

Securing Timing Channels at Runtime

Paritosh Shroff Scott F. Smith

The Johns Hopkins University

{pari,scott}@cs.jhu.edu

Abstract

We propose a general purpose runtime framework to secure timing channels. Our technique supports higher-order function invocations and computations looping on secret data, features which none of the existing approaches fully allow. We provably eliminate external and internal timing channels in both sequential and concurrent settings, in presence of deterministic as well as nondeterministic schedulers. There is a price to be paid, however – the high computation may have to be disrupted; the low computation is nevertheless guaranteed to be unaffected. We illustrate how our approach can be realized on standard computing platforms.

1. Introduction

Secure information flow analysis aims to prevent the flow of information from secure data in a program to insecure entities. A variety of covert channels [20] for insecure flow of information have been classified in the literature; Sabelfeld and Myers [37] provides a good survey. In this paper we address timing (and termination) channels.

A timing attack is a form of side-channel attack [48] wherein the attacker attempts to gain secret information by analyzing the execution time of a program; the corresponding channel of information leakage is known as a *timing channel*. Timing channels have been used to obtain secret keys of cryptographic systems by exploiting weaknesses in their implementations [17, 6], as opposed to the underlying mathematics. Various approaches and countermeasures have been proposed to eliminate or mitigate such covert attacks, falling into several categories: algorithm specific (e.g. RSA blinding, dummy extra reduction in Montgomery algorithm, quantization of RSA decryptions [6]), hardware based (e.g. fuzzy time [11], clock randomization [16], randomization of instruction set execution and/or register usage [24, 25]), information flow analysis of synchronous (clocked) hardware circuits [18]), ad hoc [13], and programming language-based [43, 1]. Language-based techniques tend to be the most general, in that they are not tied to any specific algorithm or implementation. This paper presents a new language-based approach to secure timing channels.

Timing channels reveal classified information by manifesting *high* data as timing variations in *low* observable events. Timing channels are further classified as *external* and *internal* depending on the nature of those events – *external* channels are externally observable events such as the execution time of a program and the time gap between low outputs, while *internal* channels are internally observable events like the interleavings of concurrently running threads.

A substantial body of work [43, 38, 47, 34, 3] addresses internal timing channels. However, the only known language-based technique to secure both external and internal timing channels is the static method of program transformation by cross-copying branch slices [1]. This technique has limited expressiveness in that no looping executions on high data are allowable. In fact the execution time

of computations looping on secret data is inherently tied to that secret data, and it does not seem likely to secure such computations by purely static approaches. Additionally, none of the existing approaches support higher-order functions. Higher-order function application is also dynamic in how the code to be executed is decided only at runtime, thus making it harder for static approaches to secure them as the exact code to be secured is itself not known statically. Dynamic dispatch in object-oriented languages presents a similar challenge. In fact all existing language-based approaches to eliminate timing channels are based on static analysis of programs. Timing behavior however by its very nature is a dynamic property, and runtime techniques hold significant promise – a runtime system has direct access to the running time of the program and can use this information to good advantage, for example by altering running times just enough to eliminate the insecure cases. Static techniques can only make approximate models of the timing behavior of programs, and hence must be inherently more conservative than a runtime one.

In this paper we propose a general purpose runtime framework to secure timing channels for a language which includes higher-order functions and looping computations on secret data, features which none of the existing approaches fully support – [41, 5, 10] do allow while-loops with secret guards but only under severe restrictions. We show how both external and internal timing channels are provably eliminated in both sequential and concurrent settings, in the presence of deterministic as well as nondeterministic schedulers, modulo the effect of caches [14, 2] and resource exhaustion [37] on timing behavior of programs. There is a price to be paid, however – the high computation may have to be disrupted; the low computation is nevertheless guaranteed to be unaffected. Section 2 provides an informal overview of our technique which elaborates on these issues.

We formulate three runtime systems: λ_{seq}^{sync} , the core system, which addresses external timing channels in sequential programs, and two extensions to it, $\lambda_{conc_{\mathcal{D}}}^{sync}$ and $\lambda_{conc_{\mathcal{N}}}^{sync}$, to address both external and internal timing channels in a concurrent setting, under arbitrary deterministic and nondeterministic schedulers, respectively. We prove a strong noninterference [9] result for λ_{seq}^{sync} , and scheduler-independent strong noninterference and probabilistic noninterference [38] results for $\lambda_{conc_{\mathcal{D}}}^{sync}$ and $\lambda_{conc_{\mathcal{N}}}^{sync}$, respectively. Our probabilistic noninterference result employs discrete probability distributions, as opposed to the continuous probabilities of [38], allowing for a more direct proof technique.

2. Overview

Consider the following program in Caml [22]-like syntax:

```

let len = λz. (* number of digits in z *) in
let multiplyBigInts =
  λa. λb. let alen = len a in let blen = len b in
    let i = ref 0 in let j = ref 0 in
    while (!i < alen) {
      while (!j < blen) { ...; j++; ...; i++ }; ...
    }
in
let modBigInts =
  λc. λd. let clen = len c in let dlen = len d in ... in
let h = inputBithigh() in let r = ref 5high in
let x = inputBigIntlow() in let n = inputBigIntlow() in
let xlen = len x in let nlen = len n in
outputlow("p is about to start execution");
ifp (h == 1) then
  r := modBigInts (multiplyBigInts x x) n
else
  r := modBigInts x n;
outputlow("p just finished execution");

```

This program is inspired by the square-and-multiply algorithm – an efficient algorithm for modular exponentiation of large numbers used in many cryptographic implementations including RSA [33]. For simplicity of presentation we employ the naive $\mathcal{O}(n^2)$ school-book algorithms for multiplication and modular division of big integers; the exact algorithm used is not relevant here. The program assumes the presence of arbitrary precision arithmetic¹. It leaks information because it changes its timing behavior based on the high input to the variable h . The *program point identifier* p on if_p labels this conditional branch point. The variable h holds a single high bit, that of the RSA private exponent for example, while x and n hold arbitrary precision low integers, say the message to be decrypted and RSA’s modulus, respectively. The attacker’s goal is to find the value of h by externally observing the time lag between the two low outputs across multiple program runs.

	Run 1	Run 2
h	1^{high}	0^{high}
!r, initial	5^{high}	
x, n, x_{len}, n_{len}	$i_1^{low}, i_2^{low}, 20^{low}, 15^{low}$	
branch taken at p	then	else
running time of p	2860 units	1275 units
!r, final	$(x^2 \bmod n)^{high}$	$(x \bmod n)^{high}$
h leaked?	yes	
$i_1 = 29837429873928930498, i_2 = 908028734093623$		

Figure 1. Example (1): A Set of Raw Runs

Figure 1 tabulates the traces of a set of runs of program (1) with differing high values for h , but the same low values of x and n . Recall the computational time complexities of *multiplyBigInts* and *modBigInts* are quadratic in the lengths of their inputs, so for a given x and n , the then-branch should take more than double the time of the else-branch, a fact borne out in Figure 1. (Note, the numbers denoting the running times in Figure 1, and in the discussion to follow, have been chosen to denote the expected *relative* running times of the respective branches; hence the use

¹For technical simplicity we do not have variable length arrays for implementing big integers in the formal language syntax presented in Section 3; it should be straightforward to incorporate them if desired. We also do not formally support while-loops, but we do have higher-order functions with mutable state so loops are encodable via the ‘tying the knot’ technique. Lastly, the input and output statements used in the example are not a part of our official language syntax; they can, however, be easily incorporated, as is discussed in Section 8.

	Run 1'	Run 2'
⋮	⋮	⋮
running time of p	2860 units	
!r, final	$(x^2 \bmod n)^{high}$	$(x \bmod n)^{high}$
h leaked?	no	

Figure 2. Example (1): Optimally Synchronized Runs

of ‘unit’ as the metric for measurement of time. In our formal semantics, the number of small-steps is used to express time.) Consequently, an attacker who knows the structure of program (1) and observes that the lag between the low outputs is less than half in run 2 as compared to run 1, can correctly infer that the then-branch was taken in the first run while the else in the second, that is, h was 1 in run 1, and 0 in run 2. The delay between the low outputs constitutes an external timing channel; our goal is to secure it by decoupling the time lag from the value of h .

How to secure the external timing channel? The attacker can be deceived by forcing the running time of the branching statement to be the same regardless of which branch was taken. Figure 2 tabulates runs 1’ and 2’, which are variants of runs 1 and 2 forced to synchronize in this manner. Synchronization can be implemented by setting a timer for 2860 units – the greater of the running times of the branching statement p in runs 1 and 2 – immediately after resolving the guard ($h == 1$) but just before beginning execution of the selected branch. If the execution of the selected branch finishes before the timer goes off, as is the case in run 2’, nops can be performed until the allotted time is over, thus padding the computation of p ensuring it takes exactly 2860 units. Both the then- and else-branches thus take the same time, and we say the synchronized running time (in short, the *sync time*) of p is 2860 units. Thus, with identical timing the attacker can learn nothing about the value of h by comparing runs 1’ and 2’.

The sync time of 2860 units for the branching statement p is in fact optimal, denoting the upper bound on the running time of p for $x = i_1$ and $n = i_2$. The knowledge of the structure of program (1) and the running times of p in runs 1 and 2 can be used to obtain the optimally synchronized runs 1’ and 2’. In general, however, it is impossible to always accurately determine the upper bound on the running time of a given piece of code.

Our general solution Taking cue from the above example where we used the knowledge gained from the “test” runs 1 and 2 to generate the leak-free synchronized runs 1’ and 2’, we propose a twofold solution to secure external timing channels: a) *Pre-deployment*: Estimate as accurately as possible the upper bounds on the running times of all branching statements with high guards, by a combination of rigorous testing on a variety of inputs and/or programmer input, and b) *Post-deployment*: Use the pre-deployment estimates as the sync times for the corresponding branching statements, in conjunction with a realtime fallback technique (presented below) for securely handling situations where the sync times are found to be less than needed.

How are sync times determined? The time complexities of the then- and else-branches can be approximated by a measurement-based approach – run program (1) on a variety of inputs for h , x , and n , measuring the corresponding running times in each of the runs, and then correlate the inputs with the measured running times of the corresponding branches. We know the computational time complexities of the functions *multiplyBigInts* and *modBigInts* are quadratic in the lengths of their inputs, that is, $\mathcal{O}(a_{len} * b_{len})$ and $\mathcal{O}(c_{len} * d_{len})$ respectively, implying the running times of the then- and else-branches of program (1) are in

turn quadratic in the lengths of x and n ; say the functions approximating the running times of the then- and else-branches of program (1) are $\phi_T(x_{len}, n_{len}) = 9x_{len}n_{len} + 5x_{len} + 4n_{len} + 318$ and $\phi_F(x_{len}, n_{len}) = 4x_{len}n_{len} + 3x_{len} + n_{len} + 243$ respectively, each parameterized on the lengths of x and n denoted by x_{len} and n_{len} respectively. Observe for some x_{len} and n_{len} the estimated running time of the then-branch, $\phi_T(x_{len}, n_{len})$ units, is always greater than that of the else-branch, $\phi_F(x_{len}, n_{len})$ units; hence we can use ϕ_T as the sync timer function.

One can imagine the programmer having supplied the computational complexities for the then- and else-branches, with only the constant factors in ϕ_T and ϕ_F determined empirically above. It also may be possible to automate the process of computing the sync timer functions by employing curve fitting techniques on top of empirical data. The problem of estimating time bounds has been studied extensively in the field of Worst-Case Execution Time (WCET) analysis [45]. In general any methodology which delivers the desired accuracy may be used to derive the sync timer functions – the security guarantees of our system are independent of the precision of the sync times; however the closer to optimal the sync times can be made to be, the better would be the practical usefulness of our system.

	Run 3	Run 4
h	1^{high}	0^{high}
!r, initial	5^{high}	
x, n, x_{len}, n_{len}	$i_3^{low}, i_4^{low}, 25^{low}, 20^{low}$	
branch taken at p	then	else
running time of p	4705 units	2095 units
!r, final	$(x^2 \bmod n)^{high}$	$(x \bmod n)^{high}$
h leaked?	yes	
$i_3 = 5028234720108712878120324$		
$i_4 = 62598712387561314011$		

Figure 3. Example (1): Another Set of Raw Runs

	Run 3'	Run 4'
\vdots	\vdots	\vdots
running time of p	$\phi_T(x_{len}, n_{len}) = 5023$ units	
!r, final	$(x^2 \bmod n)^{high}$	$(x \bmod n)^{high}$
h leaked?	no	

Figure 4. Example (1): Synchronized Runs, Adequate Sync Time

What if the sync time is more than the raw running time? Prior to deployment one must strive to tune the sync timer functions to be as close to optimal as possible, so as to minimize the padding of synchronized computations – besides the needless slow-down of program computation, overestimated sync times have no undesirable side-effects. Consider runs 3 and 4 tabulated in Figure 3, with the corresponding synchronized runs 3' and 4' shown in Figure 4. For $x = i_3$ and $n = i_4$, ϕ_T provides an overestimation of the raw running times for both of p's branches. Hence in Figure 4, each of the branches completes its execution before the sync timer goes off; the residual sync time is padded away with nops as before. The attacker again learns nothing of h from timing information gleaned from runs 3' and 4'. The parameterization of the sync timer functions on low input values is an important aspect of our system.

The program transformation approach by Agat [1] exhibits the same expressiveness on this program by a static technique of cross-copying slices of branches, which implicitly encodes the above parameterization in the transformed code. Agat's technique, however, would not be applicable if the functions *multiplyBigInts* and *modBigInts* were part of a library where their source code was not

available; our dynamic synchronization technique factors in only the running times and does not need the source code, and so would remain applicable to such library calls. Our technique is more general in other ways as well, a topic that will be taken up later in this section.

What if the sync time is less than the raw running time? Say the programmer had expected program (1) to be used mainly on inputs of lengths less than or equal to 32 digits for x and n ; hence, prior to deployment she had ensured that the sync timer function ϕ_T was appropriate for all inputs to x and n of lengths 32 digits or less. Recall the lengths of i_3 and i_4 in runs 3 and 4 were both less than 32 digits.

	Run 5'	Run 6'
h	0^{high}	1^{high}
!r, initial	5^{high}	
x, n, x_{len}, n_{len}	$i_5^{low}, i_6^{low}, 40^{low}, 30^{low}$	
branch taken at p	else	then
running time of p [†]	$\phi_T(x_{len}, n_{len}) = 11438$ units	
branch execution aborts?	no	yes [‡]
!r, final	$(x \bmod n)^{high}$	5^{high} [‡]
h leaked?	no	
$i_5 = 3321654873210031296870149807841203498791$		
$i_6 = 103216510606489412030654984321$		
[†] raw running time of then-branch is 12003 units, and else-branch is 5837 units.		
[‡] computation is aborted after the allotted sync time is over leaving the logically incorrect value 5^{high} in !r.		

Figure 5. Example (1): Synchronized Runs, Inadequate Sync Time

Now consider the synchronized runs 5' and 6' tabulated in Figure 5; the sync time of $\phi_T(x_{len}, n_{len}) = 11438$ units for $x = i_5$ and $n = i_6$, is less than the corresponding raw running time of the then-branch but is greater than that of the else-branch. The run 5' terminates with $(x \bmod n)^{high}$ in !r; however, if the then-branch in run 6' were allowed to complete, the attacker would observe a timing difference between runs 6' and 5', and thus learn the value of h . Hence in order to bluff the attacker the computation of the then-branch is simply *aborted* in run 6' when the allotted sync time is over, and the subsequent low output performed. Once aborted, the high computation (the portion of the program computation manipulating high data) has gone logically wrong, and the program has entered a *bluff computation* state, the purpose of which is to hide the abortion from a low observer, and to keep him from gaining information on h .

As the execution of the then-branch was terminated before the assignment to r could take place, the old and now logically incorrect value 5^{high} remains in !r at the end of run 6'; this value is to be viewed as a "placeholder" in that it is not useful but prevents information leakage. Note, only branching statements with high guards are ever synchronized in our system, implying only high branch computations can ever be aborted; hence placeholder values, which are a result of such abortions, are guaranteed to be always high.

The bluffing technique can be viewed as a generalization of the lenient execution model [7] for out-of-bounds array indices.

Bluff computation is only a fallback, not a norm Given that the programmer had not expected an input of length as large as 40 digits for x , i_5 could have been a malicious input fed by an attacker trying to covertly gain classified information on h . In fact the programmer should have disallowed invalid values greater than 32 digits for x and n in the first place by inserting

```
if ( $x_{len} > 32 \parallel n_{len} > 32$ ) then
  outputlow("invalid input"); exit 1;
```

immediately after the line ‘let $x_{len} = len\ x$ in let $n_{len} = len\ n$ in’ in program (1). This change would eliminate the need for bluff computation due to too-large an input.

In general, if the programmer is careful enough to disallow all spurious low values from reaching timing-sensitive pieces of code with running times dependent on low data, he can be quite certain the bluff computation will not be needed for them upon deployment.

Handling computations looping on high data In program (1) the value of $!r$ was never used. Now consider program (2), which is program (1) followed by the addendum below.

```

...
let  $y = !r$  in let  $y_{len-max} = n_{len}$  in
let  $h' = inputBit^{high}()$  in let  $r' = ref\ 7^{high}$  in
outputlow (“p’ is about to start execution”);
ifp’ ( $h' == 1$ ) then
   $r' := modBigInts\ (multiplyBigInts\ y\ y)\ n$ 
else
   $r' := modBigInts\ y\ n$ ;
outputlow (“p’ just finished execution”);

```

Note, y is high in program (2) as r is set under the high guard ($h == 1$) in program (1). Also note the computations of both $modBigInts\ (multiplyBigInts\ y\ y)\ n$ and $modBigInts\ y\ n$ are looping on the lengths of y and n , so the running time of p' is parametric in y and n . We use ϕ_T as the sync timer function for p' as well; however, parameterizing the sync time of p' on the length of y would reveal the secret value of h – the lag between the low outputs on each side of p' would then depend on the value of y which in turn depends on the secret bit in h . Recall the sync times of p in runs 3'–6' were solely parameterized on the low lengths of the low values x and n ; hence despite the difference in the sync times in Figures 4 and 5 no information was leaked. In order to prevent information leakage the sync time of p' must be parameterized exclusively on low values.

Note the only values possible for y in program (2) are either 5^{high} , $(x \bmod n)^{high}$, or $(x^2 \bmod n)^{high}$, implying its value can be no bigger than that of n , that is, the length of y is equal to the length of n in the worst case. Given this publicly known worst-case length for y we can conservatively parameterize the sync time of p' on $y_{len-max}$ (which is equal to n_{len}) and n_{len} . Figure 6 tabulates a set of such synchronized runs for program (2) as extensions to run 6'. The attacker again learns nothing since runtimes are uniform.

	Run 7'	Run 8'
h		1^{high}
$!r$, initial		5^{high}
x, n, x_{len}, n_{len}	$i_5^{low}, i_6^{low}, 40^{low}, 30^{low}$	
branch taken at p	then	
running time of p	$\phi_T(x_{len}, n_{len}) = 11438$ units	
branch execution aborts?	yes	
$!r$, final		$5^{high\dagger}$
$y, y_{len-max}$	$5^{high\dagger}, n_{len}$	
h'	1^{high}	0^{high}
$!r'$, initial		7^{high}
branch taken at p'	then	else
running time of p'	$\phi_T(y_{len-max}, n_{len}) = 8688$ units	
branch execution aborts?	no	
$!r'$, final	$(y^2 \bmod n)^{high} = 25^{high\dagger}$	$(y \bmod n)^{high} = 5^{high\dagger}$
h or h' leaked?	no	
\dagger high placeholder value	\ddagger garbled high data	

Figure 6. Example (2): Worst-Case Sync Time

The above example illustrates our general technique for preventing timing leaks in presence of computations looping on secret

data: parameterize the corresponding sync times over low conservative approximations of those high data.

Agat [1] does not allow looping computations with high guards. While-loops with secret guards are supported in [41, 5], but only in concurrent settings, and only if such loops are not followed by low events. Our technique has the advantage of support for high looping computations, in both sequential and concurrent settings, and without the aforementioned restriction.

Our approach of parameterizing the sync times of high looping computations with the worst-case low approximations of the corresponding high guards is analogous to the “worst-case principle” proposed by Agat and Sands [2]; the latter, however, focuses on rewriting algorithms with low termination conditions corresponding to the worst-case executions of such high loops, while we achieve a similar effect dynamically.

