THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A Formal Semantics of Patterns in XSLT and XPath

Citation for published version:

Wadler, P 2000, 'A Formal Semantics of Patterns in XSLT and XPath', *Markup Languages*, vol. 2, no. 2, pp. 183-202. <https://doi.org/10.1162/10996620052104302>

Digital Object Identifier (DOI):

[10.1162/10996620052104302](https://doi.org/10.1162/10996620052104302)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Markup Languages

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A formal semantics of patterns in XSLT

Philip Wadler
Bell Labs, Lucent Technologies

March 29, 2000

1 Introduction

Semanticists should be obstetricians, not coroners of programming languages. — John Reynolds

Programming language theorists have developed a number of techniques to formally capture the meaning of programming languages. As our epigram insists, these theories are best applied not when a language is too old for help, but when it is aborning and subject to guidance.

XSLT is a language for transforming XML documents into other XML documents [XSLT]. A key element of XSLT is the sub-language of patterns, which is used for matching and selection. The pattern language of XSLT has recently evolved into XPath [XPath], a language of selection paths and expressions that performs core functions of both XSLT and XPointer [XPointer].

This note presents a formal semantics of the pattern language from the 16 December 1998 draft of XSLT [XSL-Dec98]. The semantics is clear and concise, summarizing in one page of formulas what required about ten pages of prose to describe. With the aid of the semantics one can rigorously state and prove properties of the language; these properties helped to guide future development of the XSLT design; thus fulfilling Reynolds's prescription.

The semantics was developed using standard techniques from the programming language community, and this article provides a tutorial introduction to these techniques. While little here will be new to the language theorist, some of what is here may be of use to the markup technologist.

A formal semantics brings to light issues that can be hard to spot in an english language description. For instance, matching is a central concept in XSLT. Here is how it was described in [XSL-Dec98].

The result of the *MatchExpr* is true if, for any node in the document that contains the context of the *MatchExpr*, the result of evaluating the *SelectExpr* with that node as context contains the context of the *MatchExpr*. Otherwise the result is false.

The attempt to make this sentence precise has rendered it almost unreadable. Even so, as we shall later see, the sentence is ambiguous, so even a careful reading can leave one unsure of the intention. In contrast, the formal specification given here presents the same information in just one line, is easier to read than the english, and avoids the ambiguity. Partly as a result of insight gleaned from the semantics, the definition of match patterns in XSLT was simplified and made easier to implement.

The formal semantics is given in a style known as *denotational semantics*. There are several textbook introductions to this subject, including those by Schmidt [Schmidt] and Allison [Allison]. We will be able to get by using some of the most basic ideas from semantics. Many of the tricky corners semantics arise from possibly infinite behaviour, such as a program that may enter an infinite loop. Since in our case we deal with finite documents and finite sets of nodes, these complexities can be avoided.

The formal semantics also draws upon techniques from the functional programming community. Again, there are several fine introductions, including those by Bird [Bird] and Paulson [Paulson]. The semantics was developed and debugged by transliterating it into the functional language Haskell [Haskell], and a copy of the Haskell program may be had by contacting the author. In related work, Haskell programs for manipulating XML have been developed by Wallace and Runciman [Wallace and Runciman].

The same techniques used here can be extended to give a denotational semantics of the entire XPath language, and such a semantics has been written. However, XPath is considerably more powerful than the pattern language of the December 1998 XSLT, and the semantics is correspondingly more complex. The semantics given here seems more appropriate for a gentle introduction.

This paper is organized as follows. Section 2 reviews the relevant mathematical theory of sets and relations, showing how it can be applied to model XML documents. Section 3 fleshes out the XML data model. Section 4 gives a tutorial introduction to denotational semantics, giving two simple semantics for arithmetic expressions. Section 5 gives a denotational semantics of patterns in the 16 December 1998 draft of XSLT. Section 6 draws some lessons from this semantics. Section 7 applies the semantics to prove some propositions about patterns. Section 8 concludes.

2 Sets and relations

We review here standard operations on sets and relations, and shows how we might use these to model XML documents and operations upon them.

Here is part of your family tree, represented as an XML document.

```
<Adam>
  <Cain>
    <Enoch/>
  </Cain>
  <Abel/>
  <Seth>
    <Enosh/>
  </Seth>
</Adam>
```

Let $\{ \text{Root}, \text{Adam}, \text{Cain}, \text{Enoch}, \text{Abel}, \text{Seth}, \text{Enosh} \}$ be the set of nodes in this tree, where Root is

a distinguished root node. The key relationships in the tree can be modeled with two functions.

$$\begin{aligned}
\textit{children} & : \textit{Node} \rightarrow \textit{Set}(\textit{Node}) \\
\textit{children}(\textit{Root}) & = \{ \textit{Adam} \} \\
\textit{children}(\textit{Adam}) & = \{ \textit{Cain}, \textit{Abel}, \textit{Seth} \} \\
\textit{children}(\textit{Cain}) & = \{ \textit{Enoch} \} \\
\textit{children}(\textit{Enoch}) & = \emptyset \\
\textit{children}(\textit{Abel}) & = \emptyset \\
\textit{children}(\textit{Seth}) & = \{ \textit{Enosh} \} \\
\textit{children}(\textit{Enosh}) & = \emptyset \\
\\
\textit{parent} & : \textit{Node} \rightarrow \textit{Set}_1(\textit{Node}) \\
\textit{parent}(\textit{Root}) & = \emptyset \\
\textit{parent}(\textit{Adam}) & = \{ \textit{Root} \} \\
\textit{parent}(\textit{Cain}) & = \{ \textit{Adam} \} \\
\textit{parent}(\textit{Enoch}) & = \{ \textit{Cain} \} \\
\textit{parent}(\textit{Abel}) & = \{ \textit{Adam} \} \\
\textit{parent}(\textit{Seth}) & = \{ \textit{Adam} \} \\
\textit{parent}(\textit{Enosh}) & = \{ \textit{Seth} \}
\end{aligned}$$

We write $\textit{Set}(A)$ for the type of a set where each element is of type A , and $\textit{Set}_1(A)$ for the subtype of sets with at most one element, and $\emptyset$ for the emptyset. We write $\in$ for membership, $\subseteq$ for subset, and $\cup$ for set union. For example, $\textit{Cain} \in \textit{children}(\textit{Adam})$, $\{ \textit{Cain}, \textit{Abel} \} \subseteq \textit{children}(\textit{Adam})$, and $\{ \textit{Cain}, \textit{Abel} \} \cup \{ \textit{Seth} \} = \textit{children}(\textit{Adam})$.

