

PB138 — Transforming XML Data, XSLT

(C) 2019 Masaryk University -- Tomáš Pitner, Luděk Bártek, Adam Rambousek

Transformation Languages

XSLT

is the best known XML transformation language. The XSLT 1.0 W3C recommendation was published in 1999 together with XPath 1.0, and it has been widely implemented since then. XSLT 2.0 has become a W3C recommendation since January 2007 and implementations of the specification like Saxon 8 are already available.

Transformation Languages (2)

XQuery

is a full functional language, despite having "query" in the name. It is a de facto standard used by Microsoft, Oracle, DB2, Mark Logic, etc., is the foundation for the XRX web programming model, and has a W3C recommendation for versions 1.0. XQuery is not written in XML itself like XSLT is, so its syntax is much lighter. The language is based on XPath 2.0. XQuery programs cannot have side-effects, just like XSLT and provides almost the same capabilities (for instance: declaring variables and functions, iterating over sequences, using W3C schema types), even though the program syntax are quite different. XQuery is logic driven, using FOR, WHERE and function composition (e.g. `fn:concat("<html>", generate-body(), "</html>")`). In contrast, XSLT is data-driven (push processing model) where certain conditions of the input document trigger the execution of templates rather than the code executing in the order in which it is written.

Transformation Languages (3)

XProc

is an XML Pipeline language. The XProc 1.0 W3C Recommendation was published in May 2010.

XML document transform

is a Microsoft standard for performing simple transforms on XML documents. Primarily for creating IIS Web.config files (Config Transforms), other implementations allow it to be used for generic config files as build time (Slow Cheetah) or from the command line (CTT).

STX (Streaming Transformations for XML)

is inspired by XSLT but has been designed to allow a one-pass transformation process that never prevents streaming. Implementations are available in Java (Joost) and Perl (XML::STX).

XML Script

is an imperative scripting language inspired by Perl that uses the XML syntax. XML Script supports XPath and its proprietary DSLPath for selecting nodes from the input tree.

Transformation Languages (4)

FXT

is a functional XML transformation tool, implemented in Standard ML.

XDuce

is a typed language with a lightweight syntax, compared to XSLT. It is written in ML.

CDuce

extends XDuce to a general-purpose functional programming language, see CDuce homepage.

XACT

is a Java-based system for programming XML transformations. Notable features include XML templates as immutable values and a static analysis to ensure type safety using XML Schema types (XACT home page).

Transformation Languages (5)

XFun

is a functional language X-Fun for defining transformations between XML data trees, while providing shredding instructions. X-Fun can be understood as an extension of Frisch's XStream language with output shredding, while pattern matching is replaced by tree navigation with XPath expressions. ([1])

XStream

is a simple functional transformation language for XML documents based on CAML. XML transformations written in XStream are evaluated in streaming: when possible, parts of the output are computed and produced while the input document is still being parsed. Some transformations can thus be applied to huge XML documents which would not even fit in memory. The XStream compiler is distributed under the terms of the CeCILL free software license.

Transformation Languages (6)

Xtatic

applies methods from XDuce to C#, see Xtatic homepage. HaXml::is a library and collection of tools to write XML transformations in Haskell. Its approach is very consistent and powerful. Also see this paper about HaXml published in 1999 and this IBM developerWorks article. See also the more recent HXML and Haskell XML Toolbox (HXT), which is based on the ideas of HaXml and HXML but takes a more general approach to XML processing.

XMLambda (XMλ)

is described in a 1999 paper by Erik Meijer and Mark Shields. No implementation is available. See XMLambda home page.

Transformation Languages (7)

FleXML

is an XML processing language first implemented by Kristofer Rose. Its approach is to add actions to an XML DTD specifying processing instructions for any subset of the DTD's rules.

