Towards a Practical Secure Concurrent Language

Stefan Muller

Carnegie Mellon University

smuller@cs.cmu.edu

Stephen Chong

Harvard University

chong@seas.harvard.edu

Abstract

We demonstrate that a practical concurrent language can be extended in a natural way with information security mechanisms that provably enforce strong information security guarantees. We extend the X10 concurrent programming language with coarse-grained information-flow control. Central to X10 concurrency abstractions is the notion of a *place*: a container for data and computation. We associate a security level with each place, and restrict each place to store only data appropriate for that security level. When places interact only with other places at the same security level, then our security mechanisms impose no restrictions. When places of differing security levels interact, our information security analysis prevents potentially dangerous information flows, including information flow through covert scheduling channels. The X10 concurrency mechanisms simplify reasoning about information flow in concurrent programs.

We present a static analysis that enforces a noninterference-based extensional information security condition in a calculus that captures the key aspects of X10's place abstraction and async-finish parallelism. We extend this security analysis to support many of X10's language features, and have implemented a prototype compiler for the resulting language.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features—Concurrent programming structures; D.4.6 [*Operating Systems*]: Security and Protection—Information flow controls

Keywords Language-based security; information-flow control; X10.

1. Introduction

Enforcement of strong information security guarantees for concurrent programs poses both a challenge and an oppor-

tunity. The challenge is that, given current hardware trends towards increased parallelism, and the large number of computer systems that handle data of varying sensitivity, it is increasingly important to reason about and enforce information security guarantees in the presence of concurrency. Although progress has been made towards this end, there are not yet practical enforcement mechanisms and usable implementations. The opportunity is to adapt new and existing language abstractions for concurrency to reason precisely about information security in concurrent programs. Information security, like concurrency, is intimately connected to notions of dependency [1]. As such, there is potential for synergy between language mechanisms for concurrency, and enforcement mechanisms for information security in concurrent programs.

The X10 programming language [10] is an object-oriented language with abstractions to support fine-grained concurrency. Central to X10 concurrency abstractions is the notion of a *place*. A place is a computational unit that contains computation and data. For example, each core of a single machine or each machine within a distributed system might be represented by a different place. Data held at a place, and computation running at a place, are said to be "local" to the place. Places are first-class in X10. Multiple threads of execution, which in X10 are known as *activities*, may execute concurrently within a place. Activities at the same place share memory, and an activity may only access data at the place where it is located. Places may communicate using message passing, but X10 is designed to discourage excessive communication between places, since this reduces concurrency.

We extend X10 language abstractions for concurrency with information security mechanisms, and call the resulting language SX10 (for Secure X10). Specifically, in SX10 each place is associated with a security level, and a (completely static) security analysis ensures that each place stores only data appropriate for that security level. Thus, all computation within a place is on data at the same security level. In the case where places communicate only with other places at the same security level, then our security mechanisms do not impose any restrictions on programs. Communication between places with different security levels may pose security concerns, but because message-passing communication is used

between places, it is relatively simple to restrict such communication: the security analysis ensures that data may be sent to another place only when the security level of the destination is sufficiently high. Interaction between places may influence the scheduling of activities at a place, leading to potential covert information channels; our security analysis tracks and controls these covert channels.

We believe that this coarse-grained approach to providing information security in concurrent programs is simple, practical, and useful. All data at a place is at the same security level, which both provides simplicity of reasoning for the programmer, and allows a high degree of concurrency within a place without compromising security.

There are many highly concurrent systems that compute with data of varying sensitivity that fit naturally into such a model. The following are some examples.

- **Machine learning** A service such as Pandora processes a large amount of public data, which is then used to make recommendations to individual users based on their private usage data. All public data is at the same security level and processing it is highly parallel; data from many users are processed in parallel.

- **Social networks** Users specify that some posts are visible to all other users, and some are visible only to friends. Many users may use the system concurrently.

- **Shopping carts** An online shopping cart collects information about items ordered, which may appear in low-security logs, and credit card information, which must remain secure. Many customers may use the system concurrently.

Motivating example Reasoning about information security in the presence of concurrency can be subtle. Consider Program 1, which exhibits a timing channel.

```
1   at Low {
2       async {
3           // Activity 1
4           if (hi > 0) longComputation();
5           output ''pos'';
6       }
7       // Activity 2
8       mediumComputation(); output ''nonpos'';
9   }
```

Program 1.

Assume that memory location hi contains high-security (e.g., secret) data. Instruction at Low s indicates that s is executed at a place called Low, which we assume is allowed to handle only low-security data. Outputs at place Low should not reveal high-security information. Instruction async s creates a new activity to execute statement s, and the current activity continues with the following statement. Thus, the

if statement and subsequent output of the string ''pos'' execute concurrently with the output ''nonpos'' statement. If high-security memory location hi is positive, then Activity 1 outputs ''pos'' after a long time; otherwise it outputs ''pos'' immediately. Activity 2 computes for a medium amount of time, and outputs ''nonpos''. It is likely that the order of outputs will reveal secret information, and the program is thus insecure. This is an example of an *internal timing channel* [48], where the order of publicly observable events depends upon high-security information.

Program 1 is not an SX10 program: low-security place Low is not allowed to hold high-security data, such as that stored in memory location hi. Suppose, however, that High is a place that is permitted to hold secret information. Then Program 2 is a SX10 program and exhibits a similar timing channel. (It is correctly rejected by our security analysis.)

```
1   at Low {
2       async {
3           // Activity 1
4           at High { // activity 1 moves to High
5               if (hi > 0) longComputation();
6           }
7           // activity 1 moves back to Low
8           output ''pos'';
9       }
10      // Activity 2
11      mediumComputation(); output ''nonpos'';
12  }
```

Program 2.

In this example, Activity 1 moves to place High in order to perform computation on high-security data, before returning to place Low to perform the output of the string ''pos''.

We assume that the scheduling of activities at each place depends only on the activities at that place, an assumption that holds in the X10 2.2 runtime [14]. Nonetheless, in Program 2, the scheduling of Activity 2 at place Low depends on whether Activity 1 is running at Low, which in turn depends on how long the computation at place High takes. Thus, the scheduling of output ''pos'' and output ''nonpos'' may be influenced by high-security data. Our security analysis detects this potential information flow, and rejects the program.

Program 2 is inherently nondeterministic: it could perform the two outputs in either order. The scheduler resolves the nondeterminism, but in doing so, may reveal high-security information—a form of refinement attack [39]. One way to prevent such refinement attacks is to require that any observable behavior be deterministic [6, 49]. Our security analysis requires such observational determinism when the resolution of nondeterminism may reveal high-security information.

It is, however, possible to allow some observable nondeterminism within a secure concurrent program. Intuitively, if the resolution of the nondeterminism does not depend

on high-security information, then observable nondeterminism is secure [31]. If place P does not communicate with any other places, then, since scheduling is performed per-place, resolution of nondeterminism at P will not reveal high-security information. In some cases it is also possible to allow nondeterminism at place P even if P interacts with places of higher security levels. Consider Program 3: Activity 1 executes at place Low concurrently with Activity 2's execution at place High. The finish s instruction executes statement s, and waits until s and all activities spawned by s have terminated before continuing execution. Thus, Activity 3 and Activity 4 execute concurrently at place Low after Activities 1 and 2 have finished. Although the order of output ''B'' and output ''C'' is nondeterministic, resolution of this nondeterminism is not influenced by how long the high-security computation takes, and so does not reveal high-security information.

```
1    at Low {
2       finish {
3          async {
4             // Activity 1
5             output ''A'';
6          }
7          // Activity 2
8          at High {
9             if (hi > 0) longComputation();
10         }
11      }
12      async {
13         // Activity 3
14         output ''B'';
15      }
16      // Activity 4
17      output ''C'';
18   }
```

Program 3.

Contributions The key contribution of this paper is to demonstrate that practical and useful concurrency mechanisms can be extended in a natural way with information security mechanisms that provably enforce strong information security guarantees. We enforce coarse-grained information-flow control [39], requiring that every place can store only data at a single security level. If places interact only with other places at the same security level, then our security mechanisms do not restrict concurrency nor require determinism for security. When places of differing security interact, our information security analysis prevents potentially dangerous information flows by using X10's concurrency mechanisms to reason both about data sent between places, and about how the scheduling of activities at a place may depend on high-security information.

In Section 2 we present a calculus, based on Featherweight X10 [22], that captures key aspects of the X10 place abstraction, and its async-finish parallelism. We define a knowledge-based noninterference semantic security condition [2, 12] for this calculus in Section 3, and present a security analysis that provably enforces it. The language SX10 is the result of extending this analysis to handle many of the language features of X10. We have implemented a prototype compiler for SX10 by modifying the X10 compiler, and this is described in Section 4. We discuss related work in Section 5 and conclude in Section 6.

2. FSX10: a secure parallel calculus

In this section, we introduce the calculus FSX10, based on Featherweight X10 [22]. Like Featherweight X10, this calculus captures X10's async-finish parallelism, but adds places and interaction with the external environment via input and output instructions.

2.1 Syntax

The abstract syntax of FSX10 is presented in Figure 1. A *place* P is a container for data and activities. In FSX10, as in X10, every memory location r and every activity is associated with a place. In FSX10, however, places are simply identifiers and are not first-class values. Function $Place(\cdot)$ describes how memory locations are mapped to places: memory location r is held at place $Place(r)$, and only code executing at that place is allowed to access the location.

For simplicity, we restrict values in the calculus to integers. Expressions e consist of integer constants v, variables x, memory reads $!r$ (where r is a memory location), and total binary operations over expressions $e_1 \oplus e_2$.

Statements s are sequences of instructions. Every instruction is labeled with a program point. For example, a store instruction $r :=^p e$ has program point p. For convenience we write s^p to indicate that program point p is the program point of the first instruction of statement s. When the program point of an instruction is irrelevant, we omit it.

Instructions include no-ops (skip), selection (if e then s else s), and iteration (while e do s). Instruction $r := e$ evaluates expression e and updates memory location r with the result. Instruction let $x = e$ in s evaluates expression e to a value, and uses that value in place of variable x in the evaluation of statement s. Once defined, variable x is immutable. A variable defined at place P may be used at a different place P', which can be thought of as P sending the value of the variable to P'.

