If the optional *<id-list>* is omitted, all identifiers present in the header file will be imported, including identifiers in other header files included by the specified header file. Note that no error message will be issued for specified identifiers that are not found in the header file.

If the *<id-list>* contains the special string "#all-define", all identifiers with numeric values created in the imported header files with the #define directive will also be imported. Alternately, specific individual #define identifiers may be imported by including "#<id>" in the *<id-list>*, e.g. "#foo".

In its simplest form, the *<import-file>* specification takes the form of the name of a header file in either absolute or relative form. If the header file is specified as a relative pathname (i.e. one that does not begin at the root directory of a drive), the pathname is assumed to be relative to the directory containing the module containing the #import directive. The directory containing the imported header file is automatically added to a list of directories to search for other included files but note, however, that that search list addition is in effect only for that particular importation. If one or more other directories are needed in the include file search list for a particular importation they must be specified using the #include_path directive (described in Section 3.13) prior to the occurrence of the related import directive. Note that on Window-based computers, forward slash and backslash can be used interchangeably in both the #import directive and in the #include_path directive.

Examples

```
#import Private "ethernet.h" IPAddr  
#import C "C:/projects/MatrixMath/MatrixMath.h"
```

Two useful variations of specifying the *<import-file>* are provided that may simplify the importing process in some cases. These alternate methods are triggered by specifying either an asterisk (*) or a commercial at sign (@) in place of the header filename at the end of the pathname. The two examples below illustrate the syntax.

Examples

```
#import C "C:/projects/MatrixMath/@"  
#import C "C:/projects/MatrixMath/*"
```

The first example, using the commercial at sign, is equivalent to specifying the import file as C:/projects/MatrixMath/MatrixMath.h, i.e. it specifies the importation of the header file with the same base name as the directory name immediately preceding the final slash. Moreover, it has the additional effect of adding all of the immediate subdirectories of the last-specified directory to the include path for that importation.

The second example, using the asterisk, specifies the importation of all header files (*.h) residing in the last-specified directory. Moreover, it has the same additional effect as the commercial at sign form of adding all of the immediate subdirectories of the last-specified directory to the include path for all of those importations.

When an include file is imported, it is processed by the C preprocessor just as it is when the code is actually compiled. The include files may contain conditional sections that control what features are available and, moreover, may require certain identifiers to be defined to control the configuration. When the preprocessor is invoked a the special identifiers ZBASIC_IMPORT and ZBASIC_APP are defined that can be used to further control configuration. In particular, you may add a conditional section in a central include file to define the special configuration identifiers. An example is shown below.

```
#if defined(ZBASIC_IMPORT)
    // add necessary definitions here
#endif
```

It is important to be aware of several key issues related to importing. Firstly, importing a header file has the sole effect of making the imported identifiers available to the ZBasic code; it does not incorporate any of the code related to the include files in the application. You must add the names of the necessary source code files, object code files and/or object library files in order to make the necessary code become part of the application. The simplest way to include the actual external code is to specify the required external files (e.g. extension .c, .cpp, .o or .a) in your ZBasic project file.

The second important issue to keep in mind is that, generally speaking, the identifiers in external files are case sensitive while identifiers in ZBasic are case insensitive. This can cause problems when you need to access an external identifier that differs only in case from an existing ZBasic identifier. One way to resolve this type of problem (and identifier clashes generally) is to import the external identifiers into a ZBasic namespace. The downside of this solution is that the namespace prefix will need to be prepended to all top-level imported identifiers when used in ZBasic code. See section 3.27 describing namespaces for more information on this topic.

The third important issue arises when imported names are not legal ZBasic identifiers. This can arise, for example, because C/C++ identifiers must begin with an alphabetic character or an underscore and contain only alphabetic, numeric and underscore characters. Normally, a ZBasic identifier has the same form except that it may *not* begin with an underscore. To successfully import identifiers beginning with an underscore you must ensure that the ZBasic compiler is using the C name style (which is now the default setting).

The remaining important issue is related to importing C++ classes. Normally, during the compilation process the ZBasic compiler produces C code that is equivalent to the ZBasic code. However, it is not possible to access C++ classes, namespaces and other features specific to C++ code in C code. The solution to this is to direct the ZBasic compiler to produce C++ code instead of C code thus allowing reference to the imported C++ elements. This can be accomplished using either the `Option CodeType` directive or the `--code-type` command line option.

As a final note, it is possible to specify files to be imported using the command line options `--import` and `--import-c`, which options may appear in your project file. The command line option for specifying the include path to use while importing is `--include-path`.

Chapter 7 - Compatibility Issues

The ZBasic language and compiler were designed to be backward compatible with BasicX and to offer some of the advanced features of Visual Basic along with some capabilities not found in either VB or BasicX. A design goal was that any BasicX program that compiles error-free would be compiled without errors by the ZBasic compiler in BasicX compatibility mode. The converse may or may not be true. If you write a ZBasic program carefully avoiding those features that aren't supported in BasicX, you can expect that it will be compiled error-free by the BasicX compiler. Similarly, you can most likely write a Visual Basic program that can be successfully compiled by the ZBasic compiler if you carefully avoid those ZBasic features that aren't supported in Visual Basic as well as those VB features that aren't supported in ZBasic.

Another design goal was that, with a few unavoidable exceptions, a program that can be successfully compiled by both the BasicX compiler and ZBasic compiler (using BasicX compatibility mode) will exhibit substantially the same behavior when executed on the respective processors. The target behavior is as described in the BasicX documentation. However, in cases where the BasicX documentation is ambiguous, incorrect or incomplete, tests were developed to ascertain the actual behavior. Unless it was impossible, impractical or inadvisable to do so, the empirically determined behavior was implemented in order to maximize backward compatibility.

When porting a BasicX application to ZBasic it may be best to use BasicX compatibility mode at first. Once you get the program working in that mode you can then switch to not using compatibility mode in order to be able to utilize the advanced capabilities of ZBasic. Of course, turning off compatibility mode may cause your program to behave differently requiring some additional modifications to the program.

7.1 Known Differences and Compatibility Between ZBasic and BasicX

- ZBasic programs are not object code compatible with BasicX. Existing programs will need to be recompiled specifically for the desired ZX target.
- Some BasicX programs rely on I/O pin configuration performed by the Project | Chip dialog in the BasicX IDE. This configuration information is stored in a .prf file rather than appearing directly in the source code so it is easy to overlook. Porting such a program to ZBasic will require the addition of Option Pin or Option Port directives or the addition of equivalent configuration statements to the source code.
- Since most ZX microcontrollers run at twice the clock speed of the BX-24 your code will execute faster. This may reveal latent timing problems that existed in your BasicX application but that never before caused a problem.
- Due to the fact that the CPU runs twice as fast, the time units for `OutputCapture()` and `InputCapture()` are half as long. Depending on the particulars of your application it may work to simply apply the 0.5x scale factor to correct the values. Otherwise, you may be able to choose a different `TimerSpeed` setting to achieve the result that you need. Also, time divisors, including the serial port baud rate divisor must be twice as large to achieve the same effective rates.
- Also due to the faster clock speed, the SPI interface runs at twice the clock rate. If this is a problem, you may be able to select a different SPI clock multiplier to get satisfactory results.
- Some System Library routines are implemented differently on the ZX than they are on the BX-24. For example, on the ZX series the `ShiftIn()` and `ShiftOut()` System Library routines use Timer1 for speed regulation. This is not the case in BasicX. Also, the `X10Cmd()` routine will not return until the command has been transmitted. In BasicX, the documentation suggests that the routine returns immediately while the transmission is done in the background.
- In BasicX if the `CByte()` function is passed an `UnsignedInteger` argument that is larger than 255, the result is the low byte of the value. This is inconsistent with the handling of other types. In ZBasic, `CByte()` returns 255 when the parameter value is larger than 255 for all parameter types.

- The system-level details of ZBasic are likely to be different than those of BasicX. For example, the system information related to routine invocation, tasks, queues, etc. is probably not the same as that of BasicX so if your program relies on such information it is likely not to work correctly. Similarly, variable space may be allocated differently than it is for BasicX. If your code relies on certain variables being arranged in memory sequentially it may not work correctly. In general, any BasicX code that relies on internal implementation details needs to be examined carefully.
- By default, the ZBasic compiler omits superfluous code and does not allocate space for unused variables. There are compiler options to alter this behavior in case they are needed.
- The precedence of operators is different in ZBasic than it is in BasicX. The default operator precedence in ZBasic matches that of Visual Basic. However, in BasicX compatibility mode, the operator precedence matches that of BasicX. If you compile a BasicX program in ZBasic without using BasicX compatibility mode you may have to add parentheses to some expressions to achieve the intended effect.
- The Mod operator in ZBasic is sign-correct. In BasicX, the result of a Mod operation on a negative value is (incorrectly) a positive value.
- ZBasic supports Mod, multiplication and division operators on UnsignedLong types; BasicX does not.
- In ZBasic, enumeration types are represented internally using 16-bit values. In BasicX they are 8-bit values.
- In BasicX, all characters on an input line following a comment character are ignored. In ZBasic, an input line is examined to see if it ends with a continuation character before it is examined for comment characters. This means a line ending with a continuation character that is preceded by a comment character will not compile the same in ZBasic as in BasicX. The workaround is to delete the continuation character.

7.2 Known Differences and Compatibility Between ZBasic and Visual Basic

- VB implements run-time checks of various types including variable range overflow, array index bounds checking, stack overflow checking, etc. ZBasic checks string length only for allocated strings and fixed-length strings and has optional run-time stack overflow detection (for VM mode devices only).
- VB does not directly support UnsignedInteger and UnsignedLong data types. You can add such support to VB using classes but even then the characteristics of objects make them inherently different than using a native type.
- The VB behavior of "sign extending" hexadecimal values in the range &H8000 to &HFFFF (without the trailing ampersand) is not implemented (except in BasicX compatibility mode). This departure is implemented for compatibility with native unsigned types.
- The exponentiation operator in ZBasic is more limited in the types that can be handled as compared to VB. Also, in ZBasic the exponentiation operator is right associative compared to being left associative in VB. Right associativity for exponentiation is both more natural and more common in other languages.
- The only type designation characters recognized in ZBasic are the exclamation mark and pound sign. Both of these characters force type Single while in VB they force Single and Double respectively. ZBasic does not support type Double. Inasmuch as ZBasic may support Double types in the future, use of the pound sign type designator is strongly discouraged.

Chapter 8

Chapter 9 - The ZBasic IDE

The ZBasic integrated development environment is based on the open-source IDE called SciTE. In addition to the basic text editing capabilities, the IDE has the ability to invoke the ZBasic compiler to compile your application and help you navigate to the lines containing errors. It also has the ability to download the code into the ZX and display any output that results.

9.1 Using the Editor

The source code editor built into the IDE has features designed especially for editing source code. Before those features are described, we will cover the basic editing capabilities that should be familiar to anyone with prior Windows application usage experience.

9.1.1 Basic Editing

As you would expect, typing regular characters (e.g. alphabetic, numeric, punctuation, etc.) causes them to be inserted in the current document at the caret position. If a selection exists at the time, the characters comprising the selection will be first deleted and the typed character will be inserted in its place.

The cursor keys (Up, Down, Left, and Right arrows) move the position of the caret. If the either Shift key is held down when a cursor key is pressed, a selection will be begun if none currently exists and the selection will be extended by characters (for Left and Right) or by lines (for Up and Down). If the Ctrl key is also held down, the Left and Right cursor keys will extend the selection by a word at a time.

While a selection exists, Ctrl+C will copy the selected characters to the clipboard while Ctrl+X will cut the selected characters to the clipboard. Ctrl+V will paste the current clipboard content into the document at the caret position replacing the currently selected characters, if any.

The Home key will move the caret to the first non-space character on the line or, if the caret is already at that position, to the first column of the document. The End key moves the caret to just after the last character on the line whether or not the last character is a white space character. Ctrl+Home moves the caret to the position before the first character of the document while Ctrl+End moves the caret just after the last character of the document.

The PageUp and PageDown keys move backward a forward through the document by pages.

The Backspace and Delete keys remove the character preceding and following the caret position, respectively. Holding down the Ctrl key while pressing Backspace and Delete serve to magnify their effect, deleting to the beginning and end of a word, respectively. Similarly, holding down both the Ctrl and Shift key magnifies the effect even more; deleting to the beginning of the line and the end of the line, respectively.

The editor has the ability to undo (Ctrl+Z) and redo (Ctrl+Y) recent changes. These and other editing commands are available via menu entry as well.

9.1.2 Special Code Editing Features

The specialized “lexers” built into the editor recognize certain aspects of the code in a document based on the “extension” portion of the document’s filename. When a document having a .bas extension is loaded, the editor will apply different styles to various portions of the content. This capability, sometimes called “syntax coloring”, helps programmers more quickly identify the syntactic elements of their program and to recognize when a typing error has been made.

For example, the editor will display Basic keywords (e.g. If, Else, etc.) in a bold blue font. If you type “Ekse” when you intended “Else”, you’ll immediately see that the editor did not display the mistyped word as a keyword thereby giving you immediate feedback that you erred.

9.1.3 Expand/Collapse

Another capability of the editor that is based on its ability to recognize the structure of the code is its ability to expand and collapse parts of the code. If you have a .bas file open that contains an If-Then-Else statement, you'll notice a small minus sign in the margin just to the left of the first displayed column on the line containing the If and the line containing the Else. If you left-click either one of these minus signs, you'll note that the display collapses to hide all of the lines in the corresponding section of code. At the same time, the minus symbol changes to a plus symbol indicating that there are "invisible" lines in the document. Left-clicking on the plus sign expands the corresponding code section. The expand/collapse effect can be invoked also by typing **Ctrl+Keypad-*** as well.

9.1.4 Auto-Completion

The editor has the ability to save you some typing by presenting a list of words having a prefix that matches a sequence of one or more characters that you have already typed. This ability, called auto-completion, comes in two distinct forms that differ in the origin of the list from which the matching list is built. The first form looks for matches in a list of words collected from the current document. This is useful for quickly inserting a variable, subroutine, or function name or language keyword that is already present elsewhere in your document. The second form looks for matches in a pre-defined list of ZBasic library routines and built-in constants.

The first form of auto-completion is invoked by typing the first few characters of an identifier that you know to be in the current document and then pressing **Ctrl+Enter**. If one or more identifiers exist having those characters as a prefix, they will be presented in a popup window near the caret. If multiple entries exist, the Up and Down cursor keys will move the selection up and down the list. At any time, you may press **Enter** to replace the prefix characters that you have already typed with the entry that you have selected. If you don't wish to use any of the items in the list, you may simply continue typing or you may press the **Escape** key. In either case, the popup window with the matching word list will disappear.

The second form of auto-completion is invoked by typing the first few characters of pre-defined library function or built-in constant and then pressing **Ctrl+I**. If one or more pre-defined entities exist having those characters as a prefix, they will be presented in a popup window near the caret. As before, if multiple entries exist, the Up and Down cursor keys will move the selection up and down the list. At any time, you may press **Enter** to replace the prefix characters that you have already typed with the entry that you have selected. If you don't wish to use any of the items in the list, you may simply continue typing or you may press the **Escape** key. In either case, the popup window with the matching word list will disappear.

In either case after the popup appears you may continue typing characters of the intended word and the list will be refined to include only those entries matching the modified prefix.

9.1.5 Call Tips

Another aid that helps to more quickly compose correct code is the Call Tip feature. When you type the name of a function or subroutine that is known to the system and then type a left parenthesis, a popup window will appear that will contain information about the number and type of parameters expected by that routine. In some cases, there may be multiple ways to invoke the routine, usually because some of the parameters are optional or a given parameter may be of two or more significantly different types (e.g. Single and Integer). In these cases, the Call Tip window will display a dark rectangle containing an upward pointing arrowhead of the left side of the popup. Clicking on the rectangle will display the preceding alternative while clicking to the right of the rectangle will display the next following alternative.

The content of the parameter lists, if any, is intended to remind you of the purpose and allowable types for the parameters. In most cases, the word **Integer** implies that the expected value is the ZBasic type **Integer**. Note, however, that the compiler is often somewhat lenient in its type checking for built-in routines and it may allow, for example, an **UnsignedInteger** to be used instead. The type descriptor "AnyIntegral" means any integral type such as **Byte**, **Integer**, **Long**, etc. The descriptor "AnyNumeric" includes the same types as **AnyIntegral** but also includes the type **Single**. The descriptor "AnyType"

means just that. Note that the parameter descriptions also include the `ByVal` and `ByRef` keyword. These tell you the parameter passing convention that is used for each parameter.

In some cases, a routine requires a parameter that is an array. In ZBasic code this is indicated by the presence of left and right parentheses following the parameter name. For technical reasons, this is indicated in the Call Tip using left and right square brackets in place of the parentheses.

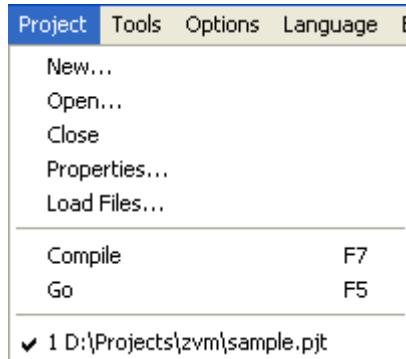
Call Tips may be added for your own routines. In addition to the file describing the system library routines (`zbasic.api`) the IDE will read a file named `ZBasicUser.api` located in the current user's "home" directory, if such a file exists. On Windows systems, this is usually the same directory that contains your "My Documents" folder. You can add information for your own Call Tips to `ZBasicUser.api` following the examples in the `zbasic.api` file. The disadvantage to using this method is that the same Call Tip information will be used for all your ZBasic projects. You can add project-specific Call Tips by creating an API file having the same name as the project file but with an extension of ".api". For example, for a ZBasic project named `C:\test\myproj.pjt` the project-specific API file is named `C:\test\myproj.api`. Although you can create and maintain this file manually, the compiler has the ability to generate the API file each time the compiler is run. See the description of the compiler option `--api` in Section 10.2 for more information.

One item that may need further discussion is the presence of a dollar sign (\$) character in the API files. This optional element causes the displayed Call Tip to have a line break at that position (the dollar sign itself is not displayed in the Call Tip text). This will usually only be used when the Call Tip text is longer than a certain number of characters, 75 for example. You may edit the `zbasic.api` file as you wish to add, move or delete the line break characters.

It is important to note that the Call Tips that are presented for a given routine will be the aggregate of the matching routines found in `zbasic.api`, your `ZBasicUser.api` file and the project-specific API file.

9.2 Project Configuration, Compiling and Downloading

The Project menu, depicted below, contains several entries for managing projects including creating a new project, loading an existing project, modifying an existing project, compiling a project and downloading the compiled code for a project.



Most of the dialogs invoked by entries on the Project menu contain a Help button that can be used to obtain more information about using the dialog. When you create a new project using the "New..." menu entry you will have the opportunity to specify the filename for the project file and the folder where it will be created. The dialog also has an option for automatically creating the first module for the new project. If you elect to do this, the filename of the module will be the same as the project filename but with a .bas extension. The newly created file will have an empty `Main()` subroutine to which you may begin adding code.

To add files to the project that is currently loaded, select the "Properties..." entry on the Project menu. The resulting dialog has a button for adding files and one for deleting files. The latter button will be enabled only when one or more entries in the project file list is selected.

The “Load Files...” menu entry is for loading constituent project files into the editor. The resulting dialog presents a list of files that are currently part of the project and you may select one or more from the list to load into the editor. Of course, you may also load files into the editor at any time using the “Open” entry on the File menu.

The “Compile” entry on the project menu will invoke the ZBasic compiler instructing it to process the project file. Any error messages and warnings issued by the compiler will appear in the Output window and the editor cursor will be positioned automatically to the first error or warning. After addressing the problem you can move to the next error using the F4 key or by selecting the “Next Message” entry on the Tools menu. Once the project compiles without errors, you may download the resulting .zxb file by choosing the “Go” entry from the Project menu.

The list on the bottom of the Project menu will contain up to 10 recently used projects. The project currently loaded will have a checkmark beside it on that list.

If you wish to have your project compiled with certain compiler command line options you may manually add those options to the project file. For example, if you wanted to enable warnings about unused parameters you might add the following line to your project file:

```
--warn=unused-param
```

Generally speaking, you'll want to add such lines at the top of the project file so that they will be in effect for all of the subsequently listed files. See Section 10.2 for more information on compiler options.

9.3 Compiling and Downloading Individual Files

In addition to being able to compile an entire project, the IDE can compile an individual file. If the currently selected document is a .bas file, you may compile that file by using the “Compile” entry on the Tools menu. This method is suitable for small programs that exist in a single file; it's not necessary to create a project for such a simple application. If the compiler detects any errors, they will be displayed in the Output window that appears near the bottom of the main application window. Note that the first error or warning message will automatically be selected and the cursor will be positioned automatically to the offending line. After you rectify the error, you can quickly move to the next error or warning by pressing the F4 key or by selecting the “Next Message” entry on the Tools menu.

When the file compiles without errors, a download file will have been created in the same directory as the source file but having a .zxb extension. To download this file to the ZX, select the “Go” entry from the Tools menu.

It is important to note that the “Go” entry on the Tools menu will attempt to download a file having the same name as the current document but with a .zxb extension. You may also download a previously compiled file at any time using the “Download...” entry on the File menu.

9.4 Setting Serial Port Options

By default, the serial port used for communicating with the ZX is COM1. The serial port to use may be changed by using the “Serial Port Options...” entry on the Options menu. The combo box on the resulting dialog will contain entries for all of the recognized serial ports on your system

The serial port options dialog has other options that you may wish to modify as well including whether or not a verification pass is performed after the download.

All of the options configurable via the serial port options dialog are stored in the User Options File. You may load this file into the editor using the “Open User Options File” entry on the Options menu.

9.5 Setting Device Options

By default, the compiler generates code for the ZX-24 device. The target device may be changed either by adding a command line option to your project file or by adding a compiler directive to the first module to be compiled (usually the one containing the Main() routine). Alternatively, you may specify the target device by using the “Device Options...” entry on the Options menu. The resulting dialog contains a combobox with entries for the supported devices.

Because of the large number of ZX and generic target devices that are supported, not all devices are initially displayed in the combobox. To add or remove devices from the combobox, click the “Edit Target List...” button on the Device Options dialog. After doing so, another dialog will appear that contains two listboxes. The listbox on the left contains the device names that will appear in the combobox on the Device Options dialog and the listbox on the right contains the remaining supported device names. Device names may be moved from one listbox to the other by selecting them (multiple selections via ctrl-click and shift click are allowed) and then clicking on the appropriate arrow button between the listboxes to effect the move. Clicking the OK button on this dialog will make your changes appear in the Device Options dialog.

The Device Options dialog can also be used to update the firmware in your device or, for some ZBasic devices, to set configuration parameters like the Program Memory EEPROM characteristics or external RAM configuration.

9.6 Setting Target Options

For generic target devices (i.e. targets that are not ZX devices), several device parameters need to be specified to allow the compiler to generate correct code for your specific target configuration. When a generic device is selected as the target device, the “Target Options...” entry on the Options menu is enabled. Selecting this entry will display the target options dialog. The dialog will be populated with either a set of default values or, if available, the values that would be seen by the compiler if it were run to compile the current project.

The Help button on the dialog may be clicked to display an HTML file describing the dialog. Note that many calculations (such as baud rate and RTC accuracy) are deferred until the Apply button is clicked. At that time, a set of matching device parameters options is generated into an edit box in the dialog (expose this by clicking the More button in the lower right corner). The set of device parameters may be selected, copied to the clipboard and pasted into the desired file.

9.7 Setting the Downloader Command

For target devices that do not contain a standard ZBasic bootloader, you must specify a downloader command if you want the IDE to perform downloads to the device. The downloader command is placed in the User Options file (load it into the IDE by selecting “Open User Options File” on the Options menu) and has one of the forms shown below for a generic Atmel AVR device.

```
command.project.<device-name>.go=<command-line>
command.project.avr.go=<command-line>
```

The first form applies only to the specified device while the second form applies to all AVR devices. The IDE looks first for the first form for a particular device and, if not found, looks for the second form.

An example downloading command using the Atmel STK500.exe utility is shown below. (Although this command appears to occupy several lines it is actually one long command line.)

```
command.project.avr.go="c:\Program Files\atmel\avr tools\stk500\stk500.exe" -
g -ccom1 -ms -I1MHz -d$(target_device) -e -if$(project.base).hex -pf -vf -
ie$(project.base).eep -pe -ve
```

There are several aspects of this example command that are useful to discuss. Firstly, note the use of the “macro” `$(target_device)`. Before executing the command, this macro is replaced with the actual target name. For example, if the selected target were the mega644p, the macro would be replaced with ATmega644P. Additionally, the macro `$(project.base)` will be replaced by the full path of the project file but with the extension removed. This makes it simple, in most cases, to refer to the code and EEPROM image files in a way that should work with any project.

An example of a device-specific download command is shown below. This example illustrates how you can define a “macro”, in this case named `atprogram`, that can be used as the downloader command. This is particularly useful if you have several devices for which you want to use the same downloader command.

```
atprogram=atprogram.exe -t avrispmk2 -i isp -cl 1MHz -d $(target_device)
erase program -f1 --verify -f $(project.base).hex --format hex program -ee -f
$(project.base).EEP --format hex

command.project.ATmega1284P.go=$atprogram
```

Some downloader applications do not accept device names like ATmega644P, e.g. it may require using m644P instead. In such a case you can provide lines in the User Options file to map the device names to alternate names. For example, consider a fictitious “Super Downloader” application that requires names like m644P, m1284P, etc. For this case we can create mappings of the device names as shown below and then use a nested macro name to yield the mapped device names. In this particular case, the prefix “sd.” was chosen to distinguish mappings for the “Super Downloader” application from mappings for other downloading applications. The actual prefix to use is unimportant as long as it is unique.

```
sd.ATmega644P=m644P
sd.ATmega1284P=m1284P
command.project.avr.go=superDownloader.exe com1 -d$(sd.$(target_device)) -
f$(project.base).hex
```

Before executing the command line, macros in it, if any, are recursively expanded. On the first pass, `$(target_device)` is replaced with the current target (say, ATmega644P). Then, on the second pass the macro `$(sd.ATmega644P)` is replaced by m644P.

9.8 Downloading Without Using DTR Signaling

Normally, the process of downloading a compiled program to a ZBasic device (one containing a ZBasic-compatible bootloader or a VM) begins with the downloader toggling the serial port’s DTR line in a particular fashion, causing the ZBasic device to enter a special *command mode*. When the device enters command mode in this fashion, it sets the baud rate of the Com1 serial channel to 115.2K baud and waits for commands from the host. Meanwhile, the downloader changes the baud rate of the PC serial port to match the 115.2K baud rate and then begins sending the device a series of commands to effect the downloading of the compiled program. Finally, when the download has completed the downloader toggles DTR again causing the device to reset and begin running the downloaded program.

The downloading procedure described above works in most situations but there are some cases where it won’t work. In particular, if a PC serial channel doesn’t support DTR or doesn’t support toggling DTR quickly enough to be recognized as the “go to command mode” signal, then the procedure above cannot be used. Additionally, if the baud rate of the serial channel cannot be changed dynamically (for example when using a radio link like BlueTooth or XBee), the procedure above cannot be used because of the required baud rate changes.

In cases where the normal DTR-triggered downloading procedure cannot be used, an alternate procedure can be used to effect a download. The alternate procedure requires a special configuration of the downloader and a matching special configuration of the ZX device. Instead of toggling DTR, the alternate procedure initiates the download procedure by sending a pre-determined character (known as the ATN character) to the device. The ZBasic device responds to the receipt of the ATN character by entering

command mode (either automatically, or by way of the ZBasic System Library subroutine `ZXCmdMode`). Once in command mode, the downloader can optionally send a command to instruct the device to switch to 115.2K baud for the download; otherwise the downloading is done at the prevailing baud rate. Finally, once the download has completed, the downloader sends a command telling the device to reset thus commencing execution of the downloaded program.

The ATN character is set in the ZBasic IDE using the dialog invoked by choosing the Serial Port Options entry on the Options menu. In the “Download Options” groupbox, you’ll find a checkbox captioned “Use ATN character”. When that checkbox is checked, two more controls become enabled. The first is an edit box with an accompanying spin control that allows you to specify the desired ATN character (the range for which is 1 to 31 decimal). Below the enabling checkbox, another checkbox captioned “Change speed for download” allows you to control whether the download is done at the prevailing speed (unchecked) or at 115.2K baud (checked). If you are using a serial channel that does not support dynamically changing the baud rate, you should uncheck the latter checkbox.

On the ZBasic device side, enabling downloading by way of the ATN character can be accomplished in one of two ways. The simplest way is to use the directive `Option AtnChar` in the first module of your program. Once your program is compiled and downloaded (which, by the way, must be initially accomplished using the traditional DTR-based method), when the device receives the designated ATN character it will automatically invoke the command mode. The second method, which should rarely be needed, is to add code to your application to detect the specified ATN character. When the ATN character is received, your code should invoke the ZBasic System Library subroutine `ZXCmdMode()` either with no parameter or with a parameter value of `False`.

It is important to note that the “Port Speed” setting in the “Execution Options” group box must be set to match the baud rate of the Com1 serial channel of the ZBasic device. The default Com1 speed is 19.2K baud but the speed can be changed using the directive `Option Com1Speed` in the first module of your program.

9.9 Updating Device Firmware

The VM-mode ZX devices have firmware that is designed to be field-upgradeable. This allows the control program to be updated with newer versions as enhancements are made and problems are fixed. An update can be installed in a VM-mode ZX using the “Device Options...” entry on the Options menu. This will bring up a dialog that will allow you to navigate to and select the specially formatted firmware update file (typically with .zvm extension). Prior to using the file’s contents to update the ZX firmware, the file is checked to ensure that it contains a valid update image. Note, particularly, that separate update files are provided for each ZX model. You must be certain that you use the correct update file for the ZX device being updated. The filenames of the update files indicate the target ZX device and the firmware version number. Additionally, the first few lines of the update file contain the same information. The update files are standard text files and may be viewed by any text file viewer.

To perform a normal firmware update the ZX must be powered up and must be connected to your PC by a serial cable. On the 24-pin ZX devices, during the update process the red LED will be illuminated continuously and the green LED will blink at a rate of about twice per second. Note that since the LEDs are also connected to pins 25 and 26, if you have circuitry connected to those pins that would be adversely affected by the LEDs being activated you will want to turn off the check box labeled “Provide visual feedback during update”.

On the other ZX devices no I/O pins are used to indicate that a firmware update is in progress.

9.10 Setting the Debug Output Limit

By default, the debug window will only retain the last 100 lines of output. You can change this by selecting the “Open User Options File” entry from the Options menu. This will display the current contents, if any, of the User Options File. If one does not exist you may insert a line like the one below.

```
debug.line.limit=250
```

This will change the limit to 250 lines. The IDE will need to be exited and restarted after manually modifying the User Options File in this manner in order for the change to take effect.

9.11 Other Configurable Items

There are many aspects of the editor that may be configured. However, discussion of the means of making such changes and the impact of the changes is beyond the scope of this document. Those interested in investigating this topic further are directed to the “ZBasic Help” entry on the Help menu.

Chapter 10 - Compiler Guide

The ZBasic compiler is a modern design that employs some advanced optimization techniques to reduce code and data size while at the same time reducing execution time. The compiler operates in several stages including syntactic analysis, semantic analysis and code generation. During each of the first two phases error messages or warnings may be generated. Even though these are logically separate passes, the resulting error output lists errors by line number to facilitate easier resolution. Error messages may be generated during the code generation phases but this should generally not occur.