Scala

is a general-purpose functional and object-oriented language with specific support for XML transformation in the form of XML pattern matching, literals, and expressions, along with standard XML libraries.[1]

LINQ to XML

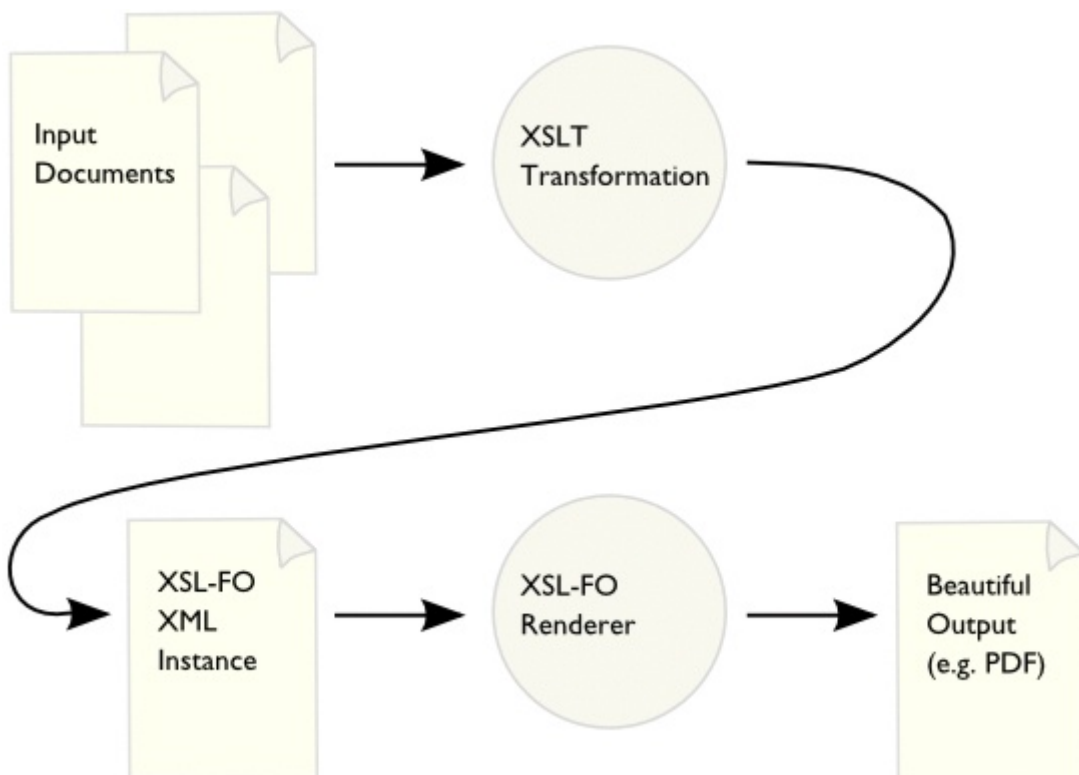
is a .NET 3.5 syntax and programming API available in C#, VB and some other .NET languages. LINQ is primarily designed as a query language, but it also supports XML transforms.

Extensible Stylesheet Language Transformation (XSLT)

- XSLT is a language for specifying transformation of XML documents on the (usually) XML outputs, or text, HTML or other output formats.
- The original application area, the transformation of XML data to [XSL:FO](#) (XSL-Formatting Objects), thus rendering XML.
- XSLT specification was therefore part of [XSL \(eXtensible Stylesheet Language\)](#).
- Later, XSL was set aside and XSLT began to be seen as a universal general description language XML → XML (txt, HTML) transformations.
- XSLT is a Turing-complete language