Instruction async s creates a new activity that starts executing statement s, and the current activity continues executing the next instruction. Instruction at P s executes statement s at place P. Note that at P s does not create a new activity: execution of the next instruction waits until s has terminated. That is, given the statement at P s; $r := 42$; skip, the assignment of 42 to memory location r will not occur until after statement s has finished execution. Instruction backat $P s$ does not appear in source programs, but is used by

Metavariables

P ranges over places
p ranges over program points
v ranges over integer constants
x ranges over program variables
r ranges over memory locations
$\oplus$ ranges over total binary integer functions

Expressions

e	$::=$	v	Integer constant
	$\mid$	x	Variable
	$\mid$	$!r$	Memory read
	$\mid$	$e \oplus e$	Binary operation

Statements

| s | $::=$ | skip^p | No-op |
| | $\mid$ | $i; s$ | Sequence |

Instructions

i	$::=$	skip^p	No-op
	$\mid$	$r :=^p e$	Memory write
	$\mid$	$\mathsf{if}^p\ e\ \mathsf{then}\ s\ \mathsf{else}\ s$	Conditional
	$\mid$	$\mathsf{while}^p\ e\ \mathsf{do}\ s$	Iteration
	$\mid$	$\mathsf{let}^p\ x = e\ \mathsf{in}\ s$	Let
	$\mid$	$\mathsf{output}^p\ e$	Output
	$\mid$	$\mathsf{input}^p\ r$	Input
	$\mid$	$\mathsf{async}^p\ s$	Asynchronous
	$\mid$	$\mathsf{at}^p\ P\ s$	At
	$\mid$	$\mathsf{finish}^p\ s$	Finish
	$\mid$	$\mathsf{backat}^p\ P$	Back at

Trees

T	$::=$	$\langle P, s \rangle$	Activity
	$\mid$	$T \parallel T$	Parallel
	$\mid$	$T \triangleright \langle P, s \rangle$	Join
	$\mid$	$\checkmark$	Done

Figure 1. FSX10 syntax

the operational semantics to track when control will return back to place P as a result of finishing an $\mathsf{at}\ P'\ s'$ instruction. Finally, instruction $\mathsf{finish}\ s$ will block until statement s, and all activities created by s, have terminated.

FSX10 programs can communicate with the external environment via input and output instructions. We assume, without loss of generality, that every place has a single communication channel. Instruction $\mathsf{output}\ e$, when executed at place P, will evaluate expression e to a value, and output that value on P's channel. Similarly, instruction $\mathsf{input}\ r$, when executed at P, will input a value from P's channel, and store the result in location r. We assume that there is always data

available for input on a channel, and thus input instructions are non-blocking.

Concurrently executing activities in FSX10 are represented using *trees*. Tree $\langle P, s \rangle$ is an activity at place P executing statement s. Tree $T_1 \parallel T_2$ represents trees T_1 and T_2 executing concurrently. Tree $\checkmark$ indicates a terminated activity, and tree $T \triangleright \langle P, s \rangle$ indicates that activity $\langle P, s \rangle$ is blocked until all activities in T have terminated.

2.2 Events, traces, and input strategies

As a program executes, it generates input and output events. Input event $i(v, P)$ is generated when an input instruction accepts value v from P's channel. Output event $o(v, P)$ is generated when an output instruction outputs value v on P's channel.

A *trace* t is a (possibly empty) sequence of input, output and location assignment events. Other events are not tracked. We write ϵ for the empty trace. We write $t\!\restriction_P$ for the subsequence of events of t that occur at place P. More formally, we have

$$\epsilon\!\restriction_P = \epsilon$$

$$(t \cdot \alpha)\!\restriction_P = \begin{cases} (t\!\restriction_P) \cdot \alpha & \text{if } Place(\alpha) = P \\ t\!\restriction_P & \text{otherwise} \end{cases}$$

where function $Place(\alpha)$ is the place at which α occurred:

$$Place(i(v, P)) = P$$
$$Place(o(v, P)) = P.$$

We model input from the external environment with *input strategies* [31]. Input strategy ω is a function from places and traces to values, such that given trace t, value $\omega(P, t\!\restriction_P)$ is the next value to input on the channel for P. Note that the choice of the next value that will be input on a channel can depend on the previous outputs of the channel. In Section 3, where we consider the security of FSX10 programs, we will be concerned with ensuring that low-security attackers are unable to learn about the inputs to high-security places.

2.3 Scheduling

Since program execution, and information security, depends on scheduling, we model the scheduler in FSX10. We explicitly refine the nondeterminism inherent in scheduling using *refiners* [31] to represent the decisions made by the scheduler. Essentially, all nondeterminism in program execution is encapsulated in a refiner; once a refiner has been chosen, program execution is deterministic.

In X10, a place represents a distinct computational node with a distinct scheduler [14]. In accordance with this model, we assume that scheduling decisions are made on a per-place basis, and the choice of which activity to run at a given place depends only on the set of activities currently executing at that place.

$$\frac{Sch(P) = chs \cdot ch \qquad \mathsf{PointsRunning}(T, P) \neq \emptyset}{ch(\mathsf{PointsRunning}(T, P)) = p \qquad (H; \omega; t; T) \xrightarrow{p} (H'; \omega; t'; T')}$$
$$\overline{(H; \omega; (P \cdot Ps, Sch); t; T) \to (H'; \omega; (Ps, Sch[P \to chs]); t'; T')}$$

IDLEPLACE

$$\frac{\mathsf{PointsRunning}(T, P) = \emptyset}{(H; \omega; (P \cdot Ps, Sch); t; T) \to (H; \omega; (Ps, Sch); t; T)}$$

$$\mathsf{PointsRunning}(\checkmark, P) = \emptyset$$

$$\mathsf{PointsRunning}(\langle P', s^p \rangle, P) = \begin{cases} \{p\} & \text{if } P' = P \\ \emptyset & \text{if } P' \neq P \end{cases}$$

$$\mathsf{PointsRunning}(T_1 \parallel T_2, P) = \mathsf{PointsRunning}(T_1, P) \cup \mathsf{PointsRunning}(T_2, P)$$

$$\mathsf{PointsRunning}(T \rhd \langle P', s \rangle, P) = \mathsf{PointsRunning}(T, P)$$

Figure 2. Program semantics $(H; \omega; R; t; T) \to (H'; \omega; R'; t'; T')$

We model these assumptions by representing a refiner R as a pair (Ps, Sch), where Ps is a stream of places indicating the order in which places take steps, and Sch is a function from places to streams chs of *scheduling functions*. A scheduling function ch takes a set of program points (representing the set of activities currently executing at the place), and returns an element of that set (representing which of the activities should be scheduled). We write $P \cdot Ps$ for a stream with first element P and remaining elements Ps.

Thus, if the refiner is $(P \cdot Ps, Sch)$, then place P will take a step next, and if $Sch(P) = ch \cdot chs$ (where ch is the first element of the stream of scheduling function, and chs is the remainder of the stream), then scheduling function ch will be used to determine which of the current activities at P will be scheduled. Note that each time a place takes a step, it may use a different scheduling function. However, the sequence of scheduling functions at a given place must be decided in advance, and may not depend on the history of computation at the place.

The use of a stream of scheduling functions per place allows our model to capture many realistic scheduling algorithms, such as round robin, shortest remaining time, and fixed priority. Scheduling algorithms that depend on the history of computation at a place (such as the work-stealing scheduling algorithm used in the X10 runtime [14, 15]) cannot be directly represented in this model. However, we believe the security guarantees still hold for the X10 runtime; we further discuss the security of the X10 scheduler in Section 4.

2.4 Operational semantics

A program configuration is a 5-tuple $(H; \omega; R; t; T)$. Heap H maps locations r to values, and is updated as the program executes. Input strategy ω is used to determine values input on channels; the input strategy does not change during execution, but we include it in the program configuration for notational convenience. Refiner R is used to determine scheduling, and is updated during execution. Trace t is the trace (of input, output, and location assignment events) pro-

duced so far by the program's execution. Tree T is the tree of currently executing activities.

The small-step operational semantics relation

$$(H; \omega; R; t; T) \to (H'; \omega; R'; t'; T')$$

describes how a program configuration changes as a result of execution. Due to the use of refiners, the operational semantics is deterministic. Inference rules for this relation are given in Figure 2.

Rule PLACE uses the refiner to select a place P to execute, and to select a scheduling function ch to schedule an activity at P. Set $\mathsf{PointsRunning}(T, P)$ is the set of program points of running activities located at P (also defined in Figure 2), which is given to scheduling function ch to select an activity. Rule IDLEPLACE handles the case where the refiner has selected place P to execute, but P does not have any currently running activities. Judgment

$$(H; \omega; t; T) \xrightarrow{p} (H'; \omega; t'; T')$$

is used to indicate that tree configuration $(H; \omega; t; T)$ executes the instruction at program point p to produce tree configuration $(H'; \omega; t'; T')$. Tree configurations are similar to program configurations, but omit the refiner, since the refiner is used only to determine which activity to execute.

Inference rules for $(H; \omega; t; T) \xrightarrow{p} (H'; \omega; t'; T')$ are given in Figure 3. Rules PARALEFT, PARARIGHT, PARALEFTDONE, PARARIGHTDONE, JOIN, and JOINDONE navigate through the tree structure to find the appropriate activity to execute. Rule SKIP1 reduces a skip statement to a terminated activity $\checkmark$. The remaining rules execute a single instruction.

Several of the rules for evaluating instructions evaluate expressions to values, using judgment $P; H; e \Downarrow v$, which is defined in Figure 4. Evaluation of expressions is standard, with the exception of memory read $!r$, which requires that memory location r is held at place P, the current place of the activity performing the read.

PARALEFT

$$\frac{(H;\,\omega;\,t;\,T_1) \xrightarrow{p} (H';\,\omega;\,t';\,T_1') \qquad T_1' \neq \checkmark}{(H;\,\omega;\,t;\,T_1\,\|\,T_2) \xrightarrow{p} (H';\,\omega;\,t';\,T_1'\,\|\,T_2)}$$

PARARIGHT

$$\frac{(H;\,\omega;\,t;\,T_2) \xrightarrow{p} (H';\,\omega;\,t';\,T_2') \qquad T_2' \neq \checkmark}{(H;\,\omega;\,t;\,T_1\,\|\,T_2) \xrightarrow{p} (H';\,\omega;\,t';\,T_1\,\|\,T_2')}$$

PARALEFTDONE

$$\frac{(H;\,\omega;\,t;\,T_1) \xrightarrow{p} (H';\,\omega;\,t';\,\checkmark)}{(H;\,\omega;\,t;\,T_1\,\|\,T_2) \xrightarrow{p} (H';\,\omega;\,t';\,T_2)}$$

PARARIGHTDONE

$$\frac{(H;\,\omega;\,t;\,T_2) \xrightarrow{p} (H';\,\omega;\,t';\,\checkmark)}{(H;\,\omega;\,t;\,T_1\,\|\,T_2) \xrightarrow{p} (H';\,\omega;\,t';\,T_1)}$$

JOIN

$$\frac{(H;\,\omega;\,t;\,T) \xrightarrow{p} (H';\,\omega;\,t';\,T') \qquad T' \neq \checkmark}{(H;\,\omega;\,t;\,T \triangleright \langle P,s\rangle) \xrightarrow{p} (H';\,\omega;\,t';\,T' \triangleright \langle P,s\rangle)}$$

JOINDONE

$$\frac{(H;\,\omega;\,t;\,T) \xrightarrow{p} (H';\,\omega;\,t';\,\checkmark)}{(H;\,\omega;\,t;\,T \triangleright \langle P,s\rangle) \xrightarrow{p} (H';\,\omega;\,t';\,\langle P,s\rangle)}$$