Placeholder values may garble high data Observe the final values of $!r'$ in runs 7' and 8' are logically incorrect due to the logical incorrectness of y . In general once a branch computation is aborted the resulting placeholder values may pollute the high computation with garbled high data – taking the analogy of program (1) to an RSA implementation, the decrypted message may be garbled. In general, an execution entering a bluff computation state may not be able to complete some high computation tasks due to garbled data, and programs must be written with this contingency in mind, not unlike how the case of “network down” needs to be accounted for programmatically. Note, any high computation task not influenced by the garbled data will nonetheless function correctly.

A Higher-Order Example Examples (1) and (2) were simple first-order programs. Now consider the following higher-order variant of program (1), where ‘.’ is a shorthand for any variable not found free in the body of the corresponding function.

```

...
ifp''' ( $h == 1$ ) then
   $f := (\lambda.. r := modBigInts\ (multiplyBigInts\ x\ x)\ n)$ 
else
   $f := (\lambda.. r := modBigInts\ x\ n)$ ;
outputlow (“p''' is about to start execution”);
(!f) ()p''';
outputlow (“p''' just finished execution”);

```

The program point identifier p''' denotes the corresponding function application statement; function application in the presence of first-class functions is a form of branching – the code to be executed next depends on the function flowing into the application site. Analogously, the running times of such function application statements depend on the exact functions flowing into them, just as the running times of branching statements depend on the actual guard flowing into them. Hence, akin to runs 1–4, program (3) will also leak the value of h by exhibiting variation in the running time of p''' based on h 's value. The synchronization technique discussed above for conditional branching statements is applicable here as well; ϕ_T could be used as the sync timer function for p''' .

We are not aware of any existing techniques for securing timing channels which support higher-order functions; we believe this is a new contribution of our approach.

Securing timing channels in presence of concurrency We now show how our technique also may be used to secure timing channels in concurrent programs. Consider the following program in Caml-like syntax, adapted from [47],

```

let  $spin = \lambda n. let\ i = ref\ 0$  in while ( $!i < n$ ) {  $i++$  } in
let  $l = ref\ 0^{low}$  in
( $if_{p''''}$  ( $h == 1$ ) then  $spin(500^{low})$  else ();  $\parallel spin(50^{low});$ )
( $output^{low}(l)$   $\parallel l := 1^{low}$ )

```

n	$\in \mathbb{N}$	<i>small-step semantics up-counter</i>
d	$\in \mathbb{N} \cup \{\infty\}$	<i>down-counter</i>
p		<i>program point identifier</i>
x		<i>variable</i>
i	$\in \mathbb{Z}$	<i>integer</i>
b	$::= \text{true} \mid \text{false}$	<i>boolean</i>
loc		<i>heap location</i>
\oplus_a	$::= + \mid - \mid * \mid /$	<i>arithmetic binary operator</i>
\oplus_r	$::= < \mid > \mid == \mid !=$	<i>relational binary operator</i>
\oplus	$::= \oplus_a \mid \oplus_r$	<i>binary operator</i>
v	$::= x \mid i \mid b \mid \lambda x. p \mid loc$	<i>value</i>
p	$::= v \mid p \oplus p \mid \text{let } x = p \text{ in } p$	<i>program (source syntax)</i>
	$\mid \text{if}_{p, \bar{v}}^v p \text{ then } p \text{ else } p \mid p(p)_{p, \bar{v}}^v$	
	$\mid \text{ref } p \mid !p \mid p := p$	
R	$::= \bullet \mid R \oplus p \mid v \oplus R \mid \text{let } x = R \text{ in } p$	<i>reduction context</i>
	$\mid \text{if}_{p, \bar{v}}^v R \text{ then } p \text{ else } p \mid R(p)_{p, \bar{v}}^v \mid v(R)_{p, \bar{v}}^v$	
	$\mid \text{ref } R \mid !R \mid R := p \mid v := R$	
s	$::= \langle p \rangle^v$	<i>synchronization construct</i>
e	$::= p \mid R[s]$	<i>expression (runtime syntax)</i>
H	$: \{loc\} \rightarrow \{\bar{v}\}$	<i>heap (memory)</i>
C	$::= (H, e)_d^n$	<i>runtime configuration</i>
ϕ^n	$: \mathbb{Z} \times \dots \times \mathbb{Z} \rightarrow \mathbb{N}$	<i>sync timer function</i>
Φ	$: \{\bar{p}\} \rightarrow \{\phi^n \vee \emptyset\}$	<i>table of sync timer functions</i>
\bar{v}	$::= i \mid b$	<i>integer/boolean value</i>

Figure 7. λ_{seq}^{sync} : Syntax Grammar

where the infix operator \parallel denotes parallel composition. The variable h holds a single high bit, while the reference variable l is shared between the program fragments to be executed concurrently.

Observe the raw running time of the then-branch in the left fragment is considerably more than that of $spin(50^{low})$ in the right fragment; the converse is true for the else-branch. Hence under most schedulers (e.g. round-robin), if h is 1 then the assignment $l := 1^{low}$ is likely to precede output $^{low}(l)$, resulting in 1 being the likely output; conversely, if h is 0 the output is likely to be 0. In summary, h can affect the *internal timing behavior* of program (4), which in turn can affect the value of l that is output to a low observer, thus indirectly leaking the value of h . Such leaks due to variations in the internal timing behavior of concurrent programs are referred to as *internal timing leaks*, with the internal timing behavior constituting the *internal timing channel*. Note, unlike external timing leaks, the attacker need not have a stop-watch to record the timings of the low outputs, the values of the low outputs themselves reveal classified information. Hence internal timing channels are easier to secure than external timing channels – only the values of the low outputs need to be decoupled from high data to secure internal timing channels, whereas both the low outputs and their timings need to be insulated from secret information to close external timing channels.

Observe that the root of the timing leak in program (4) is the asymmetry in the execution times of the then- and else-branches of the branching statement p''' . Hence our dynamic synchronization technique will also secure the internal timing channel here. In general, the technique secures all internal and external timing channels in arbitrary concurrent programs, as is proved in Sections 5 and 6.

3. The λ_{seq}^{sync} Runtime System

We now formalize the λ_{seq}^{sync} language and define its semantics.

λ_{seq}^{sync} is a sequential language with mutable state, higher-order functions, conditional branchings and let-bindings. The language syntax appears in Figure 7. Note the overbar notation denotes comma-separated sequence of zero or more items, so for example \bar{v} denotes some v_1, \dots, v_n ; the subscripted overbar notation denotes a sequence of fixed length as indicated by the subscript,

so for example $\bar{v}_k = v_1, \dots, v_k$ and $\{\overline{loc_k \mapsto v_k}\} = \{loc_1 \mapsto v_1, \dots, loc_k \mapsto v_k\}$. The runtime heap H is a partial function from heap locations to values. Program point identifiers p are used to uniquely label conditional branching and function application sites. They are intended to be automatically generated, but we embed them in program syntax for technical convenience. Also we use the terms ‘program point’ and ‘program point identifier’ interchangeably throughout this paper.

The table of sync timer functions Φ associates each program point with either a sync timer function ϕ^n , or \emptyset denoting lack of such a function – sync timer functions need to be defined only for timing-sensitive statements whose running times depend on high data. The sync timer function ϕ^n is parameterized on an n -tuple of integers and returns a natural number denoting the corresponding sync time; recall for example how the sync timer function ϕ_T in Section 2 was parameterized on a pair of integers.

The subscripts \bar{v} on the conditional branching and application sites must evaluate to low-security integral parameters during computation, and are then fed to the associated sync timer function.

The superscript v on the branching and application sites denotes the default placeholder value for the expression – in case the synchronized execution is aborted the associated default placeholder value is inserted in its place to allow the rest of the computation to proceed, thus hiding the abortion from a low observer. The default placeholder value can be any value of the same type as the computation it is intended to replace – this is ensured by our type system.

The synchronization construct $\langle p \rangle^v$ is a runtime construct denoting the synchronized computation of the enclosed program p with the default placeholder value v ; $\langle p \rangle^v$ occurs only at runtime, and so we distinguish the runtime expressions e from source programs p .

The set of free variables is defined as follows,

Definition 3.1 (Free Variables).

1. (Value).
 - (a) (Variable). $free(x) = \{x\}$; and
 - (b) (Integer, Boolean, Heap Location). $free(i) = free(b) = free(loc) = \emptyset$; and
 - (c) (Function). $free(\lambda x. p) = free(p) - \{x\}$.
2. (Program).
 - (a) (Binary Operation). $free(p \oplus p') = free(p) \cup free(p')$; and
 - (b) (Let). $free(\text{let } x = p \text{ in } p') = free(p) \cup (free(p') - \{x\})$; and
 - (c) (If). $free(\text{if}_{p, \bar{v}_k}^v p \text{ then } p_1 \text{ else } p_2) = free(p) \cup free(p_1) \cup free(p_2) \cup free(v_1) \cup \dots \cup free(v_k) \cup free(v)$; and
 - (d) (App). $free(p(p')_{p, \bar{v}_k}^v) = free(p) \cup free(p') \cup free(v_1) \cup \dots \cup free(v_k) \cup free(v)$; and
 - (e) (Ref). $free(\text{ref } p) = free(p)$; and
 - (f) (Deref). $free(!p) = free(p)$; and
 - (g) (Assign). $free(p := p') = free(p) \cup free(p')$.
3. (Synchronization Construct). $free(\langle p \rangle^v) = free(p) \cup free(v)$.
4. (Expression with a Synchronization Construct). $free(R[s]) = free(s) \cup (free(R[x]) - \{x\})$, where $R[s] = (R[x])[s/x]$.
5. (Heap). $free(H) = \bigcup_{loc \in dom(H)} free(H(loc))$.
6. (Heap, Expression). $free(H, e) = free(H) \cup free(e)$.

We write $e[e'/x]$ to denote the capture-avoiding substitution of all free occurrences of x in e with e' , and $H[e/x]$ to denote the heap H' such that $dom(H') = dom(H)$ and $\forall loc \in dom(H). H'(loc) = H(loc)[e/x]$. In addition we use $(H, e)[e'/x]$ as shorthand for $(H[e'/x], e[e'/x])$, and the multi-substitution $\mathbf{X}[v_k/x_k]$, where $\mathbf{X} ::= H \mid e$, as shorthand for

$((\mathbf{X}[v_1/x_1])[v_2/x_2]) \dots [v_k/x_k]$. Also we write “ e is a program” if $e = p$, for some p , “ e is a value” if $e = v$, for some v , and “ e has a sync construct” if there exists a R and a s , such that $e = R[s]$.

For technical ease we sometimes view functions as a set of mappings from the elements in its domain to those in its range; so for example $loc \mapsto v \in H$ is equivalent to asserting $H(loc) = v$. The complement operator \setminus on a set of mappings, say H , is defined as $H \setminus loc = \{loc' \mapsto v' \mid loc' \mapsto v' \in H \wedge loc \neq loc'\}$; the update operation is then defined as $H[loc \mapsto v] = H \setminus loc \cup \{loc \mapsto v\}$. Further for any binary relation \mathcal{R} such that $\mathcal{R} \subseteq A \times B$, $A = \{\bar{a}\}$ and $B = \{\bar{b}\}$, we use the shorthand ‘ $\bar{a}_k \mathcal{R} \bar{b}$ ’ for ‘ $\forall 1 \leq i \leq k. a_i \mathcal{R} b$ ’, and symmetrically, ‘ $a \mathcal{R} \bar{b}_k$ ’ for ‘ $\forall 1 \leq i \leq k. a \mathcal{R} b_i$ ’.

Operational Semantics Figure 8 gives the syntax-directed small-step operational semantics for λ_{seq}^{sync} . The small-step reduction relation \longrightarrow_{Φ} , parameterized by a table of sync timer functions Φ , relates heap/expression configurations $(H, e)_d$. Here d is the *down-counter* and is decremented at each step; it denotes the maximum number of small-steps available for the expression e to complete its execution. If $d = \infty$, there is an unbounded number of available steps (we consider $(\infty - 1) = \infty$). The n -step reflexive and transitive closure of \longrightarrow_{Φ} is denoted as \longrightarrow_{Φ}^n . We use the small-step counter n as our abstraction for execution time; however, our techniques are generalizable regardless of the granularity of the abstraction used.

The computation of program p starts in an initial configuration $(\emptyset, p)_{\infty}$. During the course of computation, if the execution time of a branching or application statement needs to be synchronized (i.e., the premise $\Phi(p)(\bar{i}_k) = n$ holds in the IF-SYNC and APP-SYNC rules), the corresponding code is placed in a synchronization construct $(\langle p_i \rangle^v$ in IF-SYNC, and $\langle p[v/x] \rangle^{v'}$ in APP-SYNC) with the associated default placeholder value (v and v' , respectively); the down-counter is then set to the indicated sync time (n). Since $n \in \mathbb{N}$, synchronized computation is limited to a finite number of steps by the SYNC rule – when the down-counter reaches 0, the SYNC rule is no longer applicable as -1 is not a valid down-counter.

If the computation being synchronized finishes before the sync time is over, nop’s are performed via the SYNC-PAD rule until the down-counter is 0, thus ensuring the overall synchronization takes exactly the sync time number of steps to complete from an external observer’s point of view; the computed value is then unwrapped from the synchronization construct and the down-counter is reset to ∞ via the SYNC-DONE rule, while the default placeholder value is simply thrown away. However, if the computation under synchronization does not converge to a value in the allotted sync time, then it is aborted via the SYNC-ABORT rule. The default placeholder value is inserted in place of the unfinished computation, and the down-counter is reset to ∞ indicating the synchronization has completed and non-synchronized computation can resume.

Definition 3.2 (Canonical Configuration). *A configuration $(H, e)_d$ is canonical iff either e is a program and $d = \infty$, or e has a synchronization construct and $d \neq \infty$.*

Definition 3.3 (Canonical Derivation). *A derivation $(H, e)_d \longrightarrow_{\Phi}^n (H', e')_{d'}$ is canonical iff the configuration at each node of its derivation tree is canonical.*

An important aspect of our synchronization technique is that only top-level branching and function application statements are synchronized, as indicated in the IF-SYNC and APP-SYNC rules by the infinite down-counter on the left side of the reduction relation, and the premise $d \neq \infty$ in the IF and APP rules. Hence if the a branching or function application site is visited (revisited) inside

a synchronized computation, as for example in case of nested (recursive) computations, the sync time is *not* reset at each such visit; the corresponding code is already under synchronization as part of the encapsulating computation.

The premise $\Phi(p) = \emptyset$ in the IF rule (and correspondingly in the APP rule) indicates the guard flowing into the corresponding branching site is of low security, and so its computation is not to be synchronized. A type system, to be presented in Section 3.1, ensures the table of sync timer functions Φ correctly identifies such branching and application statements based on the security levels of the guards and the functions flowing into them.

We now formally state some properties of the operational semantics.

Lemma 3.4 (Properties of Operational Semantics).

1. (Canonicity of Derivation). *If $(H, p)_{\infty} \longrightarrow_{\Phi}^n (H', e')_{d'}$ then this derivation is canonical.*
2. (Program to Program). *If $(H, p)_d \longrightarrow_{\Phi}^n (H', p')_{d'}$ then $d' = d - n$. (Hence, if $d = \infty$ then $d' = \infty$.)*
3. (A Stuck State). *For any H and p , there does not exist a Φ , H' , e' , and d' , such that $(H, p)_0 \longrightarrow_{\Phi} (H', e')_{d'}$.*
4. (Time-Fixed Computation). *If $(H, \langle p \rangle^v)_n \longrightarrow_{\Phi}^{n'} (H', e')_{d'}$ then either:*
 - (a) *For some p' , $e' = \langle p' \rangle^{v'}$ and $d' = n - n'$; or*
 - (b) *For some v' , $e' = v'$, $n' = n + 1$ and $d' = \infty$.*
5. (IF-SYNC). *If $\Phi(p)(\bar{i}_k) = n$ and $(H, \text{if}_{p, \bar{i}_k}^v b \text{ then } p_1 \text{ else } p_2)_{\infty} \longrightarrow_{\Phi}^{n'} (H', v')_{d'}$ then $n' = n + 2$ and $d' = \infty$. (Hence, if $n = 0$ then $n' = 2$.)*
6. (APP-SYNC). *If $\Phi(p)(\bar{i}_k) = n$ and $(H, (\lambda x. p)(v)_{p, \bar{i}_k}^v)_{\infty} \longrightarrow_{\Phi}^{n'} (H', v'')_{d'}$ then $n' = n + 2$ and $d' = \infty$. (Hence, if $n = 0$ then $n' = 2$.)*

Proof. 1. (Canonicity of Derivation). By induction on the derivation of $(H, p)_{\infty} \longrightarrow_{\Phi}^n (H', e')_{d'}$.

2. (Program to Program). By induction on the derivation of $(H, p)_d \longrightarrow_{\Phi}^n (H', p')_{d'}$.
3. (A Stuck State). By induction on the structure of p given the syntax-directedness of the semantics rules in Figure 8.
4. (Time-Fixed Computation). By induction on the derivation of $(H, \langle p \rangle^v)_n \longrightarrow_{\Phi}^{n'} (H', e')_{d'}$.
5. (IF-SYNC). By IF-SYNC and Lemma 3.4[4].
6. (APP-SYNC). By APP-SYNC and Lemma 3.4[4]. □

3.1 A Type System for λ_{seq}^{sync}

As discussed above, the operational semantics of λ_{seq}^{sync} relies on the table of sync timer functions Φ to indicate whether the computation of a branching or application statement is to be synchronized or not – if a sync timer function is defined then the corresponding statement is to be synchronized, otherwise not. In order to prevent timing leaks only the branching and application statements with running times dependent on high data need to be synchronized, that is, those with high guards or functions flowing into them. For simplicity we employ a basic monomorphic type system to identify such statements in this paper; more expressive polymorphic type systems [32] could be employed to better effect, as could dynamic techniques [39].