We will make extensive use of comprehension notation, as in the following example.

$$\begin{aligned}
\textit{siblings} & : \textit{Node} \rightarrow \textit{Set}(\textit{Node}) \\
\textit{siblings}(x) & = \{ z \mid y \in \textit{parent}(x), z \in \textit{children}(y) \}
\end{aligned}$$

This is read ‘the siblings of node x is the set of all nodes z such that some node y is a parent of x and z is a child of y ’. With this definition, we have $\textit{siblings}(\textit{Cain}) = \{ \textit{Cain}, \textit{Abel}, \textit{Seth} \}$. (Rather than ‘Am I my brother’s keeper?’, the apposite question here is ‘Am I my brother?’)

To further demonstrate the power of the comprehension notation, here is a definition that corrects the shortcoming of the first.

$$\begin{aligned}
\textit{properSiblings} & : \textit{Node} \rightarrow \textit{Set}(\textit{Node}) \\
\textit{properSiblings}(x) & = \{ y \mid y \in \textit{siblings}(x), x \neq y \}
\end{aligned}$$

This is read ‘the proper siblings of node x is the set of all nodes y such that y is a sibling of x , and y is distinct from x ’. With this definition, we have $\textit{properSiblings}(\textit{Cain}) = \{ \textit{Abel}, \textit{Seth} \}$.

Here is another example.

$$\begin{aligned}
\textit{grandparent} & : \textit{Node} \rightarrow \textit{Set}_1(\textit{Node}) \\
\textit{grandparent}(x) & = \{ z \mid y \in \textit{parent}(x), z \in \textit{parent}(y) \}
\end{aligned}$$

This is read ‘the grandparent of node x is the set of all nodes z such that some node y is a parent of x , and z is a parent of y ’. Thus, $\textit{grandparent}(\textit{Enosh}) = \{ \textit{Adam} \}$ and $\textit{grandparent}(\textit{Adam}) = \emptyset$. Note how this definition deals neatly with the case of a node that has no grandparent — instead of an error, one simply gets an empty set.

We represent a relation on nodes by a function $r : \textit{Node} \rightarrow \textit{Set}(\textit{Node})$. The functions $\textit{children}$, $\textit{parent}$, $\textit{siblings}$, and $\textit{grandparent}$ are all examples of relations. If r_1, r_2 are two relations, we write $r_1 ; r_2$ for their composition, defined as follows.

$$(r_1 ; r_2)(x) = \{ z \mid y \in r_1(x), z \in r_2(y) \}$$

Thus, relation $r_1 ; r_2$ relates node x to node z if relation r_1 relates x to some y , and relation r_2 relates y to some z . For example, $siblings = parent ; children$ and $grandparent = parent ; parent$.

The identity relation *self* is the one that relates a node solely to itself.

$$self(x) = \{x\}$$

Observe that the identity relation is indeed a left and right identity for composition, that is $r ; self = r = self ; r$. Further, composition is associative, so $(r_1 ; r_2) ; r_3 = r_1 ; (r_2 ; r_3)$.

If r is a relation, we define the transitive closure r^+ and the reflexive and transitive closure r^* in the usual way.

$$\begin{aligned} r^+(x) &= \{z \mid y \in r(x), z \in r^*(y)\} \\ r^*(x) &= \{x\} \cup r^+(x) \end{aligned}$$

Thus $children^+(x)$ returns all proper descendants of x (not including itself), and $children^*(x)$ returns all descendants of x (including itself), and $parent^+(x)$ returns all proper ancestors of x . (We use the same symbols as with regular expressions, because in both cases $+$ corresponds one or more repetitions, and $*$ corresponds to zero or more repetitions.)

Recall that the document order is a total order on all the nodes in a document; for element nodes, it is the order of their start tags. We write

$$first, last : Set(Node) \rightarrow Set_1(Node)$$

for the functions that take a set and return the singleton set containing just its first or last node in document order, or return the empty set if the given set is empty. For instance, $first(children(Adam)) = Cain$, $last(children(Adam)) = Seth$, and $first(children(Abel)) = \emptyset$.

Having established these preliminaries, we can now proceed to establish a more complete formal model for XML documents.

3 A data model for XML

This section presents a mathematical data model for XML. It is based on the data model given in Section 2.6 of [XSL-Dec98], which is about 4 pages long. Even this simple model revealed some small ambiguities in [XPath], which we discuss at the end of this section. For simplicity, we will ignore namespaces; although a full treatment of them along similar lines is possible.

The model we give is, unsurprisingly, quite similar to the Document Object Model (DOM), a standard interface for accessing and manipulating representations of XML trees, which works with a number of programming languages [DOM]. There are a few differences from the DOM (which we will note parenthetically).

The basic datatype is *Node*. Each *Node* is one of six kinds: root, element, attribute, text, comment, or processing instruction. We indicate this with six boolean functions; for each node, exactly one of these functions yields true.

$$\begin{aligned} isRoot & : Node \rightarrow Boolean \\ isElement & : Node \rightarrow Boolean \\ isAttribute & : Node \rightarrow Boolean \\ isText & : Node \rightarrow Boolean \\ isComment & : Node \rightarrow Boolean \\ isPI & : Node \rightarrow Boolean \end{aligned}$$

(DOM differs in that root nodes are called document nodes, and instead of six boolean functions there is a single *nodetype* function.)

