7

Unicode®

Objectives

- To become familiar with Unicode.
- To discuss the mission of the Unicode Consortium.
- To discuss the design basis of Unicode.
- To understand the three Unicode encoding forms: UTF-8, UTF-16 and UTF-32.
- To introduce characters and glyphs.
- To discuss the advantages and disadvantages of using Unicode.
- To provide a brief tour of the Unicode Consortium's Web site.

Outline

F.1 Introduction

The use of inconsistent character *encodings* (i.e., numeric values associated with characters) in the developing of global software products causes serious problems, because computers process information as numbers. For instance, the character "a" is converted to a numeric value so that a computer can manipulate that piece of data. Many countries and corporations have developed their own encoding systems that are incompatible with the encoding systems of other countries and corporations. For example, the Microsoft Windows operating system assigns the value 0xC0 to the character "A with a grave accent"; the Apple Macintosh operating system assigns that same value to an upside-down question mark. This results in the misrepresentation and possible corruption of data when data is not processed as intended.

In the absence of a widely-implemented universal character-encoding standard, global software developers had to *localize* their products extensively before distribution. Localization includes the language translation and cultural adaptation of content. The process of localization usually includes significant modifications to the source code (such as the conversion of numeric values and the underlying assumptions made by programmers), which results in increased costs and delays releasing the software. For example, some English-speaking programmers might design global software products assuming that a single character can be represented by one byte. However, when those products are localized for Asian markets, the programmer's assumptions are no longer valid; thus, the majority, if not the entirety, of the code needs to be rewritten. Localization is necessary with each release of a version. By the time a software product is localized for a particular market, a newer version, which needs to be localized as well, may be ready for distribution. As a result, it is cumbersome and costly to produce and distribute global software products in a market where there is no universal character-encoding standard.

In response to this situation, the *Unicode Standard*, an encoding standard that facilitates the production and distribution of software, was created. The Unicode Standard outlines a specification to produce consistent encoding of the world's characters and *symbols*. Software products that handle text encoded in the Unicode Standard need to be localized, but the localization process is simpler and more efficient because the numeric values need not be converted and the assumptions made by programmers about the character encoding are universal. The Unicode Standard is maintained by a nonprofit organization called the

Unicode Consortium, whose members include Apple, IBM, Microsoft, Oracle, Sun Microsystems, Sybase and many others.

When the Consortium envisioned and developed the Unicode Standard, they wanted an encoding system that was *universal*, *efficient*, *uniform* and *unambiguous*. A universal encoding system encompasses all commonly used characters. An efficient encoding system allows text files to be parsed easily. A uniform encoding system assigns fixed values to all characters. An unambiguous encoding system represents a given character in a consistent manner. These four terms are referred to as the Unicode Standard *design basis*.

F.2 Unicode Transformation Formats

Although Unicode incorporates the limited ASCII *character set* (i.e., a collection of characters), it encompasses a more comprehensive character set. In ASCII each character is represented by a byte containing 0s and 1s. One byte is capable of storing the binary numbers from 0 to 255. Each character is assigned a number between 0 and 255; thus, ASCII-based systems can support only 256 characters, a tiny fraction of world's characters. Unicode extends the ASCII character set by encoding the vast majority of the world's characters. The Unicode Standard encodes all of those characters in a uniform numerical space from 0 to 10FFFF hexadecimal. An implementation will express these numbers in one of several transformation formats, choosing the one that best fits the particular application at hand.

Three such formats are in use, called *UTF-8*, *UTF-16* and *UTF-32*, depending on the size of the units—in *bits*—being used. UTF-8, a variable-width encoding form, requires one to four bytes to express each Unicode character. UTF-8 data consists of 8-bit bytes (sequences of one, two, three or four bytes depending on the character being encoded) and is well suited for ASCII-based systems, where there is a predominance of one-byte characters (ASCII represents characters as one byte). Currently, UTF-8 is widely implemented in UNIX systems and in databases.