XSLT Original Purpose

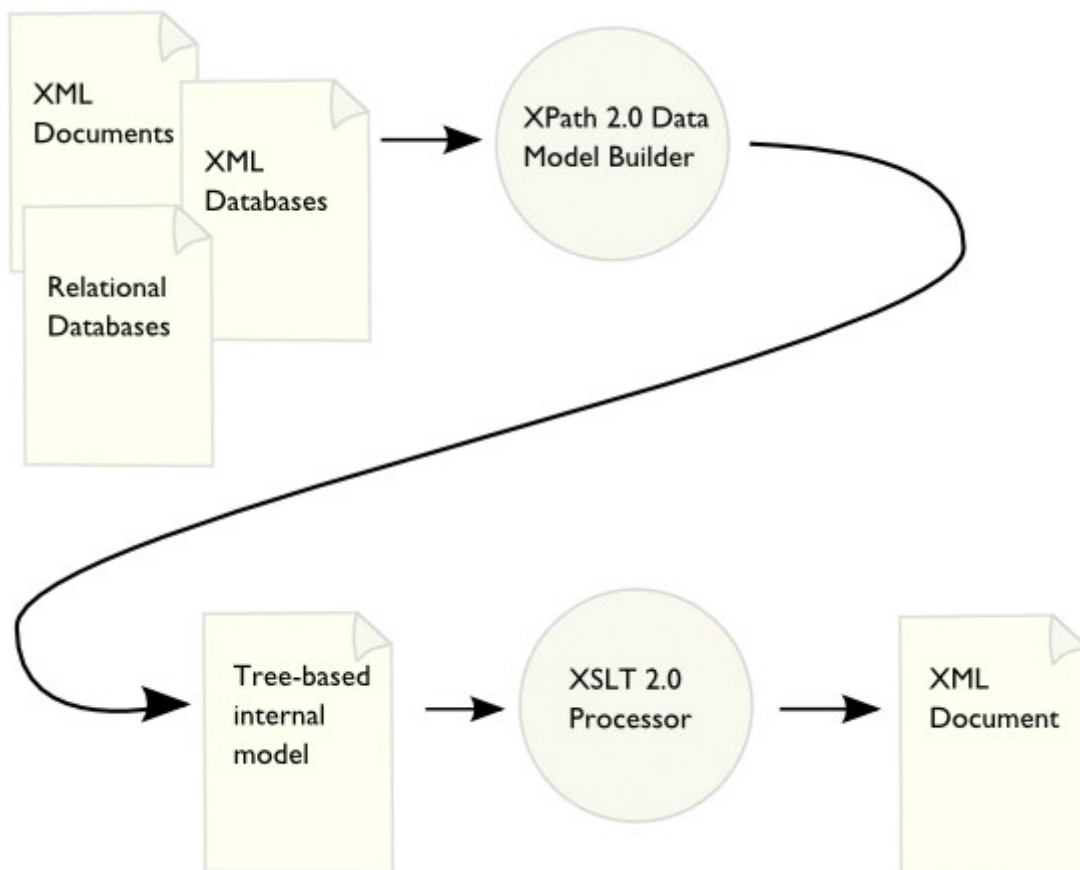
XML to XSL-Formatting Object transformation



Now the General Goal

- Now since 2.0, the goal is to enable transformations from *any* resources capable of building a XML Document Model from them.
- So XSLT 2.0 processors can operate not only on XML but on anything that can be made to look like XML:
 - relational database tables,
 - geographical information systems,
 - file systems,
 - anything from which your XSLT processor can build an XDM instance.

Now the General Goal



Versions

- The current versions are defined by the XSLT 1.0, 2.0 and 3.0 specifications:
 - [XSLT 1.0 XSL Transformations \(XSLT\) — W3C Recommendation 16 November 1999](#)
 - [XSL Transformations \(XSLT\) Version 2.0 — W3C Recommendation 23 January 2007](#)
 - [XSL Transformations \(XSLT\) Version 3.0 — W3C Last Call Working Draft 2 October 2014](#)
 - The version 1.0 is still widely used, mainly due to lacking (free) implementations of the

newer specifications. Also, for many purposes, the features of 1.0 are sufficient.

The main principles

- XSLT is a functional language, where reduction rules have the form of templates, which specify how nodes in the source document override output document.
- XSLT transformation specification is contained in a **style file (stylesheet)**, which is an XML document written in the XSLT syntax. The root element is either `xsl:stylesheet` or `xsl:transformation` (which are synonyms) where `xsl:` is a prefix for the XSL namespace.
- The XSLT style(sheet) is then processed by an **XSLT processor** and subsequently,
- XML file(s) can be transformed using that stylesheet.