We have the following functions that relate nodes to other nodes in the document.

$$\begin{aligned}
parent & : Node \rightarrow Set_1(Node) \\
children & : Node \rightarrow Set(Node) \\
attributes & : Node \rightarrow Set(Node) \\
root & : Node \rightarrow Node
\end{aligned}$$

Function *parent* returns the parent of the node, or returns the empty set if the node is the root. Function *children* returns the set of children of a node, and is empty if the node has no children. Attribute, text, comment, and PI nodes never have children. An element node may also have related attribute nodes, these are not counted as children, but instead are returned by *attributes*. Function *root* returns the root node of the document.

These functions are related by various laws. Only a root or element can have children, and only an element can have attribute nodes.

$$\begin{aligned}
children(x) \neq \emptyset & \text{ implies } isRoot(x) \vee isElement(x) \\
attributes(x) \neq \emptyset & \text{ implies } isElement(x)
\end{aligned}$$

Every child is an element, text, comment, or processing instruction. Unsurprisingly, the attributes of a node are attribute nodes, and the root of a node is the root node.

$$\begin{aligned}
y \in children(x) & \text{ implies } isElement(y) \vee isText(y) \vee isComment(y) \vee isPI(y) \\
y \in attributes(x) & \text{ implies } isAttribute(y) \\
y = root(x) & \text{ implies } isRoot(y)
\end{aligned}$$

Define the subnodes of a node to be the union of its children and attributes.

$$\begin{aligned}
subnodes & : Node \rightarrow Set(Node) \\
subnodes(x) & : children(x) \cup attributes(x)
\end{aligned}$$

Then one node is the subnode of another exactly when the parent of that node is the other.

$$y \in subnodes(x) \text{ if and only if } parent(y) = \{x\}$$

Every node in the document is descended via the subnode relation from the root.

$$x \in subnodes^*(root(x))$$

(DOM differs in that attribute nodes are considered not to have any parent.)

Other functions on a node access content, with the meaning changing for different node types.

$$\begin{aligned}
name & : Node \rightarrow String \\
value & : Node \rightarrow String
\end{aligned}$$

The meanings are best given by a table.

	<i>name</i>	<i>value</i>
<i>Root</i>	empty	value of children
<i>Element</i>	tag name	value of children
<i>Attribute</i>	attribute name	attribute value
<i>Text</i>	empty	content of text node
<i>Comment</i>	empty	content of the comment
<i>PI</i>	target	content excluding target

Here ‘empty’ denotes the empty string, and ‘value of children’ means the result of concatenating the values of the children nodes, taken in ascending document order. (DOM has a similar table, but

the name of root, text, and comment nodes is taken to be '`#document`', '`#text`', and '`#comment`', respectively, rather than empty; while the value of root and element nodes is taken to be empty, rather than the concatenation of the values of the children.)

Recall that the DTD of an XML document may declare an element to have an ID attribute; a string uniquely identifies the element in the document with that string as its ID attribute. To model this, we introduce the function *id*, which takes a string and returns the singleton set containing the corresponding element, or returns the empty set if no such element exists. (If we were dealing with multiple documents, then *id* should have an additional parameter to specify which document.) Sometimes multiple ids are contained in a single attribute. The function *split* takes a string into the set of space-separated identifiers within it.

$$\begin{aligned} id & : \text{String} \rightarrow \text{Set}_1(\text{Node}) \\ split & : \text{String} \rightarrow \text{Set}(\text{String}) \end{aligned}$$

(There is nothing in DOM that corresponds to these functions. By the way, note that XSL allows any string to be taken as an identifier, not just strings in 'idref' attributes.)

The above model already reveals some small ambiguities in [XPath], which does not say what the name of a root, text, or comment node should be. In fact, most readers of the document guess (correctly) that they should be empty, but this is not entirely obvious since the DOM makes a different definition. Indeed, when this point was brought to the attention of the XSL committee, some members argued that the DOM definition was better, although (at this writing) the definition given here appears to prevail. At any event, the act of writing a formal model revealed the ambiguity, which should now be eliminated in the next XPath draft.

We now have all the tools we need to write the denotational semantics, so the next step is to review how such semantics are constructed.

4 Denotational semantics

This section introduces the standard techniques of denotational semantics, by giving a trivial semantics for arithmetic expressions, and a simple semantics for expressions with variables. Anyone already familiar with denotational semantics may skip (or skim) this section; anyone not familiar may wish to consult standard texts like [Schmidt] or [Allison] for further details.

The starting point for *semantics* is *syntax* — we require expressions to which meaning can be attached. Fortunately, the state of the art for dealing with syntax is far superior to that for dealing with semantics, as indicated by the widespread use of formal notations such as BNF and regular expressions. Imagine what it would be like if we were forced to resort to english to describe the syntax of our programs as well as their semantics!

Much of the complication in syntax arises from the need to establish precedence of operators. For instance, here is a typical *concrete syntax* for arithmetic expressions.

$$\begin{aligned} expr & ::= \begin{array}{l} expr + term \\ | \\ expr - term \\ | \\ term \end{array} \\ term & ::= \begin{array}{l} term * primary \\ | \\ term / primary \\ | \\ primary \end{array} \\ primary & ::= \begin{array}{l} number \\ | \\ (expr) \end{array} \end{aligned}$$

This syntax specifies that addition binds less tightly than multiplication, and that all the operations are right associative, so that $1+2*3+4$ is parsed the same as $((1+(2*3))+4)$.

The purpose of the concrete syntax is to specify a mapping from strings to parse trees, and it is to the parse trees rather than the strings that we wish to attach a meaning. Therefore, we take as our starting point for semantics an *abstract syntax* that corresponds directly to the underlying parse tree. Here is an abstract syntax for expressions.

$$\begin{array}{lcl} i & : & \text{Number} \\ e & : & \text{Expr} \quad e ::= e + e \mid e - e \mid e * e \mid e / e \mid i \mid (e) \end{array}$$

Here for convenience we use e as shorthand for *expr*, *term*, and *primary*, and i as shorthand for *number*. We write *Expr* as the type of the parse trees for expressions. Numbers are sufficiently well understood that we don't bother to distinguish between the syntax and semantics, and take *Number* as both the type used to represent numbers in the parse tree and the type used for the value of expressions. If we want to be precise, we might say that *Number* corresponds to the 64 bit floating point numbers in IEEE 754 format (which is how numbers are represented in XPath).

