Harmless Advice[*]

Daniel S. Dantas
Princeton University
ddantas@cs.princeton.edu

David Walker
Princeton University
dpw@cs.princeton.edu

ABSTRACT

This paper defines an object-oriented language with *harmless* aspect-oriented advice. A piece of harmless advice is a computation that, like ordinary aspect-oriented advice, executes when control reaches a designated control-flow point. However, unlike ordinary advice, harmless advice is designed to obey a weak non-interference property. Harmless advice may change the termination behavior of computations and use I/O, but it does not otherwise influence the final result of computations that trigger it. The benefit of harmless advice is that it facilitates local reasoning about program behavior. More specifically, programmers may ignore harmless advice when reasoning about the partial correctness properties of their programs. In addition, programmers may add new pieces of harmless advice to pre-existing programs in typical "after-the-fact" aspect-oriented style without fear they will break important data invariants used by the main-line program.

In order to detect and enforce harmlessness, the paper defines a novel type and effect system related to information-flow type systems. The central technical result is that well-typed harmless advice does not interfere with the mainline computation. We have also implemented a type checker and interpreter for our system.

1. INTRODUCTION

Aspect-oriented programming languages (AOPL) such as AspectJ [16] allow programmers to specify both *what* computation to perform as well as *when* to perform it. For example, AspectJ makes it easy to implement a profiler that records statistics concerning the number of calls to each method: The *what* in this case is the computation that does the recording and the *when* is the instant of time just prior to execution of each method body. In aspect-oriented terminology, the specification of *what* to do is called *advice* and the specification of *when* to do it is called a *point cut*. A collection of point cuts and advice organized to perform a coherent task is called an *aspect*.

Many within the AOP community adhere to the tenet that aspects are most effective when the code they advise is *oblivious* to their presence. In other words, aspects are effective when a programmer is not required to annotate the advised code (henceforth, the *mainline code*) in any particular way. In fact, Filman and Friedman [12] argue that obliviousness is one of two defining characteristics of any aspect-oriented programming language. When aspect-oriented languages are oblivious, all control over when advice is applied rests with the aspect programmer as opposed to the mainline programmer. This design choice simplifies after-the-fact customization or analysis of programs using aspects. For example, obliviousness makes it trivial to implement extremely flexible profiling infrastructure. To adjust the places where profiling occurs, which might be scattered all across the code base, one need only make local changes to a single aspect. Obliviousness might be one of the reasons that aspect-oriented programming has caught on with such enthusiasm in recent years, causing major companies such as IBM and Microsoft to endorse the new paradigm and inspiring academics to create a series of conferences and workshops to study the idea.

On the other hand, obliviousness threatens conventional modularity principles and undermines a programmer's ability to reason locally about the behavior of their code. Consequently, many traditional programming language researchers believe that aspect-oriented programs are a ticking time bomb, which, if widely deployed, are bound to cause the software industry irreparable harm. One of the main problems they forsee is that while mainline code may be *syntactically* oblivious to aspects, it is not *semantically* oblivious to aspects. Aspects can reach inside modules, influence the behavior of local routines and alter local data structures. As a result, to understand the semantics of code in an aspect-oriented language such as AspectJ, programmers will have to examine all external aspects that might modify local data structures or control flow. As the number and invasiveness of aspects grows, understanding and maintaining your program may become more and more difficult.

In this work, we define a new form of *harmless* aspect-oriented advice that programmers can use to accomplish

[*] This research was supported in part by ARDA Grant no. NBCHC030106, NSF grants CCR-0238328 and CCR-0208601, and an Alfred P. Sloan Fellowship. Opinions, findings, conclusions, and recommendations expressed throughout this work are not necessarily the views of the NSF, ARDA or Sloan foundation and no official endorsement should be inferred.

nontrivial programming tasks yet also allows them to enjoy most of the local reasoning principles they have come to depend upon for program understanding, development and maintenance. Like ordinary aspect-oriented advice, harmless advice is a computation that executes whenever mainline control reaches a designated control-flow point. Unlike ordinary aspect-oriented advice, harmless advice is constrained to prevent it from *interfering* with the underlying computation. Consequently, it plays a role similar to Clifton and Leavens' notion of *observer* [5]. Since harmless advice does not interfere with the mainline computation, it can be added to a program at any point in the development cycle without fear that important program invariants will be disrupted. In addition, programmers that develop, debug or enhance mainline code can safely ignore harmless advice, if there is any present.

In principle, one could devise many variants of harmless advice depending upon exactly what it means to *interfere* with the underlying computation. At the most extreme end, changing the timing behavior of a program constitutes interference and consequently, only trivial advice is harmless. A slightly less extreme viewpoint is one taken by secure programming languages such as Jif [20] and Flow Caml [22]. These languages ignore some kinds of interference such as changes to the timing behavior and termination behavior of programs, arguing that these kinds of interference will have a minimal impact on security. However, overall, they continue to place very restrictive constraints on programs, prohibiting I/O in high security contexts, for instance. Allowing unchecked I/O would make it possible to leak secret information at too great a rate.

In our case, an appropriate balance point between useability and interference prevention is even more relaxed than in secure information-flow systems. We say that computation A does not *interfere* with computation B if A does not influence the final value produced by B. Computation A may change the timing and termination behavior of B (influencing whether or not B does indeed return a value) and it may perform I/O. In practice, of course, I/O by A may change the result eventually produced by B. However, we are willing to live with this relatively minor danger as disallowing I/O eliminates too many useful forms of advice.

Our notion of harmless, non-interfering advice continues to support many of the most common aspect-oriented applications, include the following.

- Profiling. Harmless advice can maintain its own state separate from the mainline computation to gather statistics concerning the number of times different procedures are called. When the program terminates, the harmless advice can print out the profiling statistics.

- Invariant checking and security. Harmless advice can check invariants at run-time, maintain access control tables, perform resource accounting, and terminate programs that disobey dynamic security policies.

- Program tracing and monitoring. Harmless advice can print out all sorts of debugging information including when procedures are called and what data they are passed.

- Logging and backups. Harmless advice can back up data onto persistent secondary storage or make logs of events that occurred during program execution for performance analysis, fault recovery or post-mortem security audits.

While not all applications of aspects and aspect-oriented programming can be simulated using harmless advice, we have accumulated anecdotal evidence to determine that enough important applications appear to fall into this category to make it a useful abstraction. IBM experimented with aspects in their middleware product line [6], finding them useful for such "harmless" tasks as enforcing consistency in tracing and logging concerns and for encapsulating monitoring and statistics components. We also observed a sequence of emails on the AspectJ users list [13] cataloging uses of aspects with Java projects. Many respondents specified that, in addition to some uncommon uses that they wished to highlight, they certainly used AspectJ for the common aspect-oriented concerns such as profiling, security, and monitoring mentioned above. From this anecdotal evidence, we determined that our definition of "harmlessness" encompasses many of the common uses of aspects. While it certainly may be fruitful to combine harmless advice with standard forms of advice and corresponding protection mechanisms such as Aldrich's Open Modules [1], in this paper we restrict our attention to harmless advice alone and its non-interference properties.

In the rest of this paper, we develop a theory of harmless advice following the same strategy as used in previous work by Walker, Zdancewic and Ligatti [24] (hereafter referred to as WZL). More specifically, we first develop a core calculus at an intermediate level of abstraction. The calculus contains primitive notions of point cuts, advice and a collection of static protection domains, arranged in a partial order. We define a novel type system to guarantee that code, including advice, in a low-protection domain cannot influence execution of code in a high-protection domain. Though the type system is directly inspired by information-flow type systems for security [20, 22], we take advantage of the syntactic separation between advice and code to simplify it. Also, since we are not interested in security, our type system can be somewhat less strict than information-flow type systems. We believe our simplifications make our language easier to use than existing information-flow type systems. The key technical result of our research is a proof that our type system satisfies a non-interference property. The proof adapts the syntactic technique used by Simonet and Pottier in their proof of non-interference of Flow Caml [22] to a our language involving aspects.

After developing the core language, we define a higher-level surface language that is more amenable to programming. In particular, the high-level language is *oblivious* and therefore "aspect-oriented" according to Filman and Friedman's definition [12], whereas the core language is not.[1] The high-level language allows programmers to define aspects that are collections of state, objects and advice. Each aspect operates in a separate static protection domain and does not

[1] It is neither necessary nor the slightest bit desirable for the core language to be oblivious as the syntax of the core language does not limit or constrain programmers in any way. Programmers need only concern themselves with the surface language, which is oblivious. Of course, in any aspect-oriented language, obliviousness is only a property of the source — every oblivious language is compiled into a non-oblivious language at some point.

interfere with the mainline computation or the other aspects. The semantics of the high-level language is defined via translation into the core language. We prove the translation rules directly define a type system for the surface language and we prove the translation only creates well-typed core language programs. Since the translation places each aspect in its own protection domain, which has a lower privilege than the mainline code, and the translated code is well-typed, the non-interference theorem for the core calculus guarantees that our aspects do not interfere with the mainline code.

In addition to engaging in a full meta-theoretic analysis of our language, we have implemented a type checker and interpreter for it in SML. We briefly describe our implementation towards the end of the paper, where we also have a summary of related work in this area and some conclusions.

2. CORE LANGUAGE

Our core language is a typed lambda calculus containing strings, booleans, tuples, references and simple objects. The two main features of interest in the language are labeled control-flow points and advice, both of which are slight variants of related constructs introduced by WZL.

Labels l, which are drawn from some countably infinite set, mark points in a computation at which advice may be triggered. For instance, execution of $l\,\texttt{[}e_1\texttt{]}; e_2$ proceeds by first evaluating e_1 until it reduces to a value v and at this point, any advice associated with the label l executes with v as an input. Once all advice associated with l has completed execution, control returns to the marked point and evaluation continues with e_2. Notice that a marked point $l\,\texttt{[}e_1\texttt{]}$ has type unit and that no data are returned from the triggered advice. This stands in contrast to earlier work by WZL, in which labels marked control-flow points where data exchange could occur.

Harmless advice $\{pcd.x \rightarrow e\}$ is a computation that is triggered whenever execution reaches the control-flow point described by the pointcut designator *pcd*. When advice is triggered, the value at the control-flow point is bound to x, which may be used within the body of the advice e. The advice body may have “harmless” effects (such as I/O), but it does not return any data to the mainline computation and consequently e is expected to have type unit.

Languages such as AspectJ often contain rich sublanguages for designating control-flow points. However, it is easier to study the fundamentals of labeled control-flow points and harmless advice in a setting with the simplist possible *pcd*s. Consequently, we will start our investigation in a setting where *pcd*s are simply sets of labels $\{l_1, \ldots, l_k\}$ and advice is written as $\{\{l_1, \ldots, l_k\}.x \rightarrow e\}$.

For simplicity, the core language contains a single construct $\Uparrow a$ to activate new advice a. When control reaches a label in the advice's point cut designator, the advice body will execute after any previously activated piece of advice. The following example shows how advice activation works (assuming that there is no other advice associated with label l in the environment).

$$\begin{array}{l}
\Uparrow \{\{l\}.x \rightarrow \texttt{printint x; print}\ \text{"}\texttt{ : hello }\text{"}\}; \\
\Uparrow \{\{l\}.y \rightarrow \texttt{print}\ \text{"}\texttt{world}\text{"}\}; \\
l\,\texttt{[3]}
\end{array}$$

prints `3 : hello world`

The expression $\texttt{new} : \tau$ allows programs to generate a fresh label with type τ. Labels are considered first class values, so they may be passed to functions or stored in data structures before being used to mark control-flow points. For example, we might write

$$\begin{array}{l}
\texttt{let}\ pt = \texttt{new : int in} \\
\Uparrow \{\{pt\}.x \rightarrow \texttt{print}\ \text{"}\texttt{hello }\text{"}\}; \\
\Uparrow \{\{pt\}.y \rightarrow \texttt{print}\ \text{"}\texttt{world}\text{"}\}; \\
pt\,\texttt{[3]}
\end{array}$$

to allocate a new label and use it in two pieces of advice.

2.1 Types for Enforcing Harmlessness

In order to protect the mainline computation from interference from advice, we have devised a type and effect system for the calculus we informally introduced in the previous section. The type system operates by ascribing a protection domain p to each expression in the language. These protection domains are organized in a lattice $\mathrm{L} = (\mathit{Protections}, \leq)$ where *Protections* is the set of possible protection domains and $p \leq q$ specifies that p should not interfere with q. Alternatively, one might say that data in q have higher integrity than data in p. In our examples, we often assume there are `high`, `med` and `low` protection levels with `low < med < high`.