Typical XSLT Workflow

[XSLT workflow] | [./images/XSLT_en.svg](#)

XSLT style composition

- XSLT stylesheet contains set of **templates**, represented by `xsl:template` elements.
- Templates have a **selection part** corresponding with the left-hand side of a reduction rule in a functional language and the **construction part** representing the right-hand side of such a rule.

Example: XML Source

(Wikipedia, [XSLT](#))

```
<?xml version="1.0" encoding="UTF-8"?>
<persons>
  <person username="JS1">
    <name>John</name>
    <family-name>Smith</family-name>
  </person>
  <person username="MI1">
    <name>Morka</name>
    <family-name>Ismincius</family-name>
  </person>
</persons>
```

Example: XSLT Stylesheet

(Wikipedia, [XSLT](#))

```
<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
    version="1.0">
  <xsl:output method="xml" indent="yes"/>
  <xsl:template match="/persons">
    <root>
      <xsl:apply-templates select="person"/>
    </root>
  </xsl:template>
  <xsl:template match="person">
    <name username="{@username}">
      <xsl:value-of select="name" />
    </name>
  </xsl:template>
</xsl:stylesheet>
```

Example: Resulting XML

(Wikipedia, [XSLT](#))

```
<?xml version="1.0" encoding="UTF-8"?>
<root>
  <name username="JS1">John</name>
  <name username="MI1">Morka</name>
</root>
```

XSLT Processing

1. XSLT processor (interpreter) takes the *stylesheet* (XSLT code)
2. Usually compiles it into an internal form.
3. Then it takes the nodes from the *input document*, looks for an appropriate template and select it.
4. Then it produces a *result fragment* corresponding to the construction part of the selected template.
5. Recursive takes next nodes from the input document and applies the same procedure for them.

Open-source XSLT Processors

libxslt

is a free library released under the MIT License that can be reused in commercial applications. It is based on libxml and implemented in C. It can be used at the command line via [xsltproc](#) which is included in OS X and many Linux distributions.

WebKit, Blink

used for example in the Safari and Chrome web browsers respectively, uses the libxslt library to do XSL transformations.

Saxon

an XSLT (2.0 and partial 3.0) and XQuery 3.0 processor with open-source and proprietary commercial versions for stand-alone operation and for Java, JavaScript and .NET. The open-source version does not support XSLT 3.0.

Open-source XSLT Processors (2)

Xalan

an open source XSLT 1.0 processor from the Apache Software Foundation available stand-alone and for Java and C++. Integrated into Java SE.

Web browsers

Safari, Chrome, Firefox, Opera and Internet Explorer all support XSLT 1.0. None supports XSLT 2.0 natively, although the third party products like *Saxon-CE (Saxon-Client Edition)* and *Frameless* can provide this functionality. Browsers can perform on-the-fly transformations of XML files and display the transformation output in the browser window. This is done either by embedding the XSL in the XML document or by referencing a file containing XSL instructions from the XML document. The latter may not work with Chrome because of its security model.

XMLStarlet

[XMLStarlet](#) is a command line XML toolkit which can be used to transform, query, validate, and edit XML documents and files using simple set of shell commands in similar way it is done for plain text files using grep/sed/awk/tr/diff/patch. It does not require Java. Available for Windows and Linux. Supports XSLT 1.0.

Commercial XSLT Processors

MSXML and .NET

includes an XSLT 1.0 processor. From *MSXML 4.0* it includes the command line utility msxsl.exe.

QuiXSLT

an XSLT 3.0 processor doing streaming implemented in Java by Innovimax and INRIA.

Saxon

commercial versions support the newest standards such as XSLT 3.0.

Information Resources

W3C XSLT 1.0 Recommendation

[XSLT 1.0](#) is still the most used version.

