

1

XSLT in Context

This chapter is designed to put XSLT in context. It's about the purpose of XSLT and the task it was designed to perform. It's about what kind of language it is, how it came to be that way, and how it has changed in version 2.0; and it's about how XSLT fits in with all the other technologies that you are likely to use in a typical Web-based application (including, of course, XPath, which forms a vital component of XSLT). I won't be saying very much in this chapter about what an XSLT stylesheet actually looks like or how it works: that will come later, in Chapters 2 and 3.

The chapter starts by describing the task that XSLT is designed to perform — **transformation** — and why there is the need to transform XML documents. I'll then present a trivial example of a transformation in order to explain what this means in practice.

Next, I discuss the relationship of XSLT to other standards in the growing XML family, to put its function into context and explain how it complements the other standards.

I'll describe what kind of language XSLT is, and delve a little into the history of how it came to be like that. If you're impatient you may want to skip the history and get on with using the language, but sooner or later you will ask "why on earth did they design it like that?" and at that stage I hope you will go back and read about the process by which XSLT came into being.

What Is XSLT?

XSLT (Extensible Stylesheet Language: Transformations) is a language that, according to the very first sentence in the specification (found at <http://www.w3.org/TR/xslt20/>), is primarily designed for transforming one XML document into another. However, XSLT is also capable of transforming XML to HTML and many other text-based formats, so a more general definition might be as follows:

XSLT is a language for transforming the structure and content of an XML document.

Why should you want to do that? In order to answer this question properly, we first need to remind ourselves why XML has proved such a success and generated so much excitement.

XML is a simple, standard way to interchange structured textual data between computer programs. Part of its success comes because it is also readable and writable by humans, using nothing more complicated

XSLT in Context

than a text editor, but this doesn't alter the fact that it is primarily intended for communication between software systems. As such, XML satisfies two compelling requirements:

- ❑ **Separating data from presentation:** the need to separate information (such as a weather forecast) from details of the way it is to be presented on a particular device. The early motivation for this arose from the need to deliver information not only to the traditional PC-based Web browser (which itself comes in many flavors) but also to TV sets and handheld devices, not to mention the continuing need to produce print-on-paper. Today, for many information providers an even more important driver is the opportunity to syndicate content to other organizations that can republish it with their own look-and-feel.
- ❑ **Transmitting data between applications:** the need to transmit information (such as orders and invoices) from one organization to another without investing in one-off software integration projects. As electronic commerce gathers pace, the amount of data exchanged between enterprises increases daily, and this need becomes ever more urgent.

Of course, these two ways of using XML are not mutually exclusive. An invoice can be presented on the screen as well as being input to a financial application package, and weather forecasts can be summarized, indexed, and aggregated by the recipient instead of being displayed directly. Another of the key benefits of XML is that it unifies the worlds of documents and data, providing a single way of representing structure regardless of whether the information is intended for human or machine consumption. The main point is that, whether the XML data is ultimately used by people or by a software application, it will very rarely be used directly in the form it arrives: it first has to be transformed into something else.

In order to communicate with a human reader, this something else might be a document that can be displayed or printed: for example, an HTML file, a PDF file, or even audible sound. Converting XML to HTML for display is probably still the most common application of XSLT, and it is the one I will use in most of the examples in this book. Once you have the data in HTML format, it can be displayed on any browser.

In order to transfer data between different applications, we need to be able to transform information from the data model used by one application to the model used by another. To load the data into an application, the required format might be a comma-separated-values file, a SQL script, an HTTP message, or a sequence of calls on a particular programming interface. Alternatively, it might be another XML file using a different vocabulary from the original. As XML-based electronic commerce becomes widespread, the role of XSLT in data conversion between applications also becomes ever more important. Just because everyone is using XML does not mean the need for data conversion will disappear.

There will always be multiple standards in use. As I write there is a fierce debate between the protagonists of two different XML representations of office documents: the ODF specification from the Open Office community, and the OOXML specification from Microsoft and its friends. However this gets resolved, the prospects of a single XML format for all word processor documents are remote, so there will always be a need to transform between multiple formats.

Even within the domain of a single standard, there is a need to extract information from one kind of document and insert it into another. For example, a PC manufacturer who devises a solution to a customer problem will need to extract data from the problem reports and insert it into the documents issued to field engineers so that they can recognize and fix the problem when other customers hit it. The field engineers, of course, are probably working for a different company, not for the original manufacturer. So, linking up enterprises to do e-commerce will increasingly become a case of defining how to extract and combine data from one set of XML documents to generate another set of XML documents, and XSLT is the ideal tool for the job.

During the course of this chapter, we will come back to specific examples of when XSLT should be used to transform XML. For now, I just wanted to establish a feel for the importance and usefulness of transforming XML. If you are already using XSLT, of course, this may be stale news. So let's take a look now at what XSLT version 2.0 brings to the party.

Why Version 2.0?

XSLT 1.0 came out in November 1999 and was highly successful. It was therefore almost inevitable that work would start on a version 2.0. As we will see later, the process of creating version 2.0 was far from smooth and took rather longer than some people hoped. However, XSLT 2.0 was finally published as a W3C Recommendation (that is, a final specification) in January 2007, and user reaction has been very favorable.

It's tempting to look at version 2.0 and see it as a collection of features bolted on to the language, patches to make up for the weaknesses of version 1.0. As with a new release of any other language or software package, most users will find some features here that they have been crying out for, and other additions that appear surplus to requirements.

But I think there is more to version 2.0 than just a bag of goodies; there are some underlying themes that have guided the design and the selection of features. I can identify four main themes:

- ❑ **Integration across the XML standards family:** W3C working groups do not work in isolation from each other; they spend a lot of time trying to ensure that their efforts are coordinated. A great deal of what is in XSLT 2.0 is influenced by a wider agenda of doing what is right for the whole raft of XML standards, not just for XSLT considered in isolation.
- ❑ **Extending the scope of applicability:** XSLT 1.0 was pretty good at rendering XML documents for display as HTML on screen, and for converting them to XSL Formatting Objects for print publishing. But there are many other transformation tasks for which it proved less suitable. Compared with report writers (even those from the 1980s, let alone modern data visualization tools) its data handling capabilities were very weak. The language was quite good at doing conversions of XML documents if the original markup was well designed, but much weaker at recognizing patterns in the text or markup that represent hidden structure. An important aim of XSLT 2.0 was to increase the range of applications that you can tackle using XSLT.
- ❑ **More robust software engineering:** XSLT was always designed to be used both client-side and server-side, but in many ways XSLT 1.0 optimized the language for use in the browser. However, people write large applications in XSLT, containing 100K or more lines of code, and this needs a more rigorous and robust approach to things such as error handling and type checking.
- ❑ **Tactical usability improvements:** Here we *are* into the realm of added goodies. The aim here is to achieve productivity benefits, making it easier to do things that are difficult or error-prone in version 1.0. These are probably the features that existing users will immediately recognize as the most beneficial, but in the long term the other themes probably have more strategic significance for the future of the language.

Before we discuss XSLT in more detail and have a first look at how it works, let's study a scenario that clearly demonstrates the variety of formats to which we can transform XML, using XSLT.

A Scenario: Transforming Music

As an indication of how far XML has now penetrated, Robin Cover's index of XML-based application standards at <http://xml.coverpages.org/xmlApplications.html> today runs to 594 entries. (The last one is entitled *Mind Reading Markup Language*, but as far as I can tell, all the other entries are serious.)

XSLT in Context

I'll follow just one of these 594 links, *XML and Music*, which takes us to <http://xml.coverpages.org/xmlMusic.html>. On this page we find a list of no less than 18 standards, proposals, or initiatives that use XML for marking up music.

This diversity is clearly unnecessary, and many of these initiatives are already dead or dying. Even the names of the standards are chaotic: there is a Music Markup Language, a MusicML, a MusicXML, and a MusiXML, all quite unrelated. There are at least three really serious contenders: the Music Encoding Initiative (MEI), the Standard Music Description Language (SMDL), and MusicXML. The MEI derives its inspiration from the Text Encoding Initiative, and has a particular focus on the needs of music scholars (for example, the ability to capture features found in different manuscripts of the same score), while SMDL is related to the HyTime hypermedia standards and takes into account requirements such as the need to synchronize music with video or with a lighting script (it has not been widely implemented, but it has its enthusiasts). MusicXML, by contrast, is primarily focused on the needs of composers and publishers of sheet music.

Given the variety of requirements, it's unlikely that the number of standards in use will reduce any further. The different notations were invented with different purposes in mind: a markup language used by a publisher for printing sheet music has different requirements from the one designed to let you listen to the music from a browser.

In the first edition of this book, back in 2001, I introduced the idea of using XSLT to transform music as a theoretical possibility, something to make my readers think about the range of possibilities open for the language. By the time I published the second edition, PhD students were showing that it could actually be done. Today, MusicXML is a standard part of over 80 software applications including industry leaders such as Sibelius and Finale, and XSLT is routinely used to manipulate the output. The MEI website publishes XSLT stylesheets for converting between MEI and MusicXML in either direction.




Figure 1-1

As it happens, MusicXML itself provides two ways of representing a score. In one the top-level subdivision in the XML hierarchy is by instrumental part or voice; in the other the top-level structure is the timeline of the music. XSLT stylesheets are provided to convert between the two formats.

Figure 1-1 shows some of the possibilities. You could use XSLT to:

- ❑ Convert music from one representation to another, for example from MEI to SMDL.
- ❑ Convert music from any of these representations into visual music notation, by generating the XML-based vector graphics format SVG.
- ❑ Play the music on a synthesizer, by generating a MIDI (Musical Instrument Digital Interface) file.
- ❑ Perform a musical transformation, such as transposing the music into a different key or extracting parts for different instruments or voices.
- ❑ Extract the lyrics, into HTML or into a text-only XML document.
- ❑ Capture music from non-XML formats and translate it to XML (XSLT 2.0 is especially useful here).

As you can see, XSLT is not just for converting XML documents to HTML.

How Does XSLT Transform XML?

By now you are probably wondering exactly how XSLT goes about processing an XML document in order to convert it into the required output. There are usually two aspects to this process:

1. The first stage is a structural transformation, in which the data is converted from the structure of the incoming XML document to a structure that reflects the desired output.
2. The second stage is formatting, in which the new structure is output in the required format such as HTML or PDF.

The second stage covers the ground we discussed in the previous section; the data structure that results from the first stage can be output as HTML, as a text file, or as XML. HTML output allows the information to be viewed directly in a browser by a human user or be input into any modern word processor. Plain text output allows data to be formatted in the way an existing application can accept, for example comma-separated values or one of the many text-based data interchange formats that were developed before XML arrived on the scene. Finally, XML output allows the data to be supplied to applications that accept XML directly. Typically, this will use a different vocabulary of XML tags from the original document: for example, an XSLT transformation might take the monthly sales figures as its XML input and produce a histogram as its XML output, using the XML-based SVG standard for vector graphics. Or, you could use an XSLT transformation to generate Voice XML output, for aural rendition of your data.

Information about VoiceXML can be found at <http://www.voicexml.org/>.

Let's now delve into the first stage, transformation — the stage with which XSLT is primarily concerned and which makes it possible to provide output in all of these formats. This stage might involve selecting data, aggregating and grouping it, sorting it, or performing arithmetic conversions such as changing centimeters to inches.

XSLT is not the only way of processing or transforming XML. For example, you can write applications in Java or C# that build a tree-like representation of an XML document (using the W3C-defined Document

XSLT in Context

Object Model, or other similar models such as JDOM or XOM). Your program could interrogate this tree structure to find the specific information needed. It would do so by defining a specific *sequence of steps* to be followed in order to produce the required output.