Syntax. In order to allow programmers to specify protection requirements we have augmented the syntax of the core language described in the previous section with a collection of protection annotations. The formal syntax appears below.

$$\begin{array}{lcl}
\multicolumn{3}{l}{p \in \text{Protections} \qquad l \in \text{Labels} \qquad \texttt{s} \in \text{Strings}} \\
\tau & ::= & \texttt{unit} \mid \texttt{string} \mid \texttt{bool} \mid \tau_1 \times \ldots \times \tau_n \\
 & \mid & \tau \rightarrow_p \tau \mid [m_i{:}_{p_i}\tau_i]^{1..n} \\
 & \mid & \texttt{advice}_p \mid \tau\ \texttt{label}_p \mid \tau\ \texttt{ref}_p \mid \tau\ \texttt{pcd}_p \\
v & ::= & \texttt{()} \mid \texttt{s} \mid \texttt{true} \mid \texttt{false} \mid (\vec{v}) \\
 & \mid & \lambda_p x : \tau.e \mid [m_i = \varsigma_p\, x_i.e_i]^{1..n} \mid \{v.x \rightarrow_p e\} \\
 & \mid & l \mid r \mid \{\vec{l}\}_p \\
e & ::= & v \mid x \mid e_1; e_2 \mid \texttt{print}\ e \\
 & \mid & \texttt{if}\ e_1\ \texttt{then}\ e_2\ \texttt{else}\ e_3 \\
 & \mid & (\vec{e}) \mid \texttt{split}\,(\vec{x}) = e\ \texttt{in}\ e \\
 & \mid & e\ e \mid e.m \mid \{e.x \rightarrow_p e\} \mid \Uparrow e \\
 & \mid & \texttt{new}_p : \tau \mid e\texttt{[}e\texttt{]} \\
 & \mid & \texttt{ref}_p\ e \mid\ !\,e \mid e := e \\
 & \mid & \{\vec{e}\}_p \mid e \cup_p e \mid p\texttt{<}e\texttt{>}
\end{array}$$

The values include unit values and string and boolean constants. Programmers may also use n-ary tuples. Functions are annotated with the protection domain p in which they execute. This protection domain also shows up in the type of the function. Objects are collections of methods, with each method taking a single parameter (self). Methods and object types are also annotated with protection domains. Advice values $\{v.x \rightarrow_p e\}$ are annotated with their protection domain as well. Labels l and reference locations r do not appear in initial programs; they only appear as programs execute and generate new labels and new references.

Most of the expression forms are fairly standard. For instance, in addition to values and variables, we allow ordinary expression forms for sequencing, printing strings, conditionals, tuples, function calls, and method invocations. Expressions for introducing and eliminating advice were explained in the previous section. The expressions $\texttt{new}_p : \tau$ and $\texttt{ref}_p\ e$ allocate labels that can be placed in protection domain p and references associated with protection domain p respectively. The last command p<e> is a typing coercion that changes the current protection domain to the lower protection domain p.

2.2 Typing

The main typing judgment in our system has the form $\Gamma; p \vdash e : \tau$. It states that in the context Γ, expression e has type τ and may influence computations occurring in protection domains p or lower. A related judgment $\Gamma \vdash v : \tau$ checks that value v has type τ. Since values by themselves do not have effects that influence the computations, this latter judgment is not indexed by a protection domain. The context Γ maps variables, labels and reference locations to their types. We use the notation $\Gamma, x : \tau$ to extend Γ so that it maps x to τ. Whenever we extend Γ in this way, we assume that x does not already appear in the domain of Γ. Since we also treat all terms as equivalent up to alpha-renaming of bound variables, it will always be possible to find a variable x that does not appear in Γ when we need to. Figures 1 and 2 contain the rules for typing expressions and values respectively.

The main goal of the typing relation is to guarantee that no values other than values with unit type (which have no information content) flow from a low protection domain to a high protection domain, although arbitrary data can flow in the other direction. This goal is very similar to, but not exactly the same as in, standard information flow systems such as Jif and Flow Caml. The latter systems actually do allow flow of values from low contexts to high contexts, but mark all such values with a low-protection type. Jif and Flow Caml typing rules make it impossible to use these low-protection objects in the high-protection context (without raising the protection of the context). In our system, we simply cut off the flow of low-protection values to high-protection contexts completely (aside from the unit value). We are able to do this in our setting, as there is a greater syntactic separation between high-integrity code (the main-line computation) and low-integrity code (the advice, written elsewhere) than there might be in a standard secure information-flow setting. We believe this is the right design choice for us because it simplifies the type system as we do not have to annotate basic data such as booleans, strings or tuples with information flow labels.

Most of the value typing rules are straightforward. For instance, the rule for functions $\lambda_p x : \tau.e$, states that the body of the function must be checked under the assumption that the code operates in protection domain p. The resulting type has the shape $\tau \ \rightarrow_p \ \tau'$. Checking our simple objects is similar: the type checker must verify that each method operates correctly in the declared protection domain. Labels and references are given types by the context. In the current calculus, point-cut designators are sets of labels. Unlike the other values, the rules for typing advice are fairly subtle. We will discuss these rules in a moment together with the rules for typing labeled control-flow points.

The first few expression typing rules (see Figure 2) are standard rules for type systems that track information flow. The rule for if deviates slightly from the usual rule for tracking information flow. Normally, types for booleans will contain a security level and the branches of the if will be checked at a level equal to the join of the current security level and the level of the boolean. However, in our system, any data, including booleans, manufactured by code at level p contains level p information. Consequently, the branches of the if statement may be safely checked at level p. The typing rules for function calls and method invocations require that the function or method in question be safe to run at the current protection level p.

The typing rules for references enforce the usual integrity constraint found in information-flow systems. When in protection domain p, we are allowed to dereference references in protection domain p' when p is less than or equal to p'. We are allowed to store to references in protection domain p' only if our current domain p is greater than or equal to p'.

The last rule in Figure 2 is a typing coercion that changes the protection level. It is legal for the protection level to be lowered from p to p' when no information flows back from the computation e to be executed. We prevent this information flow by constraining the result type of e to be `unit`. One might wonder whether the following dual rule, which allows one to raise the protection level is sound in our system:

$$\frac{\cdot; p' \vdash e : \tau \qquad \vdash p \leq p'}{\Gamma; p \vdash p' \texttt{>}e\texttt{<} : \tau}$$

This rule raises the protection domain for the expression e and allows information to flow out of the expression, but does not allow any information to flow in. In the context of the features we have looked at so far, this rule appears sound, but in combination with the context-sensitive advice we will introduce in Section 2.4, it is not. Fortunately, the rule does not appear useful in our application and we have omitted it.[2]

The last component of our type system involves the rules for typing advice and marking control-flow points. If we want to ensure that low-protection code cannot interfere with high-protection code by manipulating advice and control-flow labels, we must be sure that low-protection code cannot do either of the following:

1. Declare and activate high-protection advice. For instance, assume r is a high-protection reference with type $\texttt{int ref}_{\texttt{high}}$ and l is a label that has been placed in high-protection code. If we allow the expression

   $$\{l.x \rightarrow_{\texttt{high}} r := 3 + x\} \texttt{<<} e$$

   to appear in low-protection code, then this low privilege code can indirectly cause writes to the reference r.

2. Mark a control-flow point with a label that triggers high-protection advice. For instance, assume that

   $$\{l.x \rightarrow_{\texttt{high}} r := 3 + x\}$$

   is an active piece of high-protection advice which writes to the high-protection reference r. Placing the label l in low-protection code allows low-protection code to

[2] There may well be some strategy that allows us to add this rule together with the context-sensitive advice of Section 2.4. However, the naive approach does not appear to work. Rather then complicating the type structure or operational semantics for something we do not need, we leave it out.

$$\frac{}{\Gamma \vdash \texttt{()} : \texttt{unit}} \qquad \frac{}{\Gamma \vdash \texttt{s} : \texttt{string}}$$

$$\frac{}{\Gamma \vdash \texttt{true} : \texttt{bool}} \qquad \frac{}{\Gamma \vdash \texttt{false} : \texttt{bool}}$$

$$\frac{(\Gamma \vdash v_i : \tau_i)^{1 \leq i \leq n}}{\Gamma \vdash \texttt{(}\vec{v}\texttt{)} : \tau_1 \times ... \times \tau_n} \qquad \frac{\Gamma, x : \tau; p \vdash e : \tau'}{\Gamma \vdash \lambda_p x : \tau . e : \tau \ \rightarrow_p \ \tau'}$$

$$\frac{((\Gamma, x : [m_i{:}_{p_i}\tau_i]^{1..n}); p_j \vdash e_j : \tau_j)^{(1 \leq j \leq n)}}{\Gamma \vdash [m_i = \varsigma_{p_i}\, x_i . e_i]^{1..n} : [m_i{:}_{p_i}\tau_i]^{1..n}}$$

$$\frac{\Gamma \vdash v : \tau\ \texttt{pcd}_p \quad \Gamma, x : \tau; p' \vdash e : \texttt{unit} \quad \vdash p' \leq p}{\Gamma \vdash \{v.x \rightarrow_{p'} e\} : \texttt{advice}_{p'}}$$

$$\frac{\Gamma(l) = \tau\ \texttt{label}_p}{\Gamma \vdash l : \tau\ \texttt{label}_p} \qquad \frac{\Gamma(r) = \tau\ \texttt{ref}_p}{\Gamma \vdash r : \tau\ \texttt{ref}_p}$$

$$\frac{(\Gamma \vdash v_i : \tau\ \texttt{label}_{p_i})^{(1 \leq i \leq n)} \quad (\vdash p \leq p_i)^{(1 \leq i \leq n)}}{\Gamma \vdash \{\vec{l}\}_p : \tau\ \texttt{pcd}_p}$$

Figure 1: Value Typing

determine via its control-flow, when the high-protection advice will run and write to r.

In order to properly protect high-protection code in the face of these potential errors, we do the following.

1. Add protection levels to advice types (e.g., $\texttt{advice}_{\texttt{high}}$), which will allow us to prevent advice from being activated in the illegal contexts. (eg. low-protection contexts)
2. Add protection levels to label types (e.g., $\texttt{string label}_{\texttt{high}}$) which will allow us to prevent labels being placed in illegal spots. (eg. low-protection contexts)

One might hope that it would be possible to simplify the system and add protection levels to only one of the two constructs, but doing so leads to unsoundness.

Five typing rules in the middle of Figure 2 give the well-formedness conditions for advice and labels. Notice that in the rule for typing advice introduction, the protection level of the advice, and therefore the protection level the body of the advice must operate under, is connected to the protection level of the label that triggers it. Notice also that when marking a control-flow point with a label, the protection level of the label is connected to the protection level of the expression at that point. Finally, given a high-protection piece of advice, this advice cannot be launched from low-protection code. The result of these constraints is that when in a low-protection zone, there is no way to cause execution of high-protection advice.

2.3 Operational Semantics

The definition of the operational semantics for our language largely follows earlier work by WZL. In particular, we use a context-based semantics. The top-level operational judgment has the form $(S, A, p, e) \longmapsto (S', A', p, e')$ where S collects the labels l that may be used to mark control-flow points and also maps reference locations r to values.

$$\frac{\Gamma \vdash v : \tau}{\Gamma; p \vdash v : \tau} \qquad \frac{\Gamma(x) = \tau}{\Gamma; p \vdash x : \tau}$$

$$\frac{\Gamma; p \vdash e_1 : \texttt{unit} \quad \Gamma; p \vdash e_2 : \tau}{\Gamma; p \vdash e_1; e_2 : \tau}$$

$$\frac{\Gamma; p \vdash e : \texttt{string}}{\Gamma; p \vdash \texttt{print}\ e : \texttt{unit}}$$

$$\frac{\Gamma; p \vdash e_1 : \texttt{bool} \quad \Gamma; p \vdash e_2 : \tau \quad \Gamma; p \vdash e_3 : \tau}{\Gamma; p \vdash \texttt{if}\, e_1\, \texttt{then}\, e_2\, \texttt{else}\, e_3 : \tau}$$

$$\frac{(\Gamma; p \vdash e_i : \tau_i)^{1 \leq i \leq n}}{\Gamma; p \vdash \texttt{(}\vec{e}\texttt{)} : \tau_1 \times ... \times \tau_n}$$

$$\frac{\Gamma; p \vdash e_1 : \tau_1 \times ... \times \tau_n \quad \Gamma, (\vec{x} : \vec{t}); p \vdash e_2 : \tau}{\Gamma; p \vdash \texttt{split}\ \texttt{(}\vec{x}\texttt{)} = e_1\ \texttt{in}\ e_2 : \tau}$$

$$\frac{\Gamma; p \vdash e_1 : \tau_1 \ \rightarrow_p \ \tau_2 \quad \Gamma; p \vdash e_2 : \tau_1}{\Gamma; p \vdash e_1\ e_2 : \tau_2}$$

$$\frac{\Gamma; p \vdash e : [m_i{:}_{p_i}\tau_i]^{1..n} \quad 1 \leq j \leq n \quad p = p_j}{\Gamma; p \vdash e.m_j : \tau_j}$$

$$\frac{\Gamma; p \vdash e_1 : \tau\ \texttt{pcd}_{p'} \quad \Gamma, x : \tau; p'' \vdash e_2 : \texttt{unit} \quad \vdash p'' \leq p'}{\Gamma; p \vdash \{e_1.x \rightarrow_{p''} e_2\} : \texttt{advice}_{p''}}$$

$$\frac{\Gamma; p \vdash e : \texttt{advice}_{p'} \quad \vdash p' \leq p}{\Gamma; p \vdash \Uparrow e : \texttt{unit}}$$

$$\frac{\vdash p' \leq p}{\Gamma; p \vdash \texttt{new}_{p'} : \tau : \tau\ \texttt{label}_{p'}}$$

$$\frac{\Gamma; p \vdash e_1 : \tau\ \texttt{label}_p \quad \Gamma; p \vdash e_2 : \tau}{\Gamma; p \vdash e_1\, \texttt{[}e_2\texttt{]} : \texttt{unit}}$$

$$\frac{\Gamma; p \vdash e : \tau \quad \vdash p' \leq p}{\Gamma; p \vdash \texttt{ref}_{p'}\ e : \tau\ \texttt{ref}_{p'}} \qquad \frac{\Gamma; p \vdash e : \tau\ \texttt{ref}_{p'} \quad \vdash p \leq p'}{\Gamma; p \vdash !e : \tau}$$

$$\frac{\Gamma; p \vdash e_1 : \tau\ \texttt{ref}_{p'} \quad \Gamma; p \vdash e_2 : \tau \quad \vdash p' \leq p}{\Gamma; p \vdash e_1 := e_2 : \tau}$$

$$\frac{(\Gamma; p \vdash e_i : \tau\ \texttt{label}_{p_i})^{(1 \leq i \leq n)} \quad (\vdash p' \leq p_i)^{(1 \leq i \leq n)}}{\Gamma; p \vdash \{\vec{e}\}_{p'} : \tau\ \texttt{pcd}_{p'}}$$

$$\frac{\begin{array}{ll} \Gamma; p \vdash e_1 : \tau\ \texttt{pcd}_{p''} & \vdash p' \leq p'' \\ \Gamma; p \vdash e_2 : \tau\ \texttt{pcd}_{p'''} & \vdash p' \leq p''' \end{array}}{\Gamma; p \vdash e_1 \cup_{p'} e_2 : \tau\ \texttt{pcd}_{p'}}$$

$$\frac{\Gamma; p' \vdash e : \texttt{unit} \quad \vdash p' \leq p}{\Gamma; p \vdash p'\texttt{<}e\texttt{>} : \texttt{unit}}$$