Let (\mathcal{L}, \leq) be a lattice whose elements, denoted by ℓ and pc , represent *security levels*. Following Denning [8] we typically use the meta-variable pc , rather than ℓ , when considering information obtained by observing the value of the “program counter”. We write \perp for \mathcal{L} ’s least element. The type grammar and subtyping

$\frac{\text{BINOP} \quad i_1 \oplus i_2 = v}{(H, i_1 \oplus i_2)_d \longrightarrow_{\Phi} (H, v)_{d-1}}$	$\frac{\text{LET}}{(H, \text{let } x = v \text{ in } p)_d \longrightarrow_{\Phi} (H, p[v/x])_{d-1}}$	$\frac{\text{IF-SYNC} \quad i \in \{1, 2\} \quad (b_1, b_2) = (\text{true}, \text{false}) \quad \Phi(\mathbf{p})(\overline{i_k}) = n}{(H, \text{if}_{\mathbf{p}, \overline{i_k}}^v b_i \text{ then } p_1 \text{ else } p_2)_{\infty} \longrightarrow_{\Phi} (H, \langle p_i \rangle^v)_n}$
$\frac{\text{IF} \quad i \in \{1, 2\} \quad (b_1, b_2) = (\text{true}, \text{false}) \quad \Phi(\mathbf{p}) = \emptyset \vee d \neq \infty}{(H, \text{if}_{\mathbf{p}, \overline{i_k}}^v b_i \text{ then } p_1 \text{ else } p_2)_d \longrightarrow_{\Phi} (H, p_i)_{d-1}}$	$\frac{\text{APP-SYNC} \quad \Phi(\mathbf{p})(\overline{i_k}) = n}{(H, (\lambda x. p) (v)_{\mathbf{p}, \overline{i_k}}^{v'})_{\infty} \longrightarrow_{\Phi} (H, \langle p[v/x] \rangle^v)_n}$	
$\frac{\text{APP} \quad \Phi(\mathbf{p}) = \emptyset \vee d \neq \infty}{(H, (\lambda x. p) (v)_{\mathbf{p}, \overline{i_k}}^{v'})_d \longrightarrow_{\Phi} (H, p[v/x])_{d-1}}$	$\frac{\text{SYNC} \quad (H, p)_n \longrightarrow_{\Phi} (H', p')_{n-1}}{(H, \langle p \rangle^v)_n \longrightarrow_{\Phi} (H', \langle p' \rangle^v)_{n-1}}$	$\frac{\text{SYNC-PAD}}{(H, \langle v \rangle^v)_n \longrightarrow_{\Phi} (H, \langle v \rangle^v)_{n-1}}$
$\frac{\text{SYNC-DONE}}{(H, \langle v \rangle^v)_0 \longrightarrow_{\Phi} (H, v)_{\infty}}$	$\frac{\text{SYNC-ABORT} \quad p \text{ is not a value}}{(H, \langle p \rangle^v)_0 \longrightarrow_{\Phi} (H, v)_{\infty}}$	$\frac{\text{REF} \quad \text{loc is fresh}}{(H, \text{ref } v)_d \longrightarrow_{\Phi} (H \cup \{\text{loc} \mapsto v\}, \text{loc})_{d-1}}$
$\frac{\text{DEREF} \quad H(\text{loc}) = v}{(H, !\text{loc})_d \longrightarrow_{\Phi} (H, v)_{d-1}}$	$\frac{\text{ASSIGN} \quad \text{loc} \in \text{dom}(H)}{(H, \text{loc} := v)_d \longrightarrow_{\Phi} (H[\text{loc} \mapsto v], v)_{d-1}}$	$\frac{\text{CONTEXT} \quad (H, e)_d \longrightarrow_{\Phi} (H', e')_{d'}}{(H, R[e])_d \longrightarrow_{\Phi} (H', R[e'])_{d'}}$

Figure 8. $\lambda_{seq}^{\text{sync}}$: Syntax-Directed Small-Step Operational Semantics

$t ::= \text{int} \mid \text{bool} \mid \tau \rightarrow \tau \mid \tau \text{ref}$	<i>ground type</i>
$\tau ::= t^{\ell}$	<i>type</i>
$\Gamma : \{\overline{x}\} \rightarrow \{\overline{\tau}\}$	<i>type of environment</i>
$\mathcal{H} : \{\overline{\text{loc}}\} \rightarrow \{\overline{\tau}\}$	<i>type of heap</i>
$\hat{t} ::= \text{int} \mid \text{bool}$	<i>ground int/bool type</i>
$\hat{\tau} ::= \hat{t}^{\ell}$	<i>int/bool type</i>

$\ell \leq \ell' \quad \ell \leq \ell' \quad \ell' \triangleleft \tau$	$\tau_1' \leq \tau_1 \quad \tau_2 \leq \tau_2'$
$\hat{\ell}^{\ell} \leq \hat{\ell}'^{\ell'} \quad \tau \text{ref}^{\ell} \leq \tau \text{ref}^{\ell'}$	$(\tau_1 \rightarrow \tau_2)^{\ell} \leq (\tau_1' \rightarrow \tau_2')^{\ell}$

Figure 9. $\lambda_{seq}^{\text{sync}}$: Type Grammar and Subtyping Rules

rules are defined in Figure 9. The annotation ℓ on the ground type t in type τ reflects the security level of the information the corresponding expression may carry. The function *secllevel* on types returns the associated security levels: $\text{secllevel}(t^{\ell}) = \ell$, for any t and ℓ . The binary predicate $\ell \triangleleft \tau$ (read: ℓ guards τ) holds iff $\ell \leq \text{secllevel}(\tau)$. The type rules appear in Figure 10. In a type judgement $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$, pc denotes the security level of the context, while ℓ represents the least upper bound of all “low” security levels in \mathcal{L} , that is, for any $\ell' \in \mathcal{L}$, ℓ' is considered low iff $\ell' \leq \ell$, otherwise ℓ' is high. The function *type* on binary operators, used in the type rule *binop*, is defined as, $\text{type}(\oplus_a) = \text{int}$ and $\text{type}(\oplus_r) = \text{bool}$. We use the following shorthand: ‘ $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \overline{e_k} : \tau$ ’ abbreviates ‘ $\forall 1 \leq i \leq k. pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e_i : \tau$ ’, ‘ $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \overline{e_k} : \overline{\tau_k}$ ’ abbreviates ‘ $\forall 1 \leq i \leq k. pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e_i : \tau_i$ ’, and ‘ $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e_k, e'_k : \overline{\tau_k}$ ’ abbreviates ‘ $\forall 1 \leq i \leq k. pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e_i, e'_i : \tau_i$ ’.

The type system is fairly standard with much of the notation adapted from [32]. It mainly serves to identify the branching and application statements needing synchronization to secure timing channels. It does so by ensuring the existence of associated sync timer functions, with appropriate arities, for such statements via the type rules *if-sync* and *app-sync*; the nonexistence of sync timer functions for all other branching and application statements not needing synchronization is ensured by the *if* and *app* rules, respectively. As pointed out in the discussion on the operational semantics, only top-level branching and application statements, that is, only those that appear in low contexts, with high guards or func-

tions flowing into them, need to be synchronized for securing timing channels; this is captured by the premise $pc \leq \ell \wedge \ell' \not\leq \ell \wedge \Phi(\mathbf{p}) = \phi^k$ in the *if-sync* and *app-sync* rules, while the premise $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \overline{v_k} : \text{int}^{\ell}$ ensures all parameters to the corresponding sync timer function ϕ^k are low integers. The other statements, not needing synchronization, are indicated by the premise $(pc \leq \ell \wedge \ell' \leq \ell \wedge \Phi(\mathbf{p}) = \emptyset) \vee pc \not\leq \ell$ in the *if* and *app* rules; the subscripts $\overline{v_k}$ and the superscript v are in fact unused in the *if* and *app* rules.

We now formally state some properties of the type system and then prove the type safety of $\lambda_{seq}^{\text{sync}}$. We start by defining some notation. The subtyping relation between environments $\Gamma \leq \Gamma'$ holds iff $\text{dom}(\Gamma) = \text{dom}(\Gamma')$ and $\forall x \in \text{dom}(\Gamma). \Gamma(x) \leq \Gamma'(x)$. The type judgement $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (H, e) : \tau$ holds iff $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H$ and $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$ hold.

Lemma 3.5 (Properties of the Type System).

1. (*Program Counter*).
 - (a) (*Typing Invariant*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$ then $pc \triangleleft \tau$.
 - (b) (*Value*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau$ and $pc' \triangleleft \tau$ then $pc', \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau$.
 - (c) (*Bump-Up*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau$, $pc \leq pc'$ and $pc' \triangleleft \tau$ then $pc', \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau$.
 - (d) (*Bump-Down*).
 - i. (“Low” to further “Low”). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$ and $pc_{\text{sub}} \leq pc \leq \ell$ then $pc_{\text{sub}}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$.
 - ii. (“High” to lesser “High”). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau$, $pc_{\text{sub}} \leq pc$ and $pc_{\text{sub}}, pc \not\leq \ell$ then $pc_{\text{sub}}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau$.
2. (*Subtyping*).
 - (a) (*Type of Environment*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$ and $\Gamma_{\text{sub}} \leq \Gamma$ then $pc, \Gamma_{\text{sub}}, \mathcal{H} \vdash_{\Phi \ell} e : \tau$.
 - (b) (*Expression*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$ and $\tau \leq \tau'$ then $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau'$.
3. (*Heap Update*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H$ and $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \mathcal{H}(\text{loc})$ then $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H[\text{loc} \mapsto v]$.
4. (*Extension of Type of Heap*). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$ and $\mathcal{H} \subseteq \mathcal{H}'$ then $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} e : \tau$.

$$\begin{array}{c}
\text{var} \\
\frac{\Gamma(x) \leq \tau \quad pc \triangleleft \tau}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} x : \tau} \quad \text{int} \quad \frac{pc \leq \ell'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} i : \text{int}^{\ell'}} \quad \text{bool} \quad \frac{pc \leq \ell'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} b : \text{bool}^{\ell'}} \quad \text{func} \\
\frac{pc \leq \ell' \quad \ell', \Gamma[x \mapsto \tau], \mathcal{H} \vdash_{\Phi \ell} p : \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \lambda x. p : (\tau \rightarrow \tau')^{\ell'}} \\
\\
\text{location} \\
\frac{pc \leq \ell' \quad \mathcal{H}(\text{loc}) = \tau \quad \ell' \triangleleft \tau}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \text{loc} : \tau \text{ref}^{\ell'}} \quad \text{binop} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p_1, p_2 : \text{int}^{\ell'}}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p_1 \oplus p_2 : \text{type}(\oplus)^{\ell'}} \quad \text{let} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau \quad pc, \Gamma[x \mapsto \tau], \mathcal{H} \vdash_{\Phi \ell} p' : \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \text{let } x = p \text{ in } p' : \tau'} \\
\\
\text{if-sync} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \text{bool}^{\ell'} \quad pc \leq \ell \wedge \ell' \not\leq \ell \wedge \Phi(p) = \phi^k \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \overline{v_k} : \text{int}^{\ell'} \quad \ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} p_1, p_2, v : \tau}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \text{if}_{\overline{p}, \overline{v_k}}^v p \text{ then } p_1 \text{ else } p_2 : \tau} \\
\\
\text{app-sync} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : (\tau \rightarrow \tau')^{\ell'} \quad pc \leq \ell \wedge \ell' \not\leq \ell \wedge \Phi(p) = \phi^k \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p' : \tau \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \overline{v_k} : \text{int}^{\ell'} \quad \ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p(p')_{\overline{p}, \overline{v_k}}^v : \tau'} \\
\\
\text{if} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \text{bool}^{\ell'} \quad (pc \leq \ell \wedge \ell' \leq \ell \wedge \Phi(p) = \emptyset) \vee pc \not\leq \ell \quad \ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} p_1, p_2 : \tau}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \text{if}_{\overline{p}, \overline{v_k}}^v p \text{ then } p_1 \text{ else } p_2 : \tau} \\
\\
\text{app} \\
\frac{(pc \leq \ell \wedge \ell' \leq \ell \wedge \Phi(p) = \emptyset) \vee pc \not\leq \ell \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : (\tau \rightarrow \tau')^{\ell'} \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p' : \tau \quad \ell' \triangleleft \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p(p')_{\overline{p}, \overline{v_k}}^v : \tau'} \quad \text{sync} \\
\frac{pc \leq \ell \wedge \ell' \not\leq \ell \quad \ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} p, v : \tau}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \langle p \rangle^v : \tau} \\
\\
\text{ref} \\
\frac{pc \leq \ell' \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau \quad \ell' \triangleleft \tau}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \text{ref } p : \tau \text{ref}^{\ell'}} \quad \text{deref} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau \text{ref}^{\ell'} \quad \ell' \triangleleft \tau \quad \tau \leq \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} !p : \tau'} \quad \text{assign} \\
\frac{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau \text{ref}^{\ell'} \quad \ell' \triangleleft \tau \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p' : \tau_{\text{sub}} \quad \tau_{\text{sub}} \leq \tau, \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p := p' : \tau'} \\
\\
\text{sync-in-context} \\
\frac{R[s] = (R[x])[s/x] \quad pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} s : \tau \quad pc, \Gamma[x \mapsto \tau], \mathcal{H} \vdash_{\Phi \ell} R[x] : \tau'}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[s] : \tau'} \quad \text{heap} \\
\frac{\text{dom}(\mathcal{H}) = \text{dom}(H) \quad \forall \text{loc} \in \text{dom}(H). pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H(\text{loc}) : \mathcal{H}(\text{loc})}{pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H}
\end{array}$$

Figure 10. $\lambda_{\text{seq}}^{\text{sync}}$: Syntax-Directed Monomorphic Type Rules

5. (Update of Type of Environment). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (H, e) : \tau$ and $x \notin \text{free}(H, e)$ then for any τ' , $pc, \Gamma[x \mapsto \tau'], \mathcal{H} \vdash_{\Phi \ell} (H, e) : \tau$.
6. (Redex Typing). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[e_{\text{sub}}] : \tau$ then there exists a τ_{sub} such that $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e_{\text{sub}} : \tau_{\text{sub}}$.
7. (Redex Replacement). If $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[e] : \tau$, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau_{\text{sub}}$ and $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e' : \tau_{\text{sub}}$ then $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[e'] : \tau$.

Proof. 1. (Program Counter). By induction on the respective type derivations.
2. (Subtyping). By induction on the respective type derivations.
3. (Heap Update). By the type rule *heap*.
4. (Extension of Type of Heap). By induction on the derivation of $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e : \tau$.
5. (Update of Type of Environment). By induction on the derivation of $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (H, e) : \tau$.
6. (Redex Typing). By induction on the derivation of $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[e_{\text{sub}}] : \tau$.
7. (Redex Replacement). By induction on the derivation of $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[e] : \tau$. \square

Lemma 3.6 (Value Substitution in a Program). *If $pc, \Gamma[x \mapsto \tau_x], \mathcal{H} \vdash_{\Phi \ell} p : \tau$, $pc_{\text{sub}} \leq pc$ and $pc_{\text{sub}}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau_x$ then $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p[v/x] : \tau$.*

Proof. By induction on the derivation of $pc, \Gamma[x \mapsto \tau_x], \mathcal{H} \vdash_{\Phi \ell} p : \tau$. Following are the type rules from Figure 10 applicable at the root of this derivation.

1. *var.* So for some y and τ_{sub} , $e = y$ and $\Gamma[x \mapsto \tau_x](y) = \tau_{\text{sub}}$, $\tau_{\text{sub}} \leq \tau$ and $pc \triangleleft \tau$. Now there are two possible cases depending on whether x is equal to y or not:
 - (a) $x = y$. So $\tau_{\text{sub}} = \tau_x$ and $y[v/x] = v$. By Lemma 3.5[2b,1c] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau$, that is, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e[v/x] : \tau$, the requisite result.
 - (b) $x \neq y$. So $y \in \text{dom}(\Gamma)$ and $y[v/x] = y$. Then by the type rule *var* $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} y : \tau$ holds, that is, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e[v/x] : \tau$ holds, the requisite result.
2. *int.* So for some i and ℓ' , $e = i$ and $\tau = \text{int}^{\ell'}$. Now $i[v/x] = i$; then directly by *int*, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e[v/x] : \tau$ holds, the requisite result.
3. *bool, location.* Analogous to the case 2 above.
4. *func.* So for some y, p, τ_y, τ_p and ℓ' , $e = \lambda y. p$, $\tau = (\tau_y \rightarrow \tau_p)^{\ell'}$, $pc \leq \ell'$ and $\ell', \Gamma[x \mapsto \tau_x][y \mapsto \tau_y], \mathcal{H} \vdash_{\Phi \ell} p : \tau_p$. There are two possible cases depending on whether x is equal to y or not:
 - (a) $x = y$. So $(\lambda y. p)[v/x] = \lambda y. p$ and $\ell', \Gamma[y \mapsto \tau_y], \mathcal{H} \vdash_{\Phi \ell} p : \tau_p$. Then by *func*, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e[v/x] : \tau$ holds, the requisite result.

(b) $x \neq y$. So $(\lambda y. p)[v/x] = \lambda y. (p[v/x])$, and $\ell', \Gamma[y \mapsto \tau_y][x \mapsto \tau_x], \mathcal{H} \vdash_{\Phi \ell} p : \tau_p$ holds. Now given $p[v/x]$ denotes capture-avoiding substitution, $y \notin \text{free}(v)$. Then by premise and Lemma 3.5[5] $pc_{sub}, \Gamma[y \mapsto \tau_y], \mathcal{H} \vdash_{\Phi \ell} v : \tau_x$ holds. Also $pc_{sub} \leq \ell'$. Then by induction hypothesis $\ell', \Gamma[y \mapsto \tau_y], \mathcal{H} \vdash_{\Phi \ell} p[v/x] : \tau_p$ holds. Finally by *func*, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e[v/x] : \tau$ holds, the requisite result.

5. *binop, app, ref, deref, assign*. Directly by induction.

6. *let*. Similar to case 4 above.

7. *if-sync, if, app-sync*. Directly by induction given Lemma 3.5[1a]. \square

We define the type judgement $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (p, d) : \tau$ to hold iff $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} p : \tau$, $pc \leq \ell \iff d = \infty$ and $pc \not\leq \ell \iff d \neq \infty$; note the correlation between the security level of the context pc and the down-counter d – a high context bi-implies synchronized computation of the program p , while a low context bi-implies non-synchronized computation. Also, the type judgement $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (R[s], n) : \tau$ is defined to hold iff $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} R[s] : \tau$.

Main Lemma 3.7 (Type Preservation). *If $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H$, $pc_{sub} \leq pc$, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e, d) : \tau$ and $(H, e)_d \xrightarrow{\Phi} (H', e')_{d'}$ then there exists a \mathcal{H}' such that $\mathcal{H} \subseteq \mathcal{H}'$, $pc_{sub}, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} H'$ and $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (e', d') : \tau$.*

Proof. By induction on the derivation of $(H, e)_d \xrightarrow{\Phi} (H', e')_{d'}$. Following are the semantics rules from Figure 8 applicable at the root of this derivation.

1. BINOP. So for some i_1 and i_2 , $e = i_1 \oplus i_2$, $H' = H$, for $i_1 \oplus i_2 = v$, $e' = v$ and $d' = d - 1$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds. By premise and Lemma 3.5[1a], $pc \triangleleft \tau$. Then by premise with *binop* and *int/bool*, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds as well.

2. LET. So for some x, v and p , $e = (\text{let } x = v \text{ in } p)$, $H' = H$, $e' = p[v/x]$ and $d' = d - 1$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds. By premise, *let* and Lemma 3.6 (Value Substitution in a Program) the second requisite result, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$, holds as well.

3. IF-SYNC. So for some $p, \bar{i}_k, v, b, p_1, p_2$ and n , $e = \text{if } v = b \text{ then } p_1 \text{ else } p_2$, $d = \infty$, $\Phi(p)(\bar{i}_k) = n$, $H' = H$ and $d' = n$. Also, $d = \infty$ implies $pc \leq \ell$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds.

Now to prove the second requisite result we need to consider two possible cases depending on the value of b :

(a) $b = \text{true}$. So $e' = \langle p_1 \rangle^v$. By premise, *if-sync* and *sync* $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$ holds, the second requisite result.

(b) $b = \text{false}$. Analogous to the subcase 3a above.

4. IF. So for some p, \bar{i}_k, v, b, p_1 and p_2 , $e = \text{if } v = b \text{ then } p_1 \text{ else } p_2$, $H' = H$ and $d' = d - 1$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds.

Now to prove the second requisite result we need to consider two possible cases depending on the value of b :

(a) $b = \text{true}$. So $e' = p_1$. There again two possible cases depending on whether d is ∞ or not:

i. $d = \infty$. So $pc \leq \ell$ and as per the premise of IF, $\Phi(p) = \emptyset$. Then by premise, *if* and Lemma 3.5[1a,1(d)i] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$ holds, the second requisite result.

ii. $d \neq \infty$. So $pc \not\leq \ell$. Then by premise, *if* and Lemma 3.5[1a,1(d)ii] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$ holds, the second requisite result.

(b) $b = \text{false}$. Analogous to the subcase 4a above.

5. APP-SYNC. So for some $p, \bar{i}_k, v, x, p, v'$ and n , $e = (\lambda x. p)(v')_{p, \bar{i}_k}^v$, $d = \infty$, $\Phi(p)(\bar{i}_k) = n$, $H' = H$, $e' = \langle p[v'/x] \rangle^v$ and $d' = n$. Also, $d = \infty$ implies $pc \leq \ell$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds.

By *app-sync* for some τ_x and ℓ' , $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \lambda x. p : (\tau_x \rightarrow \tau)_{\ell'}$, $\ell' \not\leq \ell$, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v' : \tau_x$ and $\ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau$. Then by *func*, $pc \leq \ell'$ and $\ell', \Gamma[x \mapsto \tau_x], \mathcal{H} \vdash_{\Phi \ell} p : \tau$. By Lemma 3.6 (Value Substitution in a Program) $\ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} p[v'/x] : \tau$. Then by *sync*, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds.

6. APP. So for some p, \bar{i}_k, v, x, p and v' , $e = (\lambda x. p)(v')_{p, \bar{i}_k}^v$, $H' = H$, $e' = p[v'/x]$ and $d' = d - 1$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds.

Now to prove the second requisite result we need to consider two possible cases depending on the value of d :

(a) $d = \infty$. So $pc \leq \ell$, and as per APP $\Phi(p) = \emptyset$. Then by *app, func*, Lemma 3.6 (Value Substitution in a Program) and Lemma 3.5[1a,1(d)i] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$ holds, the second requisite result.

(b) $d \neq \infty$. So $pc \not\leq \ell$. Then by *app, func*, Lemma 3.6 (Value Substitution in a Program) and Lemma 3.5[1a,1(d)ii] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$ holds, the second requisite result.