Note that the compiler may generate “spurious” error messages due to earlier detected error conditions. Consequently, fixing the cause of some earlier error messages may result in later messages being eliminated on the next compilation.

By default, error messages are routed to the stderr device meaning that they will appear on the console unless redirected to a file. Alternately, the compiler can be directed to output the error messages to a specifically named file. The error messages are formatted in a manner to allow an Integrated Development Environment or advanced text editor to automatically position the cursor at the next error message. The default error messages format is:

```
<filename>:<line-number>:<error-message>
```

where *<filename>* is the full name of the file containing the offending line, *<line-number>* is the line that contained the error (or, in the case of continued lines, the line number of the first continued segment) and *<error-message>* is a description of the detected condition. The *<error-message>* text includes an indication of whether the detected condition is an error condition or just a warning.

There is a class of error messages that reflect the occurrence of a set of conditions that weren’t expected and shouldn’t normally occur. If the compiler generates such “internal error” messages we would like to know the circumstances that produced them so that we can rectify the problem. There is probably not much that you can do to work around the problem in such cases but it may help to turn off all optimization (see Section 10.2 for details).

The ZBasic compiler is a “console program” meaning that it has no graphical user interface. The compiler is invoked via a command line that specifies the compiler options and the files to compile. One of the compiler options allows you to specify an “arguments file” that contains a list of options and/or files. Another option allows you to specify a “project file”. This is similar to specifying an arguments file except that the use of a project file affects some filename defaults as described in subsequent sections.

10.1 Compiler Invocation

The general form of the compiler invocation is:

```
zbasic [<options>] [<files>]
```

where *<options>* represents one or more compiler options and *<files>* represents one or more files to be compiled. Both of these command line elements are optional; if neither is specified, a summary of the command line options is displayed. On the Windows platform, a file specification may contain “wild card” characters. For example, the file specification “*.bas” refers to all of the files having the .bas extension in the current directory.

Note that command line options and filenames may appear on the command line in any order and may appear multiple times. The options and files are processed left to right and when a filename is encountered the file is processed (syntactically analyzed) using the options in effect at that point in time. To understand how the intermixing of options and filenames affects the entire process, it is necessary to be aware that certain options are applied to files immediately as the files are encountered while other options are not applied until all files have been processed. If you utilize one of the options in the latter group multiple times, only the last instance will have an effect. In the tables below, the options that are immediately applied to files as they are processed are denoted with an asterisk preceding the description.

After all of the command line options and files have been processed, the compiler proceeds to the next phase – semantic analysis. If no errors are detected and unless directed otherwise, the compiler then proceeds to the final phase – code generation. It is during this last phase that the code output file, optional link map file and optional listing file are produced.

Note that some settings made by command line options may be overridden for individual modules by the presence of option directives in those modules. See Section 2.3.1 for more information on option directives.

All options are case-sensitive. Filenames may be case sensitive depending on the host operating system. Space characters are not allowed within options except within filenames (where supported by the host operating system) and possibly within strings. The command line language of the host operating system most likely requires option values containing spaces to be enclosed in quote marks.

ZBasic Compiler Options

Long Form	Short Form	Description
--alloc-str={on off default}		*Specify the use of dynamic string allocation.
--allow-conditionals		*Allow the use of conditionals in project and arguments files.
--api[(<width>)][=<file>]		Request generation of an API file for the IDE.
--args=<file>	-fa<file>	Specify an arguments file to be processed.
--arduino[=<base-dir>]		Specify Arduino compatibility mode.
--array-base=<value>		Specify the default array base value.
--call-functions[={on off}]		Control how functions may be invoked.
--called-by-list		Request generation of a “called-by” list.
--calls-list		Request generation of a “calls” list.
--code-type=<code-type>		Specify a type of generated code.
--code-limit=<value>		Specify a limit for generated code size.
--create-library=[<file>]		Request building a library.
--device-parameter=<name>,<value>		Specify a target device parameter value.
--directory[=<directory>]		*Set the “current directory”.
--entry=<subroutine>		Specify an alternate entry point routine.
--error=<file>	-fe<file>	Specify an output file for error information.
--error-format[=<fmt-spec>]		Specify a format for error messages.
--error-limit=<value>		Specify a limit for error messages.
--gcc-opt[=<options>]		Specify options to pass on to gcc.
--heap-limit=<value>		Specify the heap growth limit.
--heap-size=<value>		Specify the size of the heap area.
--help	-h	Display a summary of the invocation syntax.
--help-all	-hh	Display a longer summary of the invocation syntax.
--help-optimize	-ho	Display a summary of the optimization options with an indication of the defaults.
--help-warning	-hw	Display a summary of the warning options with an indication of the defaults.
--import=<header-file>		Specify a C++ header file from which to import identifiers such as variables and procedures.
--import-c=<header-file>		Specify a C header file from which to import identifiers such as variables and procedures.
--include-path=[<path-list>]	-I[<path-list>]	*Specify a semicolon-separated list of directories to search for include files.
--language=<language>	-l<language>	*Specify the source language variant.
--keep-files		Request retention of intermediate files.
--list[=<file>]	-fl[<file>]	Request a detailed listing file.
--main-task-stack-size=<value>		Specify the size of the task stack for Main().
--map=<file>	-fm<file>	Specify an output file for the link map.
--namespaces		Enable definition and use of namespaces.

--notice={on/off}		Enable or disable output from #notice.
--no-analyze-stack-use		Suppress the estimation of task stack size.
--no-code		Suppress the generation of code.
--no-map		Suppress the generation of a map file.
--objects		Enable object-oriented extensions.
--optimize=<opt-type>	-o<opt-type>	*Assert or deassert optimization options.
--out=<file>	-fo<file>	Specify an output file for generated code.
--overload		Enable procedure overloading.
--project=<file>	-fp<file>	Specify a project to be processed.
--strict={on/off/default}		*Specify the compilation mode.
--string-size=<value>	-s<value>	*Specify the default string size.
--target-config[=<file>]		*Specify inclusion of a target configuration file.
--target-config-force[=<file>]		
--target-device=<target>		*Specify the target for code generation.
--temp-dir=<directory>		Specify the directory for intermediate files.
--use-batch-file		Use a batch file for native mode builds.
--verbose		Display output of the build process.
--version		Display the version number of the compiler.
--warn=<warn-type>	-w<warn-type>	*Enable or disable warnings by type.
--warning-is-error		*Treat all warnings as errors.
	-D<id>[=<value>]	*Define conditional symbol with optional value.
	-U<id>	*Undefine a conditional symbol.

The options are described in detail in the next section, ordered alphabetically by the long form name followed by those options that have no long form.

10.2 Compiler Options in Detail

--alloc-str={On|Off|Default}

This option specifies whether dynamically allocated string usage should be on, off or set to the default state for the selected language. See Section 2.3.1 for more information on the default state of this option for the supported language variants.

--allow-conditionals

This option activates support for using #if conditionals in project and argument files. Because of the fact that a pound sign introduces a comment in both of these types of files, support for conditionals in them is disabled by default. Alternately, you can enable support for conditionals in specific project or argument files by adding a special comment having the format shown below as the first line of the file.

```
#!allow-conditionals
```

--api[(<width>)][=<file>]

Using this option you can request that the compiler generate an API file for use by the IDE to display “call tips” for the routines in your program. If the =<file> is not present, the output is written to a file having the same base name as the first project file specified or, if none, the first module compiled, but with the extension “.api”. For example, if the project is myproj.pjt, the default API file will be myproj.api.

The optional <width> parameter, which must be enclosed in parentheses, specifies a nominal line width for the generated API definitions. A special line break character is inserted near multiples of the specified width to allow the IDE to display long call tips on multiple lines. The width used if none is explicitly specified is 80.

--args=<file>

You can use this option to specify the name of a file that contains additional compiler options and/or filenames, one per line. The content of the file is processed before processing additional command line options. Such argument files may not be nested but a project file may be specified within an arguments file and vice versa. Argument files may contain blank lines and comment lines (beginning with a pound sign or an apostrophe).

--arduino[=<base-dir>]

This option requests Arduino compatibility mode, setting the generated code type to C++, enabling object extensions and, optionally, specifying a base directory where the Arduino compatibility files may be found.

--array-base=<value>

This option specifies base for arrays that do not explicitly specify the lower bound. The specified value must be either 0 or 1.

--call-functions[={on/off}]

This option can be used to control whether ZBasic functions may be invoked as if they were subroutines. It is only permitted for native mode devices and then only for functions that return a fundamental type (Byte, UnsignedInteger, String, etc.) The presence of this option affects only those ZBasic source files processed after the option is encountered. It may be used multiple times if desired. Omitting the equal sign and value has the effect of enabling the feature.

--called-by-list

When this option is specified a “called-by” list is appended to the map file. The called-by list will indicate, for each subroutine or function, each routine that invokes it. This information may be useful in analyzing the impact of changes in your code. If no map file is generated this option is ignored.

--calls-list

When this option is specified a “calls” list is appended to the map file. The calls list will indicate, for each subroutine or function, each routine that it invokes. This information may be useful in analyzing the impact of changes in your code. If no map file is generated this option is ignored.

--code-limit=<value>

This option requests that the compiler compare the size of the generated code with the specified limit value. The limit value is specified in decimal, optionally using the suffix K or k to denote a multiple of 1024 bytes. If the code size exceeds the limit an error message will be generated. The default code limit is 0, which value disables the code size checking.

--code-type=<code-type>

This option is used to request that the compiler generate a particular type of code. The set of acceptable values for <code-type> depends on the type of the target device. For VM devices, the only acceptable code type is zvm. For native target devices, either c (the default) or cpp may be specified. The latter code type is useful if you want to include C++ code in your application.

Example

```
--code-type=cpp
```

```
--create-library[=<file>]
```

This option, useful only for native mode devices, requests that the compiler build an object code archive (object library) instead of building an executable file. If the library filename is not specified, the library file will have the name of the project file or the first module processed with the extension replaced by .a. The compiler will also create an additional file containing ZBasic declarations for all of the public entities in the library. The declarations file will have the same name as the library with the extension replaced by .inc.

```
--device-parameter=<name>, <value>
```

This option specifies the value for a target device parameter (useful only for generic target devices). See Section 5.3 for a complete list of device parameter names and the expected values.

```
--directory[=<path>]
```

This option specifies a directory that should be made the “current directory”. If the option is given without a specific directory, the directory is inferred from the project or arguments file being processed at the time. If no project or arguments file is being processed, from which the directory can be inferred, the option is silently ignored.

Generally, a filename that has a relative path prefix will be sought relative to the current directory (but see --include-path, below). For example, since the filename ir\test.bas is a relative filename it will be expected to be in the ir sub-directory of the current directory. In contrast, a file specified with an absolute filename like c:\projects\ir\test.bas will not.

```
--entry=<subroutine>
```

This option specifies an alternate entry point for the program. By default, the entry point is the subroutine Main().

```
--error=<file>
```

This option explicitly specifies the name for the error output file. In the absence of this option, the error output is sent to stderr. Depending on your operating system, you may have the ability to redirect stderr to a file.

```
--error-format[=<format-spec>]
```

This option explicitly specifies the format that should be used for outputting error messages. If the equal sign and format specification are missing, use of the default format specification will be resumed. The error format string specifies the string that precedes the descriptive error message. The format string may contain ordinary characters and escape sequences. The supported escape sequences are described in the table below.

Error Format Specification Escape Sequences

Escape	Description
%f	The file containing the error.
%l	The line number on which the error occurred.
%%	A literal percent sign (only needed to disambiguate).

The format specification may be enclosed in matching quote marks or apostrophes but this is necessary only if the format specification contains spaces. The default format specification is "%f:%l:". This causes messages to have the format required by the IDE for error file navigation.

--error-limit=<value>

This option specifies an upper limit on the number of error messages that will be generated. When the limit is exceeded, compilation will cease. The default error limit is 100.

--gcc-opt=<options>

This option allows you to specify additional options to be passed to the gcc compiler and linker when generating code for native mode devices. The option string given will be added to the command lines after the options supplied by the ZBasic compiler and before the filenames. This position allows you to override earlier options or to add to the options or both.

--heap-limit=<value>

This option, useful only for native mode devices, specifies the heap limit, beyond which the heap will not grow. See the discussion in section 6.6 for details on the effect of this option.

--heap-size=<value>

This option, useful only for native mode devices, specifies the size of the heap and, indirectly, the heap limit. See the discussion in section 6.6 for details on the effect of this option.

--help

This option causes the compiler to output a summary of the invocation options and then exit.

--help-all

This option causes the compiler to output more comprehensive information about the invocation options and then exit.

--help-optimize

This option causes the compiler to output information about available optimization flags and default settings, and then exit.

--help-warning

This option causes the compiler to output information about available warning flags and default settings, and then exit.

--import=<header-file>
--import-c=<header-file>

These two options direct the ZBasic compiler to process the named header files and import certain identifiers (e.g. procedures, variables, classes, etc.) thus making them accessible directly from ZBasic

code. See Section 6.8, Importing Identifiers from External Modules, for more details regarding importing identifiers.

--include-path=<path-list>

With this option you can specify a list of directories in which the compiler will look for files included using the #include directive, for Program Memory data initialization files and for files imported using the #import directive. The <path-list> element consists of zero or more directory names, each separated from the next by a semicolon. If any components of the directory name contain a space, the entire list may have to be quoted depending on your computer's operating system. Quoting is neither required, nor supported, when this option occurs in a project file or arguments file. Note that the current directory can be made part of the include path in the normal fashion by using a single period to represent it.

Example

```
--include-path=..\includes;.;C:\projects\zbasic\files
```

This example specifies an include path with three components. When you use an include directive with a non-absolute path like #include "lcd.bas", the first place that the compiler will look for lcd.bas is in the includes sub-directory of the parent of the current directory. If it is not found there, the compiler will next look in the current directory because of the presence of the period. Finally, the compiler will look in the directory C:\projects\zbasic\files. If the file could not be located in any of the directories of the path, an error message will be issued.

It is important to note that the #include_path directive in ZBasic source code will override the include path specified by this option. Also, see the discussion of #include_path for information about some special "macros" that can be used in the <path-list>.

--keep-files

This option requests that the compiler not delete the intermediate files that it creates during compilation. This is only useful for native mode devices. Unless otherwise specified (using the --temp-dir option), the intermediate files will be created in a subdirectory named zbTempDir in the same directory as the project file.

--language=<language>

This option specifies the target language for the modules subsequently processed. The values that may be specified for the <language> element are shown in the table below.

Language Option Values	
Value	Description
BasicX	Compile using BasicX compatibility mode.
ZBasic	Compile using native mode (the default).

--list[=<file>]

This option requests that a listing file be generated and specifies the filename for it. For VM mode devices, the listing file is similar to an assembly language listing, giving detailed information about the code that was generated. If the equal sign and filename are omitted, the listing is output to stdout. For native mode devices, the --list option causes an actual assembly listing to be produced (--keep-files is also needed, q.v.) and in this case no filename should be specified. The listing file will have and extension of .lss.

--main-task-stack-size=<value>

This option, useful only for native mode devices, specifies the size of the stack for the `Main()` task and, indirectly, the heap limit. See the discussion in section 6.6 for details on the effect of this option.

--map=<file>

This option explicitly specifies the name for the map file. In the absence of this option, the map file name is derived from either the project file, if specified, or the first file compiled. If an earlier or later option specifies no map file should be generated this option is ignored. Currently, no map file is generated for native mode devices.

--namespaces

This option enables definition and use of namespaces. This option is implied by `--objects`.

--name-style=<name-style>

This option specifies the identifier style that will be accepted. The default name style is C and the sole alternative is Basic.

--notice={on/off}

This option enables or disables output from the `#notice` directive. By default, `#notice` output is enabled.

--no-analyze-stack-use

When this option is specified, the compiler will not perform an analysis of the tasks in your application to determine the minimum stack size required. Consequently, you must make your own determination of the proper task stack sizes by whatever means is appropriate.

--no-code

This option causes the compiler to omit the code generation step. By implication, no map file will be generated either. This may be useful if all you want is a syntax check.

--no-map

This option causes the compiler to omit map file generation.

--objects

This option enables object-oriented extensions, procedure overloading and namespaces.

--optimize=<optimization-type>[,<optimization-type>...]

This option enables or disables specific types of optimizations. In most cases, you'll want to use the default optimization settings. This option provides for unusual circumstances where more control is needed over the optimizations performed. The optimization types are described in the table below. To disable an optimization type, add the prefix `no-` to the optimization type, e.g. `no-strength-reduction`. All optimization can be turned off using `--optimize=no-optimize`.

Optimization Type	Description
constant-folding	Expressions involving constants may be evaluated at compile-time. Note, particularly, that some System Library function invocations having parameters known to be constant may be replaced by the equivalent value. In some cases the same strategy may be applied to user-defined functions.
constant-propagation	The use of expressions involving variables known to be constant may be replaced with the constant value.
expression-order	Expressions may be rearranged to facilitate additional optimizations. Such rearrangement will never be performed across parenthetical boundaries if the option <code>preserve-parens</code> is specified.
inline	Small subroutines and functions may be generated in-line instead of generating a routine invocation.
optimize	Refers to all affirmative optimization types collectively. The only optimization type not included in this group is <code>preserve-parens</code> .
preserve-parens	The presence of this options restricts the use of expression optimizations to parenthetical boundaries.
strength-reduction	Reduction-in-strength optimizations may be performed, e.g. multiplication by a power of two replaced by left shift.
string-pooling	Code size is reduced by detecting identical strings. Each string appears just once in Program Memory but may be referred to in multiple places.
unreachable-code	Code that cannot possibly be executed may be eliminated.
unreferenced-code	Routines that are not used are not included in the generated code.
unreferenced-vars	Variables that are not referenced or are eliminated by optimization may not be allocated space.
use-identities	Expression complexity may be reduced by applying algebraic or logical identities. For example, <code>i * j</code> can be replaced by <code>i</code> if <code>j</code> is known to have the value 1.
useless-code	Code that is known to have no useful effect may be eliminated. For example, assigning a value to a local variable that is never used is considered useless.

The invocation option `--help-optimize` displays similar information and also indicates which optimizations are on by default.

--out=<file>

This option explicitly specifies the name for the file for the generated code. In the absence of this option, the output file name is derived from either the project file, if specified, or the first file compiled. If an earlier or later option specifies that no code should be generated this option is ignored.

--overload

This option enables procedure overloading. This option is implied by `--objects`.

--project=<file>

This option specifies the name of a project file. A project file is similar to an arguments file in that it may contain additional compiler options and filenames that are processed before processing subsequent command line options. Additionally, however, several other file names (e.g. map file, output file) are derived from the first-specified project's filename in the absence of other overriding options. Project files may contain blank lines and comment lines (beginning with a pound sign or an apostrophe). Project files may not be nested but an arguments file may be specified within a project file and vice versa.

```
--strict={On|Off|Default}
```

This option specifies whether strict syntax mode should be on, off or the default state for the selected language. See Section 2.3.1 for more information on the default state of this option for the supported language variants and the effects of strict mode.

```
--string-size=<value>
```

This option specifies the default string length, in decimal, for statically allocated strings in modules subsequently processed. It may be overridden in any module by the Option StringSize directive. See Section 2.3.1 for more information on how this value is used.

```
--target-CPU=<target>
```

This option is deprecated, use --target-device instead.

```
--target-config[=<config file>]  
--target-config-force[=<config file>]
```

The first form of this option is silently ignored unless the target device is a generic device. The second form forces processing of the option irrespective of the current target device. In both cases, if the option appears without the equal sign and filename the name of the configuration file is inferred by taking the full path of the project file (which must have been previously specified) and replacing the extension with .cfg. The specified or inferred file is processed as if it had been specified with the --args option, i.e. the file should contain additional command line options. Usually, the command line options in the configuration file will be limited to --device-parameter options but they are not restricted to being only those. The first form of this option is useful because it will be ignored when compiling for a ZX device (for which --device-parameter options are not allowed) but will be processed when compiling for a generic target device (for which --device-parameter options must be specified). The second form is useful when you want to specify the target device in the configuration file as well.

```
--target-device=<target>
```

This option specifies the target for which code should be generated. If this option is used, it must appear before any modules are compiled (the default target is zx24). The target device may alternately be specified in the first module compiled using the Option TargetDevice directive. Doing so will override the specification on the command line. A complete list of supported device names is given in Appendix B.

```
--temp-dir=<directory>
```

For native mode devices, the compiler generates several intermediate files in the process of compiling an application. Normally, those files are created in a temporary subdirectory of the directory containing the project file and then the directory and its content are deleted when the compilation is done. The name of the temporary directory is selected to avoid conflicting with any existing files and directories. The temporary files can be retained using --keep-files option. In this case the directory in which they will be created will be zxTempDir. A separate subdirectory will be created in this temporary directory for each project compiled. If the default name of the temporary directory is unsuitable, you may specify a different directory using the --temp-dir option. If the specified directory is relative, it's interpreted as being relative to the directory containing the project file.

--use-batch-file

For native mode devices, this option causes the compiler to generate a batch file containing commands to perform the back-end compile-link process instead of building using a makefile.

--verbose

This option, useful only for native mode devices, causes the output from the final build process to be sent to stderr. In the absence of this option, the output from the build process is captured in a file named build.log created in the temporary directory. (Use the --keep-files option to prevent the build log from being deleted after the build is complete.)

--version

This option causes the version number of the compiler to be sent to stdout. The compiler will then exit.

--warn=<warning-type>[,<warning-type>...]

This option enables or disables specific types of warnings. The warning types are described in the table below. To disable a warning type, add the prefix no- to the warning type, e.g. no-unused-param. All warnings may be disabled en masse using --warn=no-warnings.

Warning Type	Number	Description
array-bounds	9	Warn about constant indices on arrays being outside of the valid range.
calltask-byref	1	Warn about invoking a task that uses ByRef parameters.
case-overlap	10	Warn about the same value appearing in more than one case expression or range in a Select Case statement.
data-range	3	Warn about data values exceeding the capacity for the specified type.
for-loop-termination	11	Warn about a For loop that may not terminate properly.
hidden-data	5	Warn about data definitions hiding definitions at outer levels.
never-returns	8	Warn that a routine will never return (automatically suppressed for the entry routine and routines invoked using CallTask).
questionable-code	4	Warn about questionable coding practices.
structure-compare	12	Warn about comparison of structures containing allocated strings.
task-stack-size	13	Warn about insufficient task stack size.
undefined-variable	2	Warn about the use of a variable before a value is assigned to it or when at least one path through a function does not set the return value.
unused-parameter	6	Warn about an unused parameter.
unused-procedure	14	Warn about an unused procedure.
useless-code	7	Warn about code that will never be executed or has no effect.
warnings	1000	Refers to all warning types collectively.

The invocation option --help-warning displays similar information and also indicates which warning types are on by default.

--warning-is-error

This option causes the compiler to treat all warnings as errors.

-D<id>[=<value>]

This option defines an identifier that may be used in a conditional construct (see Section 3.9). If no value is specified, the identifier is considered to be an integral identifier and is assigned the value 1. The value may be specified as a decimal value or, when prefixed by &H or 0, as a hexadecimal value. Otherwise, the value is considered to be a string value. The string value may optionally be enclosed in matching quote marks or apostrophes but this is only necessary if the value contains white space or begins with a character sequence that would otherwise indicate that an integral value is being specified. Note that depending on your operating system, using the ampersand or other special characters on the command line may require special quoting or escape characters.

If the identifier already exists, an error message will be displayed. The `-U` option may be used to avoid this situation. Also note that identifiers defined in this fashion are also available to be used in program code somewhat as if they were defined by a `Const` definition. The primary difference is that integral valued identifiers have a universal type that can be used just like a literal integral value.

Identifiers defined by this means are visible in all modules processed after the appearance of the option. Note that within a particular module the identifier may be undefined and re-defined. However, this has effect only within that particular module; the remaining modules will see the value as originally defined.

`-U<id>`

This option causes the compiler to remove the specified identifier (as defined by `-D`) from the internal symbol table if the identifier exists. If the identifier doesn't exist, the option is silently ignored.

10.3 Error and Warning Messages

The compiler will output messages for detected error conditions and warnings to `stderr` (by default, the console). The message output may be redirected to a file using the I/O redirection capabilities of the operating system or by using a compiler option to specify the error output file (see `--error` above). Error messages relating to problems with compiler options will not be affected by the presence of the error output specification.

For code-related errors and warnings, the message will contain a reference to the filename and line number of the code corresponding to the error or warning. Although the line number given will generally be correct, code involving line continuations may lead to the line number reported being different than the actual physical line number where the offending code appears.

The general format of the error and warning messages may be modified to some extent using the `--error-format` option described in the preceding section. This may be useful if you are using the compiler with a different IDE that expects to see error message in a slightly different format for its “jump to the next error” function.

10.3.1 Controlling Warnings

The types of warnings that the compiler will issue may be controlled from the command line using the `--warn` option. Most of the warning types are enabled by default. It is desirable in some cases to allow the compiler to generate a certain type of warning message except for specific code sequences. This can be accomplished using the `#pragma warning` compiler directive in your program, the syntax of which is shown below.

```
#pragma warning( <warning id> : <disposition> )
```

The `<warning id>` element is either a numeric value (included as part of the warning message) or one of the mnemonic warning identifiers used with the `-warn` option. The `<disposition>` element must be one of the keywords `On`, `Off` or `Default`, the effect of which is to enable the warning, disable the warning or set it to the default state.

An example of the use of the warning directive to disable a warning for useless code is shown below.

```
Dim b as Byte
Dim i as Integer
Sub Main()
    b = 3
#pragma warning(7 : Off)
    If (b > 5) Then
        i = 100
    ElseIf (b <= 5) Then
        i = 10
    End If
#pragma warning(7 : On)
    b = 25
End Sub
```

Because of the immediately preceding assignment, the compiler can deduce the Boolean value of the conditions in both the `If` and `ElseIf` statements and will issue a “useless-code” warning indicating that those conditions are always false. The first warning directive disables that warning and the second one unconditionally enables it. Often, it is undesirable to unconditionally enable a warning after disabling it as was done in this example. To avoid enabling a warning that was not previously enabled, two variations of the warning directive are provided to save the current set of enabled warnings (`push`) and restore them afterward (`pop`) instead of unconditionally re-enabling a warning. The syntax for the `push` and `pop` warning directives is illustrated in the modified example below.

```
Dim b as Byte
Dim i as Integer
Sub Main()
    b = 3
#pragma warning(push)
#pragma warning(7 : Off)
    If (b > 5) Then
        i = 100
    ElseIf (b <= 5) Then
        i = 10
    End If
#pragma warning(pop)
    b = 25
End Sub
```

Note that it is permissible to combine two or more warning directives by separating the warning specifications using a semicolon. For example:

```
#pragma warning(push; 7:Off; unused-param:On)
```

10.3.2 Internal Errors

The compiler may, in unusual circumstances, generate an Internal Error message. This will happen when it encounters a situation where inconsistencies in internal data structures arise that prevent continued operation. If you encounter an internal error or a program fault, please report it so that it can be investigated to see why the problem occurred. If possible, construct the smallest possible program that still exhibits the problem and submit that program along with the error report. In some cases, you may be

able to work around the problem by turning off all optimization although this is not guaranteed to be helpful. Still, it may be worth a try to get you past the obstacle while the problem is being addressed.

Problem reports may be emailed to support@zbasic.net or they may be reported via the Support Forum at <http://www.zbasic.net/forum>.

Chapter 11 - Downloader Utility

Included as part of the ZBasic package is a standalone command line utility that can download code to the ZX or a generic target device with a ZBasic-compatible bootloader. Like the compiler, it is a console application meaning that it has no graphical user interface. The invocation syntax for the downloader is:

```
zload [<options>] <code-file>
```

where *<code-file>* is a file containing object code created by the compiler, usually having a .zxb extension. The code file is in an industry standard format and contains checksums to help detect communication problems.

After the code file is downloaded, the ZX is reset so that it begins running the user program.

The available options for the downloader are described in the table below. All options are case sensitive. For some options, an accompanying hexadecimal value can be specified using a 0x prefix.

zload Options

Option	Description
-h	Display an invocation syntax summary and then exit.
-c<port>[:<speed>]	Specify the serial port to use instead of the default port. <port> must be a decimal value between 1 and 99. If the optional decimal <speed> value is given, the port is opened at that baud rate; otherwise the speed is 19.2K baud.
-b<run-speed>:<dl-speed>	Specify the baud rates for run mode and download mode.
-d<family>	Specify the device family for downloading to a generic target device. The supported families are mega, ATmega, xmega, and ATxmega.
-f<frequency>	Specify the operating frequency for a generic target device.
-v	After downloading, perform a verification pass.
-m[<log file>]	After downloading, remain connected and display any received characters on the console. If the optional log file is specified, output received from the ZX will also be written to the specified file.
-s[<term-char>]	Terminate the monitor mode (-m) if a character with the value specified by <term-char> is received. The value may be specified in decimal or hexadecimal. If not specified, the termination character is EOT (04).
-a	Enter ATN test mode. Used to verify correct connection to the ZX.
-a+	Specify that the ATN signal should idle in the high state (+V).
-ar	Specify that RTS should be used for ATN signalling.
-a<period>:<duty>	Specify the period and duty cycle of the DTR/RTS toggling for ATN. The period is expressed in milliseconds (5 to 25) and the duty cycle in percent (1 to 99).
-a<ATN char>	Specify use of the given ATN character (decimal or hexadecimal) whose value must be 1-31.
-A<ATN char>	Same as -a but without changing to 115.2K baud during downloading.
-i	Retrieve an identification string from the attached device.
-u	Download a specially formatted file to update the control program of the ZX. See Section 11.1 for more details.
-U	Same as -u except that no visual feedback is provided (using the LEDs on the 24-pin ZX devices) during the update.
-e	Use a special “emergency update” method. See Section 11.1 for more details.
-z<id>:<value>	This option is used to send configuration information to the device.

Some example invocations of zload are:

Download the file test.zxb using COM2.

```
zload -c2 test.zxb
```

Download the file test.zxb using COM2 with verification.

```
zload -c2 -v test.zxb
```

Download the file `test.zxb` using the default COM port, go to monitor mode until character code 8 is received.

```
zload -s0x08 -m test.zxb
```

Download using the ATN character &H14 at 57.6K baud on COM2, constant speed (no speed change).

```
zload -c2:57600 -A0x14
```

Download to a generic mega-based target device running at 16MHz using 19200 baud for both downloading and running after the download.

```
zload -c2 -dmega -f16000000 -b19200:19200 -m test.zxb
```

Perform ATN testing using COM2.

```
zload -c2 -a
```

The default serial port may be specified by setting an environment variable named `ZLOAD_PORT`. The value of this environment variable should be a series of decimal digits specifying the serial port number (1-99). The method used to set an environment variable varies depending on the operating system. Consult your OS documentation for more information.

Example

```
ZLOAD_PORT=2
```

If no default serial port is specified and the `-c` option is not specified, COM1 is used by default. Source code is provided in the distribution for the `zload` utility. It is a fairly simple program written in C and C++ targeted to the Windows platform. The utility may be able to be ported to another operating system by a skilled programmer who is familiar with programming for both Windows and the target platform.

11.1 Firmware Updates

The firmware of VM-mode ZX devices (as opposed to native mode ZX devices) is designed to be field-upgradeable. This allows the control program to be updated with newer versions as enhancements are made and problems are fixed. An update can be installed in the ZX using the `-u` option of the `zload` command line utility and specifying the name of a file to install. The file must contain specially formatted data, the integrity of which is verified before downloading to the ZX. It is important to note that the separate update files are provided for the various models. Be sure that you have the correct update file for your ZX device.

