

Advanced Programming in Visual Basic .NET™

Rex Jaeschke

Advanced Programming in Visual Basic .NET

© 2002, 2009 Rex Jaeschke. All rights reserved.

Edition: 2.0

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means whatsoever, except in the case of brief quotations embodied in critical reviews and articles.

The information in this book is subject to change without notice and should not be construed as a commitment by the author or the publisher. Although every precaution has been taken in the preparation of this book, the author and the publisher assume no responsibility for errors or omissions.

Java and JavaScript are trademarks of Sun Microsystems.

.NET, Visual Basic, Visual C#, Visual C++, Visual Studio, and JScript are trademarks of Microsoft.

The training materials associated with this book are available for license. Interested parties should contact the author at the address below.

Please address comments, corrections, and questions to the author:

Rex Jaeschke
2051 Swans Neck Way
Reston, VA 20191-4023
+1 (703) 860-0091
+1 (703) 860-3008 (fax)
www.RexJaeschke.com
rex@RexJaeschke.com

Preface	v
Reader Assumptions	v
Presentation Style	v
Exercises and Solutions	vi
The Status of Visual Basic .NET	vi
1. Threads.....	1
1.1 Introduction	1
1.2 Creating Threads.....	2
1.3 Synchronized Statements	5
1.4 Other Forms of Synchronization.....	10
1.5 Managing Threads	14
1.6 Thread-Local Storage	15
1.7 Atomicity and Interlocked Operations	17
2. Object Serialization	21
2.1 Introduction	21
2.2 Serializing Objects that Contain References.....	24
2.3 Handling Multiple References	26
2.4 Customized Serialization.....	29
2.5 Identifying the Fields to be Serialized.....	33
2.6 Serialization Format.....	36
2.7 Type Packaging	38
3. Sockets	41
3.1 Introduction	41
3.2 Server-Side Sockets.....	41
3.3 Client-Side Sockets.....	45
3.4 Serialization over Sockets	48
4. Cloning Objects	53
4.1 Copying by Constructor	53
4.2 Class Cloning	54
4.3 The Procedure Clone.....	55
4.4 Using Object.MemberwiseClone	56
4.5 Cloning Arrays.....	59
4.6 Cloning and Derived Classes	60
4.7 Creation without Construction	62
5. Attributes	65
5.1 Introduction	65
5.2 Predefined .NET Attributes.....	66
5.3 StructLayout and FieldOffset.....	70
5.4 DllImport	72
5.5 CLSCompliant.....	75
5.6 Obsolete.....	76
5.7 Custom Attributes.....	77
Annex A. Operator Precedence	79
Annex B. Visual Basic .NET Keywords	83
Index	85

Preface

This text covers a number of more advanced Visual Basic .NET topics. The material is not hardware or operating system-specific.

Reader Assumptions

To fully understand and exploit the material, you should be conversant with the following concepts and the syntax required to express them in Visual Basic .NET:

- Basic Language Elements
- Looping and Testing
- Procedures
- References
- Strings and Arrays
- Classes
- Inheritance
- Interfaces
- Exception Handling
- Input and Output
- Namespaces

Presentation Style

The approach used in this book is different from that used in many other books and training courses. Having developed and delivered programming-language training courses for more than 15 years, I have found that the best approach for my students is an incremental one. I avoid introducing things that are unnecessary at any given time, thus making examples small, simple, and well focused. Specifically, I introduce the basic elements and constructs of the language using procedural programming examples. Once those fundamentals have been mastered, I move on to object-oriented concepts and syntax. Then follow the more advanced language and library topics. Many books on object-oriented languages use objects, inheritance, exception handling, GUI, and numerous non-trivial library facilities in the first few examples, and certainly in the first chapter. I do not care for this approach, either as a reader or as an educator. Instead, I prefer the student to have an excellent chance of understanding and absorbing small amounts of new material, and reinforcing it with lab sessions, as they progress. The intent here is to eliminate any chance of their being overwhelmed, provided, of course, they meet the prerequisites.

Different styles work for different teachers and different students. I do not suggest that my approach is better than is any other; I simply know that my approach works well, and has formed the basis of my successful seminar business for more than a decade.

Exercises and Solutions

The programs shown in the text are available electronically in a directory tree named **Source**, where each chapter has its own subdirectory. By convention, the names of Visual Basic .NET source files end in “.vb”.