SKIP1

$$\frac{}{(H;\,\omega;\,t;\,\langle P,\mathsf{skip}^p\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\checkmark)}$$

SKIP2

$$\frac{}{(H;\,\omega;\,t;\,\langle P,\mathsf{skip}^p;s\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,s\rangle)}$$

WRITE

$$\frac{P;H;e \Downarrow v \qquad Place(r) = P}{(H;\,\omega;\,t;\,\langle P,r :=^p e;s\rangle) \xrightarrow{p} (H[r \mapsto v];\,\omega;\,t;\,\langle P,s\rangle)}$$

LET

$$\frac{P;H;e \Downarrow v \qquad s_1' = s_1\{v/x\}}{(H;\,\omega;\,t;\,\langle P,\mathsf{let}^p\, x = e\,\mathsf{in}\,s_1;s_2\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,s_1'\bullet s_2\rangle)}$$

ASYNC

$$\frac{}{(H;\,\omega;\,t;\,\langle P,\mathsf{async}^p\,s_1;s_2\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,s_1\rangle\,\|\,\langle P,s_2\rangle)}$$

FINISH

$$\frac{}{(H;\,\omega;\,t;\,\langle P,\mathsf{finish}^p\,s_1;s_2\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,s_1\rangle \triangleright \langle P,s_2\rangle)}$$

AT

$$\frac{}{(H;\,\omega;\,t;\,\langle P_1,\mathsf{at}^p\,P_2\,s_1;s_2\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P_2,s_1\bullet(\mathsf{backat}\,P_1;s_2)\rangle)}$$

BACKAT

$$\frac{}{(H;\,\omega;\,t;\,\langle P_2,\mathsf{backat}^p\,P_1;s\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P_1,s\rangle)}$$

OUTPUT

$$\frac{P;H;e \Downarrow v}{(H;\,\omega;\,t;\,\langle P,\mathsf{output}^p\,e;s\rangle) \xrightarrow{p} (H;\,\omega;\,t \cdot o(v,P);\,\langle P,s\rangle)}$$

INPUT

$$\frac{Place(r) = P \qquad \omega(P,t{\restriction}_P) = v}{(H;\,\omega;\,t;\,\langle P,\mathsf{input}^p\,r;s\rangle) \xrightarrow{p} (H[r \mapsto v];\,\omega;\,t \cdot i(v,P);\,\langle P,s\rangle)}$$

IF1

$$\frac{P;H;e \Downarrow v \qquad v \neq 0}{(H;\,\omega;\,t;\,\langle P,\mathsf{if}^p\,e\,\mathsf{then}\,s_1\,\mathsf{else}\,s_2;s_3\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,s_1\bullet s_3\rangle)}$$

IF2

$$\frac{P;H;e \Downarrow v \qquad v = 0}{(H;\,\omega;\,t;\,\langle P,\mathsf{if}^p\,e\,\mathsf{then}\,s_1\,\mathsf{else}\,s_2;s_3\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,s_2\bullet s_3\rangle)}$$

WHILE

$$\frac{}{(H;\,\omega;\,t;\,\langle P,(\mathsf{while}^p\,e\,\mathsf{do}\,s_1);s_2\rangle) \xrightarrow{p} (H;\,\omega;\,t;\,\langle P,(\mathsf{if}\,e\,\mathsf{then}\,(s_1\bullet\mathsf{while}\,e\,\mathsf{do}\,s_1;\mathsf{skip})\,\mathsf{else}\,\mathsf{skip});s_2\rangle)}$$

Figure 3. Tree and statement semantics $(H;\,\omega;\,t;\,T) \xrightarrow{p} (H';\,\omega;\,t';\,T')$

$$\frac{\text{CONST}}{P; H; v \Downarrow v} \qquad \frac{\text{READ} \quad H(r) = v \qquad Place(r) = P}{P; H; !r \Downarrow v}$$

$$\frac{\text{OP} \quad P; H; e_1 \Downarrow v_1 \qquad P; H; e_2 \Downarrow v_2}{P; H; e_1 \oplus e_2 \Downarrow v_1 \oplus v_2}$$

Figure 4. Expression semantics $P; H; e \Downarrow v$

Rule SKIP2 handles the instruction skip—it is a no-op. Rule WRITE executes write instruction $r :=^p e$ by evaluating expression e to a value v and updating the heap to map location r to v. Note that location r must be stored at the place at which the activity is executing: $Place(r) = P$.

Rule LET executes let instruction let $x = e$ in s by evaluating expression e to value v, and substituting uses of variable x in s with v using capture-avoiding substitution $s\{v/x\}$. The rule uses the operation $\bullet$ to "stitch together" two statements into a single statement. This operation is defined recursively as follows.

$$(i; s_1) \bullet s_2 = i; (s_1 \bullet s_2)$$
$$\text{skip} \bullet s_2 = s_2$$

Instruction async s_1 creates a new activity to execute s_1, and the current activity continues with the next statement. Thus, rule ASYNC executes the activity $\langle P, \text{async } s_1; s_2 \rangle$ by reducing it to the tree $\langle P, s_1 \rangle \,||\, \langle P, s_2 \rangle$.

Statement finish $s_1; s_2$ executes s_1, and waits until all activities spawned by s_1 have terminated before executing s_2. Rule FINISH transforms activity $\langle P, \text{finish } s_1; s_2 \rangle$ to the tree $\langle P, s_1 \rangle \rhd \langle P, s_2 \rangle$.

Statement at $P \, s_1; s_2$ executes statement s_1 at place P, and then executes s_2 at the original place. Rule AT thus transforms activity $\langle P', \text{at } P \, s_1; s_2 \rangle$ to an activity at place P: $\langle P, s_1 \bullet (\text{backat } P'; s_2) \rangle$. We insert the instruction backat P' to let us know both that execution of s_2 will be at place P', and that the movement of the activity to P' is the result of returning from a previous at instruction. Rule BACKAT for statement backat P' simply changes the place of the activity back to place P'.

Rule OUTPUT evaluates output instruction output e by evaluating e to value v, and appending event $o(v, P)$ to the program's trace, where P is the current place of the activity. Similarly, rule INPUT evaluates input instruction input r by inputting value v from P's communication channel, updating the heap to map location r to v, and appending event $i(v, P)$ to the program's trace. The value to input is determined by input strategy ω, and is equal to $\omega(P, t\restriction_P)$, where P is the current place of the activity, and $t\restriction_P$ is the program's trace so far restricted to events occurring at place P.

Rules IF1 and IF2 handle the conditional instruction if e then s_1 else s_2 by reducing it to s_1 if e evaluates to a nonzero value, and reducing it to s_2 otherwise. Rule WHILE handles a while e do s_1 instruction by unrolling it into a conditional instruction.

2.5 Program execution

A program is an activity $\langle P, s \rangle$, that is, a statement s that is intended to start execution at place P. Program execution depends on an input strategy ω and a refiner R. The initial configuration of a program is $(H_{init}; \omega; R; \epsilon; \langle P, s \rangle)$, where H_{init} is a distinguished heap and ϵ is the empty trace.

For program $\langle P, s \rangle$, input strategy ω, and refiner R, we write

$$(\langle P, s \rangle, \omega, R) \text{ emits } t$$

to indicate that program execution can produce trace t. That is, there is some heap H', refiner R' and tree T' such that

$$(H_{init}; \omega; R; \epsilon; \langle P, s \rangle) \rightarrow^* (H'; \omega; R'; t; T')$$

where $\rightarrow^*$ is the reflexive transitive closure of the small-step relation $\rightarrow$.

3. Security

We are interested in enforcing strong information security in concurrent programs. Towards that end, in this section, we define a noninterference-based [12] definition of security for FSX10, and present a type system that enforces security while allowing many useful and highly concurrent programs.

3.1 Defining security

Intuitively, we want to ensure that a consumer of low-security information from a FSX10 program does not learn anything about high-security information. In our setting, consumers of low-security information are entities that can observe the communication channel of low-security places, and the high-security information that needs to be protected are the values input at high-security places.

We assume that there is a set of security levels $\mathcal{L}$ with a partial order $\sqsubseteq$ that describes relative restrictiveness of the security levels. We further assume that every place P is associated with a security level, denoted $\mathcal{L}(P)$. Intuitively, place P will be allowed to store and handle only data of security level $\mathcal{L}(P)$ and lower, and to send values to, and invoke code on, only places P' such that $\mathcal{L}(P) \sqsubseteq \mathcal{L}(P')$.

For a given security level $\ell \in \mathcal{L}$, a *low-security place* is any place P such that the level of the place is less than or equal to ℓ, that is, $\mathcal{L}(P) \sqsubseteq \ell$. All other places are *high-security places*, i.e., P is a high-security place if $\mathcal{L}(P) \not\sqsubseteq \ell$.

We define a semantic security condition based on *attacker knowledge* [2]. An attacker observes the communication channels of low-security places. The knowledge of an attacker is the set of input strategies that are consistent with the attacker's observations: the smaller the set, the more accurate the attacker's knowledge. The semantic security condition will require that at all times, the attacker's knowledge includes all possible input strategies for high-security places.

That is, all possible input strategies for high-security places are consistent with the attacker's observations. For ease of presentation, we will use a slightly weaker semantic security condition, a *progress-insensitive* condition [3] that also allows the attacker to learn not only the input strategies for low security places, but also whether low-security output is produced.

Trace equivalence Given an attacker with security level $\ell \in \mathcal{L}$ (i.e., who can observe communication channels of places P such that $\mathcal{L}(P) \sqsubseteq \ell$), two executions of a program look the same to the attacker if the trace of inputs and outputs at low-security places are the same in both executions. We define this formally via ℓ-equivalence of traces.

Definition 1 (ℓ-equivalence of traces)**.** Let $\ell \in \mathcal{L}$. Traces t_0 and t_1 are ℓ-*equivalent*, written $t_0 \sim_\ell t_1$, if $t_0 \restriction_\ell = t_1 \restriction_\ell$, where

$$\epsilon \restriction_\ell = \epsilon$$
$$(t \cdot \alpha) \restriction_\ell = \begin{cases} (t \restriction_\ell) \cdot \alpha & \text{if } \mathcal{L}(Place(\alpha)) \sqsubseteq \ell \\ t \restriction_\ell & \text{otherwise} \end{cases}$$

Attacker knowledge For a given execution of a program, starting from program configuration $(H_{init}; \omega; R; \epsilon; \langle P, s \rangle)$, that produces trace t, the knowledge of an attacker with security level ℓ, written $k(\langle P, s \rangle, R, t, \ell)$, is the set of input strategies that could have produced a trace that is equivalent to what the attacker observed.