Figure 2: Expression Typing

The meta-variable A represents an advice store, which is a list of advice. The current protection level of the code is p. The protection level does not influence execution of the code, and could be omitted, but is useful to consider in our noninterference proof. Most of the real work is done by the auxiliary relation $(S, A, p, e) \longmapsto_\beta (S', A', p, e')$. The additional syntactic categories are given below.

$$\begin{array}{lcl}
S & ::= & \cdot \;|\; S, r = e \;|\; S, l \\
A & ::= & \cdot \;|\; A, \{v.x \rightarrow_p e\}
\end{array}$$

$$\begin{array}{lcl}
E & ::= & E; e \;|\; \texttt{print}\, E \;|\; \texttt{if}\, E\, \texttt{then}\, e_2\, \texttt{else}\, e_3 \\
 & | & (v_i, ..., v_i, E, e_{i+2}, ..., e_n) \\
 & | & \texttt{split}\, (\vec{x}) = E\, \texttt{in}\, e \\
 & | & E\ e \;|\; v\ E \;|\; E.m \\
 & | & \{E.x \rightarrow_p e\} \;|\; \Uparrow E \\
 & | & E\texttt{[}e\texttt{]} \;|\; l\texttt{[}E\texttt{]} \\
 & | & \texttt{ref}_p\, E \;|\; !\, E \;|\; E := e \;|\; r := E \\
 & | & \{v_1, \ldots, v_i, E, e_{i+2}, \ldots, e_n\}_p \;|\; E \cup_p e \;|\; v \cup_p E
\end{array}$$

The definitions of these relations can be found in Figure 3. Notice that the rule for marked control-flow points depends upon an auxiliary function $\mathcal{A}\llbracket A \rrbracket_{l\,\texttt{[}v\texttt{]}} = e$. This function selects all advice in A that is triggered by the label l and combines their bodies to form the expression e. The advice composition function can be found in Figure 4. Finally, an abstract machine configuration (S, A, p, e) is well-typed if it satisfies the judgement $\vdash (S, A, p, e)\ \texttt{ok}$ specified in Figure 5.

2.4 Context-Sensitive Advice

The advice defined in previous sections could not analyze the call stack from which it was activated. Programming languages such as AspectJ allow this flexibility via special pointcut designators such as CFlow. In this section, we describe a fully general facility for analysis of information on the current call stack. Our new mechanism is inspired by earlier work by WZL, but is more general and fits better with the functional programming paradigm. The following definitions describe the syntactic extensions to our calculus:

$$\begin{array}{rcl}
\tau & ::= & ... \;|\; \texttt{stack} \\
 & & \\
v & ::= & ... \;|\; \cdot \;|\; l\,\texttt{[}v\texttt{]} :: v \\
 & & \\
e & ::= & ... \;|\; \texttt{stack()} \;|\; \texttt{store}\, e\texttt{[}e\texttt{]}\, \texttt{in}\, e \\
 & | & \texttt{case}\, e\, \texttt{of}\, (pat \Rightarrow e \;|\; \_ \Rightarrow e) \\
 & & \\
pat & ::= & \texttt{nil} \;|\; e\,\texttt{[}x\texttt{]} :: pat \\
 & | & \_ :: pat \;|\; x \\
 & & \\
vpat & ::= & \texttt{nil} \;|\; \{\vec{l}\}_p\,\texttt{[}x\texttt{]} :: vpat \\
 & | & \_ :: vpat \;|\; x \\
 & & \\
E & ::= & ... \;|\; \texttt{store}\, E\,\texttt{[}e\texttt{]}\, \texttt{in}\, e \;|\; \texttt{store}\, l\,\texttt{[}E\texttt{]}\, \texttt{in}\, e \\
 & | & \texttt{store}\, l\,\texttt{[}v\texttt{]}\, \texttt{in}\, E \\
 & | & \texttt{case}\, E\, \texttt{of}\, (pat \Rightarrow e \;|\; \_ \Rightarrow e) \\
 & | & \texttt{case}\, v\, \texttt{of}\, (Epat \Rightarrow e \;|\; \_ \Rightarrow e) \\
 & & \\
Epat & ::= & ... \;|\; E\,\texttt{[}x\texttt{]} :: pat \;|\; \{\vec{l}\}_p\,\texttt{[}x\texttt{]} :: Epat \\
 & | & \_ :: Epat \\
 & & \\
F & ::= & ... \;|\; \texttt{[]} \;|\; E[F] \;|\; p < F >
\end{array}$$

$$\frac{(S, A, p, e) \longmapsto_\beta (S', A', p, e')}{(S, A, p, e) \longmapsto (S', A', p, e')}$$

$$\frac{(S, A, p, e) \longmapsto (S', A', p, e)}{(S, A, p, E[e]) \longmapsto (S', A', p, E[e'])}$$

$$\frac{(S, A, p', e) \longmapsto (S', A', p', e')}{(S, A, p, p'\texttt{<}e\texttt{>}) \longmapsto (S', A', p, p'\texttt{<}e'\texttt{>})}$$

$$(S, A, p, (); e) \longmapsto_\beta (S, A, p, e)$$

$$(S, A, p, \texttt{print}\ s) \longmapsto_\beta (S, A, p, ())$$

$$(S, A, p, \texttt{if}\, true\, \texttt{then}\, e_1\, \texttt{else}\, e_2) \longmapsto_\beta (S, A, p, e_1)$$

$$(S, A, p, \texttt{if}\, false\, \texttt{then}\, e_1\, \texttt{else}\, e_2) \longmapsto_\beta (S, A, p, e_2)$$

$$(S, A, p, \texttt{split}\, (\vec{x}) = (\vec{v})\, \texttt{in}\ e) \longmapsto_\beta (S, A, p, e\{\vec{v}/\vec{x}\})$$

$$(S, A, p, \lambda_p x : t.e\ v) \longmapsto_\beta (S, A, p, e\{v/x\})$$

$$\begin{array}{l}(S, A, p, [m_i = \varsigma_{p_i}\, x_i.e_i]^{1..n}.m_j) \longmapsto_\beta \\ \qquad (S, A, p, e_j\{[m_i = \varsigma_{p_i}\, x_i.e_i]^{1..n}/x_j\})\end{array}$$

$$(S, A, p, \Uparrow \{v.x \rightarrow_{p'} e_1\}) \longmapsto_\beta (S, (A, \{v.x \rightarrow_{p'} e_1\}), p, ())$$

$$(l \notin S) \quad (S, A, p, \texttt{new}_{p'} : \tau) \longmapsto_\beta ((S, l), A, p, l)$$

$$\frac{l \in S \quad \mathcal{A}\llbracket A \rrbracket_{l\,\texttt{[}v\texttt{]}} = e}{(S, A, p, l\,\texttt{[}v\texttt{]}) \longmapsto_\beta (S, A, p, e)}$$

$$(r \notin S) \quad (S, A, p, \texttt{ref}_{p'}\ v) \longmapsto_\beta ((S, r = v), A, p, r)$$

$$(S, A, p, !\, r) \longmapsto_\beta (S, A, p, S(r))$$

$$(S, A, p, r := v) \longmapsto_\beta ((S, r = v), A, p, v)$$

$$(S, A, p, \{\vec{l_1}\}_{p'} \cup_{p''} \{\vec{l_2}\}_{p'''}) \longmapsto_\beta (S, A, p, \{\vec{l_1}\ \vec{l_2}\}_{p''})$$

$$(S, A, p, p'\texttt{<()>}) \longmapsto_\beta (S, A, p, ())$$

Figure 3: Operational Semantics

$$\frac{}{\mathcal{A}\llbracket \cdot \rrbracket_{l\,\texttt{[}v\texttt{]}} = ()}$$

$$\frac{l\,\texttt{[}v\texttt{]} \models v' \quad \mathcal{A}\llbracket A \rrbracket_{l\,\texttt{[}v\texttt{]}} = e}{\mathcal{A}\llbracket \{v'.x \rightarrow_p e'\}, A \rrbracket_{l\,\texttt{[}v\texttt{]}} = p\texttt{<}e'\{v/x\}\texttt{>}; e}$$

$$\frac{l\,\texttt{[}v\texttt{]} \not\models v' \quad \mathcal{A}\llbracket A \rrbracket_{l\,\texttt{[}v\texttt{]}} = e}{\mathcal{A}\llbracket \{v'.x \rightarrow_p e'\}, A \rrbracket_C = e}$$

$$\frac{l \in \{\vec{l}\}_p}{l\,\texttt{[}v\texttt{]} \models \{\vec{l}\}_p}$$

Figure 4: Aspect Composition

$$\frac{\begin{array}{c} dom(S) = dom(\Gamma) \\ \forall r \in dom(S).\ \Gamma(r) = \tau\ \mathtt{ref}_p \quad \Gamma \vdash S(r) : \tau\ \mathtt{for\ some}\ p, \tau \\ \forall l \in dom(S).\ \Gamma(r) = \tau\ \mathtt{ref}_p\ \mathtt{for\ some}\ p, \tau \end{array}}{\vdash S : \Gamma}$$

$$\frac{}{\Gamma \vdash \cdot\ \mathtt{ok}} \qquad \frac{\Gamma \vdash a : \mathtt{advice}_p\ \mathtt{for\ some}\ p \quad \Gamma \vdash A\ \mathtt{ok}}{\Gamma \vdash A, a\ \mathtt{ok}}$$

$$\frac{\vdash S : \Gamma \quad \Gamma \vdash A\ \mathtt{ok} \quad \Gamma; p \vdash e : \tau\ \mathtt{for\ some}\ \tau}{\vdash (S, A, p, e)\ \mathtt{ok}}$$

Figure 5: Abstract Machine Judgement

In order to program with context-sensitive advice, programmers grab the current stack using the `stack()` command. Data is explicitly allocated on the stack using the command $\mathtt{store}\ e_1\mathtt{[}e_2\mathtt{]}\ \mathtt{in}\ e_3$, where e_1 is a label and e_2 represents a value associated with the label. e_2 is typically used to store the value passed into the control flow point marked by the label. The `store` command evaluates e_1 to a label l and e_2 to a value v_2, places $l\,\mathtt{[}v_2\mathtt{]}$ on the stack, evaluates e_3 to a value v_3 and finally removes $l\,\mathtt{[}v_2\mathtt{]}$ from the stack and returns v_3. The programmer may examine a stack data structure using the $\mathtt{case}\ e\ \mathtt{of}\ (pat \Rightarrow e \ \mid\ \_ \Rightarrow e)$ command, which matches the stack e against the pattern pat. If there is a match, the first branch is executed; otherwise, the second branch is executed. There are patterns that match the empty stack (e.g., $\cdot$), patterns that match a stack starting with any label in a particular set (e.g., $\{\vec{l}\}_p\,\mathtt{[}x\mathtt{]} :: pat$) where x is bound to the value associated with the label on the top of the stack if it is in the label set, patterns that match a stack starting with anything at all (e.g., $\_ :: pat$), and patterns involving stack variables (e.g., x).

The typing rules for these extensions appear in Figure 6. There are three sets of rules in this figure. The first two extend the value typing and expression typing relations respectively. The last set of rules gives types to patterns where the type of a pattern is a context Γ that describes the types of the variables bound within the pattern.

The rules for evaluating these new expressions appear in Figure 7. Again, there are three sets of rules. The first defines a new set of top-level evaluation rules, and the second adds additional β-evaluation rules. Notice that the top-level rule for evaluating the stack primitive uses an auxiliary function $\mathcal{S}(F)$ that extracts the current stack of values from F contexts, which contains evaluation context E's, and $p\mathtt{<}F\mathtt{>}$ contexts. Here, we use the notation $st@X$ to append the object X to the bottom of the stack st. The last set of rules conclude in judgments with the form $st \models vpat \Rightarrow sub$. These rules describe the circumstances under which a stack st matches an (evaluated) pattern $vpat$ and generates a substitution of values for variables sub.