Each chapter contains exercises, some of which have the character * following their number. For each exercise so marked, a solution is provided in a directory tree named **Labs**, in which each chapter has its own subdirectory.¹ Exercises that are not so marked have no general solution and require experimentation or research in an implementation's documentation. Exercises having solutions contain a statement of the form “(See lab directory xx.)”, which indicates the corresponding solution or test file in the **Labs** subdirectory.

You are strongly encouraged to solve all exercises in one section before continuing to the next. Also, invent your own exercises as you go and be inquisitive; don't be afraid to experiment. Try to understand why the compiler gives you each error or why a program fails at run time.

The Status of Visual Basic .NET

Microsoft announced the Visual Basic .NET language and .NET platform in July 2000. At that time, Microsoft stated that some subset of the library and runtime environment would be submitted for standardization, to Ecma International, a standards organization. In September 2000, Ecma committee TC39, which was responsible for standardization of the scripting language ECMAScript (known commercially as JavaScript and JScript), agreed to take on this new work. As a result, TC39 was split into three Task Groups: TG1 (ECMAScript), TG2 (C#), and TG3 (Common Library Infrastructure [CLI]). This author served as project editor of the standards produced by TG2 and TG3.

The submission to Ecma was sponsored by Hewlett-Packard, Intel, and Microsoft. A number of other companies also agreed to participate in the standards work, which began in November 2000. That work was completed in September 2001, and the CLI (and C#) standard was adopted by Ecma in December 2001. The standard for CLI is ECMA-335. (This standard can be downloaded free of charge from www.ecma-international.org.)

In January 2002, a 6-month review and ballot period began within ISO/IEC JTC 1 to determine if this specification should be adopted as an ISO/IEC standard as well. After a ballot-resolution meeting in September 2002, this specification was approved unanimously, and it was forwarded to ISO for publication. The designation of the CLI standard is ISO/IEC 23271. (See www.iso.org.)

A revision of the Ecma version of this standard started in January 2003, and was completed in March 2005. Ecma adopted it in June 2005, and ISO/IEC JTC 1 adopted it in April 2006.

In 2008, Ecma transferred the Task Groups TG2 (C#) and TG3 (CLI) from Technical Committee TC39 to TC49.

The Visual Basic .NET language itself has not been standardized.

Rex Jaeschke, September 2009

¹ The solutions are only available to licensees of these materials when they are used in formal training scenarios.

1. Threads

Visual Basic .NET supports the ability to create multiple threads of execution within a single application. In this chapter, we'll see how threads are created and synchronized.¹ We'll also see how shared variables can be guarded against compromise during concurrent operations.

1.1 Introduction

A *thread* is an individual stream of execution as seen by the processor, and each thread has its own register and stack context. The run-time environment executes only one thread at a time. The execution of a thread is interrupted when it needs resources that are not available, it is waiting for an operation such as an I/O to complete, or if it uses up its processor time slice. When the processor changes from executing one thread to another, this is called *context switching*. By executing another thread when one thread becomes blocked, the system allows processor idle time to be reduced. This is called *multitasking*.

When an application is executed, the system is told where on disk to get instructions and shared data. A set of virtual memory locations, collectively called an address space is allocated to that application, as are various system resources. This runtime context is called a *process*. However, before a process can do any work, it must have at least one thread. When each process is created, it is automatically given one thread, called the *primary thread*. However, this thread has no more capability than other threads created for that process; it just happened to be the first thread created for that process. The number of threads in a process can vary at runtime, under application control. Any thread can create other threads; however, a creating thread does not in any sense own the threads it creates; all threads in a process belong to the process as a whole.

The work done by a process can be broken into subtasks with each being executed by a different thread. This is called *multithreading*. Each thread in a process shares the same address space and process resources. When the last remaining thread in a process terminates, the parent process terminates.

Why have more than one thread in a process? If a process has only one thread, it executes serially. When the thread is blocked, the system is idle if no other process has an active thread waiting. This may be unavoidable if the subtasks of the process must be performed serially; however, this is not the case with many processes. Consider a process that has multiple options. A user selects some option, which results in lots of computations using data in memory or a file and the generation of a report. By spawning off a new thread to perform this work, a process can continue accepting new requests for work without waiting for the previous option to complete. And by specifying thread priorities, a process can allow less-critical threads to run only when more-critical threads are blocked.