7. REF. So for some v and fresh loc , $e = \text{ref } v$, $H' = H \cup \{loc \mapsto v\}$, $e' = loc$ and $d' = d - 1$. By *ref* for some τ_v and ℓ' , $\tau = \tau_v \text{ref } \ell'$, $pc \leq \ell'$, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau_v$ and $\ell' \triangleleft \tau_v$. By Lemma 3.5[1b] $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau_v$; then by premise, Lemma 3.5[4] and *heap*, for $\mathcal{H}' = \mathcal{H} \cup \{loc \mapsto \tau_v\}$, $pc_{sub}, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} H'$, the first requisite result, holds, and by *location*, $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds.

8. DEREf. So for some loc and v , $e = !loc$, $H(loc) = v$, $H' = H$, $e' = v$ and $d' = d - 1$.

Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds.

By *deref* and *location* for some τ_v and ℓ' , $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} loc : \tau_v \text{ref } \ell'$, $pc \leq \ell'$, $\mathcal{H}(loc) = \tau_v$, $\ell' \triangleleft \tau_v$ and $\tau_v \leq \tau$. By *heap* $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau_v$. Then by Lemma 3.5[1b,2b] $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds.

9. ASSIGN. So for some loc and v , $e = loc := v$, $loc \in \text{dom}(H)$, $H' = H[loc \mapsto v]$, $e' = v$ and $d' = d - 1$.

By premise and *assign* and *location* for some τ_{loc} , ℓ' and τ_v , $\mathcal{H}(loc) = \tau_{loc}$, $pc \leq \ell'$, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} loc : \tau_{loc} \text{ref } \ell'$, $\ell' \triangleleft \tau_{loc}$, $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau_v$ and $\tau_v \leq \tau_{loc}, \tau$. Then by premise and Lemma 3.5[1b,2b] $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} v : \tau_{loc}$ holds.

Then by Lemma 3.5[3] $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds. And by Lemma 3.5[2b] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds.

10. SYNC. So for some p, v, n and p' , $e = \langle p \rangle^v$, $d = n$, $(H, p)_n \xrightarrow{\Phi} (H', p')_{n-1}$, $e' = \langle p' \rangle^v$ and $d' = n - 1$.

By *sync* $pc \leq \ell, \ell', \ell' \not\leq \ell$ and $\ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} p, v : \tau$. Then $pc_{sub} \leq \ell'$ and $\ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} (p, n) : \tau$. By induction hypothesis there exists a \mathcal{H}' such that $\mathcal{H} \subseteq \mathcal{H}'$, $pc_{sub}, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} H'$, the first requisite result, and $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (p', n - 1) : \tau$. Then by *sync*, $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds.

11. SYNC-PAD. So for some v, v' and $n, e = \langle v \rangle^{v'}, d = n, H' = H, e' = e$ and $d' = n - 1$. Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$ and $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$, the requisite results, hold.
12. SYNC-DONE. So for some v and $v', e = \langle v \rangle^{v'}, d = 0, H' = H, e' = v$ and $d' = \infty$.
Directly by premise $pc_{sub}, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H'$, the first requisite result, holds.
By *sync* $pc \leq \ell, \ell', \ell' \not\leq \ell$ and $\ell', \Gamma, \mathcal{H} \vdash_{\Phi \ell} v, v' : \tau$.
By Lemma 3.5[1a] $\ell' \triangleleft \tau$, implying $pc \triangleleft \tau$. Then by Lemma 3.5[1b] $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (e', d') : \tau$, the second requisite result, holds.
13. SYNC-ABORT. Similar to case 12 above.
14. CONTEXT. So for some R and $e_{sub}, e = R[e_{sub}], (H, e_{sub})_d \longrightarrow_{\Phi} (H', e'_{sub})_{d'}$ and $e' = R[e'_{sub}]$.
By Lemma 3.5[6] there exists a τ_{sub} such that there exists a τ_{sub} such that $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} e_{sub} : \tau_{sub}$. By induction hypothesis, there exists a \mathcal{H}' such that $\mathcal{H} \subseteq \mathcal{H}', pc_{sub}, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} H'$ and $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (e'_{sub}, d') : \tau_{sub}$. Then by Lemma 3.5[4,7] $pc, \Gamma, \mathcal{H}' \vdash_{\Phi \ell} (e', d') : \tau$ holds, the second requisite result. \square

Main Lemma 3.8 (Progress). *If $pc_{sub}, \emptyset, \mathcal{H} \vdash_{\Phi \ell} H, pc_{sub} \leq pc$ and $pc, \emptyset, \mathcal{H} \vdash_{\Phi \ell} (e, d) : \tau$ then either e is a value, or e is a program and $d = 0$, or there exists a H', e' and d' such that $(H, e)_d \longrightarrow_{\Phi} (H', e')_{d'}$.*

Proof. By induction on the derivation of $pc, \emptyset, \mathcal{H} \vdash_{\Phi \ell} (e, d) : \tau$. \square

Definition 3.9 (Convergence). *A configuration C given a table of sync timer functions Φ converges to a value v in n steps iff $C \longrightarrow_{\Phi}^n (H, v)_d$, for some H and d .*

Definition 3.10 (Divergence). *A configuration C given a table of sync timer functions Φ diverges iff for some $C', C \longrightarrow_{\Phi} C'$ and C' , given Φ , in turn diverges.*

Definition 3.11. *A program p given a table of sync timer functions Φ converges to value v in n steps (diverges) iff the configuration $(\emptyset, p)_{\infty}$ given Φ converges to the value v in n steps (diverges).*

Definition 3.12 (Well-Formed Program). *A program p given a table of sync timer functions Φ is said to be well-formed iff $\perp, \emptyset, \emptyset \vdash_{\Phi \ell} p : \tau$ holds, for some ℓ and τ .*

Theorem 3.13 (Type Safety). *A well-formed program either converges to a value or diverges.*

Proof. By Lemma 3.4[2] and Main Lemmas 3.7 (Type Preservation) and 3.8 (Progress). \square

4. Security Properties of λ_{seq}^{sync}

We now formally establish the security of λ_{seq}^{sync} . We show λ_{seq}^{sync} is secure with respect to external timing channels by proving a strong noninterference result [9] between the high data and the running time of a program; the term *strong noninterference* is used to indicate the robustness of the result in presence of termination channels. We prove this result by showing how executions of two programs, which differ only in high values, are strongly bisimilar. Strong bisimilarity is an isomorphism of expressions at all intermediate steps of execution; the term *strong bisimilarity* indicates the lock-step nature of such executions: they must match up step-for-step. The bisimulation relation, to be defined, requires all low values to be identical, while allowing high values to differ. This is a much tighter alignment than is used in standard noninterference

results, as the prevention of timing leaks requires a tighter correspondence.

Let $\mu : \{\overline{loc}\} \rightarrow \{\overline{loc}\}$ be an injective function from heap locations (of one run) to the heap locations (of the other run), denoting the correspondence between the supposed bisimilar locations of the respective heaps. Note the fresh heap locations generated by the REF rule at intermediate steps during bisimulation need not be identical; hence the need for the following binary relation, $\dot{=}_{\mu}$, on programs which establishes equality of low values in bisimilar programs, except for the heap locations which are required to be isomorphic as per μ . A program p_1 is μ -equal to a program p_2 , that is, $p_1 \dot{=}_{\mu} p_2$, iff for $p_1 = loc$ we have $p_2 = \mu(loc)$, for some loc ; for all other cases of p_1 the relation is homomorphic.

Definition 4.1 (μ -Equality Relation).

1. (Values). $v_1 \dot{=}_{\mu} v_2$ iff either,
 - (a) $v_1 = v_2$; or
 - (b) $v_1 = \lambda x. p_1, v_2 = \lambda x. p_2$ and $p_1 \dot{=}_{\mu} p_2$, for some x ; or
 - (c) $v_1 = loc$ and $v_2 = \mu(loc)$, for some loc .
2. (Programs). $p_1 \dot{=}_{\mu} p_2$ iff either,
 - (a) (Binary Operation). $p_1 = p'_1 \oplus p''_1, p_2 = p'_2 \oplus p''_2, p'_1 \dot{=}_{\mu} p'_2$ and $p''_1 \dot{=}_{\mu} p''_2$; or
 - (b) (Let). $p_1 = (\text{let } x = p'_1 \text{ in } p''_1), p_2 = (\text{let } x = p'_2 \text{ in } p''_2), p'_1 \dot{=}_{\mu} p'_2$ and $p''_1 \dot{=}_{\mu} p''_2$; or
 - (c) (If). $p_1 = \text{if}_{p, v_k}^{v'_1} p'_1 \text{ then } p''_1 \text{ else } p'''_1, p_2 = \text{if}_{p, v_k}^{v'_2} p'_2 \text{ then } p''_2 \text{ else } p'''_2, p'_1 \dot{=}_{\mu} p'_2, p''_1 \dot{=}_{\mu} p''_2, p'''_1 \dot{=}_{\mu} p'''_2, v'_1 \dot{=}_{\mu} v'_2$ and $\overline{v_k} \dot{=}_{\mu} v'_k$; or
 - (d) (App). $p_1 = p'_1 (p''_1)_{p, v_k}^{v'_1}, p_2 = p'_2 (p''_2)_{p, v_k}^{v'_2}, p'_1 \dot{=}_{\mu} p'_2, p''_1 \dot{=}_{\mu} p''_2, v'_1 \dot{=}_{\mu} v'_2$ and $\overline{v_k} \dot{=}_{\mu} v'_k$; or
 - (e) (Ref). $p_1 = \text{ref } p'_1, p_2 = \text{ref } p'_2$ and $p'_1 \dot{=}_{\mu} p'_2$; or
 - (f) (Deref). $p_1 = !p'_1, p_2 = !p'_2$ and $p'_1 \dot{=}_{\mu} p'_2$; or
 - (g) (Assign). $p_1 = (p'_1 := p''_1), p_2 = (p'_2 := p''_2), p'_1 \dot{=}_{\mu} p'_2$ and $p''_1 \dot{=}_{\mu} p''_2$.

The notation $\mu \circ \mu'$ (read as “ μ composed with μ' ”) denotes function composition.

Lemma 4.2 (Properties of μ -Equality Relation).

1. (Symmetry). If $p_1 \dot{=}_{\mu} p_2$ then $p_2 \dot{=}_{\mu^{-1}} p_1$.
2. (Reflexivity). For all p and $\mu, p \dot{=}_{\mu} p$.
3. (Transitivity). If $p_1 \dot{=}_{\mu} p_2$ and $p_2 \dot{=}_{\mu'} p_3$ then $p_1 \dot{=}_{\mu' \circ \mu} p_3$.
4. (Structure of Programs). If $p_1 \dot{=}_{\mu} p_2$ then p_1 and p_2 have the same outermost structure.
5. (Substitution). If $p_1 \dot{=}_{\mu} p_2$ and $v_1 \dot{=}_{\mu} v_2$ then $p_1[v_1/x] \dot{=}_{\mu} p_2[v_2/x]$.
6. (Extension of μ). If $p_1 \dot{=}_{\mu} p_2$ and $\mu \subseteq \mu'$ then $p_1 \dot{=}_{\mu'} p_2$ holds.

Proof. 1. (Symmetry). By induction on the derivation of $p_1 \dot{=}_{\mu} p_2$, noting μ is an injection.

2. (Reflexivity). By induction on the structure of p , given Definition 4.1, in particular Condition 1a.

3. (Transitivity). By induction on the pair of the derivations of $p_1 \dot{=}_{\mu} p_2$ and $p_2 \dot{=}_{\mu'} p_3$, given Definition 4.1, in particular Condition 1c.

4. (Structure of Programs). By induction on the derivation of $p_1 \dot{=}_{\mu} p_2$.

5. (Substitution). By induction on the derivation of $p_1 \dot{=}_{\mu} p_2$.

6. (Extension of μ). By induction on the derivation of $p_1 \dot{=}_{\mu} p_2$. \square

We now define the bisimulation relation. We write the type judgement $\vdash pc, \Gamma, \mathcal{H} \vdash_{\Phi\ell}^{low} p : \tau$ to denote $pc, \Gamma, \mathcal{H} \vdash_{\Phi\ell} p : \tau$ and $secllevel(\tau) \leq \ell$, and $\vdash pc, \Gamma, \mathcal{H} \vdash_{\Phi\ell}^{high} p : \tau$ for $pc, \Gamma, \mathcal{H} \vdash_{\Phi\ell} p : \tau$ and $secllevel(\tau) \not\leq \ell$.

Definition 4.3 (λ_{seq}^{sync} : Bisimulation Relation).

1. (Programs). $(\mathcal{H}_1, p_1)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, p_2)_{\infty}$ iff there exists a $p'_1, p'_2, \bar{x}_k, \bar{v}_k, \bar{v}'_k$ and $\bar{\tau}_k$ such that,
 - (a) $p_1 = p'_1[\bar{v}_k/\bar{x}_k]$ and $p_2 = p'_2[\bar{v}'_k/\bar{x}_k]$; and
 - (b) $p'_1 \dot{=}_{\mu} p'_2$; and
 - (c) $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell}^{high} \bar{v}_k : \bar{\tau}_k$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell}^{high} \bar{v}'_k : \bar{\tau}_k$; and
 - (d) $\Gamma' = \Gamma \cup \{\bar{x}_k \mapsto \bar{\tau}_k\}$; and
 - (e) $\perp, \Gamma', \mathcal{H}_1 \vdash_{\Phi\ell} p'_1 : \tau$ and $\perp, \Gamma', \mathcal{H}_2 \vdash_{\Phi\ell} p'_2 : \tau$.
2. (Expressions with Synchronization Constructs). $(\mathcal{H}_1, R_1[s_1])_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, R_2[s_2])_n$ iff there exists a x and τ' such that,
 - (a) $R_1[s_1] = (R_1[x])[s_1/x]$ and $R_2[s_2] = (R_2[x])[s_2/x]$; and
 - (b) $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell} s_1 : \tau'$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} s_2 : \tau'$; and
 - (c) $(\mathcal{H}_1, R_1[x])_{\infty} \sim_{\Phi\ell\mu}^{\Gamma \cup \{x \mapsto \tau'\} \tau} (\mathcal{H}_2, R_2[x])_{\infty}$.
3. (Heaps). $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2)$ iff
 - (a) $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell} H_1$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} H_2$; and
 - (b) For all $loc \in dom(\mu)$, there exists a τ such that, $\mathcal{H}_1(loc) = \tau = \mathcal{H}_2(\mu(loc))$ and $(\mathcal{H}_1, H_1(loc))_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2(\mu(loc)))_{\infty}$.
4. (Heaps, Expressions). $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$ iff $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$.

The values \bar{v}_k and \bar{v}'_k in case 1 above denote the high values inside programs p_1 and p_2 , respectively, which may be different; hence they are elided in case 1a before establishing μ -equality of the residual programs in case 1b. Analogously, as the code in a synchronization construct is high, the synchronization code is removed in case 2a, before relating the enclosing programs for bisimilarity in case 2c. Furthermore, the down-counters of bisimilar configurations are required to be identical.

Lemma 4.4 (λ_{seq}^{sync} : Properties of Bisimulation Relation).

1. (Symmetry). If $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$ then $(\mathcal{H}_2, H_2, e_2)_{d_2} \sim_{\Phi\ell\mu^{-1}}^{\Gamma\tau} (\mathcal{H}_1, H_1, e_1)_{d_1}$.
2. (Reflexivity).
 - (a) (Programs). If $\perp, \Gamma, \mathcal{H} \vdash_{\Phi\ell} p : \tau$ then for all μ , $(\mathcal{H}, p)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}, p)_{\infty}$.
 - (b) (Expressions with Synchronization Constructs). If $\perp, \Gamma, \mathcal{H} \vdash_{\Phi\ell} R[s] : \tau$ then for all n and μ , $(\mathcal{H}, R[s])_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}, R[s])_n$.
 - (c) (Heaps). If $\perp, \Gamma, \mathcal{H} \vdash_{\Phi\ell} H$ and $dom(\mathcal{H}) = \{\overline{loc_k}\}$ then for all μ , such that $\mu \subseteq \{\overline{loc_k} \mapsto \overline{loc_k}\}$, $(\mathcal{H}, H) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}, H)$ holds.
 - (d) (Heaps, Expressions). If $(\mathcal{H}, e)_d \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}, e)_d$ and $(\mathcal{H}, H) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}, H)$ then $(\mathcal{H}, H, e)_d \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}, H, e)_d$.
3. (Transitivity). If $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$ and $(\mathcal{H}_2, H_2, e_2)_{d_2} \sim_{\Phi\ell\mu'}^{\Gamma\tau} (\mathcal{H}_3, H_3, e_3)_{d_3}$ then $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell(\mu' \circ \mu)}^{\Gamma\tau} (\mathcal{H}_3, H_3, e_3)_{d_3}$.
4. (Isomorphism). If $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$ then e_1 and e_2 are either both programs, or both have a sync construct.
5. (Programs). If $(\mathcal{H}_1, p_1)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, p_2)_{\infty}$ then $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell} p_1 : \tau$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} p_2 : \tau$.

6. (Expressions with Synchronization Constructs). If $(\mathcal{H}_1, R_1[s_1])_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, R_2[s_2])_n$ then $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell} R_1[s_1] : \tau$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} R_2[s_2] : \tau$.
7. (Synchronization Constructs). If $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell} s_1 : \tau$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} s_2 : \tau$ then for all n and μ , $(\mathcal{H}_1, s_1)_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, s_2)_n$.
8. (Heaps). If $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2)$ then $dom(\mathcal{H}_1) = dom(\mathcal{H}_2)$, $dom(\mu) \subseteq dom(H_1)$ and $range(\mu) \subseteq dom(H_2)$.
9. (Extension of μ). If $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$ and $\mu \subseteq \mu'$ then $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu'}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$.
10. (Extension of Heaps).
 - (a) (Expressions). If $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$, $\mathcal{H}_1 \subseteq \mathcal{H}'_1$ and $\mathcal{H}_2 \subseteq \mathcal{H}'_2$ then $(\mathcal{H}'_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}'_2, e_2)_{d_2}$.
 - (b) (Heaps). If $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2)$, $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $H_1 \subseteq H'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $H_2 \subseteq H'_2$, $\perp, \Gamma, \mathcal{H}'_1 \vdash_{\Phi\ell} H'_1$ and $\perp, \Gamma, \mathcal{H}'_2 \vdash_{\Phi\ell} H'_2$ then $(\mathcal{H}'_1, H'_1) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}'_2, H'_2)$.
11. (Values).
 - (a) (Structural Congruence). If $(\mathcal{H}_1, v_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and e_2 is a value.
 - (b) (“High” Values). If $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell}^{high} v_1 : \tau$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell}^{high} v_2 : \tau$ then for all μ , $(\mathcal{H}_1, v_1)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, v_2)_{\infty}$.
 - (c) (“Low” Int/Bool Type). If $(\mathcal{H}_1, \hat{v})_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{\infty}$ and $secllevel(\hat{\tau}) \leq \ell$ then $e_2 = \hat{v}$.
12. (Heap Update with “High” Value). If $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2)$ and $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell}^{high} v : \mathcal{H}_1(loc)$ then $(\mathcal{H}_1, H_1[loc \mapsto v]) \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, H_2)$.
13. (Binary Operation).
 - (a) (Structural Congruence). If $(\mathcal{H}_1, i_1 \oplus i'_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some i_2 and i'_2 , $e_2 = i_2 \oplus i'_2$.
 - (b) (BINOP). If $(\mathcal{H}_1, i_1 \oplus i'_1)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, i_2 \oplus i'_2)_{\infty}$, $i_1 \oplus i'_1 = v_1$ and $i_2 \oplus i'_2 = v_2$ then $(\mathcal{H}_1, v_1)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, v_2)_{\infty}$.
14. (Let).
 - (a) (Structural Congruence). If $(\mathcal{H}_1, \text{let } x = v_1 \text{ in } p_1)_{d_1} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some v_2 and p_2 , $e_2 = (\text{let } x = v_2 \text{ in } p_2)$.
 - (b) (LET). If $(\mathcal{H}_1, \text{let } x = v_1 \text{ in } p_1)_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, \text{let } x = v_2 \text{ in } p_2)_{\infty}$ then $(\mathcal{H}_1, p_1[v_1/x])_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, p_2[v_2/x])_{\infty}$.
15. (If).
 - (a) (Structural Congruence). If $(\mathcal{H}_1, \text{if}_{p, i_k}^{v_1} b_1 \text{ then } p_{1T} \text{ else } p_{1F})_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some v_2, b_2, p_{2T} and p_{2F} , $e_2 = \text{if}_{p, i_k}^{v_2} b_2 \text{ then } p_{2T} \text{ else } p_{2F}$.
 - (b) (IF-SYNC). If $(\mathcal{H}_1, \text{if}_{p, i_k}^{v_1} b_1 \text{ then } p_{1T} \text{ else } p_{1F})_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, \text{if}_{p, i_k}^{v_2} b_2 \text{ then } p_{2T} \text{ else } p_{2F})_{\infty}$ and $\Phi(p)(i_k) = n$ then $(\mathcal{H}_1, \langle p_{1T} \rangle^{v_1})_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, \langle p_{2T} \rangle^{v_2})_n$, $(\mathcal{H}_1, \langle p_{1F} \rangle^{v_1})_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, \langle p_{2F} \rangle^{v_2})_n$, $(\mathcal{H}_1, \langle p_{1T} \rangle^{v_1})_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, \langle p_{2T} \rangle^{v_2})_n$, $(\mathcal{H}_1, \langle p_{1F} \rangle^{v_1})_n \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, \langle p_{2F} \rangle^{v_2})_n$, $(\mathcal{H}_1, p_{1T})_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, p_{2T})_{\infty}$ and $(\mathcal{H}_1, p_{1F})_{\infty} \sim_{\Phi\ell\mu}^{\Gamma\tau} (\mathcal{H}_2, p_{2F})_{\infty}$.