To perform a normal firmware update the ZX must be powered up and must be connected to your PC by a serial cable. An example invocation of `zload` to perform a normal firmware update is shown below. The example shows the use of the `-c` option to also specify the serial port to use. The file `zx24_1-10-2.zvm` is an example of the specially formatted update file for the ZX-24.

```
zload -c2 -u zx24_1-10-2.zvm
```

On the 24-pin ZX devices, during the update process the red LED will be illuminated continuously and the green LED will blink at a rate of about twice per second. Note that since the LEDs are also connected to pins 25 and 26, if you have circuitry connected to these pins that would be adversely affected by the LEDs being activated you can suppress the activation of the LEDs by using the `-U` option instead of using `-u`.

It is important to ensure that the update process, once begun, is allowed to run to completion. Powering down the ZX or resetting it during the update may leave it in an unusable state. It is possible that the ZX may no longer be able to properly interact with the `zload` program to effect a subsequent complete update. In such a case, you may be able to use the special emergency update procedure described below.

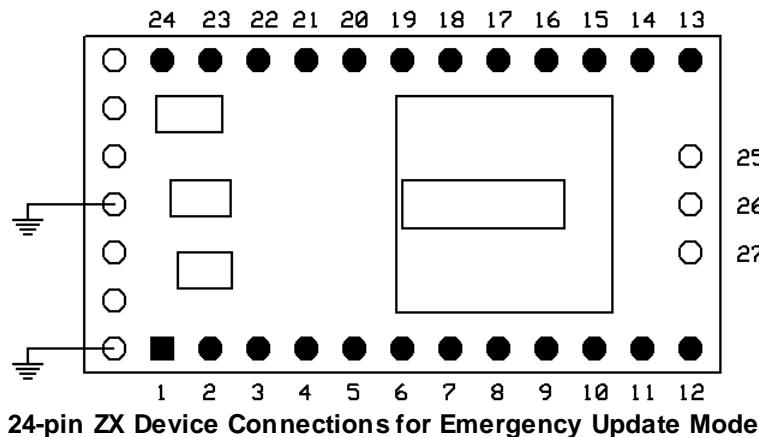
11.1.1 Emergency Update Procedure

To prepare a 24-pin VM mode ZX device for an emergency firmware update you must ground the center terminal and the bottom terminal on the left end of the ZX device as shown in the diagram below. This can usually be done with a temporary connection like a test wire with alligator clips or similar connectors. When the ZX is powered up with these connections it enters the emergency update mode instead of operating normally.

Finally, the zload command must be invoked in emergency update mode, an example of which is shown below.

```
zload -c2 -e zvm_1-12-0.zvm
```

When the emergency update completes successfully the ZX may be powered off and the grounding jumpers may be removed. When it is powered up again it will begin running the newly installed control program normally. Note that the emergency mode may be used any time but because of the extra connections that are required it is primarily intended for extraordinary circumstances when the normal update mode cannot be employed.



The emergency update procedure for other ZX devices is similar to that described above. To prepare for the emergency firmware update, consult the table below for the specific device and ground the indicated pin or pins.

Device	Pins to Ground
ZX-40, ZX-40a, ZX-40p, ZX-40s	5, 7 (B.4, B.6)
ZX-44, ZX-44a, ZX-44p, ZX-44s	2, 44 (B.4, B.6)
ZX-1281	25 (D.0)
ZX-1280	43 (D.0)

When the ZX powers up and detects that these pins are grounded, it will enter the emergency update mode and await an emergency firmware download on the serial port. After the firmware update is downloaded using the zload command with the -e option the device should be powered off and the grounding jumpers should be removed. When it is powered up again it will begin running the newly installed control program normally.

11.2 Device Configuration

The device configuration option, -z, can be used to set configuration parameters of the device (held in the reserved system Persistent Memory). The supported elements are given in the table below.

Configuration Element	Syntax
EEPROM type and size	-z0:<EEPROM type>
External RAM configuration	-z1:<ext RAM cfg>
ATN character	-z3:<ATN char>
Default Com1 speed (VM devices only)	-z4:<Com1 speed index>
Sign-on flag	-z5:<sign-on flag>
Byte value	-zb:<addr>:<value>
Word value	-zw:<addr>:<value>
Double word value	-zd:<addr>:<value>

In general, the values for the parameters to the configuration option may be given in decimal or hexadecimal form. A prefix of 0x indicates hexadecimal form.

-z0:<EEPROM type>

This configuration option is only useful for devices that require an external EEPROM and where you can choose which EEPROM to use, e.g. some ZX-40 or ZX-44 series devices. The <EEPROM type> element is a 16-bit composite value that specifies characteristics of the attached EEPROM. The least significant byte specifies the “page size” of the EEPROM in terms of 16-byte blocks. For example, a page size of 64 bytes would be specified using the value 4 while a page size of 128 bytes would be specified using the value 8. The least significant bit of the most significant byte specifies whether the EEPROM requires full-page writes. If the EEPROM supports writing less than a full page, this bit should be off.

The example configuration specifications below all use the hexadecimal form. Note that the values specified in the examples are shown with 4 hexadecimal digits for clarity but leading zero digits may be omitted.

-z0:0x0004

This specifies an EEPROM with a 64-byte page that supports partial page writes. This is the configuration to use for the Atmel AT25256A. This is the default configuration, ZX devices come pre-configured for this EEPROM.

-z0:0x0104

This specifies an EEPROM with a 64-byte page that requires full page writes. This is the configuration value to use for the Atmel AT25HP256.

-z0:0x0008

This specifies an EEPROM with a 128-byte page that supports partial page writes. This is the configuration value to use for the ST Microelectronics M95512.

-z0:0x0108

This specifies an EEPROM with a 128-byte page that requires full page writes. This is the configuration value to use for the Atmel AT25HP512.

-z1:<ext RAM cfg>

The value for the external RAM configuration is described earlier in this document. See the section describing Option ExtRamConfig.

-z3:<ATN char>

The value of the ATN character configuration must be in the range 0 to 31 (&H00 to &H1F). The value of zero indicates that no ATN character is being specified.

`-z4:<Com1 speed index>`

This configuration value (used only on VM-based devices) specifies the default baud rate for Com1 according to the table below. This configuration value is not supported in VM versions prior to v3.0.4.

Index Value	Default Com1 Baud Rate
0	Not specified
1	1200
2	2400
3	4800
4	9600
5	19200
6	38400
7	57600
8	115200
9	230400

The index value of zero is indicates that the standard default value of 19.2K baud should be used. Index values other than those listed in the table above will produce undefined results.

`-z5:<sign-on flag>`

This configuration value controls whether the “sign on” message will be displayed when the ZX begins running. A value of zero will suppress the sign on message, any other value will allow it.

`-zb:<addr>:<value>`
`-zw:<addr>:<value>`
`-zd:<addr>:<value>`

These forms allow you to write an 8-bit, 16-bit or 32-bit value to a specified address in Persistent Memory. These forms can be used to set values in both the system portion (addresses 0 to 31) and in the user portion.

It is important to note that whenever the ZX firmware is updated the default configuration values (in the system portion of Persistent Memory) are restored. Consequently, you must reconfigure a device after updating the firmware.

11.3 Downloader API

In some cases, it may be useful to invoke the downloading functions from another program. For this purpose, a downloader API is provided in the form of a dynamic link library (DLL). The source code for the DLL and the DLL itself may be found in the `src\zload` subdirectory of the directory where ZBasic is installed. Two sample applications, one in C and one in Visual Basic, using the downloader DLL are also available in the same directory sub-tree.

Appendix A - Reserved Words

The list below enumerates the words that are reserved for use as keywords in ZBasic or reserved for possible future use. The reserved words may not be used as identifiers in user-written programs.

Standard Reserved Words

alias	defobj	longtabledata	rset
and	defsng	longtabledatarw	seek
array	defstr	longvectordata	select
as	defvar	longvectordatarw	set
attribute	dim	loop	sgn
based	do	lset	single
bit	doevents	me	singletabledata
boolean	double	mod	singletabledatarw
boundedstring	each	module	singlevectordata
byref	else	new	singlevectordatarw
byte	elseif	next	sizeof
bytealign	end	nibble	spc
bytetabledata	enum	not	static
bytetabledatarw	eqv	open	step
bytevectordata	erase	option	stop
bytevectordatarw	exit	or	string
byval	for	persistent	stringtabledata
call	function	persistentboolean	stringvectordata
calltask	get	persistentbyte	sub
case	gosub	persistentinteger	tab
ccur	goto	persistentlong	then
cdate	if	persistentsingle	to
cdec	imp	persistentstring	type
close	imports	persistentstring	ubound
const	in	persistentstructure	union
currency	input	persistenttype	unlock
cvar	int	preserve	unsignedinteger
cverr	integer	print	unsignedlong
date	integertabledata	private	until
declare	integertabledatarw	progmem	variant
defbool	integervectordata	public	vb_name
defbyte	integervectordatarw	put	version
defcur	is	redim	volatile
defdate	lbound	register	wend
defdbl	let	rem	while
defdec	like	resume	with
defint	lock	return	write
deflng	long		xor

When the object-oriented extensions are enabled (using `Option Objects`) several additional keywords are reserved as enumerated in the table below.

Object-Oriented Reserved Words

abstract	extends	includes
class	final	protected

Appendix B - Supported Target Devices

This section gives the names of target devices supported by the ZBasic compiler, which fall into two general categories: ZX devices and generic devices. The device names may be used with the `Option TargetDevice` directive and with the `--target-device` compiler command line option. In both cases, the device names are not case sensitive. Note that a special license for the ZBasic compiler is required in order to generate code for generic devices.

ZX Devices

ZX24	ZX24a	ZX24p	ZX24n	ZX24r	ZX24s	ZX24t	ZX24x
ZX40	ZX40a	ZX40p	ZX40n	ZX40r	ZX40s	ZX40t	ZX24u
ZX44	ZX44a	ZX44p	ZX44n	ZX44r	ZX44s	ZX44t	
ZX1281	ZX1281n	ZX1280	ZX1280n				ZX32a4
ZX328n	ZX3281	ZX32n	ZX321		ZX128a4u	ZX128a1	
ZX24e	ZX24ae	ZX24pe	ZX24pu	ZX24nu	ZX24ru	ZX24su	
ZX24xu	ZX128e	ZX128ne	ZX1281e	ZX1281ne	ZX328nu		

Generic Devices

Atmel AVR Tiny Series

ATTiny2313	ATTiny2313A	ATTiny4313	ATTiny24	ATTiny24A	ATTiny44
ATTiny44A	ATTiny441	ATTiny84	ATTiny84A	ATTiny841	ATTiny87
ATTiny167	ATTiny48	ATTiny88	ATTiny828	ATTiny1634	

Atmel AVR Mega Series

ATmega48	ATmega48A	ATmega48P	ATmega48PA	ATmega48PB	ATmega8
ATmega8A	ATmega88	ATmega88A	ATmega88PA	ATmega88PB	ATmega168
ATmega168A	ATmega168P	ATmega168PA	ATmega168PB	ATmega328P	ATmega328PB
ATmega16	ATmega16A	ATmega164A	ATmega164P	ATmega164PA	ATmega32
ATmega32A	ATmega324P	ATmega324PA	ATmega644	ATmega644A	ATmega644P
ATmega644PA	ATmega1284P	ATmega64	ATmega64A	ATmega128	ATmega128A
ATmega1281	ATmega2561	ATmega1280	ATmega2560		
ATmega8515	ATmega8535	ATmega161	ATmega162	ATmega163	ATmega165
ATmega165A	ATmega165P	ATmega165PA	ATmega323	ATmega325	ATmega325P
ATmega3250	ATmega3250P	ATmega645	ATmega645A	ATmega645P	ATmega6450
ATmega6450A	ATmega6450P				

Atmel AVR Xmega Series

ATxmega64A1	ATxmega128A1	ATxmega64A1U	ATxmega128A1U	ATxmega64A3	ATxmega128A3
ATxmega192A3	ATxmega256A3	ATxmega64A3U	ATxmega128A3U	ATxmega192A3U	ATxmega256A3U
ATxmega16A4	ATxmega32A4	ATxmega16A4U	ATxmega32A4U	ATxmega64A4U	ATxmega128A4U
ATxmega256A4U	ATxmega64D3	ATxmega128D3	ATxmega192D3	ATxmega256D3	ATxmega384D3
ATxmega256A3B	ATxmega16D4	ATxmega32D4	ATxmega64D4	ATxmega128D4	ATxmega256A3BU

Atmel AVR LCD Devices

ATmega169	ATmega169A	ATmega169P	ATmega169PA	ATmega329	ATmega329P
ATmega329PA	ATmega649	ATmega649A	ATmega649P	ATmega3290	ATmega3290P
ATmega6490	ATmega6490A	ATmega6490P			

Atmel AVR USB Devices

ATmega8U2	ATmega16U2	ATmega32U2	ATmega16U4	ATmega32U4	AT90USB82
AT90USB162	AT90USB646	AT90USB647	AT90USB1286	AT90USB1287	

Atmel AVR CAN Devices

AT90CAN32	AT90CAN64	AT90CAN128	
-----------	-----------	------------	--

Appendix C - Pre-Defined Structures

The following pre-defined structures are available for most ZBasic devices based on Atmel ATmega CPUs. However, due to irregularities in the arrangement of registers on some mega CPUs they are not available. In particular, the predefined structures are not available for ZBasic devices based on the mega32 or the mega128.

One or more of the pre-defined structures may be added to your application with the `Option Include` directive using the structure names below.

ATTiny- and ATmega-based ZBasic Devices

```
' For accessing I/O ports:  
Structure Port_t  
    Public Volatile Pin as Byte  
    Public Volatile DDR as Byte  
    Public Volatile Port as Byte  
End Structure  
  
' For accessing 8-bit timers (if present in standard form):  
Structure Timer8_t  
    Public Volatile TCCRA as Byte  
    Public Volatile TCCRB as Byte  
    Public Volatile CNT as Byte  
    Public Volatile ICR as Byte  
    Public Volatile OCRA as Byte  
    Public Volatile OCRB as Byte  
End Structure  
  
' For accessing 16-bit timers (if present in standard form):  
Structure Timer16_t  
    Public Volatile TCCRA as Byte  
    Public Volatile TCCRB as Byte  
    Public Volatile TCCRC as Byte  
    Public Volatile unused_1 as Byte  
    Public Volatile CNT as UnsignedInteger  
    Public Volatile ICR as UnsignedInteger  
    Public Volatile OCRA as UnsignedInteger  
    Public Volatile OCRB as UnsignedInteger  
    Public Volatile OCRC as UnsignedInteger (if applicable)  
End Structure  
  
' For accessing USARTs (if present in standard form):  
Structure USART_t  
    Public Volatile UCSRA as Byte  
    Public Volatile UCSRB as Byte  
    Public Volatile UCSRC as Byte  
    Public Volatile unused_1 as Byte  
    Public Volatile UBRR as UnsignedInteger  
    Public Volatile UDR as Byte  
End Structure  
  
' For representing Microtime:  
Structure Microtime_t  
    Public Volatile timerTicks as Byte  
    Public Volatile fastTicks UnsignedLong  
End Structure
```

ATxmega-based ZBasic Devices

```
' For accessing I/O ports:  
Structure Port_t  
    Public Volatile Dir as Byte  
    Public Volatile DirSet as Byte  
    Public Volatile DirClr as Byte  
    Public Volatile DirTgl as Byte  
    Public Volatile Out as Byte  
    Public Volatile OutSet as Byte  
    Public Volatile OutClr as Byte  
    Public Volatile OutTgl as Byte  
    Public Volatile In as Byte  
    Public Volatile IntCtrl as Byte  
    Public Volatile Int0Mask as Byte  
    Public Volatile Int1Mask as Byte  
    Public Volatile IntFlags as Byte  
    Public Volatile unused(1 to 3) as Byte  
    Public Volatile Pin0Ctrl as Byte  
    Public Volatile Pin1Ctrl as Byte  
    Public Volatile Pin2Ctrl as Byte  
    Public Volatile Pin3Ctrl as Byte  
    Public Volatile Pin4Ctrl as Byte  
    Public Volatile Pin5Ctrl as Byte  
    Public Volatile Pin6Ctrl as Byte  
    Public Volatile Pin7Ctrl as Byte  
End Structure  
  
' For accessing 16-bit timers (Type 0):  
Structure TC0_t  
    Public Volatile CtrlA as Byte  
    Public Volatile CtrlB as Byte  
    Public Volatile CtrlC as Byte  
    Public Volatile CtrlD as Byte  
    Public Volatile CtrlE as Byte  
    Public Volatile unused1 as Byte  
    Public Volatile IntCtrlA as Byte  
    Public Volatile IntCtrlB as Byte  
    Public Volatile CtrlFCtrl as Byte  
    Public Volatile CtrlFSet as Byte  
    Public Volatile CtrlGCtrl as Byte  
    Public Volatile CtrlGSet as Byte  
    Public Volatile IntFlags as Byte  
    Public Volatile unused2(1 to 2) as Byte  
    Public Volatile Temp as Byte  
    Public Volatile unused3(1 to 16) as Byte  
    Public Volatile Cnt as UnsignedInteger  
    Public Volatile unused4(1 to 4) as Byte  
    Public Volatile Per as UnsignedInteger  
    Public Volatile CCA as UnsignedInteger  
    Public Volatile CCB as UnsignedInteger  
    Public Volatile CCC as UnsignedInteger  
    Public Volatile CCD as UnsignedInteger  
    Public Volatile unused5(1 to 6) as Byte  
    Public Volatile PerBuf as UnsignedInteger  
    Public Volatile CCABuf as UnsignedInteger  
    Public Volatile CCBBuf as UnsignedInteger  
    Public Volatile CCCBuf as UnsignedInteger  
    Public Volatile CCDBuf as UnsignedInteger  
End Structure
```

```

' For accessing 16-bit timers (Type 1):
Structure TC1_t
    Public Volatile CtrlA as Byte
    Public Volatile CtrlB as Byte
    Public Volatile CtrlC as Byte
    Public Volatile CtrlD as Byte
    Public Volatile CtrlE as Byte
    Public Volatile unused1 as Byte
    Public Volatile IntCtrlA as Byte
    Public Volatile IntCtrlB as Byte
    Public Volatile CtrlFClr as Byte
    Public Volatile CtrlFSet as Byte
    Public Volatile CtrlGClr as Byte
    Public Volatile CtrlGSet as Byte
    Public Volatile IntFlags as Byte
    Public Volatile unused2(1 to 2) as Byte
    Public Volatile Temp as Byte
    Public Volatile unused3(1 to 16) as Byte
    Public Volatile Cnt as UnsignedInteger
    Public Volatile unused4(1 to 4) as Byte
    Public Volatile Per as UnsignedInteger
    Public Volatile CCA as UnsignedInteger
    Public Volatile CCB as UnsignedInteger
    Public Volatile unused5(1 to 10) as Byte
    Public Volatile PerBuf as UnsignedInteger
    Public Volatile CCABuf as UnsignedInteger
    Public Volatile CCBBuf as UnsignedInteger
End Structure

' For accessing USARTs:
Structure USART_t
    Public Volatile Data as Byte
    Public Volatile Status as Byte
    Public Volatile unused1 as Byte
    Public Volatile CtrlA as Byte
    Public Volatile CtrlB as Byte
    Public Volatile CtrlC as Byte
    Public Volatile BaudCtrlA as Byte
    Public Volatile BaudCtrlB as Byte
End Structure

' For representing Microtime:
Structure Microtime_t
    Public Volatile timerTicks as UnsignedInteger
    Public Volatile fastTicks UnsignedLong
End Structure

```


Appendix D - ZX-24 Series Hardware Reference

The ZX-24 series devices comprise an Atmel AVR ATmega microcontroller along with some support circuitry on a 24-pin module that is simple to connect for operation. The primary difference between the members of the ZX-24 series devices is specific ATmega or ATxmega microcontroller used and the amount of RAM, EEPROM and Program Memory available for use by your programs. The table below summarizes the differences.

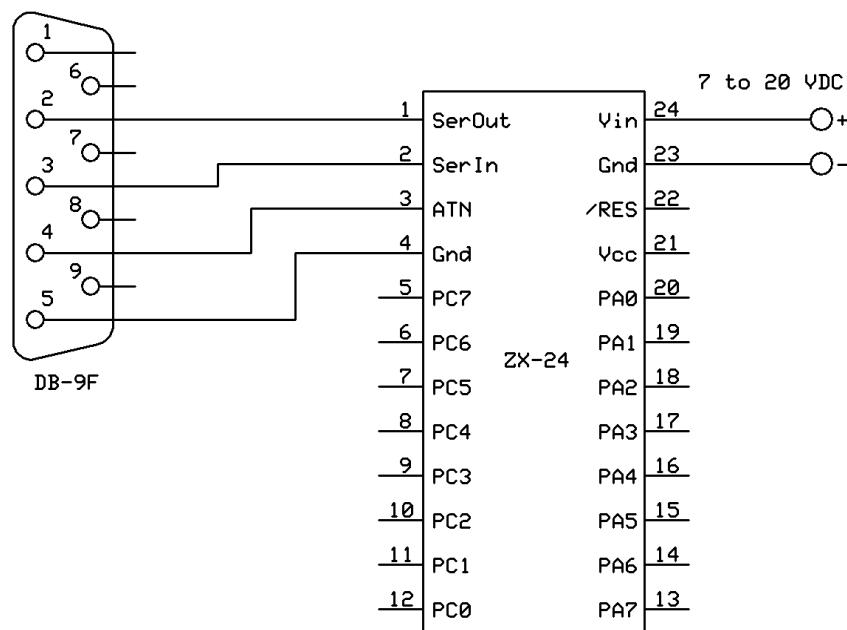
ZX-24 Series Devices

Device	Base AVR Microcontroller	Mode	RAM	Persistent Memory	Program Memory	Operating Frequency	Operating Voltage
ZX-24	ATmega32	VM	1536	992	32K	14.7MHz	3.6-5.5V
ZX-24a	ATmega644	VM	3584	2016	32K	14.7MHz	3.6-5.5V
ZX-24p	ATmega644P	VM	3584	2016	32K	14.7MHz	3.6-5.5V
ZX-24n	ATmega644P	Native	4096	2016	62K	14.7MHz	3.6-5.5V
ZX-24r	ATmega1284P	VM	15744	4064	60K	14.7MHz	3.6-5.5V
ZX-24s	ATmega1284P	Native	16384	4064	124K	14.7MHz	3.6-5.5V
ZX-24t	ATmega1284P	Native	16384	4064	124K	7.37MHz	2.3-5.5V
ZX-24x	ATxmega32A4	Native	4096	1024	32K	29.5MHz	2.6-3.6V
ZX-24u	ATxmega128A4U	Native	8192	2048	128K	29.5MHz	2.6-3.6V

The ZX-24 series devices have other resources available to your program including a high-speed serial port, analog-to-digital converters, timers and other sub-systems that may be accessed using routines in the ZBasic System Library. Alternately, some of the resources may be accessed directly using built-in registers. See Section 3.7.1 for more information on this topic.

D.1 External Connections

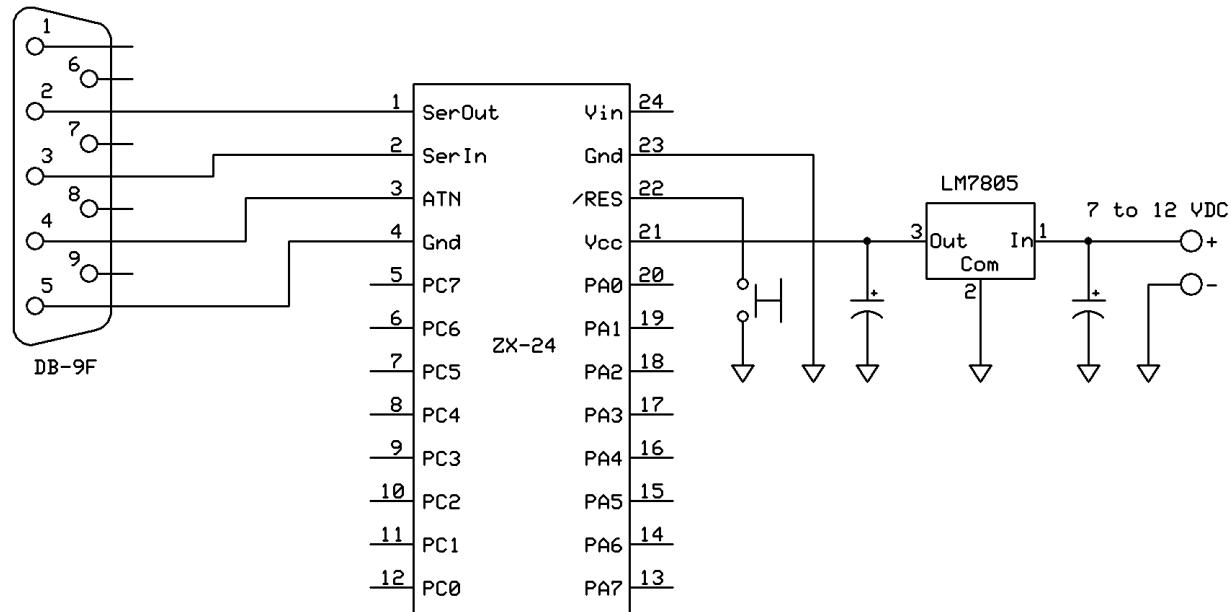
The ZX-24 series devices can be hooked up in several different ways depending on your requirements. The simplest method is depicted below. The connector on the left is the serial connection to your PC for downloading code and transmitting/receiving data on serial channel 1. The connection to pin 3 of the ZX-24 series is only needed for downloading. You may wish to add a jumper in that line so that it may be disconnected for "normal" operation. Some operating systems, notably Windows XP, toggle the DTR line (pin 4 of the DB-9) on the serial ports during boot up. This will cause the ZX to reset on every positive-going transition of the DTR line.



Simple ZX-24 Series Interconnection

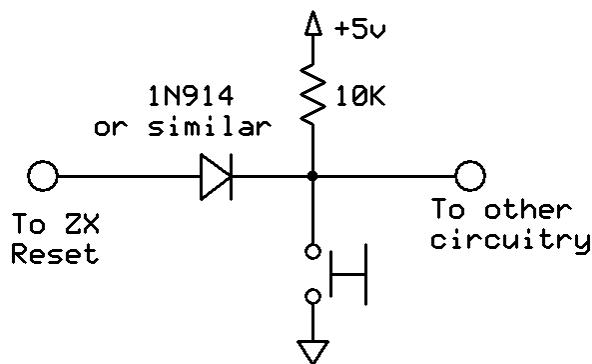
By default, the serial connection to the ZX-24 series devices is 19.2K baud. You can operate the serial channel at higher and lower speeds by explicitly opening the COM1 serial channel. See the description of OpenCom() in the ZBasic System Library Reference manual.

The simple configuration shown above is limited by the capacity of the on-board regulator. Although it will work for simple systems, many users will benefit from powering the ZX-24 series device from an externally regulated source. A suggested connection for such a configuration is shown below. (Note that a different regulator must be used for the ATxmega-based ZX devices; one that provides 2.6-3.6V.) The output of the external regulator can also be used to power additional external circuitry up to the capacity of the regulator. Depending on the power dissipated by the regulator, it may need to be mounted on a heat sink to prevent it from getting too hot.



ZX-24 Series Device Interconnection Using an External Regulator

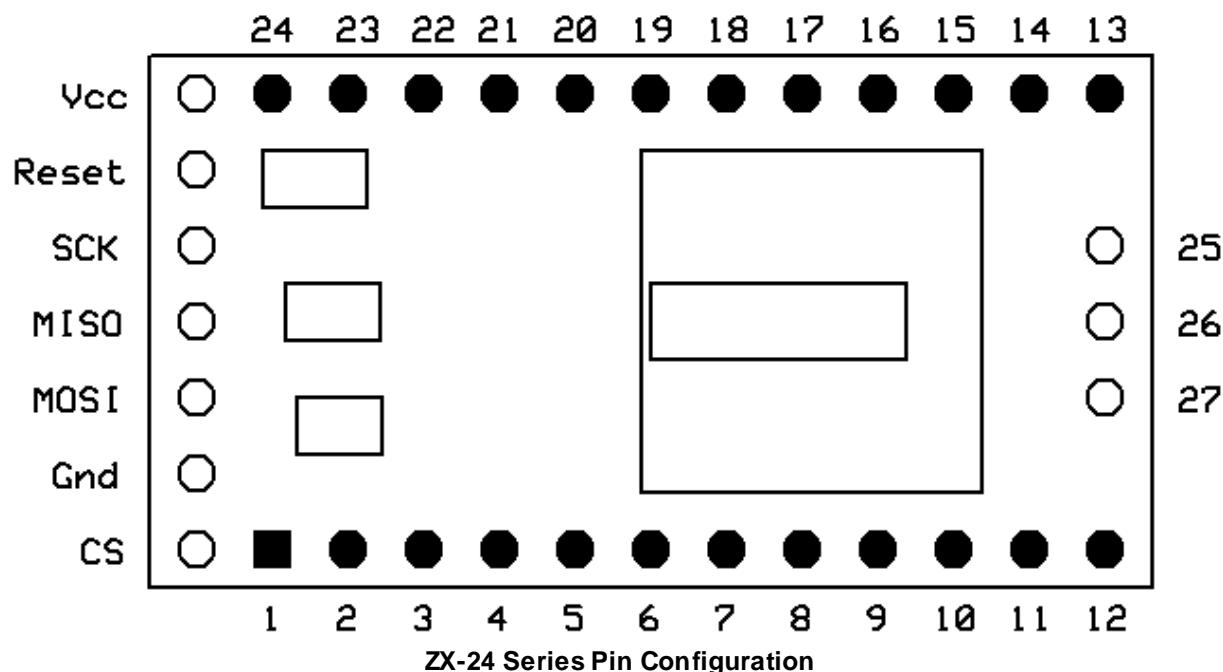
This configuration also shows how to connect a reset switch. A reset switch is sometimes useful, particularly during debugging, to get the ZX to start running its program from the beginning. If the reset is also connected to other external circuitry, you may wish to isolate the ZX from the remaining circuitry so that the on-board resets that occur during downloading are not also applied to the external devices. A simple isolation circuit employing a small signal diode is shown below. You can also implement an isolation circuit using a discrete transistor or an open-collector gate.



Reset Signal Isolation

D.2 Pin Configuration

The ZX-24 series devices have 24 male pins, numbered 1 through 24 in a standard DIP configuration. Additionally, there are 10 other connections that may be made to the ZX-24 series circuit board for expansion and other special purposes. In the diagram below, the male pins are represented by filled circles while the open circles represent solderable connections. The filled square indicates pin 1. The remaining rectangles on the diagram are stylized representations of some of the components on the board and are shown for orientation purposes only. Pins 1 to 24 are referred to as the standard pins while the remaining connections are referred to as the expansion pins.



D.2.1 Standard Pins

The table below summarizes the functions of the standard ZX-24 series pins. A detailed description of each pin is given following the table. Note that several of the external pins connect to more than one pin of the microcontroller chip. The multiple connections allow such pins to serve more than one function. See the descriptions below and the Atmel documentation for more information.