What is XSLT? on XML.COM

<http://www.xml.com/pub/a/2000/08/holman/index.html>

Mulberrytech.com XSLT Quick Reference (2xA4, PDF)

<http://www.mulberrytech.com/quickref/XSLTquickref.pdf>

Dr. Pawson XSLT FAQ

<http://www.dpawson.co.uk/xsl/xslfaq.html>

Zvon XSLT Tutorial

<http://zvon.org/xxl/XSLTutorial/Books/Book1/index.html>

- [Safari online XSLT Reference](#)

XSLT Syntax

Basic XSLT Elements

(From http://en.wikipedia.org/wiki/XSLT_elements)

xsl:stylesheet

(or `xsl:transform`) is the top-level element. Occurs only once in a stylesheet document. The attribute `version` specifies which XSLT version is being used. The NS declaration `xmlns:xsl` specifies the URL, which is always `http://www.w3.org/1999/XSL/Transform` regardless of the XSLT version.

xsl:output

Child element of stylesheet. It describes how data will be returned. The attribute `method` designates what kind of data is returned (such as `xml`, `text`, `html`). The attribute `omit-xml-declaration` indicates if the initial `<?xml` heading should be included. The attribute `encoding` designates the encoding used for output.

xsl:template

Specifies processing templates “match” is when the template should be used. “name” gives the template a name which `xsl:call-template` can use to call this template.

Declarations in xsl:stylesheet

xsl:param

parameter declarations (and their implicit value). Such parameters can then be set when calling XSLT processing, e.g. `java net.sf.saxon.Transform -o outfile.xml infile.xml style.xml -Dparam=paramvalue`

xsl:variable

similarly to parameters, it declares and initializes variables. They however cannot be set from outside. It should also be noted that XSLT (without processor-specific extension) is a pure

functional language, i.e. applications of templates do not have side effects → variables (or parameters) can be assigned just once, then just read!

XSLT Templates

Template (`xsl:template`) is a specification *which node* should be rewritten and how (into what):

- Which node is to be processed (i.e. rewritten), is described in the attribute `match`.
- The resulting fragment (into *what* it is rewritten) is stated in the body of the template.
- After processing of the source node, the processing continues at the nodes selected by `xsl:apply-templates select="_<xpath expression>_"`.

The template can also be explicitly *named* (named template) using the `name` attribute, in which case it can be called directly / explicitly using `xsl:call-template`.

Modularization

xsl:import

Retrieves another XSLT file addressed by the `href` (URI of the file). The templates in the linking (originating) stylesheet have priority over the imported ones.

xsl:include

Similarly, but works as a textual (verbatim) include, so no prioritization of the linking stylesheet is done.

Processing modes

- A template can specify a (processing) *mode* in which it can be activated.
- The mode is indicated using the `mode` attribute at the `xsl:template` element.
- Processing starts in *no mode* and
- can be switched into another mode by using the attribute `mode` in the `xsl:apply-templates` or `xsl:call-templates`.

Where to use processing modes?

Motivation

Modes allow a parallel set of templates with the same patterns match, but used for different purposes, for example:

- one set of templates for generating *table of contents* (index) from the document
- one for the *full text* of the document itself

Transformation Process in Detail

1. First, the processor selects the root (document node) as the *current node*, i.e. the node corresponding to the XPath expression /
2. Then, it finds the template matching it. If it is found, the template is used for processing this node.
3. Otherwise, a default (implicit) template is used to process the document node.
4. The processing recursively continues at nodes selected by the template that has been used for processing in the previous step.

Template Priority

- If there are two or more template matching, then ambiguity occurs and an error is emitted.
- This situation can be avoided by distinguishing the templates by setting their priorities using their priority attribute. The priority can be an integer, greater number means higher priority.
- Implicit templates have lower priority than explicit ones.

Template Invocation

Direct

xsl:call-template (possibly using xsl:with-params)

With implicit node selection

by applying a matching template (again may be with parameters) using xsl:apply-templates without explicit selection of nodes for further processing. Then, all child elements of the current (context) node will be selected and processed. Equivalent to xsl:apply-templates select="*".

With explicit node selection

when using xsl:apply-templates with explicit selection of nodes for further processing by using the select attribute.

- In general, the preferred way is to avoid *direct invocation* because the other ways correspond better to the functional nature of the XSLT language.