The variable-width UTF-16 encoding form expresses Unicode characters in units of 16 bits (i.e., as two adjacent bytes, or a short integer in many machines). Most characters of Unicode are expressed in a single 16-bit unit. However, characters with values above FFFF hexadecimal are expressed with an ordered pair of 16-bit units called *surrogates*. Surrogates are 16-bit integers in the range D800 through DFFF, which are used solely for the purpose of "escaping" into higher numbered characters. Approximately one million characters can be expressed in this manner. Although a surrogate pair requires 32 bits to represent characters, it is space-efficient to use these 16-bit units. Surrogates are rare characters in current implementations. Many string-handling implementations are written in terms of UTF-16. [*Note*: Details and sample code for UTF-16 handling are available on the Unicode Consortium Web site at **`www.unicode.org`**.]

Implementations that require significant use of rare characters or entire scripts encoded above FFFF hexadecimal should use UTF-32, a 32-bit, fixed-width encoding form that usually requires twice as much memory as UTF-16 encoded characters. The major advantage of the fixed-width UTF-32 encoding form is that it expresses all characters uniformly, so it is easy to handle in arrays.

There are few guidelines that state when to use a particular encoding form. The best encoding form to use depends on computer systems and business protocols, not on the data itself. Typically, the UTF-8 encoding form should be used where computer systems and

business protocols require data to be handled in 8-bit units, particularly in legacy systems being upgraded, because it often simplifies changes to existing programs. For this reason, UTF-8 has become the encoding form of choice on the Internet. Likewise, UTF-16 is the encoding form of choice on Microsoft Windows applications. UTF-32 is likely to become more widely used in the future as more characters are encoded with values above FFFF hexadecimal. Also, UTF-32 requires less sophisticated handling than UTF-16 in the presence of surrogate pairs. Figure F.1 shows the different ways in which the three encoding forms handle character encoding.

F.3 Characters and Glyphs

The Unicode Standard consists of *characters*, written components (i.e., alphabetic letters, numerals, punctuation marks, accent marks, etc.) that can be represented by numeric values. Examples of characters include: U+0041 LATIN CAPITAL LETTER A. In the first character representation, U+*yyyy* is a *code value*, in which U+ refers to Unicode code values, as opposed to other hexadecimal values. The *yyyy* represents a four-digit hexadecimal number of an encoded character. Code values are bit combinations that represent encoded characters. Characters are represented with *glyphs*, various shapes, fonts and sizes for displaying characters. There are no code values for glyphs in the Unicode Standard. Examples of glyphs are shown in Fig. F.2.

The Unicode Standard encompasses the alphabets, ideographs, syllabaries, punctuation marks, *diacritics*, mathematical operators and so on. that comprose the written languages and scripts of the world. A diacritic is a special mark added to a character to distinguish it from another letter or to indicate an accent (e.g., in Spanish, the tilde "~" above the character "n"). Currently, Unicode provides code values for 94,140 character representations, with more than 880,000 code values reserved for future expansion.

Character	UTF-8	UTF-16	UTF-32
LATIN CAPITAL LETTER A	0x41	0x0041	0x00000041
GREEK CAPITAL LETTER ALPHA	0xCD 0x91	0x0391	0x00000391
CJK UNIFIED IDEOGRAPH-4E95	0xE4 0xBA 0x95	0x4E95	0x00004E95
OLD ITALIC LETTER A	0xF0 0x80 0x83 0x80	0xDC00 0xDF00	0x00010300

Fig. F.1 Correlation between the three encoding forms.

Fig. F.2 Various glyphs of the character A.

F.4 Advantages/Disadvantages of Unicode

The Unicode Standard has several significant advantages that promote its use. One is the impact it has on the performance of the international economy. Unicode standardizes the characters for the world's writing systems to a uniform model that promotes transferring and sharing data. Programs developed using such a schema maintain their accuracy because each character has a single definition (i.e., *a* is always U+0061, % is always U+0025). This enables corporations to manage the high demands of international markets by processing different writing systems at the same time. Also, all characters can be managed in an identical manner, thus avoiding any confusion caused by different character-code architectures. Moreover, managing data in a consistent manner eliminates data corruption, because data can be sorted, searched and manipulated via a consistent process.