So, how is using XSLT to perform transformations on XML better than writing custom applications? Well, the design of XSLT is based on a recognition that these programs are all very similar, and it should therefore be possible to describe what they do using a high-level *declarative* language rather than writing each program from scratch in Java or C#. The required transformation can be expressed as a set of rules. These rules are based on defining what output should be generated when particular patterns occur in the input. The language is declarative in the sense that you describe the transformation you require, rather than providing a sequence of procedural instructions to achieve it. XSLT describes the required transformation and then relies on the XSLT processor to decide the most efficient way to go about it.

XSLT still relies on an XML parser — it might be a DOM parser or a SAX-compliant one, or one of the new breed of “pull parsers” — to convert the XML document into a tree structure. It is the structure of this tree representation of the document that XSLT manipulates, not the document itself. If you are familiar with the DOM, then you will be happy with the idea of treating every item in an XML document (elements, attributes, processing instructions, and so on) as a node in a tree. With XSLT we have a high-level language that can navigate around a node tree, select specific nodes, and perform complex manipulations on these nodes.

The XSLT tree model is similar in concept to the DOM, but it is not the same. The full XSLT processing model is discussed in Chapter 2.

The description of XSLT given thus far (a declarative language that can navigate to and select specific data and then manipulate that data) may strike you as being similar to that of the standard database query language, SQL. Let’s take a closer look at this comparison.

XSLT and SQL: An Analogy

In a relational database, the data consists of a set of tables. By themselves, the tables are not of much use, the data might as well be stored in flat files in comma-separated values format. The power of a relational database doesn’t come from its data structure; it comes from the language that processes the data, SQL. In the same way, XML on its own just defines a data structure. It’s a bit richer than the tables of the relational model, but by itself it doesn’t actually do anything very useful. It’s when we get a high-level language expressly designed to manipulate the data structure that we start to find we’ve got something interesting on our hands, and for XML data the main language that does that is XSLT.

Superficially, SQL and XSLT are very different languages. But if you look below the surface, they actually have a lot in common. For starters, in order to process specific data, be it in a relational database or an XML document, the processing language must incorporate a declarative query syntax for selecting the data that needs to be processed. In SQL, that’s the `SELECT` statement. In XSLT, the equivalent is the *XPath expression*.

The XPath expression language forms an essential part of XSLT, though it is actually defined in a separate W3C Recommendation (<http://www.w3.org/TR/xpath>) because it can also be used independently of XSLT (the relationship between XPath and XSLT is discussed further on page 22).

The XPath syntax is designed to retrieve nodes from an XML document, based on a path through the XML document or the context in which the node appears. It allows access to specific nodes, while preserving the hierarchy and structure of the document. XSLT instructions are then used to manipulate the results of these queries, for example by rearranging selected nodes and constructing new nodes.

There are further similarities between XSLT and SQL:

- ❑ Both languages augment the basic query facilities with useful additions for performing arithmetic, string manipulation, and comparison operations.
- ❑ Both languages supplement the declarative query syntax with semiprocedural facilities for describing the processing to be carried out, and they also provide hooks to escape into conventional programming languages where the algorithms start to get too complex.
- ❑ Both languages have an important property called *closure*, which means that the output has the same data structure as the input. For SQL, this structure is tables, for XSLT it is trees — the tree representation of XML documents. The closure property is extremely valuable because it means operations performed using the language can be combined end-to-end to define bigger, more complex operations: you just take the output of one operation and make it the input of the next operation. In SQL you can do this by defining views or subqueries; in XSLT you can do it by passing your data through a series of stylesheets, or by capturing the tree produced by one transformation phase in a variable and using that variable as the input of another transformation phase. This last feature is new in XSLT 2.0, though most XSLT 1.0 processors offered a similar capability as a language extension.

In the real world, of course, XSLT and SQL have to coexist. There are many possible relationships, but typically, data is stored in relational databases and transmitted between systems in XML. The two languages don't fit together as comfortably as one would like, because the data models are so different. But XSLT transformations can play an important role in bridging the divide. All of the major relational database vendors have released extensions that allow XML data to be stored and manipulated directly within what is still nominally a relational database.

Before we look at a simple working example of an XSLT transformation, we should briefly discuss a few of the XSLT processors that are available to effect these transformations.

XSLT Processors

The job of an XSLT processor is to apply an XSLT stylesheet to an XML source document and produce a result document.

With XSLT 1.0, there are quite a few good XSLT processors to choose from, and many of them can be downloaded free of charge (but do read the licensing conditions).

If you're using Microsoft technology, there is a choice of two products. The most widely used option is MSXML (Google for "Download MSXML" to find it — the current version is MSXML 6). Usually, I'm not Microsoft's greatest fan, but with this processor it's generally agreed that they have done an excellent job. This product comes as standard with Internet Explorer and is therefore the preferred choice for running transformations in the browser. For the .NET environment, however, Microsoft developed a new processor. This doesn't have a product name of its own, other than its package name within the .NET framework, which is `System.Xml.Xsl`. This processor was initially said to be significantly slower than the MSXML product, but that seems to be no longer true, and if you're writing your application using .NET technologies such as C# and ASP.NET, it's probably the one that you'll find more convenient.

In the Java world, there's a choice of Xalan-J product. There's my own Saxon product (version 6.5.5 is the version that supports XSLT 1.0) available from <http://saxon.sf.net/>, and there's the Xalan-J product which is bundled with Sun's Java JDK software from JDK 1.4 onwards. There's also an Oracle implementation, and a new processor from Intel.

XSLT in Context

For C and C++ the most popular processor is the `libxslt` engine (<http://xmlsoft.org/XSLT/>).

Most of these products are XSLT interpreters, but there are two well-known XSLT compilers: XSLTC, which is distributed as part of the Xalan-J package mentioned earlier, and Gregor, from Jacek Ambroziak (<http://www.ambrosoft.com/gregor.html>). At the top end of the market, for those with a large budget, you can also get hardware-accelerated engines from IBM and Intel.

With XSLT 2.0, at the time of writing, the choice is more limited. There are currently three XSLT 2.0 processors available:

- ❑ My own Saxon product. This is available in two variants, corresponding to the two conformance levels defined in the W3C specification: Saxon-B 9.x (<http://saxon.sf.net/>) is an open-source implementation of a basic XSLT 2.0 processor, and Saxon-SA 9.x (<http://www.saxonica.com/>) is a commercial implementation of a schema-aware XSLT processor. Both variants are available for both the Java and .NET platforms. You can run most of the examples in this book, using a basic XSLT processor, but Chapter 4 and Chapter 11 focus on the additional capability of XSLT when used with XML Schemas, with examples that will only run with a basic XSLT processor.
- ❑ Saxon as delivered offers a command line interface and a Java or .NET API. A number of toolsets provide graphical user interfaces on top: the simplest is the open-source Kernow product (kernowforsaxon.sourceforge.net); other options are the comprehensive Stylus Studio development environment (www.stylusstudio.com) and the oXygen XML editor (www.oxygenxml.com), both of which support step-by-step Saxon debugging.
- ❑ Altova (www.altova.com) have released their own XSLT 2.0 processor, which is available either as a freestanding COM component with a command line interface and Windows API, or built in to the popular XML Spy development environment. As a standalone component the product is free (but not open source), and offers both schema-aware and non-schema-aware processing. The level of conformance of the Altova product has been improving with each release, and nearly all the examples in this book run successfully with the 2008 edition of this product.
- ❑ Gestalt is an open-source XSLT 2.0 processor written in the Eiffel language by Colin Adams (<http://sourceforge.net/projects/gestalt>). The product is work in progress, and although the software is under active development, the documentation is a little sketchy and this implementation is unlikely to appeal to you unless you are something of an Eiffel enthusiast. We have not tested the examples in this book against the Gestalt processor.

By reading papers published at research conferences and blog postings by employees, you can make an educated guess that implementations are underway within IBM, Oracle and Microsoft, but none of these companies has announced any official product plans or release dates at the time of writing — and of course, not everything those companies do in their research labs sees the light of day as a product.

I think it's a little unlikely that there will be quite as many XSLT 2.0 processors as there are for XSLT 1.0 (there is bound to be some shakeout in a maturing market), but I'm confident there will be four or five, which should be enough.

Meanwhile, you can use Altova or Saxon, and those are the primary vehicles I will be using for all the examples in this book.

A Simple XSLT Stylesheet

We're now ready to take a look at an example of using XSLT to transform a very simple XML document.

Example: A “Hello, world!” XSLT Stylesheet

Kernighan and Ritchie in their classic *The C Programming Language* (Prentice-Hall, 1988) originated the idea of presenting a trivial but complete program right at the beginning of the book, and ever since then the `Hello world` program has been an honored tradition. Of course, a complete description of how this example works is not possible until all the concepts have been defined, and so if you feel I’m not explaining it fully, don’t worry — the explanations will come later.

Input

What kind of transformation would we like to do? Let’s try transforming the following XML document.

```
<?xml version="1.0" encoding="iso-8859-1"?>
<?xml-stylesheet type="text/xsl" href="hello.xsl"?>
<greeting>Hello, world!</greeting>
```

This document is available as file `hello.xml` in the download directory for this chapter.

A simple node-tree representation of this document is shown in Figure 1-2.

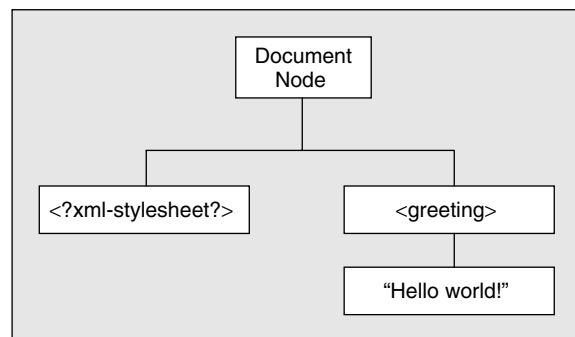


Figure 1-2

There are four nodes in this tree: a document node that represents the document as a whole; an `<?xml-stylesheet?>` processing instruction that identifies the stylesheet to be used; the `<greeting>` element; and the text within the `<greeting>` element.

The document node in the XSLT model performs the same function as the document node in the DOM model (it was called the root node in XSLT 1.0, but the nomenclature has been brought into line with the DOM). The XML declaration is not visible to the XSLT processor and, therefore, is not included in the tree.

I’ve deliberately made it easy by including an `<?xml-stylesheet?>` processing instruction in the source XML file. Many XSLT processors will use this to identify the stylesheet if you don’t specify a different stylesheet to use. The `href` attribute gives the relative URI of the default stylesheet for this document.

Output

Our required output is the following HTML, which will simply change the browser title to "Today's Greeting" and display whatever greeting is in the source XML file:

```
<html>
<head>
  <title>Today's greeting</title>
</head>
<body>
  <p>Hello, world!</p>
</body>
</html>
```

XSLT Stylesheet

Without any more ado, here's the XSLT stylesheet `hello.xsl` to effect the transformation. This is an XSLT 1.0 stylesheet, hence the `<<version = "1.0">` in the `<xsl:stylesheet>` element.

```
<?xml version="1.0" encoding="iso-8859-1"?>
<xsl:stylesheet
  version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

  <xsl:template match="/">
    <html>
      <head>
        <title>Today's greeting</title>
      </head>
      <body>
        <p><xsl:value-of select="greeting"/></p>
      </body>
    </html>
  </xsl:template>
</xsl:stylesheet>
```

Running the Stylesheet

There are a number of ways you can run this stylesheet, which we'll look at in the following sections. It may be worth trying several of the different approaches, to find out which you are most comfortable with — familiarity with the tools will help you master the examples in the rest of the book, as well as making you a more proficient XSLT developer.

Whichever approach you use, the first stage is to download the sample code for this book from the Wrox web site (www.wrox.com), as explained in the Introduction. The code is organized by chapter, so you'll find the two files `hello.xml` and `hello.xsl` in the directory `ch01`.

If you're a serious XSLT developer then your company may well have invested in an XML development environment such as XML Spy, Stylus Studio, or oXygen. All these products

have the capability to run XSLT 2.0 stylesheets: XML Spy uses Altova's XSLT 2.0 processor (though it can also be configured to run Saxon), while Stylus Studio and oXygen both use Saxon. We won't cover the use of these development environments in this book.