Definition 2 (Attacker knowledge)**.** For any $\ell \in \mathcal{L}$, program $\langle P, s \rangle$, trace t, and refiner R, the attacker's knowledge is:

$$k(\langle P, s \rangle, R, t, \ell) =$$
$$\{\omega \mid \exists t'. (\langle P, s \rangle, \omega, R) \text{ emits } t' \wedge t \sim_\ell t'\}$$

We define what information an attacker with security level ℓ is permitted to learn about input strategies by defining ℓ-equivalence of input strategies. Intuitively, if two strategies are ℓ-equivalent, then they provide the exact same inputs for all low-security places, and an attacker with security level ℓ should not be able to distinguish them.

Definition 3 (ℓ-equivalence of input strategies)**.** Let $\ell \in \mathcal{L}$. Input strategies ω_0 and ω_1 are ℓ-*equivalent*, written $\omega_0 \sim_\ell \omega_1$, if for all places P such that $\mathcal{L}(P) \sqsubseteq \ell$, and for all traces t, we have $\omega_0(P, t \restriction_P) = \omega_1(P, t \restriction_P)$.

Relation $\sim_\ell$ is an equivalence relation, and we write $[\omega]_\ell$ for the equivalence class of ω under the relation $\sim_\ell$.

Given a program configuration $(H_{init}; \omega; R; \epsilon; \langle P, s \rangle)$ that produces trace t, *progress knowledge* [3] is the set of input strategies that could have produced a trace that is ℓ-equivalent to t, *and* could produce at least one more observable event. We will use progress knowledge as a lower

bound on the allowed knowledge of an attacker. That is, we will explicitly allow the attacker to learn whether a program will produce another observable event. This means that the attacker may be permitted to learn the termination behavior of statements that depend on high-security information.

Definition 4 (Progress knowledge)**.** For any $\ell \in \mathcal{L}$, program $\langle P, s \rangle$, trace t, and refiner R, progress knowledge is:

$$k^+(\langle P, s \rangle, R, t, \ell) =$$
$$\{\omega \mid \exists t', \alpha. (\langle P, s \rangle, \omega, R) \text{ emits } (t' \cdot \alpha)$$
$$\wedge t \sim_\ell t' \wedge \mathcal{L}(Place(\alpha)) \sqsubseteq \ell\}$$

Our security condition requires that, for all attackers, and all executions, for each event the attacker can observe, the attacker learns no more than the input strategy for low-security places, and the fact that another event was produced.

Definition 5 (Security)**.** Program $\langle P, s \rangle$ is secure if for all $\ell \in \mathcal{L}$, traces $t \cdot \alpha$, refiners R, and input strategies ω such that

$$(\langle P, s \rangle, \omega, R) \text{ emits } (t \cdot \alpha)$$

we have

$$k(\langle P, s \rangle, R, t \cdot \alpha, \ell) \quad \supseteq \quad [\omega]_\ell \cap k^+(\langle P, s \rangle, R, t, \ell).$$

Recall that the attacker's knowledge is set of input strategies that are consistent with the attacker's observations: a smaller set means more precise knowledge. Security requires that there are lower bounds to the precision of the attacker's knowledge. That is, there is information that the attacker is not permitted to learn. Thus, security requires that attacker's knowledge is a superset of the knowledge it is permitted to learn.

According to this definition of security, Program 2 from the Introduction is insecure (assuming that memory location hi is initialized from an input from place High), since there exists a refiner and a strategy that will produce a trace that allows an observer of low-security outputs to learn something about the high-security input strategy. Indeed, our definition of security rules out internal timing channels [48], in which the order of low-security events (here, input, output, and accesses to memory locations) depends upon high-security information. Program 3 does not exhibit an internal timing channel, and is secure.

This definition of security is *progress insensitive* [3], as it permits the attacker to learn that program execution makes progress, and produces another observable output. This definition can be strengthened in a straightforward way to a *progress sensitive* security condition. While the type system of Section 3.2 enforces progress insensitive security, it can be modified using standard techniques (to conservatively reason about termination of loops) to enforce progress sensitive security [28]. We refrain from doing so to simplify the presentation of the type system.

3.2 Enforcing security

We enforce security using a security type system. The type system ensures that each place P stores and handles only data input from places P' such that $\mathcal{L}(P') \sqsubseteq \mathcal{L}(P)$. However, as noted in the Introduction, it is possible for the scheduling of activities at place P to be influenced by input from a place P' such that $\mathcal{L}(P') \not\sqsubseteq \mathcal{L}(P)$. Our type system tracks and controls information flow through this covert channel through program point contexts.

A *program point context* Δ is a function from program points to security levels such that $\Delta(p)$ is an upper bound on the level of information that may influence the scheduling of program point p. More precisely, it is an upper bound on information that may affect the presence or absence of activities that may run concurrently with p at the same place.

Each program point is statically associated with a place, and we write $Place(p)$ for the place at which program point p will execute. Intuitively, since program point p is executed at place $Place(p)$, and $Place(p)$ handles data at security level $\mathcal{L}(Place(p))$, we would expect that $\Delta(p)$ is at least as restrictive as $\mathcal{L}(Place(p))$. Indeed, the type system ensures for all p that $\mathcal{L}(Place(p)) \sqsubseteq \Delta(p)$.

It is often the case that $\mathcal{L}(Place(p))$ is also an upper bound of $\Delta(p)$. That is, the scheduling of p does not depend on any high-security information. However, if p may happen immediately after a computation at a high-security place finishes (as with the output "pos" instruction in Program 2), or in parallel with another program point at the same place whose scheduling depends on high-security information, then it is possible that $\Delta(p) \not\sqsubseteq \mathcal{L}(Place(p))$. In that case, in order to ensure that the scheduling decision at place $Place(p)$ does not leak high-security information, we require *observational determinism* [49] at $Place(p)$ during the scheduling of p. That is, for each memory location stored at $Place(p)$, there are no data races on that location, and the order of input and output at $Place(p)$ is determined.

Finally, *variable context* Γ maps program variables to the place at which the variable was created. The type system uses the variable context to ensure that if variable x was declared at place P, then x is used only at places P' such that $\mathcal{L}(P) \sqsubseteq \mathcal{L}(P')$.

May-happen-in-parallel analysis The type system relies on the results of a may-happen-in-parallel analysis, such as the one presented by Lee and Palsberg for Featherweight X10 [22]. The async-finish parallelism of X10 is amenable to a precise may-happen-in-parallel analysis. We write $\mathsf{MHPP}(p)$ for the set of program points that may happen in parallel with program point p at the same place (i.e., at $Place(p)$).

Typing expressions Judgment $p; \Gamma; \Delta \vdash e$ indicates that expression e, occurring at program point p is well typed under variable context Γ and program point context Δ. Inference rules for this judgment are given in Figure 5. Constants

$$noWrite(r, p) = \forall p' \in \mathsf{MHPP}(p).\text{ instruction at } p' \text{ does not write to } r$$

$$noReadWrite(r, p) = noWrite(r, p)\ \wedge$$
$$\forall p' \in \mathsf{MHPP}(p).\text{ instruction at } p' \text{ does not read } r.$$

$$noIO(p) = \forall p' \in \mathsf{MHPP}(p).\text{ instruction at } p' \text{ does not perform input or output.}$$

Figure 6. Predicate definitions

v are always well typed, and the use of variable x is well typed if the level of the place at which x is defined ($\mathcal{L}(\Gamma(x))$) is less than or equal to the level of the place at which x is used ($\mathcal{L}(Place(p))$). Expression $e_1 \oplus e_2$ is well typed if both e_1 and e_2 are well typed.

There are two different rules for reading memory location r. The first rule, rule T-READNONDET, handles the case where the scheduling of the expression's execution at place $Place(p)$ is influenced by information at most at level $\mathcal{L}(Place(p))$. In that case, there are no restrictions on when the read may occur: it may occur concurrently with activities at the same place that write to the location since the resolution of the data race will not be a covert information channel. (The existence of a data race may, however, be undesirable in terms of program functionality.)

The second rule, rule T-READDET, applies when the scheduling of the expression may be influenced by information that is not allowed to flow to level $\mathcal{L}(Place(p))$. In that case, the read is required to be observationally deterministic: predicate $noWrite(r, p)$ must hold, implying that the read of memory location r at program point p must not execute concurrently with any statement that may write to r. Predicate $noWrite(r, p)$ is defined in Figure 6.

Typing statements Judgment $\Gamma; \Delta \vdash s : \ell$ indicates that statement s is well typed in variable context Γ and program point context Δ, and that security level ℓ is an upper bound on the security level of information that may influence the scheduling of the last program point of s. Inference rules for the judgment are given in Figure 7.

Every inference rule for a statement s^p includes the premise $\forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)$. Intuitively, the set $\mathsf{MHPA}(p)$ is the set of program points that may influence the presence or absence of activities running in parallel with p at the same place. Assuming that $Place(p) = P$, $\mathsf{MHPA}(p)$ contains the program points of backat P instructions that may happen in parallel with p, and the set of program points immediately following an $\mathsf{at}^{p'}\ P'\ s'$ instruction, where $Place(p') = P$ and p' may happen in parallel with p. The set $\mathsf{MHPA}(p)$ is a subset of the program points that may happen in parallel with p, and can easily be computed from the results of a may-happen-in-parallel analysis. Given this definition, the premise above requires that $\Delta(p)$, the up-

T-CONST
$$\frac{}{p;\Gamma;\Delta \vdash v}$$

T-VAR
$$\frac{\mathcal{L}(\Gamma(x)) \sqsubseteq \mathcal{L}(Place(p))}{p;\Gamma;\Delta \vdash x}$$

T-OP
$$\frac{p;\Gamma;\Delta \vdash e_1 \qquad p;\Gamma;\Delta \vdash e_2}{p;\Gamma;\Delta \vdash e_1 \oplus e_2}$$

T-READNONDET
$$\frac{\Delta(p) \sqsubseteq \mathcal{L}(Place(p))}{p;\Gamma;\Delta \vdash\, !r}$$

T-READDET
$$\frac{\Delta(p) \not\sqsubseteq \mathcal{L}(Place(p)) \qquad noWrite(r,p)}{p;\Gamma;\Delta \vdash\, !r}$$

Figure 5. Expression typing judgment $p;\Gamma;\Delta \vdash e$

T-SKIP1
$$\frac{\forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)}{\Gamma;\Delta \vdash \mathsf{skip}^p : \Delta(p)}$$

T-SKIP2
$$\frac{\Gamma;\Delta \vdash s : \ell \qquad \Delta(p) \sqsubseteq \Delta(p_1) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)}{\Gamma;\Delta \vdash \mathsf{skip}^p; s^{p_1} : \ell}$$

T-LET
$$\frac{\begin{array}{c}p;\Gamma;\Delta \vdash e \qquad \Gamma[x \mapsto Place(p)];\Delta \vdash s_1 : \ell_1 \\ \Gamma;\Delta \vdash s_2 : \ell_2 \qquad \Delta(p) \sqsubseteq \Delta(p_1) \qquad \ell_1 \sqsubseteq \Delta(p_2) \\ \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{let}^p\ x = e\ \mathsf{in}\ s_1^{p_1}; s_2^{p_2} : \ell_2}$$