Once a thread has been dispatched, another thread can be used to service keyboard or mouse input. For example, the user might decide that a previous request is not the way to go after all, and wishes to abort the first thread. This can be done by selecting the appropriate option on a pull-down menu and having one thread stop the other.

¹ It is important to note that Visual Basic .NET does not support synchronization of threads in different applications.

Another example involves a print spooler. Its job is to keep a printer busy as much as possible and to service print requests from users. The users would be very unhappy if the spooler waits until a job had completed printing before it started accepting new requests. Of course, it could periodically stop printing to see if any new requests were pending (this is called *polling*), but that wastes time if there are no requests. And if the time interval between polls is too long, there is a delay in servicing requests. If it is too short, the thread spends too much time polling. Why not have the spooler have two threads: one to send work to the printer, the other to deal with requests from users. Each runs independent of the other and when a thread runs out of work, it either terminates itself or goes into an efficient state of hibernation.

When dealing with concurrently executing threads, we must understand two important aspects: atomicity and reentrancy.

An *atomic* variable or object is one that can be accessed as a whole even in the presence of asynchronous operations that access the same variable or object. For example, if one thread is updating an atomic variable or object while another thread reads its contents, the logical integrity of those contents cannot be compromised—the read will get either the old or the new value, never part of each. Normally, the only things that can be accessed atomically are those having types supported atomically in hardware, such as bytes and words. All the primitive types in Visual Basic .NET except `Date`, `Decimal`, `Double` and `Long` are guaranteed to be atomic. Clearly, a `Point` object is not atomic; it has two parts, an x- and a y-coordinate, and a writer of a `Point`'s value could be interrupted by a reader to that `Point`, resulting in the reader getting the new x and old y, or vice-versa. Similarly, arrays cannot be accessed atomically. Since most objects cannot be accessed atomically, we must use some form of synchronization to ensure that only one thread at a time can operate on certain objects. For this reason, Visual Basic .NET assigns each object, array, and class a synchronization lock.

A *reentrant* procedure is one that can safely be executed in parallel by multiple threads of execution. When a thread begins executing a procedure, all data allocated in that procedure comes either from the stack or the heap. In any event, it's unique to that invocation. If another thread begins executing that same procedure while the first thread is still working there, each thread's data will be kept separate. However, if that procedure accesses variables or files that are shared between threads, it must use some form of synchronization.

1.2 Creating Threads

In the following example (see directory `Th01`), the primary thread creates two other threads, and the three threads run in parallel without synchronization. No data is shared between the threads and the process terminates when the last thread terminates:

```
Imports System.Threading

Public Class Th01
    Private loopStart As Integer
    Private loopEnd As Integer
    Private dispFrequency As Integer
```



```

Public Sub New(ByVal startValue As Integer, ByVal endValue As Integer, _
    ByVal frequency As Integer)
    loopStart = startValue
    loopEnd = endValue
    dispFrequency = frequency
End Sub
'1
Public Sub ThreadEntryPoint()
'2
    Dim threadName As String = Thread.CurrentThread.Name
    Dim i As Integer = loopStart
    While i <= loopEnd
        If i Mod dispFrequency = 0 Then
            Console.WriteLine("{0}: i = {1,10}", threadName, i)
        End If
        i += 1
    End While
    Console.WriteLine("{0} thread terminating", threadName)
End Sub]
Public Shared Sub Main()
'3a
    Dim o1 As New Th01(0, 1000000, 200000)
'3b
    Dim t1 As New Thread(AddressOf o1.ThreadEntryPoint)
'3c
    t1.Name = "t1"
'4a
    Dim o2 As New Th01(-1000000, 0, 200000)
'4b
    Dim t2 As New Thread(AddressOf o2.ThreadEntryPoint)
'4c
    t2.Name = "t2"
'5
    t1.Start()
'6
    t2.Start()
    Console.WriteLine("Primary thread terminating")
End Sub
End Class

```

Let's begin by looking at the first executable statement in the program, that in case 3a. Here we create an object having the user-defined type `Th01`. That class has a constructor, an instance procedure, and three fields. We call the constructor passing it a start and end count, and an increment amount, which it stores for later use in controlling a loop.

2. Object Serialization

Most useful applications depend on information of a more permanent nature than that generated during a single execution. For example, applications that access an inventory, typically query (and possibly update) one of more related data files. The lives of such “master files” transcend that of the execution of any of the applications that use them. Other applications involve the communication of messages between separate applications, often referred to as *client* and *server*. While the life of a message is often much shorter than that of a database record, both cases involve the use of some data format external to the applications that manipulate them.