- (c) (IF). If $(\mathcal{H}_1, \text{if}_{p, \bar{i}_k}^{v_1} b_1 \text{ then } p_{1_T} \text{ else } p_{1_F})_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, \text{if}_{p, \bar{i}_k}^{v_2} b_2 \text{ then } p_{2_T} \text{ else } p_{2_F})_\infty$ and $\Phi(p) = \emptyset$
then $b_1 = b_2$, $(\mathcal{H}_1, p_{1_T})_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, p_{2_T})_\infty$ and
 $(\mathcal{H}_1, p_{1_F})_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, p_{2_F})_\infty$.
16. (App).
(a) (Structural Congruence). If $(\mathcal{H}_1, (\lambda x_1. p_1) (v_{1_{arg}})_{p, \bar{i}_k}^{v_1})_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some $x_2, p_2, v_{2_{arg}}$ and $v_2, e_2 = (\lambda x_2. p_2) (v_{2_{arg}})_{p, \bar{i}_k}^{v_2}$.
(b) (APP-SYNC). If $(\mathcal{H}_1, (\lambda x_1. p_1) (v_{1_{arg}})_{p, \bar{i}_k}^{v_1})_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, (\lambda x_2. p_2) (v_{2_{arg}})_{p, \bar{i}_k}^{v_2})_\infty$ and $\Phi(p)(\bar{i}_k) = n$ then
 $(\mathcal{H}_1, \langle p_1[v_{1_{arg}}/x_1] \rangle_{v_1}^n) \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, \langle p_2[v_{2_{arg}}/x_2] \rangle_{v_2}^n)$.
(c) (APP). If $(\mathcal{H}_1, (\lambda x_1. p_1) (v_{1_{arg}})_{p, \bar{i}_k}^{v_1})_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, (\lambda x_2. p_2) (v_{2_{arg}})_{p, \bar{i}_k}^{v_2})_\infty$ and $\Phi(p) = \emptyset$ then $x_1 = x_2$
and $(\mathcal{H}_1, p_1[v_{1_{arg}}/x_1])_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, p_2[v_{2_{arg}}/x_2])_\infty$.
17. (Ref).
(a) (Structural Congruence). If $(\mathcal{H}_1, \text{ref } v_1)_{d_1} \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some $v_2, e_2 = \text{ref } v_2$.
(b) (REF). If $(\mathcal{H}_1, H_1, \text{ref } v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2, \text{ref } v_2)_\infty$
then, for some τ' and ℓ' , $\tau = \tau' \text{ref } \ell'$ and for all loc_1
and loc_2 such that $loc_1 \notin \text{dom}(H_1)$ and $loc_2 \notin \text{dom}(H_2)$,
 $(\mathcal{H}'_1, H'_1, loc_1)_\infty \sim_{\Phi \ell \mu'}^{\Gamma \tau'}$ $(\mathcal{H}'_2, H'_2, loc_2)_\infty$ holds, such that
 $H'_1 = H_1 \cup \{loc_1 \mapsto v_1\}$, $H'_2 = H_2 \cup \{loc_2 \mapsto v_2\}$,
 $\mathcal{H}'_1 = \mathcal{H}_1 \cup \{loc_1 \mapsto \tau'\}$, $\mathcal{H}'_2 = \mathcal{H}_2 \cup \{loc_2 \mapsto \tau'\}$, and
 $\mu' = \mu \cup \{loc_1 \mapsto loc_2\}$.
18. (Deref).
(a) (Structural Congruence). If $(\mathcal{H}_1, H_1, !loc_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau}$
 $(\mathcal{H}_2, H_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some $loc_2 \in \text{dom}(H_2)$,
 $e_2 = !loc_2$.
(b) (DEREF). If $(\mathcal{H}_1, H_1, !loc_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2, !loc_2)_\infty$
then $(\mathcal{H}_1, H_1, H_1(loc_1))_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2, H_2(loc_2))_\infty$.
19. (Assign).
(a) (Structural Congruence). If $(\mathcal{H}_1, H_1, loc_1 := v_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau}$
 $(\mathcal{H}_2, H_2, e_2)_{d_2}$ then $d_1 = d_2 = \infty$ and for some $loc_2 \in \text{dom}(H_2)$
and $v_2, e_2 = (loc_2 := v_2)$.
(b) (ASSIGN). If $(\mathcal{H}_1, H_1, loc_1 := v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, H_2, loc_2 := v_2)_\infty$ then $(\mathcal{H}_1, H_1[loc_1 \mapsto v_1], v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, H_2[loc_2 \mapsto v_2], v_2)_\infty$.
20. (Synchronization Construct).
(a) (Structural Congruence). If $(\mathcal{H}_1, \langle p_1 \rangle_{v_1})_{d_1} \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, e_2)_{d_2}$ then for some n, p_2 and $v_2, d_1 = d_2 = n$,
 $e_2 = \langle p_2 \rangle_{v_2}$ and $(\mathcal{H}_1, v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, v_2)_\infty$.
(b) (Sub-Structural Bisimilarities). If $(\mathcal{H}_1, \langle v_1 \rangle_{v_1})_{n_1} \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, \langle v_2 \rangle_{v_2})_{n_2}$ then $(\mathcal{H}_1, v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, v_2)_\infty$,
 $(\mathcal{H}_1, v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, v_2')_\infty$, $(\mathcal{H}_1, v_1')_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, v_2)_\infty$ and $(\mathcal{H}_1, v_1')_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, v_2')_\infty$.
21. (Context).
(a) (Structural Congruence). If $(\mathcal{H}_1, R_1[e_{1_{sub}}])_{d_1} \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, e_2)_{d_2}$ then for some $R_2, e_{2_{sub}}$ and $\tau_{sub}, e_2 = R_2[e_{2_{sub}}]$
and $(\mathcal{H}_1, e_{1_{sub}})_{d_1} \sim_{\Phi \ell \mu}^{\Gamma \tau_{sub}} (\mathcal{H}_2, e_{2_{sub}})_{d_2}$.
(b) (CONTEXT). If $(\mathcal{H}_1, H_1, R_1[e_{1_{sub}}])_{d_1} \sim_{\Phi \ell \mu}^{\Gamma \tau}$
 $(\mathcal{H}_2, H_2, R_2[e_{2_{sub}}])_{d_2}$, $(\mathcal{H}_1, H_1, e_{1_{sub}})_{d_1} \sim_{\Phi \ell \mu}^{\Gamma \tau_{sub}}$
- $(\mathcal{H}_2, H_2, e_{2_{sub}})_{d_2}$, $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$
and $(\mathcal{H}'_1, H'_1, e'_{1_{sub}})_{d'_1} \sim_{\Phi \ell \mu'}^{\Gamma \tau_{sub}} (\mathcal{H}'_2, H'_2, e'_{2_{sub}})_{d'_2}$ then
 $(\mathcal{H}'_1, H'_1, R_1[e'_{1_{sub}}])_{d'_1} \sim_{\Phi \ell \mu'}^{\Gamma \tau} (\mathcal{H}'_2, H'_2, R_2[e'_{2_{sub}}])_{d'_2}$.
- Proof.* 1. (Symmetry). By Definition 4.3 given Lemma 4.2[1].
2. (Reflexivity).
(a) (Programs). By Definition 4.3[1] and Lemma 4.2[2].
(b) (Expressions with Synchronization Constructs). By the type rule *sync-in-context*, Lemma 4.4[2a] and Definition 4.3[2].
(c) (Heaps). By Definition 4.3[3], the type rule *heap* and Lemma 4.4[2a].
(d) (Heaps, Expressions). By Definition 4.3[4].
3. (Transitivity). By Definition 4.3 and Lemma 4.2[3].
4. (Isomorphism). By Definition 4.3[1,2].
5. (Programs). By Definition 4.3[1c,1e] and Lemma 3.6 (Value Substitution in a Program).
6. (Expressions with Synchronization Constructs). By Definition 4.3[2], Lemma 4.4[5], and the type rule *sync-in-context*.
7. (Synchronization Constructs). By Definition 4.3[2].
8. (Heaps). By Definition 4.3[3] and the type rule *heap*.
9. (Extension of μ). By Definition 4.3[1,2] and Lemma 4.2[6].
10. (Extension of Heaps).
(a) (Expressions). By Definition 4.3[1,2] and Lemma 3.5[4].
(b) (Heaps). By Definition 4.3[3] and Lemma 4.4[10a].
11. (Values).
(a) (Structural Congruence). By Definition 4.3[1], in particular Condition 1b, and Lemma 4.2[4].
(b) (“High” Values). By Definition 4.3[1], in particular Condition 1c.
(c) (“Low” Int/Bool Type). By Definition 4.3[1] and Definition 4.1[1a].
12. (Heap Update with “High” Value). The premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2)$, given Definition 4.3[3a], implies $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} H_1$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi \ell} H_2$. We know $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} v : \mathcal{H}_1(loc)$; then by Lemma 3.5[3] $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} H_1[loc \mapsto v]$ holds. Now there are two possible cases:
(a) $loc \notin \text{dom}(\mu)$. By Definition 4.3[3], given the premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2)$ and knowing $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} H_1[loc \mapsto v]$, we get $(\mathcal{H}_1, H_1[loc \mapsto v]) \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2)$, the requisite result.
(b) $loc \in \text{dom}(\mu)$. By premise we know $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2)$ and $\text{selevel}(\mathcal{H}_1(loc)) \not\leq \ell$. Let $H_1(loc) = v_1$ and $H_2(\mu(loc)) = v_2$. Then as per Definition 4.3[3b], for $\mathcal{H}_1(loc) = \tau$, $(\mathcal{H}_1, v_1)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, v_2)_\infty$ holds; then as per Lemma 4.4[5] $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi \ell}^{high} v_2 : \tau$. Now knowing $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell}^{high} v : \tau$, by Lemma 4.4[11b] $(\mathcal{H}_1, v)_\infty \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, v_2)_\infty$ holds. Then by Definition 4.3[3] we get $(\mathcal{H}_1, H_1[loc \mapsto v]) \sim_{\Phi \ell \mu}^{\Gamma \tau} (\mathcal{H}_2, H_2)$, the requisite result.
13. (Binary Operation).
(a) (Structural Congruence). By Definition 4.3[1], in particular Condition 1b, and Lemma 4.2[4].
(b) (BINOP). By Definition 4.3[1].
14. (Let).
(a) (Structural Congruence). By Definition 4.3[1], in particular Condition 1b, and Lemma 4.2[4].
(b) (LET). By Definition 4.3[1], in particular Condition 1b, Definition 4.1[2b], Lemma 3.6 (Value Substitution in a Program) and Lemma 4.2[5].

15. (If).
 - (a) (*Structural Congruence*). By Definition 4.3[1], in particular Condition 1b, and Lemma 4.2[4].
 - (b) (IF-SYNC). By Definition 4.3[1,2], the type rules *if-sync* and *sync* and Definition 4.1[2c].
 - (c) (IF). By Definition 4.3[1], the type rule *if*, and Definition 4.1[2c,1a].
16. (App).
 - (a) (*Structural Congruence*). By Definition 4.3[1], in particular Condition 1b, and Lemma 4.2[4].
 - (b) (APP-SYNC). By Definition 4.3[1,2], the type rules *app-sync*, *func* and *sync*, Definition 4.1[2d] and Lemma 3.6 (Value Substitution in a Program).
 - (c) (APP). By Definition 4.3[1], the type rule *app*, Definition 4.1[2d,1b], Lemma 4.2[5] and Lemma 3.6 (Value Substitution in a Program).
17. (Ref).
 - (a) (*Structural Congruence*). By Definition 4.3[1], in particular Condition 1b, and Lemma 4.2[4].
 - (b) (REF). By Definition 4.3[1,3] and the type rules *ref* and *heap*.
18. (Deref).
 - (a) (*Structural Congruence*). By Definition 4.3[4,1,3], Lemma 4.2[4] and the premise $dom(\mathcal{H}) = dom(H)$ of the type rule *heap*.
 - (b) (DEREF). By Definition 4.3[1,3] and the type rules *deref* and *heap*.
19. (Assign).
 - (a) (*Structural Congruence*). By Definition 4.3[4,1,3], Lemma 4.2[4] and the premise $dom(\mathcal{H}) = dom(H)$ of the type rule *heap*.
 - (b) (ASSIGN). By Definition 4.3[1,3], and the type rules *assign* and *heap*.
20. (Synchronization Construct).
 - (a) (*Structural Congruence*). By Definition 4.3[2], in particular Condition 2b, the type rule *sync*, Lemma 3.5[1a] and Lemma 4.4[11b].
 - (b) (*Sub-Structural Bisimilarities*). By Definition 4.3[2], in particular Condition 2b, the type rule *sync* and Lemma 4.4[11b].
21. (Context).
 - (a) (*Structural Congruence*). By induction on the structure of R_1 given Lemma 4.2[4] and Definition 4.3[1,2].
 - (b) (CONTEXT). By induction on the structures of R_1 and R_2 , given Lemma 4.4[9,10a] and Definition 4.3[1,2].

□

The following lemma states that heap bisimilarity is preserved in computations under high contexts; ℓ' in the following lemma statement denotes such a high context.

Lemma 4.5 (Preservation of Heap Bisimilarity under a High Context, 1-step). *If $(H_1, p_1)_{n_1} \longrightarrow_{\Phi} (H'_1, p'_1)_{n'_1}$, $\ell' \not\leq \ell$, $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} p_1 : \tau_1$ and $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$ then there exists a \mathcal{H}'_1 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$ and $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$.*

Proof. By induction on the derivation of $(H_1, p_1)_{n_1} \longrightarrow_{\Phi} (H'_1, p'_1)_{n'_1}$. Following are the possible syntax-directed semantics rules from Figure 8 that are applicable at the root of this derivation.

1. BINOP, LET, IF, APP and Deref. These do not modify the heap, that is, $H'_1 = H_1$ for in the case of all these rules.

By premise and Definition 4.3[3a] $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} H_1$ and $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi \ell} H_2$ hold; also $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} (p_1, n_1) : \tau_1$ holds. Then by Main Lemma 3.7 (Type Preservation) there exists a \mathcal{H}'_1 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\perp, \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} H'_1$ and $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} (p'_1, n'_1) : \tau_1$, that is, $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$, the first requisite result.

We know $H_1 = H'_1$. Then by premise and Lemma 4.4[10b] we get $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$, the second requisite result.

2. REF. So for some $v_1, p_1 = \text{ref } v_1$ and for some fresh loc_1 , $(H_1, p_1)_{n_1} \longrightarrow_{\Phi} (H'_1, p'_1)_{n'_1}$ such that $H'_1 = H_1 \cup \{loc_1 \mapsto v_1\}$, $p'_1 = loc_1$ and $n'_1 = n_1 - 1$.

As in case 1 the first requisite result, that is, there exists a \mathcal{H}'_1 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$ and $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$, and $\perp, \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} H'_1$ hold.

We know $H_1 \subseteq H'_1$. Then by premise and Lemma 4.4[10b] we get $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$, the second requisite result.

3. ASSIGN. So for some loc_1 and $v_1, p_1 = (loc_1 := v_1)$, $loc_1 \in dom(H_1)$ and $(H_1, p_1)_{n_1} \longrightarrow_{\Phi} (H'_1, p'_1)_{n'_1}$ such that $H'_1 = H_1[loc_1 \mapsto v_1]$, $p'_1 = v_1$ and $n'_1 = n_1 - 1$.

As in case 1 the first requisite result, that is, there exists a \mathcal{H}'_1 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$ and $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$, and $\perp, \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} H'_1$ hold.

By assumption, premise, Lemma 3.5[1a] and the type rules *assign* and *location* for some $\tau_{1_{sub}}, \ell' \triangleleft \mathcal{H}_1(loc_1)$, $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} v_1 : \tau_{1_{sub}}$ and $\tau_{1_{sub}} \leq \mathcal{H}_1(loc_1), \tau_1$. Then by Lemma 3.5[2b,1a,1b] $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell}^{high} v_1 : \mathcal{H}_1(loc_1)$ holds. By Lemma 4.4[12] we get $(\mathcal{H}_1, H_1[loc_1 \mapsto v_1]) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$, that is, $(\mathcal{H}_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$.

We know $H_1 \subseteq H'_1$ and $\perp, \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} H'_1$. Then by Lemma 4.4[10b] we get $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$, the second requisite result.

4. CONTEXT. So for some $R_1, p_{1_{sub}}$ and $p'_{1_{sub}}, p_1 = R_1[p_{1_{sub}}]$, and for some $p'_{1_{sub}}, (H_1, p_{1_{sub}})_{n_1} \longrightarrow_{\Phi} (H'_1, p'_{1_{sub}})_{n'_1}$ and $(H_1, p_1)_{n_1} \longrightarrow_{\Phi} (H'_1, p'_1)_{n'_1}$ such that $p'_1 = R_1[p'_{1_{sub}}]$. Further by premise we know $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} p_1 : \tau_1$; then by Lemma 3.5[6] there exists a $\tau_{1_{sub}}$ such that $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} p_{1_{sub}} : \tau_{1_{sub}}$.

By induction hypothesis there exists a \mathcal{H}'_1 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_{1_{sub}} : \tau_{1_{sub}}$ and $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$, the second requisite result.

Now we know $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} R_1[p_{1_{sub}}] : \tau_1$, $\ell', \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} p_{1_{sub}} : \tau_{1_{sub}}$, $\mathcal{H}_1 \subseteq \mathcal{H}'_1$ and $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_{1_{sub}} : \tau_{1_{sub}}$; then by Lemma 3.5[4,7] $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} R_1[p'_{1_{sub}}] : \tau_1$, that is, $\ell', \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$ holds, the first requisite result.

□

Lemma 4.6 (Preservation of Heap Bisimilarity under High Context, n -steps). *If $(H_1, p_1)_{n_1} \xrightarrow{n}_{\Phi} (H'_1, p'_1)_{n'_1}$, $(H_2, p_2)_{n_2} \xrightarrow{n}_{\Phi} (H'_2, p'_2)_{n'_2}$, $\ell'_1, \ell'_2 \not\leq \ell$, $\ell'_1, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} p_1 : \tau_1$, $\ell'_2, \Gamma, \mathcal{H}_2 \vdash_{\Phi \ell} p_2 : \tau_2$ and $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$, then there exists a \mathcal{H}'_1 and \mathcal{H}'_2 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\ell'_1, \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$, $\ell'_2, \Gamma, \mathcal{H}'_2 \vdash_{\Phi \ell} p'_2 : \tau_2$ and $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}'_2, H'_2)$.*

Proof. By premise we know $(H_1, p_1)_{n_1} \xrightarrow{n}_{\Phi} (H'_1, p'_1)_{n'_1}$, $\ell'_1 \not\leq \ell$, $\ell'_1, \Gamma, \mathcal{H}_1 \vdash_{\Phi \ell} p_1 : \tau_1$ and $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$; by induction on the length of the above computation given Lemma 4.5 there exists a \mathcal{H}'_1 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\ell'_1, \Gamma, \mathcal{H}'_1 \vdash_{\Phi \ell} p'_1 : \tau_1$, the first requisite result, and $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\Gamma} (\mathcal{H}_2, H_2)$. By Lemma 4.4[1] for $\mu' = \mu^{-1}$, $(\mathcal{H}_2, H_2) \sim_{\Phi \ell \mu'}^{\Gamma} (\mathcal{H}'_1, H'_1)$ holds.