For the most part, it is relatively straightforward to reassure oneself that these extensions will not disrupt the noninterference properties that our language possesses. However, there is one major subtlety to consider: the `stack()` primitive. In order for this primitive to be safe, it must be the case that whenever it is activated in a high-level context, there is no low-level data on the stack, which could influence execution in that high-level context. Fortunately, this is indeed the case. The only way to switch protection levels

$$\boxed{\Gamma \vdash v : \tau}$$

$$\frac{}{\Gamma \vdash \cdot : \mathtt{stack}}$$

$$\frac{\Gamma \vdash l : \tau\ \mathtt{label}_p \quad \Gamma \vdash v_1 : \tau \quad \Gamma \vdash v_2 : \mathtt{stack}}{\Gamma \vdash l\,\mathtt{[}v_1\mathtt{]} :: v_2 : \mathtt{stack}}$$

$$\boxed{\Gamma; p \vdash e : \tau}$$

$$\frac{}{\Gamma; p \vdash \mathtt{stack()} : \mathtt{stack}}$$

$$\frac{\Gamma; p \vdash e_1 : \tau'\ \mathtt{label}_{p'} \quad \Gamma; p \vdash e_2 : \tau' \quad \Gamma; p \vdash e_3 : \tau}{\Gamma; p \vdash \mathtt{store}\ e_1\mathtt{[}e_2\mathtt{]}\ \mathtt{in}\ e_3 : \tau}$$

$$\frac{\begin{array}{c} \Gamma; p \vdash e_1 : \mathtt{stack} \\ \Gamma; p \vdash pat \Rightarrow \Gamma' \qquad \Gamma, \Gamma'; p \vdash e_2 : \tau \qquad \Gamma; p \vdash e_3 : \tau \end{array}}{\Gamma; p \vdash \mathtt{case}\ e_1\ \mathtt{of}\ (pat \Rightarrow e_2 \ \mid\ \_ \Rightarrow e_3) : \tau}$$

$$\boxed{\Gamma; p \vdash pat \Rightarrow \Gamma}$$

$$\frac{}{\Gamma; p \vdash \mathtt{nil} \Rightarrow \cdot} \qquad \frac{\Gamma; p \vdash e : \tau\ \mathtt{pcd}_{p'} \quad \Gamma; p \vdash pat \Rightarrow \Gamma'}{\Gamma; p \vdash e\mathtt{[}x\mathtt{]} :: pat \Rightarrow \Gamma', x : \tau}$$

$$\frac{\Gamma; p \vdash pat \Rightarrow \Gamma'}{\Gamma; p \vdash \_ :: pat \Rightarrow \Gamma'} \qquad \frac{}{\Gamma; p \vdash x \Rightarrow \cdot, x : \mathtt{stack}}$$

Figure 6: Advanced Point-cut Designator Typing

from one evaluation context to the next is via the context p<E>, which lowers the protection level. Consequently, any use of the `stack()` command is done in the context that looks like p_1<$E_1[p_2$<$E_2[p_3$<E_3>$]$>$]$> where $p_3 \leq p_2 \leq p_1$. So while a low-level expression can read high-level data via the `stack()` command and subsequent `scase` expressions, the opposite is not possible. We are safe.

2.5 Core Language Meta-theory

To prove noninterference, we use the technique developed by Simonet and Pottier [22]. In order to do so, we initially assume the collection of protection domains has been divided into two groups, the high protection domains (H) and the low protection domains (L). The low-protection group is a downward-closed subset of protection domains and the high-protection group contains all other protection domains. The goal is to prove that low-protection code cannot interfere with the behavior of high-protection code, no matter how aspects, references or labels are used. Overall, our proof may be broken down into five steps:

- Define a new language `Core2` that simulates execution of two original (henceforth referred to as `Core1`) programs.
- Show `Core2` is a correct simulation of `Core1` programs via Soundness and Completeness theorems.
- Prove `Core2` is a safe language via the standard Progress and Preservation theorems.
- Show that the previous step implies well-typed `Core2` programs simulate `Core1` programs that produce indistinguishable results.
- Put the theorems above together to prove the final noninterference result for `Core1`.

A diagram of the proof may be seen in Figure 8.

Each of the following subsections sketches one portion of the proof.

2.5.1 *Defining* `Core2`

We begin by defining a new language (`Core2`) that simulates execution of two of our original programs. The main syntactic difference between `Core1` expressions and `Core2` expressions is the brackets expression, p<$e_1|e_2$>. Here p is a low-protection label and the e_i are `Core1` expressions. These brackets expressions encapsulate all differences between the two `Core1` expressions that are being simulated. For instance, the `Core2` expression

```
p<print ``hi'' | print ``bi''>;x+3
```

represents the two `Core1` programs

```
p<print ``hi''>;x+3
```

```
p<print ``bi''>;x+3
```

The typing rule for the bracket expression requires that the two subexpressions have low protection and release no information into the surrounding high-protection context.

$$\frac{\begin{array}{c}\Gamma; p' \vdash_2 e_1 : \texttt{unit} \qquad \Gamma; p' \vdash_2 e_2 : \texttt{unit}\\ p \in H \quad p' \in L \quad \vdash_2 p' \leq p\end{array}}{\Gamma; p \vdash_2 p'\texttt{<}e_1\,|\,e_2\texttt{>} : \texttt{unit}}$$

$\boxed{(S, A, p, e) \longmapsto_{top} (S', A', p, e')}$

$$\frac{(S, A, p, e) \longmapsto (S', A', p, e')}{(S, A, p, e) \longmapsto_{top} (S', A', p, e')}$$

$$(S, A, p, F[\texttt{stack()}]) \longmapsto_{top} (S, A, p, F[\mathcal{S}(F)])$$

where :

$$\begin{array}{rcl}
\mathcal{S}(\texttt{[]}) & = & \cdot \\
\mathcal{S}(\texttt{store}\ l\texttt{[}v\texttt{]}\ \texttt{in}\ F) & = & \mathcal{S}(F)\ @\ (l\texttt{[}v\texttt{]}) \\
\mathcal{S}(p < F >) & = & \mathcal{S}(F) \\
\mathcal{S}(E[F]) & = & \mathcal{S}(F) \\
 & & \text{when } E \neq \texttt{store}\ l\texttt{[}v\texttt{]}\ \texttt{in}\ F
\end{array}$$

$\boxed{(S, A, p, e) \longmapsto_{\beta} (S, A, p, e)}$

$$\frac{}{(S, A, p, \texttt{store}\ v_1\texttt{[}\tau\texttt{]}\ \texttt{in}\ v_2) \longmapsto_{\beta} (S, A, p, v_2)}$$

$$\frac{v \models vpat \Rightarrow sub}{\begin{array}{c}(S, A, p, \texttt{case}\ v\ \texttt{of}\ (vpat \Rightarrow e_1 \ \mid\ \_ \Rightarrow e_2)) \longmapsto_{\beta}\\ (S, A, p, sub(e_1))\end{array}}$$

$$\frac{v \not\models vpat \Rightarrow sub}{\begin{array}{c}(S, A, p, \texttt{case}\ v\ \texttt{of}\ (vpat \Rightarrow e_1 \ \mid\ \_ \Rightarrow e_2)) \longmapsto_{\beta}\\ (S, A, p, e_2)\end{array}}$$

$\boxed{v \models vpat \Rightarrow sub}$

$$\frac{}{\cdot \models \texttt{nil} \Rightarrow \cdot}$$

$$\frac{l \in \{\vec{l}\}_p \quad v_2 \models vpat \Rightarrow sub}{l\texttt{[}v_1\texttt{]} :: v_2 \models \{\vec{l}\}_p\texttt{[}x\texttt{]} :: vpat \Rightarrow sub, \{v_1/x\}}$$

$$\frac{v_2 \models vpat \Rightarrow sub}{l\texttt{[}v_1\texttt{]} :: v_2 \models \_ :: vpat \Rightarrow sub}$$

$$\frac{}{v \models x \Rightarrow \{v/x\}}$$

Figure 7: Advanced Point-cut Designator Evaluation

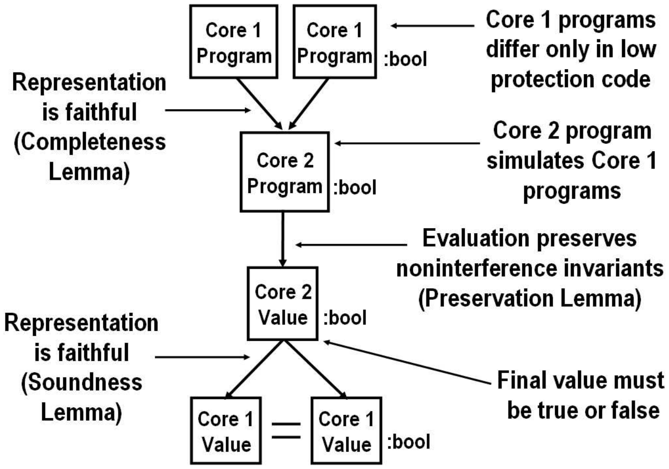
Core 1 Program
Core 1 Program
:bool
Core 1 programs differ only in low protection code
Representation is faithful (Completeness Lemma)
Core 2 Program
:bool
Core 2 program simulates Core 1 programs
Evaluation preserves noninterference invariants (Preservation Lemma)
Core 2 Value
:bool
Representation is faithful (Soundness Lemma)
Final value must be true or false
Core 1 Value
=
Core 1 Value
:bool

Figure 8: Noninterference Proof Diagram

To express the operational semantics of `Core2` we add need to add similar bracket constructs for the contents of the reference/label store S and the aspect store A.

$$\begin{array}{rcl}
v^2 & ::= & v \mid \texttt{<}v\,|\,v\texttt{>} \mid \texttt{<}v\,|\,void\texttt{>} \mid \texttt{<}void\,|\,v\texttt{>} \\
\tau^2 & ::= & \tau \mid \texttt{<}\tau\,|\,void\texttt{>} \mid \texttt{<}void\,|\,\tau\texttt{>} \\
a^2 & ::= & v \mid \texttt{<}v\,|\,void\texttt{>} \mid \texttt{<}void\,|\,v\texttt{>} \\
S & ::= & \cdot \mid S, r = v^2 \mid S, l \to \tau^2 \\
A & ::= & \cdot \mid A, a^2
\end{array}$$

The *void* marker indicates that the appropriate element is not present in the program. For example, if advice a is activated in only the left instance but not the right instance of simultaneously executing `Core1` programs, the aspect store of the `Core2` program that simulates them will contain $\texttt{<}a\,|\,void\texttt{>}$.

To relate `Core1` to `Core2`, we define the projection function $|\ |_i$ where i ∈ 1,2. $|p\texttt{<}e_1|e_2\texttt{>}|_i$ is $p\texttt{<}e_i\texttt{>}$ and $|\ |_i$ is a homomorphism on all other expressions. Since $p\texttt{<}e_1|e_2\texttt{>}$ in `Core2` simulates the simultaneous execution of two low-protection original `Core1` expressions, the projection function extracts one of these two executions.

The `Core2` machine state (S, A, p, e) symbolizes the current state of the two simultaneously executing `Core1` programs where the i-th projection is the state of the i-th `Core1` program:

$$|(S, A, p, e)|_i = (|S|_i, |A|_i, p, |e|_i)$$

The projection function for the reference/label store and the aspect store is similar to the one for expressions.

The definition of the complete operational semantics of `Core2` is not too difficult, but it does involve a substantial amount of notation. To avoid overwhelming the reader with details in this short report, we omit the full definition. Intuitively, the main ideas are as follows:

- Ordinary `Core1` expressions embedded within `Core2` expressions operate as `Core1` expressions normally do.

- To evaluate inside the brackets expression $p\texttt{<}e_1|e_2\texttt{>}$, the semantics nondeterministically chooses one of e_1 or e_2 to execute. Operations on references or aspects executed in e_1 use the "left-hand" component of the reference store or aspect store; operations on references or aspects executed in e_2 use the "right-hand" component of the reference store or aspect store.

- To evaluate the expression $p\texttt{<()|()>}$, we simply throw away the brackets, returning the unit value `()`. Since `()` carries no information, no information is transmitted between low- and high-protection contexts.

- Whenever values v_1 and v_2 are not both `()`, the expression $p\texttt{<}v_1|v_2\texttt{>}$ is stuck. Fortunately, such an expression is ill-typed and never arises from evaluation of a well-typed program.

The judgment form for execution of top-level `Core2` expressions has the same shape as the judgment form for `Core1` expressions. The index 2 on the arrow distinguishes the two judgements:

$$(S, A, p, e) \longmapsto^*_{2,top} (S', A', p, e')$$

2.5.2 *Relating* `Core1` *and* `Core2`

Once `Core2` has been defined, it is necessary to show that it accurately simulates two `Core1` programs. Two theorems, one concerning the *soundness* of `Core2`execution relative to `Core1` and the other concerning the *completeness* of `Core2` relative to `Core1` help to establish the proper correspondence.