Using the Browser

The easiest way to run this example is simply to load the XML file `hello.xml` into any recent version of Internet Explorer or Firefox (right-click the file and select `Open With...`). The browser will recognize the `<?xml-stylesheet?>` processing instruction and will use this to fetch the stylesheet and execute it. The result is a display like the one in Figure 1-3.

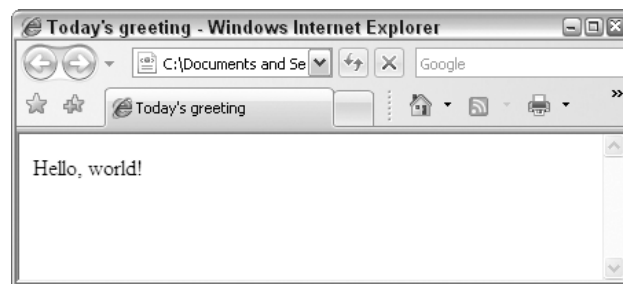


Figure 1-3

We can run this example in the browser because it is actually an XSLT 1.0 stylesheet. Most modern browsers support XSLT 1.0 processing in this way, but at the time of writing none yet supports XSLT 2.0.

Using Saxon from Kernow

We'll be running most of the XSLT 2.0 examples in this book using Saxon, and one of the easiest ways to run Saxon is by using a graphical front-end called Kernow. You can download Kernow from <http://kernowforsaxon.sourceforge.net/>. It includes the Saxon JAR files, so the only prerequisite is that Java itself is installed. You will need Java Standard Edition 6 or later, which you can get from <http://java.sun.com/>. (Saxon itself works with JDK 1.4 or later, the requirement for Java SE 6 comes from Kernow.)

To install Kernow, the only thing you need to do is to unzip the download file into a suitable directory, and if you are running Windows you can then start the product by double-clicking the `Run.bat` file. In the dialog box that appears, select the "Single File" tab, and browse to the source and stylesheet files. Then click Run to run the transformation. The output is as shown in Figure 1-4.

Note that this time, the result is shown in terms of raw HTML. If you want to see what the HTML looks like in a browser, you can always save it to a file and then open the file in your browser. But Kernow is designed for developers, not for end users, and as a developer you need to see the HTML that you have generated. I wouldn't advise anyone to do serious XSLT development by testing directly in a browser, because it's hard to see what's gone wrong when you make a mistake.

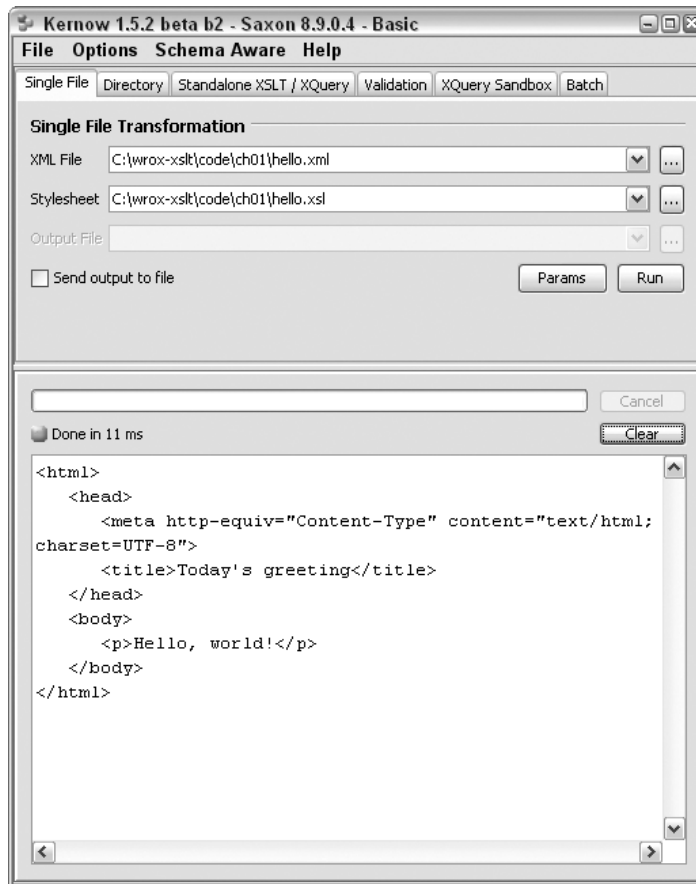


Figure 1-4

Using Saxon on Java from the Command Line

If you're a professional software developer then you're probably not averse to running utilities from the command line, which is the natural way to test your stylesheets when you use Saxon "out of the box." The steps are as follows:

1. Ensure you have Java installed on your machine. Saxon will work with JDK 1.4 or later. Saxon is pure Java code, and so it will run on any platform that supports Java, but I will usually assume that you are using a Windows machine.
2. Download the Saxon-B processor from <http://sf.net/saxon>. Choose the most recent version of Saxon-B for Java, and unzip the download file into a suitable directory, for example `c:\saxon`.
3. Bring up an MSDOS-style console window (use Start ⇨ Run, and type "cmd").

4. Type the following at the command prompt:

```
java -jar c:\saxon\saxon9.jar -a -s:hello.xml
```

5. Admire the HTML displayed on the standard output.

If you want to view the output using your browser, simply save the command line output as an HTML file, in the following manner:

```
java -jar c:\saxon\saxon8.jar -a -s:hello.xml -o:hello.html
```

(Using the command prompt in Windows isn't much fun. If you're a Unix fan, install Cygwin. Most editors have the ability to invoke a command line processor that is generally much more usable than the one provided by the operating system. I use UltraEdit, but there are many other choices available.)

Using Saxon on .NET

The Saxon XSLT processor also runs on the .NET platform. You'll need to make sure that .NET (either version 1.1 or 2.0) is installed — if it's not already present on your machine, it's a free download from Microsoft. Download the latest version of Saxon-B for .NET from <http://sf.net/saxon>, and unzip the download file into a directory such as `c:\saxon`.

You can then run it using the command line:

```
c:\saxon\Transform -a -s:hello.xml -o:hello.html
```

If you want to develop applications with this product, you'll need to install Saxon in the Global Assembly Cache, but for running from the command line, this isn't necessary.

Using the Altova XSLT processor

You can obtain the freestanding AltovaXML product from www.altova.com. The current version at the time of writing is AltovaXML 2008. This includes separate XSLT 1.0 and 2.0 processors, and an XQuery engine. The product installs by default into `c:\Program Files\Altova\AltovaXML2008`. Assuming this directory is on your PATH, you can then run our example transformation with the command line:

```
AltovaXML -in hello.xml -xslt2 hello.xsl -out hello.html
```

Using XMLSpy

If you have installed XMLSpy, use File ⇨ Open to open the stylesheet file `hello.xsl`, then F10 (or the XSL Transformation icon) to run a transformation. In the next dialog box, select the source file `hello.xml`, and click OK. The output is displayed as shown in Figure 1-5. You can switch between the Text and Browser tabs at the bottom of the output window to see either the generated HTML or the result as it will appear in the browser.

XSLT in Context

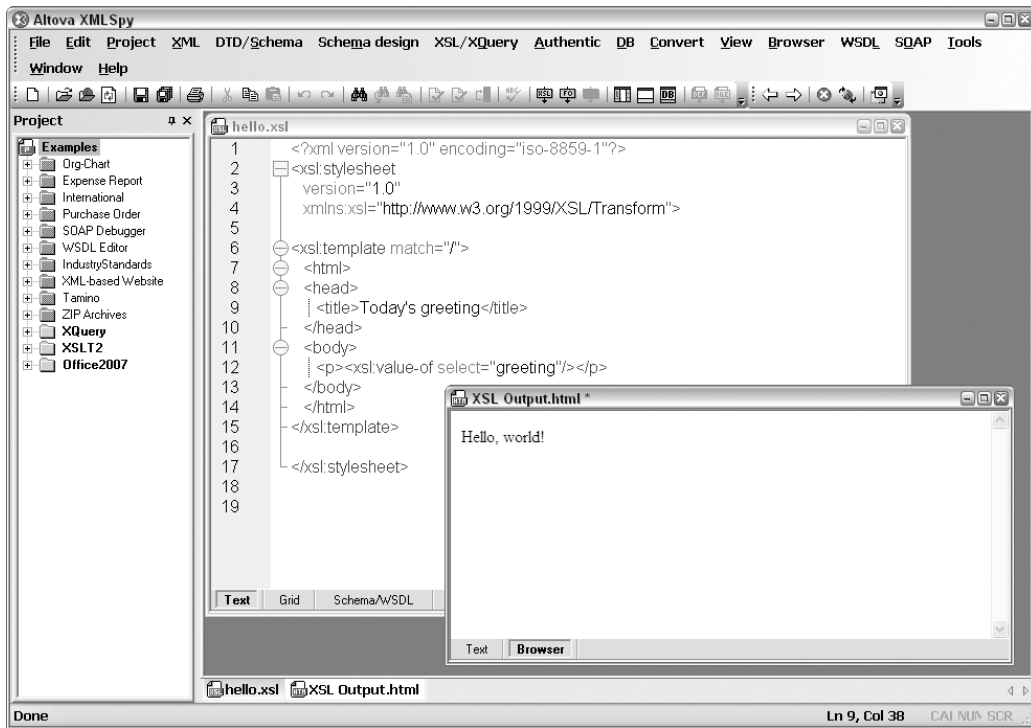


Figure 1-5

Using Stylus Studio

Stylus Studio uses its own built-in XSLT processor by default. This only supports XSLT 1.0, but that's fine for this example.

First open the stylesheet using File ⇨ Open. The next step is to create a scenario (that is a transformation task). Click the "... " button next to the "Create Scenario" drop-down box. In the resulting dialog box (shown in Figure 1-6), define a name for the scenario (so that next time you run the task you can simply select it from the drop-down) and browse to the source XML file in the "Source XML URL" box. Then click the green triangle icon (with hover text "Preview Result") to run the transformation. The output appears in a new window, and you can again switch between a source HTML view and a browser view, or an Explorer-style tree view of the result tree if you prefer.

If you want to use XSLT 2.0 with Stylus Studio, choose the Processor tab in the Create Scenario dialog, and choose the most recent version of Saxon from the list of processors offered (one of the nice features of Stylus is that you can use it to test that your code is portable across a range of different processors). You can then click the Settings... button to select from a range of Saxon-specific options (for example, you can choose whether or not to use schema-aware processing.) These correspond to options that are available on the Saxon command line.