Further by premise we know $(H_2, p_2)_{n_2} \xrightarrow{\Gamma} (H'_2, p'_2)_{n'_2}$, $\ell'_2 \not\leq \ell$ and $\ell'_2, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} p_2 : \tau_2$, and from above we know $(\mathcal{H}_2, H_2) \sim_{\Phi\ell\mu'}^{\Gamma} (\mathcal{H}'_1, H'_1)$; then by induction on the length of the above computation given Lemma 4.5 there exists a \mathcal{H}'_2 such that $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\ell'_2, \Gamma, \mathcal{H}'_2 \vdash_{\Phi\ell} p'_2 : \tau_2$, the second requisite result, and $(\mathcal{H}'_2, H'_2) \sim_{\Phi\ell\mu'}^{\Gamma} (\mathcal{H}'_1, H'_1)$ hold. By Lemma 4.4[1] $(\mathcal{H}'_1, H'_1) \sim_{\Phi\ell\mu}^{\Gamma} (\mathcal{H}'_2, H'_2)$ holds, the third requisite result. \square

As discussed above, only high subcomputations under high contexts are ever synchronized in our system; hence the following corollary to the above lemma states the preservation of heap bisimilarity in synchronized computations.

Corollary 4.7 (Preservation of Heap Bisimilarity in Synchronized Computations). *If $(H_1, s_1)_{n_1} \xrightarrow{\Phi} (H'_1, s'_1)_{n'_1}$, $(H_2, s_2)_{n_2} \xrightarrow{\Phi} (H'_2, s'_2)_{n'_2}$, $\perp, \Gamma, \mathcal{H}_1 \vdash_{\Phi\ell} s_1 : \tau_1$, $\perp, \Gamma, \mathcal{H}_2 \vdash_{\Phi\ell} s_2 : \tau_2$ and $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\Gamma} (\mathcal{H}_2, H_2)$, then there exists a \mathcal{H}'_1 and \mathcal{H}'_2 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\perp, \Gamma, \mathcal{H}'_1 \vdash_{\Phi\ell} s'_1 : \tau_1$, $\perp, \Gamma, \mathcal{H}'_2 \vdash_{\Phi\ell} s'_2 : \tau_2$ and $(\mathcal{H}'_1, H'_1) \sim_{\Phi\ell\mu}^{\Gamma} (\mathcal{H}'_2, H'_2)$.*

Proof. By *sync*, *SYNC* and Lemma 4.6. \square

The following single-step strong bisimulation lemma states: given a pair of bisimilar configurations, if the first configuration takes a step, then the second configuration also takes a step, and bisimilarity is preserved in the resulting configurations.

Lemma 4.8 (λ_{seq}^{sync} : Strong Bisimulation, 1-step). *If $(H_1, e_1)_{d_1} \xrightarrow{\Phi} (H'_1, e'_1)_{d'_1}$ and $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau}$ $(\mathcal{H}_2, H_2, e_2)_{d_2}$ then for some H_2, e'_2 and d'_2 , $(H_2, e_2)_{d_2} \xrightarrow{\Phi} (H'_2, e'_2)_{d'_2}$ holds, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu'}^{\emptyset\tau}$ $(\mathcal{H}'_2, H'_2, e'_2)_{d'_2}$.*

Proof. By induction on the derivation of $(H_1, e_1)_{d_1} \xrightarrow{\Phi} (H'_1, e'_1)_{d'_1}$. Following are the possible syntax-directed semantics rules from Figure 8 that are applicable at the root of this derivation.

1. BINOP. So for some i_1, i'_1 and v'_1 , $e_1 = i_1 \oplus i'_1$, $i_1 \oplus i'_1 = v'_1$ and $(H_1, e_1)_{d_1} \xrightarrow{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1$, $e'_1 = v'_1$ and $d'_1 = d_1 - 1$.
By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[13a] $d_1 = d_2 = \infty$ and for some i_2 and i'_2 , $e_2 = i_2 \oplus i'_2$. By BINOP and letting $i_2 \oplus i'_2 = v'_2$, we get $(H_2, e_2)_{d_2} \xrightarrow{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2$, $e'_2 = v'_2$ and $d'_2 = d_2 - 1 = \infty$. Then by Lemma 4.4[13b] $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$, the second requisite result.
2. LET. So for some x, v_1 and p_1 , $e_1 = (\text{let } x = v_1 \text{ in } p_1)$ and $(H_1, e_1)_{d_1} \xrightarrow{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1$, $e'_1 = p_1[v_1/x]$ and $d'_1 = d_1 - 1$.
By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[14a] $d_1 = d_2 = \infty$ and for some v_2 and p_2 , $e_2 = (\text{let } x = v_2 \text{ in } p_2)$. By LET $(H_2, e_2)_{d_2} \xrightarrow{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2$, $e'_2 = e_2 \text{next}[v_2/x]$, and

$d'_2 = d_2 - 1 = \infty$. Then by Lemma 4.4[14b] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

3. IF-SYNC. So for some $p, \bar{i}_k, v_1, b_1, p_{1_T}$ and p_{1_F} , $e_1 = \text{if}_{p, i}^{v_1} b_1$ then p_{1_T} else p_{1_F} and for some $n, \Phi(p)(\bar{i}_k) = n$.
By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[15a] $d_1 = d_2 = \infty$ and for some v_2, b_2, p_{2_T} and p_{2_F} , $e_2 = \text{if}_{p, \bar{i}_k}^{v_2} b_2$ then p_{2_T} else p_{2_F} . Now there are two possible cases depending on the value of b_1 :

(a) $b_1 = \text{true}$. As per IF-SYNC $(H_1, e_1)_{d_1} \xrightarrow{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1$, $e'_1 = \langle p_{1_T} \rangle^{v_1}$ and $d'_1 = n$. There are again two possible cases depending on the value of b_2 :

i. $b_2 = \text{true}$. By IF-SYNC $(H_2, e_2)_{d_2} \xrightarrow{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2$, $e'_2 = \langle p_{2_T} \rangle^{v_2}$, and $d'_2 = n$.

We know $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$; then by Lemma 4.4[15b] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

ii. $b_2 = \text{false}$. Analogous to the subsubcase 3(a)i above. By IF-SYNC $(H_2, e_2)_{d_2} \xrightarrow{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2$, $e'_2 = \langle p_{2_F} \rangle^{v_2}$ and $d'_2 = n$.

We know $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$; then by Lemma 4.4[15b] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

(b) $b_2 = \text{false}$. Analogous to the subcase 3a above.

4. APP-SYNC. (Similar to case 3 above.) So for some $p, \bar{i}_k, v_1, x_1, p_1$ and $v_{1_{arg}}$, $e_1 = (\lambda x_1. p_1)(v_{1_{arg}})_{p, \bar{i}_k}^{v_1}$ and for some $n, \Phi(p)(\bar{i}_k) = n$.

By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[16a] $d_1 = d_2 = \infty$ and for some v_2, x_2, p_2 and $v_{2_{arg}}$, $e_2 = (\lambda x_2. p_2)(v_{2_{arg}})_{p, \bar{i}_k}^{v_2}$.

Then as per APP-SYNC $(H_1, e_1)_{d_1} \xrightarrow{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1$, $e'_1 = \langle p_1[v_{1_{arg}}/x_1] \rangle^{v_1}$ and $d'_1 = n$; and $(H_2, e_2)_{d_2} \xrightarrow{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2$, $e'_2 = \langle p_2[v_{2_{arg}}/x_2] \rangle^{v_2}$, and $d'_2 = n$.

We know $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e_2)_{d_2}$; then by Lemma 4.4[16b] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu}^{\emptyset\tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

5. IF. So for some $p, \bar{i}_k, v_1, b_1, p_{1T}$ and $p_{1F}, e_1 = \text{if}_{\bar{p}, \bar{i}_k}^{v_1} b_1$ then p_{1T} else p_{1F} .

By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[15a] $d_1 = d_2 = \infty$ and for some v_2, b_2, p_{2T} and $p_{2F}, e_2 = \text{if}_{\bar{p}, \bar{i}_k}^{v_2} b_2$ then p_{2T} else p_{2F} . Now there are two possible cases depending on the value of b_1 :

(a) $b_1 = \text{true}$. As per IF $(H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1, e'_1 = p_{1T}, d'_1 = d_1 - 1 = \infty$ and $\Phi(p) = \emptyset$. Then by Lemma 4.4[15c] $b_1 = b_2 = \text{true}$. By IF $(H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2, e'_2 = p_{2T}$ and $d'_2 = d_2 - 1 = \infty$. Then by Lemma 4.4[15c] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

(b) $b_1 = \text{false}$. Analogous to the subcase 5a.

6. APP. (Similar to case 5 above.) So for some $p, \bar{i}_k, v_1, x_1, p_1$ and $v_{1\text{arg}}, e_1 = (\lambda x_1. p_1) (v_{1\text{arg}})_{\bar{p}, \bar{i}_k}^{v_1}$.

By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[16a] $d_1 = d_2 = \infty$ and for some v_2, x_2, p_2 and $v_{2\text{arg}}, e_2 = (\lambda x_2. p_2) (v_{2\text{arg}})_{\bar{p}, \bar{i}_k}^{v_2}$.

As per APP $(H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1, e'_1 = p_1[v_{1\text{arg}}/x_1], d'_1 = d_1 - 1 = \infty$ and $\Phi(p) = \emptyset$; and $(H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2, e'_2 = p_2[v_{2\text{arg}}/x_2]$ and $d'_2 = d_2 - 1 = \infty$. Then by Lemma 4.4[16c] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

7. REF. So for some $v_1, e_1 = \text{ref } v_1$, and for some fresh $loc_1, (H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1 \cup \{loc_1 \mapsto v_1\}, e'_1 = loc_1$ and $d'_1 = d_1 - 1$.

By premise $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ holds; then by Lemma 4.4[17a] $d_1 = d_2 = \infty$ and for some $v_2, e_2 = \text{ref } v_2$. By REF for some fresh $loc_2, (H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2 \cup \{loc_2 \mapsto v_2\}, e'_2 = loc_2$ and $d'_2 = d_2 - 1 = \infty$. Now loc_1 and loc_2 are fresh implying $loc_1 \notin \text{dom}(H_1)$ and $loc_2 \notin \text{dom}(H_2)$. We know by premise $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$; then by Lemma 4.4[17b] for some τ' and ℓ' , $\tau = \tau' \text{ref}^{\ell'}$ and $(\mathcal{H}'_1, H'_1, loc_1)_{d'_1} \sim_{\Phi \ell' \mu'}^{\emptyset \tau} (\mathcal{H}'_2, H'_2, loc_2)_{d'_2}$, that is, $(\mathcal{H}'_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell' \mu'}^{\emptyset \tau} (\mathcal{H}'_2, H'_2, e'_2)_{d'_2}$, the second requisite result, where $\mathcal{H}'_1 = \mathcal{H}_1 \cup \{loc_1 \mapsto \tau'\}, \mathcal{H}'_2 = \mathcal{H}_2 \cup \{loc_2 \mapsto \tau'\}$ and $\mu' = \mu \cup \{loc_1 \mapsto loc_2\}$.

8. Deref. So for some $loc_1, e_1 = !loc_1$, and for $H_1(loc_1) = v_1, (H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1, e'_1 = v_1$, and $d'_1 = d_1 - 1$.

By premise $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$ holds; then by Lemma 4.4[18a] $d_1 = d_2 = \infty$ and for some $loc_2 \in \text{dom}(H_2), e_2 = !loc_2$. By Deref for $H_2(loc_2) = v_2, (H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first

requisite result, such that $H'_2 = H_2, e'_2 = v_2$ and $d'_2 = d_2 - 1 = \infty$. Then by Lemma 4.4[18b] $(\mathcal{H}_1, H_1, H_1(loc_1))_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, H_2(loc_2))_{d_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$ holds, the second requisite result.

9. ASSIGN. So for some loc_1 and $v_1, e_1 = (loc_1 := v_1), loc_1 \in \text{dom}(H_1)$ and $(H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1[loc_1 \mapsto v_1], e'_1 = v_1$ and $d'_1 = d_1 - 1$.

By premise $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$ holds; then by Lemma 4.4[19a] $d_1 = d_2 = \infty$ and for some $loc_2 \in \text{dom}(H_2)$, and $v_2, e_2 = (loc_2 := v_2)$. By ASSIGN $(H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2[loc_2 \mapsto v_2], e'_2 = v_2$, and $d'_2 = d_2 - 1 = \infty$. Then by Lemma 4.4[19b] $(\mathcal{H}_1, H_1[loc_1 \mapsto v_1], v_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2[loc_2 \mapsto v_2], v_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$ holds, the second requisite result.

10. SYNC. So for some $s_1, p_1, v_1, e'_1, s'_1, p'_1$ and $n, e_1 = s_1 = \langle p_1 \rangle^{v_1}, d_1 = n, (H_1, p_1)_n \rightarrow_{\Phi} (H'_1, p'_1)_{n-1}$ and $(H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $e'_1 = s'_1 = \langle p'_1 \rangle^{v_1}$ and $d'_1 = d_1 - 1$; hence $d_1 > 0$, that is, $n > 0$.

By premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ hold; then by Lemma 4.4[20a] for some p_2, v_2 and $s_2, e_2 = s_2 = \langle p_2 \rangle^{v_2}$ and $d_1 = d_2 = n$, and by Definition 4.3[3a] $\perp, \emptyset, \mathcal{H}_2 \vdash_{\Phi \ell} H_2$.

We know $(\mathcal{H}_1, s_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, s_2)_{d_2}$; then by Lemma 4.4[6] $\perp, \emptyset, \mathcal{H}_1 \vdash_{\Phi \ell} s_1 : \tau$ and $\perp, \emptyset, \mathcal{H}_2 \vdash_{\Phi \ell} s_2 : \tau$; hence $\perp, \emptyset, \mathcal{H}_2 \vdash_{\Phi \ell} (H_2, s_2)_n : \tau$ holds. By Main Lemma 3.8 (Progress) there exists a H'_2, e'_2 and d'_2 such that $(H_2, s_2)_n \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, that is, $(H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result.

By SYNC $d'_2 = n - 1$, that is, $d'_2 = d_2 - 1$, and for some s'_2 and $p'_2, e'_2 = s'_2 = \langle p'_2 \rangle^{v_2}$. By Corollary 4.7 (Preservation of Heap Bisimilarity under Synchronized Computation) there exists a \mathcal{H}'_1 and \mathcal{H}'_2 such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1, \mathcal{H}_2 \subseteq \mathcal{H}'_2, \perp, \emptyset, \mathcal{H}'_1 \vdash_{\Phi \ell} e'_1 : \tau, \perp, \emptyset, \mathcal{H}'_2 \vdash_{\Phi \ell} e'_2 : \tau$ and $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}'_2, H'_2)$.

By Lemma 4.4[7] $(\mathcal{H}'_1, s'_1)_{n-1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}'_2, s'_2)_{n-1}$, that is, $(\mathcal{H}'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}'_2, e'_2)_{d'_2}$. Knowing $(\mathcal{H}'_1, H'_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}'_2, H'_2)$ by Definition 4.3[4] we get $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

11. SYNC-PAD. Similar to the case 10 above.

12. SYNC-DONE. So for some v'_1 and $v_1, e_1 = \langle v'_1 \rangle^{v_1}, d_1 = 0$ and $(H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $H'_1 = H_1, e'_1 = v'_1$ and $d'_1 = \infty$.

By premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ hold; then by Lemma 4.4[20a] for some p_2 , and $v_2, e_2 = \langle p_2 \rangle^{v_2}$ and $d_1 = d_2 = 0$. Now there are two possible cases depending on whether p_2 is a value or not:

(a) $p_2 = v'_2$, for some v'_2 . By SYNC-DONE $(H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2, e'_2 = v'_2$ and $d'_2 = \infty$. Then by Lemma 4.4[20b] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is,

$(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

- (b) p_2 is not a value. By SYNC-ABORT $(H_2, e_2)_{d_2} \longrightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $H'_2 = H_2$, $e'_2 = v_2$ and $d'_2 = \infty$. Then by Lemma 4.4[20b] $(\mathcal{H}_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e'_2)_{d'_2}$ holds, which in conjunction with the premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ implies $(\mathcal{H}_1, H_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e'_2)_{d'_2}$, that is, $(\mathcal{H}_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H'_2, e'_2)_{d'_2}$, the second requisite result.

13. SYNC-ABORT. Analogous to the case 12 above given Lemma 4.4[20b].

14. CONTEXT. So for some $R_1, e_{1_{sub}}$ and $e'_{1_{sub}}, e_1 = R_1[e_{1_{sub}}]$ and for some $e'_{1_{sub}}, (H_1, e_{1_{sub}})_{d_1} \longrightarrow_{\Phi} (H'_1, e'_{1_{sub}})_{d'_1}$ and $(H_1, e_1)_{d_1} \longrightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ such that $e'_1 = R_1[e'_{1_{sub}}]$.

By premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, e_2)_{d_2}$ hold; then by Lemma 4.4[21a] for some $R_2, e_{2_{sub}}$ and $\tau_{sub}, e_2 = R_2[e_{2_{sub}}]$ and $(\mathcal{H}_1, e_{1_{sub}})_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau_{sub}} (\mathcal{H}_2, e_{2_{sub}})_{d_2}$.

Knowing $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\emptyset} (\mathcal{H}_2, H_2)$, by Definition 4.3[4] we get $(\mathcal{H}_1, H_1, e_{1_{sub}})_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau_{sub}} (\mathcal{H}_2, H_2, e_{2_{sub}})_{d_2}$. By assumption we know $(H_1, e_{1_{sub}})_{d_1} \longrightarrow_{\Phi} (H'_1, e'_{1_{sub}})_{d'_1}$. Then by induction hypothesis for some $H'_2, e'_{2_{sub}}$ and $d'_2, (H_2, e_{2_{sub}})_{d_2} \longrightarrow_{\Phi} (H'_2, e'_{2_{sub}})_{d'_2}$, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' , such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1, \mathcal{H}_2 \subseteq \mathcal{H}'_2, \mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, e'_{1_{sub}})_{d'_1} \sim_{\Phi \ell \mu'}^{\emptyset \tau} (\mathcal{H}'_2, H'_2, e'_{2_{sub}})_{d'_2}$.

Then by CONTEXT $(H_2, e_2)_{d_2} \longrightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, the first requisite result, such that $e'_2 = R_2[e'_{2_{sub}}]$. Finally by Lemma 4.4[21b] $(\mathcal{H}'_1, H'_1, R_1[e'_{1_{sub}}])_{d'_1} \sim_{\Phi \ell \mu'}^{\emptyset \tau} (\mathcal{H}'_2, H'_2, R_2[e'_{2_{sub}}])_{d'_2}$, that is, $(\mathcal{H}'_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu'}^{\emptyset \tau} (\mathcal{H}'_2, H'_2, e'_2)_{d'_2}$ holds, the second requisite result. \square

The following n -step strong bisimulation lemma states: given a pair of bisimilar configurations, if the first configuration takes n steps, then the second configuration also take n steps, and bisimilarity is preserved in the resulting configurations.

Lemma 4.9 (λ_{seq}^{sync} : Strong Bisimulation, n -steps). *If $(H_1, e_1)_{d_1} \longrightarrow_{\Phi}^n (H'_1, e'_1)_{d'_1}$ and $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \tau} (\mathcal{H}_2, H_2, e_2)_{d_2}$ then for some H'_2, e'_2 and $d'_2, (H_2, e_2)_{d_2} \longrightarrow_{\Phi}^n (H'_2, e'_2)_{d'_2}$ holds, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1, \mathcal{H}_2 \subseteq \mathcal{H}'_2, \mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, e'_1)_{d'_1} \sim_{\Phi \ell \mu'}^{\emptyset \tau} (\mathcal{H}'_2, H'_2, e'_2)_{d'_2}$.*

Proof. By induction on the derivation of $(H_1, e_1)_{d_1} \longrightarrow_{\Phi}^n (H'_1, e'_1)_{d'_1}$ given Lemma 4.8 (λ_{seq}^{sync} : Strong Bisimulation, 1-step). \square

The strong bisimulation lemma then entails the following strong noninterference property. It states: if one run of a program converges to a low integral or boolean value in a certain number of steps, then a second run of that program, but possibly differing in

high values from the first run, also converges to the same value and in the same number of steps as the first one.