Standard ZX-24 Series Pins

Pin	ATmega	ATxmega
1	Serial Output	Serial Output
2	Serial Input	Serial Input
3	ATN	ATN
4	Ground (common with pin 22)	Ground (common with pin 22)
5	Port C, bit 7 and Port B, bit 1	Port C, bit 7
6	Port C, bit 6 and Port D, bit 2	Port C, bit 6
7	Port C, bit 5	Port C, bit 5 and Port B, bit 0
8	Port C, bit 4	Port C, bit 4 and Port B, bit 2
9	Port C, bit 3	Port C, bit 3 and Port B, bit 3
10	Port C, bit 2	Port C, bit 2
11	Port C, bit 1 and Port D, bit 3	Port C, bit 1
12	Port C, bit 0 and Port D, bit 6	Port C, bit 0
13	Port A, bit 7	Port A, bit 7
14	Port A, bit 6	Port A, bit 6
15	Port A, bit 5	Port A, bit 5
16	Port A, bit 4	Port A, bit 4

17	Port A, bit 3 and Port B, bit 0	Port A, bit 3 and Port E, bit 1
18	Port A, bit 2 and Port B, bit 2	Port A, bit 2 and Port E, bit 2
19	Port A, bit 1 and Port B, bit 3	Port A, bit 1 and Port E, bit 3
20	Port A, bit 0	Port A, bit 0
21	+V out or +V in (regulated)	+V out or +V in (regulated)
22	Reset (in and out)	Reset (in and out)
23	Ground (common with pin 4)	Ground (common with pin 4)
24	+V in (unregulated)	+V in (unregulated)

Detailed Pin Descriptions

Pin 1, Serial Output

This is the output pin for serial channel 1 and Debug.Print. The voltage swing is approximately 0V to the supply voltage (+5V if using the on-board regulator) and can supply a maximum of 25mA of current. Note that the voltage swing does not meet the specifications for the RS-232 standard but it will be properly recognized by most modern serial interfaces.

This pin should be connected to pin 2 of a DB-9F serial connector.

Pin 2, Serial Input

This is the input pin for serial channel 1. It will tolerate a voltage swing from -15V to +15V allowing it to be connected to an RS-232 output from a serial device.

This pin should be connected to pin 3 of a DB-9F serial connector.

Pin 3, ATN

This input is used for signaling the ZX to enter the download mode. If you don't need to download, no connection is required but you may wish to ground this pin to eliminate the possibility of spurious resets caused by electrical noise.

For downloading, this pin should be connected to pin 4 of a DB-9F serial connector. Note that some operating systems (particularly Windows XP) tend to toggle the DTR serial port line, to which this input is normally connected, during the booting process. Every positive transition on this input will cause the ZX to reset.

Pin 4, Ground

This is the ground connection for the serial port and is common with pin 23. This pin should be connected to pin 5 of a DB-9F serial connector.

Pin 5, Port C Bit 7 and Port B Bit 1 (ATmega)

In addition to being a general input or output pin, on an ATmega a signal may be applied to this pin to serve as the clock source for Timer 1. Port C Bit 7 must be an input for this to work properly. Alternately, if no external signal is connected, Port C Bit 7 may be an output and serve as the clock source for the timer.

Pin 6, Port C Bit 6 and Port D Bit 2 (ATmega)

In addition to being a general input or output pin, on the ATmega a signal may be applied to this pin to serve as external Interrupt 0. Port C Bit 6 must be an input for this to work properly. Alternately, if no external signal is connected, Port C Bit 6 may be an output and serve as the source for the interrupt.

Pins 7-10, Port C Bits 5-2 (ATmega)

These pins may be used as general inputs or outputs.

Pin 11, Port C Bit 1 and Port D bit 3 (ATmega)

In addition to being a general input or output pin, on an ATmega a signal may be applied to this pin to serve as external Interrupt 1. Port C Bit 1 must be an input for this to work properly. Alternately, if no external signal is connected, Port C Bit 1 may be an output and serve as the source for the interrupt. Additionally, this pin serves as the SDA line when using I2C channel 0.

Pin 12, Port C Bit 0 and Port D bit 6 (ATmega)

In addition to being a general input or output pin, a signal may be applied to this pin to serve as the input for the InputCapture() routine. Port C Bit 0 must be an input for this to work properly. Additionally, this pin serves as the SCL line when using I2C channel 0.

Pins 13-17, Port A Bits 7-3

In addition to being general digital input or output pins, these pins may be used for analog input to the analog-to-digital converter. The analog voltage may be read using the GetADC() routine.

Pin 18, Port A Bit 2 and Port B Bit 2 (ATmega)

In addition to being a general digital input or output pin this pin may be used for analog input to the analog-to-digital converter. Alternately, on an ATmega a signal may be applied to this pin to serve as external Interrupt 2. Port A Bit 2 must be an input for this to work properly. Moreover, if no external signal is connected, Port A Bit 2 may be an output and serve as the source for the interrupt. A second alternate use for this pin is as the analog comparator input. See the Atmel documentation for the AIN0 input for more details.

Pin 19, Port A Bit 1 and Port B Bit 3 (ATmega)

In addition to being a general digital input or output pin this pin may be used for analog input to the analog-to-digital converter. On an ATmega an alternate use for this pin is as the analog comparator input AIN1. See the Atmel documentation for the AIN1 input for more details.

Pins 20, Port A Bit 0

In addition to being general a digital input or output pin, this pin may be used for analog input to the analog-to-digital converter. The analog voltage may be read using the GetADC() routine.

Pin 21, Regulated Voltage In or Regulated Voltage Out

If the on-board regulator is being used (power being fed to pin 24) this pin will output a regulated voltage of about +5 volts (ATmega) or +3.3V (ATxmega). The capacity of the on-board regulator is approximately 100mA. About 60mA of this capacity is used by the on-board devices with all pins in the input state. If any of the pins are outputs, the source current of each output pin must be added to this usage. The balance of the regulator's capacity (very little, in most cases) may be used by external circuitry.

Except for the simplest configurations, it is advisable not to use the on-board regulator and, instead, supply a regulated voltage of 3.6 to 5.5 volts (ATmega) or 2.6 to 3.6V (ATxmega) to pin 21. In this case, no connection should be made to pin 24.

Pin 22, Reset

You may apply an active low signal to this pin to reset the processor. Note, however, that the on-board reset circuitry will also pull this line low via an open collector output each time the ATN input makes a positive transition (notably during downloading). If you don't want your external circuitry to receive this reset signal, you must isolate the on-board reset from external devices using a diode, a transistor or an open collector gate. An example circuit is shown in Section D.1.

Pin 23, Ground

This pin, common with pin 4, serves as the reference for the power supply and all I/O pins.

Pin 24, Unregulated Voltage In

To use the on-board regulator, you may supply an unregulated (but filtered) voltage source from 7 to 20 volts DC to this pin. The power supply must be capable of supplying at least 150mA. See the description of pin 21 for more information about the on-board regulator and the limitations of using it. It is important to note that the maximum input voltage must be derated above an ambient temperature of 45°C according to the formula $V_{in} \leq 5 + (125 - T_{amb}) / 5.2$. For example, with an ambient temperature of 75°C the input voltage must be kept below 14.6V. For operating in high ambient temperatures it is recommended to use an external regulator with an appropriate heatsink.

Additional Notes:

Each of the I/O pins (5-20, 25-27) can source 20mA or sink 40mA (at Vcc=5V). However, the total source current and the total sink current for all pins may not exceed 200mA. This is further limited by the capacity of the on-board regulator if it is being used. At Vcc=3V, the maximum source and sink currents are one-half of the 5V rating.

All of the I/O pins (5-20, 25-27) have protection diodes built in. If the signals that you connect to these pins go above V+ or below ground, you must include a current limiting resistor to keep the pin current below 20mA.

D.2.2 Expansion Pins

The table below summarizes the functions of the ZX-24 series expansion pins. These "pins" are actually holes on the circuit board that require a soldered or mechanical clip connection to use. A detailed description of each pin is given following the table.

Expansion Pins

Pin	ATmega	ATxmega
25	Port D Bit 7, Red LED	Port E Bit 0, Red LED
26	Port D Bit 5, Green LED	Port D Bit 0, Green LED
27	Port D Bit 4, OutputCapture pin	Port D Bit 1, OutputCapture pin
Vcc	Power, common with pin 21	Power, common with pin 21
Reset	Reset, common with pin 22	Reset, common with pin 22
SCK	SPI Clock	SPI Clock
MISO	SPI Master In Slave Out	SPI Master In Slave Out
MOSI	SPI Master Out Slave In	SPI Master Out Slave In
Gnd	Ground, common with pins 4 and 23	Ground, common with pins 4 and 23
CS	On-board EEPROM Chip Select or Slave Select for SPI	Slave Select for SPI

Detailed Pin Descriptions

Pin 25, Port D Bit 7 (ATmega) or Port E Bit 0 (ATxmega)

This pin drives the red LED and may also be used for PWM generation based on Timer 2. See the Atmel documentation for more details on the latter function.

Pin 26, Port D Bit 5 (ATmega) or Port D Bit 0 (ATxmega)

This pin drives the green LED and may also be used for PWM generation based on Timer 1. See the Atmel documentation for more details on the latter function.

Pin 27, Port D Bit 4 (ATmega) or Port D Bit 1 (ATxmega)

This pin is the output point for the standard OutputCapture() routine. It may also be used for PWM generation based on Timer 1. See the Atmel documentation for more details on the latter function.

SCK, MISO, MOSI

These pins are the inputs and outputs for the SPI bus by which external SPI devices may be connected to the ZX-24 series device.

CS

For ZX devices with an on-board EEPROM, this pin is connected to the chip select pin for that device and to the SS (slave select) pin of the ATmega. For ZX devices without an on-board EEPROM, this pin is connected only to the SS (slave select) pin of the ATmega or ATxmega.

Vcc, Reset, Gnd, CS

These pins are used in the manufacturing and testing processes.

Appendix E - ZX-40 Series Hardware Reference

The ZX-40 series devices are the 40-pin DIP package versions of the respective Atmel AVR ATmega microcontrollers that have been programmed with the ZX control firmware (for VM mode devices) or a bootloader (for native mode devices). They offer the same capabilities as the ZX-24 series plus access to 7 additional I/O pins. The primary difference between the members of the ZX-40 series devices is specific ATmega microcontroller used and the amount of RAM, EEPROM and Program Memory available for use by your programs. The table below summarizes the differences. The entries marked with an asterisk require an external serial EEPROM for program storage.

ZX-40 Series Devices

Device	AVR Microcontroller	Mode	RAM	Persistent Memory	Program Memory	Operating Frequency	Operating Voltage
ZX-40*	ATmega32	VM	1536	992	32K	14.7MHz	3.6-5.5V
ZX-40a*	ATmega644	VM	3584	2016	32K	14.7MHz	3.6-5.5V
ZX-40p*	ATmega644P	VM	3584	2016	32K	14.7MHz	3.6-5.5V
ZX-40n	ATmega644P	Native	4096	2016	62K	14.7MHz	3.6-5.5V
ZX-40r	ATmega1284P	VM	15744	4064	60K	14.7MHz	3.6-5.5V
ZX-40s	ATmega1284P	Native	16384	4064	124K	14.7MHz	3.6-5.5V
ZX-40t	ATmega1284P	Native	16384	4064	124K	7.37MHz	2.3-5.5V

In order to use the ZX-40 series devices you must add several additional components as described below. In each of the diagrams presented below, only a portion of the device's 40 pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

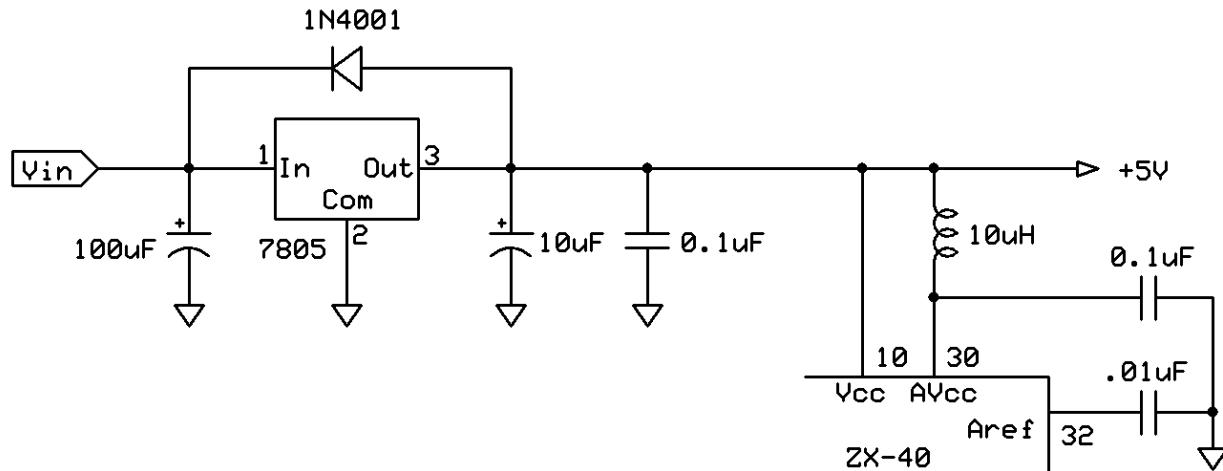
E.1 ZX-40 Series Specifications

The electrical specifications of the ZX-40 series devices are exactly those of the underlying ATmega microcontrollers. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

E.2 ZX-40 Series Required External Components

Power Source

The ZX-40 series devices need a regulated voltage source capable of providing at least 200mA of current. The voltage typically used is 5 volts but the device will operate between 4.5 volts and 5.5 volts except for the ZX-40t which can operate as low as 2.75 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



ZX-40 Series Power Source

If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the two capacitors on the right side of the diagram. In this case, pin 30 would be connected directly to the power source (as are pins 10 and 30) and pin 32 can be left open. Although not shown on this diagram, the ground pins of the ZX-40 (pins 11 and 31) must be connected to the common ground of the system.

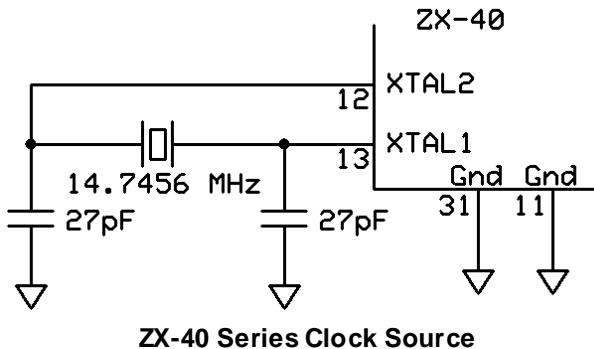
Clock Source

The ZX-40 series devices require a clock source running at 14.7456MHz except for the ZX-40t which required a 7.3728MHz clock. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 12 and 13 along with the necessary capacitors.

Recommended Crystals

Package	14.7456MHz	7.3728MHz
HC-49/US	ECS-147.4-20-4X (Digi-Key #X1102-ND)	ECS-73-20-4X (Digi-Key #X1084-ND)
HC-49/U	ECS-147.4-20-1X (Digi-Key #X1032-ND)	ECS-73-20-1X (Digi-Key #X978-ND)
Cylinder	ECS-147.4-18-9X (Digi-Key #XC-1382-ND)	

The capacitor value to use depends somewhat on the layout (e.g. stray capacitance). Values in the range of 22pF to 27pF have been used successfully in various situations.

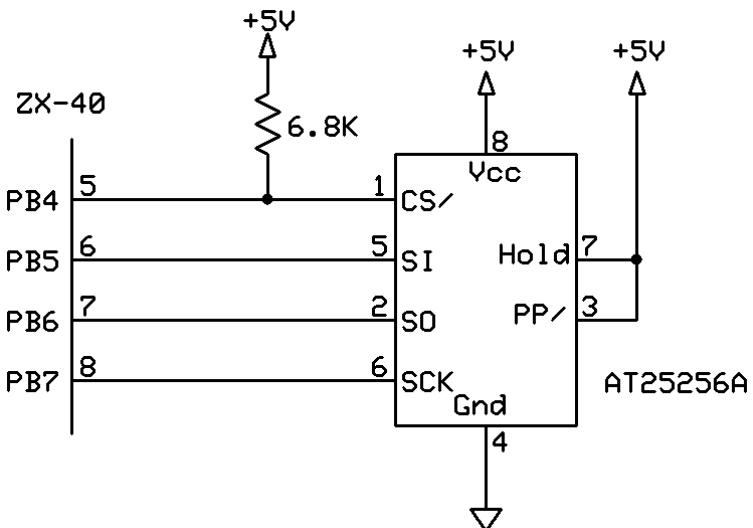


If you already have a clock source at the required frequency you can feed that signal directly to pin 13. In this case there would be no connection to pin 12.

Program Memory

Some ZX-40 series devices require a serial EEPROM in which to store your program's code. (The ZX-40n, ZX-40r, ZX-40s and ZX-40t use internal Flash memory for Program Memory.) The recommended device to use is the Atmel AT25256A (Digi-Key #AT25256A-10PI-2.7-ND). Note, particularly, that the older part without the A suffix will not work due to its slower speed. An EEPROM capable of operation at 7.5MHz or more is required. An alternate part is the ON Semiconductor CAT25256LI-G (Digi-Key #ON Semiconductor CAT25256LI-G-ND).

The recommended connection to the ZX-40 series is shown below. The value of the pullup resistor is non-critical, anything from 4.7K to 22K should work well. Higher values will reduce power consumption somewhat.



ZX-40 Series Program Memory (VM models only)

With firmware versions v1.1 and later, the Atmel 25HP512 (Digi-Key #AT25HP512-10PI2.7-ND) can also be used for applications that require larger EEPROM space. The electrical connections for this device are identical to those shown above. However, since this device has different characteristics than the AT25256A, special configuration is required. See Section 11.2 for more information.

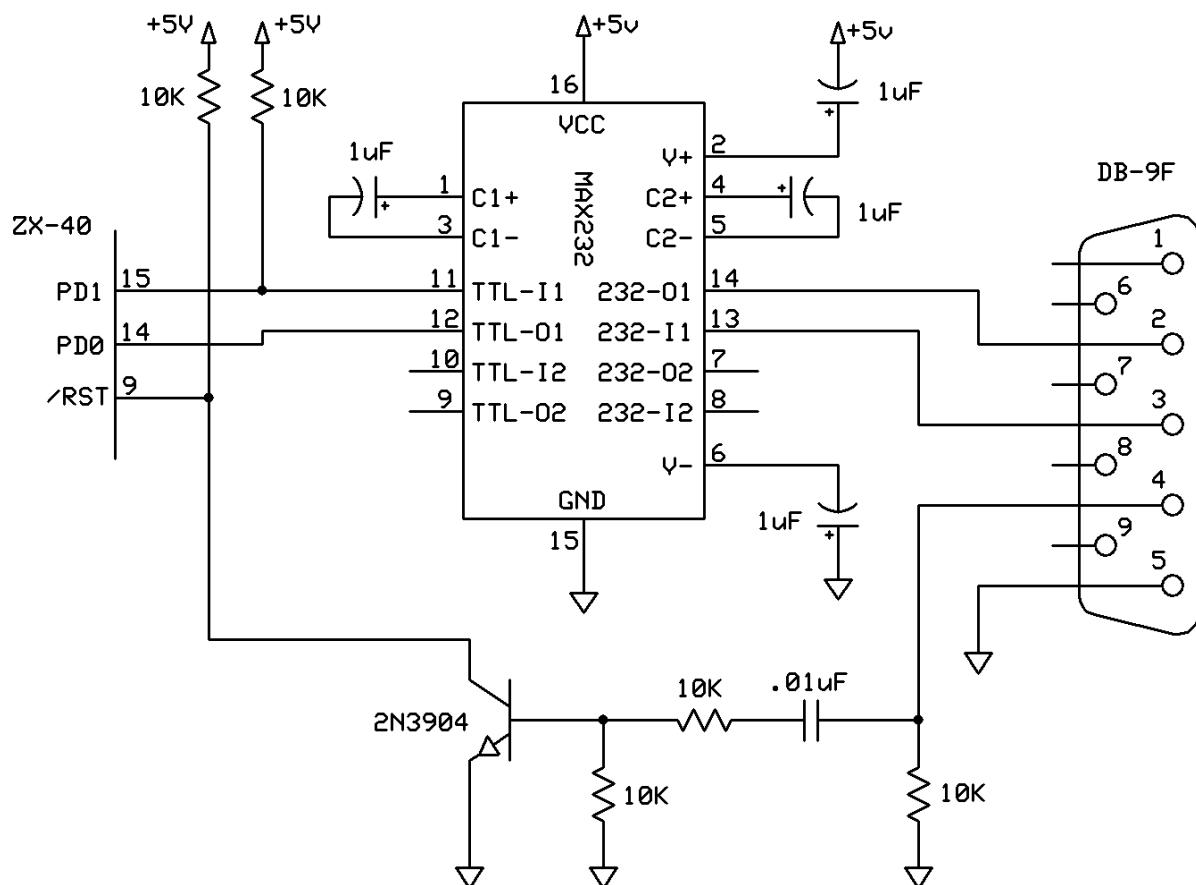
Serial Interface

The ZX-40 series devices require a serial interface for downloading code into Program Memory, performing field updates of the control program, and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-5 volt signals of the ZX to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 9 of the ZX. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 12V from the 5 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies.

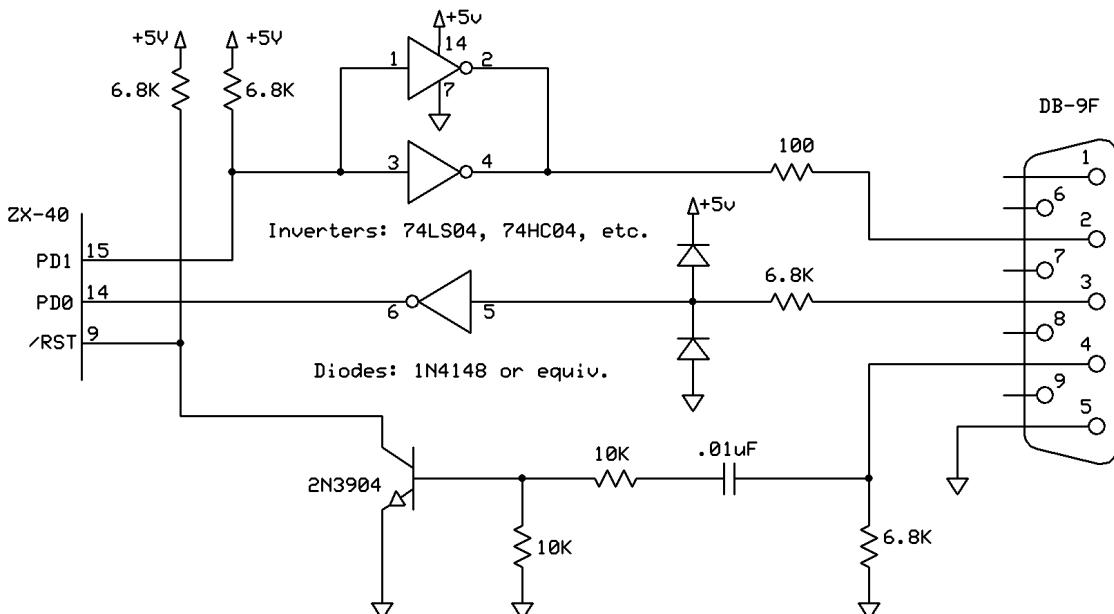
The schematics below contain several references to +5V. For ZX-40 series devices that can operate at a lower voltage, e.g. 3.3V, these references should be to the actual operating voltage.



Recommended ZX-40 Series Serial Interface Circuit

The pullup resistor shown on the ZX's transmit output (pin 15) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

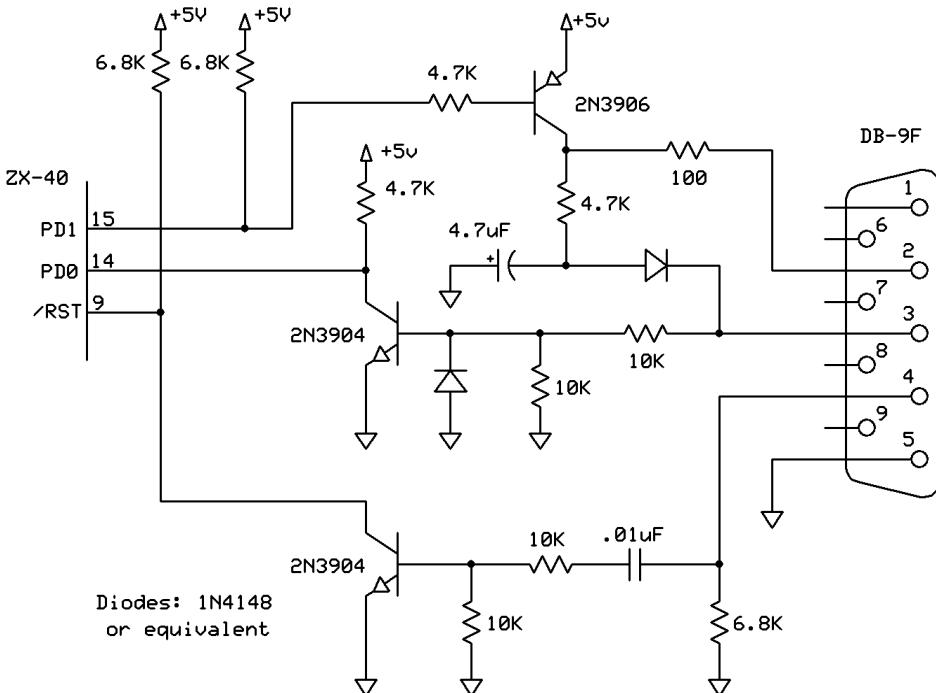
There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.



Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

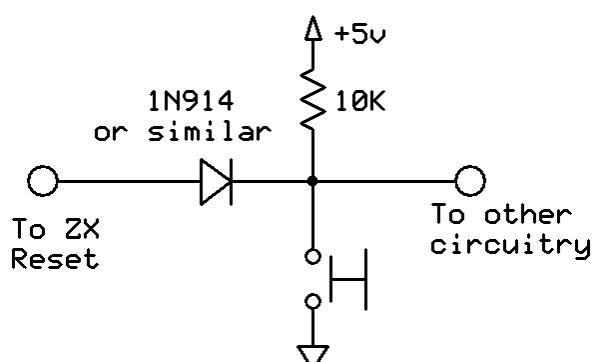
The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Alternate Serial Interface Circuit #2

Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 9 of the ZX and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix F - ZX-44 Series Hardware Reference

The ZX-44 series devices are the 44-pin TQFP package version of the Atmel AVR ATmega microcontroller that has been programmed with the ZX control firmware (VM version) or bootloader (native mode version). They offer the same capabilities as the corresponding ZX-24 series devices plus access to 7 additional I/O pins. The primary difference between the members of the ZX-44 series devices is specific ATmega microcontroller used and the amount of RAM, EEPROM and Program Memory available for use by your programs. The table below summarizes the differences. The entries marked with an asterisk require an external serial EEPROM for program storage.

ZX-44 Series Devices

Device	AVR Microcontroller	Mode	RAM	Persistent Memory	Program Memory	Operating Frequency	Operating Voltage
ZX-44*	ATmega32	VM	1536	992	32K	14.7MHz	3.6-5.5V
ZX-44a*	ATmega644	VM	3584	2016	32K	14.7MHz	3.6-5.5V
ZX-44p*	ATmega644P	VM	3584	2016	32K	14.7MHz	3.6-5.5V
ZX-44n	ATmega644P	Native	4096	2016	62K	14.7MHz	3.6-5.5V
ZX-44r	ATmega1284P	VM	15744	4064	60K	14.7MHz	3.6-5.5V
ZX-44s	ATmega1284P	Native	16384	4064	124K	14.7MHz	3.6-5.5V
ZX-44t	ATmega1284P	Native	16384	4064	124K	7.37MHz	2.3-5.5V

In order to use the ZX-44 series device you must add several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

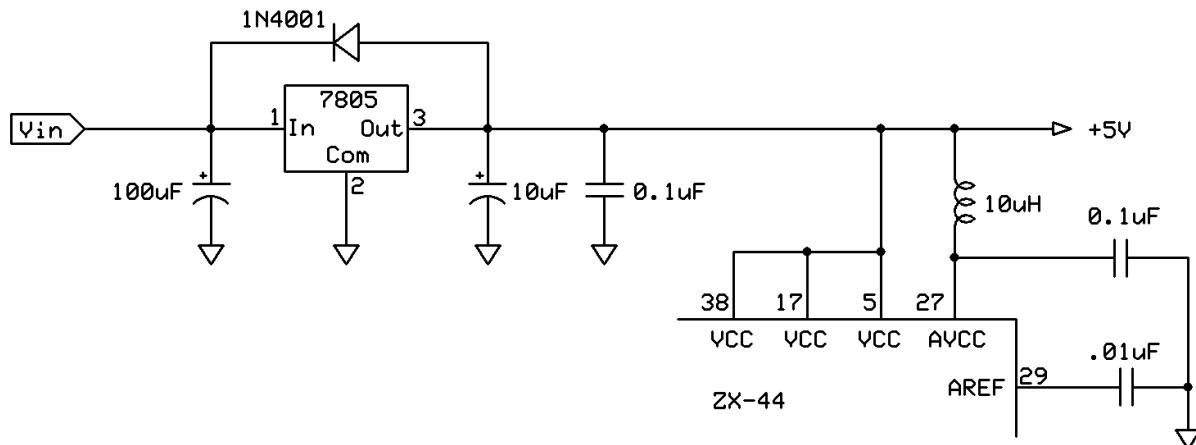
F.1 ZX-44 Series Specifications

The electrical specifications of the ZX-44 series devices are exactly those of the underlying ATmega microcontroller. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

F.2 ZX-44 Series Required External Components

Power Source

The ZX-44 series devices need a regulated voltage source capable of providing at least 200mA of current. The voltage typically used is 5 volts but the device will operate between 4.5 volts and 5.5 volts except for the ZX-44t which can operate in the range 2.75 to 5.5 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



ZX-44 Series Power Source

If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the two capacitors on the right side of the diagram. In this case, pin 27 would be connected directly to the power source (same as pin 5, etc.) and pin 29 can be left open. Although not shown on this diagram, the ground pins of the ZX-44 (pins 6, 18, 28 and 39) must be connected to the common ground of the system.

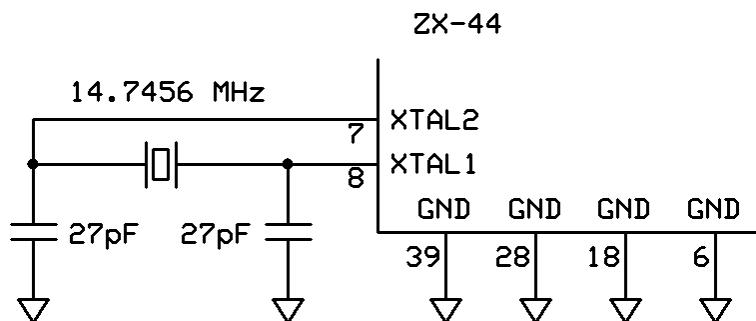
Clock Source

The ZX-44 series devices require a clock source running at 14.7456MHz except for the ZX-44t which required a 7.3728MHz clock. The easiest way to provide this clock source is to connect a crystal of the correct frequency to pins 7 and 8 along with the necessary capacitors.