The soundness theorem states that if a `Core2` expression takes a step, then the two corresponding `Core1` programs (the projections of the `Core2` expression) must each take the same respective steps. The proof of this theorem requires, among other things, an auxiliary lemma that establishes a soundness result for aspect composition.

Lemma 2.1 (Aspect Composition Soundness Lemma). *For* $i \in \{1, 2\}$, *if* $\mathcal{A}[\![A]\!]_{l\,\texttt{[}v\texttt{]}} = e$, *then* $\mathcal{A}[\![|A|_i]\!]_{l\,\texttt{[}|v|_i\texttt{]}} = |e|_i$.

Proof. By induction on the structure of the definition of $\mathcal{A}[\![A]\!]_{l\,\texttt{[}v\texttt{]}} = e$. □

Theorem 2.1 (Soundness). *For* $i \in \{1, 2\}$,
if $(S, A, p, e) \longmapsto^*_{2,top} (S', A', p, e')$
then $|(S, A, p, e)|_i \longmapsto^*_{top} |(S', A', p, e')|_i$

Proof. By induction on the structure of the operational judgment $(S, A, p, e) \longmapsto^*_{2,top} (S', A', p, e')$, with use of Lemma 2.1 □

The completeness theorem states that if two `Core1` programs step to values, then the representation in `Core2` that simulates them simultaneously must step to a value. The completeness theorem requires an auxiliary lemma stating that a `Core2` program is only stuck when one of its corresponding `Core1` programs are stuck.

Lemma 2.2 (Completeness Stuck Lemma). *Assume* (S, A, p, e) *is stuck. Then* $|(S, A, p, e)|_i$ *is stuck for some* $i \in \{1, 2\}$.

Proof. Proof by induction on the structure of e. □

Theorem 2.2 (Completeness). *Assume* $|(S, A, p, e)|_i \longmapsto^*_{top} (S'_i, A'_i, p, v_i)$ *for all* $i \in 1, 2$ *then there exists* (S', A', p, v) *such that* $(S, A, p, e) \longmapsto^*_{2,top} (S', A', p, v)$

Proof. If (S, A, p, e) yields an infinite reduction sequence, then $|(S, A, p, e)|_i$ yields a infinite reduction sequence by Theorem 2.1. If (S, A, p, e) is stuck, then $|(S, A, p, e)|_i$ is stuck by Lemma 2.2. Therefore, (S, A, p, e) reduces to a successful configuration. □

2.5.3 Safety of `Core2`

To continue we prove that the type system of `Core2` is sound with respect to our operational semantics using Progress and Preservation theorems. This strategy requires that we extend the typing relation to cover all of the run-time terms in the language as well as the other elements of the abstract machine (*i.e.,* the code store and aspect store). A `Core2` configuration (S, A, p, e) is well-typed if it satisfies the judgement $\vdash_2 (S, A, p, e)\ \texttt{ok}$ specified in Figure 9. Note that if the stores and the expression contain brackets, the protection domains associated with the brackets must be low.

Part of the proof of Progress involves defining the canonical forms of each type. It is important to notice here that the brackets expression is not a value and therefore the only values with type `bool`, for instance, are `true` and `false`. This fact comes into play later in the noninterference proof. The following lemma gives the rest of canonical forms.

Lemma 2.3 (Canonical Forms). *Suppose* $\cdot \vdash_2 v : \tau$ *is a closed, well-formed value.*

- *If* $\tau = \texttt{unit}$, *then* $v = ()$.
- *If* $\tau = \texttt{string}$, *then* $v = \texttt{s}$.
- *If* $\tau = \texttt{bool}$, *then* $v = \texttt{true}$ *or* $\texttt{false}$.
- *If* $\tau = \tau_1 \times ... \times \tau_n$, *then* $v = (\vec{v})$.
- *If* $\tau = \tau_1 \ \rightarrow_p\ \tau_2$, *then* $v = \lambda_p x : \tau_1.e$.
- *If* $\tau = [m_i{:}_{p_i}\tau_i]^{1..n}$, *then* $v = [m_i = \varsigma_p\, x_i.e_i]^{1..n}$.
- *If* $\tau = \texttt{advice}_p$, *then* $v = \{v.x \rightarrow_p e\}$.
- *If* $\tau = \tau\ \texttt{label}_p$, *then* $v = l$.
- *If* $\tau = \tau\ \texttt{ref}_p$, *then* $v = r$.
- *If* $\tau = \tau\ \texttt{pcd}_p$, *then* $v = \{\vec{l}\}_p$.
- *If* $\tau = \texttt{stack}$, *then* $v = \cdot$ *or* $l\,[v_1] :: v_2$.

Proof. By induction on the structure of $\Gamma \vdash_2 v : \tau$, using the fact that v is a value. □

We now state the standard Progress and Preservation lemmas.

Theorem 2.3 (Progress). *If* $\vdash_2 (S, A, p, e)\ \texttt{ok}$ *then either* e *is a value, or there exists* (S', A', p, e') *such that* $(S, A, p, e) \longmapsto_{2,top} (S', A', p, e')$.

Proof. By induction on the structure of the typing judgment $\vdash_2 (S, A, p, e)\ \texttt{ok}$. □

Theorem 2.4 (Preservation). *If* $\vdash_2 (S, A, p, e)\ \texttt{ok}$ *and* $(S, A, p, e) \longmapsto_{2,top} (S', A', p, e')$ *then* $\vdash_2 (S', A', p, e')\ \texttt{ok}$.

Proof. By induction on the structure of the operational judgment $(S, A, p, e) \longmapsto_{2,top} (S', A', p, e')$. □

$$\boxed{\vdash_2 (S, A, p, e)\ \texttt{ok}}$$

$$\frac{\vdash_2 S : \Gamma \quad \Gamma \vdash_2 A\ \texttt{ok} \quad \Gamma; p \vdash_2 e : \tau\ \texttt{for some}\ \tau \quad p \in H}{\vdash_2 (S, A, p, e)\ \texttt{ok}}$$

$$\boxed{\vdash_2 S : \Gamma}$$

$$\frac{\begin{array}{c} dom(S) = dom(\Gamma) \\ \forall r \in dom(S).\ \Gamma \vdash_2 S(r) \Rightarrow \Gamma(r) \\ \forall l \in dom(S).\ \vdash_2 S(l) \Rightarrow \Gamma(l) \end{array}}{\vdash_2 S : \Gamma}$$

$$\boxed{\Gamma \vdash_2 v^2 \Rightarrow \tau\ \texttt{ref}_p}$$

$$\frac{\Gamma \vdash_2 v : \tau}{\Gamma \vdash_2 v \Rightarrow \tau\ \texttt{ref}_p} \qquad \frac{\Gamma \vdash_2 v_1 : \tau \quad \Gamma \vdash_2 v_2 : \tau \quad p \in L}{\Gamma \vdash_2 \texttt{<}v_1|v_2\texttt{>} \Rightarrow \tau\ \texttt{ref}_p}$$

$$\frac{\Gamma \vdash_2 v : \tau \quad p \in L}{\Gamma \vdash_2 \texttt{<}v|void\texttt{>} \Rightarrow \tau\ \texttt{ref}_p} \qquad \frac{\Gamma \vdash_2 v : \tau \quad p \in L}{\Gamma \vdash_2 \texttt{<}void|v\texttt{>} \Rightarrow \tau\ \texttt{ref}_p}$$

$$\boxed{\vdash_2 \tau^2 \Rightarrow \tau\ \texttt{label}_p}$$

$$\frac{}{\vdash_2 \tau \Rightarrow \tau\ \texttt{label}_p}$$

$$\frac{p \in L}{\vdash_2 \texttt{<}\tau|void\texttt{>} \Rightarrow \tau\ \texttt{label}_p} \qquad \frac{p \in L}{\vdash_2 \texttt{<}void|\tau\texttt{>} \Rightarrow \tau\ \texttt{label}_p}$$

$$\boxed{\Gamma \vdash_2 A\ \texttt{ok}}$$

$$\frac{}{\Gamma \vdash_2 \cdot\ \texttt{ok}} \qquad \frac{\Gamma \vdash_2 A\ \texttt{ok} \quad \Gamma \vdash_2 a^2 \Rightarrow \texttt{advice}_p}{\Gamma \vdash_2 A, a^2\ \texttt{ok}}$$

$$\boxed{\Gamma \vdash_2 a^2 \Rightarrow \texttt{advice}_p}$$

$$\frac{\Gamma \vdash_2 a : \texttt{advice}_p}{\Gamma \vdash_2 a \Rightarrow \texttt{advice}_p}$$

$$\frac{\Gamma \vdash_2 a : \texttt{advice}_p \quad p \in L}{\Gamma \vdash_2 \texttt{<}a|void\texttt{>} \Rightarrow \texttt{advice}_p} \qquad \frac{\Gamma \vdash_2 a : \texttt{advice}_p \quad p \in L}{\Gamma \vdash_2 \texttt{<}void|a\texttt{>} \Rightarrow \texttt{advice}_p}$$

Figure 9: `Core2` Abstract Machine Judgement

2.5.4 Well-typed `Core2` programs produce indistinguishable `Core1` results

Most of the difficult work has now been done. We merely need to apply lemmas and theorems we have already proven to get our first powerful result: If a high-protection `Core2` expression steps to a boolean value, then the corresponding `Core1` projections (which differ only in low protection code) step to equal values. In other words, no low-production code (be it aspect-oriented features or otherwise) has influenced execution of high-protection expressions.

LEMMA 2.4 (EQUIVALENT EXECUTION IN `Core2`). *If* $\mathit{high} \in H$ *and* $\cdot; \mathit{high} \vdash_2 e : \texttt{bool}$ *and* $(\cdot, \cdot, \mathit{high}, e) \longmapsto^*_{2,top} (S, A, \mathit{high}, v)$ *then* $|v|_1 = |v|_2$.

PROOF. By Theorem 2.4, $\vdash_2 S : \Gamma$ and $\Gamma \vdash_2 v : \texttt{bool}$. By Lemma 2.3, v is either $\texttt{true}$ or $\texttt{false}$. $|\texttt{true}|_1 = |\texttt{true}|_2$ and $|\texttt{false}|_1 = |\texttt{false}|_2$. □

2.5.5 Putting it all together: Noninterference

Finally, for the noninterference proof, we assume a high-protection `Core1` expression e steps to a value. We add a low-protection expression $\mathit{low}\texttt{<}e'\texttt{>}$ where $\mathit{low} \in L$ to e so that e with the low-protection code and e alone are executed simultaneously and their resulting values compared. This is achieved by constructing the `Core2` expression $\mathit{low}\texttt{<}e'|()\texttt{>}; e$ where the left projection is e with the low-protection code and the right projection steps to e alone. Using the soundness, completeness, preservation theorems, and the equivalent execution in `Core2` lemma, we show that both e with the added low-protection code and e alone step to the same value. Therefore the low-protection code did not interfere with execution.

THEOREM 2.5 (NONINTERFERENCE). *If* $\mathit{high} \in H$ *and* $\mathit{low} \in L$ *and* $\vdash \mathit{low} \leq \mathit{high}$ *and* e *is a core language expression where* $\cdot; \mathit{high} \vdash e : \texttt{bool}$ *and* $\cdot; \mathit{low} \vdash e' : \texttt{unit}$ *and* $(\cdot, \cdot, \mathit{high}, \mathit{low}\texttt{<}e'\texttt{>}; e) \longmapsto^*_{top} (S_1, A_1, \mathit{high}, v_1)$ *and* $(\cdot, \cdot, \mathit{high}, \mathit{low}\texttt{<()>}; e) \longmapsto^*_{top} (S_2, A_2, \mathit{high}, v_2)$ *then* $v_1 = v_2$.

PROOF. Construct the `Core2` expression $\mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e$, such that $|\mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e|_1 = \mathit{low}\texttt{<}e'\texttt{>}; e$ and $|\mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e|_2 = \mathit{low}\texttt{<()>}; e$.
Therefore, $|(\cdot, \cdot, \mathit{high}, \mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e)|_1 \longmapsto^*_{2,top} (S_1, A_1, \mathit{high}, v_1)$ and $|(\cdot, \cdot, \mathit{high}, \mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e)|_2 \longmapsto^*_{2,top} (S_2, A_2, \mathit{high}, v_2)$.
By Theorem 2.2, $(\cdot, \cdot, \mathit{high}, \mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e) \longmapsto^*_{2,top} (S, A, \mathit{high}, v)$ for some S, A, and v.
By Theorem 2.1, for $i \in \{1, 2\}$, $|(\cdot, \cdot, \mathit{high}, \mathit{low}\texttt{<}e'\,\texttt{|}\texttt{()}\texttt{>}; e)|_i \longmapsto^*_{2,top} |(S, A, \mathit{high}, v)|_i$.
Therefore, $|v|_1 = v_1$ and $|v|_2 = v_2$.
By Lemma 2.4, $|v|_1 = |v|_2$.
Therefore, $v_1 = v_2$. □

3. SOURCE LANGUAGE

Our core calculus is intended to be used as a semantic intermediate language rather than as a source-level programming language of its own. The main reason for this is that core calculus sits at a convenient level of abstraction for formulating a semantics, but programmers would almost certainly complain that it is inconvenient to have to mark control-flow labels in code, to allocate values on the stack by hand, and to deal with the low-level core calculus notion of advice. In addition, the core calculus does not actually define a policy concerning whether or not advice can interfere with each other or the mainline computation. Rather, it defines a way for a programmer (or compiler) to assign different protection levels to code and a mechanism (the type system) that can check that there is no interference between the appropriate protection levels.