Another advantage of the Unicode Standard is *portability* (i.e., the ability to execute software on disparate computers or with disparate operating systems). Most operating systems, databases, programming languages and Web browsers currently support, or are planning to support, Unicode. Additionally, Unicode includes more characters than any other character set in common use (although it does not yet include all of the world's characters.

A disadvantage of the Unicode Standard is the amount of memory required by UTF-16 and UTF-32. ASCII character sets are 8 bits in length, so they require less storage than the default 16-bit Unicode character set. However, the *double-byte character set (DBCS)* and the *multi-byte character set (MBCS)* that encode Asian characters (ideographs) require two to four bytes, respectively. In such instances, the UTF-16 or the UTF-32 encoding forms may be used with little hindrance on memory and performance.

F.5 Unicode Consortium's Web Site

If you would like to learn more about the Unicode Standard, visit **`www.unicode.org`**. This site provides a wealth of information about the Unicode Standard. Currently, the home page is organized into various sections: *New to Unicode*, *General Information*, *The Consortium*, *The Unicode Standard*, *Work in Progress* and *For Members*.

The *New to Unicode* section consists of two subsections: ***What is Unicode?*** and ***How to Use this Site***. The first subsection provides a technical introduction to Unicode by describing design principles, character interpretations and assignments, text processing and Unicode conformance. This subsection is recommended reading for anyone new to Unicode. Also, this subsection provides a list of related links that provide the reader with additional information about Unicode. The **How to Use this Site** subsection contains information about using and navigating the site as well hyperlinks to additional resources.

The *General Information* section contains six subsections: ***Where is my Character?***, ***Display Problems?***, ***Useful Resources***, ***Enabled Products***, ***Mail Lists*** and ***Conferences***. The main areas covered in this section include a link to the Unicode code charts (a complete listing of code values) assembled by the Unicode Consortium as well as a detailed outline on how to locate an encoded character in the code chart. Also, the section contains advice on how to configure different operating systems and Web browsers so that the Unicode characters can be viewed properly. Moreover, from this section, the user can navigate to other sites that provide information on various topics, such as fonts, linguistics and such other standards as the *Armenian Standards Page* and the *Chinese GB 18030 Encoding Standard*.

The Consortium section consists of five subsections: ***Who we are***, ***Our Members***, ***How to Join***, ***Press Info*** and ***Contact Us***. This section provides a list of the current Unicode Consortium members as well as information on how to become a member. Privileges for each member type—*full*, *associate*, *specialist* and *individual*—and the fees assessed to each member are listed here.

The Unicode Standard section consists of nine subsections: ***Start Here***, ***Latest Version***, ***Technical Reports***, ***Code Charts***, ***Unicode Data***, ***Updates & Errata***, ***Unicode Policies***, ***Glossary*** and ***Technical FAQ***. This section describes the updates applied to the latest version of the Unicode Standard and categorizes all defined encoding. The user can learn how the latest version has been modified to encompass more features and capabilities. For instance, one enhancement of Version 3.1 is that it contains additional encoded characters. Also, if users are unfamiliar with vocabulary terms used by the Unicode Consortium, they can navigate to the **Glossary** subsection.

The *Work in Progress* section consists of three subsections: ***Calendar of Meetings***, ***Proposed Characters*** and ***Submitting Proposals***. This section presents the user with a catalog of the recent characters included into the Unicode Standard scheme as well as those characters being considered for inclusion. If users determine that a character has been overlooked, then they can submit a written proposal for the inclusion of that character. The **Submitting Proposals** subsection contains strict guidelines that must be adhered to when submitting written proposals.

The *For Members* section consists of two subsections: ***Member Resources*** and ***Working Documents***. These subsections are password protected; only consortium members can access these links.