Theorem 4.10 (λ_{seq}^{sync} : Strong Noninterference). *If $\perp, \{\overline{x_k} \mapsto \overline{\tau_k}\}, \emptyset \vdash_{\Phi \ell}^{low} p : \hat{\tau}, \perp, \emptyset, \emptyset \vdash_{\Phi \ell}^{high} v_k, v'_k : \tau_k, p_1 = p[v_k/x_k], p_2 = p[v'_k/x_k]$, and p_1 given Φ converges to a value \hat{v} in n steps, then p_2 given Φ converges to the same value \hat{v} in n steps.*

Proof. By premise, for some H_1 and $d_1, (\emptyset, p_1)_{\infty} \longrightarrow_{\Phi}^n (H_1, \hat{v})_{d_1}$ and Definition 4.3[1,3,4] $(\emptyset, \emptyset, p_1)_{\infty} \sim_{\Phi \ell \emptyset}^{\emptyset \hat{\tau}} (\emptyset, \emptyset, p_2)_{\infty}$. Then by Lemma 4.9 (λ_{seq}^{sync} : Strong Bisimulation, n -steps) for some H_2, e_2 and $d_2, (\emptyset, p_2)_{\infty} \longrightarrow_{\Phi}^n (H_2, e_2)_{d_2}$, and there exists a $\mathcal{H}_1, \mathcal{H}_2$ and μ such that $\emptyset \subseteq \mathcal{H}_1, \emptyset \subseteq \mathcal{H}_2, \emptyset \subseteq \mu$ and $(\mathcal{H}_1, H_1, \hat{v})_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \hat{\tau}} (\mathcal{H}_2, H_2, e_2)_{d_2}$. By Definition 4.3[4,1] $(\mathcal{H}_1, \hat{v})_{d_1} \sim_{\Phi \ell \mu}^{\emptyset \hat{\tau}} (\mathcal{H}_2, e_2)_{d_2}$ and $d_1 = d_2 = \infty$; then by Lemma 4.4[11c] $e_2 = \hat{v}$. That is, p_2 given Φ converges to the value \hat{v} in n steps, the requisite result. \square

The following corollary to the above lemma states the strong noninterference property in presence of divergence.

Corollary 4.11 (λ_{seq}^{sync} : Strong Noninterference under Divergence). *If $\perp, \{\overline{x_k} \mapsto \overline{\tau_k}\}, \emptyset \vdash_{\Phi \ell}^{low} p : \hat{\tau}, \perp, \emptyset, \emptyset \vdash_{\Phi \ell}^{high} v_k, v'_k : \tau_k, p_1 = p[v_k/x_k], p_2 = p[v'_k/x_k]$, and p_1 given Φ diverges, then p_2 given Φ diverges as well.*

Proof. By Lemma 3.6 (Value Substitution in a Program), Theorem 3.13 (Type Safety) and Theorem 4.10 (λ_{seq}^{sync} : Strong Noninterference). \square

The combination of Theorem 4.10 and Corollary 4.11 imply the security of timing and termination channels in λ_{seq}^{sync} .

Definition 4.12 (Secure Program). *A program p given a table of sync timer functions Φ is secure iff for $free(p) = \{\overline{x_k}\}$ and for all $\overline{v_k}$ and $\overline{v'_k}$, either both $p[\overline{v_k}/\overline{x_k}]$ and $p[\overline{v'_k}/\overline{x_k}]$ given Φ converge to the same value in the same number of steps, or both diverge.*

Definition 4.13 (Secure-Typed Program). *A program p given a table of sync timer functions Φ is secure-typed iff $\perp, \{\overline{x_k} \mapsto \overline{int}^{\ell_k}\}, \emptyset \vdash_{\Phi \ell}^{low} p : \hat{\tau}$, where $free(p) = \{\overline{x_k}\}, \ell_k \preceq \ell$ and some $\hat{\tau}$.*

The following corollary to the above noninterference lemmas directly states the security guarantee of $\lambda_{seq}^{sync} - \lambda_{seq}^{sync}$ is secure with respect to external timing channels.

Theorem 4.14 (Security of λ_{seq}^{sync}). *Secure-typed programs are secure.*

Proof. By Theorem 4.10 (λ_{seq}^{sync} : Strong Noninterference) and Corollary 4.11 (λ_{seq}^{sync} : Strong Noninterference under Divergence). \square

5. The $\lambda_{conc \mathcal{D}}^{sync}$ Runtime System

We now show how the results of the previous section can directly generalize to the case of concurrent execution, with a deterministic scheduler in this section, and with a nondeterministic scheduler in the following section. Here we formalize the $\lambda_{conc \mathcal{D}}^{sync}$ system as an extension to λ_{seq}^{sync} . The syntax for $\lambda_{conc \mathcal{D}}^{sync}$ appears in Figure 11 and its small-step operational semantics in Figure 12, as addendums to Figures 7 and 8 respectively. The trace t denotes the scheduling history – the interleavings thus far – of a pool of threads; it represents the internal timing behavior of concurrent programs, as discussed in Section 2. The scheduler \mathcal{D} denotes an arbitrary deterministic

t	::= $\epsilon \mid t, n$	<i>trace</i>
\hat{H}	::= $\{\overline{loc}\} \rightarrow \{\overline{v}\}$	<i>heap with only integer/boolean values</i>
\mathcal{D}	: $\mathbb{N} \times \{\overline{t}\} \times \{\overline{H}\} \rightarrow \mathbb{N}$	<i>deterministic scheduler</i>
ζ	::= (e, d)	<i>thread</i>
T	::= $\overline{\zeta}$	<i>pool of threads</i>
C	::= $(H, T)_t$	<i>concurrent runtime configuration</i>
\mathbb{L}	::= $\{\overline{loc}\}$	<i>set of "low integer/boolean" heap locations</i>
Υ	::= $\overline{\tau}$	<i>pool of types</i>
$\hat{\mathcal{H}}$::= $\{\overline{loc}\} \rightarrow \{\overline{\tau}\}$	<i>type of heap with only int/bool types</i>

Figure 11. $\lambda_{conc\mathcal{D}}^{sync}$: Syntax and Type Grammar

THREAD	$\frac{\zeta = (e, d) \quad (H, e)_d \longrightarrow_{\Phi} (H', e')_{d'} \quad \zeta' = (e', d')}{(H, \zeta) \longrightarrow_{\Phi} (H', \zeta')}$
THREAD-POOL-NEXT-SCHEDULE	$\frac{T[i] = \zeta_i \quad (H, \zeta_i) \longrightarrow_{\Phi} (H', \zeta'_i)}{(H, T)_t \longrightarrow_{\Phi i} (H', T[i \mapsto \zeta'_i])_{t,i}}$
THREAD-POOL- \mathcal{D}	$\frac{\hat{H}_{low} = H \upharpoonright \mathbb{L} \quad \mathcal{D}(k, t, \hat{H}_{low}) = i \quad (H, T)_t \longrightarrow_{\Phi i} (H', T')_{t'}}{(H, T)_t \longrightarrow_{\Phi \mathcal{L}\mathcal{D}} (H', T')_{t'}} \quad T = k$

Figure 12. $\lambda_{conc\mathcal{D}}^{sync}$: Syntax-Directed Small-Step Operational Semantics

scheduler (possibly chosen by an attacker), and is parameterized over a 3-tuple: the number of threads in the thread pool, the trace so far, and a low portion of the heap with only integer or boolean values; it returns an index into the thread pool for the next scheduled thread. Each thread ζ is a pair of an expression e and a down-counter d .

The small-step reduction relations $\longrightarrow_{\Phi i}$ and $\longrightarrow_{\Phi \mathcal{L}\mathcal{D}}$, for a table of sync timer functions Φ , the next scheduled thread i , a set of low integer/boolean heap locations \mathbb{L} (to be fed to the scheduler), and a deterministic scheduler \mathcal{D} , are each defined over configurations $(H, T)_t$, where the heap H is shared amongst all threads in the thread pool T ; the n -step reflexive and transitive closure of $\longrightarrow_{\Phi \mathcal{L}\mathcal{D}}$ is denoted as $\longrightarrow_{\Phi \mathcal{L}\mathcal{D}}^n$. The down-counter in a thread is decremented via the THREAD rule only when the thread is under execution, meaning only the cumulative scheduled time of a thread, and not the wall-clock time, is counted towards the synchronized execution time of a branching statement in that thread. The restricted-to operator \upharpoonright is defined as: $H \upharpoonright \mathbb{L} = H'$ iff $H' \subseteq H$ and $dom(H) = \mathbb{L}$, and the indexing operation on a thread pool T is defined as: $T[i] = \zeta_i$ iff $T = \overline{\zeta_k}$, for some $\overline{\zeta_k}$, and $1 \leq i \leq k$. The thread update operation $T[i \mapsto \zeta'_i]$, where i denotes the index of the thread to be updated to ζ'_i , is defined to be equal to T' iff $T = \zeta_1, \dots, \zeta_i, \dots, \zeta_k$, for some $\overline{\zeta_k}$, and $T' = \zeta_1, \dots, \zeta'_i, \dots, \zeta_k$; while the number of threads in a thread pool is defined as: $|T| = k$ iff $T = \overline{\zeta_k}$, for some $\overline{\zeta_k}$. For technical ease the scheduler is invoked after each step, via the rule THREAD-POOL- \mathcal{D} ; more liberal scheduling policies are subsumed by this approach.

$\lambda_{conc\mathcal{D}}^{sync}$ does not support dynamic thread creation, or explicit thread yielding. The former can be incorporated using the thread pool partitioning technique of [34]; while the latter is directly encodable in our system. A thread ζ wanting to yield control can do so by simply setting a specific location in the portion of the heap passed to the scheduler – \hat{H}_{low} in the rule THREAD-POOL- \mathcal{D} – thus

informing the scheduler of its intent to yield. As only low heap locations are observable to the scheduler, yielding is disallowed under high guards – yielding has direct impact on the interleavings of threads, and hence must be prohibited in high contexts to prevent internal timing leaks; [34] enforces the same restriction.

5.1 Security Properties of $\lambda_{conc\mathcal{D}}^{sync}$

We now formally establish the security guarantee of $\lambda_{conc\mathcal{D}}^{sync}$. We show $\lambda_{conc\mathcal{D}}^{sync}$ is secure with respect to both external and internal timing channels by proving a strong noninterference result between high data and the external and internal timing behavior of concurrent programs. We demonstrate how executions of two pools of threads, which differ only in high values, are strongly bisimilar. Strong bisimilarity implies isomorphism of the pools of threads at all intermediate steps of execution; as for λ_{seq}^{sync} , the term *strong bisimilarity* indicates the lock-stepness of bisimilar executions. The bisimulation relation, defined below, essentially requires all low values in corresponding threads as well as the interleavings of the concurrent executions to be identical, while allowing high values to differ; it is an addendum to λ_{seq}^{sync} 's bisimulation relation (Definition 4.3). We define the predicate $secllevel(\mathcal{H}) \leq \ell$ to hold iff $\forall \tau \in range(\mathcal{H}). secllevel(\tau) \leq \ell$.

Definition 5.1 ($\lambda_{conc\mathcal{D}}^{sync}$: Bisimulation Relation).

1. (*Threads*). $(\mathcal{H}_1, \zeta_1) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, \zeta_2)$ iff $\zeta_1 = (e_1, d_1)$, $\zeta_2 = (e_2, d_2)$ and $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, e_2)_{d_2}$.
2. (*Pools of Threads*). $(\mathcal{H}_1, T_1)_{t_1} \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, T_2)_{t_2}$ iff $t_1 = t_2$ and for some $\overline{\zeta_k}, \overline{\zeta'_k}$ and $\overline{\tau_k}$,
 - (a) $T_1 = \overline{\zeta_k}$, $T_2 = \overline{\zeta'_k}$, and $\Upsilon = \overline{\tau_k}$; and
 - (b) $\forall 1 \leq i \leq k. (\mathcal{H}_1, \zeta_i) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, \zeta'_i)$.
3. (*Heaps*). $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\Upsilon} (\mathcal{H}_2, H_2)$ iff for some $\hat{\mathcal{H}}$,
 - (a) $\mathcal{H}_1 \upharpoonright \mathbb{L} = \mathcal{H}_2 \upharpoonright \mathbb{L} = \hat{\mathcal{H}}$, and $secllevel(\hat{\mathcal{H}}) \leq \ell$; and
 - (b) $\forall loc \in \mathbb{L}. loc \mapsto loc \in \mu$, and $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, H_2)$.
4. (*Heaps, Threads*). $(\mathcal{H}_1, H_1, \zeta_1) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, H_2, \zeta_2)$ iff $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, \zeta_1) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, \zeta_2)$.
5. (*Heaps, Pools of Threads*). $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi \ell \mu \mathbb{L}}^{\Upsilon} (\mathcal{H}_2, H_2, T_2)_{t_2}$ iff $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\Upsilon} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, T_1)_{t_1} \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, T_2)_{t_2}$.

As discussed in Section 2, the trace of a concurrent execution constitutes an internal timing channel, and hence must exhibit low behavior; accordingly the case 2 above requires the traces of bisimilar configurations, t_1 and t_2 , to be identical. Further the corresponding threads in bisimilar thread pools are required to be bisimilar. The portions of bisimilar heaps to be passed to the scheduler, as indicated by \mathbb{L} , are restricted to contain only low integral or boolean values by the condition 3a – low function values may have high values wrapped inside them as per Definition 4.3[1], while low heap locations may store high values as implied by the type rule *location*; hence neither is guaranteed to be independent of high data, and thus cannot be passed to the scheduler without potentially violating noninterference. The other two parameters to the scheduler, the number of threads in the thread pool and the trace, are both low, thus guaranteeing the scheduler behavior is low as well. Further the above bisimulation relation is scheduler-independent, akin to the one in [38].

Lemma 5.2 ($\lambda_{conc\mathcal{D}}^{sync}$: Properties of Bisimulation Relation).

1. (*Symmetry*).
 - (a) (*Heaps, Threads*). If $(\mathcal{H}_1, H_1, \zeta_1) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_2, H_2, \zeta_2)$ then $(\mathcal{H}_2, H_2, \zeta_2) \sim_{\Phi \ell \mu}^{\Upsilon} (\mathcal{H}_1, H_1, \zeta_1)$.

(b) (Heaps, Pools of Threads). If $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi\ell\mu\mathbb{L}}^{\Gamma\Upsilon}$
 $(\mathcal{H}_2, H_2, T_2)_{t_2}$ then $(\mathcal{H}_2, H_2, T_2)_{t_2} \sim_{\Phi\ell\mu^{-1}\mathbb{L}}^{\Gamma\Upsilon}$
 $(\mathcal{H}_1, H_1, T_1)_{t_1}$.

2. (Reflexivity).

(a) (Threads). If $\zeta = (e, d)$ and $(\mathcal{H}, e)_d \sim_{\Phi\ell\mu}^{\Gamma\tau}$ $(\mathcal{H}, e)_d$ then
 $(\mathcal{H}, \zeta) \sim_{\Phi\ell\mu}^{\Gamma\tau}$ (\mathcal{H}, ζ) .

(b) (Pool of Threads). If $T = \overline{\zeta}_k$, $\Upsilon = \overline{\tau}_k$ and $\forall i \leq k$.
 $(\mathcal{H}, \zeta_i) \sim_{\Phi\ell\mu}^{\Gamma\tau_i}$ (\mathcal{H}, ζ_i) then for any t , $(\mathcal{H}, T)_t \sim_{\Phi\ell\mu}^{\Gamma\Upsilon}$
 $(\mathcal{H}, T)_t$.

(c) (Heaps). If $(\mathcal{H}, H) \sim_{\Phi\ell\mu}^{\Gamma}$ (\mathcal{H}, H) , $\mathcal{H}[\mathbb{L} = \hat{\mathcal{H}}$,
 $\text{secl}(\hat{\mathcal{H}}) \leq \ell$ and $\forall \text{loc} \in \mathbb{L}. \text{loc} \mapsto \text{loc} \in \mu$, then
 $(\mathcal{H}, H) \sim_{\Phi\ell\mu\mathbb{L}}^{\Gamma}$ (\mathcal{H}, H) .

(d) (Heaps, Threads). If $\zeta = (e, d)$, $(\mathcal{H}, e)_d \sim_{\Phi\ell\mu}^{\Gamma\tau}$ $(\mathcal{H}, e)_d$
and $(\mathcal{H}, H) \sim_{\Phi\ell\mu}^{\Gamma}$ (\mathcal{H}, H) then $(\mathcal{H}, H, \zeta) \sim_{\Phi\ell\mu}^{\Gamma\tau}$ (\mathcal{H}, H, ζ) .

(e) (Heaps, Pools of Threads). If $(\mathcal{H}, T)_t \sim_{\Phi\ell\mu}^{\Gamma\Upsilon}$ $(\mathcal{H}, T)_t$
and $(\mathcal{H}, H) \sim_{\Phi\ell\mu\mathbb{L}}^{\Gamma}$ (\mathcal{H}, H) then $(\mathcal{H}, H, T)_t \sim_{\Phi\ell\mu\mathbb{L}}^{\Gamma\Upsilon}$
 $(\mathcal{H}, H, T)_t$.

3. (Transitivity).

(a) (Heaps, Threads). If $(\mathcal{H}_1, H_1, \zeta_1) \sim_{\Phi\ell\mu}^{\Gamma\tau}$ $(\mathcal{H}_2, H_2, \zeta_2)$ and
 $(\mathcal{H}_2, H_2, \zeta_2) \sim_{\Phi\ell\mu'}^{\Gamma\tau}$ $(\mathcal{H}_3, H_3, \zeta_3)$ then for $\mu'' = \mu' \circ \mu$,
 $(\mathcal{H}_1, H_1, \zeta_1) \sim_{\Phi\ell\mu''}^{\Gamma\tau}$ $(\mathcal{H}_3, H_3, \zeta_3)$.

(b) (Heaps, Pools of Threads). If $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi\ell\mu\mathbb{L}}^{\Gamma\Upsilon}$
 $(\mathcal{H}_2, H_2, T_2)_{t_2}$ and $(\mathcal{H}_2, H_2, T_2)_{t_2} \sim_{\Phi\ell\mu'\mathbb{L}}^{\Gamma\Upsilon}$
 $(\mathcal{H}_3, H_3, T_3)_{t_3}$ then for $\mu'' = \mu' \circ \mu$,
 $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi\ell\mu''\mathbb{L}}^{\Gamma\Upsilon}$ $(\mathcal{H}_3, H_3, T_3)_{t_3}$.

4. (Extension of Types of Heaps and μ).

(a) (Threads). If $(\mathcal{H}_1, \zeta_1) \sim_{\Phi\ell\mu}^{\Gamma\tau}$ (\mathcal{H}_2, ζ_2) , $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq$
 \mathcal{H}'_2 and $\mu \subseteq \mu'$ then $(\mathcal{H}'_1, \zeta_1) \sim_{\Phi\ell\mu'}^{\Gamma\tau}$ $(\mathcal{H}'_2, \zeta_2)$.

(b) (Pools of Threads). If $(\mathcal{H}_1, T_1)_{t_1} \sim_{\Phi\ell\mu}^{\Gamma\Upsilon}$ $(\mathcal{H}_2, T_2)_{t_2}$, $\mathcal{H}_1 \subseteq$
 \mathcal{H}'_1 , $\mathcal{H}_2 \subseteq \mathcal{H}'_2$ and $\mu \subseteq \mu'$ then $(\mathcal{H}'_1, T_1)_{t_1} \sim_{\Phi\ell\mu'}^{\Gamma\Upsilon}$
 $(\mathcal{H}'_2, T_2)_{t_2}$.

5. (Equality of Int/Bool Typed “Low” Sub-Heaps). If
 $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu\mathbb{L}}^{\Gamma}$ (\mathcal{H}_2, H_2) then $H_1[\mathbb{L} = H_2[\mathbb{L}$.

Proof. 1. (Symmetry).

(a) (Heaps, Threads). By Definition 5.1[4,1] and
Lemma 4.4[1].

(b) (Heaps, Pools of Threads). By Definition 5.1[5,2] and
Lemma 5.2[1a].

2. (Reflexivity). By Definition 5.1.

3. (Transitivity).

(a) (Heaps, Threads). By Definition 5.1[4,1] and
Lemma 4.4[3].