Recommended Crystals

Package	14.7456MHz	7.3728MHz
HC-49/US	ECS-147.4-20-4X (Digi-Key #X1102-ND)	ECS-73-20-4X (Digi-Key #X1084-ND)
HC-49/U	ECS-147.4-20-1X (Digi-Key #X1032-ND)	ECS-73-20-1X (Digi-Key #X978-ND)
Cylinder	ECS-147.4-18-9X (Digi-Key #XC-1382-ND)	

The capacitor value to use depends somewhat on the layout (e.g. stray capacitance). Values in the range of 22pF to 27pF have been used successfully in various situations.



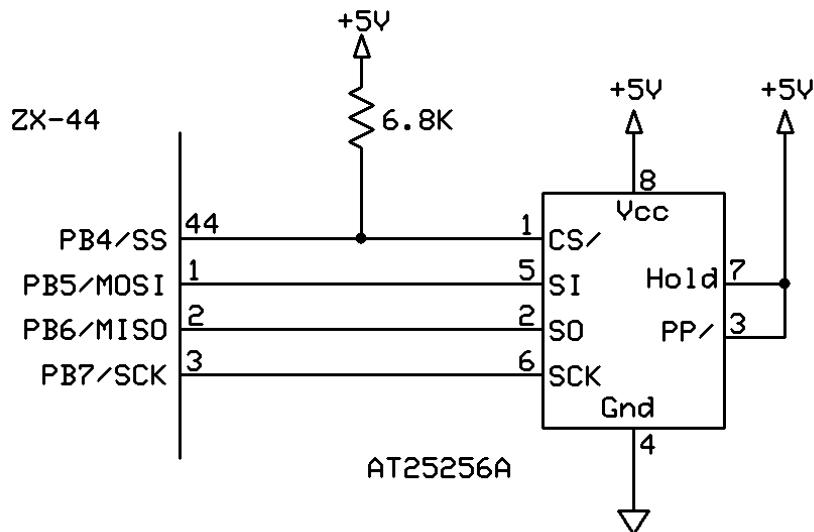
ZX-44 Series Clock Source

If you already have a clock source at the required frequency you can feed that signal directly to pin 8. In this case there would be no connection to pin 7.

Program Memory

Some ZX-44 series devices also require a serial EEPROM in which to store your program's code. (The ZX-44n, ZX-44r, ZX-44s and ZX-44t use internal Flash memory for Program Memory.) The recommended device to use is the Atmel AT25256A (Digi-Key #AT25256A-10PI-2.7-ND). Note, particularly, that the older part without the A suffix will not work due to its slower speed. An EEPROM capable of operation at 7.5MHz or more is required. An alternate part is the ON Semiconductor CAT25256LI-G (Digi-Key #ON Semiconductor CAT25256LI-G-ND).

The recommended connection to the ZX is shown below. The value of the pullup resistor is non-critical. Anything from 4.7K to 22K should work well. Higher values will reduce power consumption somewhat.



ZX-44 Series Program Memory (some VM devices only)

With firmware versions v1.1 and later, the Atmel 25HP512 (Digi-Key #AT25HP512-10PI2.7-ND) can also be used for applications that require larger EEPROM space. The electrical connections for this device are identical to those shown above. However, since this device has different characteristics than the AT25256A, special configuration is required. See Section 11.2 for more information.

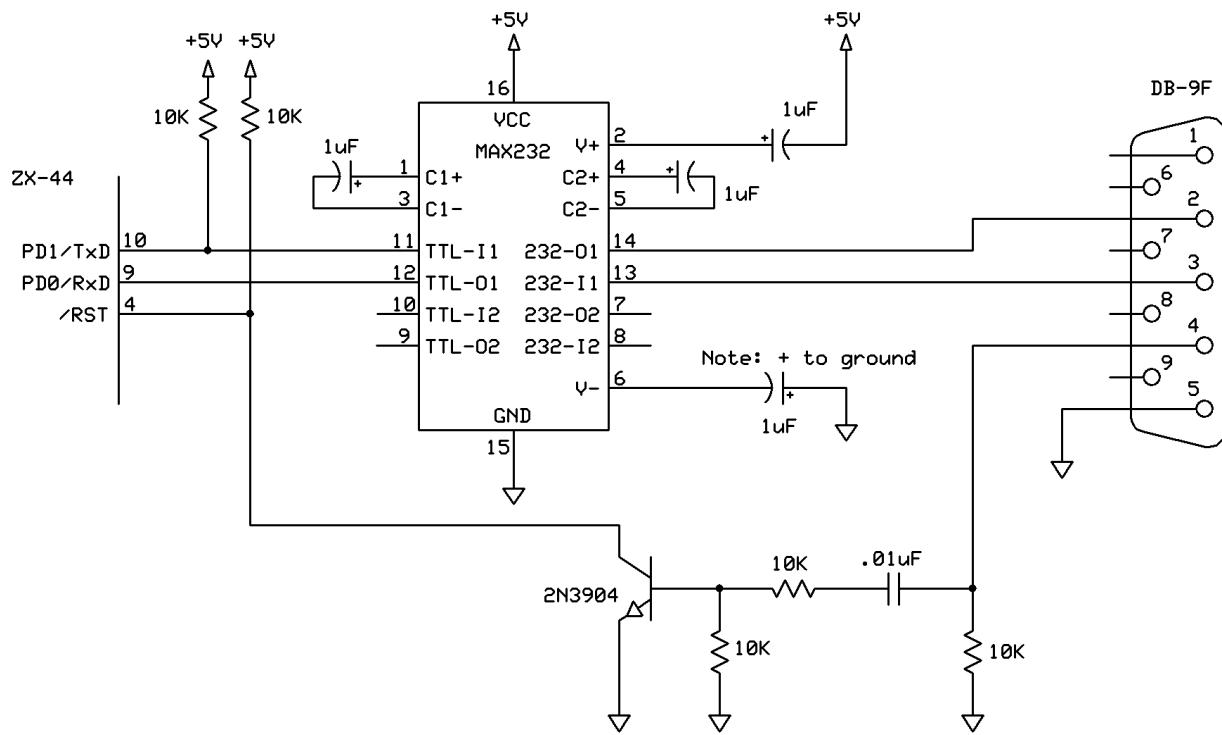
Serial Interface

The ZX-44 series devices require a serial interface for downloading code into Program Memory, performing field updates of the control program, and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-5 volt signals of the ZX to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 4 of the ZX. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 12V from the 5 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies.

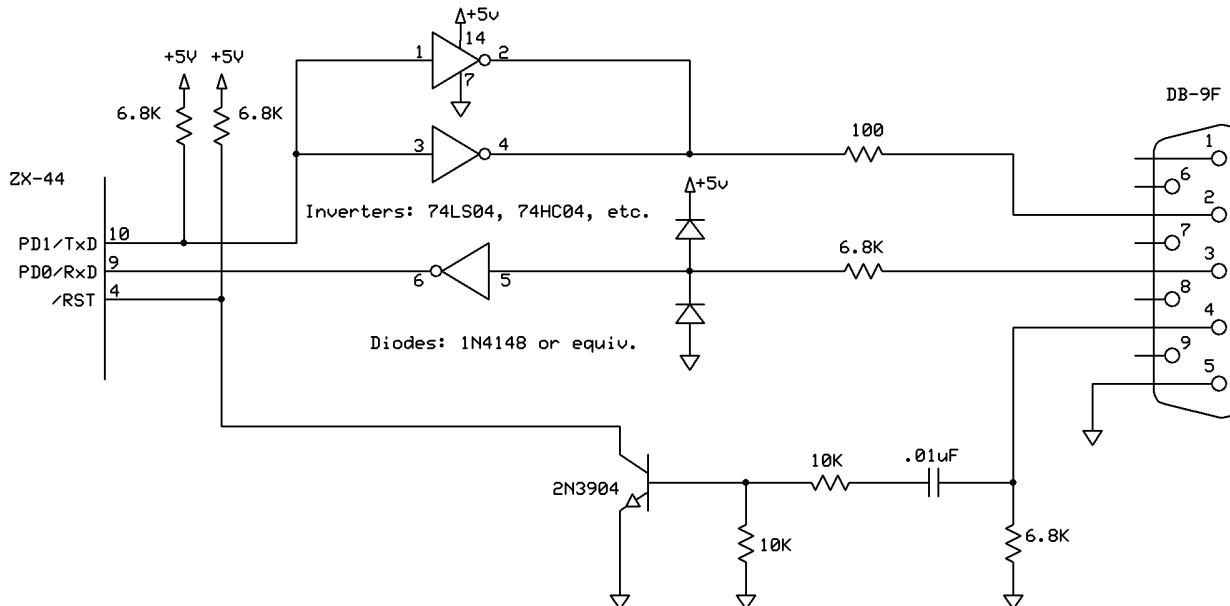
The schematics below contain several references to +5V. For ZX-44 series devices that can operate at a lower voltage, e.g. 3.3V, these references should be to the actual operating voltage.



Recommended ZX-44 Series Serial Interface Circuit

The pullup resistor shown on the ZX's transmit output (pin 10) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

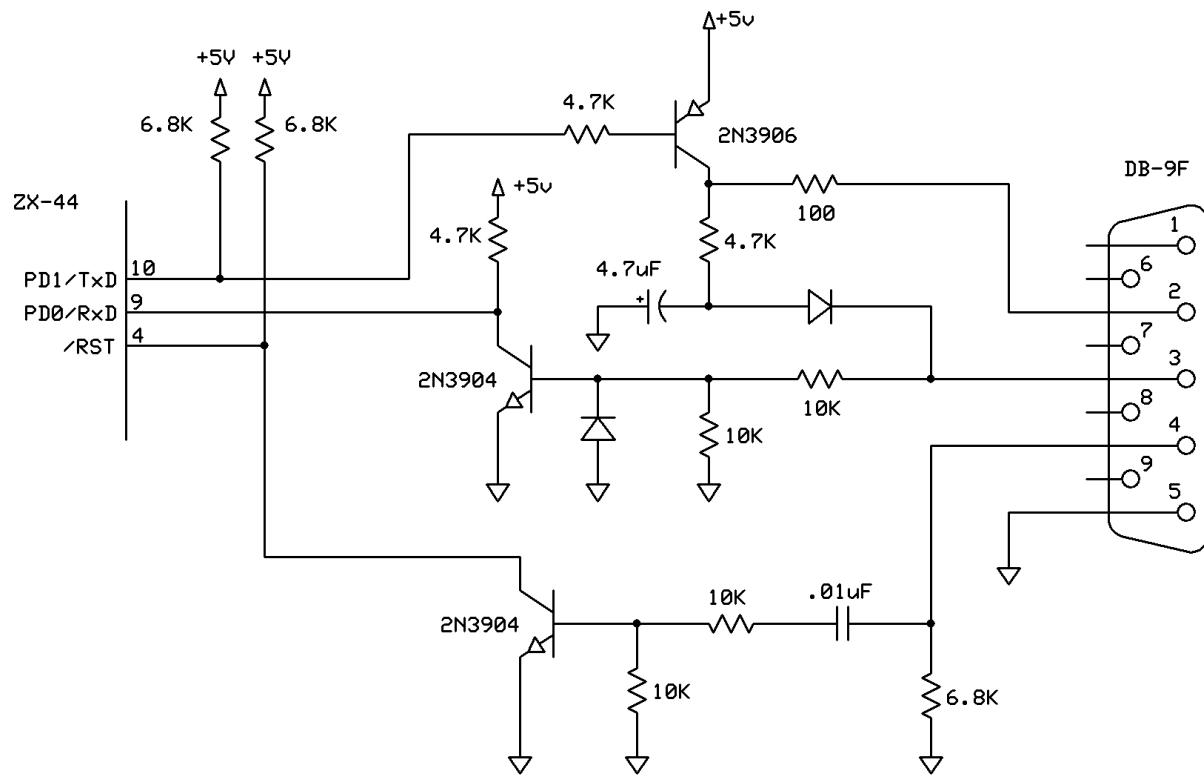
There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.



Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

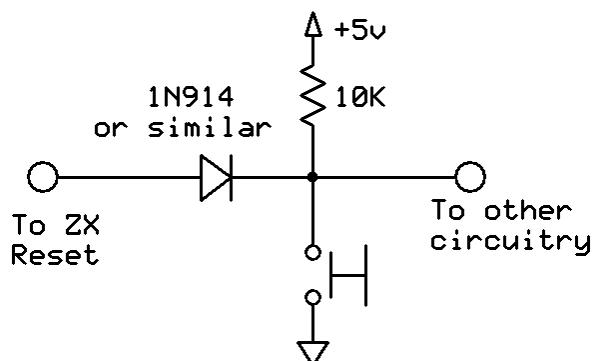
The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Alternate Serial Interface Circuit #2

Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 4 of the ZX and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix G - ZX-1281 Series Hardware Reference

The ZX-1281 series devices are the 64-pin TQFP package version of the Atmel AVR ATmega1281 microcontroller that has been programmed with the ZX control firmware. In order to use the ZX-1281 you must add several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

G.1 ZX-1281 Series Specifications

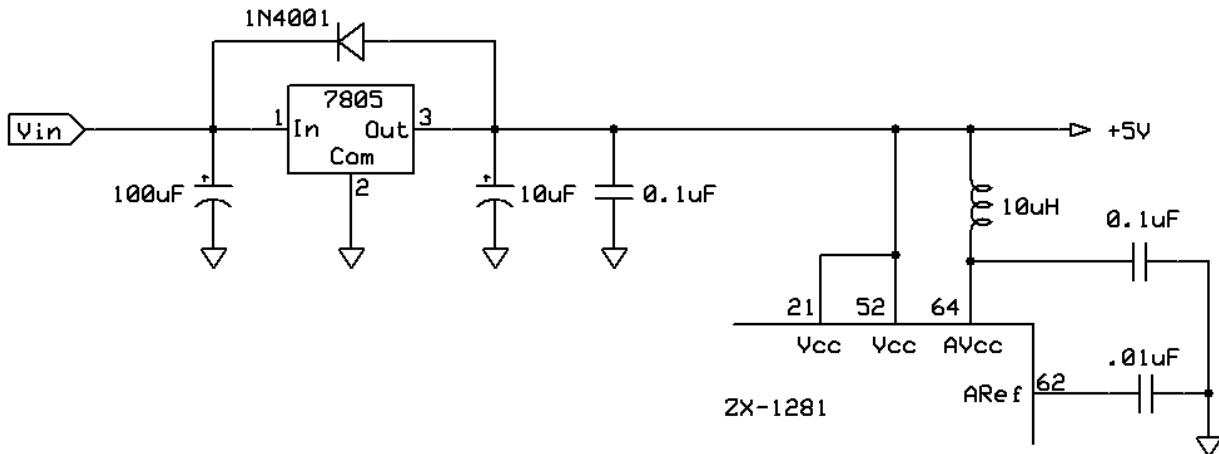
The electrical specifications of the ZX-1281 series devices are exactly those of the ATmega1281. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

G.2 ZX-1281 Series Required External Components

The circuits described below represent the minimum external circuitry required to operate the ZX-1281 series devices. Depending on your application, you may need additional circuitry to take advantage of the capabilities of the ZX-1281 series device. In contrast to some other ZX devices, these devices do not require an external EEPROM for program storage. Rather, the compiled user program is stored in the ZX's internal Flash memory. The maximum user program size is 60K bytes for the ZX-1281 and 124K for the ZX-1281n.

Power Source

The ZX-1281 series devices need a regulated voltage source capable of providing at least 200mA of current. The voltage typically used is 5 volts but the ZX-1281 will operate between 4.5 volts and 5.5 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



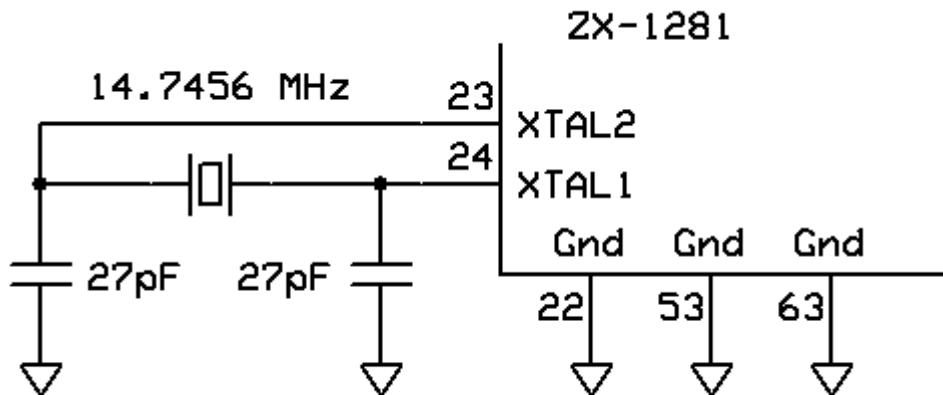
ZX-1281 Series Power Source

If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the two capacitors on the right side of the diagram. In this case, pin 64 would be connected directly to the power source (same as pin 21, etc.) and pin 62 can be left open. Although not shown on this diagram, the ground pins of the ZX-1281 (pins 22, 53, and 63) must be connected to the common ground of the system.

Clock Source

The ZX-1281 series devices require a clock source running at 14.7456MHz. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 23 and 24 along with the necessary capacitors.

The recommended crystal to use is ECS-147.4-20-4 (Digi-Key #X175-ND). Alternate crystals are the ECS-147.4-20-1 (Digi-Key #X142-ND) in a slightly larger package and the ECS-147.4-18-9 (Digi-Key #XC-957ND) in a cylindrical package. The capacitor value to use depends somewhat on the layout (e.g. stray capacitance). Values in the range of 22pF to 27pF have been used successfully in various situations.



ZX-1281 Series Clock Source

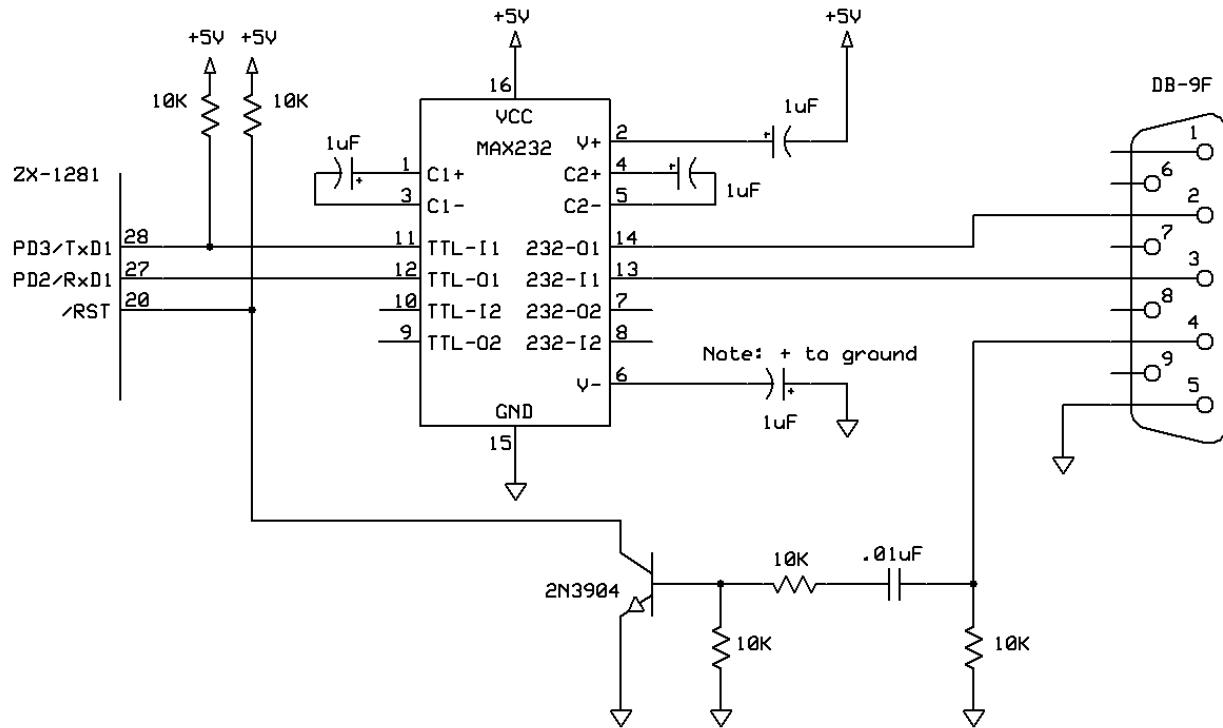
If you already have a clock source at the required frequency you can feed that signal directly to pin 24. In this case there would be no connection to pin 23.

Serial Interface

The ZX-1281 series devices require a serial interface for downloading code into Program Memory, performing field updates of the control program, and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-5 volt signals of the ZX to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 20 of the ZX-1281. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

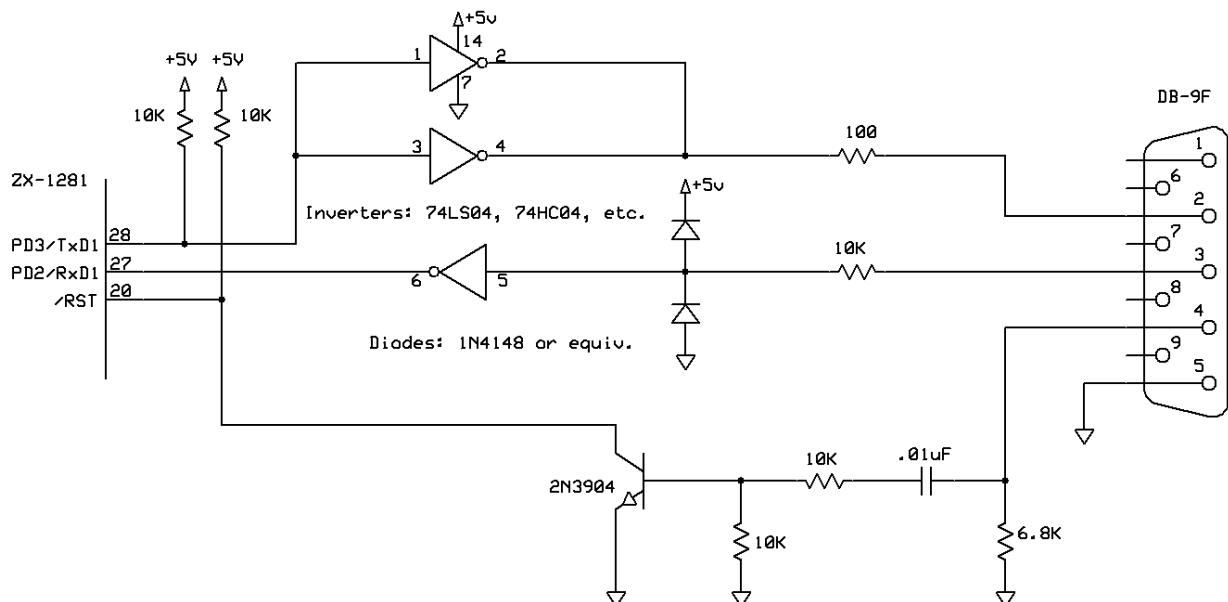
One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 12V from the 5 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies.



Recommended ZX-1281 Series Serial Interface Circuit

The pullup resistor shown on the ZX's Com1 transmit output (pin 28) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.

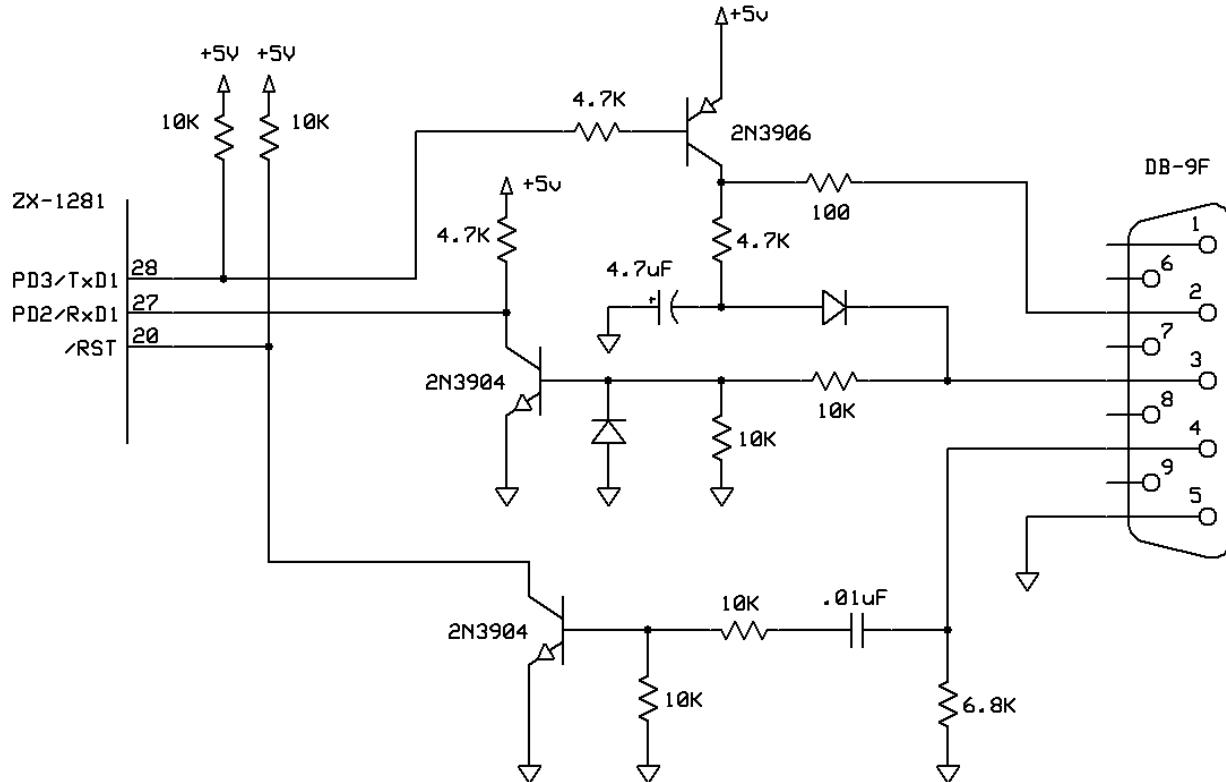


Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not

meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

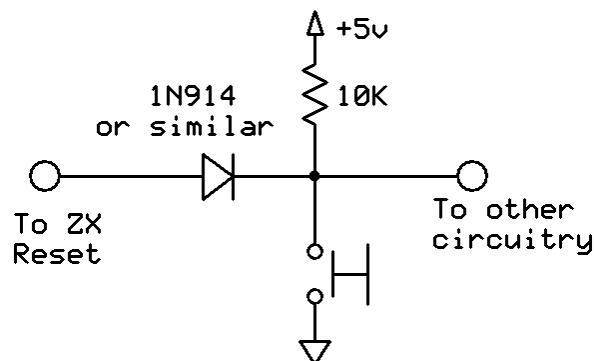
The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Alternate Serial Interface Circuit #2

Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 20 of the ZX and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix H - ZX-1280 Series Hardware Reference

The ZX-1280 series devices utilize the 100-pin TQFP package version of the Atmel AVR ATmega1280 microcontroller that has been programmed with the ZX control firmware. In order to use the ZX-1280 or ZX-1280n you must add several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

H.1 ZX-1280 Series Specifications

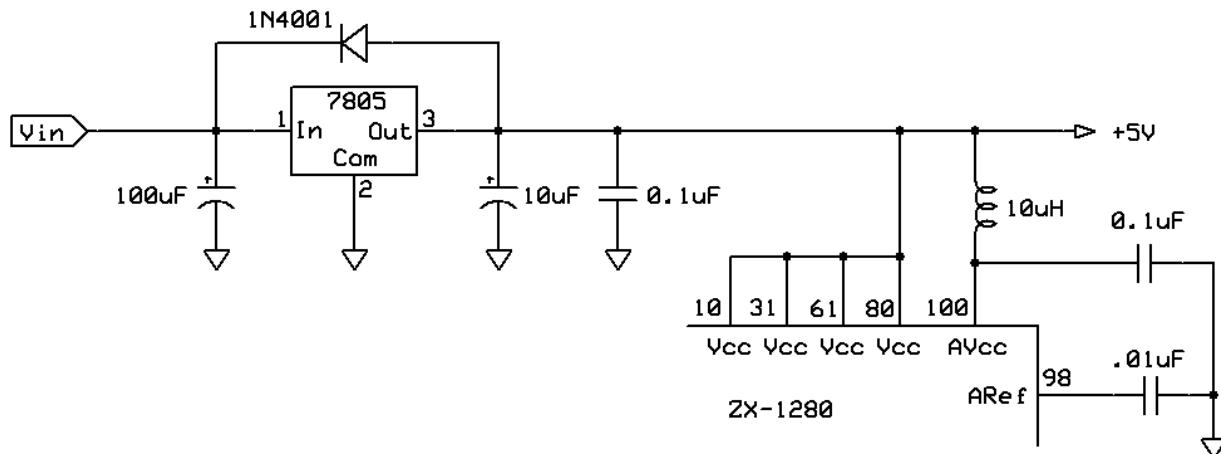
The electrical specifications of the ZX-1280 series devices are exactly those of the ATmega1280. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

H.2 ZX-1280 Series Required External Components

The circuits described below represent the minimum external circuitry required to operate the ZX-1280 or ZX-1280n. Depending on your application, you may need additional circuitry to take advantage of the capabilities of the ZX-1280 series device. In contrast to some other ZX devices, the ZX-1280 series does not require an external EEPROM for program storage. Rather, the compiled user program is stored in the ZX's internal Flash memory. The maximum user program size is 60K bytes for the ZX-1280 and 124K for the ZX-1280n.

Power Source

The ZX-1280 series devices need a regulated voltage source capable of providing at least 200mA of current. The voltage typically used is 5 volts but the ZX-1280 will operate between 4.5 volts and 5.5 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.

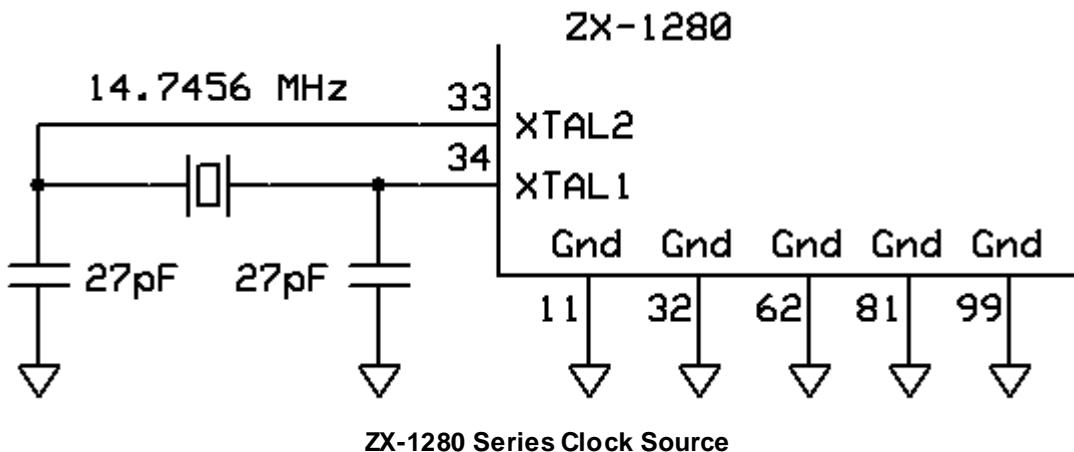


If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the two capacitors on the right side of the diagram. In this case, pin 64 would be connected directly to the power source (same as pin 21, etc.) and pin 62 can be left open. Although not shown on this diagram, the ground pins of the ZX-1280 (pins 11, 32, 62, 81 and 99) must be connected to the common ground of the system.