F.6 Using Unicode

Visual Studio .NET uses Unicode UTF-16 encoding to represent all characters. Figure F.3 uses Visual Basic to display the text "Welcome to Unicode!" in eight different languages: English, French, German, Japanese, Portuguese, Russian, Spanish and Traditional Chinese. [*Note*: The Unicode Consortium's Web site contains a link to code charts that lists the 16-bit Unicode code values.]

The first welcome message (lines 13–16) contains the hexadecimal codes for the English text. The **Code Charts** page on the Unicode Consortium Web site contains a document that lists the code values for the **Basic Latin** *block* (or category), which includes the English alphabet. The hexadecimal codes in lines 13–14 equate to "Welcome." When using Unicode characters in Visual Basic, the format **`&H`***yyyy* is used, where *yyyy* represents the hexadecimal Unicode encoding. For example, the letter "W" (in "Welcome") is denoted by **`&H57`**. [*Note*: The actual code for the letter "W" is **`&H0057`**, but Visual Studio removes the two zeros.] Line 15 contains the hexadecimal for the *space* character (**`&H20`**). The hexadecimal value for the word "to" is on line 15 and the word "Unicode" is on line 14. "Unicode" is not encoded because it is a registered trademark and has no equivalent translation in most languages. Line 16 also contains the **`&H21`** notation for the exclamation mark (**`!`**).

The remaining welcome messages (lines 18–61) contain the hexadecimal codes for the other seven languages. The code values used for the French, German, Portuguese and Spanish text are located in the **Basic Latin** block, the code values used for the Traditional Chinese text are located in the **CJK Unified Ideographs** block, the code values used for

the Russian text are located in the **Cyrillic** block and the code values used for the Japanese text are located in the **Hiragana** block.

```
' Fig. F.3: Unicode.vb
' Using Unicode encoding.

Public Class FrmUnicode
   Inherits System.Windows.Forms.Form

   ' Visual Studio .NET generated code

   Private Sub Form1_Load(ByVal sender As System.Object, _
      ByVal e As System.EventArgs) Handles MyBase.Load

      'English
      lblEnglish.Text = ChrW(&H57) & ChrW(&H65) & ChrW(&H6C) & _
         ChrW(&H63) & ChrW(&H6F) & ChrW(&H6D) & ChrW(&H65) & _
         ChrW(&H20) & ChrW(&H74) & ChrW(&H6F) & ChrW(&H20) & _
         "Unicode" & ChrW(&H21)

      ' French
      lblFrench.Text = ChrW(&H42) & ChrW(&H69) & ChrW(&H65) & _
         ChrW(&H6E) & ChrW(&H76) & ChrW(&H65) & ChrW(&H6E) & _
         ChrW(&H75) & ChrW(&H65) & ChrW(&H20) & ChrW(&H61) & _
         ChrW(&H75) & ChrW(&H20) & "Unicode" & ChrW(&H21)

      ' German
      lblGerman.Text = ChrW(&H57) & ChrW(&H69) & ChrW(&H6C) & _
         ChrW(&H6B) & ChrW(&H6F) & ChrW(&H6D) & ChrW(&H6D) & _
         ChrW(&H65) & ChrW(&H6E) & ChrW(&H20) & ChrW(&H7A) & _
         ChrW(&H75) & ChrW(&H20) & "Unicode" & ChrW(&H21)

      ' Japanese
      lblJapanese.Text = "Unicode" & ChrW(&H3078) & _
         ChrW(&H3087) & ChrW(&H3045) & ChrW(&H3053) & _
         ChrW(&H305D) & ChrW(&H21)

      ' Portuguese
      lblPortuguese.Text = ChrW(&H53) & ChrW(&HE9) & ChrW(&H6A) & _
         ChrW(&H61) & ChrW(&H20) & ChrW(&H42) & _
         ChrW(&H65) & ChrW(&H6D) & ChrW(&H76) & _
         ChrW(&H69) & ChrW(&H6E) & ChrW(&H64) & _
         ChrW(&H6F) & ChrW(&H20) & "Unicode" & ChrW(&H21)

      ' Russian
      lblRussian.Text = ChrW(&H414) & ChrW(&H43E) & ChrW(&H431) & _
         ChrW(&H440) & ChrW(&H43E) & ChrW(&H20) & _
         ChrW(&H43F) & ChrW(&H43E) & ChrW(&H436) & _
         ChrW(&H430) & ChrW(&H43B) & ChrW(&H43E) & _
         ChrW(&H432) & ChrW(&H430) & ChrW(&H442) & _
         ChrW(&H44A) & ChrW(&H20) & ChrW(&H432) & _
         ChrW(&H20) & "Unicode" & ChrW(&H21)