Outputting text nodes

Either:

- Type in text directly (as a literal) to output (body part) of the template. Be careful with the whitespace characters (spaces, CR/LF) as everything gets into the output.
- When the whitespace handling is important, eg. no unnecessary whitespaces should be produced, use the special element xsl:text. Inside of it, whitespaces are always maintained!

Implicit/Default templates

- Purpose: provide at least some standard "fallback" way to process basic structures such as traverse the document tree structure
- to "save typing" frequently used templates (ignoring comments and PI).
- These default templates have low priority and can be overridden by specifying explicit templates the same (or overlapping) match clause.
- The following default templates are implicitly "embedded" in each correct XSLT processor.

Default tree (do-nothing) traversal

```
<xsl:template match="*/">
  <xsl:apply-templates/>
</xsl:template>
```

- Selects any element and the root.
- Produces nothing for it but
- traverses its all child elements.

Default tree (do-nothing) traversal for specified mode

```
<xsl:template match="*/" mode="...">
  <xsl:apply-templates mode="..."/>
</xsl:template>
```

- Does the same but only for the specified mode.

Copy text nodes and attributes

```
<xsl:template match="text()|@">
  <xsl:value-of select="."/>
</xsl:template>
```

- Copies text nodes and attributes to the result

Ignore PIs and comments

```
<xsl:template match="processing-instruction()|comment()" />
```

- Ignores (does not include the results of the PI and comments)

Generating values programmatically

- Not only elements, attributes and texts from the source are copied to the output.
- All can be programmatically dynamically generated.

Generation of element with calculated attribute value

Objective

Generate the output of a predetermined element (with pre-known name), but with attributes with values with calculated during transformation.

Solution

Use the normal procedure - literal result element - attributes and values specified as the attribute value templates (AVT)

Example

Input

```
<link ref="a_link_href">
  ...
</link>
```

Template

```
<xsl:template match="link">
  <a href="#{@ref}"> ... </a>
</xsl: template>
```

Explanation

- Transforms the link to a (possibly HTML) a element, the href attribute value is composed of # and the value of the original ref attribute.

Output

```
<a href="#a_link_href"> ... </a>
```

Generating with calculated element- or attribute name

Objective

Generate the output element whose name, attributes and content is NOT known in advance when writing the style. So it must be determined (calculated) in runtime (when transforming).

Solution

Use a template to component `xsl:element`

Example

Input

```
<generate element-name="elt_name"> ... </generate>
```

Template

```
<xsl:template match="generate">
  <xsl:element name="{@element-name}">
    <xsl:attribute name="id">ID1</xsl:attribute>
  </xsl:element>
</xsl:template>
```

Result

Creates an element with the name `elt_name`, equipping it with the attribute `id="ID1"`. Also the attribute name could be generated if we wished so.

XSLT Conditional processing

Objective

To influence the output based on a condition.

Solution

Use branching in the template - either

- `xsl:if` for single *then/else* branches or
- multiway `xsl:choose` / `xsl:when` / `xsl:otherwise`

Example `xsl:if`

Input

```
<bread price="50"> ... </bread>
```

Template

```
<xsl:template match="bread">
  <p>
    <xsl:if test="@price > 30">
      <span class="expensive">Expensive </span>
    </xsl:if>bread - price <xsl:value-of select="@price"/> CZK</p>
</xsl:template>
```

Result

Creates an element p with a record about the bread. If the bread was expensive, also the "Expensive" indication is produced.

Example xsl:choose

Input

```
<bread price="12"> ... </bread>
<bread price="19"> ... </bread>
<bread price="30"> ... </bread>
```

Template

```
<xsl:template match="bread">
  <xsl:choose>
    <xsl:when test="@price > 30">
      <span class="expensive">Expensive</span>
    </xsl:when>
    <xsl:when test="@price < 10">
      <span class="strangely-cheap">Suspiciously cheap</span>
    </xsl:when>
    <xsl:otherwise>
      <span class="normal-price">Normal</span>
    </xsl:otherwise>
  </xsl:choose> bread - price <xsl:value-of select="@price"/> CZK
</xsl:template>
```

Result

Filters out the two extreme price level — normal prices remain for xsl:otherwise.