Clock Source

The ZX-1280 series devices require a clock source running at 14.7456MHz. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 33 and 34 along with the necessary capacitors.

The recommended crystal to use is ECS-147.4-20-4 (Digi-Key #X175-ND). Alternate crystals are the ECS-147.4-20-1 (Digi-Key #X142-ND) in a slightly larger package and the ECS-147.4-18-9 (Digi-Key #XC-957ND) in a cylindrical package. The capacitor value to use depends somewhat on the layout (e.g. stray capacitance). Values in the range of 22pF to 27pF have been used successfully in various situations.



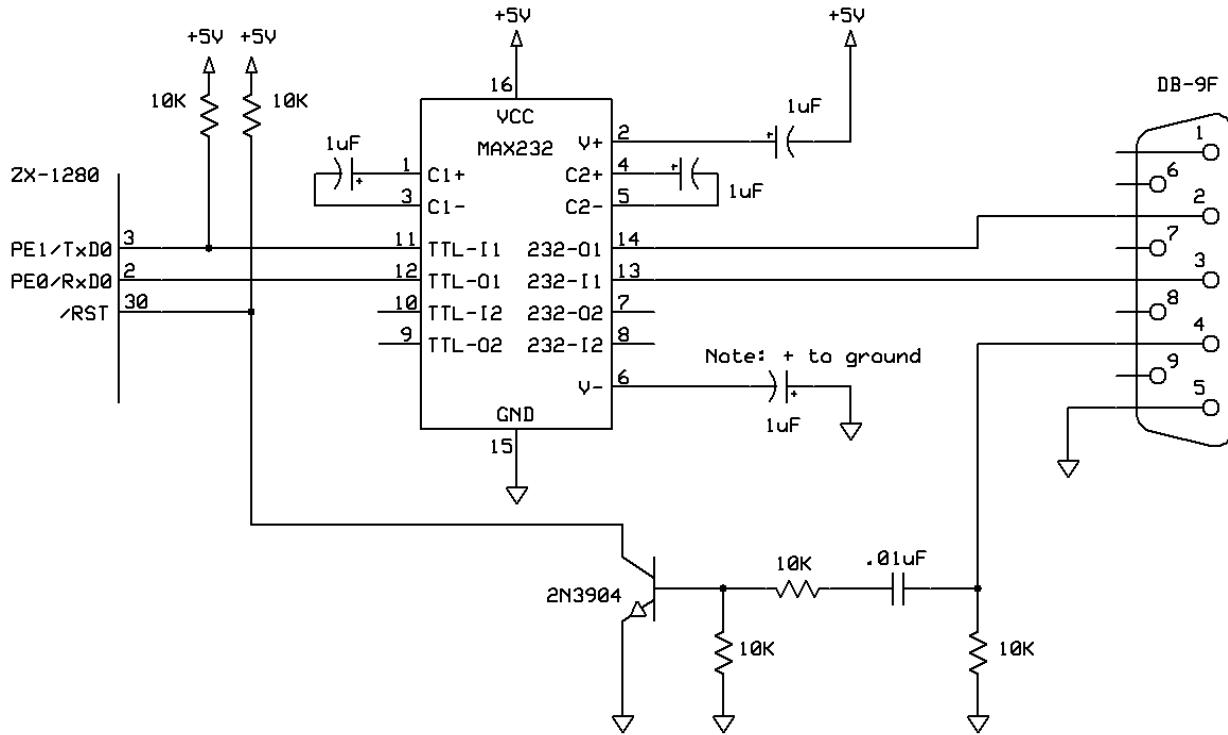
If you already have a clock source at the required frequency you can feed that signal directly to pin 34. In this case there would be no connection to pin 33.

Serial Interface

The ZX-1280 series devices require a serial interface for downloading code into Program Memory, performing field updates of the control program, and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-5 volt signals of the ZX to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 30 of the ZX-1280. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

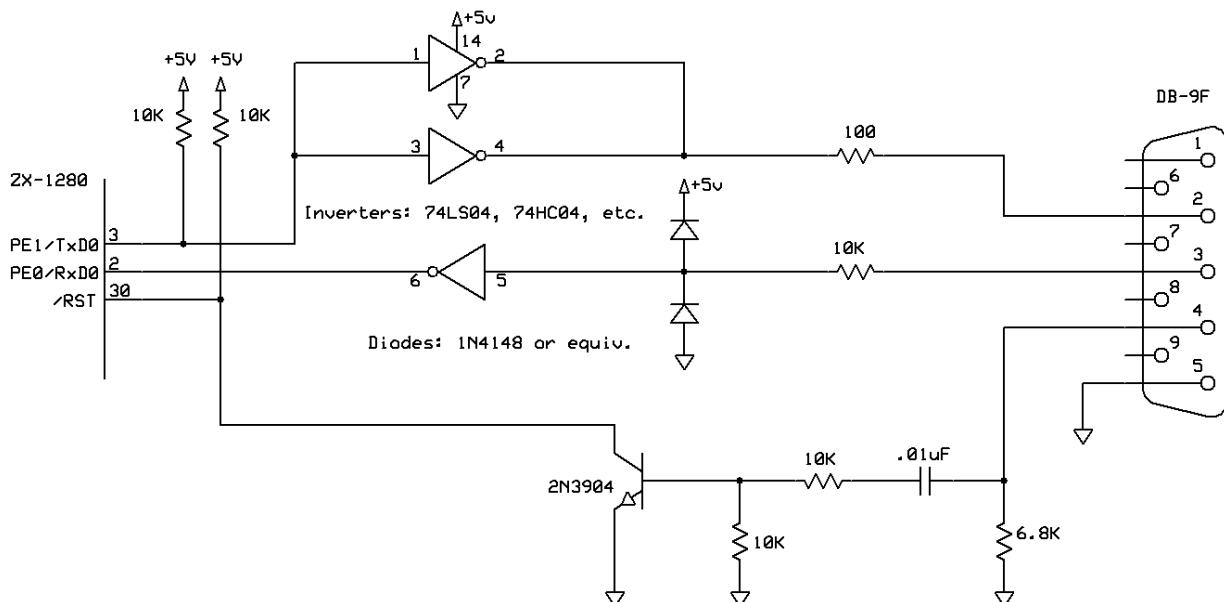
One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 12V from the 5 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies.



Recommended ZX-1280 Series Serial Interface Circuit

The pullup resistor shown on the ZX's Com1 transmit output (pin 28) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.

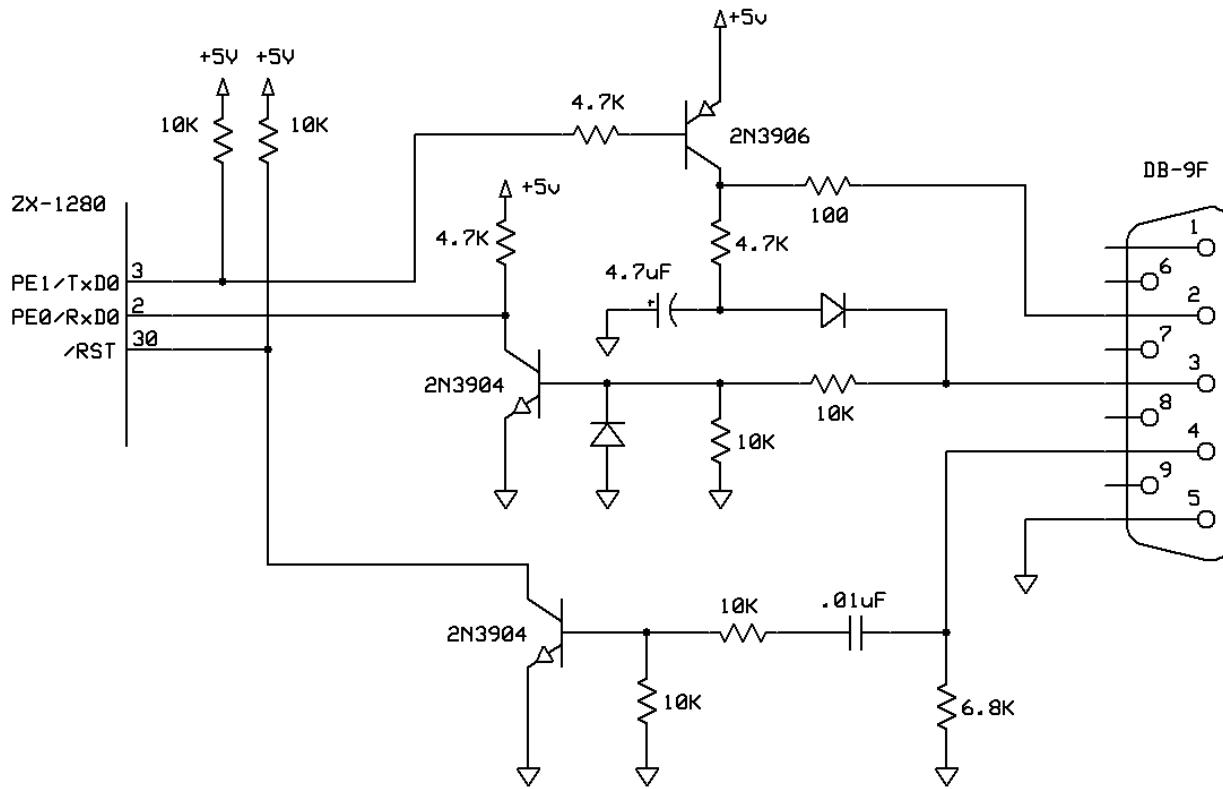


Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not

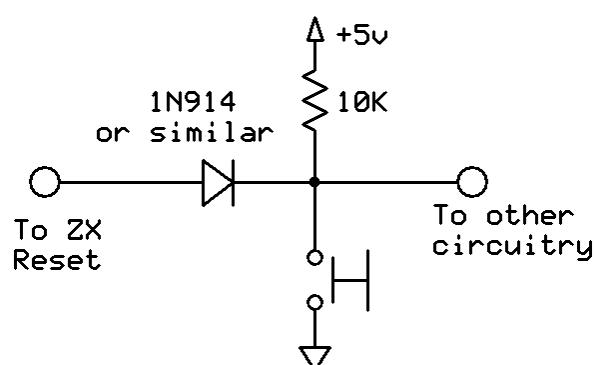
meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 20 of the ZX and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix I - ZX-328n and ZX-328I Hardware Reference

The ZX-328n and ZX-328I devices utilize the 28-pin DIP package version of the Atmel AVR ATmega328P microcontroller that has been programmed with a special ZX control program. In order to use the ZX-328n or ZX-328I (hereafter, collectively referred to as simply ZX-328) you must supply several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

I.1 ZX-328 Series Specifications

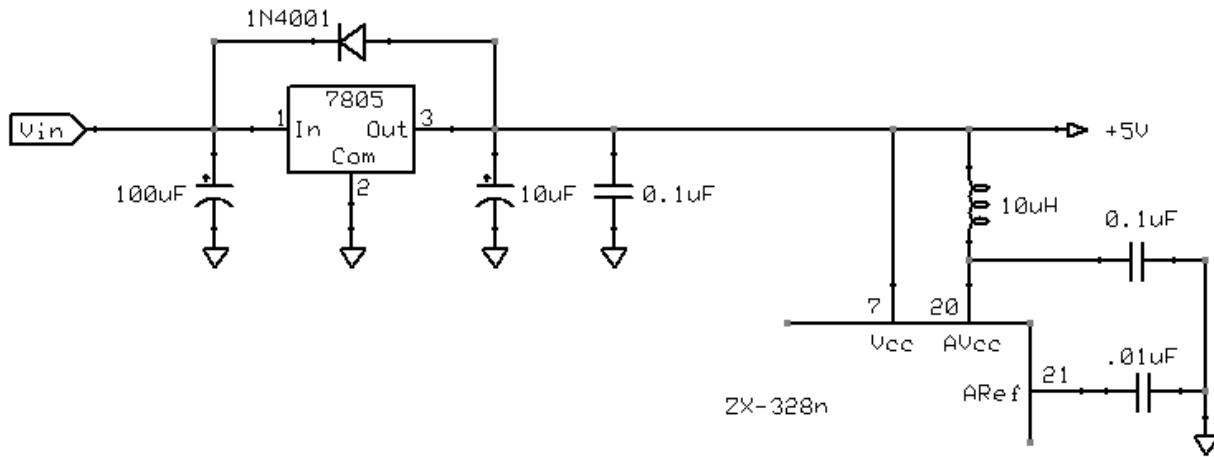
The electrical specifications of the ZX-328 are exactly those of the ATmega328P. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

I.2 ZX-328 Series Required External Components

The circuits described below represent the minimum external circuitry required to operate the ZX-328. Depending on your application, you may need additional circuitry to take advantage of the capabilities of the ZX-328 series device. In contrast to some other ZX devices, the ZX-328 devices do not require an external EEPROM for program storage. Rather, the compiled user program is stored in the ZX's internal Flash memory. The maximum user program size is 30K.

Power Source

The ZX-328 needs a regulated voltage source capable of providing at least 200mA of current. The ZX-328n will operate reliably with a supply voltage between 4.5 volts and 5.5 volts while the ZX-328I supply voltage can be as low as 2.75 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



ZX-328 Power Source

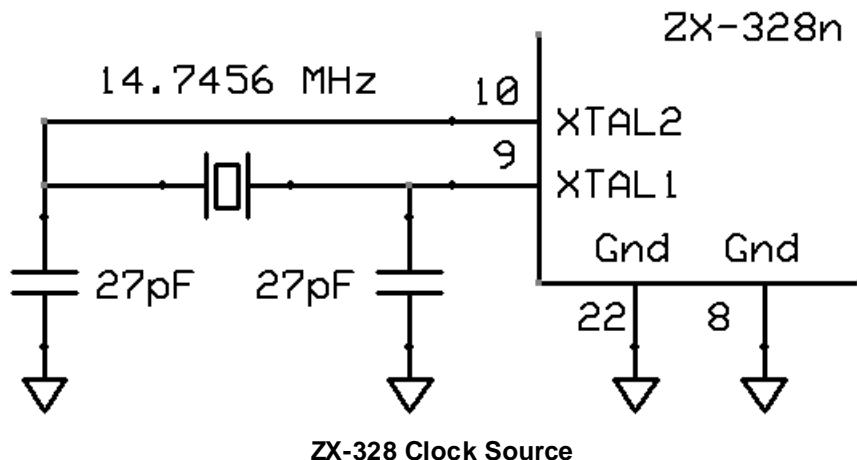
If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the two capacitors on the right side of the diagram. In this case, pin 20 would be connected directly to the power source (same as pin 7) and pin 21 can be left open. Although not shown on this diagram, the ground pins of the ZX-328 (pins 8 and 22) must be connected to the common ground of the system.

Clock Source

The ZX-328n device requires a clock source running at 14.7456MHz while the ZX-32I requires a frequency of 7.3728MHz. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 9 and 10 along with the necessary capacitors. Some recommended crystals are shown in the table below.

Recommended Crystals

Package	ZX-328n	ZX-32I
HC-49/US	ECS-147.4-20-4X (Digi-Key #X1102-ND)	ECS-73-20-4X (Digi-Key #X1084-ND)
HC-49/U	ECS-147.4-20-1X (Digi-Key #X1032-ND)	ECS-73-20-1X (Digi-Key #X978-ND)
Cylinder	ECS-147.4-18-9X (Digi-Key #XC-1382-ND)	



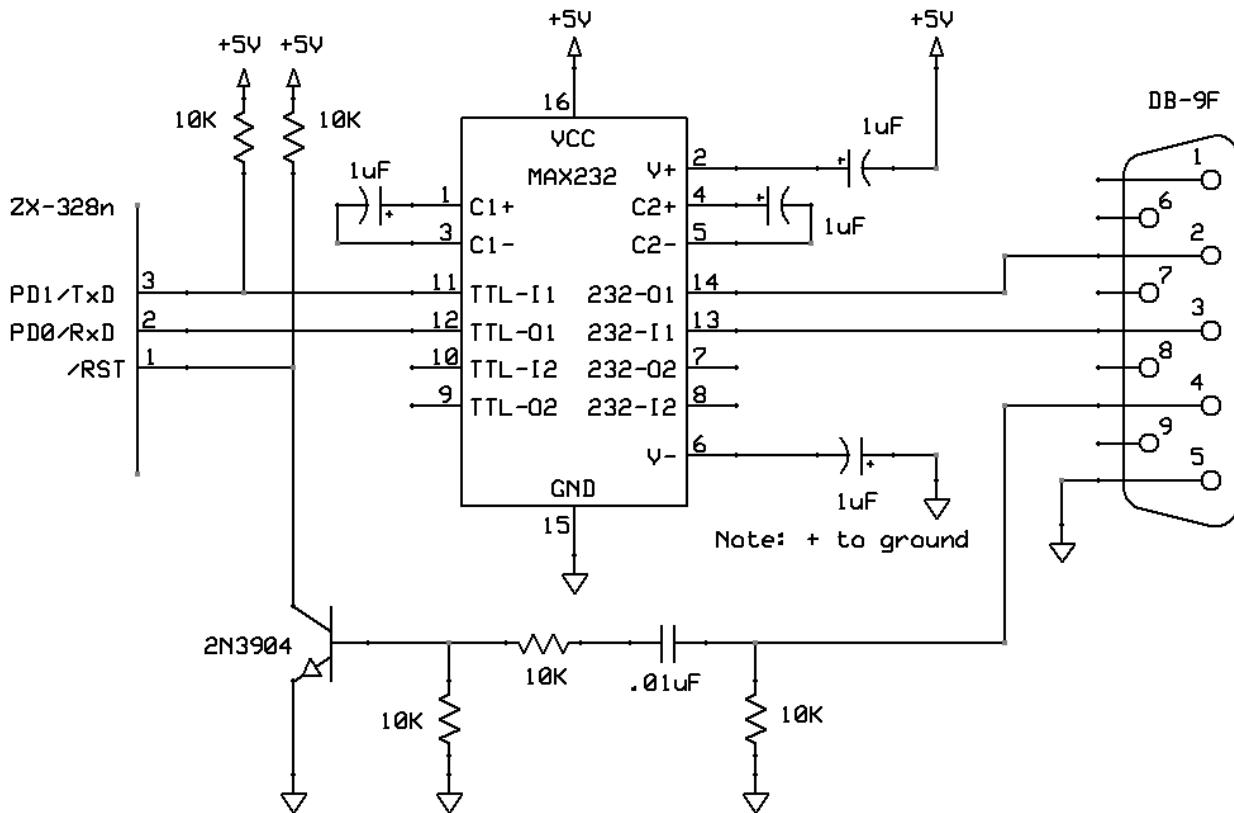
If you already have a clock source at the required frequency you can feed that signal directly to pin 10. In this case there would be no connection to pin 9.

Serial Interface

The ZX-328 device requires a serial interface for downloading code into Program Memory and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-5 volt signals of the ZX-328 to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX-328 to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 1 of the ZX-328. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX-328 will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX-328 receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

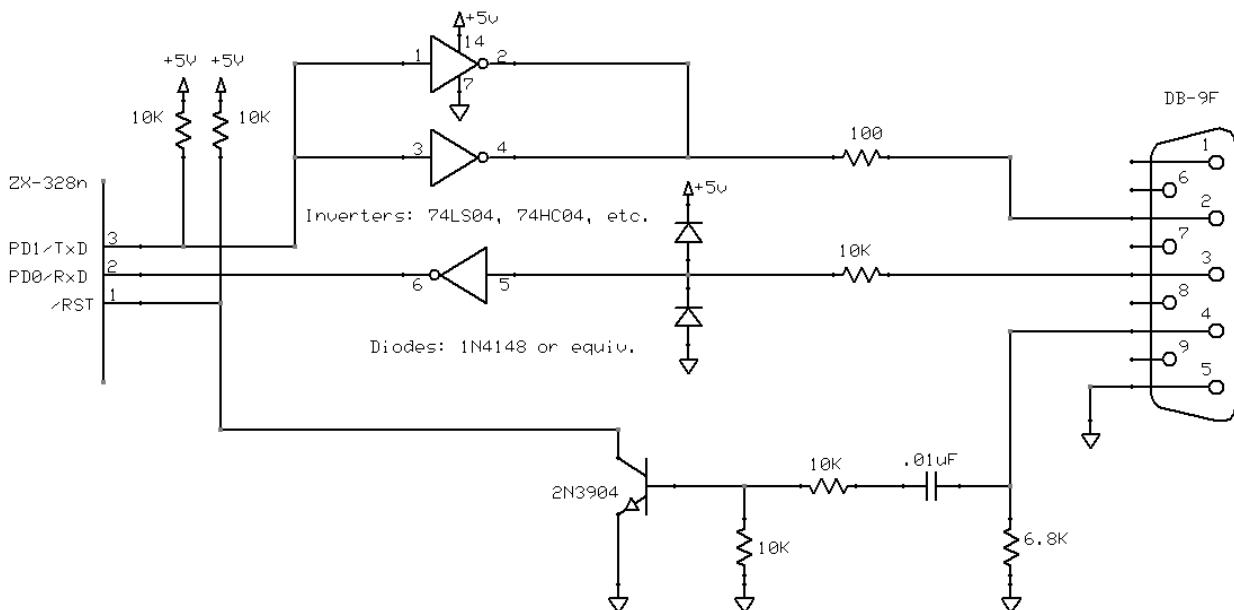
One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 12V from the 5 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies.



Recommended ZX-328 Serial Interface Circuit

The pullup resistor shown on the ZX's Com1 transmit output (pin 3) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

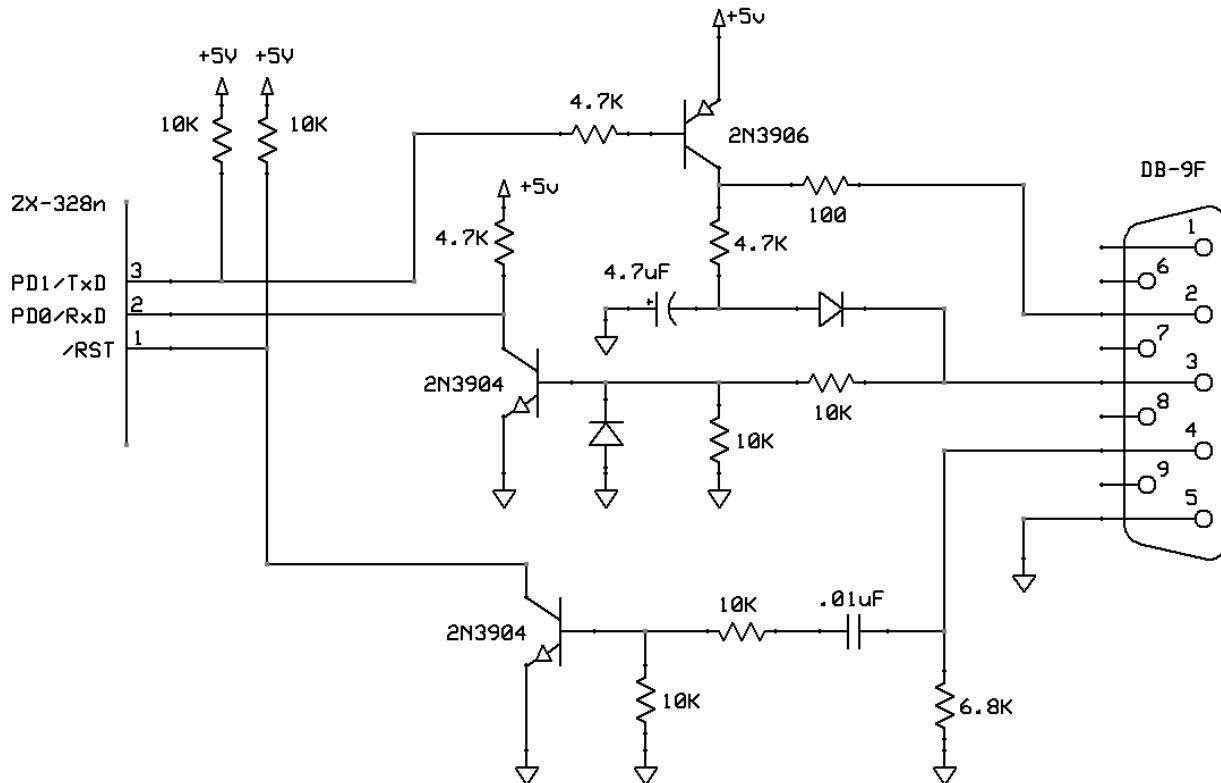
There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.



Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

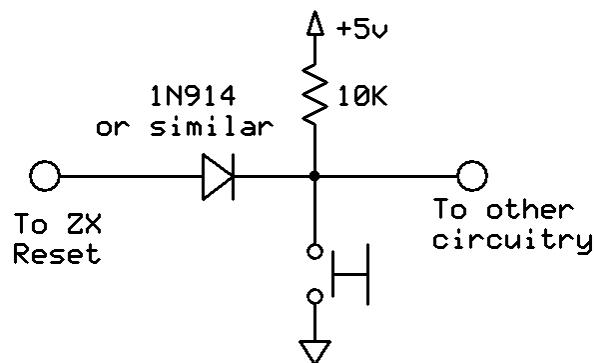
The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Alternate Serial Interface Circuit #2

Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 1 of the ZX-328 and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix J - ZX-32n and ZX-32I Hardware Reference

The ZX-32n and ZX-32I devices utilize the 32-pin TQFP package version of the Atmel AVR ATmega328P microcontroller that has been programmed with a special ZX control program. In order to use the ZX-32n or ZX-32I (hereafter, collectively referred to as simply ZX-328) you must supply several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

J.1 ZX-32 Series Specifications

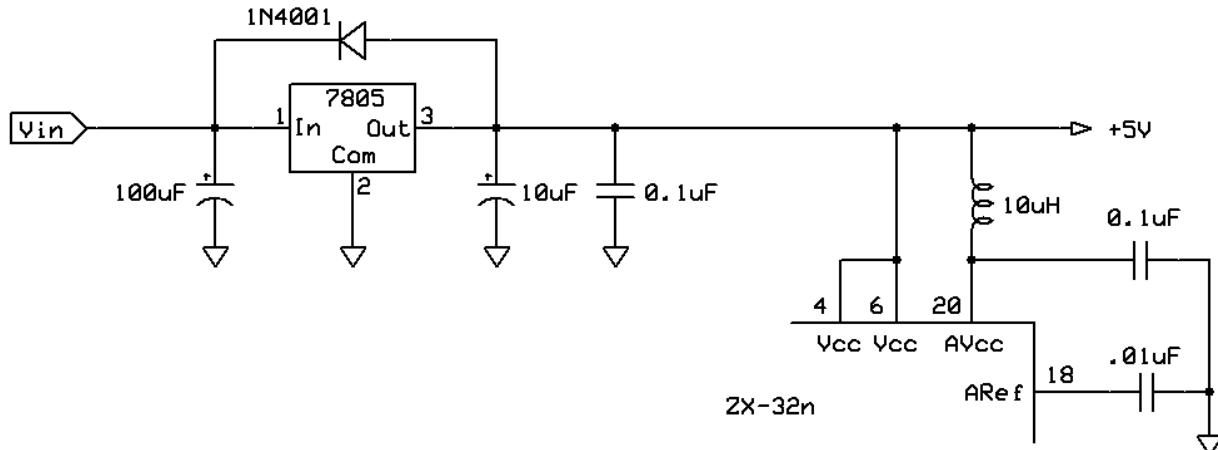
The electrical specifications of the ZX-32 are exactly those of the ATmega328P. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

J.2 ZX-32 Series Required External Components

The circuits described below represent the minimum external circuitry required to operate the ZX-32. Depending on your application, you may need additional circuitry to take advantage of the capabilities of the ZX-32 series device. In contrast to some other ZX devices, the ZX-32 devices do not require an external EEPROM for program storage. Rather, the compiled user program is stored in the ZX's internal Flash memory. The maximum user program size is 30K.

Power Source

The ZX-32 needs a regulated voltage source capable of providing at least 200mA of current. The ZX-32n will operate reliably with a supply voltage between 4.5 volts and 5.5 volts while the ZX-32I supply voltage can be as low as 2.75 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



ZX-32 Power Source

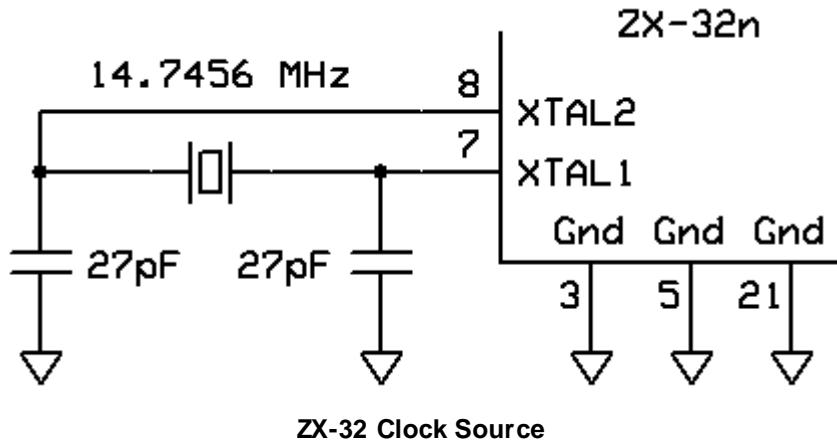
If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the two capacitors on the right side of the diagram. In this case, pin 18 would be connected directly to the power source (same as pins 4 and 6) and pin 20 can be left open. Although not shown on this diagram, the ground pins of the ZX-32 (pins 3, 5 and 21) must be connected to the common ground of the system.

Clock Source

The ZX-32n device requires a clock source running at 14.7456MHz while the ZX-32I requires a frequency of 7.3728MHz. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 9 and 10 along with the necessary capacitors. Some recommended crystals are shown in the table below.