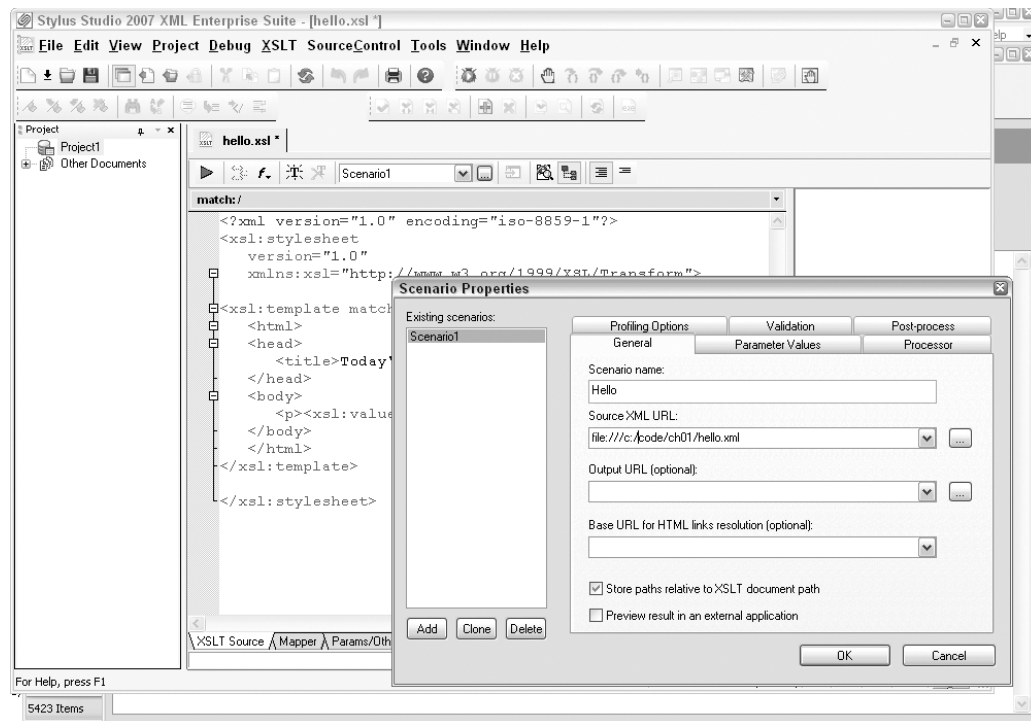


Figure 1-6

How It Works

If you've succeeded in running this example, or even if you just want to get on with reading the book, you'll want to know how it works. Let's dissect it.

```
<?xml version="1.0" encoding="iso-8859-1"?>
```

This is just the standard XML heading. The interesting point is that an XSLT stylesheet is itself an XML document. I'll have more to say about this later in the chapter. I've used `iso-8859-1` character encoding (which is the official name for the character set that Microsoft sometimes calls "ANSI") because in Western Europe and North America it's the character set that most text editors support. If you've got a text editor that supports UTF-8 or some other character encoding, feel free to use that instead.

```

<xsl:stylesheet
  version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

```

This is the standard XSLT 1.0 heading. In XML terms it's an element start tag, and it identifies the document as a stylesheet. The `xmlns:xsl` attribute is an XML namespace declaration, which indicates that the prefix `xsl` is going to be used for elements defined in the W3C XSLT specification. XSLT makes extensive use of XML namespaces, and all the element names

XSLT in Context

defined in the standard are prefixed with this namespace to avoid any clash with names used in your source document. The `version` attribute indicates that the stylesheet is designed to work with an XSLT 1.0 processor.

Let's move on.

```
<xsl:template match="/">
```

An `<xsl:template>` element defines a template rule to be triggered when a particular part of the source document is being processed. The attribute `<match = "/">` indicates that this particular rule is triggered right at the start of processing the source document. Here `</>` is a pattern that identifies the *document node* of the document: an XML document has a hierarchic structure, and in the same way as Unix uses the special filename `</>` to indicate the root of a hierarchic file store, XPath uses `</>` to represent the root of the XML content hierarchy.

```
<html>
<head>
  <title>Today's greeting</title>
</head>
<body>
  <p><xsl:value-of select="greeting" /></p>
</body>
</html>
```

Once this rule is triggered, the body of the template says what output to generate. Most of the template body here is a sequence of HTML elements and text to be copied into the output file. There's one exception: an `<xsl:value-of>` element, which we recognize as an XSLT instruction because it uses the namespace prefix `xsl`. This particular instruction copies the textual content of a node in the source document to the output document. The `select` attribute of the element specifies the node for which the value should be evaluated. The XPath expression `<greeting>` means "find the set of all `<greeting>` elements that are children of the node that this template rule is currently processing." In this case, this means the `<greeting>` element that's the outermost element of the source document. The `<xsl:value-of>` instruction then extracts the text of this element and copies it to the output at the relevant place — in other words, within the generated `<p>` element.

All that remains is to finish what we started.

```
</xsl:template>
</xsl:stylesheet>
```

Why would you want to place today's greeting in a separate XML file and display it using a stylesheet? One reason is that you might want to show the greeting in different ways, depending on the context; for example, it might be shown differently on a different device, or the greeting might depend on the time of day. In this case, you could write a different stylesheet to transform the same source document in a different way. This raises the question of how a stylesheet gets selected at runtime. There is no single answer to this question; it depends on the product you are using.

With Saxon, we used the `-a` option to process the XML document using the stylesheet specified in its `<?xml-stylesheet?>` processing instruction. Instead, we could simply have specified the stylesheet on the command line:

```
java -jar c:\saxon\saxon8.jar -s:hello.xml -xsl:hello.xsl -o:hello.html
```

Having looked at a very simple XSLT 1.0 stylesheet, let's now look at a stylesheet that uses features that are new in XSLT 2.0.

An XSLT 2.0 Stylesheet

This stylesheet is very short, but it manages to use four or five new XSLT 2.0 and XPath 2.0 features within the space of a few lines. I wrote it in response to a user enquiry raised on the `xsl-list` at <http://www.mulberrytech.com/> (an excellent place for meeting other XSLT developers with widely varying levels of experience); so it's a real problem, not an invention. The XSLT 1.0 solution to this problem is about 60 lines of code.

Example: Tabulating Word Frequencies

The problem is simply stated: given any XML document, produce a list of the words that appear in its text, giving the number of times each word appears, together with its frequency.

Input

The input can be any XML document. I will use the text of Shakespeare's *Othello* as an example; this is provided as `othello.xml` in the download files for this book.

Output

The required output is an XML file that lists words in decreasing order of frequency. If you run the transformation using Kernow, the output appears as shown in Figure 1-7.

Stylesheet

Here is the stylesheet that produces this output. You can find it in `wordcount.xsl`.

```
<?xml version="1.0" encoding="iso-8859-1"?>
<xsl:stylesheet
  version="2.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:output method="xml" indent="yes"/>

<xsl:template match="/">
  <wordcount>
    <xsl:for-each-group group-by="." select="
      for $w in //text()/tokenize(., '\W+')[.!=''] return lower-case($w) ">
    <xsl:sort select="count(current-group())" order="descending"/>
    <word word="{current-grouping-key()}"
      frequency="{count(current-group())}"/>
  </wordcount>
</template>
</xsl:stylesheet>
```

XSLT in Context

```
</xsl:for-each-group>
</wordcount>
</xsl:template>
</xsl:stylesheet>
```

Let's see how this works.

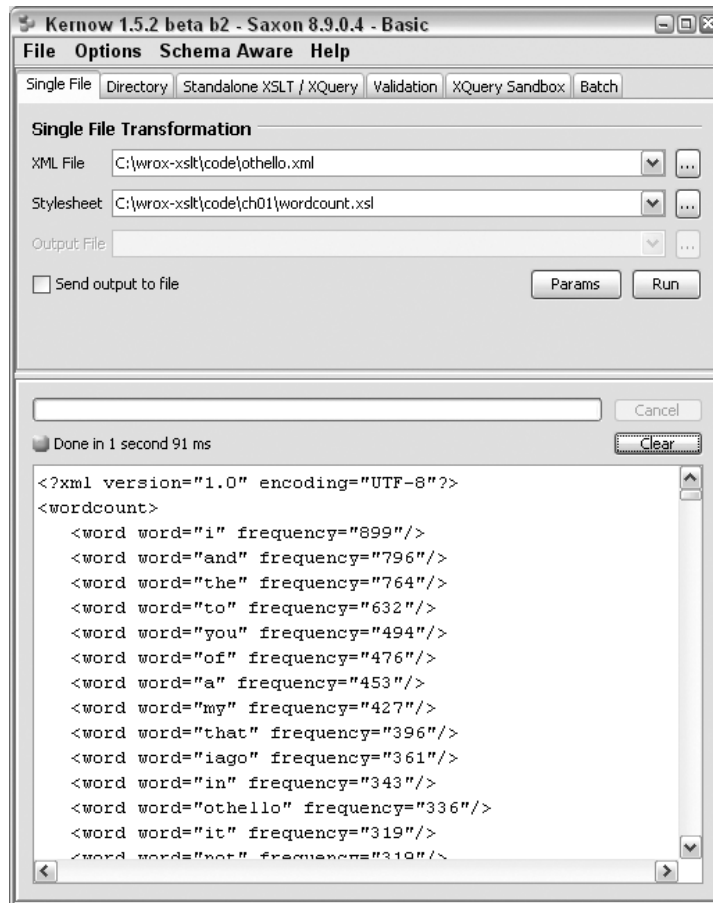


Figure 1-7

The `<xsl:stylesheet>` element introduces the XSLT namespace, as before, and tells us that this stylesheet is designed to be used with an XSLT 2.0 processor.

The `<xsl:output>` element asks for the XML output of the stylesheet to be indented, which makes it much easier for humans to read.

There is one `<xsl:template>` element, as before, which defines the code to be executed when the document node of the source document is encountered. This generates a `<wordcount>` element in the result, and within this it puts the word frequencies.

To understand the `<xsl:for-each-group>` instruction, which is new in XSLT 2.0, we first need to look at its `select` attribute. This contains the XPath 2.0 expression

```
for $w in //text()/tokenize(., '\W+')[.!=''] return lower-case($w)
```

This first selects `<<//text()>`, the set of all text nodes in the input tree. It then tokenizes each of these text nodes, that is, it splits it into a sequence of substrings. The tokenizing is done by applying the regular expression `<<\W+>`. Regular expressions are new in XPath 2.0 and XSLT 2.0, though they will be very familiar to users of other languages such as Perl. They provide the language with greatly enhanced text handling capability. This particular expression, `<<\W+>`, matches any sequence of one-or-more “non-word” characters, a convenient category that includes spaces, punctuation marks, and other separators. So the result of calling the `tokenize()` function is a sequence of strings containing the words that appear in the text. Because there are text nodes that contain nothing of interest, the result also includes some zero-length tokens, and we filter these out by applying the predicate `<<[.!='']>`

The XPath `<<for>` expression now applies the function `lower-case()` to each of the strings in this sequence, producing the lower-case equivalent of the word. (Almost everything in this XPath expression is new in XPath 2.0: the `lower-case()` function, the `tokenize()` function, the `<<for>` expression, and indeed the ability to manipulate a sequence of strings.)

The XSLT stylesheet now takes this sequence of strings and applies the `<xsl:for-each-group>` instruction to it. This processes the body of the `<xsl:for-each-group>` instruction once for each group of selected items, where a group is identified as those items that have a common value for a grouping key. In this case the grouping key is written as `<<group-by=".">`, which means that the values (the words) are grouped on their own value. (In another application, we might have chosen to group them by their length, or by their initial letter.) So, the body of the instruction is executed once for each distinct word, and the `<xsl:sort>` instruction tells us to sort the groups in descending order of the size of the groups (that is, the number of times each word appears). For each of the groups, we output a `<word>` element with two attributes: one attribute is the value we used as the grouping key; the other is the number of items in the group.

Don't worry if this example seemed a bit bewildering: it uses many concepts that haven't been explained yet. The purpose was to give you a feeling for some of the new features in XSLT 2.0 and XPath 2.0, which will all be explained in much greater detail elsewhere in this book.

Having dipped our toes briefly into some XSLT code, I'd now like to take a closer look at the relationship between XSLT and XPath and other XML-related technologies.

The Place of XSLT in the XML Family

XSLT is published by the World Wide Web Consortium (W3C) and fits into the XML family of standards, most of which are also developed by W3C. In this section I will try to explain the sometimes-confusing relationship of XSLT to other related standards and specifications.

XSLT and XSL Formatting Objects

XSLT started life as part of a bigger language called *XSL (Extensible Stylesheet Language)*. As the name implies, XSL was (and is) intended to define the formatting and presentation of XML documents for display on screen, on paper, or in the spoken word. As the development of XSL proceeded, it became clear that this was usually a two-stage process: first a structural transformation, in which elements are selected, grouped and reordered; and then a formatting process in which the resulting elements are rendered as ink on paper, or pixels on the screen. It was recognized that these two stages were quite independent, so XSL was split into two parts: XSLT for defining transformations; and “the rest” — which is still officially called XSL, though most people prefer to call it *XSL-FO (XSL Formatting Objects)* — for the formatting stage.