All information about precedence and associativity of operators has been removed from the abstract syntax above. Some authors of denotational semantics would go even further and remove the case for parentheses, which serves only to identify grouping and have no role in the underlying parse tree.

The denotational semantics of expressions is a function $\mathcal{E}$ from expressions to numbers.

$$\begin{array}{lcl} \mathcal{E} & : & \text{Expr} \rightarrow \text{Number} \\ \mathcal{E}[e_1 + e_2] & = & \mathcal{E}[e_1] + \mathcal{E}[e_2] \\ \mathcal{E}[e_1 - e_2] & = & \mathcal{E}[e_1] - \mathcal{E}[e_2] \\ \mathcal{E}[e_1 * e_2] & = & \mathcal{E}[e_1] \times \mathcal{E}[e_2] \\ \mathcal{E}[e_1 / e_2] & = & \mathcal{E}[e_1] \div \mathcal{E}[e_2] \\ \mathcal{E}[i] & = & i \\ \mathcal{E}[(e)] & = & \mathcal{E}[e] \end{array}$$

We write $\mathcal{E}[e]$ to denote the meaning of the expression e , where by convention syntax is distinguished from semantics by writing syntactic entities inside special semantic brackets. The meaning of each expression is given in terms of the meanings of its subexpressions. Thus, the first line says that the value of the expression $e_1 + e_2$ is computed by finding the value of expression e_1 and the value of expression e_2 and adding the results.

The semantics is pretty trivial, amounting to little more than a change of symbols. But that is because we have chosen to explain arithmetic in terms of arithmetic, so there is little useful work to do.

For a less trivial example, consider a semantics for expressions augmented with variable names. Here is the new abstract syntax.

$$\begin{array}{lcl} i & : & \text{Number} \\ n & : & \text{Name} \\ e & : & \text{Expr} \quad e ::= n \mid \text{let } n = e \text{ in } e \mid \\ & & e + e \mid e - e \mid e * e \mid e / e \mid i \mid (e) \end{array}$$

The grammar for expressions has two new lines, to define variables and bindings. For example, `let a=1+2 in a*a` evaluates to 9.

The syntax is given in terms of an environment, which maps variables to their value. We let ρ range over environments.

$$\rho : \text{Environment} = \text{Name} \rightarrow \text{Number}$$

We need two operations on environments. The first takes an environment and a name and returns the corresponding number. The second takes a name, a number, and an environment, and returns a new environment that is the same as the first, but also maps the given name to the given number.

$$\begin{aligned}
\text{lookup} & : (Name, Environment) \rightarrow Number \\
\text{lookup}(v, \rho) & = \rho(v) \\
\\
\text{extend} & : (Name, Number, Environment) \rightarrow Environment \\
\text{extend}(v, i, \rho) & = \rho_1 \text{ where } \rho_1(v_1) = \text{if } (v_1 = v) \text{ then } i \text{ else } \rho(v_1)
\end{aligned}$$

These definitions exhibit the utility of *higher-order* functions. Both *lookup* and *extend* take a function (ρ) as an argument, and *extend* also returns a function (ρ_1) as its result. For example, let ρ_0 be the environment that maps every name to the IEEE 754 floating point number NaN (Not a Number). Then

$$\rho_1 = \text{extend}(\mathbf{b}, 4, \text{extend}(\mathbf{a}, 3, \rho_0))$$

is the environment such that $\text{lookup}(\mathbf{a}, \rho_1) = 3$. According to the definition, a new binding may *shadow* the old binding of a variable. For example, if $\rho_2 = \text{extend}(\mathbf{a}, 5, \rho_1)$ then $\text{lookup}(\mathbf{a}, \rho_2) = 5$, not 3.

The augmented denotational semantics takes an expression and an environment and returns a number.

$$\begin{aligned}
\mathcal{E} & : Expr \rightarrow Environment \rightarrow Number \\
\mathcal{E}[\mathbf{v}]\rho & = \text{lookup}(\mathbf{v}, \rho) \\
\mathcal{E}[\mathbf{let } v = e_1 \text{ in } e_2]\rho & = \mathcal{E}[e_2](\text{extend}(\mathbf{v}, \mathcal{E}[e_1]\rho, \rho)) \\
\mathcal{E}[e_1 + e_2]\rho & = \mathcal{E}[e_1]\rho + \mathcal{E}[e_2]\rho \\
\mathcal{E}[e_1 - e_2]\rho & = \mathcal{E}[e_1]\rho - \mathcal{E}[e_2]\rho \\
\mathcal{E}[e_1 * e_2]\rho & = \mathcal{E}[e_1]\rho \times \mathcal{E}[e_2]\rho \\
\mathcal{E}[e_1 / e_2]\rho & = \mathcal{E}[e_1]\rho \div \mathcal{E}[e_2]\rho \\
\mathcal{E}[\mathbf{i}]\rho & = i \\
\mathcal{E}[(e)]\rho & = \mathcal{E}[e]\rho
\end{aligned}$$

Now we write $\mathcal{E}[e]\rho$ for the meaning of expression e in environment ρ . Again, the meaning of each expression is given in terms of the meaning of its components. The first line says that the value of name n in environment ρ is given by looking up the name in the environment. The second line says that the value of the expression $\mathbf{let } n = e_1 \text{ in } e_2$ in environment ρ is given by evaluating e_1 in environment ρ , extending environment ρ by binding name n to the value just computed, then evaluating e_2 in the new environment. The rest is much as before, passing the environment unchanged to each subexpression. (If you find this paragraph hard to read, you may find it easier to just look at the formulas — that is the point of using a formalism!)

By convention, the semantic function $\mathcal{E}$ is given a higher-order type: it takes an expression, and returns a function from environments to values. Thus, for instance, the meaning of $\mathcal{E}[\mathbf{a+b}]$ is the function $f : Environment \rightarrow Number$ such that $f(\rho) = \text{lookup}(\mathbf{a}, \rho) + \text{lookup}(\mathbf{b}, \rho)$. Also by convention the parentheses may be dropped around the second argument, so we write $\mathcal{E}[e]\rho$ rather than $\mathcal{E}[e](\rho)$.