Recommended Crystals

Package	ZX-32n	ZX-32I
HC-49/US	ECS-147.4-20-4X (Digi-Key #X1102-ND)	ECS-73-20-4X (Digi-Key #X1084-ND)
HC-49/U	ECS-147.4-20-1X (Digi-Key #X1032-ND)	ECS-73-20-1X (Digi-Key #X978-ND)
Cylinder	ECS-147.4-18-9X (Digi-Key #XC-1382-ND)	



ZX-32 Clock Source

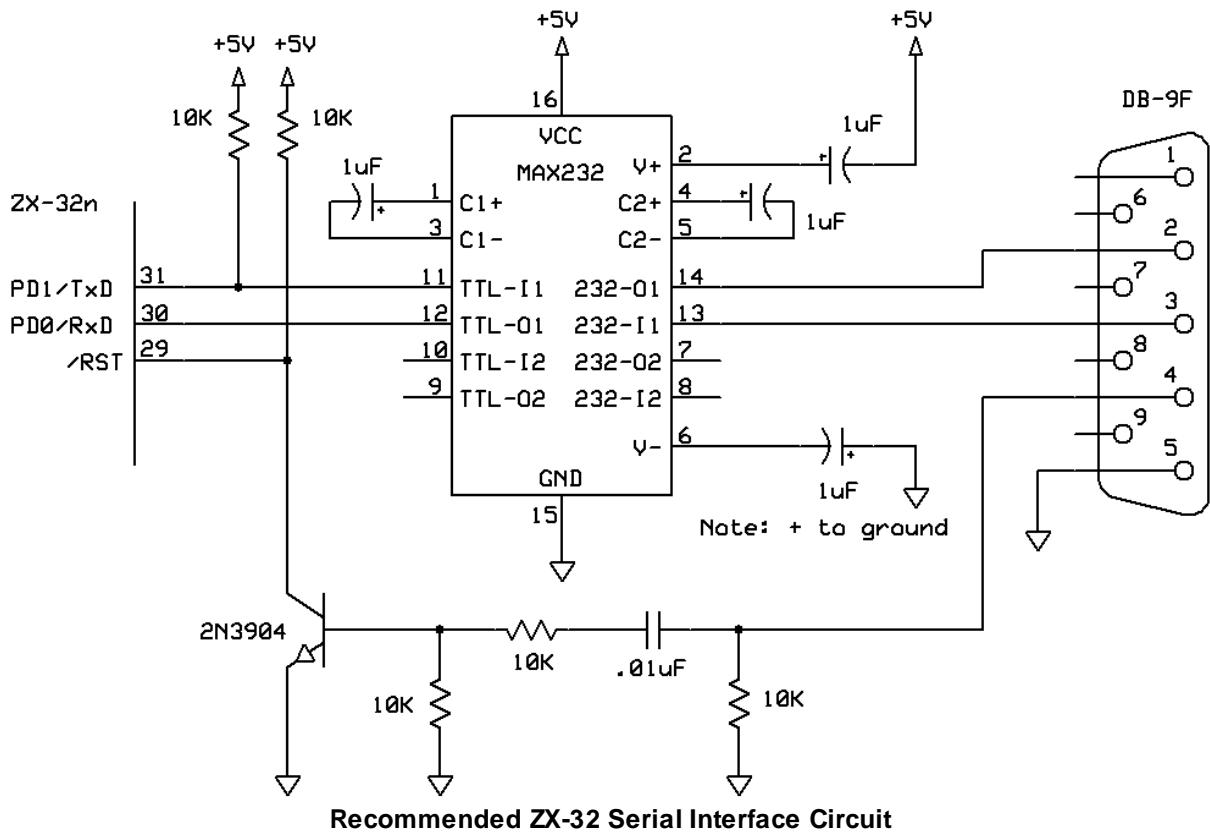
If you already have a clock source at the required frequency you can feed that signal directly to pin 8. In this case there would be no connection to pin 7.

Serial Interface

The ZX-32 device requires a serial interface for downloading code into Program Memory and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-5 volt signals of the ZX-32 to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX-32 to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 29 of the ZX-32. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX-32 will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX-32 receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

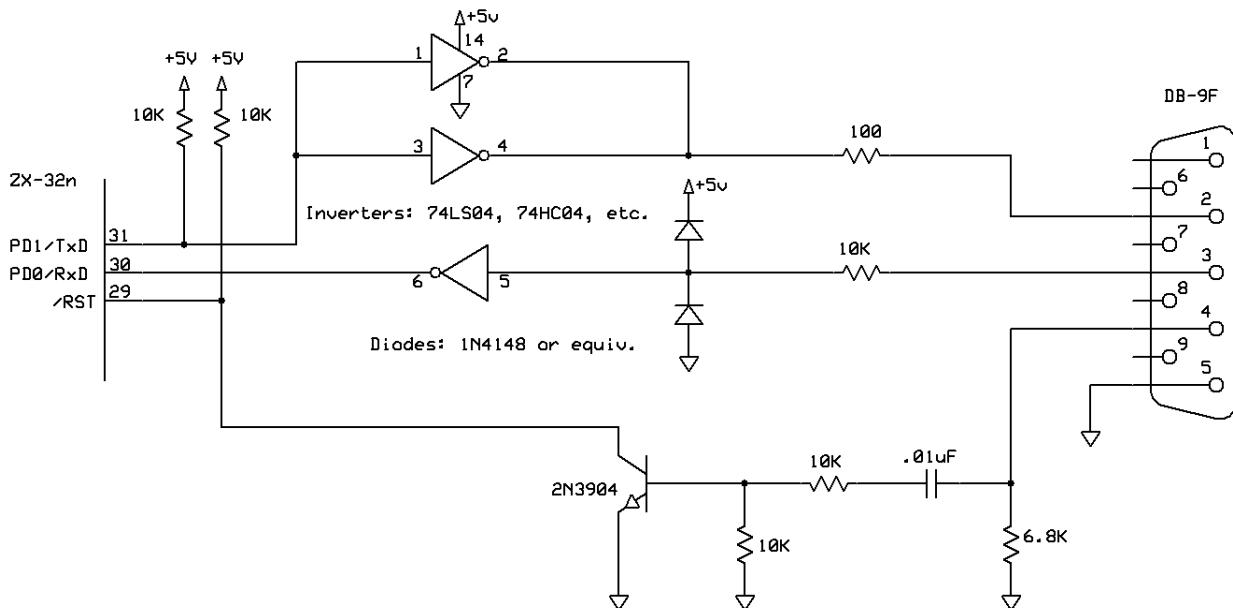
One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 12V from the 5 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies.



Recommended ZX-32 Serial Interface Circuit

The pullup resistor shown on the ZX's Com1 transmit output (pin 31) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

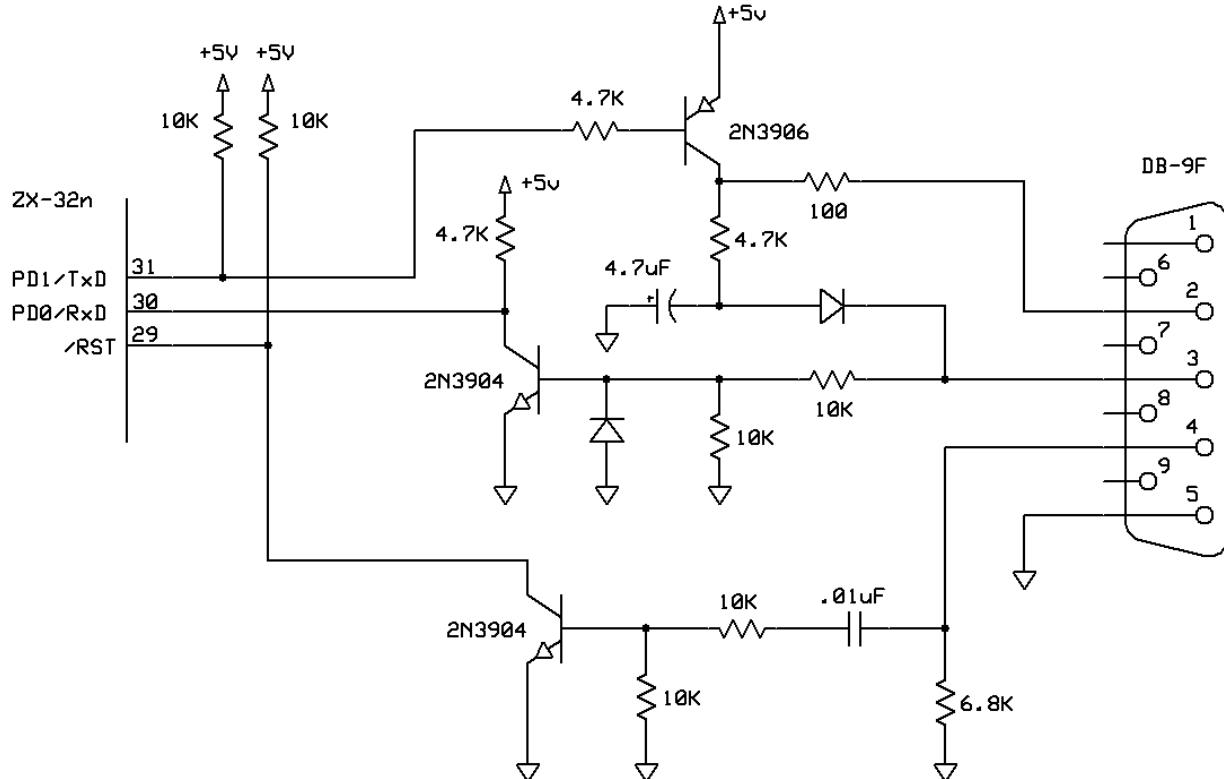
There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.



Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

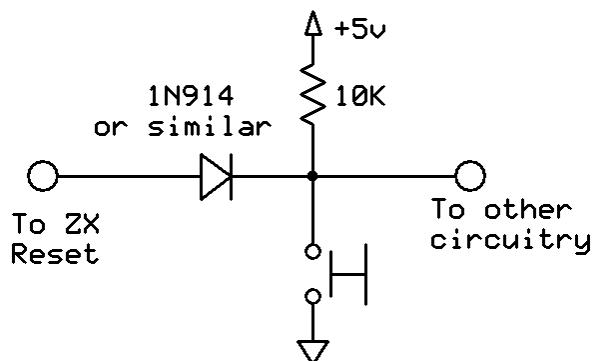
The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Alternate Serial Interface Circuit #2

Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 29 of the ZX-32 and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix K - ZX-32a4 and ZX-128a4u Hardware Reference

The ZX-32a4 and ZX-128a4u devices utilize the 44-pin TQFP package version of the Atmel AVR ATxmega32A4 and ATxmega128A4U microcontrollers that have been programmed with a special ZX control program. In order to use the ZX-32a4 or ZX-128a4u you must supply several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

K.1 ZX-32a4 and ZX-128a4u Specifications

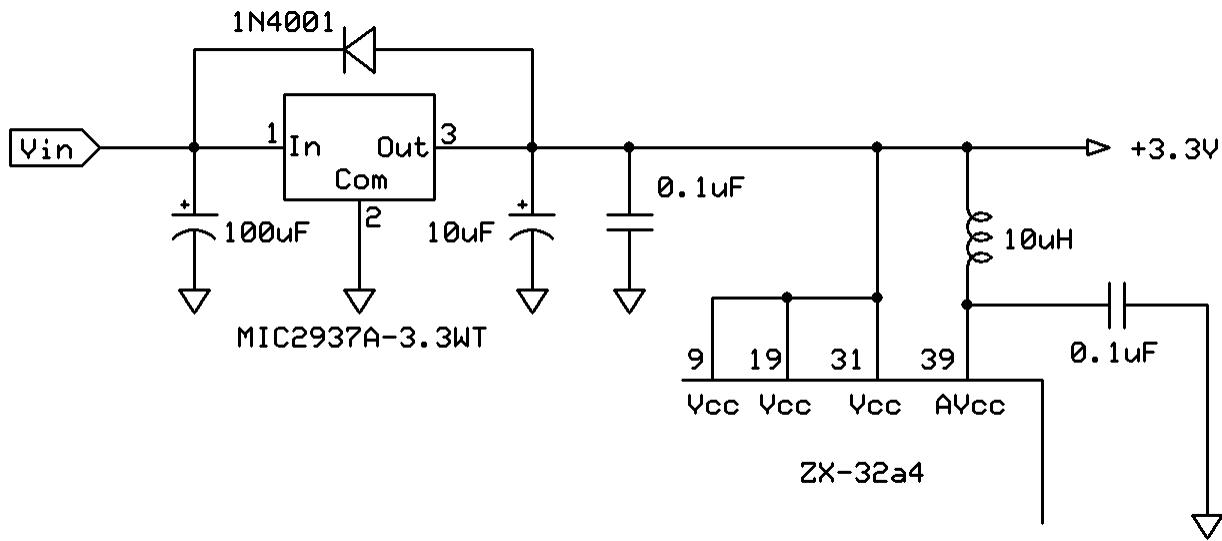
The electrical specifications of the ZX-32a4 and ZX-128a4u are exactly those of the ATxmega32A4 and ATxmega128A4U, respectively. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

K.2 ZX-32a4 and ZX-128a4u Required External Components

The circuits described below represent the minimum external circuitry required to operate the ZX-32a4 and ZX-128a4u. Depending on your application, you may need additional circuitry to take advantage of the capabilities of the ZX-32a4 and ZX-128a4u series devices. In contrast to some other ZX devices, the ZX-32a4 and ZX-128a4u devices do not require an external EEPROM for program storage. Rather, the compiled user program is stored in the ZX's internal Flash memory. The maximum user program size is 32K for the ZX-32a4 and 128K for the ZX-128a4u.

Power Source

The ZX-32a4 and ZX-128a4u need a regulated voltage source capable of providing at least 200mA of current. The ZX-32a4 and ZX-128a4u will operate reliably with a supply voltage between 2.6 volts and 3.6 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



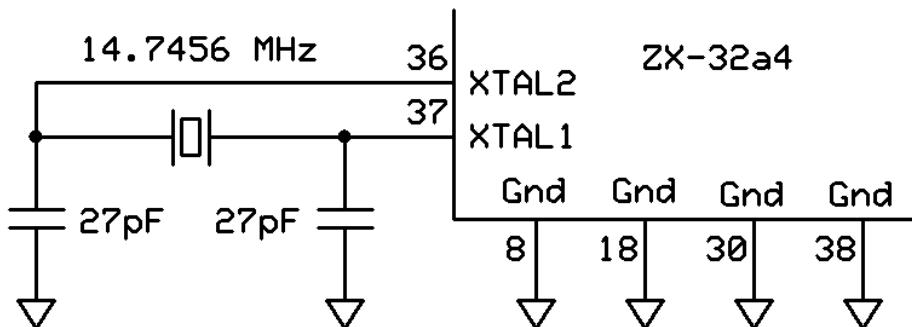
ZX-32a4 and ZX-128a4u Power Source

If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the capacitor on the right side of the diagram. In this case, pin 39 would be connected directly to the power source (same as pins 9, 19 and 31). Although not shown on this diagram, the ground pins of the chip (pins 8, 18, 30 and 38) must be connected to the common ground of the system.

Clock Source

The ZX-32a4n and ZX-128a4u devices require a clock source running at 14.7456MHz. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 36 and 37 along with the necessary capacitors. Some recommended crystals are shown in the table below.

Recommended Crystals	
Package	Crystal
HC-49/US	ECS-147.4-20-4X (Digi-Key #X1102-ND)
HC-49/U	ECS-147.4-20-1X (Digi-Key #X1032-ND)
Cylinder	ECS-147.4-18-9X (Digi-Key #XC-1382-ND)



ZX-32a4 and ZX-128a4u Clock Source

If you already have a clock source at the required frequency you can feed that signal directly to pin 37. In this case there would be no connection to pin 36.

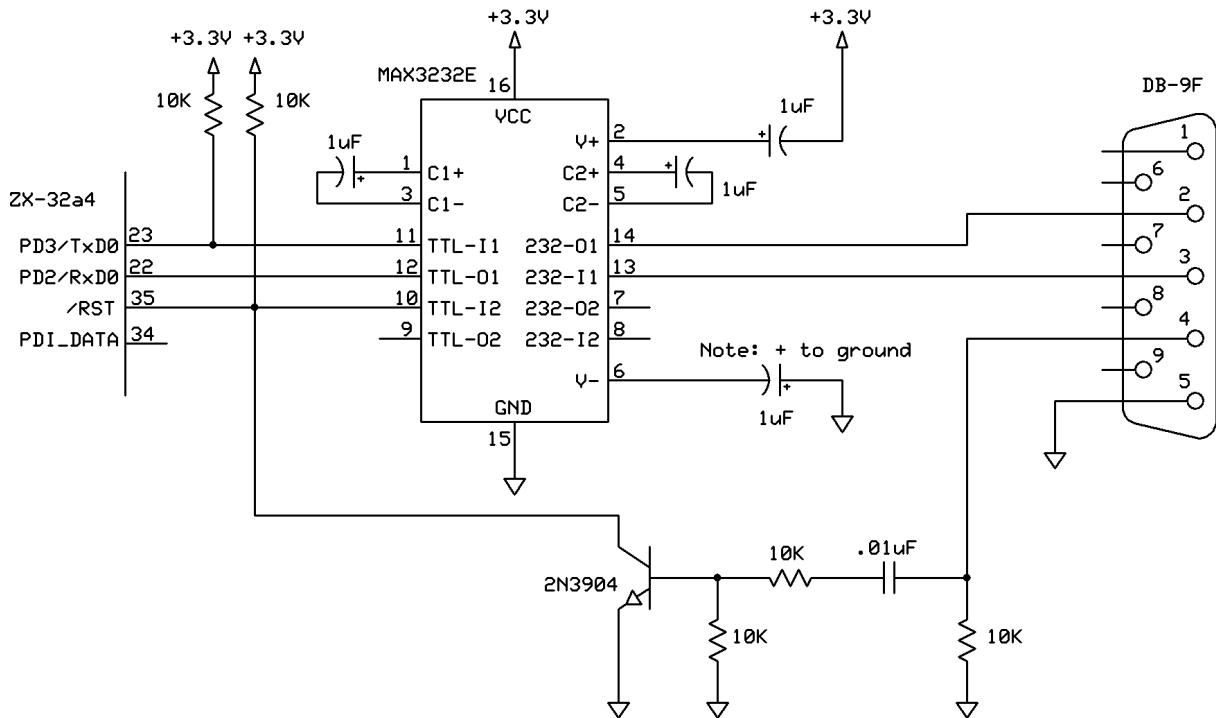
Serial Interface

The ZX-32a4 and ZX-128a4u devices require a serial interface for downloading code into Program Memory and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX3232E chip functions as an RS-232 level converter that translates the 0-3.3 volt signals of the chip to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the chip to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 35 of the chip. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the device will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the chip receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

One advantage to using this recommended circuit is that the MAX3232E chip includes a DC-DC converter that produces nominally +/- 10V from a 3.3 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX3232E while the negative voltage is available on pin 6. Consult the datasheet for the MAX3232E device for information on the current capacity of these supplies.

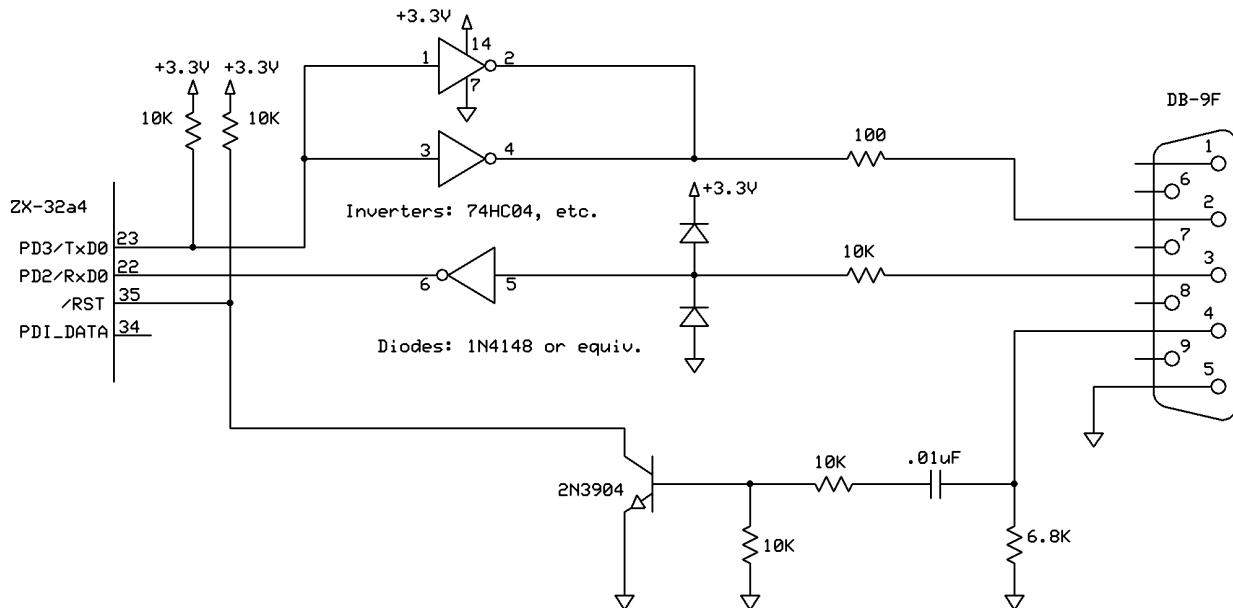
The PDI_DATA pin of the chip (pin 34) is used for factory programming and should be left unconnected.



Recommended ZX-32a4 and ZX-128a4u Serial Interface Circuit

The pullup resistor shown on the ZX's Com1 transmit output (pin 23) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.

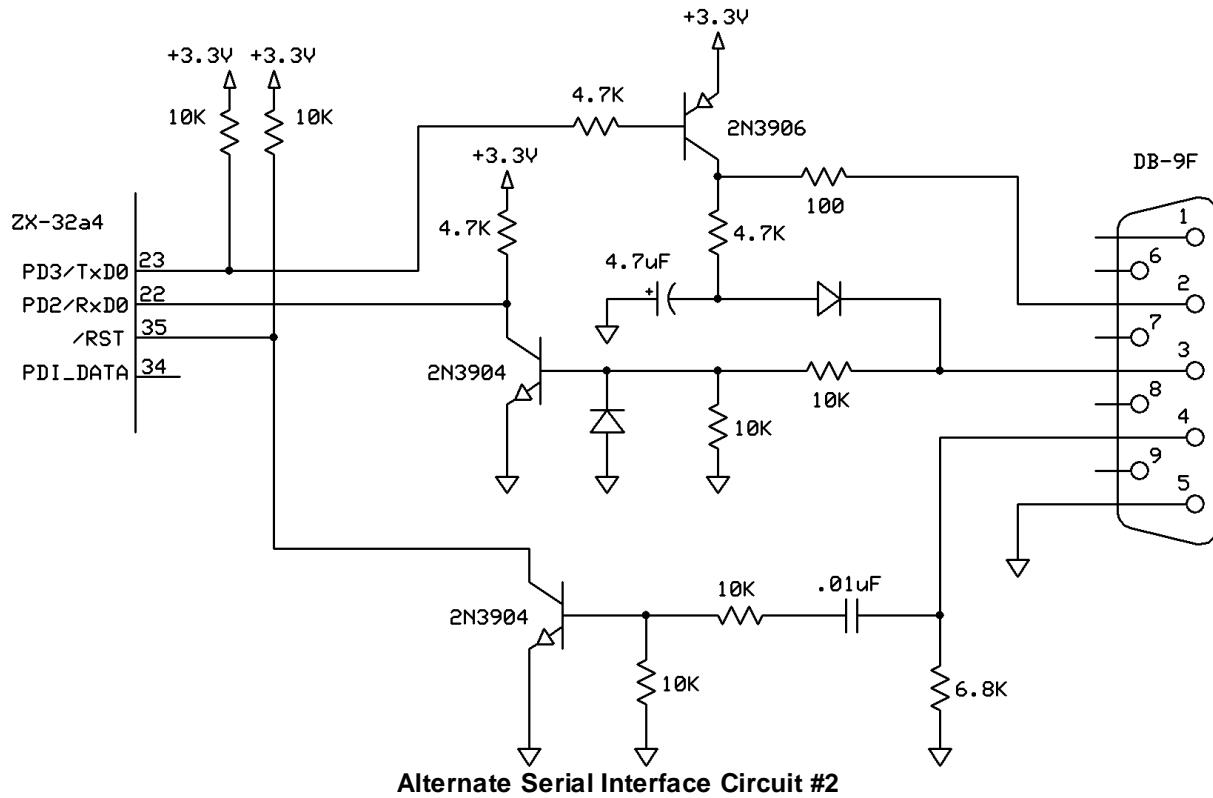


Alternate Serial Interface Circuit #1

The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not

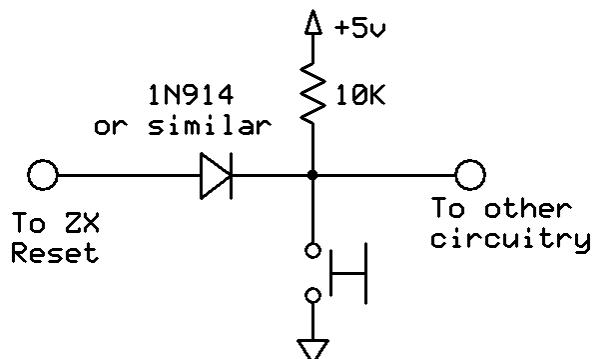
meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 35 of the chip and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix L - ZX-128a1 Hardware Reference

The ZX-128a1 device utilizes the 100-pin TQFP package version of the Atmel AVR ATxmega128A1 microcontroller that has been programmed with a special ZX control program. In order to use the ZX-128a1 you must supply several additional components as described below. In each of the diagrams presented below, only a portion of the ZX pins is shown. Those that are not germane to the circuit being discussed are omitted for clarity.

L.1 ZX-128a1 Specifications

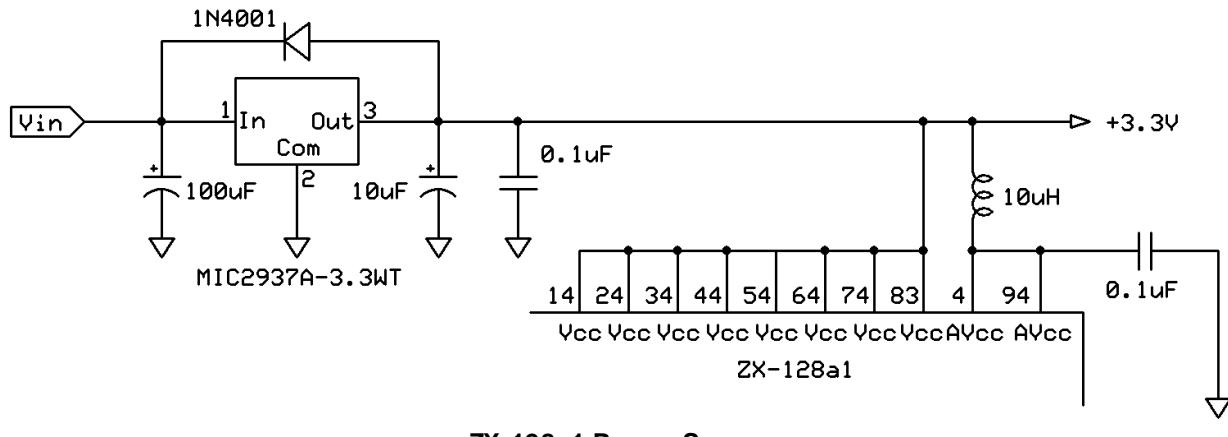
The electrical specifications of the ZX-128a1 are exactly those of the ATxmega128A1. Rather than reproducing them here the reader is directed to the datasheet published by Atmel. It can be obtained from the Atmel website <http://www.atmel.com> or from the ZBasic website <http://www.zbasic.net>.

L.2 ZX-128a1 Required External Components

The circuits described below represent the minimum external circuitry required to operate the ZX-128a1. Depending on your application, you may need additional circuitry to take advantage of the capabilities of the ZX-128a1 series device. In contrast to some other ZX devices, the ZX-128a1 devices do not require an external EEPROM for program storage. Rather, the compiled user program is stored in the ZX's internal Flash memory. The maximum user program size is 128K.

Power Source

The ZX-128a1 needs a regulated voltage source capable of providing at least 200mA of current. The ZX-128a1n will operate reliably with a supply voltage between 2.6 volts and 3.6 volts. A recommended circuit is shown below. A suitable heatsink will probably be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.

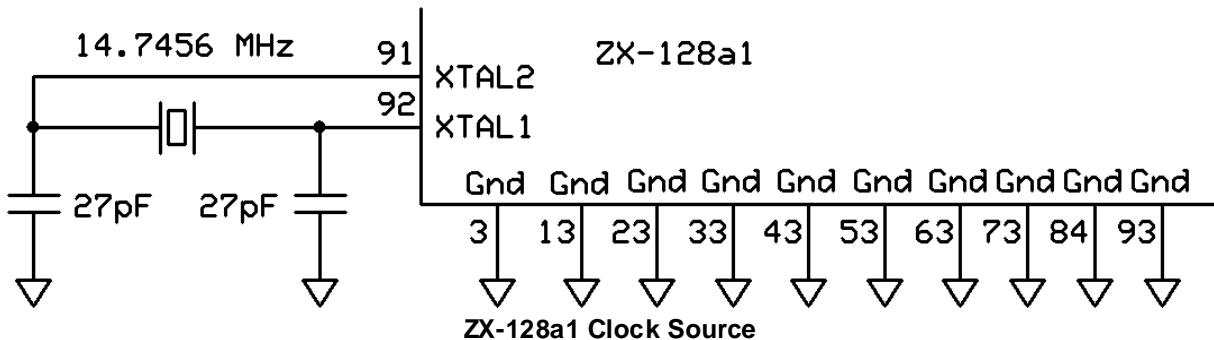


If you do not plan to use the Analog-to-Digital converter channels you can eliminate the inductor and the capacitor on the right side of the diagram. In this case, pins 4 and 94 would be connected directly to the power source (same as pins 4, 24, etc.). Although not shown on this diagram, the ground pins of the ZX-128a1 (pins 3, 13, 23, 33, 43, 53, 63, 73, 84, 93) must be connected to the common ground of the system.

Clock Source

The ZX-128a1n device requires a clock source running at 14.7456MHz. The easiest way to provide this clock source is to connect a crystal of that frequency to pins 91 and 92 along with the necessary capacitors. Some recommended crystals are shown in the table below.

Recommended Crystals	
Package	Crystal
HC-49/US	ECS-147.4-20-4X (Digi-Key #X1102-ND)
HC-49/U	ECS-147.4-20-1X (Digi-Key #X1032-ND)
Cylinder	ECS-147.4-18-9X (Digi-Key #XC-1382-ND)



If you already have a clock source at the required frequency you can feed that signal directly to pin 92. In this case there would be no connection to pin 91.