XSL-FO is nothing more than another XML vocabulary, in which the objects described are areas of the printed page and their properties. Since this is just another XML vocabulary, XSLT needs no special capabilities to generate this as its output. XSL-FO is outside the scope of this book. It’s a big subject. XSL-FO provides wonderful facilities to achieve high-quality typographical output of your documents. However, for many people translating documents into HTML for presentation by a standard browser is quite good enough, and that can be achieved using XSLT alone, or if necessary, by using XSLT in conjunction with Cascading Style Sheets (CSS or CSS2), which I shall return to shortly.

It’s best to avoid the term XSL, because it’s used with so many different meanings. It’s the proper name for XSL Formatting Objects, but many people use it to mean XSLT. It’s also used in older Microsoft documents to refer to their obsolete WD-xsl language, which was issued as part of Internet Explorer 4 before XSLT was standardized in 1999.

XSLT and XPath

Halfway through the development of XSLT 1.0, it was recognized that there was a significant overlap between the expression syntax in XSLT for selecting parts of a document and the XPointer language being developed for linking from one document to another. To avoid having two separate but overlapping expression languages, the two committees decided to join forces and define a single language, *XPath*, which would serve both purposes. XPath 1.0 was published on the same day as XSLT 1.0, November 16, 1999.

XPath acts as a sublanguage within an XSLT stylesheet. An XPath expression may be used for numerical calculations or string manipulations, or for testing Boolean conditions, but its most characteristic use (and the one that gives it its name) is to identify parts of the input document to be processed. For example, the following instruction outputs the average price of all the books in the input document:

```
<xsl:value-of select="avg(//book/@price)"/>
```

Here, the `<xsl:value-of>` element is an instruction defined in the XSLT standard, which causes a value to be written to the output document. The `select` attribute contains an XPath expression, which calculates the value to be written: specifically, the average value of the `price` attributes on all the `<book>` elements. (The `avg()` function too is new in XPath 2.0.)

Following its publication, the XPath specification increasingly took on a life of its own, separate from XSLT. Several DOM implementations (including Microsoft’s) allowed you to select nodes within a DOM tree structure, using a method such as `selectNodes(xpath)`, and this feature is now included in the current version of the standard, DOM3. Subsets of XPath are used within the XML Schema language and in XForms for defining validation conditions, and bindings of XPath to other languages such as

Perl are multiplying. Perhaps most important of all, the designers of XQuery decided to make their language a pure superset of XPath. The language has also proved interesting to academics, and a number of papers have been published analyzing its semantics, which provides the basis for optimized implementations.

XSLT and XML Namespaces

XSLT is designed on the basis that *XML namespaces* are an essential part of the XML standard. So when the XSLT standard refers to an XML document, it always means an XML document that conforms to the XML Namespaces specification, which can be found at <http://www.w3.org/TR/REC-xml-names>.

Namespaces play an important role in XSLT. Their purpose is to allow you to mix tags from two different vocabularies in the same XML document. We've already seen how a stylesheet can mix elements from the target vocabulary (for example HTML or XSL-FO) with elements that act as XSLT instructions. Here's a quick reminder of how namespaces work:

- ❑ Namespaces are identified by a Uniform Resource Identifier (URI). This can take a number of forms. One form is the familiar URL, for example `http://www.wrox.com/namespace`. Another form, not fully standardized but being used in some XML vocabularies, is a URN, for example `urn:biztalk-org:biztalk:biztalk_1`. The detailed form of the URI doesn't matter, but it is a good idea to choose one that will be unique. One good way of achieving this is to use the domain name of your own website. But don't let this confuse you into thinking that there must be something on the website for the URI to point to. The namespace URI is simply a string that you have chosen to be different from other people's namespace URIs; it doesn't need to point to anything.
- ❑ The latest version, XML Namespaces 1.1, allows you to use an International Resource Identifier (IRI) rather than a URI. The main difference is that this permits characters from any alphabet (for example, Chinese); it is no longer confined to ASCII. In practice, most XML parsers have always allowed you to use any characters you like in a namespace URI.
- ❑ Since namespace URIs are often rather long and use special characters such as `</>`, they are not used in full as part of the element and attribute names. Instead, each namespace used in a document can be given a short nickname, and this nickname is used as a prefix of the element and attribute names. It doesn't matter what prefix you choose, because the real name of the element or attribute is determined only by its namespace URI and its local name (the part of the name after the prefix). For example, all my examples use the prefix `xs1` to refer to the namespace URI `http://www.w3.org/1999/XSL/Transform`, but you could equally well use the prefix `xs1t`, so long as you use it consistently.
- ❑ For element names, you can also declare a default namespace URI, which is to be associated with unprefixes element names. The default namespace URI, however, does not apply to unprefixes attribute names.

A namespace prefix is declared using a special pseudo-attribute within any element start tag, with the form:

```
xmlns:prefix = "namespace-URI"
```

This declares a namespace prefix, which can be used for the name of that element, for its attributes, and for any element or attribute name contained in that element. The default namespace, which is used for elements having no prefix (but not for attributes), is similarly declared using a pseudo-attribute:

```
xmlns = "namespace-URI"
```

XSLT in Context

XML Namespaces 1.1 became a Recommendation on February 4, 2004, and the XSLT 2.0 specification makes provision for XSLT processors to work with this version, though it isn't required. Apart from the largely cosmetic change from URIs to IRIs mentioned earlier, the main innovation is the ability to undeclare a namespace, using syntax of the form `<<xmlns:prefix=" " >>`. This is particularly intended for applications like SOAP messaging, where an XML payload document is wrapped in an XML envelope for transmission. Without namespace undeclarations, there is a tendency for namespaces used in the SOAP envelope to stick to the payload XML when this is removed from the envelope, which can cause problems — for example, it can invalidate a digital signature attached to the document.

XSLT and CSS

Why are there two stylesheet languages, XSL (that is, XSLT plus XSL Formatting Objects) as well as Cascading Style Sheets (CSS and CSS2)?

It's only fair to say that in an ideal world there would be a single language in this role, and that the reason there are two is that no one was able to invent something that achieved the simplicity and economy of CSS for doing simple things, combined with the power of XSL for doing more complex things.

CSS is mainly used for rendering HTML, but it can also be used for rendering XML directly, by defining the display characteristics of each XML element. However, it has serious limitations. It cannot reorder the elements in the source document, it cannot add text or images, it cannot decide which elements should be displayed and which omitted, neither can it calculate totals or averages or sequence numbers. In other words, it can only be used when the structure of the source document is already very close to the final display form.

Having said this, CSS is simple to write, and it is very economical in machine resources. It doesn't reorder the document, and so it doesn't need to build a tree representation of the document in memory, and it can start displaying the document as soon as the first text is received over the network. Perhaps, most important of all, CSS is very simple for HTML authors to write, without any programming skills. In comparison, XSLT is far more powerful, but it also consumes a lot more memory and processor power, as well as training budget.

It's often appropriate to use both tools together. Use XSLT to create a representation of the document that is close to its final form, in that it contains the right text in the right order, and then use CSS to add the finishing touches, by selecting font sizes, colors, and so on. Typically, you would do the XSLT processing on the server and the CSS processing on the client (in the browser); so, another advantage of this approach is that you reduce the amount of data sent down the line, which should improve response time for your users and postpone the next expensive bandwidth increase.

XSLT and XML Schemas

One of the biggest changes in XSLT 2.0, and one of the most controversial, is the integration of XSLT with the XML Schema language. XML Schema provides a replacement for DTDs as a way of specifying the structural constraints that apply to a class of documents; unlike DTDs, an XML Schema can regulate the content of the text as well as the nesting of the elements and attributes. Many of the industry vocabularies being used to define XML interchange standards are specified using XML Schema definitions. For example, several of the XML vocabularies for describing music, which I alluded to earlier in the chapter, have an XML Schema to define their rules, and this schema can be used to check the conformance of individual documents to the standard in question.

When you write a stylesheet, you need to make assumptions about the structure of the input documents it is designed to process and the structure of the result documents it is designed to produce. With

XSLT 1.0, these assumptions were implicit; there was no formal way of stating the assumptions in the stylesheet itself. As a result, if you try applying a stylesheet to the wrong kind of input document, the result will generally be garbage.

The idea of linking XSLT and XML Schema was driven by two main considerations:

- ❑ There should, in principle, be software engineering benefits if a program (and a stylesheet is indeed a program) makes explicit assertions about its expected inputs and outputs. These assertions can lead to better and faster detection of errors, often enabling errors to be reported at compile time that otherwise would only be reported the first time the stylesheet was applied to some test data that happened to exercise a particular part of the code.
- ❑ The more information that's available to an XSLT processor at compile time, the more potential it has to generate optimal code, giving faster execution and better use of memory.

So why the controversy? It's mainly because XML Schema itself is less than universally popular. It's an extremely complex specification that's very hard to read, and when you discover what it says, it appears to be full of rules that seem artificial and inconsistent. It manages at the same time to be specified in very formal language, and yet to have a worryingly high number of bugs that have been fixed through published errata. Although there are good books that present XML Schema in a more readable way, they achieve this by glossing over the complications, which means that the error messages you get when you do something wrong can be extremely obscure. As a result, there has been a significant amount of support for an alternative schema language, Relax NG, which as it happens was co-developed by the designer of XSLT and XPath, James Clark, and is widely regarded as a much more elegant approach.

The XSL and XQuery Working Groups responded to these concerns by ensuring that support for XML Schema was optional, both for implementors and for users. This has largely silenced the objections.

The signs are that XML Schema is here to stay, whether people like it or not. It has the backing of all the major software vendors such as IBM, Oracle, and Microsoft, and it has been adopted by most of the larger user organizations and industries. And like so many things that the IT world has adopted as standards, it may be imperfect but it does actually work. Meanwhile, to simplify the situation rather cruelly, Relax NG is taking the role of the Apple Mac: the choice of the cognoscenti who judge a design by its intrinsic quality rather than by cold-blooded cost-benefit analysis.

As I've already mentioned, W3C is not an organization that likes to let a thousand flowers bloom. It is not a loose umbrella organization in which each working group is free to do its own thing. There are strong processes that ensure the working groups cooperate and strive to reconcile their differences. There is therefore a determination to make all the specifications work properly together, and the message was that if XML Schema had its problems, then people should work together to get them fixed. XSLT and XML Schema come from the same stable, so they were expected to work together. And now that the specs are finished and products are out, I think users are starting to discover that they can work together beneficially.

Chapter 4 provides an overview of how stylesheets and schemas are integrated in XSLT 2.0, and Chapter 19 provides a worked example of an application that uses this capability. When I first developed this application for the book (which I did at the same time as I developed the underlying support in Saxon), I was pleasantly surprised to see that I really was getting benefits from the integration. At the simplest level, I really liked the immediate feedback you get when a stylesheet generates output that does not conform to the schema for the result document, with error messages that point straight to the offending line in the stylesheet. This makes for a much faster debugging cycle than does the old approach of putting the finished output file through a schema validator as a completely separate operation.

XSLT in Context

XSLT and XQuery

XQuery is a separate specification from W3C, designed to allow data in XML documents to be queried. It can operate on single documents, or on collections containing millions of documents held in an XML database.

Functionally, XQuery offers a subset of the capabilities of XSLT. You could regard it as XSLT without the template rules, and without some of the extra features such as the ability to do grouping, or to format dates and times, or to import modules and selectively override them. It would be a mistake, however, to think that being a smaller language makes XQuery a poor relation. The relative simplicity of XQuery does indeed make it harder to write large and complex applications, but it does bring two significant advantages: the language is easier to learn, especially for those coming from a SQL background, and it is easier to optimize, especially when running against gigabytes of data preloaded and preindexed in an XML database.

XQuery has XPath 2.0 as a subset. This makes it very much a member of the same family as XSLT. The two languages have a great deal in common, most importantly their type system. There are no formal facilities in the W3C specifications that allow XSLT and XQuery to be mixed in a single application, but because the processing models are so closely aligned, many implementations allow one language to be called from the other. Saxon in fact implements both languages as different surface syntaxes for the same underlying processing engine.