The semantics we have given answers some subtle questions about the meaning of expressions. For instance, consider the expression

`let a=3 in (let a=4 in a+a)+a`

Is the value of this expression 11 or 12? In a language like C, the value of the similar expression `(a=4, a+a)+a` is 12. But it should be clear from looking at the semantics that the value of the above expression is 11, because the rightmost `a` is evaluated in a different environment than the other two occurrences of `a`. (To learn how to give a denotational semantics for an imperative language like C, consult any standard text on semantics, such as [Schmidt] or [Allison].)

This concludes our semantics tutorial. We are now ready to look at a denotational semantics for part of XSL.

5 A semantics for patterns

This section presents the semantics of matching and selection in the 16 December 1998 draft of XSLT [XSL-Dec98]. What we describe here corresponds to Section 2.6 of that draft. That section is 11 pages long (including two pages for concrete syntax), as compared to the one page of text and one page of formulas given here (excluding concrete but including abstract syntax).

The abstract syntax of patterns and qualifiers is as follows.

$$\begin{aligned}
 n & : \text{Name} \\
 s & : \text{String} \\
 p & : \text{Pattern} \quad p ::= p_1 | p_2 \mid /p \mid //p \mid p_1/p_2 \mid p_1//p_2 \mid p[q] \mid \\
 & \quad n \mid * \mid @n \mid @* \mid \text{text}() \mid \text{comment}() \mid \text{pi}(n) \mid \text{pi}() \mid \text{id}(p) \mid \text{id}(s) \mid \\
 & \quad \text{ancestor}(p) \mid \text{ancestor-or-self}(p) \mid . \mid \dots \\
 q & : \text{Qualifier} \quad q ::= q_1 \text{ and } q_2 \mid q_1 \text{ or } q_2 \mid \text{not}(q) \mid (q) \\
 & \quad \text{first-of-type}() \mid \text{last-of-type}() \\
 & \quad \text{first-of-any}() \mid \text{last-of-any}() \mid p=s \mid p
 \end{aligned}$$

The abstract syntax is much more permissive than the corresponding concrete syntax. For instance, in the concrete syntax the pattern union operator $p_1 | p_2$ is constrained to appear at the top level of a pattern, while here we allow it at any level in the syntax tree. As we will see, there is no problem whatsoever in giving a syntax for the more general language, and the XSL document does not explain why the syntax is restricted, though presumably the designers believed the restrictions make the language easier and more efficient to implement.

The semantics is given in Figure 1. A pattern may be used either for matching or selection. We write $\mathcal{M}[p]x$ to denote whether the pattern p matches against the node x , we write $\mathcal{S}[p]x$ to denote the set of nodes selected by the pattern p when the context node is x , and we write $\mathcal{Q}[q]x$ to denote whether the qualifier q is satisfied when the context node is x . Thus, the semantic functions have the following types.

$$\begin{aligned}
 \mathcal{M} & : \text{Pattern} \rightarrow \text{Node} \rightarrow \text{Boolean} \\
 \mathcal{S} & : \text{Pattern} \rightarrow \text{Node} \rightarrow \text{Set}(\text{Node}) \\
 \mathcal{Q} & : \text{Qualifier} \rightarrow \text{Node} \rightarrow \text{Boolean}
 \end{aligned}$$

Again, the semantics of each pattern is defined in terms of the semantics of its subpatterns, and similarly for quantifiers.

We give a reading for a few lines in the definition. The first line of $\mathcal{S}$ says that the pattern $p_1 | p_2$ with current node x selects the union of the nodes selected by p_1 and p_2 , each with current node x . The second line says that the pattern $/p$ with current node x selects the same nodes as selected by the pattern p with current node $\text{root}(x)$, the root of the document containing node x . The fourth line says that the pattern p_1/p_2 with current node x is the set of all nodes x_2 selected as follows: let x_1 be any node selected by p_1 with current node x , and then let x_2 be any node selected by p_2 with current node x_1 . The seventh line says that the name n with current node x selects all subnodes of x that are elements and have the name n . The eighth line says that the pattern $*$ with current node x selects all the subnodes of x that are elements. The ninth line says that the pattern $@n$ with current node x selects all subnodes of x that are attributes and have the name n . The penultimate line says that the pattern $.$ with current node x selects the set of nodes containing just x . The readings for the other formulas are similar. (Again, if you find this paragraph hard to read, you may prefer to just look at the formulas — and, again, that is the point!)

If p is a pattern, then $\mathcal{S}[p] : \text{Node} \rightarrow \text{Set}(\text{Node})$ is a relation. In other words, the meaning of a pattern is a relation between the current node and all the nodes the pattern selects. Many of the

operations on patterns can be re-expressed directly as operations on relations. In particular, it is easy to check that p_1/p_2 is relational composition and $.$ is the identity relation.

$$\begin{aligned}\mathcal{S}[p_1/p_2] &= \mathcal{S}[p_1] ; \mathcal{S}[p_2] \\ \mathcal{S}[.] &= \textit{self}\end{aligned}$$

It follows immediately that the patterns $p/.$ and p and $./p$ are all equivalent.

The semantics clarifies a number of points that may not be obvious on a first reading of the XSL draft specification [XSL-Dec98]. We'll explore some of these in the next section, and see how the formal semantics aided the development of XSL.

6 Lessons from the semantics

Now that we've defined the semantics, what can we do with it? In this section, we make use of the semantics to illustrate some subtle points of the XSL design. And we will show how the semantics helped us to obey Reynold's dictum — rather than merely cataloguing a dead design, insight from the semantics contributed to the growth of a living language.

We begin by examining some key features of the structure of the semantics. For purpose of symmetry, it is written entirely using the *subnodes* function defined in Section 3, and never uses *children* or *attributes* directly. Note that *subnodes* does not appear in the definition of p_1/p_2 , but instead appears in the definition of n , $*$, $@n$, $@*$, `text()`, `comment()`, `pi()`, and `pi(n)`. It is very tempting to give a different definition where *subnodes* appears only once in the definition of p_1/p_2 and does not appear in the definitions of n , $*$, $@n$, $@*$, and the like. However, that definition does not work — there is then no possible way to define the $.$ pattern! One consequence of this choice is that the element selection patterns n and $*$ always step down to a subnode, and therefore it is crucial that our data model defines a root node which has the document element as a subnode.