T-ASYNC
$$\frac{\begin{array}{c}\Gamma;\Delta \vdash s_1 : \ell_1 \qquad \Gamma;\Delta \vdash s_2 : \ell_2 \qquad \Delta(p) \sqsubseteq \Delta(p_1) \\ \Delta(p) \sqsubseteq \Delta(p_2) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{async}^p\ s_1^{p_1}; s_2^{p_2} : \ell_2}$$

T-FINISH
$$\frac{\begin{array}{c}\Gamma;\Delta \vdash s_1 : \ell_1 \quad \Gamma;\Delta \vdash s_2 : \ell_2 \quad \Delta(p) \sqsubseteq \Delta(p_1) \quad \mathcal{L}(Place(p_2)) \sqsubseteq \Delta(p_2) \\ \forall p' \in \mathsf{MHPP}(p).\ \Delta(p') \sqsubseteq \Delta(p_2) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{finish}^p\ s_1^{p_1}; s_2^{p_2} : \ell_2}$$

T-AT
$$\frac{\begin{array}{c}\Gamma;\Delta \vdash s_1 : \ell_1 \qquad \Gamma;\Delta \vdash s_2 : \ell_2 \qquad \Delta(p) \sqsubseteq \mathcal{L}(P) \\ \mathcal{L}(P) \sqsubseteq \Delta(p_1) \qquad \ell_1 \sqsubseteq \Delta(p_2) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{at}^p\ P\ s_1^{p_1}; s_2^{p_2} : \ell_2}$$

T-BACKAT
$$\frac{\Gamma;\Delta \vdash s : \ell \qquad \Delta(p) \sqsubseteq \Delta(p_1) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)}{\Gamma;\Delta \vdash \mathsf{backat}^p\ P; s^{p_1} : \ell}$$

T-WRITENONDET
$$\frac{\begin{array}{c}\Delta(p) \sqsubseteq \mathcal{L}(Place(p)) \\ \Gamma;\Delta \vdash s : \ell \qquad p;\Gamma;\Delta \vdash e \qquad \Delta(p) \sqsubseteq \Delta(p_1) \\ \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash r :=^p e; s^{p_1} : \ell}$$

T-WRITEDET
$$\frac{\begin{array}{c}\Delta(p) \not\sqsubseteq \mathcal{L}(Place(p)) \qquad noReadWrite(r,p) \\ \Gamma;\Delta \vdash s : \ell \qquad p;\Gamma;\Delta \vdash e \qquad \Delta(p) \sqsubseteq \Delta(p_1) \\ \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash r :=^p e; s^{p_1} : \ell}$$

T-OUTPUTNONDET
$$\frac{\begin{array}{c}\Delta(p) \sqsubseteq \mathcal{L}(Place(p)) \\ \Gamma;\Delta \vdash s : \ell \qquad p;\Gamma;\Delta \vdash e \qquad \Delta(p) \sqsubseteq \Delta(p_1) \\ \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{output}^p\ e; s^{p_1} : \ell}$$

T-OUTPUTDET
$$\frac{\begin{array}{c}\Delta(p) \not\sqsubseteq \mathcal{L}(Place(p)) \qquad noIO(p) \\ \Gamma;\Delta \vdash s : \ell \qquad p;\Gamma;\Delta \vdash e \qquad \Delta(p) \sqsubseteq \Delta(p_1) \\ \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{output}^p\ e; s^{p_1} : \ell}$$

T-INPUTNONDET
$$\frac{\begin{array}{c}\Delta(p) \sqsubseteq \mathcal{L}(Place(p)) \qquad \Gamma;\Delta \vdash s : \ell \\ \Delta(p) \sqsubseteq \Delta(p_1) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{input}^p\ r; s^{p_1} : \ell}$$

T-INPUTDET
$$\frac{\begin{array}{c}\Delta(p) \not\sqsubseteq \mathcal{L}(Place(p)) \qquad noReadWrite(r,p) \qquad noIO(p) \\ \Gamma;\Delta \vdash s : \ell \qquad \Delta(p) \sqsubseteq \Delta(p_1) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{input}^p\ r; s^{p_1} : \ell}$$

T-IF
$$\frac{\begin{array}{c}\Gamma;\Delta \vdash s_1 : \ell_1 \qquad \Gamma;\Delta \vdash s_2 : \ell_2 \qquad \Gamma;\Delta \vdash s_3 : \ell_3 \\ p;\Gamma;\Delta \vdash e \qquad \Delta(p) \sqsubseteq \Delta(p_1) \qquad \Delta(p) \sqsubseteq \Delta(p_2) \\ \ell_1 \sqsubseteq \Delta(p_3) \qquad \ell_2 \sqsubseteq \Delta(p_3) \qquad \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{if}^p\ e\ \mathsf{then}\ s_1^{p_1}\ \mathsf{else}\ s_2^{p_2}; s_3^{p_3} : \ell_3}$$

T-WHILE
$$\frac{\begin{array}{c}\Gamma;\Delta \vdash s_1 : \ell_1 \qquad \Gamma;\Delta \vdash s_2 : \ell_2 \qquad p_1;\Gamma;\Delta \vdash e \\ \Delta(p) \sqsubseteq \Delta(p_1) \qquad \ell_1 \sqsubseteq \Delta(p) \qquad \ell_1 \sqsubseteq \Delta(p_2) \\ \forall p' \in \mathsf{MHPA}(p).\ \Delta(p') \sqsubseteq \Delta(p)\end{array}}{\Gamma;\Delta \vdash \mathsf{while}^p\ e\ \mathsf{do}\ s_1^{p_1}; s_2^{p_2} : \ell_2}$$

Figure 7. Statement typing judgment $\Gamma;\Delta \vdash s : \ell$

per bound on the scheduling of s, is at least as restrictive as the scheduling of any program point that may influence the presence or absence of activities running in parallel with p at the same place.

Also, almost all inference rules for statements ensure that if program point p executes after p' (for example, because they are in sequence), then $\Delta(p') \sqsubseteq \Delta(p)$. The intuition here is that if information at level $\Delta(p')$ may influence the scheduling of p', and p follows in sequence after p', then information at level $\Delta(p')$ may influence the scheduling of p. For example, the typing rule for $\mathrm{if}^p\ e$ then $s_1^{p_1}$ else $s_2^{p_2}$; $s_3^{p_3}$ requires that $\Delta(p) \sqsubseteq \Delta(p_1)$ and $\Delta(p) \sqsubseteq \Delta(p_2)$, since the execution of s_1 and s_2 will occur only after the evaluation of the conditional guard. Similarly, since the execution of s_3 will follow the execution of either s_1 or s_2, the rule requires that $\ell_1 \sqsubseteq \Delta(p_3)$ and $\ell_2 \sqsubseteq \Delta(p_3)$, where ℓ_1 and ℓ_2 are upper bounds of the scheduling of the last program points of s_1 and s_2 respectively.

We discuss only the inference rules that have premises in addition to those common to all rules.

Statement $\mathrm{let}^p\ x = e$ in $s_1^{p_1}$; $s_2^{p_2}$ declares a variable x, and allows x to be used in the scope of statement s_1. Rule T-LET thus allows s_1 to be typed with a variable context that maps variable x to the place at which it was defined: $Place(p)$.

Statement $\mathrm{finish}^p\ s_1^{p_1}$; $s_2^{p_2}$ executes statement s_1, and waits until all activities spawned by s_1 have finished before executing s_2. Rule T-FINISH requires that $\Delta(p) \sqsubseteq \Delta(p_1)$ (since p_1 is executed after p) but notably does not require either $\Delta(p) \sqsubseteq \Delta(p_2)$ or $\ell_1 \sqsubseteq \Delta(p_2)$, despite the fact that p_2 is executed after p and p_1. The intuition is that because the scheduling behavior at place $P = Place(p)$ depends only on the current activities at P, by the time that p_2 is scheduled, program points p and p_1 (and all activities spawned by s_1) have finished execution, and do not influence the scheduling of p_2. In Program 3 in the Introduction, this reasoning is what permits us to conclude that the scheduling of output "B" and output "C" do not depend on high-security computation.

```
1   at Low {
2       async {
3           // Activity 1
4           mediumComputation(); output "nonpos";
5       }
6       finish {
7           // Activity 2
8           at High {
9               if (hi > 0) longComputation();
10          }
11      }
12      // Activity 3
13      output "pos";
14  }
```

Program 4.

However, it may be possible that scheduling of activities spawned by s_1 indirectly influences the scheduling of p_2. Consider Program 4, which contains a finish s_1; s_2 statement where $s_2 =$ output "pos", and s_1 invokes computation at high-security place High. There is an additional activity that executes concurrently with the finish statement: mediumComputation(); output "nonpos". The scheduling of this activity relative to s_2 will depend on the high-security computation. Indeed, this program is equivalent to Program 2, and both are insecure. Thus, typing rule T-FINISH requires that $\Delta(p_2)$ is at least as restrictive as $\Delta(p')$ for any program point p' that may execute in parallel with p at the same place. This ensures that insecure Program 4, and others like it, are rejected by the type system.

Statement $\mathrm{at}^p\ P\ s_1^{p_1}$; $s_2^{p_2}$ executes s_1 at place P, and then executes s_2 back at place $Place(p)$. Rule T-AT requires that the upper bound on the scheduling of the at instruction is permitted to flow to the level of place P ($\Delta(p) \sqsubseteq \mathcal{L}(P)$). Thus the type system restricts the creation of an activity at place P to reveal only information that is allowed to flow to level $\mathcal{L}(P)$. Also, because statement $s_1^{p_1}$ is executing at place P, information at level $\mathcal{L}(P)$ will influence the scheduling of p_1: $\mathcal{L}(P) \sqsubseteq \Delta(p_1)$. Finally, because statement $s_2^{p_2}$ is executed only after s_1, the scheduling of p_2 depends on when the last statement of s_1 is scheduled: $\ell_1 \sqsubseteq \Delta(p_2)$ where ℓ_1 is an upper bound on the scheduling of the last program point of s_1.

Similar to the typing rules for reading memory locations, there are two rules for writing memory locations: T-WRITENONDET and T-WRITEDET. As with the rules for reading memory, the first is for the case where the scheduling of the write is not influenced by high-security information, and there are thus no restrictions on when the write may occur. Rule T-WRITEDET applies when the scheduling of the write may be influenced by high-security information, and requires observational determinism via the predicate $noReadWrite(r, p)$, defined in Figure 6, which ensures that no reads or writes to the same memory location may happen in parallel.

The rules for input and output are similar to the rules for reading and writing memory locations: if the scheduling of input or output may depend on high-security information, the input or output must be observationally deterministic, which is achieved for output by requiring that there is no other input or output at that place that may happen in parallel (see predicate $noIO(p)$, defined in Figure 6). Since an input instruction writes to a memory location r, rule T-INPUTDET requires both that no input or output may happen at the place in parallel, and that no reads or writes to r may happen in parallel.