In order to show how the core calculus can be used, we define a simple source language and show how to translate it into the core calculus. This source language consists of a sequence of ordinary declarations, *aspects*, which are collections of advice declarations and ordinary declarations, and a mainline program. The translation from the source into the core places the state and code for each aspect into its own protection domain. The mainline code and initial declarations get their own protection domain, which sits above the protection domains for the aspects in the security lattice. Consequently, the translation specifies the noninterference policy that we wish to enforce, namely that no aspect interferes with any other aspect and that no aspect interferes with the mainline computation. The syntax of the source language appears below.

$$
\begin{array}{rcl}
\tau & ::= & \texttt{unit} \mid \texttt{string} \mid \texttt{bool} \\
 & \mid & [m_i{:}_{p_i}\tau_i \;\rightarrow_{p_i}\; \tau_i]^{1..n} \mid \tau\,\texttt{ref}_p \mid \texttt{stack} \\
 & & \\
v & ::= & \texttt{()} \mid \texttt{s} \mid \texttt{true} \mid \texttt{false} \\
 & & \\
e & ::= & v \mid x \mid e; e \mid \texttt{print}\, e \\
 & \mid & \texttt{if}\; e\; \texttt{then}\; e\; \texttt{else}\; e \\
 & \mid & \texttt{let}\; ds\; \texttt{in}\; e \\
 & \mid & e.m(e) \\
 & \mid & !\, e \mid e := e \\
 & \mid & \texttt{case}\; e\; \texttt{of}\; (\mathit{pat} \Rightarrow e \mid \_ \Rightarrow e)
\end{array}
$$

$$
\begin{array}{rcl}
\mathit{pat} & ::= & \texttt{nil} \mid \{o.m_i\}^{1..n}\,\texttt{[}x, y, n\texttt{]} :: \mathit{pat} \mid \_ :: \mathit{pat} \mid x \\
 & & \\
d & ::= & (\texttt{string}\, x = e) \\
 & \mid & (\texttt{bool}\, x = e) \\
 & \mid & (\texttt{ref}\, x = e) \\
 & \mid & (\texttt{object}\, o = [m_i : \tau_i \;\rightarrow\; \tau_i' = \varsigma\, x_i.\lambda y_i.e_i]^{1..n}) \\
 & & \\
ds & ::= & . \mid d\; ds \\
 & & \\
a & ::= & (\texttt{before}\; \{o.m_i\}^{1..n}(x, y, s, n) = e) \\
 & \mid & (\texttt{after}\; \{o.m_i\}^{1..n}(x, y, s, n) = e) \\
 & & \\
as & ::= & . \mid d\; as \mid a\; as \\
 & & \\
\mathit{aspcts} & ::= & . \mid p : \{as\}\; \mathit{aspcts} \\
 & & \\
\mathit{prog} & ::= & ds\;\; \mathit{aspcts}\;\; e
\end{array}
$$

The types of the source language objects are a restricted form of the internal language types. In particular, source language object types are the composition of a core language object and function type. Also, since programmers in the source language do not explicitly manipulate labels, there are no label types in the source language.

Most of the source language expressions and values mimic the core language expressions and values, although there are a few differences. For instance, none of the run-time-only values such as labels, reference locations, or stack values need appear in the collection of source values as the source language is not executed directly.[3] Also, for convenience, we allow a local let declaration in expressions, which programmers can use to allocate values with basic type, references or objects. Note that we use the meta-variable o to stand for program variables bound to objects. We use the meta-variable x to stand for any kind of program variable.

The source language case expressions analyze stack values in a similar way to the target, only the patterns are slightly different, reflecting a particular compilation strategy. More specifically, when compiling a method, we will allocate automatically on the stack the label corresponding to the method on top of the stack and a tuple containing a pointer to self, a pointer to the method argument, and a string corresponding to the name of the method that was called. Consequently, the patterns that match stack frames have the form $\{o.m_i\}^{1..n}[x, y, n]$, where $\{o.m_i\}^{1..n}$ is checked against the label, and x, y, and n are bound to self, the argument and the string respectively. The string can be used when printing out debugging information, profiling information, etc.

Advice in the source language is either before advice that runs before a method call or after advice that runs after the method call. Similar to the source-language stack patterns, when the advice is triggered, x is bound to self, y is bound to the method argument, and n is bound to a string corresponding to the method name. The variable s is bound to the stack at the point the advice is triggered. In the source language, programmers do not explicitly allocate their own data on the stack, nor do they explicitly grab the current stack. Code for performing these actions is emitted at specific points during the translation from source into core.

Finally, as mentioned above, a whole source-language program (*prog*) is a collection of declarations (*ds*) together with a collection of aspects (*aspcts*) and a mainline computation (e). The protection level of the mainline code is *main*. Each aspect is given a distinct name p, which will also serve as its protection domain when translated into the core calculus. We assume the translation program operates in the presence of a security lattice in which $p \leq main$ for all aspects p in the program. Figure 10 displays the protection domains for the mainline code and several aspects. Aspects are simply collections of local declarations and advice (*as*).

As an example of the basic features of our source language, consider the code in Figure 11, where we take the liberty of assuming our language has been augmented with integers. It declares a math object which has a internal integer state that can be modified with the set, add, and sub methods. We write a tracer aspect that prints informative messages before and after the add and subtract methods are executed. The mainline computation performs a series of arithmetic operations on the math object.

3.1 Translation to Core Language

The translation from source into core is defined by a series of 5 mutually recursive judgments. The translation judgments are generally parameterized by a typing context involving a point-cut context (P), which contains a collection of declarations that can be used in source-level point-cuts, a standard type context (Γ), which maps source variables to types, and a protection level/aspect name (p). The

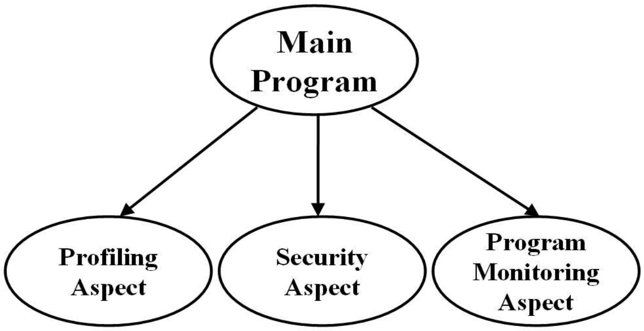

Figure 10: Organization of Protection Domains

```
ref r = 0
object math = [
  get:unit->int = ςx.λy.!r
  set:int->unit = ςx.λy.r:=y
  add:int->int = ςx.λy.
                  let z = y + x.get() in
                  x.set(z); z
  sub:int->int = ςx.λy.
                  let z = y - x.get() in
                  x.set(z); z
]
tracer: {
  before {math.add,math.sub}(x,y,s,n) =
    print "entering "; print n;
    print " with arg "; print (itos y)
  after {math.add,math.sub}(x,y,s,n) =
    print " and leaving\n"
}
let x = math.add(math.add(1)) in
math.sub(3 - x)
```

Figure 11: Source Language Example

[3] "Execution" of the source occurs by translation of the source into the core and then execution of the resulting core program.

point-cut context P contains declarations of the form $o.m : (\tau_{self}, \tau_{arg}, \tau_{res}, p)$. These declarations say that an object named o with method m has been declared and may be advised. The object has the type τ_{self} and the method takes an argument with type τ_{arg} and returns a result with type τ_{res}. The object inhabits protection domain p.

The form of the translation judgments are as follows.

- The judgment $P; \Gamma \vdash v : \tau \stackrel{\mathsf{val}}{\Longrightarrow} v'$ describes the translation from source language values v with type τ to core language values v' with type τ.

- The judgment $P; \Gamma; p \vdash e : \tau \stackrel{\mathsf{exp}}{\Longrightarrow} e'$ describes the translation from source language expressions e with type τ to core language expressions e' with type τ.

- The judgement $\mathtt{split}(\Theta, e)$ is used by the stack case operation. What is extracted from the core language stack is a tuple containing the object, the argument of the method, and the name of the method. The `split` function extracts the individual elements from these tuples.

- The judgement $\Gamma \dashv \Theta \Rightarrow \Gamma'$ takes a context for individual elements pulled from the stack–the object, the argument of the method, and the name of the method and returns a context containing a tuple of those individual elements. This new context with tuples is what is actually generated by the pattern translation described in the next section. This judgement is used in the proof of translation type safety.

- The judgment $P; p \vdash pat \stackrel{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma; \Theta$ describes the translation from source language patterns pat to core language patterns pat' binding variables described by Γ. Notice that the context Γ returned describes individual elements–the object, the argument to the method, and the name of the method. It is modified by Θ by the judgement $\Gamma \dashv \Theta \Rightarrow \Gamma'$ to generate the new context containing tuples that the core language pattern pat' actually generates. Later, the `split` command in the stack case translation will be used to extract the individual elements from the tuples.

- The judgment $P; \Gamma; p \vdash as; aspcts; e : \tau \stackrel{\mathsf{dec}}{\Longrightarrow} e'$ describes the translation of declarations as, aspects $aspcts$ and mainline code e. The scope of the declarations as includes both $aspcts$ and e. Mainline code e has type τ and the expression e' that results from the translation has type τ as well.

- The judgment $\vdash ds\ aspcts\ e \stackrel{\mathsf{prog}}{\Longrightarrow} e'$ translates a whole program $prog$ with a mainline computation producing values of type τ into a core language expression e' with type τ.

The definition of these judgments may be found in Figures 12 and 13. Throughout the translation we use the abbreviation $\mathtt{let}\, x = e_1 \,\mathtt{in}\, e_2$ to stand for $(\lambda_p x{:}\tau.e_2)\ e_1$ for some appropriate type τ and protection p, which can be determined from the context.

Most of the translation is rather mundane. The interesting cases involve object declarations and advice. Object declarations are translated by first allocating two sets of labels, one set for the control flow points at the beginning of methods, and one for the control flow points at the end of methods. In a rather severe abuse of notation, we bind these new labels to variables with the names "$om_{i,\mathtt{pre}}$" and "$om_{i,\mathtt{post}}$." During the translation, we maintain the invariant that whenever $o.m : (\tau_{self}, \tau_{arg}, \tau_{res}, p)$ appears in the context P, the translated term is well typed in a context including the variables $om_{i,\mathtt{pre}}$ with type $(\tau_{self}, \tau_{arg}, \mathtt{string})\ \mathtt{label}_p$ and $om_{i,\mathtt{post}}$ with type $(\tau_{self}, \tau_{res}, \mathtt{string})\ \mathtt{label}_p$. In the body of each method of an object, the translation first allocates onto the stack the $om_{i,\mathtt{pre}}$ label with a tuple containing self, the argument of the method, and a string name corresponding to the source object and method name[4]. Then we mark the following control-flow point with the $om_{i,\mathtt{pre}}$ label for the method, passing a tuple including self, the argument and the string to the advice. Next comes the body of the method and finally, the $om_{i,\mathtt{post}}$ label including self, the result and the string.

Before and after advice are translated similarly although before advice is triggered by the $om_{i,\mathtt{pre}}$ label whereas after advice is triggered by $om_{i,\mathtt{post}}$ label. In both cases, the first action inside the advice body involves extracting the components (self, method argument or result, and string name) from the advice argument z. Next, the translated advice grabs the current stack and binds it to the variable s. Finally, the advice executes the translated body. After declaring the advice, the translated code immediately activates it, placing it after any previously encountered advice.

3.2 Translation Meta-theory

An important property of the translation is that it produces well-typed core language expressions. Define $\mathcal{T}(o.m : (\tau_{self}, \tau_{arg}, \tau_{res}, p))$ to be the context $om_{\mathtt{pre}} : (\tau_{self} \times \tau_{arg} \times \mathtt{string})\ \mathtt{label}_p, om_{\mathtt{post}} : (\tau_{self} \times \tau_{res} \times \mathtt{string})\ \mathtt{label}_p$ and let $\mathcal{T}(P)$ be the point-wise extension of the previous $\mathcal{T}$ function.

Lemma 3.1 (Translation Split Lemma). *If* $\Gamma, \Gamma'; p \vdash e : \tau$ *and* $\Gamma' \dashv \Theta \Rightarrow \Gamma''$*, then* $\Gamma, \Gamma''; p \vdash \mathtt{split}(\Theta, e) : \tau$

Proof. By induction on the structure of $\Gamma \dashv \Theta \Rightarrow \Gamma'$. □

Lemma 3.2 (Translation Type Safety Lemmas).

- *If* $P; p \vdash pat \stackrel{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma; \Theta$ *and* $\Gamma \dashv \Theta \Rightarrow \Gamma'$ *then* $\mathcal{T}(P); p \vdash pat' \Rightarrow \Gamma'$.

- *If* $P; \Gamma \vdash v : \tau \stackrel{\mathsf{val}}{\Longrightarrow} v'$*, then* $\mathcal{T}(P), \Gamma \vdash v' : \tau$.

- *If* $P; \Gamma; p \vdash e : \tau \stackrel{\mathsf{exp}}{\Longrightarrow} e'$*, then* $\mathcal{T}(P), \Gamma; p \vdash e' : \tau$.