There are some applications for which XSLT is definitely better suited, particularly document publishing. There are others where XQuery is the only sensible choice, notably searching for data in large XML databases. There's a third class of applications, especially message conversion, where either language will do the job, and where the choice is largely a matter of personal preference. My advice would be to use XQuery if it's a very small application and XSLT if it's bigger, largely because in my experience it's easier to write XSLT code that's adaptable to change and reusable in different applications.

The History of XSL

Like most of the XML family of standards, XSLT was developed by the World Wide Web Consortium (W3C), a coalition of companies orchestrated by Tim Berners-Lee, the inventor of the Web. There is an interesting page on the history of XSL, and styling proposals generally, at <http://www.w3.org/Style/History/>.

Writing history is a tricky business. Sharon Adler, the chair of the XSL Working Group, tells me that her recollections of events are very different from the way I describe them. This just goes to show that the documentary record is a very crude snapshot of what people were actually thinking and talking about. Unfortunately, though, it's all that we've got.

Prehistory

HTML was originally conceived by Berners-Lee (www.w3.org/Markup/draft-ietf-iiir-html-01.txt) as a set of tags to mark the logical structure of a document; headings, paragraphs, links, quotes, code sections, and the like. Soon, people wanted more control over how the document looked; they wanted to achieve the same control over the appearance of the delivered publication as they had with printing and paper. So, HTML acquired more and more tags and attributes to control presentation; fonts, margins, tables, colors, and all the rest that followed. As it evolved, the documents being published became more and more browser-dependent, and it was seen that the original goals of simplicity and universality were starting to slip away.

The remedy was widely seen as separation of content from presentation. This was not a new concept; it had been well developed through the 1980s in the development of *Standard Generalized Markup Language* (SGML).

Just as XML was derived as a greatly simplified subset of SGML, so XSLT has its origins in an SGML-based standard called *DSSSL* (*Document Style Semantics and Specification Language*). DSSSL (pronounced *Disssel*) was developed primarily to fill the need for a standard device-independent language to define the output rendition of SGML documents, particularly for high-quality typographical presentation. SGML was around for a long time before DSSSL appeared in the early 1990s, but until then the output side had been handled using proprietary and often extremely expensive tools, geared toward driving equally expensive phototypesetters, so that the technology was really taken up only by the big publishing houses.

Michael Sperberg-McQueen and Robert F. Goldstein presented an influential paper at the WWW '94 conference in Chicago under the title *A Manifesto for Adding SGML Intelligence to the World-Wide Web*. You can find it at <http://tigger.uic.edu/~cmsmcq/htmlmax.html>.

The authors presented a set of requirements for a stylesheet language, which is as good a statement as any of the aims that the XSL designers were trying to meet. As with other proposals from around that time, the concept of a separate transformation language had not yet appeared, and a great deal of the paper is devoted to the rendition capabilities of the language. There are many formative ideas, however, including the concept of fallback processing to cope with situations where particular features are not available in the current environment.

It is worth quoting some extracts from the paper here:

Ideally, the stylesheet language should be declarative, not procedural, and should allow stylesheets to exploit the structure of SGML documents to the fullest. Styles must be able to vary with the structural location of the element: paragraphs within notes may be formatted differently from paragraphs in the main text. Styles must be able to vary with the attribute values of the element in question: a quotation of type "display" may need to be formatted differently from a quotation of type "inline"...

At the same time, the language has to be reasonably easy to interpret in a procedural way: implementing the stylesheet language should not become the major challenge in implementing a Web client.

The semantics should be additive: It should be possible for users to create new stylesheets by adding new specifications to some existing (possibly standard) stylesheet. This should not require copying the entire base stylesheet; instead, the user should be able to store locally just the user's own changes to the standard stylesheet, and they should be added in at browse time. This is particularly important to support local modifications of standard DTDs.

Syntactically, the stylesheet language must be very simple, preferably trivial to parse. One obvious possibility: formulate the stylesheet language as an SGML DTD, so that each stylesheet will be an SGML document. Since the browser already knows how to parse SGML, no extra effort will be needed.

We recommend strongly that a subset of DSSSL be used to formulate stylesheets for use on the World Wide Web; with the completion of the standards work on DSSSL, there is no reason for any community to invent their own style-sheet language from scratch. The full DSSSL standard may well be too demanding to implement in its entirety, but even if that proves true, it provides only an argument for defining a subset of DSSSL that must be supported, not an argument for rolling our own. Unlike home-brew specifications, a subset of a standard comes with an automatically predefined growth path. We expect to work on the formulation of a usable, implementable subset of DSSSL for use in WWW stylesheets, and invite all interested parties to join in the effort.

XSLT in Context

In late 1995, a W3C-sponsored workshop on stylesheet languages was held in Paris. In view of the subsequent role of James Clark as editor of the XSLT Recommendation, it is interesting to read the notes of his contribution on the goals of DSSSL, which can be found at http://www.w3.org/Style/951106_Workshop/report1.html#clark.

Here are a few selected paragraphs from these notes:

DSSSL contains both a transformation language and a formatting language. Originally the transformation was needed to make certain kinds of styles possible (such as tables of contents). The query language now takes care of that, but the transformation language survives because it is useful in its own right.

The language is strictly declarative, which is achieved by adopting a functional subset of Scheme. Interactive stylesheet editors must be possible.

A DSSSL stylesheet very precisely describes a function from SGML to a flow object tree. It allows partial stylesheets to be combined ("cascaded" as in CSS): some rule may override some other rule, based on implicit and explicit priorities, but there is no blending between conflicting styles.

James Clark closed his talk with the remark:

Creating a good, extensible style language is hard!

One suspects that the effort of editing the XSLT 1.0 Recommendation didn't cause him to change his mind.

The First XSL Proposal

Following these early discussions, the W3C set up a formal activity to create a stylesheet language proposal. The remit for this group specified that it should be based on DSSSL.

As an output of this activity came the first formal proposal for XSL, dated August 27, 1997. Entitled *A Proposal for XSL*, it lists 11 authors: James Clark (who works for himself), five from Microsoft, three from Imso Corporation, one from ArborText, and one (Henry Thompson) from the University of Edinburgh. The document can be found at <http://www.w3.org/TR/NOTE-XSL.html>.

The section describing the purpose of the language is worth reading.

XSL is a stylesheet language designed for the Web community. It provides functionality beyond CSS (e.g. element reordering). We expect that CSS will be used to display simply structured XML documents and XSL will be used where more powerful formatting capabilities are required or for formatting highly structured information such as XML structured data or XML documents that contain structured data.

Web authors create content at three different levels of sophistication given as follows:

- markup: relies solely on a declarative syntax*
- script: additionally uses code "snippets" for more complex behaviors*
- program: uses a full programming language*

XSL is intended to be accessible to the "markup" level user by providing a declarative solution to most data description and rendering requirements. Less common tasks are accommodated through a graceful escape to a familiar scripting environment. This approach is familiar to the Web publishing community as it is modeled after the HTML/JavaScript environment.

The powerful capabilities provided by XSL allow:

- ❑ *formatting of source elements based on ancestry/descendancy, position, and uniqueness*
- ❑ *the creation of formatting constructs including generated text and graphics*
- ❑ *the definition of reusable formatting macros*
- ❑ *writing-direction independent stylesheets*
- ❑ *extensible set of formatting objects*

The authors then explained carefully why they had felt it necessary to diverge from DSSSL and described why a separate language from CSS (Cascading Style Sheets) was thought necessary.

They then stated some design principles:

- ❑ XSL should be straightforwardly usable over the Internet.
- ❑ XSL should be expressed in XML syntax.
- ❑ XSL should provide a declarative language to do all common formatting tasks.
- ❑ XSL should provide an “escape” into a scripting language to accommodate more sophisticated formatting tasks and to allow for extensibility and completeness.
- ❑ XSL will be a subset of DSSSL with the proposed amendment. (*As XSL was no longer a subset of DSSSL, they cannily proposed amending DSSSL so it would become a superset of XSL.*)
- ❑ A mechanical mapping of a CSS stylesheet into an XSL stylesheet should be possible.
- ❑ XSL should be informed by user experience with the FOSI stylesheet language.
- ❑ The number of optional features in XSL should be kept to a minimum.
- ❑ XSL stylesheets should be human-legible and reasonably clear.
- ❑ The XSL design should be prepared quickly.
- ❑ XSL stylesheets shall be easy to create.
- ❑ Terseness in XSL markup is of minimal importance.

As a requirements statement, this doesn’t rank among the best. It doesn’t read like the kind of list you get when you talk to users and find out what they need. It’s much more the kind of list that designers write when they know what they want to produce, including a few political concessions to the people who might raise objections. But if you want to understand why XSLT became the language it did, this list is certainly evidence of the thinking.

The language described in this first proposal contains many of the key concepts of XSLT as it finally emerged, but the syntax is virtually unrecognizable. It was already clear that the language should be based on templates that handled nodes in the source document matching a defined pattern, and that the language should be free of side effects, to allow “progressive rendering and handling of large documents.” I’ll explore the significance of this requirement in more detail on page 34, and discuss its implications on the way stylesheets are designed in Chapter 17. The basic idea is that if a stylesheet is expressed as a collection of completely independent operations, each of which has no external effect other than generating part of the output from its input (for example, it cannot update global variables), then it becomes possible to generate any part of the output independently if that particular part of the input changes. Whether the XSLT language actually achieves this objective is still an open question.

The first Working Draft of XSL (not to be confused with the Proposal) was published on August 18, 1998, and the language started to take shape, gradually converging on

XSLT in Context

the final form it took in the November 16, 1999 Recommendation through a series of Working Drafts, each of which made radical changes, but kept the original design principles intact.

A Recommendation is the most definitive of documents produced by the W3C. It's not technically a standard, because standards can only be published by government-approved standards organizations. But I will often refer to it loosely as "the standard" in this book.

Saxon

At this point it might be a good idea to clarify how I got involved in the story. In 1998 I was working for the British computer manufacturer ICL, a part of Fujitsu. Fujitsu, in Japan, had developed an object database system, later marketed as Jasmine, and I was using this technology to build content management applications for large publishers. We developed a few successful large applications, but found that it was too complex for people who wanted something in six weeks rather than six months. So I was asked to look at what we could do with XML, which was just appearing on the horizon.

I came to the conclusion that XML looked like a good thing, but that there wasn't any software. So I developed the very first early versions of Saxon to provide a proof-of-concept demonstration. At that stage Saxon was just a Java library, not an XSLT processor, but as the XSL standards developed I found that my own ideas were converging more and more with what the W3C working group was doing, and I started implementing the language as it was being specified. ICL had decided that its marketing resources were spread thinly over too many products, and so the management took the imaginative decision to make the technology available as open source. Seventeen days after the XSLT 1.0 specification was published in November 1999, I announced the first conformant implementation. And on the day it was published, I started work on the first edition of this book.

When the book was published, the XSL Working Group invited me to join and participate in the development of XSLT 1.1. Initially, being based in the United Kingdom and with limited time available for the work, my involvement was fairly sporadic. But early in 2001 I changed jobs and joined Software AG, which wanted me to take a full role in the W3C work. The following year James Clark pulled out of the Working Group, and I stepped into his shoes as editor.

The reason I'm explaining this sequence of events is that I hope it will help you to understand the viewpoint from which this book is written. When I wrote the first edition I was an outsider, and I felt completely free to criticize the specification when I felt it necessary. I have tried to retain an objective approach in the present edition, but as editor of the language spec it is much more difficult to be impartial. I've tried to keep a balance: it wouldn't be fair to use the book as a platform to push my views over those of my colleagues of the working group, but at the same time, I've made no effort to be defensive about decisions that I would have made differently if they had been left to me.