These architectural points are not obvious from the english draft of the specification. Discussions reveal that people often think of p_1/p_2 as performing a subnode step; that they don't think of n and $*$ as performing a subnode step; or that they don't understand why the root is different from the document element. Wallace and Runciman [Wallace and Runciman] have even published a set of Haskell functions for XSL-style pattern matching that makes the opposite choice, where the operation corresponding to p_1/p_2 does perform the subnode step, and hence no operation corresponding to $.$ can be defined.

Another subtlety is that there are two quite distinct ways in which a pattern is converted to a boolean — the definition of $\mathcal{M}[p]x$ is quite different from the definition of $\mathcal{Q}[p]x$. The latter states that a pattern p acting as a qualifier is satisfied in context node x if selecting from pattern p with context node x yields a non-empty set of nodes. The former is discussed in detail next.

Recall the description of matching in [XSL-Dec98], quoted in the introduction.

The result of the *MatchExpr* is true if, for any node in the document that contains the context of the *MatchExpr*, the result of evaluating the *SelectExpr* with that node as context contains the context of the *MatchExpr*. Otherwise the result is false.

This is summarized in the one line describing match patterns, which says that a pattern p matches in a context x if there is *any* node x_1 in the document such that selecting with pattern p in context x_1 yields a set that contains the original node x . (Since it is allowed to mention variables, the above sentence is perhaps a little clearer than the XML document, and arguably the one-line formula is clearer still.)

Note that the original paragraph is ambiguous. Should the phrase 'any node in the document that contains the context of the *MatchExpr*' be read as 'any node in (the document that contains the context of the *MatchExpr*)', meaning any node in the document whatsoever? Or should it be read

as (any node in the document) that contains the context of the *MatchExpr*', meaning any node in the document that is an ancestor of (i.e., that contains) the context node of the *MatchExpr*? As we can see from the formal specification the intended meaning is the former, since *subnodes*^{*}(*root*(*x*)) generates all nodes in the document containing the context node *x*. The other interpretation could be formalized by replacing that expression by *parent*^{*}(*x*), which generates all ancestors of *x*.

Matching against a pattern as defined in this way can potentially be quite expensive. For instance, consider the following document, where the *id* attribute is declared in the DTD to be unique identifiers as used in the *id*() pattern.

```
<doc>
  <cite>ref1</cite>
  <cite>ref2</cite>
  <reference id='ref1'/>reference one</reference>
  <reference id='ref2'/>reference two</reference>
  <reference id='ref3'/>reference three</reference>
</doc>
```

The pattern *id*(*cite*) will match against any node that is cited in the document. But this means that when checking a reference node for a match, one must scan the entire document to check for a citation; the document may be rather larger than seven lines, so this could be expensive. To avoid repeatedly scanning the document, one might be clever and try to build a reverse index, showing which nodes might be referenced by which other pieces of text, but this won't be easy. (Especially since XSL does not guarantee that *id* will be contained in an *idref* attribute, so the entire text must be indexed.) Furthermore, users are unlikely to exploit this sort of pattern, so the effort of implementation is not commensurate with the benefit. Indeed, it turns out that neither XT nor Microsoft IE5 properly implemented this aspect of the December 1998 XSL draft.

Partly as a result of writing the formal semantics, this problem came to light. It was proposed to solve the problem by restricting the patterns that could be used as match patterns to a subset of those that could be used as select patterns. In particular, only *id*(*s*) can be used as a match pattern, yielding nodes identified by a literal string *s*, but *id*(*p*) cannot. Other restrictions were to outlaw the use of *ancestor*(*p*), *ancestor-or-self*(*p*), and *..* in match patterns, all of which were thought to be of little use, hard to implement, and confusing.

One consequence of these restrictions is that the ambiguity noted previously becomes irrelevant: matching from any node in the document becomes equivalent to matching from any ancestor of the context node! This isn't just an idle guess, in the next section we will see how we can state and prove this as a formal proposition.

One benefit of a formal semantics is that it becomes easier to prove equivalences between patterns. We have already observed that *p*/. and *./p* and *p* are all equivalent patterns. More interestingly, one can also show that *p*₁/*id*(*p*₂) is equivalent to *id*(*p*₁/*p*₂). Once this equivalence was formulated, it was decided that it was confusing to have two forms of expression that do the same thing, so patterns were further restricted to eliminate the first form in favor of the second.

7 Some propositions about patterns

When it comes to formalism, the proof of the pudding is in, well, the proof. Here we state and prove some propositions about patterns referred to in the preceding section.

We previously asserted that for a restricted notion of match pattern, matching against any node in the document is equivalent to matching against any ancestor. To formally state this, we first

define a second semantics of match patterns, which matches against ancestors rather than against any node.

$$\begin{aligned} \mathcal{M}' & : \text{Pattern} \times \text{Node} \rightarrow \text{Boolean} \\ \mathcal{M}'[p]x & = x \in \{x_2 \mid x_1 \in \text{parent}^*(x), x_2 \in \mathcal{S}[p](x_1)\} \end{aligned}$$

Then we have a proposition asserting the two semantics are equivalent for the restricted patterns.

Proposition. For any pattern p (not containing the forms $\text{id}(p)$, $\text{ancestor}(p)$, $\text{ancestor-or-self}(p)$, or $\dots$) and any node x we have

$$\mathcal{M}[p]x = \mathcal{M}'[p]x.$$

The proof requires two lemmas. With the restricted patterns removed, every remaining pattern is either absolute (that is ignores the context node) or descending (that is, any nodes it yields are descendants of the context node).