- *If* $P; \Gamma; p \vdash as; aspcts; e : \tau \stackrel{\mathsf{dec}}{\Longrightarrow} e'$*, then* $\mathcal{T}(P), \Gamma; p \vdash e' : \tau$.

Proof. By induction on the structure of $P; p \vdash pat \stackrel{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma; \Theta$, $P; \Gamma \vdash v : \tau \stackrel{\mathsf{val}}{\Longrightarrow} v'$, $P; \Gamma; p \vdash e : \tau \stackrel{\mathsf{exp}}{\Longrightarrow} e'$, and $P; \Gamma; p \vdash as; aspcts; e : \tau \stackrel{\mathsf{dec}}{\Longrightarrow} e'$, with use of Lemma 3.1. □

Using these lemmas, we can now prove translation type safety.

Theorem 3.1 (Translation Type Safety). *If* $\vdash ds\ aspcts\ e \stackrel{\mathsf{prog}}{\Longrightarrow} e'$*, then* $.; main \vdash e' : \tau$ *for some* τ.

Proof. Straightforward use of Lemma 3.2. □

[4] Again, there is an abuse of notation here. We assume that we may write "$o.m$" for the string equivalent of the object name and method.

$\boxed{P; \Gamma \vdash v : \tau \overset{\mathsf{val}}{\Longrightarrow} v'}$

$$\frac{}{P; \Gamma \vdash () : \texttt{unit} \overset{\mathsf{val}}{\Longrightarrow} ()} \qquad \frac{}{P; \Gamma \vdash s : \texttt{string} \overset{\mathsf{val}}{\Longrightarrow} s}$$

$$\frac{}{P; \Gamma \vdash \texttt{true} : \texttt{bool} \overset{\mathsf{val}}{\Longrightarrow} \texttt{true}} \qquad \frac{}{P; \Gamma \vdash \texttt{false} : \texttt{bool} \overset{\mathsf{val}}{\Longrightarrow} \texttt{false}}$$

$\boxed{P; \Gamma; p \vdash e : \tau \overset{\mathsf{exp}}{\Longrightarrow} e'}$

$$\frac{P; \Gamma \vdash v : \tau \overset{\mathsf{val}}{\Longrightarrow} v'}{P; \Gamma; p \vdash v : \tau \overset{\mathsf{exp}}{\Longrightarrow} v'} \qquad \frac{\Gamma(x) = \tau}{P; \Gamma; p \vdash x : \tau \overset{\mathsf{exp}}{\Longrightarrow} x} \qquad \frac{P; \Gamma; p \vdash e_1 : \texttt{unit} \overset{\mathsf{exp}}{\Longrightarrow} e_1' \quad P; \Gamma; p \vdash e_2 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_2'}{P; \Gamma; p \vdash e_1; e_2 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_1'; e_2'}$$

$$\frac{P; \Gamma; p \vdash e : \texttt{string} \overset{\mathsf{exp}}{\Longrightarrow} e'}{P; \Gamma; p \vdash \texttt{print}\ e : \texttt{unit} \overset{\mathsf{exp}}{\Longrightarrow} \texttt{print}\ e'} \qquad \frac{P; \Gamma; p \vdash e_1 : \texttt{bool} \overset{\mathsf{exp}}{\Longrightarrow} e_1' \quad P; \Gamma; p \vdash e_2 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_2' \quad P; \Gamma; p \vdash e_3 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_3'}{P; \Gamma; p \vdash \texttt{if}\, e_1\, \texttt{then}\, e_2\, \texttt{else}\, e_3 : \tau \overset{\mathsf{exp}}{\Longrightarrow} \texttt{if}\, e_1'\, \texttt{then}\, e_2'\, \texttt{else}\, e_3'}$$

$$\frac{P; \Gamma; p \vdash ds; .; e : \tau \overset{\mathsf{dec}}{\Longrightarrow} e'}{P; \Gamma; p \vdash \texttt{let}\ ds\ \texttt{in}\ e : \tau \overset{\mathsf{exp}}{\Longrightarrow} e'} \qquad \frac{\begin{array}{ll} P; \Gamma; p \vdash e_1 : [m_i{:}_{p_i}\tau_i]^{1..n} \overset{\mathsf{exp}}{\Longrightarrow} e_1' & \tau_j = \tau \ \to_p \ \tau' \\ P; \Gamma; p \vdash e_2 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_2' & \vdash p_j = p \end{array}}{P; \Gamma; p \vdash e_1.m_j(e_2) : \tau' \overset{\mathsf{exp}}{\Longrightarrow} e_1'.m_j\ e_2'}$$

$$\frac{P; \Gamma; p \vdash e : \tau\ \texttt{ref}_{p'} \overset{\mathsf{exp}}{\Longrightarrow} e' \quad \vdash p \leq p'}{P; \Gamma; p \vdash !\ e : \tau \overset{\mathsf{exp}}{\Longrightarrow} !\ e'} \qquad \frac{P; \Gamma; p \vdash e_1 : \tau\ \texttt{ref}_{p'} \overset{\mathsf{exp}}{\Longrightarrow} e_1' \quad P; \Gamma; p \vdash e_2 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_2' \quad \vdash p' \leq p}{P; \Gamma; p \vdash e_1 := e_2 : \texttt{unit} \overset{\mathsf{exp}}{\Longrightarrow} e_1' := e_2'}$$

$$\frac{P; \Gamma; p \vdash e_1 : \texttt{stack} \overset{\mathsf{exp}}{\Longrightarrow} e_1' \quad P; p \vdash pat \overset{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma'; \Theta \quad P; \Gamma, \Gamma'; p \vdash e_2 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_2' \quad P; \Gamma; p \vdash e_3 : \tau \overset{\mathsf{exp}}{\Longrightarrow} e_3'}{P; \Gamma; p \vdash \texttt{case}\ e_1\ \texttt{of}\ (pat \Rightarrow e_2 \ \mid\ \_ \Rightarrow e_3) : \tau \overset{\mathsf{exp}}{\Longrightarrow} \texttt{case}\ e_1'\ \texttt{of}\ (pat' \Rightarrow \texttt{split}(\Theta, e_2') \ \mid\ \_ \Rightarrow e_3')}$$

$\boxed{\texttt{split}(\Theta, e)}$

$$\texttt{split}(\cdot, e) = e \qquad \texttt{split}(a \to (x, y, z), \Theta) = \texttt{split}(\Theta, \texttt{split}\ (x, y, z) = a\ \texttt{in}\ e)$$

$\boxed{\Gamma \dashv \Theta \Rightarrow \Gamma'}$

$$\frac{}{\Gamma \dashv \cdot \Rightarrow \Gamma} \qquad \frac{\Gamma \dashv \Theta \Rightarrow \Gamma'}{\Gamma, x : \tau, y : \tau', z : \tau'' \dashv \Theta, z \to (x, y, z) \Rightarrow \Gamma', z : (\tau \times \tau' \times \tau'')}$$

$\boxed{P; p \vdash pat \overset{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma'; \Theta}$

$$\frac{(P(o.m_i) = (\tau_{self}, \tau_{arg}, \tau_{res}, p_i))^{(1 \leq i \leq n)} \quad P; p \vdash pat \overset{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma'; \Theta}{\begin{array}{c} P; p \vdash \{o.\vec{m}\}\, [x, y, n]\ ::\ pat \overset{\mathsf{pat}}{\Longrightarrow} \{o\vec{m}_{\mathtt{pre}}\}_p\, [z]\ ::\ pat' : \\ (\Gamma', x : \tau_{self}, y : \tau_{arg}, n : \texttt{string}; \Theta, z \to (x, y, n)) \end{array}}$$

$$\frac{}{P; p \vdash nil \overset{\mathsf{pat}}{\Longrightarrow} nil \dashv \cdot; \cdot} \qquad \frac{P; p \vdash pat \overset{\mathsf{pat}}{\Longrightarrow} pat' \dashv \Gamma'; \Theta}{P; p \vdash \_ :: pat \overset{\mathsf{pat}}{\Longrightarrow} \_ :: pat' \dashv \Gamma'; \Theta} \qquad \frac{}{P; p \vdash x \overset{\mathsf{pat}}{\Longrightarrow} x \dashv \cdot, (x : \texttt{stack}); \cdot}$$

Figure 12: Translation: Part 1

$$\boxed{P;\Gamma;p \vdash as;aspcts;e:\tau \xRightarrow{\textsf{dec}} e'}$$

$$\frac{P;\Gamma;p' \vdash as;.;\texttt{()}:\texttt{unit} \xRightarrow{\textsf{as}} e' \quad P;\Gamma;p \vdash .;aspcts;e:\tau \xRightarrow{\textsf{dec}} e'' \quad \vdash p' \leq p}{P;\Gamma;p \vdash .;p':\{as\}\ aspcts;e:\tau \xRightarrow{\textsf{dec}} p'\texttt{<}e'\texttt{>};e''} \qquad \frac{P;\Gamma;p \vdash e:\tau \xRightarrow{\textsf{exp}} e'}{P;\Gamma;p \vdash .;.;e:\tau \xRightarrow{\textsf{dec}} e'}$$

$$\frac{\begin{array}{c}P;\Gamma;p \vdash e_1:\texttt{string} \xRightarrow{\textsf{exp}} e_1' \\ P;\Gamma,x:\texttt{string};p \vdash as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} e_2'\end{array}}{\begin{array}{c}P;\Gamma;p \vdash \texttt{(string}\,x=e_1\texttt{)}\,as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} \\ \texttt{let}\ x=e_1'\ \texttt{in}\ e_2'\end{array}} \quad \frac{\begin{array}{c}P;\Gamma;p \vdash e_1:\texttt{bool} \xRightarrow{\textsf{exp}} e_1' \\ P;\Gamma,x:bool;p \vdash as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} e_2'\end{array}}{\begin{array}{c}P;\Gamma;p \vdash \texttt{(bool}\,x=e_1\texttt{)}\,as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} \\ \texttt{let}\ x=e_1'\ \texttt{in}\ e_2'\end{array}}$$

$$\frac{\begin{array}{c}(P;(\Gamma,x:\tau_{self},y:\tau_i);p \vdash e_i:\tau_i' \xRightarrow{\textsf{exp}} e_i')^{1\leq i\leq n} \\ (P,(o.m_i:(\tau_{self},\tau_i,\tau_i',p))^{1..n});(\Gamma,o:\tau_{self});p \vdash as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} e_2'\end{array}}{\begin{array}{l}\quad P;\Gamma;p \vdash \texttt{(object}\,o=[m_i:\tau_i\ \rightarrow\ \tau_i'=\varsigma\,x_i.\lambda y_i.e_i]^{1..n}\texttt{)}\,as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} \\ \texttt{let}\ om_{1,\texttt{pre}}=\texttt{new}_p:(\tau_{self},\tau_1,\texttt{string})\ \texttt{in}\ \ldots\ \texttt{let}\ om_{n,\texttt{pre}}=\texttt{new}_p:(\tau_{self},\tau_n,\texttt{string})\ \texttt{in} \\ \texttt{let}\ om_{1,\texttt{post}}=\texttt{new}_p:(\tau_{self},\tau_1',\texttt{string})\ \texttt{in}\ \ldots\ \texttt{let}\ om_{n,\texttt{post}}=\texttt{new}_p:(\tau_{self},\tau_n',\texttt{string})\ \texttt{in} \\ \texttt{let}\ o=[m_i=\varsigma_p\,x_i.\lambda_p y_i:\tau_i.\quad \begin{array}[t]{l}\texttt{store}\ om_{i,\texttt{pre}}\texttt{[(}x_i,y_i,\text{“}o.m_i''\text{”}\texttt{)]}\ \texttt{in} \\ om_{i,\texttt{pre}}\texttt{[(}x_i,y_i,\text{“}o.m_i\text{“}\texttt{)]}; \\ \texttt{let}\ res_i=e_i'\ \texttt{in} \\ om_{i,\texttt{post}}\texttt{[(}x_i,res_i,\text{“}o.m_i\text{“}\texttt{)]}; \\ res_i\end{array} \\ \qquad\qquad]^{1..n}\ \texttt{in}\ e_2'\end{array}}$$

$$\texttt{where}\quad \tau_{self}\ =\ [m_i{:}_p\tau_i\ \rightarrow_p\ \tau_i']^{1..n}$$

$$\frac{P;\Gamma;p \vdash e_1:\tau \xRightarrow{\textsf{exp}} e_1' \quad P;\Gamma,x:\tau\ \texttt{ref}_p;p \vdash as;aspcts;e_2:\tau' \xRightarrow{\textsf{dec}} e_2'}{P;\Gamma;p \vdash \texttt{(ref}\,x=e_1\texttt{)}\,as;aspcts;e_2:\tau' \xRightarrow{\textsf{dec}} \texttt{let}\ x=\texttt{ref}_p\ e_1'\ \texttt{in}\ e_2'}$$

$$\frac{\begin{array}{c}(P(o.m_i)=(\tau_{self},\tau_{arg},\tau_{res},p_i))^{(1\leq i\leq n)} \\ P;(\Gamma,x:\tau_{self},y:\tau_{arg},s:\texttt{stack},n:\texttt{string});p \vdash e_1:\texttt{unit} \xRightarrow{\textsf{exp}} e_1' \quad P;\Gamma;p \vdash as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} e_2'\end{array}}{\begin{array}{l}P;\Gamma;p \vdash \texttt{(before}\ \{o.\vec{m}\}(x,y,s,n)=e_1\texttt{)}\,as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} \\ \quad \Uparrow\{\{\vec{om_{\texttt{pre}}}\}_p.z \rightarrow_p \quad \begin{array}[t]{l}\texttt{split}\ (x,y,n)=z\ \texttt{in} \\ \texttt{let}\ s=\texttt{stack()}\ \texttt{in} \\ e_1'\};e_2'\end{array}\end{array}}$$

$$\frac{\begin{array}{c}(P(o.m_i)=(\tau_{self},\tau_{arg},\tau_{res},p_i))^{(1\leq i\leq n)} \\ P;(\Gamma,x:\tau_{self},y:\tau_{res},s:\texttt{stack},n:\texttt{string});p \vdash e_1:\texttt{unit} \xRightarrow{\textsf{exp}} e_1' \quad P;\Gamma;p \vdash as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} e_2'\end{array}}{\begin{array}{l}P;\Gamma;p \vdash \texttt{(after}\ \{o.\vec{m}\}(x,y,s,n)=e_1\texttt{)}\,as;aspcts;e_2:\tau \xRightarrow{\textsf{dec}} \\ \quad \Uparrow\{\{\vec{om_{\texttt{post}}}\}_p.z \rightarrow_p \quad \begin{array}[t]{l}\texttt{split}\ (x,y,n)=z\ \texttt{in} \\ \texttt{let}\ s=\texttt{stack()}\ \texttt{in} \\ e_1'\};e_2'\end{array}\end{array}}$$

$$\boxed{\vdash ds\ \ aspcts\ \ e \xRightarrow{\textsf{prog}} e'}$$

$$\frac{.;.;main \vdash ds;aspcts;e:\tau \xRightarrow{\textsf{dec}} e'}{\vdash ds\ \ aspcts\ \ e \xRightarrow{\textsf{prog}} e'}$$