Software AG continued to support my involvement in the W3C work (on the XQuery group as well as the XSL group), along with the development of Saxon and the writing of this book, through till February 2004, at which point I left to set up my own company, Saxonica.

Beyond XSLT 1.0

After XSLT 1.0 was published, the XSL Working Group responsible for the language decided to split the requirements for enhancements into two categories: XSLT 1.1 would standardize a small number of

urgent features that vendors had already found necessary to add to their products as extensions, while XSLT 2.0 would handle the more strategic requirements that needed further research.

A working draft of XSLT 1.1 was published on December 12, 2000. It described three main enhancements to the XSLT 1.0 specification: the ability to produce multiple output documents, the ability to use temporary trees to create a multi-pass transformation, and standard bindings to extension functions written in Java or ECMAScript.

For a number of reasons XSLT 1.1 never got past the working draft stage. This was partly because of controversy surrounding the Java language bindings, but more particularly because it was becoming clearer that XSLT 2.0 would be a fairly radical revision of the language, and the Working Group didn't want to do anything in 1.1 that would get in the way of achieving the 2.0 goals. There were feelings, for example, that the facility for temporary trees might prejudice the ability to support sequences in 2.0, a fear which as it happens proved largely unfounded.

Convergence with XQuery

By the time work on XSLT 2.0 was starting, the separate XQuery Working Group in W3C had created a draft of its own language.

While the XSL Working Group had identified the need for a transformation language to support a self-contained part of the formatting process, XQuery originated from the need to search large quantities of XML documents stored in a database.

Work on an XML Query Language had started as early as 1998. A workshop was held in December 1998, and you can find all 66 position papers presented at this workshop at <http://www.w3.org/TandS/QL/QL98/pp.html>. It's interesting to see how the participants saw the relationship with XSL, as it was then known. The Microsoft position paper states the belief that a query language could be developed as an extension of XSLT, but in this it is almost alone. Many of the participants came from a database background, with ideas firmly rooted in the tradition of SQL and object database languages such as OQL, and to these people, XSL didn't look remotely like a query language. But in the light of subsequent events, it's interesting to read the position paper from the XSL Working Group, which states in its summary:

1. The query language should use XSL patterns as the basis for information retrieval.
2. The query language should use XSL templates as the basis for materializing query results.
3. The query language should be at least as expressive as XSL is, currently.
4. Development of the pattern and transformation languages should remain in the XSL Working Group.
5. A coordination group should ensure either that a single query language satisfies all working group requirements or that all W3C query languages share an underlying query model.

(Remember that XPath had not yet been identified as a separate language, and that the expressions that later became XPath were then known as patterns.)

This offer to coordinate, and the strong desire to ensure consistency among the different W3C specifications, can be seen as directly leading to the subsequent collaboration between the two working groups to define XPath 2.0.

The XQuery group started meeting in September 1999. The first published requirements document was published the following January (<http://www.w3.org/TR/2000/WD-xmlquery-req-20000131>). It

XSLT in Context

included a commitment to compatibility with XML Schema, and a rather cautiously worded promise to “take into consideration the expressibility and search facilities of XPath when formulating its algebra and query syntax.” July 2000 saw a revised requirements document that included a selection of queries that the language must be able to express. The first externally visible draft of the XQuery language was published in February 2001 (see <http://www.w3.org/TR/2001/WD-xquery-20010215/>), and it was at this stage that the collaboration between the two working groups began in earnest.

The close cooperation between the teams developing the two languages contrasts strangely with the somewhat adversarial position adopted by parts of the user community. XSLT users were quick to point out that XSLT 1.0 satisfied every single requirement in the first XQuery requirements document, and could solve all the use cases published in the second version in August 2000. At the same time, users on the XQuery side of the fence have often been dismissive about XSLT, complaining about its verbose syntax and sometimes arcane semantics. Even today, when the similarities of the two languages at a deep level are clearly apparent, there is little overlap between their user communities: I find that most users of the XQuery engine in Saxon have no XSLT experience. The difference between XSLT and XQuery is in many ways a difference of style rather than substance, but users often feel strongly about style.

The Development of XSLT 2.0 and XPath 2.0

The requirements for XSLT 2.0 and XPath 2.0 were published on February 14, 2001. In the case of the XPath 2.0 requirements, the document was written jointly by the two working groups. You can find the documents at the following URLs:

<http://www.w3.org/TR/2001/WD-xslt20req-20010214>

<http://www.w3.org/TR/2001/WD-xpath20req-20010214>

Broadly, the requirements fall into three categories:

- ❑ Features that are obviously missing from the current standards and that would make users’ lives much easier, for example, facilities for grouping related nodes, extra string-handling and numeric functions, and the ability to read text files as well as XML documents.
- ❑ Changes desired by the XML Query Working Group. The difficulty at this stage was that the Query group did not just want additions to the XPath language; they wanted fundamental changes to its semantics. Many members of the XQuery group felt they could not live with some of the arbitrariness of the way XPath handled data types generally, and node-sets in particular, for example the fact that `<a = 1>` tests whether there is some `<a>` that equals one, whereas `<a - 1 = 0>` tests whether the first `<a>` equals one.
- ❑ Features designed to exploit and integrate with XML Schema. The W3C XML Schema specification had reached an advanced stage (it became a Candidate Recommendation on October 20, 2000), and implementations were starting to appear in products. The thinking was that if the schema specified that a particular element contains a number or a date (for example), then it ought to be possible to use this knowledge when comparing or sorting dates within a stylesheet.

The development of XSLT 2.0, culminating in the Recommendation of January 23, 2007, proved to be a long drawn out process. There were early delays getting agreement with the XQuery group on the details of XPath 2.0. This took a long time firstly, because of the number of people involved; secondly, because of the very different places where people were coming from (the database community and the document community have historically been completely isolated from each other, and it took a lot of talking

before people started to understand each others' positions); and finally, because of the sheer technical difficulty of finding a workable design that offered the right balance between backwards compatibility and rigorous, consistent semantics. Later, when the specification appeared to be all but finished, it still took a couple of years to get through the public reviews demanded by the W3C process, which generated thousands of detailed comments.

But that's all history now. Let's look next at the essential characteristics of XSLT 2.0 as a language.

XSLT 2.0 as a Language

What are the most significant characteristics of XSLT as a language, which distinguish it from other languages? In this section I shall pick five of the most striking features: the fact that it is written in XML syntax, the fact that it is a language free of side effects, the fact that processing is described as a set of pattern-matching rules, the fact that it has a type system based on XML Schema, and the fact that it is a two-language system in which one language (XPath) is embedded in another (XSLT).

Use of XML Syntax

As we've seen, the use of SGML syntax for stylesheets was proposed as long ago as 1994, and it seems that this idea gradually became the accepted wisdom. It's difficult to trace exactly what the overriding arguments were, and when you find yourself writing something like:

```
<xsl:variable name="y">
  <xsl:call-template name="f">
    <xsl:with-param name="x"/>
  </xsl:call-template>
</xsl:variable>
```

to express what in other languages would be written as $\langle y = f(x) \rangle$, then you may find yourself wondering how such a decision came to be made.

The most obvious arguments for expressing XSLT stylesheets in XML are perhaps as follows:

- ❑ There is already an XML parser in the browser; so it keeps the footprint small if this can be reused.
- ❑ Everyone had got fed up with the syntactic inconsistencies between HTML/XML and CSS and didn't want the same thing to happen again.
- ❑ The Lisp-like syntax of DSSSL was widely seen as a barrier to its adoption; so it would be better to have a syntax that was already familiar in the target community.
- ❑ Many existing popular template languages (including simple ASP and JSP pages) are expressed as an outline of the output document with embedded instructions; so this is a familiar concept.
- ❑ The lexical apparatus is reusable, for example Unicode support, character and entity references, whitespace handling, namespaces.
- ❑ Visual development tools remove the inconvenience of typing lots of angle brackets.
- ❑ It's often useful to have a stylesheet as the input or output of a transformation; so it's a benefit if a stylesheet can read and write other stylesheets.

XSLT in Context

In my experience, the most pervasive argument is the last one: it's surprising how often complex applications construct or modify stylesheets on the fly. But like it or not, the XML-based syntax is now an intrinsic feature of the language that has both benefits and drawbacks. It does make the language verbose, but in the end, the number of keystrokes has very little bearing on the ease or difficulty of solving particular transformation problems.

In XSLT 2.0, the long-windedness of the language has been reduced considerably by increasing the expressiveness of the non-XML part of the syntax, namely XPath expressions. Many computations that required five lines of XSLT code in 1.0 can now be expressed in a single XPath expression. Two constructs in particular led to this simplification: the conditional expression (`if..then..else`) in XPath 2.0; and the ability to define a function in XSLT (using `<xsl:function>`) that can be called directly from an XPath expression. To take the example discussed earlier, if you replace the template `<f>` by a user-written function `<f>`, you can replace the five lines in the example with:

```
<xsl:variable name="y" select="f($x)"/>
```

The decision to base the XSLT syntax on XML has proved its worth in several ways that I would not have predicted in advance:

- ❑ It has proved very easy to extend the syntax. Adding new elements and attributes is trivial; there is no risk of introducing parsing difficulties when doing so, and it is easy to manage backwards compatibility. (In contrast, extending XQuery's non-XML syntax without introducing parsing ambiguities is a highly delicate operation.)
- ❑ The separation of XML parsing from XSLT processing leads to good error reporting and recovery in the compiler. It makes it much easier to report the location of an error with precision and to report many errors in one run of the compiler. This leads to a faster development cycle.
- ❑ It makes it easier to maintain stylistic consistency between different constructs in the language. The discipline of defining the language through elements and attributes creates a constrained vocabulary with which the language designers must work, and these constraints impose a certain consistency of design.

No Side Effects

The idea that XSL should be a declarative language free of side effects appears repeatedly in the early statements about the goals and design principles of the language, but no one ever seems to explain *why*: what would be the user benefit?

A function or procedure in a programming language is said to have side effects if it makes changes to its environment; for example, if it can update a global variable that another function or procedure can read, or if it can write messages to a log file, or prompt the user. If functions have side effects, it becomes important to call them the right number of times and in the correct order. Functions that have no side effects (sometimes called pure functions) can be called any number of times and in any order. It doesn't matter how many times you evaluate the area of a triangle, you will always get the same answer; but if the function to calculate the area has a side effect such as changing the size of the triangle, or if you don't know whether it has side effects or not, then it becomes important to call it once only.

I expand further on this concept in the section on Computational Stylesheets in Chapter 17, page 985.

It is possible to find hints at the reason why this was considered desirable in the statements that the language should be equally suitable for batch or interactive use, and that it should be capable of *progressive rendering*. There is a concern that when you download a large XML document, you won't be able to see anything on your screen until the last byte has been received from the server. Equally, if a small change were made to the XML document, it would be nice to be able to determine the change needed to the screen display, without recalculating the whole thing from scratch. If a language has side effects, then the order of execution of the statements in the language has to be defined, or the final result becomes unpredictable. Without side effects, the statements can be executed in any order, which means it is possible, in principle, to process the parts of a stylesheet selectively and independently.

What it means in practice to be free of side effects is that you cannot update the value of a variable. This restriction is something many users find very frustrating at first, and a big price to pay for these rather remote benefits. But as you get the feel of the language and learn to think about using it the way it was designed to be used, rather than the way you are familiar with from other languages, you will find you stop thinking about this as a restriction. In fact, one of the benefits is that it eliminates a whole class of bugs from your code. I shall come back to this subject in Chapter 17, where I outline some of the common design patterns for XSLT stylesheets and, in particular, describe how to use recursive code to handle situations where in the past you would probably have used updateable variables to keep track of the current state.

Rule-Based

The dominant feature of a typical XSLT stylesheet is that it consists of a set of template rules, each of which describes how a particular element type or other construct should be processed. The rules are not arranged in any particular order; they don't have to match the order of the input or the order of the output, and in fact there are very few clues as to what ordering or nesting of elements the stylesheet author expects to encounter in the source document. It is this that makes XSLT a declarative language, because you specify what output should be produced when particular patterns occur in the input, as distinct from a procedural program where you have to say what tasks to perform in what order.