Typing trees Judgment $\Gamma; \Delta \vdash T$ means that tree T is well typed in variable context Γ and program point context Δ. Inference rules for the judgment are given in Figure 8. The rules require that all activities in the tree are well typed.

T-ACTIVITY
$$\frac{\Gamma; \Delta \vdash s : \ell \quad \mathcal{L}(P) \sqsubseteq \Delta(p)}{\Gamma; \Delta \vdash \langle P, s^p \rangle}$$

T-PARA
$$\frac{\Gamma; \Delta \vdash T_1 \quad \Gamma; \Delta \vdash T_2}{\Gamma; \Delta \vdash T_1 \| T_2}$$

T-JOIN
$$\frac{\Gamma; \Delta \vdash T \quad \Gamma; \Delta \vdash \langle P, s^p \rangle \quad \forall p' \in \mathsf{MHPP}(p).\ \Delta(p') \sqsubseteq \Delta(p)}{\Gamma; \Delta \vdash T \triangleright \langle P, s^p \rangle}$$

T-DONE
$$\frac{}{\Gamma; \Delta \vdash \checkmark}$$

Figure 8. Tree typing judgment $\Gamma; \Delta \vdash T$

Also, the rule for tree $T \triangleright \langle P, s^p \rangle$, T-JOIN, requires that $\Delta(p)$ is at least as restrictive as $\Delta(p')$ for any program point p' that may execute in parallel with p at place P, for similar reasons to the typing rule for finish statements, T-FINISH.

Soundness of type system The type system enforces security. That is, if a program is well typed, then it is secure.

Theorem 1. If $\langle P, s \rangle$ is a program such that $\Gamma; \Delta \vdash \langle P, s \rangle$ for some variable context Γ and program point context Δ, then $\langle P, s \rangle$ is secure according to Definition 5.

We present a brief sketch of the proof here. A more detailed proof appears in the companion technical report [29].

Outline of Proof. The proof uses a technique similar to that of Terauchi [45]. We first introduce the concept of an *erased configuration*. A configuration m erases to a configuration m' at security level ℓ if m', when executed, performs no computation at places with security level higher than ℓ but m and m' otherwise agree. Erased programs are defined similarly, with erased configurations containing erased programs.

Suppose we have a well-typed program $\langle P, s \rangle$, some security level ℓ, and two ℓ-equivalent input strategies ω_1 and ω_2. First, we erase the program $\langle P, s \rangle$ to program $\langle P, s' \rangle$ at level ℓ and consider side-by-side executions of these two programs with the same input strategy. Suppose the original program with input strategy ω_1 produces trace t_1. Then the erased program with input strategy ω_1 can produce a trace t_1' that is ℓ-equivalent. Similarly, if the original program with input strategy ω_2 produces trace t_2, then the erased program with input strategy ω_2 can produce a trace t_2' that is ℓ-equivalent.

Second, we consider the executions of the erased program with strategy ω_1 and strategy ω_2 that produced traces t_1' and t_2' respectively. Since the erased program performs no computation at high-security places, either t_1' is a prefix of t_2', or vice versa. Combining this with the previous result, if $(\langle P, s \rangle, \omega_1, R)$ emits t_1 and $(\langle P, s \rangle, \omega_2, R)$ emits t_2, then either the low-security events of t_1 are a prefix of the low-security events of t_2, or vice versa.

Knowledge-based security can then be shown as follows. Let ω be an input strategy, R a refiner, and ℓ a security level. Suppose that $(\langle P, s \rangle, \omega, R)$ emits $t \cdot \alpha$. Let ω' be another input strategy such that $\omega \sim_\ell \omega'$ and $(\langle P, s \rangle, \omega', R)$ emits $t' \cdot \alpha'$ such that $t \sim_\ell t'$ and $\mathcal{L}(Place(\alpha')) \sqsubseteq \ell$. The above result implies that either $\alpha = \alpha'$ or $\mathcal{L}(Place(\alpha)) \not\sqsubseteq \ell$. In either case, $t \cdot \alpha \sim_\ell t' \cdot \alpha'$, and so the inclusion required by Definition 5 is proven. □

4. SX10 prototype implementation

We have extended the principles of the security analysis of Section 3 to handle many of the language features of X10. The resulting language, called SX10, is a subset of X10. We have implemented a prototype compiler for SX10 by extending the open-source X10 compiler (version 2.1.2), which is implemented using the Polyglot extensible compiler framework [30], and is included in the X10 distribution. Our extension comprises approximately 8,500 lines of non-comment non-blank lines of Java code.

We do not modify the X10 run-time system: SX10 programs run using the standard X10 run-time system. We thus do not provide a performance comparison of SX10 with X10 or with other secure concurrent systems. Such a performance comparison is not directly useful, as it would evaluate the efficiency of the X10 runtime, not our enforcement technique, which is entirely static.

In this section, we describe how we extend the analysis to handle additional language features of X10 and present some example SX10 programs.

May-happen-in-parallel analysis We have implemented the may-happen-in-parallel (MHP) analysis of Lee and Palsberg [22] for SX10, which is a straightforward exercise. However, for additional precision in our security analysis, we implemented a *place-sensitive* MHP analysis. In our calculus FSX10, for every program point it is possible to statically determine which place the program point would execute on. In SX10, however, code for a given class may be executed at more than one place, since objects of the same class may reside at different places. Thus, if an activity at place P is executing code from program point p, our place-sensitive MHP analysis conservatively approximates the set $\mathsf{MHP}(p, P)$ such that if $(p', P') \in \mathsf{MHP}(p, P)$ then an activity at place P' may concurrently be executing code from program point p'.

Places We assume that all places are statically known, and that a security level is associated with each place. A configuration file specifies the set of security levels $\mathcal{L}$, the ordering $\sqsubseteq$ over the levels, and maps places to levels. Our prototype implementation does not currently support first-class places. If places are computed dynamically, then the choice of the place at which to execute a computation could be a covert channel, and would thus require the security analysis to track and control information flow through this channel. In this respect, first-class places are similar to first-class security levels (e.g., [13, 50]), and the security analysis could be extended to handle first-class places using similar techniques, such as dependent type systems.

As in FSX10, we restrict at statements to allow place P to invoke code on place P' only if $\mathcal{L}(P) \sqsubseteq \mathcal{L}(P')$.

In addition to at statements, X10 has at expressions: at P e evaluates expression e at place P. We allow at expressions, but only from place P to place P' where $\mathcal{L}(P) = \mathcal{L}(P')$. If the security level of the places differed, then either data would be sent from a high-security place to a low-security place, or a high-security place would invoke code on a low-security place. Either way, a potentially dangerous information flow occurs, and must be ruled out.

Concurrency mechanisms The X10 async and finish statements are restricted similarly to their counterparts in FSX10. X10 provides additional synchronization mechanisms, including clocks (a form of synchronization barrier), futures, and atomic blocks. Our prototype implementation does not currently support these additional mechanisms. However, they can be incorporated in a straightforward manner by extending the MHP analysis to reason about them. Once the MHP analysis supports these constructs, our security analysis can be extended to add constraints similar to those for async and finish statements.

Objects Fields of objects can be mutable locations, and we enforce restrictions similar to those of other memory locations: we require determinism on accesses when scheduling may be influenced by high-security information. If an object is sent in a message from one place to another, the X10 runtime will create a copy of the object, thus ensuring that if an activity at a place attempts to update a field of an object, the memory location is local to the place. When objects are copied to send to another place, we impose restrictions similar to the use of variables: a copy of an object created at place P may be sent to place P' only if $\mathcal{L}(P) \sqsubseteq \mathcal{L}(P')$.

Control-flow constructs X10 has much richer control-flow constructs than the calculus FSX10. We support local-control-flow constructs, such as for loops and switch statements. We support dynamic dispatch of methods, using class information to conservatively over-approximate the set of possible callees at a method call site. We do not currently support exceptions, although they can be incorporated by extending the MHP analysis. Note that exceptions interact in an interesting way with the concurrency mechanisms, due to X10's *rooted exception model* [41].

Input and output The security analysis for SX10 restricts input and output from the system to enforce strong information security guarantees. We currently require methods that perform communication with the external environment to be explicitly annotated as such, but it is straightforward to infer where such methods are used, for example, detecting method calls to objects of classes x10.io.Printer, x10.io.Reader, etc. (Fields x10.io.Console.OUT and x10.io.Console.IN are instances of Printer and Reader, respectively.)

Arrays Our implementation supports local arrays, since these are simply objects of the class x10.array.Array[T], and elements of the array are stored at a single place. We do not currently support distributed arrays, which store elements over multiple places. Adding support for distributed arrays would require support for first-class places.

X10 runtime scheduler The X10 runtime scheduler uses a work-stealing algorithm to schedule activities within a place. This requires that threads maintain a double-ended queue of pending activities, and idle threads may steal activities from busy threads. Because the state of the queues may be influenced by which activities were or were not running at the place in the past, such work-stealing algorithms cannot be represented in our scheduling model in FSX10, which requires that scheduling functions do not depend on the history of computation at a place.

The type system relies on this requirement only in the rule for a finish $s_1 ; s_2$ statement, which allows the program point context of the first program point of s_2 to be lower than program point context of the last program point of s_1 when there are no other activities running at the place. However, in that situation there are no other activities to schedule other than s_1 and activities spawned by s_1: activity s_2 will not start execution until it is the only activity at the place. In that case, the state of the work queues for threads will be independent of the history of the computation up to that point. Thus, we expect the security guarantee to transfer to the actual scheduler used by the X10 runtime.

It is future work to extend the model of schedulers in FSX10 to include such work-stealing schedulers.

Improvements to Analysis In implementing the SX10 compiler, we add an optimization that allows the type system to be more permissive without compromising security. If statement at P s is executed at place P' and no statement at P' is dependent on the termination of s, then the scheduling of activities at P' is independent of when s terminates, and we do not need to increase the program point context of statements that may happen in parallel. This corresponds to having a more precise definition of the set MHPA(p).
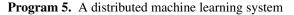
4.1 Example programs

Distributed Machine Learning Consider a music recommendation service, such as Pandora. Here, a large database of music information exists: the Music Genome Project. The service would like to process this data—perhaps running machine-learning algorithms on it—and then combine it with data from individual users to produce recommendations for users. We assume the database of music information is public, but the personal data from users is secret, and should not influence the results observed by other users.

We assume that the processing of the public data can be performed in parallel. We will process this data at a number of places pub0 through pubn, all with the same low security level L. Results from the processing will be sent back to the coordinator place, and collated into a value called results. We use the results of the processing of public data to compute

recommendations for each user. We assume that each user U_i has its own place Pi, with a unique security level H_i such that $L \sqsubseteq H_i$.