Figure 13: Translation: Part II

4. IMPLEMENTATION

We have implemented an interpreter for our language in Standard ML. There are three major elements to the implementation: the translator, the typechecker, and the evaluator. Each component has been implemented as specified in this report.

The translator parses and converts a high-level source program to a low-level core language program according to the translation rules defined in section 3.1. As we have explained, the source language program is oblivious and the programmer can write aspects independently of the mainline program. The translator will automatically insert the labels that are necessary to properly trigger the harmless advice.

The typechecker verifies that a core language program satisfies the static semantics of our language as presented in section 2.2. The protection domain lattice that it uses is created during the translation. It contains the *main* protection domain as well as separate protection domains p_i for each aspect the programmer declares. The lattice is organized so that for all i, $p_i \leq main$.

Finally, the evaluator implements the operational semantics defined in section 2.3. Due to the progress and preservation theorems of the core language, once a core language program has passed the typechecker, the evaluator will never get "stuck"

5. RELATED WORK

Over the last several years, a number of researchers have begun to build semantic foundations for aspect-oriented programming paradigms [25, 9, 14, 15, 19, 23, 24]. This foundational work provides a starting point from which one can begin to analyze the properties of aspect-oriented programs, develop principled new programming features, study verification techniques and derive useful type systems. In this paper, our semantic foundations were derived directly from earlier work by Walker, Zdancewic and Ligatti [24]. The main novelty with respect to this earlier research is the development of a type system for ensuring that aspects do not interfere with each other or the mainline computation.

Clifton and Leavens [5] proposed techniques for Hoare-style reasoning about aspect-oriented programs using *assistants* and *observers*. Their notion of observers is similar to our conception of harmless advice — observers do not interfere with the mainline computation. However, the details of our type and effect system are entirely different from their Hoare logic. One point of interest is that Clifton and Leavens mention that it is not clear whether their model can "accommodate dynamic context join points like CFlow." Our analysis of our stack operations, which are sufficient for coding up CFlow-like primitives, indicates that harmless advice can indeed safely use these primitives and avoid interfering with the mainline computation or each other.

Several authors have looked specifically at techniques for explicitly combining several pieces of advice and detecting interference between them. For instance, Bauer, Ligatti and Walker [3] introduced a calculus that included several different kinds of aspect combinators (parallel conjunction and disjunction; sequenced conjunction and disjunction) and used a type and effect system to prevent interference between them. The technical machinery used here was extremely complicated and quite different from the current work. In contrast to our work here, they did not concern themselves with the effects these aspects would have on the mainline computation. Recently, Bauer, Ligatti and Walker have completed the implementation of a general-purpose, higher-order language for composing aspects in the context of Java [4].

In similar work, Douence, Fradet and Südholt [10, 11] analyze aspects defined by recursion together with parallel and sequencing combinators. They develop a number of formal laws for reasoning about their combinators and an algorithm that is able to detect *strong independence*. Two pieces of advice are strongly independent when they do not interfere with each other regardless of the contents of the advice bodies or the contents of the programs they are applied to. In other words, strong independence is determined exclusively by analysis of the point cut designators of the two pieces of advice and consequently it is orthogonal to our analysis which (mostly) ignores the point cuts and examines the advice bodies instead. It would be interesting to explore how to put these two different ideas together.

Krishnamurthi, Fisler and Greenberg [17] tackle the more general problem of verifying aspect-oriented programs. Given a set of properties a program must satisfy, specified in a temporal logic, and a set of point cut designators, they verify programs using model checking. Their approach to verification is partly modular since as long as the set of point cuts does not change and the underlying program remains fixed, it is not necessary to reanalyze the underlying program as advice definitions are edited. Aside from the fact that we are both interested in modular checking of aspect-oriented programs, there is not too much similarity between the techniques. In terms of trade-offs, our approach is lighter weight (temporal specifications of properties are unnecessary) and more modular (changing point cut designators used by advice does not necessitate re-type checking the mainline program), but checks a much coarser-grained property (we guarantee that *all* functional properties of advice are preserved) and applies in fewer situations (for instance, we cannot handle around advice). It would be interesting to explore how these two different ideas could be combined in a single system.

Another interesting line of current research involves finding ways to add aspect-oriented programming features to languages with module systems, or vice-versa. The goal of this research is often quite similar to our own work: To find mechanisms to protect the internals of a module from outside interference by advice. However, the techniques used and resulting properties are quite different. One of the first systems to combine aspects and modules effectively was Lieberherr, Lorenz and Ovlinger's *Aspectual Collaborations* [18, 21]. Their proposal allows module programmers to choose the join points (i.e., control-flow points) that they will expose to external advice. External advice cannot intercept control-flow points that have not been exposed. In this sense, Aspectual Collaborations are not completely oblivious – programmers must make choices about which control-flow points to expose upfront, during program design. Aspectual collaborations do enjoy a number of important properties including strong encapsulation, type safety and the possibility of separately compiling and checking module definitions. Aldrich [2] has proposed another model for combining aspects and modules called *Open Modules*. The central novelty of this proposal is a special module sealing operator that hides internal control-flow points from external advice. Aldrich has used logical relations to show that sealed modules have a powerful implementation-independence property [1]. In an earlier report [7], we suggested augmenting these proposals with access-control specifications in the module interfaces that allow programmers to specify whether or not data at join points may be read or written. The current report differs from this previous research as it does not suggest that visibility of the interception points be limited; instead, we suggest limiting the capabilities of advice. However, once again, it seems quite likely that one could design a powerful system that combines both ideas.

6. CONCLUSIONS

In this paper, we have investigated the idea of *harmless advice*: aspect-oriented advice that does not interfere with the mainline computation. While strictly less powerful than ordinary advice, we believe that harmless advice can be used in many contexts including security monitoring, profiling, logging, and for some debugging tasks. Harmless advice has the advantage that it may be added to a program after-the-fact, in the typical aspect-oriented style, yet programmers do not have to worry about it corrupting important mainline data invariants.

There are a number of directions for future research, several of which we are currently working on. Our first priority is to extend the calculus to include additional features common to object-oriented languages. In particular, we are interested in adding Java-style classes to the language and granting our aspect language the capability to externally extend classes while maintaining appropriate noninterference properties. In addition, together with Washburn and Weirich [8], we are investigating how to extend the calculus with polymorphic functions, polymorphic advice and type analysis to support safe but flexible type-directed aspect-oriented programming.

Acknowledgments

Stimulating discussions on aspects and related topics with Geoff Washburn, Stephanie Weirich and Steve Zdancewic helped form some of the ideas in this paper.

This research was supported in part by ARDA Grant no. NBCHC030106, National Science Foundation CAREER grant No. CCR-0238328 and an Alfred P. Sloan Fellowship. This work does not necessarily reflect the opinions or policy of the federal government or Sloan foundation and no official endorsement should be inferred.

7. REFERENCES

[1] J. Aldrich. Open modules: A proposal for modular reasoning in aspect-oriented programming. In *Workshop on foundations of aspect-oriented languages*, Mar. 2004.

[2] J. Aldrich. Open modules: Reconciling extensibility and information hiding. In *Proceedings of the Software Engineering Properties of Languages for Aspect Technologies*, Mar. 2004.

[3] L. Bauer, J. Ligatti, and D. Walker. Types and effects for non-interfering program monitors. In *International Symposium on Software Security*, Tokyo, Japan, Nov. 2002.

[4] L. Bauer, J. Ligatti, and D. Walker. A language and system for composing security policies. Technical Report TR-699-04, Princeton University, Jan. 2004.

[5] C. Clifton and G. T. Leavens. Assistants and observers: A proposal for modular aspect-oriented reasoning. In *Foundations of Aspect Languages*, Apr. 2002.

[6] A. Colyer and A. Clement. Large-scale aosd for middleware. In *Proceedings of the 3rd international conference on Aspect-oriented software development*, pages 56–65. ACM Press, 2004.

[7] D. S. Dantas and D. Walker. Aspects, information hiding and modularity. Technical Report TR-696-04, Princeton University, Nov. 2003.

[8] D. S. Dantas, D. Walker, G. Washburn, and S. Weirich. Analyzing polymorphic advice. Technical Report TR-717-04, Princeton University, Dec. 2004.

[9] R. Douence, O. Motelet, and M. Südholt. A formal definition of crosscuts. In *Third International Conference on Metalevel architectures and separation of crosscutting concerns*, volume 2192 of *Lecture Notes in Computer Science*, pages 170–186, Berlin, Sept. 2001. Springer-Verlag.

[10] R. Douence, O. Motelet, and M. Südholt. Detection and resolution of aspect interactions. Technical Report 4435, INRIA, Apr. 2002.

[11] R. Douence, O. Motelet, and M. Südholt. Composition, reuse and interaction analysis of stateful aspects. In *Conference on Aspect-Oriented Software Development*, pages 141–150, Mar. 2004.

[12] R. E. Filman and D. P. Friedman. Aspect-oriented programming is quantification and obliviousness. In *Workshop on Advanced Separation of Concerns*, Oct. 2000.

[13] List of main users. AspectJ Users List: aspectj-users@eclipse.org, June 2004. Requires subscription to access archives.

[14] R. Jagadeesan, A. Jeffrey, and J. Riely. A calculus of typed aspect-oriented programs. Unpublished manuscript., 2003.

[15] R. Jagadeesan, A. Jeffrey, and J. Riely. A calculus of untyped aspect-oriented programs. In *European Conference on Object-Oriented Programming*, Darmstadt, Germany, July 2003.

[16] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. Griswold. An overview of AspectJ. In *European Conference on Object-oriented Programming*. Springer-Verlag, 2001.

[17] S. Krishnamurthi, K. Fisler, and M. Greenberg. Verifying aspect advice modularly. In *Foundations of Software Engineering*, Oct.-Nov. 2004.

[18] K. J. Lieberherr, D. Lorenz, and J. Ovlinger. Aspectual collaborations – combining modules and aspects. *The Computer Journal*, 46(5):542–565, September 2003.

[19] H. Masuhara, G. Kiczales, and C. Dutchyn. Compilation semantics of aspect-oriented programs. In G. T. Leavens and R. Cytron, editors, *Foundations of Aspect-Oriented Languages Workshop*, pages 17–25, Apr. 2002.

[20] A. Myers and B. Liskov. Jflow: Practical mostly-static information flow control. In *Twenty-Sixth ACM Symposium on Principles of Programming Languages*, pages 226–241, Jan. 1998.

[21] J. Ovlinger. *Modular Programming with Aspectual Collaborations.* PhD thesis, Northeastern University, 2003.

[22] F. Pottier and V. Simonet. Information flow inference for ML. *ACM Transactions on Programming Languages and Systems*, 25(1):117–158, Jan. 2003.

[23] D. B. Tucker and S. Krishnamurthi. Pointcuts and advice in higher-order languages. In *Proceedings of the 2nd International Conference on Aspect-Oriented Software Development*, pages 158–167, 2003.

[24] D. Walker, S. Zdancewic, and J. Ligatti. A theory of aspects. In *ACM International Conference on Functional Programming*, Uppsala, Sweden, Aug. 2003.

[25] M. Wand, G. Kiczales, and C. Dutchyn. A semantics for advice and dynamic join points in aspect-oriented programming. In G. T. Leavens and R. Cytron, editors, *Foundations of Aspect-Oriented Languages Workshop*,

pages 17–25, Apr. 2002. Iowa State University University technical report 02-06.