```

Fig. F.3 Windows application demonstrating Unicode encoding (part 1 of 2).

```
        ' Spanish
        lblSpanish.Text = ChrW(&H42) & ChrW(&H69) & ChrW(&H65) & _
           ChrW(&H6E) & ChrW(&H76) & ChrW(&H65) & _
           ChrW(&H6E) & ChrW(&H69) & ChrW(&H64) & _
           ChrW(&H61) & ChrW(&H20) & ChrW(&H61) & _
           ChrW(&H20) & "Unicode" & ChrW(&H21)

        ' Traditional Chinese
        lblChinese.Text = ChrW(&H6B22) & ChrW(&H8FCE) & _
           ChrW(&H4F7F) & ChrW(&H7528) & ChrW(&H20) & _
           "Unicode" & ChrW(&H21)
    End Sub

End Class
```

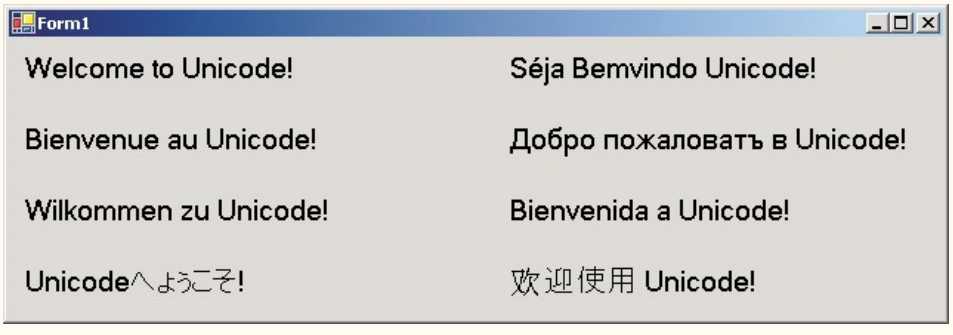

Fig. F.3 Windows application demonstrating Unicode encoding (part 2 of 2).

[*Note*: To render the Asian characters in a Windows application, you would need to install the proper language files on your computer. To do this, open the **Regional Options** dialog from the **Control Panel** (**Start > Settings > Control Panel**). At the bottom of the **General** tab is a list of languages. Check the **Japanese** and the **Traditional Chinese** checkboxes and press **Apply**. Follow the directions of the install wizard to install the languages. For additional assistance, visit **`www.unicode.org/help/display_problems.html`**.]

F.7 Character Ranges

The Unicode Standard assigns code values, which range from **`0000`** (**Basic Latin**) to **`E007F`** (***Tags***), to the written characters of the world. Currently, there are code values for 94,140 characters. To simplify the search for a character and its associated code value, the Unicode Standard generally groups code values by *script* and function (i.e., Latin characters are grouped in a block, mathematical operators are grouped in another block, etc.). As a rule, a script is a single writing system that is used for multiple languages (e.g., the Latin script is used for English, French, Spanish, etc.). The **Code Charts** page on the Unicode Consortium Web site lists all the defined blocks and their respective code values. Figure F.4 lists some blocks (scripts) from the Web site and their range of code values.