In this chapter, we'll see how Visual Basic .NET objects and variables of simple types can be converted into some external form suitable for use in file storage or for transmission during inter-application communication. The process of converting to some external form is known as *serialization* while that of converting back again is known as *deserialization*.

2.1 Introduction

Consider the following example (see directory Sr01) which writes a number of values of a variety of object and simple types to a disk file, closes that file, and then reads those values back into memory again:

```
Imports System.IO
Imports System.Runtime.Serialization.Formatters.Binary
Public Module Sr01
    Public Sub Main()
        Dim intArray() As Integer = {10, 20, 30}
        Dim floatArray(,) As Single = { _
            {1.2F, 2.4F}, _
            {3.5F, 6.8F}, _
            {8.4F, 9.7F} _
        }
        Dim dt As DateTime = DateTime.Now
        Console.WriteLine("dt      >{0}<", dt)
'1
        Dim formatter As New BinaryFormatter()
```

In case 1, we define a variable of type `BinaryFormatter`. Objects of this type allow the serialization and deserialization of any object or an entire graph¹ of connected objects in some binary format. (We'll see in §2.6 how alternate formats can be used.)

¹ The term *graph* as used here refers to a kind of nonlinear data structure. For example, an `Employee` type might contain a name field, which, in turn, is made up of a first name, an initial, and a last name. The type might also contain several date objects, representing date of birth, date of hire, and date of last review. The graph designated by a reference to an instance of such a type involves all of that instance's dependent data tree starting at the root reference.

```

    ' Serialize data to a f.

'2
    Dim f As Stream = File.Open("Sr01.ser", FileMode.Create)

'3a
    formatter.Serialize(f, "Hello")
'3b
    formatter.Serialize(f, intArray)
'3c
    formatter.Serialize(f, floatArray)
'3d
    formatter.Serialize(f, True)
'3e
    formatter.Serialize(f, dt)
'3f
    formatter.Serialize(f, 1000)
'3g
    formatter.Serialize(f, "X"c)
'3h
    formatter.Serialize(f, 1.23456F)
'4
    f.Close()

```

In case 2, we create a new file having the name shown. The suffix “.ser” has no special meaning; it is simply a local convention that signifies a serialized data file. Cases 3a through 3h each result in an object being serialized to that file. In the case of the string, each character is written. In the case of the arrays, all elements are written. In the case of the DateTime, all data contained within that type and any associated dependencies are written. In the case of the simple type values, they are first boxed and the corresponding Objects are written. So `Serialize` need only be defined to accept an argument of type `Object`.

```

    ' Deserialize data from a f.

'5
    f = File.Open("Sr01.ser", FileMode.Open)

'6a
    Dim s As String = CStr(formatter.Deserialize(f))
    Console.WriteLine("String >{0}<", s)

```

We retrieve the serialized data by calling the procedure `Deserialize`, as shown in case 6a. Since that procedure returns a value of type `Object`, we need to convert that to the appropriate type.

```

'6b
    Dim newArray() As Integer = CType(formatter.Deserialize(f), _
        Integer())
    Console.WriteLine("newIntArray:")
    Dim i As Integer = 0
    While i < newArray.Length
        Console.Write(" {0}", newArray(i))
        i += 1
    End While

    Console.WriteLine()
'6c
    Dim newArray(,) As Single = _
        CType(formatter.Deserialize(f), Single(,))
    Console.WriteLine("newSingleArray:")
    i = 0
    While i < 3
        Dim j As Integer = 0
        While j < 2
            Console.Write(" {0}", newArray(i, j))
            j += 1
        End While
        Console.WriteLine()
        i += 1
    End While
'6d
    Dim b As Boolean = CBool(formatter.Deserialize(f))
    Console.WriteLine("Boolean >{0}<", b)
'6e
    Dim newDT As DateTime = CDate(formatter.Deserialize(f))
    Console.WriteLine("newDT >{0}<", newDT)
'6f
    Dim v As Integer = CInt(formatter.Deserialize(f))
    Console.WriteLine("Integer >{0}<", v)
'6g
    Dim c As Char = CChar(formatter.Deserialize(f))
    Console.WriteLine("Char >{0}<", c)
'6h
    Dim sn As Single = CSng(formatter.Deserialize(f))
    Console.WriteLine("Single >{0}<", sn)
'7
    f.Close()
End Sub
End Module

```

When executed on the date and time shown, the output produced is:

4. Cloning Objects

Unlike most other languages, Visual Basic .NET provides no way to have an expression of some object type; the best we can do is to have an expression whose value is really a reference to that type. As a result, Visual Basic .NET provides no linguistic way to copy an object. Assuming we can justify it, how then can we make a complete copy of an object?