A sketch of the code for this system appears in Program 5. Public data is supplied in the array pubData and we assume that place P_i already holds the data for user U_i.

```
1   at coordinator {
2     finish {
3       async at pub0 {
4         val res = processPublicData(pubData(0));
5         at coordinator { addResult(res); }
6       }
7       ...
8       async at pubn {
9         val res = processPublicData(pubData(n));
10        at coordinator { addResult(res); }
11      }
12    }
13
14    async { at P0 { processPrivateData(results); } }
15    ...
16    async { at Pn { processPrivateData(results); } }
17  }
```

Program 5. A distributed machine learning system

Note that no additional synchronization is required to make this program secure, and the secure program is in fact allowed to be highly concurrent. The translation of this program from X10 to SX10 does, however, require significant code duplication due to the lack of support for distributed arrays and first-class places. Natural X10 programming style would use a loop over places to execute a block of code at each place, rather than duplicate the code as we do here.

Online Shopping Following Tsai et al. [47], we consider a server running a shopping website. We assume that the server is concerned with keeping credit card data secure. Our example program models a multithreaded server accepting input from two web forms. On the first form, a user enters the item number they wish to purchase. This form is submitted along with the user's unique customer ID, which persists through the user's session. When this form is processed, the user's order is both saved on the server and output to a log for inventory purposes. The user is then presented with the next form, which requests his or her credit card number. This form is submitted with the same customer ID, and the credit card number is sent to an external service for processing.

This example contains two security levels. The customer ID and order are considered low-security, and the log is considered low-security output. The customer's credit card number is high-security, and so the action of exporting it should occur at a high place. We would like to ensure that no data from either the second form or the credit card processing service can leak to the log. The code for handling one customer's purchase is shown in Program 6.

Simpson's Rule Our final example program demonstrates that in programs, or sections of programs, in which all data is at the same security level, our analysis requires very few changes to the code for compilation. Thus, when a program operates on homogenous data, our security analysis does not significantly impact usability. The code for this example was taken from the Simpson's Rule example available on the X10 website[1]. The original program consists of approximately 200 non-blank non-comment lines of X10 code. Converting this program to SX10 required modifying seven lines of code, most of them trivially, and adding ten. The changes were as follows:

- Five statements producing console output were annotated as required by SX10.

- The original program uses all available places. Since SX10 requires static places, Place.MAX_PLACES, which in X10 is set to the number of places, was replaced with a new (arbitrary) constant, set to four for the purposes of this example. Identifiers representing these four places were declared.

- The code to start computation at each place was duplicated, since SX10 does not support loops over places. This required five additional lines of code.

Note that neither the number of lines modified nor the number added necessarily scales with the size of the code. Most required modifications were to input or output statements, and the number of lines of code added was proportional to the number of places used, not the size of the program.

Discussion of Example Programs The example programs demonstrate that it is possible to write realistic, highly concurrent programs in SX10. Note that the first two examples contain a high degree of nondeterminism. The order in which blocks of data are processed in Program 5 and the order in which entries are written to the log and credit card service in Program 6 are nondeterministic. This is secure because the resolution of this nondeterminism can in no way reveal high-security information. As will be discussed in Section 5, some previous security-type systems for noninterference in concurrent programs would rule out this nondeterminism and require additional synchronization and overhead.

The third example demonstrates that our analysis does not significantly prohibit the compilation and execution of programs that operate on a single security level.

The biggest restriction in SX10 is the lack of first-class places. As we gain more experience writing SX10 programs, we will identify and address further challenges to practical and secure concurrent programming.

5. Related work

This work seeks to provide strong language-based information security guarantees for concurrent programs. We discuss

---

[1] http://x10-lang.org/

```
1   async {
2       at lowform {
3           val custID = Int.parse(Console.IN.readLine().trim());
4           // Input item number (low security)
5           items(custID) = Int.parse(Console.IN.readLine().trim());
6           at log Console.OUT.println(custID + "\t" + items(custID));
7           at highform {
8               // Input credit card number (high security)
9               val card = Int.parse(Console.IN.readLine().trim());
10              at cc Console.OUT.println(card + "\t" + costs(items(custID)));
11          }
12      }
13  }
```

Program 6. An online shopping cart

related work, focusing on recent work that controls information flow in concurrent settings.

Observational determinism Our security analysis ensures that if the scheduling of input, output, or memory accesses may leak sensitive information, then the order of such instructions must be deterministic. This approach is inspired by Zdancewic and Myers [49], who propose (following McLean [27] and Roscoe [33]) that there should be no nondeterminism (including thread scheduling) observable by a low-security observer. They present a semantic security condition that, for each observable memory location, requires determinism of the sequence of updates to that location. Huisman et al. [19] point out that this semantic security condition may reveal more information than intended, and propose that the sequence of updates to *all* observable memory locations should be deterministic. Terauchi [45] presents a type system that enforces such a semantic security condition in a shared-memory setting using fractional capabilities.

Mantel et al. [26] present semantic conditions that allow composition of concurrent programs. The semantic conditions use assume-guarantee reasoning to ensure that the composed program is free of data races, and thus is observationally deterministic.

Requiring observational determinism throughout a program is, however, overly restrictive. O'Neill et al. [31] note that low-observable nondeterminism is acceptable so long as its resolution depends only on low-observable information. We thus allow nondeterminism in the scheduling of activities at a place, provided that the resolution of the nondeterminism cannot leak sensitive information. Our model assumes that scheduling of activities at a place depends only on the activities at that place, and our security analysis exploits this assumption to allow non-determinism where possible.

Recent work on deterministic concurrency (e.g., [7, 46]) highlights functional benefits of determinism, and also allows some nondeterminism when it is safe to do so [8].

Scheduler independence Sabelfeld and Sands [38, 40] argue that the definition of security in multi-threaded programs should be *scheduler independent*, since the scheduler is typically outside of the language specification, and violations of scheduler assumptions may lead to vulnerabilities. By contrast, Boudol and Castellani [9] present a type system for schedulers and threads, and show that well-typed schedulers and threads satisfy a definition of security [4].

Barthe et al. [5] have developed a framework for security of multi-threaded programs that allows programs to be written without knowledge of the scheduler, i.e., in a scheduler-independent manner. Mechanisms to interact with the scheduler and secure timing channels are introduced during compilation, and enable a security-aware scheduler to enforce strong information security guarantees.

Mantel and Sabelfeld [24] show a scheduler-independent security property in a multi-threaded while language. Russo and Sabelfeld [35] suggest a model in which threads may increase and decrease their security levels and permitted scheduling decisions depend on the security levels of threads. Mantel and Sudbrock [25] prove a security property for programs consisting of threads with assigned security levels when these are run under any scheduler in a class of robust schedulers. Robust schedulers, such as round-robin schedulers, have the property that the probability that a particular low thread will be selected to run from among all low threads remains the same if high threads are removed. Our assumptions about scheduling in the X10 runtime imply that schedulers for places are robust, in that scheduling of activities at a place cannot depend upon the existence or non-existence of activities at higher-security places.

We do not provide scheduler independence. Our type-system and security proof assume that scheduling at a place depends only on the activities currently executing at a place. While this assumption enables greater concurrency while preserving security, it perhaps violates an abstraction boundary, as it makes assumptions about the behavior of the X10 runtime that are not necessarily intended as part of the runtime's specification.

Dynamic enforcement of concurrent information security
Tsai et al. [47] extend work of Li and Zdancewic [23] and
Russo and Sabelfeld [34] to encode information-flow con-
trol in Haskell with support for concurrency and side-effects.
However, their mechanism relies on co-operative (i.e., non-
preemptive) scheduling, which may not be suitable for mod-
ern operating systems.

Stefan et al. [44] present a dynamic information-flow con-
trol system that eliminates termination and internal timing
channels, and mitigates external timing channels, without re-
lying on co-operative scheduling. Implemented as a Haskell
library, their technique requires that the security level of a
thread A that waits on a forked thread B must be at least as
restrictive as the information that influences the control flow
of thread B. A similar restriction is true of our static mech-
anism: the security level of a program point p that occurs
after execution of program point p' is at least as restrictive
as information that influences the scheduling of p'. However,
our static analysis allows us to lower the security level of p
in a particular situation: for a finish $s_1 ; s_2$ statement, the pro-
gram point context of the first program point of s_2 can be
lower than program point context of the last program point
of s_1 if there are no other activities running at the place. The
dynamic nature of the system of Stefan et al. allow them to
be more precise than our static analysis in certain situations,
highlighting the incomparability of static and dynamic flow-
sensitive security [36].

Le Guernic [21] uses a hybrid execution monitor (which
combines static and dynamic analyses) to enforce a strong
security condition. The enforcement mechanism (similar to
the type system of Smith and Volpano [43]) is restrictive:
while loops cannot have high-security guards, and while
loops are not permitted in branches of if commands with
high-security guards. These restrictions are severe enough
to rule out many useful programs.

Process calculi   Focardi et al. [11] establish a link be-
tween language-based security for imperative programs, and
process-algebraic frameworks of security properties. How-
ever, they consider only sequential imperative programs,
and do not explore concurrency. Honda et al. [18] present
a security-type system for the π-calculus (further devel-
oped by Honda and Yoshida [16, 17]) to address internal
timing and progress channels. In their type system, chan-
nels are assigned security levels, and may be given linear
types. Linear channels must statically have a single send
and receive, which enables precise reasoning about syn-
chronization between processes. Non-linear channels may
have non-deterministic behavior, and processes cannot send
low-security outputs after receiving high-security input on a
non-linear channel, as the resolution of the non-determinism
may be a covert channel. Pottier [32] presents a simpler type
system (without linear channel types, and with a simpler
proof) that also prevents low-security outputs after receiving
high-security input.

Kobayashi [20] presents a type system for π-calculus that
allows low-security output after high-security synchroniza-
tion for a variety of synchronization mechanisms. It extends
the idea of linear channels by using types to describe the
channel usage, which permits precise reasoning about the
information flow resulting from synchronization.

Security-type systems   Other work concerned with in-
formation security in concurrent systems have also used
security-type systems to enforce strong semantic security
conditions (e.g., [5, 9, 37, 40, 42, 43, 45, 48, 49]). Some
of these previous security-type systems are overly restric-
tive on synchronization between threads, either disallowing
low-security output after synchronization with high-security
threads or activities, or disallowing nondeterminism even
when resolution of the nondeterminism is not influenced by
high-security information.

The key difference in this work is that we integrate
information-security guarantees with modern concurrency
abstractions (i.e., X10 places). In doing so, we reason about
information at coarser granularity than previous work, which
simplifies reasoning about information flow (and thus, we
believe, leads to increased practicality). Our type system
does not track information flow on a per-location basis, but
rather focuses on tracking how interaction with high-security
places affects the scheduling of program points at a place.

6.   Conclusion

We have extended the X10 concurrent programming lan-
guage with coarse-grained information-flow control. The re-
sulting language, SX10, provides information security for
concurrent programs. Each place is associated with a secu-
rity level, and may only handle data that is appropriate for
the security level. We believe this language provides a bet-
ter intuition for information flow than previous methods for
controlling information flow, and will allow programmers to
write secure programs more effectively.