Script	Range of Code Values
Arabic	U+0600–U+06FF
Basic Latin	U+0000–U+007F
Bengali (India)	U+0980–U+09FF
Cherokee (Native America)	U+13A0–U+13FF
CJK Unified Ideographs (East Asia)	U+4E00–U+9FAF
Cyrillic (Russia and Eastern Europe)	U+0400–U+04FF
Ethiopic	U+1200–U+137F
Greek	U+0370–U+03FF
Hangul Jamo (Korea)	U+1100–U+11FF
Hebrew	U+0590–U+05FF
Hiragana (Japan)	U+3040–U+309F
Khmer (Cambodia)	U+1780–U+17FF
Lao (Laos)	U+0E80–U+0EFF
Mongolian	U+1800–U+18AF
Myanmar	U+1000–U+109F
Ogham (Ireland)	U+1680–U+169F
Runic (Germany and Scandinavia)	U+16A0–U+16FF
Sinhala (Sri Lanka)	U+0D80–U+0DFF
Telugu (India)	U+0C00–U+0C7F
Thai	U+0E00–U+0E7F

Fig. F.4 Some character ranges.

SUMMARY

- Before Unicode, software developers were plagued by the use of inconsistent character encoding (i.e., numeric values for characters). Most countries and organizations had their own encoding systems, which were incompatible. A good example is the individual encoding systems on the Windows and Macintosh platforms.
- Computers process data by converting characters to numeric values. For instance, the character "a" is converted to a numeric value so that a computer can manipulate that piece of data.
- Without Unicode, localization of global software requires significant modifications to the source code, which results in increased cost and delays in releasing the product.
- Localization is necessary with each release of a version. By the time a software product is localized for a particular market, a newer version, which needs to be localized as well, is ready for distribution. As a result, it is cumbersome and costly to produce and distribute global software products in a market where there is no universal character-encoding standard.
- The Unicode Consortium developed the Unicode Standard in response to the serious problems created by multiple character encodings and the use of those encodings.
- The Unicode Standard facilitates the production and distribution of localized software. It outlines a specification for the consistent encoding of the world's characters and symbols.

- Software products that handle text encoded in the Unicode Standard need to be localized, but the localization process is simpler and more efficient because the numeric values need not be converted.
- The Unicode Standard is designed to be universal, efficient, uniform and unambiguous.
- A universal encoding system encompasses all commonly used characters; an efficient encoding system parses text files easily; a uniform encoding system assigns fixed values to all characters; and an unambiguous encoding system represents the same character for any given value.
- Unicode extends the limited ASCII character set to include all the major characters of the world.
- Unicode makes use of three Unicode Transformation Formats (UTF): UTF-8, UTF-16 and UTF-32, each of which may be appropriate for use in different contexts.
- UTF-8 data consists of 8-bit bytes (sequences of one, two, three or four bytes depending on the character being encoded) and is well suited for ASCII-based systems, where there is a predominance of one-byte characters (ASCII represents characters as one byte).
- UTF-8 is a variable-width encoding form that is more compact for text involving mostly Latin characters and ASCII punctuation.
- UTF-16 is the default encoding form of the Unicode Standard. It is a variable-width encoding form that uses 16-bit code units instead of bytes. Most characters are represented by a single unit, but some characters require surrogate pairs.
- Surrogates are 16-bit integers in the range D800 through DFFF, which are used solely for the purpose of "escaping" into higher numbered characters.
- Without surrogate pairs, the UTF-16 encoding form can only encompass 65,000 characters, but with the surrogate pairs, this is expanded to include over a million characters.
- UTF-32 is a 32-bit encoding form. The major advantage of the fixed-width encoding form is that it uniformly expresses all characters, so that they are easy to handle in arrays and so forth.
- The Unicode Standard consists of characters. A character is any written component that can be represented by a numeric value.
- Characters are represented with glyphs (various shapes, fonts and sizes for displaying characters).
- Code values are bit combinations that represent encoded characters. The Unicode notation for a code value is U+*yyyy*, in which U+ refers to the Unicode code values, as opposed to other hexadecimal values. The *yyyy* represents a four-digit hexadecimal number.
- Currently, the Unicode Standard provides code values for 94,140 character representations.
- An advantage of the Unicode Standard is its impact on the overall performance of the international economy. Applications that conform to an encoding standard can be processed easily by computers anywhere.
- Another advantage of the Unicode Standard is its portability. Applications written in Unicode can be easily transferred to different operating systems, databases, Web browsers and so on. Most companies currently support, or are planning to support, Unicode.
- To obtain more information about the Unicode Standard and the Unicode Consortium, visit **`www.unicode.org`**. It contains a link to the code charts, which contain the 16-bit code values for the currently encoded characters.
- The Unicode Standard has become the default encoding system for XML and any language derived from XML, such as XHTML.
- The Visual Basic .NET IDE uses Unicode UTF-16 encoding to represent all characters.
- In the marking up of Visual Basic documents, the entity reference **`&H`***yyyy* is used, where *yyyy* represents the hexadecimal code value.