Lemma. If p is a pattern of the form $/p_1$ or $\text{id}(s)$, then

$$\mathcal{S}[p]x = \mathcal{S}[p](\text{root}(x))$$

Lemma. For any pattern p (not containing the forms $/p$, $\text{id}(s)$, $\text{ancestor}(p)$, $\text{ancestor-or-self}(p)$, or $\dots$) and any node x we have

$$\mathcal{S}[p]x \subseteq \text{subnodes}^*(x)$$

The first lemma is proved by case analysis on the two possibilities, the second lemma is proved by case analysis on the remaining possibilities and structural induction. The proposition then follows easily, since $y \in \text{subnodes}^*(x)$ is equivalent to $x \in \text{parent}^*(y)$. QED.

We skip the details of the proof — further explanation of the relevant techniques can be found in any good text on functional programming or denotational semantics, such as [Bird], [Paulson], [Schmidt], and [Allison]. However, just to give the flavor, the next example does present one simple proof in detail.

As we mentioned earlier, the patterns $p_1/\text{id}(p_2)$ and $\text{id}(p_1/p_2)$ are equivalent. Here we give the proof in some detail. We begin with the special case where p_2 is taken to be the $\cdot$ pattern.

Proposition. For any pattern p and for any node x , we have

$$\mathcal{S}[p/\text{id}(\cdot)]x = \mathcal{S}[\text{id}(p)]x$$

The proof is a simple calculation, expanding definitions and simplifying.

$$\begin{aligned} & \mathcal{S}[p/\text{id}(\cdot)]x \\ = & \text{definition of } \mathcal{S}[p_1/p_2] \\ & \{x_2 \mid x_1 \in \mathcal{S}[p]x, x_2 \in \mathcal{S}[\text{id}(\cdot)]x_1\} \\ = & \text{definition of } \mathcal{S}[\text{id}(p)] \\ & \{x_2 \mid x_1 \in \mathcal{S}[p]x, x_2 \in \{x_4 \mid x_3 \in \mathcal{S}[\cdot]x_1, s \in \text{split}(\text{value}(x_3)), x_4 \in \text{id}(s)\}\} \\ = & \text{definition of } \mathcal{S}[\cdot] \\ & \{x_2 \mid x_1 \in \mathcal{S}[p]x, x_2 \in \{x_4 \mid x_3 \in \{x_1\}, s \in \text{split}(\text{value}(x_3)), x_4 \in \text{id}(s)\}\} \\ = & \text{simplify} \\ & \{x_2 \mid x_1 \in \mathcal{S}[p]x, x_2 \in \{x_4 \mid s \in \text{split}(\text{value}(x_1)), x_4 \in \text{id}(s)\}\} \\ = & \text{simplify} \\ & \{x_4 \mid x_1 \in \mathcal{S}[p]x, s \in \text{split}(\text{value}(x_1)), x_4 \in \text{id}(s)\} \\ = & \text{definition of } \mathcal{S}[\text{id}(p)] \\ & \mathcal{S}[\text{id}(p)]x \end{aligned}$$

(This style of proof is widely used, see [Bird] for examples of its power.)

Using the fact that p_1/p_2 corresponds to relational composition, we can rewrite the above proposition as follows.

$$\mathcal{S}[\text{id}(p)] = \mathcal{S}[p] ; \mathcal{S}[\text{id}(.)]$$

Using the fact that relational composition is associative, it is then easy to prove the more general result.

Proposition. For any patterns p_1 and p_2 and for any node x , we have

$$\mathcal{S}[p_1/\text{id}(p_2)]x = \mathcal{S}[\text{id}(p_1/p_2)]x$$

Here is the calculation, spelled out in some detail.

$$\begin{aligned} & \mathcal{S}[p_1/\text{id}(p_2)] \\ = & \quad / \text{ is relational composition} \\ & \mathcal{S}[p_1] ; \mathcal{S}[\text{id}(p_2)] \\ = & \quad \text{previous proposition} \\ & \mathcal{S}[p_1] ; (\mathcal{S}[p_2] ; \mathcal{S}[\text{id}(.)]) \\ = & \quad \text{relational composition is associative} \\ & (\mathcal{S}[p_1] ; \mathcal{S}[p_2]) ; \mathcal{S}[\text{id}(.)] \\ = & \quad / \text{ is relational composition} \\ & \mathcal{S}[p_1/p_2] ; \mathcal{S}[\text{id}(.)] \\ = & \quad \text{previous proposition} \\ & \mathcal{S}[\text{id}(p_1/p_2)] \end{aligned}$$

QED.

8 Conclusions

We have seen that a formal semantics can be more concise and readable than english; that it lets one formulate and prove precise statements about the meaning of programs; and that such statements can aid further development of a language.

The techniques described in this paper have been applied to derive a denotational semantics for the entire XPath language. Unfortunately, that semantics is significantly more complex than the one given here. Will the insight gleaned from the new semantics aid in further simplification of XPath? Only time (and the workings of W3C committees) will tell.

Acknowledgements James Clark provided much useful feedback on the formal semantics, as well as providing an excellent informal semantics in the XSLT draft documents. Catherine Lyons made several suggestions that improved the presentation of this paper. Matias Giovannini and Dan Connolly spotted a number of typos.

References

- [Allison] Lloyd Allison, *A Practical Introduction to Denotational Semantics*, Cambridge University Press, 1987.
- [Bird] Richard Bird, *Introduction to Functional Programming, 2nd edition*, Addison-Wesley, 1998.