This rule-based structure is very like CSS, but with the major difference that both the patterns (the description of which nodes a rule applies to), and the actions (the description of what happens when the rule is matched) are much richer in functionality.

Example: Displaying a Poem

Let's see how we can use the rule-based approach to format a poem. Again, we haven't introduced all the concepts yet, and so I won't try to explain every detail of how this works, but it's useful to see what the template rules actually look like in practice.

Input

Let's take this poem as our XML source. The source file is called `poem.xml`, and the stylesheet is `poem.xsl`.

```
<poem>
  <author>Rupert Brooke</author>
  <date>1912</date>
  <title>Song</title>
```

```
<stanza>
  <line>And suddenly the wind comes soft,</line>
  <line>And Spring is here again;</line>
  <line>And the hawthorn quickens with buds of green</line>
  <line>And my heart with buds of pain.</line>
</stanza>
<stanza>
  <line>My heart all Winter lay so numb,</line>
  <line>The earth so dead and frore,</line>
  <line>That I never thought the Spring would come again</line>
  <line>Or my heart wake any more.</line>
</stanza>
<stanza>
  <line>But Winter's broken and earth has woken,</line>
  <line>And the small birds cry again;</line>
  <line>And the hawthorn hedge puts forth its buds,</line>
  <line>And my heart puts forth its pain.</line>
</stanza>
</poem>
```

Output

Let's write a stylesheet such that this document appears in the browser, as shown in Figure 1-8.

Stylesheet

It starts with the standard header.

```
<xsl:stylesheet
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  version="1.0">
```

Now we write one template rule for each element type in the source document. The rule for the `<poem>` element creates the skeleton of the HTML output, defining the ordering of the elements in the output (which doesn't have to be the same as the input order). The `<xsl:value-of>` instruction inserts the value of the selected element at this point in the output. The `<xsl:apply-templates>` instructions cause the selected child elements to be processed, each using its own template rule.

```
<xsl:template match="poem">
  <html>
  <head>
    <title><xsl:value-of select="title"/></title>
  </head>
  <body>
    <xsl:apply-templates select="title"/>
    <xsl:apply-templates select="author"/>
    <xsl:apply-templates select="stanza"/>
    <xsl:apply-templates select="date"/>
  </body>
  </html>
</xsl:template>
```

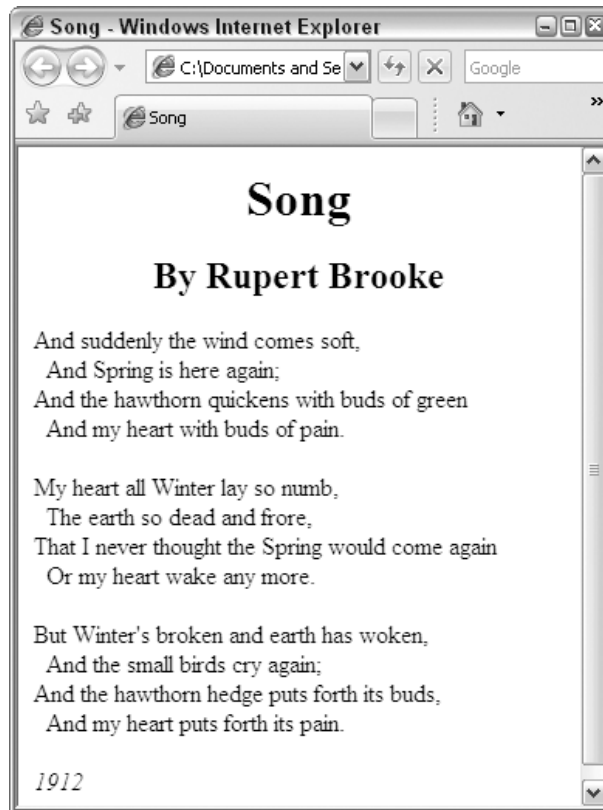


Figure 1-8

In XSLT 2.0 we could replace the four `<xsl:apply-templates>` instructions with one, written as follows:

```
<xsl:apply-templates select="title, author, stanza, date"/>
```

This takes advantage of the fact that the type system for the language now supports ordered sequences. The `«,»` operator performs list concatenation and is used here to form a list containing the `<title>`, `<author>`, `<stanza>`, and `<date>` elements in that order. Note that this includes all the `<stanza>` elements, so in general this will be a sequence containing more than four items.

The template rules for the `<title>`, `<author>`, and `<date>` elements are very simple; they take the content of the element (denoted by `«select = ". »`), and surround it within appropriate HTML tags to define its display style.

```
<xsl:template match="title">
  <div align="center">
    <h1><xsl:value-of select="."/;></h1>
  </div>
</xsl:template>
```

XSLT in Context

```
<xsl:template match="author">
  <div align="center">
    <h2>By <xsl:value-of select="." /></h2>
  </div>
</xsl:template>

<xsl:template match="date">
  <p><i><xsl:value-of select="." /></i></p>
</xsl:template>
```

The template rule for the `<stanza>` element puts each stanza into an HTML paragraph, and then invokes processing of the lines within the stanza, as defined by the template rule for lines:

```
<xsl:template match="stanza">
  <p><xsl:apply-templates select="line" /></p>
</xsl:template>
```

The rule for `<line>` elements is a little more complex: if the position of the line within the stanza is an even number, it precedes the line with two non-breaking-space characters (` `). The `<xsl:if>` instruction tests a boolean condition, which in this case calls the `position()` function to determine the relative position of the current line. It then outputs the contents of the line, followed by an empty HTML `
` element to end the line.

```
<xsl:template match="line">
  <xsl:if test="position() mod 2 = 0">&#160;&#160;</xsl:if>
  <xsl:value-of select="." /><br/>
</xsl:template>
```

And to finish off, we close the `<xsl:stylesheet>` element:

```
</xsl:stylesheet>
```

Although template rules are a characteristic feature of the XSLT language, we'll see that this is not the only way of writing a stylesheet. In Chapter 17, I will describe four different design patterns for XSLT stylesheets, only one of which makes extensive use of template rules. In fact, the `Hello World` stylesheet I presented earlier in this chapter doesn't make any real use of template rules; it fits into the design pattern I call *fill-in-the-blanks*, because the stylesheet essentially contains the fixed part of the output with embedded instructions saying where to get the data to put in the variable parts.

Types Based on XML Schema

I have described three characteristics of the XSLT language (the use of XML syntax, the principle of no side-effects, and the rule-based processing model) that were essential features of XSLT 1.0 and that have been retained essentially unchanged in XSLT 2.0. The fourth characteristic is new in XSLT 2.0, and creates a fundamental change in the nature of XSLT as a language. This is the adoption of a type system based on XML Schema.

There are two aspects to the type system of any programming language. The first is the set of types that are supported (for example, integers, strings, lists, tuples), together with the mechanisms for creating user-defined types. The second aspect is the rules that the language enforces to ensure type-correctness.

XSLT 1.0 had a very small set of types (integers, booleans, strings, node-sets, and result tree fragments), and the rules it applied were what is often called “weak typing”: this means that the processor would always attempt to convert the value supplied in an expression or function call to the type that was required in that context. This makes for a very happy-go-lucky environment: if you supply an integer where a string is expected, or vice versa, nothing will break.

XSLT 2.0 has changed both aspects of the type system. There is now a much richer set of types available (and this set is user-extensible), and the rules for type checking are stricter.

We will look at the implications of this in greater detail in Chapters 4 and 5.

A Two-Language System: XSLT and XPath

The final characteristic of XSLT that I want to explore is its use of XPath as an embedded sublanguage. This is not unlike the way SQL is often embedded in other languages, and it is done for the same reason — it allows XPath to be used in many different contexts.

The fact that XPath was designed to be embedded in other language has some specific consequences:

- ❑ An embedded language does not need to have every conceivable piece of functionality. In the language of computer science, it does not need to be computationally complete. In more practical terms, it can be restricted to being able to access variables but not to declare them, to call functions but not to define them, to navigate around nodes in a tree but not to create new nodes.
- ❑ An embedded language can depend on a context established by the host language in which it is embedded. If an embedded language is to be well integrated with a host language, then they should share information so that the user does not need to declare things twice, once for each language. The information that XPath shares with its host language is called the context. This can be divided into information that’s available at compile time (the static context), and information that’s not available until runtime (the dynamic context). Both aspects of the XPath context are described in Chapter 7 of this book.
- ❑ An embedded language can be called by its host language, but cannot make calls in the other direction. XSLT can invoke XPath, but not the other way around. This means that there is incomplete composability of expressions when the XSLT language is considered as a whole. However, XPath can invoke XSLT indirectly by means of function calls. Some functionality, notable conditional and iteration constructs, are duplicated in XSLT and XPath.

The syntax of XPath has some unusual features, which reflect the fact that it amalgamates ideas from a number of different sources. One can identify three different syntactic styles within XPath expressions:

- ❑ **Conventional programming expressions:** This allows the same kind of expressions, infix operators, and function calls as many other programming languages; an example is an expression such as `<<$x + 1 = round($y) mod 3>`. Such expressions trace their roots via programming languages such as Algol and Fortran back to the notations of elementary mathematics.
- ❑ **Path expressions:** These perform hierarchic selection of a node within a tree, an example is `<</a/b//c>`. These expressions can be seen as a generalization of the syntax used by operating systems to identify files within a hierarchic filestore.
- ❑ **Predicate logic:** This includes the `<<for>`, `<<some>` and `<<every>` expressions, for example `<<for $i in //item[@price > 30] return $i/@code>`. These expressions, which are new in XPath 2.0, derive from the tradition of database query languages (SQL, the object database language OQL, and precursors to XQuery), which can be seen as adaptations of the notation of mathematical symbolic logic.

XSLT in Context

Some other factors that have influenced the design of the XPath syntax include:

- ❑ A decision that XPath should have no reserved words. This means that any name that is legal as an XML element name (which includes names such as `<<and>>` and `<<for>>`) should be legal in a path expression, without any need for escaping. As a result, all names used with some other role, for example function names, variable names, operator names, and keywords such as `<<for>>`, have to be recognizable by their context.
- ❑ In both the original applications for XPath (that is, XSLT and XPointer), the language was designed to be embedded within the attributes of an XML document. It therefore has no mechanisms of its own for character escaping, relying instead on the mechanisms available at the XML level (such as numeric character references and entity references). This also made the designers reluctant to use symbols such as `<&&&>`, which would require heavy escaping. This principle has been abandoned in XPath 2.0 with the introduction of the operators `<<<>` and `<>>>`; however, these operators are not likely to be used very often.
- ❑ There was originally an expectation that XPath expressions (especially in an XPointer environment) would often be used as fragment identifiers in a URI. This usage of XPointer never really took off, but it meant there was a reluctance to use special characters such as `<&#>`, `<&%>`, and `<&?>` that have special significance in URIs.

Despite its disparate syntactic roots and its lexical quirks, XPath has managed to integrate these different kinds of expression surprisingly well. In particular, it has retained full composability within itself, so any kind of expression can be nested inside any other.

Summary

This introductory chapter answered the following questions about XSLT:

- ❑ What kind of language is it?
- ❑ How is it used?
- ❑ Where does it fit into the XML family?
- ❑ Where does it come from and why was it designed the way it is?

You now know that XSLT is a declarative high-level language designed for transforming the structure of XML documents; that it has two major applications: data conversion and presentation; and that it can be used at a number of different points in the overall application architecture, including at data capture time, at delivery time on the server, and at display time on the browser. You also have some idea why XSLT has developed in the way it has.

Now it's time to start taking an in-depth look inside the language to see how it does this job. In the next chapter, we look at the way transformation is carried out by treating the input and output as tree structures, and using patterns to match particular nodes in the input tree and define what nodes should be added to the result tree when the pattern is matched.