TERMINOLOGY

&H*yyyy* notation
ASCII
block
character
character set
code value
diacritic
double-byte character set (DBCS)
efficient (Unicode design basis)
encode
entity reference
glyph
hexadecimal notation
localization
multi-byte character set (MBCS)
portability
script
surrogate
symbol
unambiguous (Unicode design basis)
Unicode Consortium
Unicode design basis
Unicode Standard
Unicode Transformation Format (UTF)
uniform (Unicode design basis)
universal (Unicode design basis)
UTF-8
UTF-16
UTF-32

SELF-REVIEW EXERCISES

F.1 Fill in the blanks in each of the following.

a) Global software developers had to _____ their products to a specific market before distribution.
b) The Unicode Standard is an _____ standard that facilitates the uniform production and distribution of software products.
c) The four design basis that constitute the Unicode Standard are: _____, _____, _____ and _____.
d) A _____ is the smallest written component the can be represented with a numeric value.
e) Software that can execute on different operating systems is said to be _____.
f) Of the three encoding forms, _____ is currently supported by Internet Explorer 5.5 and Netscape Communicator 6.

F.2 State whether each of the following is *true* or *false*. If *false*, explain why.

a) The Unicode Standard encompasses all the world's characters.
b) A Unicode code value is represented as U+*yyyy*, where *yyyy* represents a number in binary notation.
c) A diacritic is a character with a special mark that emphasizes an accent.
d) Unicode is portable.
e) When designing Visual Basic programs, the entity reference is denoted by **#U+***yyyy*.

ANSWERS TO SELF-REVIEW EXERCISES

F.1 a) localize. b) encoding. c) universal, efficient, uniform, unambiguous. d) character. e) portable. f) UTF-8.

F.2 a) False. It encompasses the majority of the world's characters. b) False. The *yyyy* represents a hexadecimal number. c) False. A diacritic is a special mark added to a character to distinguish it from another letter or to indicate an accent. d) True. e) False. The entity reference is denoted by **&H***yyyy*.

EXERCISES

F.3 Navigate to the Unicode Consortium Web site (**www.unicode.org**) and write the hexadecimal code values for the following characters. In which block are they located?

a) Latin letter ‘Z.’
b) Latin letter ‘n’ with the ‘tilde (~).’
c) Greek letter ‘delta.’
d) Mathematical operator ‘less than or equal to.’
e) Punctuation symbol ‘open quote (“).’

F.4 Describe the Unicode Standard design basis.

F.5 Define the following terms:
a) code value.
b) surrogates.
c) Unicode Standard.
d) UTF-8.
e) UTF-16.
f) UTF-32.

F.6 Describe a scenario where it is optimal to store your data in UTF-16 format.

F.7 Using the Unicode Standard code values, create a program that prints your first and last name. If you know other writing systems, print your first and last name in those as well. Use a **`Label`** to display your name.

F.8 Write an ASP.NET program that prints “Welcome to Unicode!” in English, French, German, Japanese, Portuguese, Russian, Spanish and Traditional Chinese. Use the code values provided in Fig. F.3. In ASP.NET, a code value is represented the same way as in a Windows application (**`&H`***yyyy*, where *yyyy* is a four-digit hexadecimal number).