Loops

Input

```
<grocery>
  <bread price="12"> ... </bread>
  <bread price="19"> ... </bread>
  <bread price="30"> ... </bread>
</grocery>
```

Template

```
<xsl:template match="grocery">
  <xsl:for-each select="bread">
    <p> bread - price <xsl:value-of select="@price"/> CZK </p>
  </xsl:for-each>
</xsl:template>
```

Result

Creates series of elements p with bread prices.

Caution

Construction xsl:for-each typically has procedural nature, which is *generally not recommended* for XSLT as it namely gives minimum flexibility to iterate through the contents of a set of nodes — we must know its exact structures beforehand. The style is also more difficult to modify if the structure changes (eg. new or altered element names).

Template calls and parameters

Named template declaration

<xsl:template name="_thistemplatename_>. The template may contain declarations of parameters: <xsl:param name="_parametername_"/> (parameter type is not specified — i.e. dynamic typing)

Template call

using <xsl:call-template name="_atemplatename_"/> The call can also specify the parameters (if they were declared at the template definition): <xsl:with-param name="_parametername_" select="_parametervalueexpression_"/> or <xsl:with-param name="_parametername_">_parametervalue_</xsl:with-param>

Default parameter value

can also be specified using <xsl:param name="_parametername_" select="_defaultvalueexpression_"/>

Automatic (generated) numbering

- Achieved by using `xsl:number` element
- For either (or both): *counting elements* in input to allow automatic numbering — for example to number book chapters sequentially, or *formatting numbers*, eg. writing them in Arabic or

Roman numbers. Resembles part of the internationalization support seen in java.text.

- The autonumbering can also be *multilevel* eg. (sub)chapter numbers like 1.1 etc.

Example

- Plenty of variants shown in [XSLT Cookbook - Recipe of the Day](#)

Example style from this article

```
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
  <xsl:output method="text"/>
  <xsl:strip-space elements="*" />
  <xsl:template match="group">
    <xsl:text>Group </xsl:text>
    <xsl:number count="group" level="multiple" />
    <xsl:text>&#xa;</xsl:text>
    <xsl:apply-templates/>
  </xsl:template>
  <xsl:template match="person">
    <xsl:number count="group" level="multiple" format="1.1.1." />
    <xsl:number count="person" level="single" format="1 " />
    <xsl:value-of select="@name" />
    <xsl:text>&#xa;</xsl:text>
  </xsl:template>
</xsl:stylesheet>
```

Example (cont.)

Applied to (shortened)

```
<people>
  <group>
    <person name="Al Zehtoonney" age="33" sex="m" smoker="no" />
    <person name="Brad York" age="38" sex="m" smoker="yes" />
  </group>
  <group>
    <person name="Greg Sutter" age="40" sex="m" smoker="no" />
    <person name="Harry Rogers" age="37" sex="m" smoker="no" />
  </group>
  <group>
    <person name="John Quincy" age="43" sex="m" smoker="yes" />
    <person name="Kent Peterson" age="31" sex="m" smoker="no" />
  </group>
  <person name="John Frank" age="24" sex="m" smoker="no" />
</people>
```

Example (cont. 2)

Result

```
Group 1
1.1 Al Zehtoney
1.2 Brad York
Group 2
2.1 Greg Sutter
2.2 Harry Rogers
Group 2.1
2.1.1 John Quincy
2.1.2 Kent Peterson
2.3 John Frank
```

Namespace Handling

- XSLT allows to select and produce nodes (elements, attributes) in namespaces.
- However, it has some pitfalls, see [Namespaces in XSLT issues](#)
- [Multiple namespaces in XSLT](#)
- [What can be done with Namespaces in XSLT 2.0](#)
- [Avoid Namespace mistakes in XSLT](#) at Developerworks

Where to do XSLT?