- [DOM] Lauren Wood, editor, *Document Object Model (DOM)*, Version 1.0, W3C Recommendation, 1 October 1998. <http://www.w3.org/TR/REC-DOM-Level-1/>
- [Haskell] The Haskell home page, www.haskell.org.
- [Paulson] Larry Paulson, *ML for the Working Programmer, 2nd edition*, Cambridge University Press, 1998.
- [Schmidt] David A. Schmidt, *The Structure of Typed Programming Languages*, MIT Press, Cambridge, MA, 367 pages, 1994.
- [Wallace and Runciman] Malcolm Wallace and Colin Runciman, Haskell and XML: Generic Combinators or Type-Based Translation? *4'th International Conference on Functional Programming (ICFP 99)*, Paris, ACM Press, September 1999.
- [XML] Tim Bray, Jean Paoli, and C. M. Sperberg-McQueen, editors, *Extensible Markup Language (XML) 1.0*, W3C Recommendation, 10 February 1998. <http://www.w3.org/TR/1998/REC-xml>
- [XPointer] Steve DeRose, editor, *XML Pointer Language (XPointer)*, W3C Working Draft. <http://www.w3.org/TR/WD-xptr>
- [XSL-Dec98] James Clark and Stephen Deach, editors, *Extensible Stylesheet Language (XSL)*, W3C Working Draft, 16 December 1998. <http://www.w3.org/TR/1998/WD-xsl-19981216>
- [XSLT] James Clark, editor, *XSL Transformations (XSLT)*, W3C Working Draft. <http://www.w3.org/TR/WD-xslt>
- [XPath] James Clark and Steve DeRose, editors, *XSL Path Language (XPath)*, W3C Working Draft, 9 July 1999. <http://www.w3.org/1999/07/WD-xpath-19990709>

$\mathcal{M}$	: $Pattern \rightarrow Node \rightarrow Boolean$
$\mathcal{M}[[p]]x$	= $x \in \{x_2 \mid x_1 \in \text{subnodes}^*(\text{root}(x)), x_2 \in \mathcal{S}[[p]](x_1)\}$
$\mathcal{S}$	: $Pattern \rightarrow Node \rightarrow Set(Node)$
$\mathcal{S}[[p_1 \mid p_2]]x$	= $\mathcal{S}[[p_1]]x \cup \mathcal{S}[[p_2]]x$
$\mathcal{S}[/math>[/math>p]]x$	= $\mathcal{S}[[p]](\text{root}(x))$
$\mathcal{S}[/math>[/math>p]]x$	= $\{x_2 \mid x_1 \in \text{subnodes}^*(\text{root}(x)), x_2 \in \mathcal{S}[[p]]x_1\}$
$\mathcal{S}[[p_1/p_2]]x$	= $\{x_2 \mid x_1 \in \mathcal{S}[[p_1]]x, x_2 \in \mathcal{S}[[p_2]]x_1\}$
$\mathcal{S}[[p_1//p_2]]x$	= $\{x_3 \mid x_1 \in \mathcal{S}[[p_1]]x, x_2 \in \text{subnodes}^*(x_1), x_3 \in \mathcal{S}[[p_2]]x_2\}$
$\mathcal{S}[[p[q]]]x$	= $\{x_1 \mid x_1 \in \mathcal{S}[[p]]x, \mathcal{Q}[[q]]x_1\}$
$\mathcal{S}[[n]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isElement}(x_1), \text{name}(x_1) = n\}$
$\mathcal{S}[[*]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isElement}(x_1)\}$
$\mathcal{S}[[@n]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isAttribute}(x_1), \text{name}(x_1) = n\}$
$\mathcal{S}[[@*]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isAttribute}(x_1)\}$
$\mathcal{S}[[\text{text}()]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isText}(x_1)\}$
$\mathcal{S}[[\text{comment}()]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isComment}(x_1)\}$
$\mathcal{S}[[\text{pi}(n)]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isPI}(x_1), \text{name}(x_1) = n\}$
$\mathcal{S}[[\text{pi}()]]x$	= $\{x_1 \mid x_1 \in \text{subnodes}(x), \text{isPI}(x_1)\}$
$\mathcal{S}[[\text{id}(p)]]x$	= $\{x_2 \mid x_1 \in \mathcal{S}[[p]]x, s \in \text{split}(\text{value}(x_1)), x_2 \in \text{id}(s)\}$
$\mathcal{S}[[\text{id}(s)]]x$	= $\{x_1 \mid s_1 \in \text{split}(s), x_1 \in \text{id}(s_1)\}$
$\mathcal{S}[[\text{ancestor}(p)]]x$	= $\text{last}(\{x_1 \mid x_1 \in \text{parent}^+(x), \mathcal{M}[[p]]x_1\})$
$\mathcal{S}[[\text{ancestor-or-self}(p)]]x$	= $\text{last}(\{x_1 \mid x_1 \in \text{parent}^*(x), \mathcal{M}[[p]]x_1\})$
$\mathcal{S}[[.]]x$	= $\{x\}$
$\mathcal{S}[[..]]x$	= $\text{parent}(x)$
$\mathcal{Q}$	: $Qualifier \rightarrow Node \rightarrow Boolean$
$\mathcal{Q}[[q_1 \text{ and } q_2]]x$	= $\mathcal{Q}[[q_1]]x \wedge \mathcal{Q}[[q_2]]x$
$\mathcal{Q}[[q_1 \text{ or } q_2]]x$	= $\mathcal{Q}[[q_1]]x \vee \mathcal{Q}[[q_2]]x$
$\mathcal{Q}[[\text{not}(q)]]x$	= $\neg \mathcal{Q}[[q]]x$
$\mathcal{Q}[[q]]x$	= $\mathcal{Q}[[q]]x$
$\mathcal{Q}[[\text{first-of-type}()]]x$	= $x \in \text{first}(\{x_1 \mid x_1 \in \text{siblingElements}(x), \text{name}(x_1) = \text{name}(x)\})$
$\mathcal{Q}[[\text{last-of-type}()]]x$	= $x \in \text{last}(\{x_1 \mid x_1 \in \text{siblingElements}(x), \text{name}(x_1) = \text{name}(x)\})$
$\mathcal{Q}[[\text{first-of-any}()]]x$	= $x \in \text{first}(\text{siblingElements}(x))$
$\mathcal{Q}[[\text{last-of-any}()]]x$	= $x \in \text{last}(\text{siblingElements}(x))$
$\mathcal{Q}[[p=s]]x$	= $\{x_1 \mid x_1 \in \mathcal{S}[[p]]x, \text{value}(x_1) = s\} \neq \emptyset$
$\mathcal{Q}[[p]]x$	= $\mathcal{S}[[p]]x \neq \emptyset$
$\text{siblingElements}(x)$	= $\{x_2 \mid x_1 \in \text{parent}(x), x_2 \in \text{subnodes}(x_1), \text{isElement}(x_2)\}$

Figure 1: Semantics of patterns in December 1998 draft of XSLT