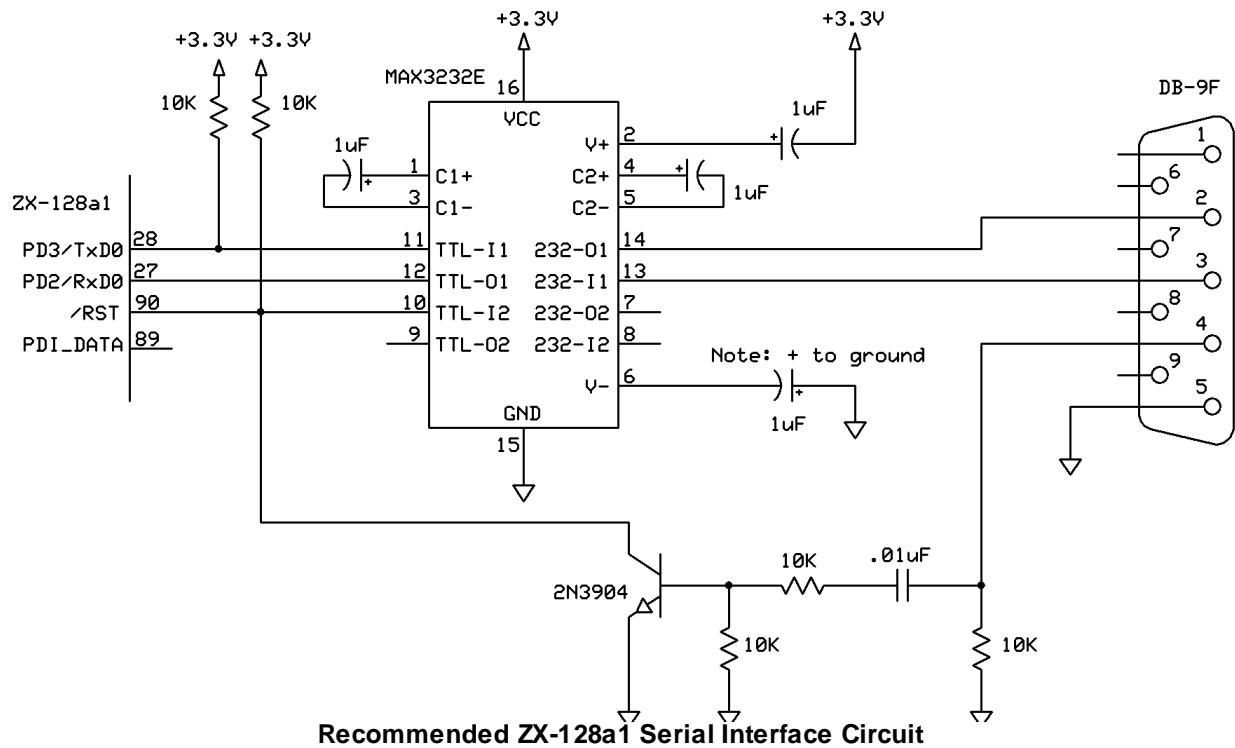
Serial Interface

The ZX-128a1 device requires a serial interface for downloading code into Program Memory and for your program's use via Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX3232E chip functions as an RS-232 level converter that translates the 0-3.3 volt signals of the ZX-128a1 to the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the ZX-128a1 to go into download mode. The latter circuitry is that portion connected between pin 4 of the DB-9 serial connector and pin 90 of the ZX-128a1. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the ZX-128a1 will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the ZX-128a1 receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

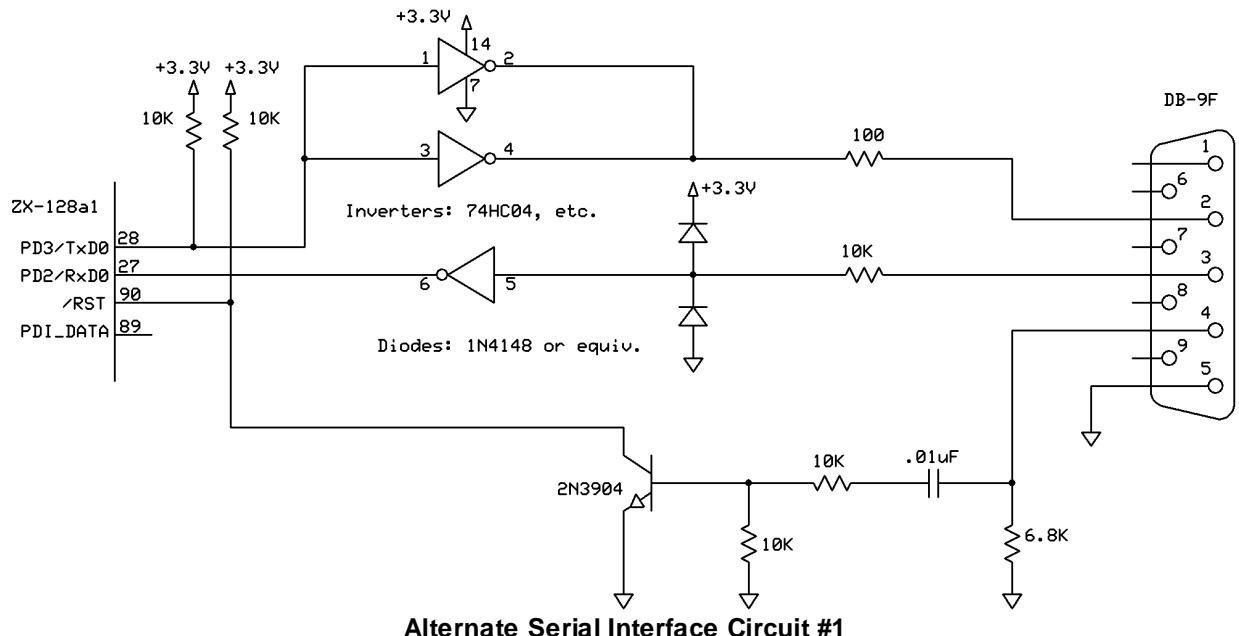
One advantage to using this recommended circuit is that the MAX3232E chip includes a DC-DC converter that produces nominally +/- 10V from a 3.3 volt source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX3232E while the negative voltage is available on pin 6. Consult the datasheet for the MAX3232E device for information on the current capacity of these supplies.

The PDI_DATA pin of the ZX-128a1 (pin 89) is used for factory programming and should be left unconnected.



The pullup resistor shown on the ZX's Com1 transmit output (pin 28) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

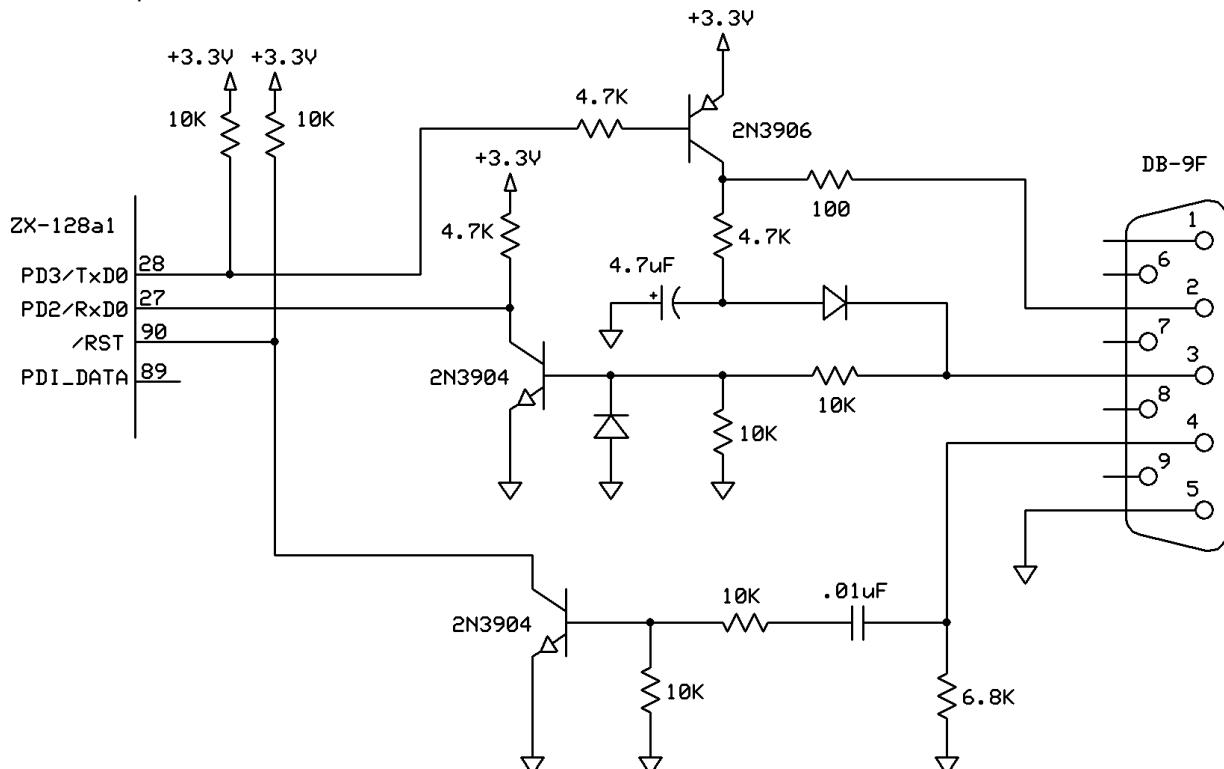
There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.



The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal.

Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

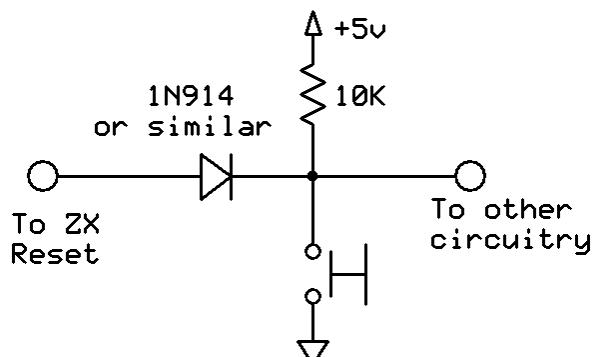
The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Alternate Serial Interface Circuit #2

Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between pin 90 of the ZX-128a1 and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



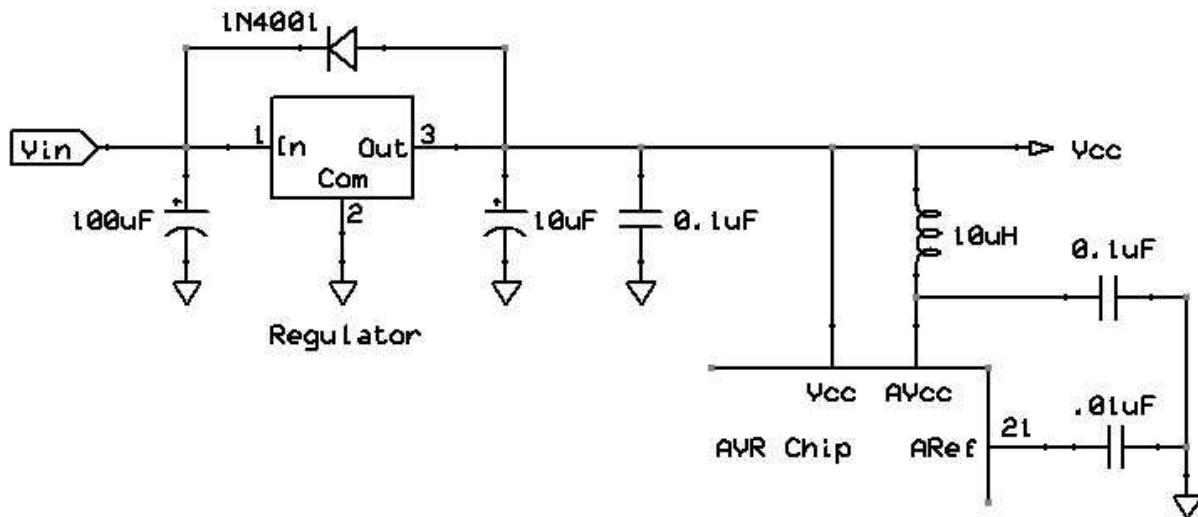
Reset Signal Isolation

Appendix M - External Components for Generic Target Devices

The circuits described below represent the minimum external circuitry required to operate a generic target device. The requirements vary depending on the device, e.g. some generic target devices have on-board serial hardware (i.e. a UART), on-board analog-to-digital-converter, etc. and some don't.

Power Source

The AVR chip needs a regulated voltage source capable of providing current to both the AVR and surrounding circuitry. See the datasheet for the particular target device for the minimum and maximum acceptable operating voltages. A stylized power supply circuit is shown below. Note that a particular AVR devices may not have an AVcc or ARef pin. Generally, a suitable heatsink will be required in most cases to keep the regulator IC below its maximum operating temperature. Consult the regulator datasheet for more information.



Power Supply Circuit

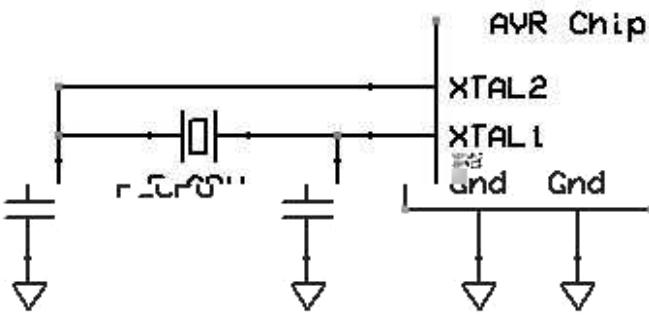
If the target AVR does not have analog circuitry (analog-to-digital converter and/or analog comparator) it will not have an AVcc or Aref pin. Some AVR devices do not have an Aref pin even though they have analog circuitry. If it does have an AVcc pin but you do not plan to use the on-board analog elements you can eliminate the components shown connected to AVcc and Aref. In this case, AVcc would be connected to Vcc and Aref can be left unconnected. Although not shown on this diagram, the AVR chips have one or more ground pins, all of which must be connected to the common ground of the system.

Clock Source

The AVR chip requires a clock source running at a frequency suitable for your application (denoted as F_{CPU} in the generic target parameters, see section 5.3). Most AVR chips have internal oscillators that can be used in some applications. However, the internal oscillator frequency may not be accurate enough or stable enough (over the operating temperature range) to meet the requirements of your application. Generally speaking, the internal oscillator is not accurate enough to support reliable serial communication where the baud rate must be within 2% or less of the desired rate for reliable communication.

If the AVR supports it, the best way to provide the required clock source is to connect a crystal of the desired resonant frequency to the Xtal1 and Xtal2 pins along with the necessary load capacitors (the value of which depends on the crystal specifications) as shown in the circuit diagram below.

Another option is to feed an externally generated clock signal into one of the Xtal pins (check the data sheet for the correct pin).



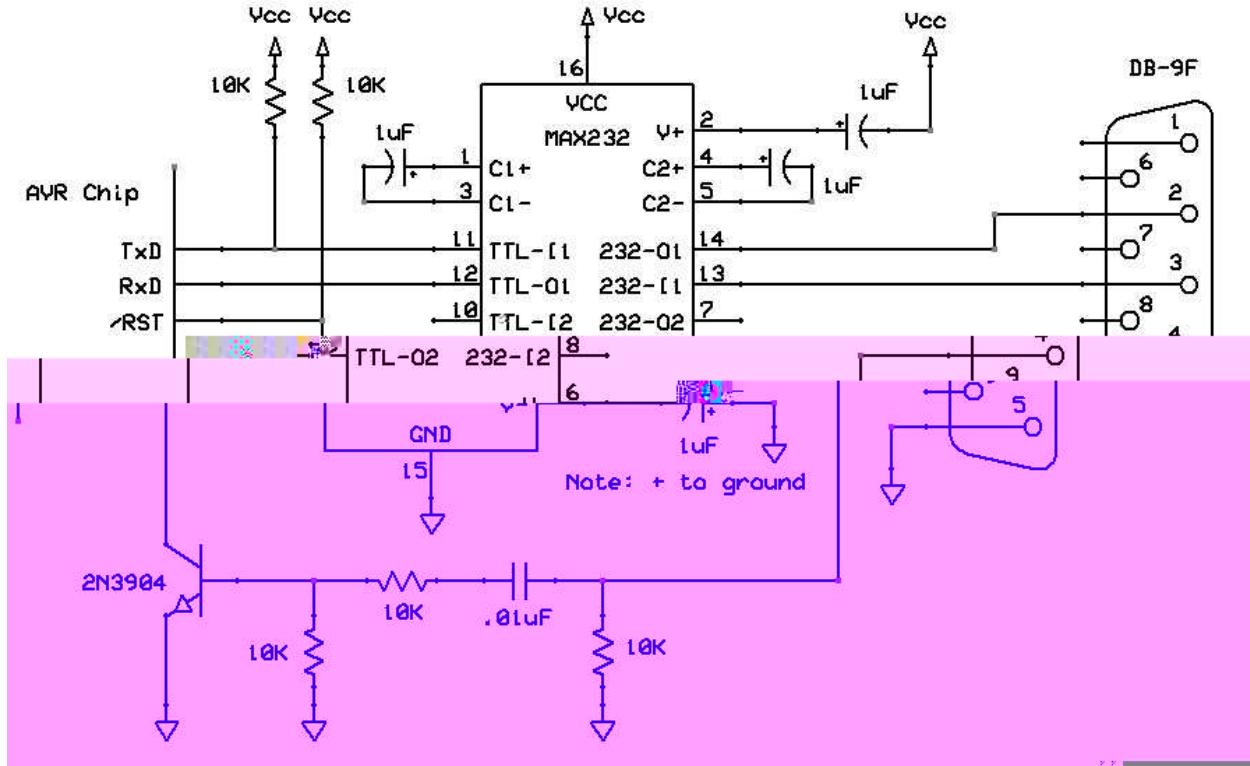
Clock Source

Serial Interface

In many cases, you'll want to provide serial interface circuitry to convert the AVR serial channel signals to/from RS-232 levels. Note that not all AVR chips have on-board serial hardware (UART) but if it does and you elect to install the ZBasic bootloader (or other bootloader) you'll probably need the serial interface circuitry for serial channel Com1. The recommended serial interface circuitry is shown below. This circuit has two sub components. The MAX232 chip functions as an RS-232 level converter that translates the 0-Vcc signals of the AVR to/from the standard RS-232 voltage levels.

The second component is the ATN circuitry that is used to signal the AVR to go into download mode when using the ZBasic bootloader. The ATN circuitry is that portion connected between pin 4 of the DB-9 serial connector and the /RST pin of the AVR. If you do not intend to use the ZBasic bootloader, the ATN circuitry can be omitted. You may want to include a jumper in the ATN circuitry so that it can be disconnected when downloads are not required. With the ATN circuitry connected, the AVR will receive a reset pulse on every positive transition of pin 4 (DTR) of the serial connector. If alternate component values are chosen you must ensure that the AVR receives a reset pulse of at least 2uS on every positive transition of the DTR signal on pin 4 of the serial connector.

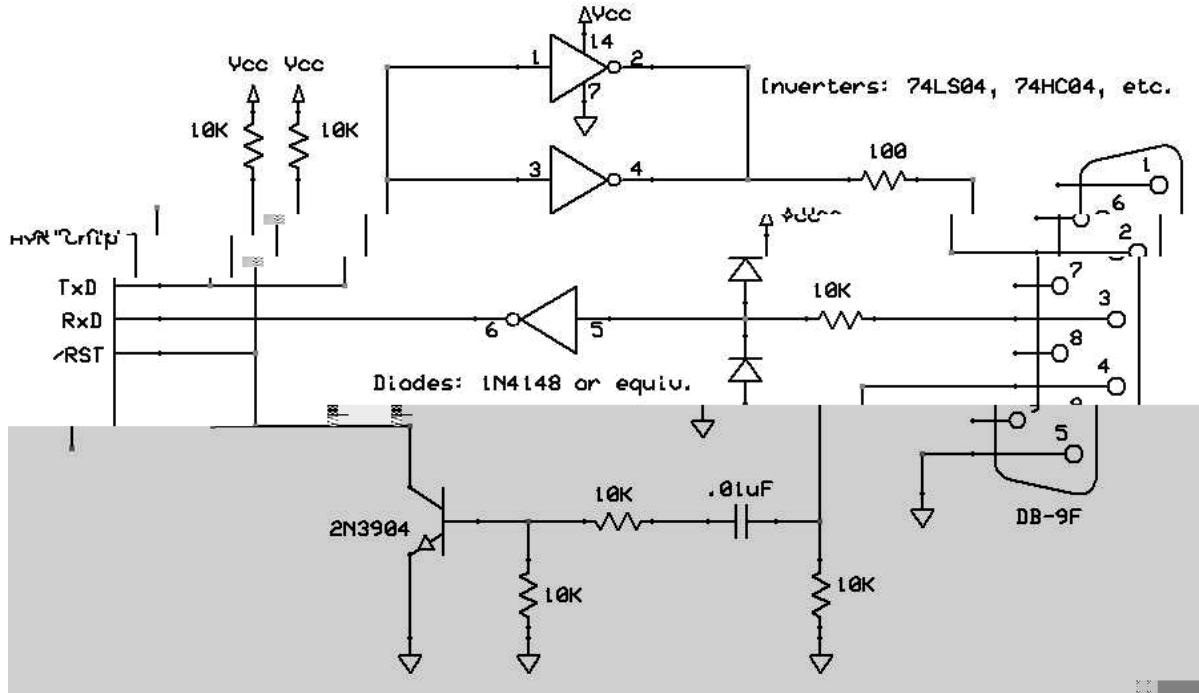
One advantage to using this recommended circuit is that the MAX232 chip includes a DC-DC converter that produces nominally +/- 10V from the positive Vcc source. This is particularly useful if your application requires one or more of these voltages. For example, some LCD devices require a low-current negative supply for their backlight circuitry. The positive voltage is available on pin 2 of the MAX232 while the negative voltage is available on pin 6. Consult the datasheet for the MAX232 device for information on the current capacity of these supplies. Also note that the MAX232 may not operate at the voltage level that you've chosen for your AVR. There are alternate RS-232 level converter chips that operate at different or wider voltage ranges. For example, the MAX3232E can operate from 3.3V to 5V.



Recommended Serial Interface Circuit

The pullup resistor shown on the AVR's transmit output (TxO) is needed to ensure that the serial output line stays in the idle state during reset cycles. Its value is non-critical – anything from 4.7K to 22K should work fine. The pullup resistor on the reset input isn't technically required since the chip itself has an internal pullup. The external pullup could be eliminated or could be made larger if desired.

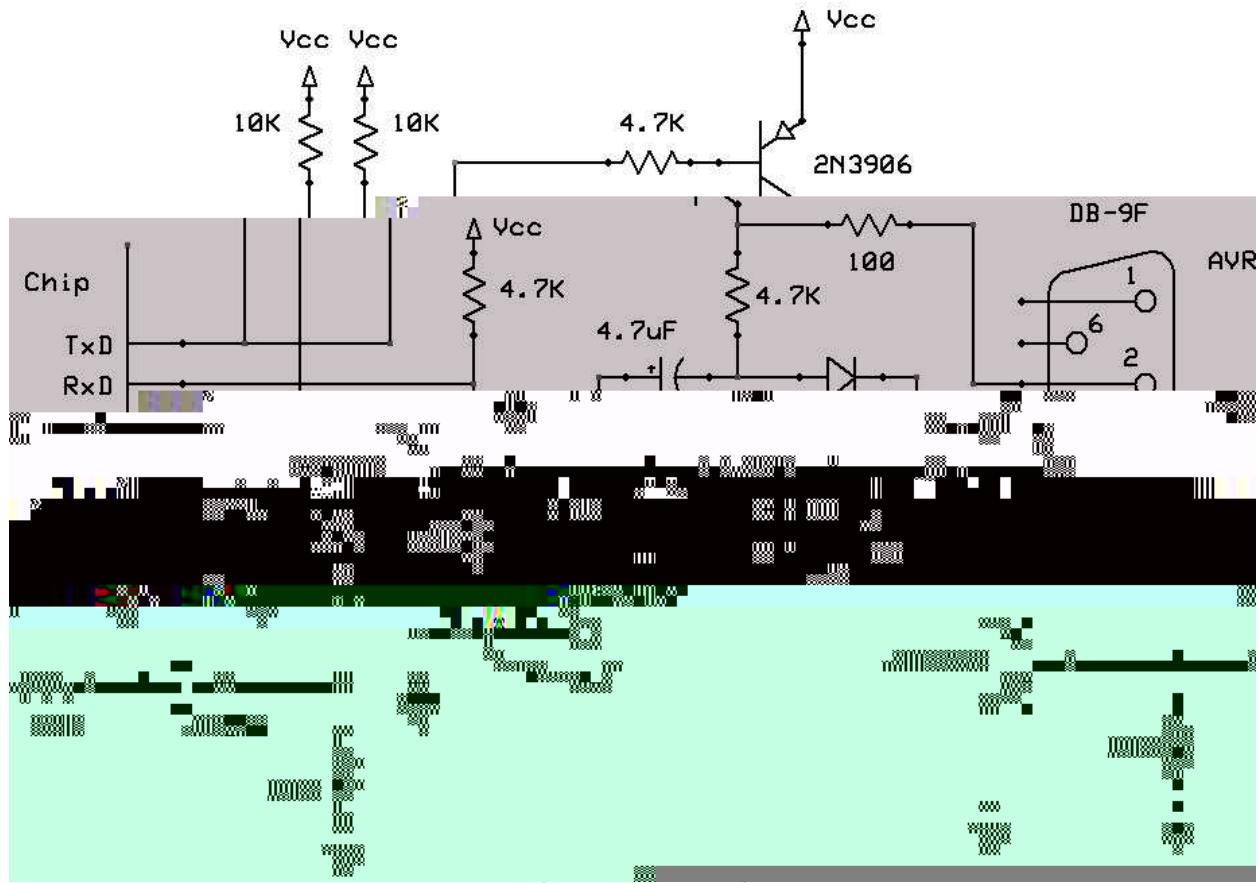
There are other serial interface circuits that would provide the required functionality, two of which are shown below. Both alternate circuits use the same ATN circuitry as the recommended circuit above.



Alternate Serial Interface Circuit #1

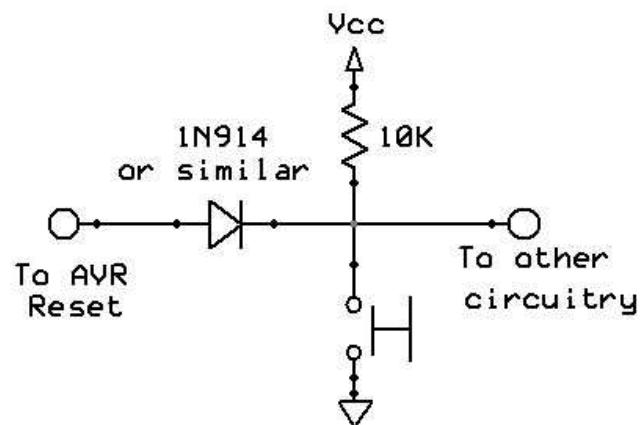
The first alternate above is similar to the circuitry used on the ZX-24 series devices. The advantage to this circuit is its simplicity. The disadvantage is that the voltage levels of the serial output signal do not meet the RS-232 standard. Even so, most serial receivers are able to properly interpret the output signal. Note that if the inverters used have protection diodes on their inputs the external diodes shown are not necessary.

The second alternate serial interface is shown in the diagram below. This circuit produces a serial output signal that meets the RS-232 standard by using the negative voltage on the receive line to provide a reference for the output level. Although the parts count is higher, the circuit uses inexpensive and non-critical components.



Reset Circuit

You may also want to incorporate a reset button in your circuitry. If so, a normally open, momentary contact switch may be connected directly between the /RST pin of the AVR and ground. If other circuitry in your application also needs the reset signal, you may want to isolate that circuitry from the ATN resets (if using the ZBasic bootloader) by using a buffer gate or a diode. A suggested diode isolation circuit for the reset line is shown below.



Reset Signal Isolation

Appendix N - Processor Resource Description Files

The resource description files are XML files that contain information about some of the resources provided by a particular processor, e.g. registers and interrupt vectors. They are used by the compiler to determine the resources that are available for a particular target processor. They are also used by the ZBasic IDE to provide keyword information related to register names. In most cases, a single XML file contains resource information for multiple processors. When necessary, the resource descriptions contain attribute values to identify the processors to which the resource descriptions apply.

The excerpt from an XML resource description file below illustrates how the information for several processors is represented in a single XML file. The first four lines of the XML excerpt define attributes that a particular processor may or may not have. For example, ATTR_TIMER3 indicates the presence of Timer3. Following the attribute definitions (which may or may not be present in a particular XML file), one or more target processors are defined. For example, the target M164 is defined and associated with the name "mega164" and having the attribute ATTR_USART1. Multiple attributes may be given for a target as illustrated with the M1284P target.

Excerpt from an XML Resource Description File

```
<Attribute>ATTR_USART1</Attribute>
<Attribute>ATTR_PRR1</Attribute>
<Attribute>ATTR_TIMER3</Attribute>
<Attribute>ATTR_COMPAT_M32</Attribute>
<Target name="mega164" attr="ATTR_USART1">M164</Target>
<Target name="mega324" attr="ATTR_USART1">M324</Target>
<Target name="mega644" attr="ATTR_COMPAT_M32">M644</Target>
<Target name="mega644P" attr="ATTR_USART1,ATTR_COMPAT_M32">M644P</Target>
<Target name="mega1284P" attr="ATTR_USART1,ATTR_PRR1,ATTR_TIMER3">M1284P</Target>
<Registers>
    <Register8 address="0x0021">DDRA</Register8>
    <Register16 address="0x0086">ICR1</Register16>
    <Register16 address="0x0096" attr="ATTR_TIMER3">ICR3</Register16>
    <Register8 address="0x005b" target="M1284P">RAMPZ</Register8>
</Registers>
<Vectors>
    <Vector index="01" name="INT0_vect">INT0</Vector>
    <Vector index="28" name="USART1_RX_vect" attr="ATTR_USART1">USART1_RX</Vector>
</Vectors>
```

The excerpt further describes several registers using the `<Register8>` tag for 8-bit registers and the `<Register16>` tag for 16-bit registers. The register name given, e.g. DDRA, is the name used with the Register. prefix in ZBasic application code. The first two register descriptors apply to all targets while the third applies only to targets with the ATTR_TIMER3 attribute and the fourth applies only to the M1284P target. Note that a register description may be restricted to multiple targets by specifying multiple attribute values or multiple target identifiers, e.g. `target="M644P,M1284P"`.

The section defining the interrupt vectors is similar to the register description section in the way that attribute values and target names are used to restrict the applicability to certain processors. The name given in the vector descriptions, e.g. INT0, is the name used when defining an interrupt service routine in a ZBasic application.

The table below lists the target devices described in the various XML resource description files. Note that for the mega and tiny series, target devices with A, P and PA suffixes are included along with devices with no suffix.

XML Files for Various Target Devices

XML File	Target Devices
xml/avr/megaX.xml	mega8, mega16, mega32
xml/avr/mega128.xml	mega64, mega128
xml/avr/megaX0.xml	mega640, mega1280, mega2560
xml/avr/megaX1.xml	mega1281, mega2561
xml/avr/megaX3.xml	mega163, mega323
xml/avr/megaX4.xml	mega164, mega324, mega644, mega1284
xml/avr/megaX5.xml	mega165, mega325, mega645, mega3250, mega6450
xml/avr/megaX8.xml	mega48, mega88, mega168, mega328
xml/avr/megaX9.xml	mega169, mega329, mega649, mega3290, mega6490
xml/avr/mega16X.xml	mega161, mega162
xml/avr/mega85x5.xml	mega8515, mega8535
xml/avr/AT90CAN.xml	AT90CAN32, AT90CAN64, AT90CAN128
xml/avr/AT90USB.xml	AT90USB646, AT90USB647, AT90USB1286, AT90USB1287
xml/avr/megaXU2.xml	mega8U2, mega16U2, AT90USB82, AT90USB162
xml/avr/megaXU4.xml	mega16U4, mega32U4
xml/avr/tinyX4.xml	tiny24, tiny44, tiny84
xml/avr/tinyX41.xml	tiny841
xml/avr/tinyX7.xml	tiny87, tiny167
xml/avr/tinyX8.xml	tiny48, tiny88
xml/avr/tinyX13.xml	tiny2313, tiny4313
xml/avr/tiny828.xml	tiny828
xml/avr/tiny1634.xml	tiny1634
xml/avr/xmegaA1.xml	xmega64A1, xmega128A1, xmega192A1, xmega256A1
xml/avr/xmegaA3.xml	xmega64A3, xmega128A3, xmega192A3, xmega256A3
xml/avr/xmegaA3B.xml	xmega256A3B
xml/avr/xmegaA4.xml	xmega16A4, xmega32A4, xmega64A4, xmega128A4
xml/avr/xmegaD3.xml	xmega64D3, xmega128D3, xmega192D3, xmega256D3, xmega384D3
xml/avr/xmegaD4.xml	xmega16D4, xmega32D4, xmega64D4, xmega256D4
xml/avr/xmegaA1U.xml	xmega64A1U, xmega128A1U
xml/avr/xmegaA3U.xml	xmega64A3U, xmega128A3U, xmega192A3U, xmega256A3U
xml/avr/xmegaA3BU.xml	xmega256A3BU
xml/avr/xmegaA4U.xml	xmega16A4U, xmega32A4U, xmega64A4U, xmega128A4U