- Online (just for fun)
- In all XML professional editors and many programmers' IDE such as *NetBeans*
- Command-line tools, such as `xsltproc` or `xmlstarlet`
- From within Java programs using Java Core API (`javax.xml.transform` package)
- Using specialized tools programmatically (via API) or command-line, such as *Saxon*
- Similarly for other languages, almost all now have XML/XSLT support

Online Tools

Good for simple try-and-see:

- [Online XSLT engine @Freeformatter.com](#) plus a demo XML input/XSLT/XML output
- [W3Schools online XSLT engine](#) maybe even better :-)

With XSLT 2.0 support:

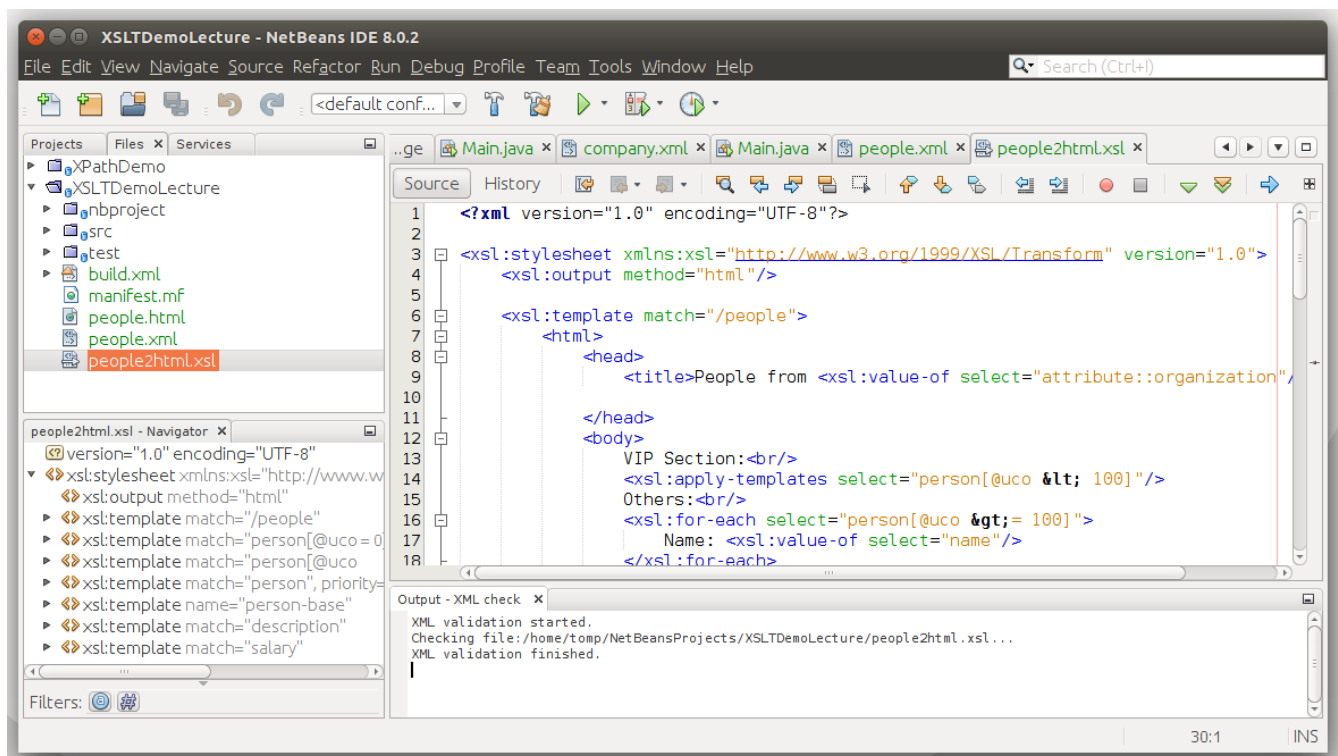
- <http://xslttest.appspot.com/>

XSLT in NetBeans - steps

All recent NetBeans version allow to launch XSLT just by:

- Opening an XSLT (or source XML) file
- Click on the blue right arrow on the toolbar right
- Specify source XML, XSLT, and output files
- Run the transformation
- Inspect the resulting file directly in the IDE

XSLT in NetBeans - screenshot



Using XSLT in Java (Core API)

- See [Using the XSLT Processor for Java](#) from Oracle