The security analysis benefits from X10's abstractions
for concurrency: potentially dangerous information flows
correspond to interactions between places, which are rela-
tively easy to detect, since communication between places
is by message passing. Interaction between places may re-
sult in the scheduling of activities at a place being influ-
enced by high-security information. Through a may-happen-
in-parallel analysis for X10 [22], our security analysis will
determine when this situation may arise, and requires *obser-
vational determinism* [49] to hold, which prevents activity
scheduling from being a covert information channel. In the
absence of interaction between places with different secu-
rity levels, our security mechanism places no restrictions on
the concurrent program. While some restrictions on concur-
rency necessarily remain, this allows a large class of useful
programs to be written without burdensome synchronization
between threads for the purposes of security.

This work highlights the opportunity for synergy between mechanisms for concurrency and mechanisms for information security: both rely on reasoning about dependencies within a program. We believe it is a promising step towards languages and tools for building secure concurrent systems.

Acknowledgments

References

[1] M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *Conference Record of the Twenty-Sixth Annual ACM Symposium on Principles of Programming Languages*, pages 147–160, New York, NY, USA, 1999. ACM Press.

[2] A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 207–221. IEEE Computer Society, 2007.

[3] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-insensitive noninterference leaks more than just a bit. In *Proceedings of the 13th European Symposium on Research in Computer Security*, Oct. 2008.

[4] G. Barthe and L. P. Nieto. Formally verifying information flow type systems for concurrent and thread systems. In *Proceedings of the 2004 ACM workshop on Formal methods in security engineering*, pages 13–22, New York, NY, USA, 2004. ACM.

[5] G. Barthe, T. Rezk, A. Russo, and A. Sabelfeld. Security of multithreaded programs by compilation. *ACM Transactions on Information and System Security*, 13(3):21:1–21:32, July 2010.

[6] R. L. Bocchino, V. S. Adve, S. V. Adve, and M. Snir. Parallel programming must be deterministic by default. In *Proceedings of the First USENIX Workshop on Hot Topics in Parallelism*, 2009.

[7] R. L. Bocchino, Jr., V. S. Adve, D. Dig, S. V. Adve, S. Heumann, R. Komuravelli, J. Overbey, P. Simmons, H. Sung, and M. Vakilian. A type and effect system for deterministic parallel Java. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications*, pages 97–116, New York, NY, USA, 2009. ACM.

[8] R. L. Bocchino, Jr., S. Heumann, N. Honarmand, S. V. Adve, V. S. Adve, A. Welc, and T. Shpeisman. Safe nondeterminism in a deterministic-by-default parallel language. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 535–548, New York, NY, USA, 2011. ACM.

[9] G. Boudol and I. Castellani. Non-interference for concurrent programs and thread systems. *Theoretical Computer Science*, 281(1):109–130, June 2002.

[10] P. Charles, C. Grothoff, V. Saraswat, C. Donawa, A. Kielstra, K. Ebcioglu, C. von Praun, and V. Sarkar. X10: an object-oriented approach to non-uniform cluster computing. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications*, pages 519–538, New York, NY, USA, 2005. ACM.

[11] R. Focardi, S. Rossi, and A. Sabelfeld. Bridging language-based and process calculi security. In *Foundations of Software Science and Computation Structure*, volume 3441 of *Lecture Notes in Computer Science*, pages 299–315, Edinburgh, UK, Apr. 2005. Springer-Verlag.

[12] J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society, Apr. 1982.

[13] R. Grabowski and L. Beringer. Noninterference with dynamic security domains and policies. In *13th Asian Computing Science Conference, Focusing on Information Security and Privacy*, 2009.

[14] D. Grove, O. Tardieu, D. Cunningham, B. Herta, I. Peshansky, and V. Saraswat. A performance model for x10 applications. In *Proceedings of The First X10 Workshop*, 2011.

[15] Y. Guo, R. Barik, R. Raman, and V. Sarkar. Work-first and help-first scheduling policies for async-finish task parallelism. In *Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing*, pages 1–12, Washington, DC, USA, 2009. IEEE Computer Society.

[16] K. Honda and N. Yoshida. A uniform type structure for secure information flow. In *Conference Record of the Twenty-Ninth Annual ACM Symposium on Principles of Programming Languages*, pages 81–92, New York, NY, USA, Jan. 2002. ACM Press.

[17] K. Honda and N. Yoshida. Noninterference through flow analysis. *Journal of Functional Programming*, 15(2):293–349, Mar. 2005.

[18] K. Honda, V. Vasconcelos, and N. Yoshida. Secure information flow as typed process behaviour. In *Proceedings of the Ninth European Symposium on Programming*, volume 1782 of *Lecture Notes in Computer Science*, pages 180–199. Springer, 2000.

[19] M. Huisman, P. Worah, and K. Sunesen. A temporal logic characterisation of observational determinism. In *Proceedings of the 19th IEEE Workshop on Computer Security Foundations*, 2006.

[20] N. Kobayashi. Type-based information flow analysis for the pi-calculus. *Acta Informatica*, 42(4-5):291–347, 2005.

[21] G. Le Guernic. Automaton-based Confidentiality Monitoring of Concurrent Programs. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium*, pages 218–232, 2007.

[22] J. K. Lee and J. Palsberg. Featherweight X10: a core calculus for async-finish parallelism. In *Proceedings of the 15th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, pages 25–36, New York, NY, USA, January 2010. ACM.

[23] P. Li and S. Zdancewic. Encoding information flow in Haskell. In *Proceedings of the 19th IEEE Workshop on Computer Security Foundations*, pages 16–27, Washington, DC, USA, 2006. IEEE Computer Society.

[24] H. Mantel and A. Sabelfeld. A generic approach to the security of multi-threaded programs. In *Proceedings of the 14th IEEE Computer Security Foundations Workshop*, page 126, Washington, DC, USA, 2001. IEEE Computer Society.

[25] H. Mantel and H. Sudbrock. Flexible scheduler-independent security. In *Proceedings of the 15th European Conference on Research in Computer Security*, pages 116–133, Berlin, Heidelberg, 2010. Springer-Verlag.

[26] H. Mantel, D. Sands, and H. Sudbrock. Assumptions and guarantees for compositional noninterference. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium (CSF)*, pages 218–232. IEEE Computer Society, 2011.

[27] J. McLean. Proving noninterference and functional correctness using traces. *Journal of Computer Security*, 1(1):37–58, 1992.

[28] S. Moore, A. Askarov, and S. Chong. Precise enforcement of progress-sensitive security. In *Proceedings of the 19th ACM Conference on Computer and Communications Security*, New York, NY, USA, 2012. ACM Press.

[29] S. Muller and S. Chong. Towards a practical secure concurrent language. Technical Report TR-05-12, Harvard University, 2012.

[30] N. Nystrom, M. R. Clarkson, and A. C. Myers. Polyglot: An extensible compiler framework for java. In *In 12th International Conference on Compiler Construction*, pages 138–152. Springer-Verlag, 2003.

[31] K. R. O'Neill, M. R. Clarkson, and S. Chong. Information-flow security for interactive programs. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop*, pages 190–201. IEEE Computer Society, June 2006.

[32] F. Pottier. A simple view of type-secure information flow in the π-calculus. In *Proceedings of the 15th IEEE Computer Security Foundations Workshop*, pages 320–330, June 2002.

[33] A. W. Roscoe. CSP and determinism in security modelling. In *Proceedings of the 1995 IEEE Symposium on Security and Privacy*, pages 114–127, Washington, DC, USA, 1995. IEEE Computer Society.

[34] A. Russo and A. Sabelfeld. Security for multithreaded programs under cooperative scheduling. In *Proceedings of Andrei Ershov International Conference on Perspectives of System Informatics*, volume 4378 of *Lecture Notes in Computer Science*, pages 474–480. Springer-Verlag, 2006.

[35] A. Russo and A. Sabelfeld. Securing interaction between threads and the scheduler. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop*, pages 177–189, 2006.

[36] A. Russo and A. Sabelfeld. Dynamic vs. static flow-sensitive security analysis. In *Proceedings of the IEEE Computer Security Foundations Symposium*, 2010.

[37] A. Sabelfeld. The impact of synchronisation on secure information flow in concurrent programs. In *Proceedings of Andrei Ershov 4th International Conference on Perspectives of System Informatics*, volume 2244 of *Lecture Notes in Computer Science*, pages 225–239. Springer-Verlag, 2002.

[38] A. Sabelfeld. Confidentiality for multithreaded programs via bisimulation. In *Proceedings of the Andrei Ershov International Conference on Perspectives of System Informatics*, volume 2890 of *Lecture Notes in Computer Science*, pages 260–273. Springer-Verlag, 2003.

[39] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003.

[40] A. Sabelfeld and D. Sands. Probabilistic noninterference for multi-threaded programs. In *Proceedings of the 13th IEEE Computer Security Foundations Workshop*, pages 200–214. IEEE Computer Society, July 2000.

[41] V. Saraswat, B. Bloom, I. Peshansky, O. Tardieu, and D. Grove. *X10 Language Specification: Version 2.1.2*, Feb. 2011. Available at `http://x10.sourceforge.net/\-documentation/\-languagespec/\-x10-212.pdf`.

[42] G. Smith. A new type system for secure information flow. In *Proceedings of the Proceedings of the 14th IEEE Computer Security Foundations Workshop*, pages 115–125. IEEE Computer Society, June 2001.

[43] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of the Twenty-Fifth Annual ACM Symposium on Principles of Programming Languages*, pages 355–364, New York, NY, USA, Jan. 1998. ACM Press.

[44] D. Stefan, A. Russo, P. Buiras, A. Levy, J. C. Mitchell, and D. Mazières. Addressing covert termination and timing channels in concurrent information flow systems. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming*, New York, NY, USA, June 2012. ACM Press.

[45] T. Terauchi. A type system for observational determinism. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium*, pages 287–300, June 2008.

[46] T. Terauchi and A. Aiken. A capability calculus for concurrency and determinism. *ACM Transactions on Programming Languages and Systems*, 30(5):27:1–27:30, Sept. 2008.

[47] T. Tsai, A. Russo, and J. Hughes. A library for secure multi-threaded information flow in haskell. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium*, pages 187–202, Washington, DC, USA, 2007. IEEE Computer Society.

[48] D. Volpano and G. Smith. Probabilistic noninterference in a concurrent language. In *Proceedings of the 11th IEEE Computer Security Foundations Workshop*, pages 34–45, Washington, DC, USA, 1998. IEEE Computer Society.

[49] S. Zdancewic and A. C. Myers. Observational determinism for concurrent program security. In *Proceedings of the 16th IEEE Computer Security Foundations Workshop*, pages 29–43, Pacific Grove, California, June 2003. IEEE Computer Society.

[50] L. Zheng and A. C. Myers. Dynamic security labels and noninterference. In *Formal Aspects in Security and Trust*, Toulouse, France, Aug. 2004.