4.1 Copying by Constructor

Consider the following code fragment (see directory Cn01):

```
Public Class Point
    ' ...
    Public Sub New(ByVal xorigin As Integer, ByVal yorigin As Integer)
        X = xorigin
        Y = yorigin
    End Sub

    Public Sub New(ByVal p As Point)
        X = p.X
        Y = p.Y
    End Sub
    ' ...
End Class
```

The constructor creates a new Point using the contents of an existing one. (In C++, this kind of constructor is so special it has the name *copy constructor*; however, this term is not used in the context of Visual Basic .NET.)

```
Public Module Cn01
    Public Sub Main()
'1
        Dim p1 As New Point(3, 5)
'2
        Console.WriteLine("p1: {0}", p1)
'3
        Dim p2 As New Point(p1)
'4
        p1.Move(9, 11)
'5
        Console.WriteLine("p1: {0}", p1)
'6
        Console.WriteLine("p2: {0}", p2)
    End Sub
End Module
```

Point p2 is a copy of p1, so when the coordinate values of p1 are changed in case 4, we see from the following output that those of p2 are not:

```
p1: (3,5)
p1: (9,11)
p2: (3,5)
```

While this approach works fine, we won't find standard library classes having this form of constructor, however.

4.2 Class Cloning

Those few standard library classes that do support some sort of copy mechanism achieve it by what is referred to as *cloning*. Consider the following example (see directory Cn02) that uses the library class `ArrayList`, a type that supports vectors:

```
Imports System.Collections

Public Module Cn02
    Public Sub Main()
        Dim al1 As New ArrayList()
'1
        al1.Add("Red")
        al1.Add("Blue")
        al1.Add("Green")
        al1.Add("Yellow")
'2
        PrintEntries("al1", al1)
'3
        Dim al2 As ArrayList = CType(al1.Clone(), ArrayList)
'4
        PrintEntries("al2", al2)
'5
        al1.Remove("Blue")
        al1.Add("Black")
        al1.RemoveAt(0)
        al1.Insert(0, "Brown")
'6
        PrintEntries("al1", al1)
'7
        PrintEntries("al2", al2)
    End Sub
```

```

Private Sub PrintEntries(ByVal s As String, ByVal aList As ArrayList)
    Console.Write("{0: ", s)
    Dim o As Object

    For Each o In aList
        Console.Write(vbTab + "{0}", o)
    Next
    Console.WriteLine()
End Sub
End Module

```

The output produced is:

```

a1:    Red    Blue   Green  Yellow
a2:    Red    Blue   Green  Yellow
a1:    Brown  Green  Yellow  Black
a2:    Red    Blue   Green  Yellow

```

ArrayList `a1` consists of a set of 4 strings specifying different colors. We then make a complete copy of that object in case 3 by calling the procedure `ArrayList.Clone`, so the output produced by cases 2 and 4 is the same. Then we modify `a1` by removing the second element, by adding a new element to the end, and by changing the first element's value. When we compare the output produced by cases 6 and 7, we see that the changes applied to `a1` have no affect on `a2`. Specifically, the internal reference inside of `a2` points to its own private copy of elements, not to the same set as `a1`. This is referred to as a *deep copy*, whereas simply making both ArrayLists' internal reference point to the same set of values (as would be the case with the assignment `a2 = a1`) is called a *shallow copy*.

If we wish to make copies of our own objects, we would do well to model the copy process on the library cloning machinery.

4.3 The Procedure Clone

The key to cloning is to implement the standard interface `ICloneable`, which, in turn, requires us to define a procedure called `Clone` that takes no arguments and has a return type of `Object`. For example, consider the following program fragment (see directory `Cn03`):

```

Public Class Point
    Implements ICloneable

    '...
    Public Sub New(ByVal xorigin As Integer, ByVal yorigin As Integer)
        X = xorigin
        Y = yorigin
    End Sub

```