(b) (Heaps, Pools of Threads). By Definition 5.1[5,2] and
Lemma 5.2[3a].

4. (Extension of Types of Heaps and μ).

(a) (Threads). By Definition 5.1[1] and Lemma 4.4[10a,9].

(b) (Pools of Threads). By Definition 5.1[2] and
Lemma 5.2[4a].

5. (Equality of Int/Bool Typed “Low” Sub-Heaps). By Defini-
tion 5.1[3], Definition 4.3[3] and Lemma 4.4[11c]. \square

The following single-step strong bisimulation lemma states:
given a pair of bisimilar configurations, if the first configuration
steps, then the second configuration also steps, and bisimilarity
is preserved in the resulting configurations. For the most part the

lemma follows directly from Lemma 4.8 (λ_{seq}^{sync} : Strong Bisimu-
lation, 1-step); the only exception being the isomorphism of the
traces in case 3 below, which is a direct consequence of low sched-
uler behavior.

Lemma 5.3 (λ_{conc}^{sync} : Strong Bisimulation, 1-step).

1. (Threads). If $(H_1, \zeta_1) \rightarrow_{\Phi} (H'_1, \zeta'_1)$ and $(\mathcal{H}_1, H_1, \zeta_1) \sim_{\Phi\ell\mu}^{\emptyset\tau}$
 $(\mathcal{H}_2, H_2, \zeta_2)$ then for some H'_2 and ζ'_2 , $(H_2, \zeta_2) \rightarrow_{\Phi} (H'_2, \zeta'_2)$
holds, and there exists a \mathcal{H}'_1 , \mathcal{H}'_2 and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$,
 $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, \zeta'_1) \sim_{\Phi\ell\mu'}^{\emptyset\tau}$ $(\mathcal{H}'_2, H'_2, \zeta'_2)$.
2. (Pools of Threads with Next Schedule). If $(H_1, T_1)_{t_1} \rightarrow_{\Phi i}$
 $(H'_1, T'_1)_{t'_1}$ and $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi\ell\mu\mathbb{L}}^{\emptyset\Upsilon}$ $(\mathcal{H}_2, H_2, T_2)_{t_2}$ then
for some H'_2 , T'_2 and t'_2 , $(H_2, T_2)_{t_2} \rightarrow_{\Phi i} (H'_2, T'_2)_{t'_2}$ holds,
and there exists a \mathcal{H}'_1 , \mathcal{H}'_2 and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq$
 \mathcal{H}'_2 , $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, T'_1)_{t'_1} \sim_{\Phi\ell\mu'\mathbb{L}}^{\emptyset\Upsilon}$ $(\mathcal{H}'_2, H'_2, T'_2)_{t'_2}$.
3. (Pools of Threads with Deterministic Scheduler). If
 $(H_1, T_1)_{t_1} \rightarrow_{\Phi LD} (H'_1, T'_1)_{t'_1}$ and $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi\ell\mu\mathbb{L}}^{\emptyset\Upsilon}$
 $(\mathcal{H}_2, H_2, T_2)_{t_2}$ then for some H'_2 , T'_2 and t'_2 ,
 $(H_2, T_2)_{t_2} \rightarrow_{\Phi LD} (H'_2, T'_2)_{t'_2}$ holds, and there exists a
 \mathcal{H}'_1 , \mathcal{H}'_2 and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and
 $(\mathcal{H}'_1, H'_1, T'_1)_{t'_1} \sim_{\Phi\ell\mu'\mathbb{L}}^{\emptyset\Upsilon}$ $(\mathcal{H}'_2, H'_2, T'_2)_{t'_2}$.

Proof. 1. (Threads). By premise $(H_1, \zeta_1) \rightarrow_{\Phi} (H'_1, \zeta'_1)$; as
per THREAD for some e_1, d_1, e'_1 and d'_1 , $\zeta_1 = (e_1, d_1)$,
 $(H_1, e_1)_{d_1} \rightarrow_{\Phi} (H'_1, e'_1)_{d'_1}$ and $\zeta'_1 = (e'_1, d'_1)$.

By premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu}^{\emptyset}$ (\mathcal{H}_2, H_2) and $(\mathcal{H}_1, \zeta_1) \sim_{\Phi\ell\mu}^{\emptyset\tau}$
 (\mathcal{H}_2, T_2) hold, and by Definition 5.1[1] for some e_2 and d_2 ,
 $\zeta_2 = (e_2, d_2)$ and $(\mathcal{H}_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau}$ $(\mathcal{H}_2, e_2)_{d_2}$. Then by
Definition 4.3[4] $(\mathcal{H}_1, H_1, e_1)_{d_1} \sim_{\Phi\ell\mu}^{\emptyset\tau}$ $(\mathcal{H}_2, H_2, e_2)_{d_2}$.

By Lemma 4.8 (λ_{seq}^{sync} : Strong Bisimulation, 1-step) for some
 H'_2 , e'_2 and d'_2 , $(H_2, e_2)_{d_2} \rightarrow_{\Phi} (H'_2, e'_2)_{d'_2}$, and there exists
a \mathcal{H}'_1 , \mathcal{H}'_2 and μ' , such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$
and $(\mathcal{H}'_1, H'_1, e'_1)_{d'_1} \sim_{\Phi\ell\mu'}^{\emptyset\tau}$ $(\mathcal{H}'_2, H'_2, e'_2)_{d'_2}$. By THREAD for
 $\zeta'_2 = (e'_2, d'_2)$, $(H_2, \zeta_2) \rightarrow_{\Phi} (H'_2, \zeta'_2)$, the first requisite
result. Then by Definition 4.3[4] and Definition 5.1[1,4] we get
 $(\mathcal{H}'_1, H'_1, \zeta'_1) \sim_{\Phi\ell\mu'}^{\emptyset\tau}$ $(\mathcal{H}'_2, H'_2, \zeta'_2)$, the second requisite result.

2. (Pools of Threads with Next Schedule). By premise
 $(H_1, T_1)_{t_1} \rightarrow_{\Phi i} (H'_1, T'_1)_{t'_1}$; as per THREAD-
POOL_NEXT-SCHEDULE for some ζ_{1k} and ζ'_{1i} , $T_1 = \overline{\zeta}_{1k}$,
 $(H, \zeta_1) \rightarrow_{\Phi} (H'_1, \zeta'_{1i})$, $T'_1 = T_1[i \mapsto \zeta'_{1i}]$ and $t'_1 = t_1, i$.
By premise $(\mathcal{H}_1, H_1) \sim_{\Phi\ell\mu\mathbb{L}}^{\emptyset}$ (\mathcal{H}_2, H_2) and $(\mathcal{H}_1, T_1)_{t_1} \sim_{\Phi\ell\mu}^{\emptyset\Upsilon}$
 $(\mathcal{H}_2, T_2)_{t_2}$ hold. Then by Definition 5.1[2] $t_1 = t_2$
and for some ζ_{2k} and τ_k , $T_2 = \overline{\zeta}_{2k}$, $\Upsilon = \overline{\tau}_k$
and $(\mathcal{H}_1, \zeta_{1i}) \sim_{\Phi\ell\mu}^{\emptyset\tau_i}$ $(\mathcal{H}_2, \zeta_{2i})$. By Definition 5.1[4]
 $(\mathcal{H}_1, H_1, \zeta_{1i}) \sim_{\Phi\ell\mu}^{\emptyset\tau_i}$ $(\mathcal{H}_2, H_2, \zeta_{2i})$.
By Lemma 5.3[1] for some H'_2 and ζ'_{2i} , $(H_2, \zeta_{2i}) \rightarrow_{\Phi}$
 (H'_2, ζ'_{2i}) , and there exists a \mathcal{H}'_1 , \mathcal{H}'_2 and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$,
 $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, \zeta'_{1i}) \sim_{\Phi\ell\mu'}^{\emptyset\tau_i}$ $(\mathcal{H}'_2, H'_2, \zeta'_{2i})$. By
THREAD-POOL_NEXT-SCHEDULE for $T'_2 = T_2[i \mapsto \zeta'_{2i}]$ and
 $t'_2 = t_2, i$, $(H_2, T_2)_{t_2} \rightarrow_{\Phi i} (H'_2, T'_2)_{t'_2}$ holds, the first req-
uisite result. Further knowing $t_1 = t_2$ and $t'_1 = t_1, i$ implies
 $t'_1 = t'_2$. Finally by Definition 5.1[4,2,5] and Lemma 5.2[4] we
get $(\mathcal{H}'_1, H'_1, T'_1)_{t'_1} \sim_{\Phi\ell\mu'\mathbb{L}}^{\emptyset\Upsilon}$ $(\mathcal{H}'_2, H'_2, T'_2)_{t'_2}$, the second requi-
site result.

3. (Pools of Threads with Deterministic Scheduler). By premise
 $(H_1, T_1)_{t_1} \rightarrow_{\Phi LD} (H'_1, T'_1)_{t'_1}$; as per THREAD-POOL_D

for some $\overline{\zeta_{1k}}$, \hat{H}_{low} and i , $T_1 = \overline{\zeta_{1k}}$, $\hat{H}_{low} = H_1 \upharpoonright \mathbb{L}$, $\mathcal{D}(k, t_1, \hat{H}_{low}) = i$ and $(H_1, T_1)_{t_1} \xrightarrow{\Phi i} (H'_1, T'_1)_{t'_1}$.

By premise $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset} (\mathcal{H}_2, H_2)$ and $(\mathcal{H}_1, T_1)_{t_1} \sim_{\Phi \ell \mu}^{\emptyset \Upsilon} (\mathcal{H}_2, T_2)_{t_2}$ hold. Then by Lemma 5.2[5] $H_1 \upharpoonright \mathbb{L} = H_2 \upharpoonright \mathbb{L} = \hat{H}_{low}$, and by Definition 5.1[2] $t_1 = t_2$ and for some $\overline{\zeta_{2k}}$, $T_2 = \overline{\zeta_{2k}}$. So $\mathcal{D}(k, t_2, \hat{H}_{low}) = \mathcal{D}(k, t_1, \hat{H}_{low}) = i$.

Now we know $(H_1, T_1)_{t_1} \xrightarrow{\Phi i} (H'_1, T'_1)_{t'_1}$ and $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}_2, H_2, T_2)_{t_2}$; by Lemma 5.3[2] for some H'_2, T'_2 and t'_2 , $(H_2, T_2)_{t_2} \xrightarrow{\Phi i} (H'_2, T'_2)_{t'_2}$, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, T'_1)_{t'_1} \sim_{\Phi \ell \mu' \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}'_2, H'_2, T'_2)_{t'_2}$, the second requisite result. Then by $\text{THREAD-POOL-}\mathcal{D}$ $(H_2, T_2)_{t_2} \xrightarrow{\Phi \mathbb{L} \mathcal{D}} (H'_2, T'_2)_{t'_2}$ holds, the first requisite result. \square

The following n -step strong bisimulation lemma then states: given a pair of bisimilar configurations, if the first pool of threads takes n steps with a certain sequence of interleavings, then the second thread pool also exhibits the same external and internal timing behavior, that is, it also takes n steps, and with the same sequence of interleavings, respectively, as the first one; further bisimilarity is preserved in the resulting configurations.

Lemma 5.4 ($\lambda_{conc \mathcal{D}}^{sync}$: Strong Bisimulation, n -steps). *If $(H_1, T_1)_{t_1} \xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}} (H'_1, T'_1)_{t'_1}$ and $(\mathcal{H}_1, H_1, T_1)_{t_1} \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}_2, H_2, T_2)_{t_2}$ then for some H'_2, T'_2 and t'_2 , $(H_2, T_2)_{t_2} \xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}} (H'_2, T'_2)_{t'_2}$ holds, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, T'_1)_{t'_1} \sim_{\Phi \ell \mu' \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}'_2, H'_2, T'_2)_{t'_2}$.*

Proof. By induction on the derivation of $(H_1, T_1)_{t_1} \xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}} (H'_1, T'_1)_{t'_1}$ given Lemma 5.3[3]. \square

Definition 5.5 (Convergence under a Deterministic Scheduler). *A thread i in a configuration C , given a table of sync timer functions Φ , a set of heap locations \mathbb{L} , and a deterministic scheduler \mathcal{D} , converges to a value v in n steps iff $C \xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}} (H, T)_t$ for some H, T and t , and $T[i] = (v, d)$, for some d .*

The following lemma, entailed by the strong bisimulation lemma, then formalizes the property of strong noninterference exhibited by $\lambda_{conc \mathcal{D}}^{sync}$. It states: if a thread in one run of a concurrent program converges to a low integral or boolean value in a certain number of steps, then the same thread in a second run of that concurrent program, but possibly differing in high values from the first run, also converges to the same value and in the same number of steps as the first one.

We define the type judgement $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} (H, T) : \Upsilon$, where $T = \overline{\zeta_k}$ and $\Upsilon = \overline{\tau_k}$, to hold iff $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} H$ and $\forall 1 \leq i \leq k$. $pc, \Gamma, \mathcal{H} \vdash_{\Phi \ell} \zeta_i : \tau_i$. The indexing operation on a type pool is defined as on a thread pool. The multi-substitution $\zeta[\overline{v_k/x_k}]$, where $\zeta = (e, d)$, is shorthand for ζ' such that $\zeta' = (e[\overline{v_k/x_k}], d)$, and the multi-substitution $T[\overline{v_k/x_k}]$, where $T = \overline{\zeta_g}$, is shorthand for T' such that $T' = \overline{\zeta'_g}$ and $\forall 1 \leq i \leq g$. $\zeta'_i = \zeta_i[\overline{v_k/x_k}]$; while $(H, T)[\overline{v_k/x_k}]$ abbreviates $(H[\overline{v_k/x_k}], T[\overline{v_k/x_k}])$.

Theorem 5.6 ($\lambda_{conc \mathcal{D}}^{sync}$: Strong Noninterference). *If $\perp, \{x_k \mapsto \tau_k\}, \mathcal{H} \vdash_{\Phi \ell} (H, T) : \Upsilon$, $\Upsilon[i] = \hat{\tau}$, $\mathcal{H}[\mathbb{L} = \hat{\mathcal{H}}]$,*

$\delta ::= \overline{n}$ *discrete distribution of thread indices*
 $\mathcal{N} : \mathbb{N} \times \{\bar{t}\} \times \{\bar{H}\} \rightarrow \{\bar{\delta}\}$ *nondeterministic scheduler*

Figure 13. $\lambda_{conc \mathcal{N}}^{sync}$: Syntax Grammar

$$\begin{array}{c} \text{THREAD-POOL-}\mathcal{N} \\ \frac{|T| = k \quad \hat{H}_{low} = H \upharpoonright \mathbb{L}}{\mathcal{N}(k, t, \hat{H}_{low}) = \delta \quad i \in \delta \quad (H, T)_t \xrightarrow{\Phi i} (H', T')_{t'}} \\ (H, T)_t \xrightarrow{\Phi \mathbb{L} \mathcal{N}} (H', T')_{t'} \end{array}$$

Figure 14. $\lambda_{conc \mathcal{N}}^{sync}$: Syntax-Directed Small-Step Operational Semantics

$\text{selevel}(\hat{\tau}), \text{selevel}(\hat{\mathcal{H}}) \leq \ell, \perp, \emptyset, \emptyset \vdash_{\Phi \ell}^{\text{high}} \overline{v_k, v'_k : \tau_k}$, $(H_1, T_1) = (H, T)[\overline{v_k/x_k}]$, $(H_2, T_2) = (H, T)[\overline{v'_k/x_k}]$ and the thread i in $(H_1, T_1)_\epsilon$ given Φ, \mathbb{L} and a deterministic scheduler \mathcal{D} converges to a value \hat{v} in n steps, then the thread i in $(H_2, T_2)_\epsilon$ given Φ, \mathbb{L} and \mathcal{D} converges to the same value \hat{v} in n steps.

Proof. By premise for some H'_1, T'_1, t'_1 and d'_1 , $(H_1, T_1)_\epsilon \xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}} (H'_1, T'_1)_{t'_1}$ and $T'_1[i] = (\hat{v}, d'_1)$. Let $\mathbb{L} = \{\overline{loc_k}\}$, for some $\overline{loc_k}$, and $\mu = \{\overline{loc_k \mapsto loc_k}\}$; then by Definition 5.1 and 4.3[1,3] $(\mathcal{H}, H_1, T_1)_\epsilon \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}, H_2, T_2)_\epsilon$. By Lemma 5.4 ($\lambda_{conc \mathcal{D}}^{sync}$: Strong Bisimulation, n -steps) for some H'_2, T'_2 and t'_2 , $(H_2, T_2)_\epsilon \xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}} (H'_2, T'_2)_{t'_2}$ holds, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' such that $\mathcal{H} \subseteq \mathcal{H}'_1$, $\mathcal{H} \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, H'_1, T'_1)_{t'_1} \sim_{\Phi \ell \mu' \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}'_2, H'_2, T'_2)_{t'_2}$. By Definition 5.1[5] $(\mathcal{H}'_1, T'_1)_{t'_1} \sim_{\Phi \ell \mu' \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}'_2, T'_2)_{t'_2}$ holds. Then by Definition 5.1[2,1] and Definition 4.3[1], for some e'_2 and d'_2 such that $T'_2[i] = (e'_2, d'_2)$, $(\mathcal{H}'_1, \hat{v})_{d'_1} \sim_{\Phi \ell \mu'}^{\emptyset \Upsilon} (\mathcal{H}'_2, e'_2)_{d'_2}$ and $d'_1 = d'_2 = \infty$; by Lemma 4.4[11c] $e'_2 = \hat{v}$. That is, the thread i in $(H_2, T_2)_\epsilon$ given Φ, \mathbb{L} and \mathcal{D} converges to the value \hat{v} in n steps, the requisite result. \square

6. The $\lambda_{conc \mathcal{N}}^{sync}$ Runtime System

We now extend $\lambda_{conc \mathcal{D}}^{sync}$ to $\lambda_{conc \mathcal{N}}^{sync}$, which incorporates nondeterminism into the scheduler. The language syntax for $\lambda_{conc \mathcal{N}}^{sync}$ appears in Figure 13, and its small-step operational semantics in Figure 14, as addendums to Figures 11 and 12 respectively. The scheduler \mathcal{N} denotes an arbitrary nondeterministic scheduler (possibly chosen by an attacker), and is parameterized over the same 3-tuple as the deterministic scheduler \mathcal{D} in $\lambda_{conc \mathcal{D}}^{sync}$; however, unlike \mathcal{D} , it returns a discrete distribution δ of indices into the thread pool denoting the probability of each thread to be scheduled next – an index can occur multiple times in the distribution δ , and the more times an index occurs, the greater the likelihood of the corresponding thread to be chosen. The operational semantics rule $\text{THREAD-POOL-}\mathcal{N}$ randomly picks an index from δ choosing the next scheduled thread; we write $i \in \delta$ iff $\delta = \dots i \dots$. The small-step reduction relation $\xrightarrow{\Phi \mathbb{L} \mathcal{N}}$ and its n -step version $\xrightarrow{n_{\Phi \mathbb{L} \mathcal{N}}}$ are defined analogous to the reduction relations $\xrightarrow{\Phi \mathbb{L} \mathcal{D}}$ and $\xrightarrow{n_{\Phi \mathbb{L} \mathcal{D}}}$, respectively, for $\lambda_{conc \mathcal{D}}^{sync}$.

Our nondeterministic scheduler \mathcal{N} is similar to the one in [38], but the latter assigns continuous (real-number) likelihoods to the threads being scheduled, while \mathcal{N} employs discrete probabilities; our discrete probabilities can however be made arbitrarily precise as the distribution δ may be arbitrarily large. We take a discrete probability approach as it allows for a simpler formalism.

\mathbf{C}	::= $\mathbf{C} \mid \overline{\mathbf{C}}$	discrete distribution of runtime configurations
\mathbb{H}	::= $\mathcal{H} \mid \overline{\mathbb{H}}$	discrete distribution of types of heaps
\mathbf{u}	::= $\mu \mid \overline{\mathbf{u}}$	discrete distribution of μ 's
Δ	::= \overline{v}	discrete distribution of integer/boolean values

Figure 15. $\lambda_{conc\mathcal{N}}^{sync}$ (All-Paths-In-Parallel): Grammar (for proof of probabilistic noninterference only)

THREAD-POOL- \mathcal{N} -ALL-PATHS	
$\mathbf{T} = \overline{\zeta}_k$	$\mathbf{C} = (\mathbb{H}, \mathbf{T})_{\mathbf{t}}$
$\hat{\mathbf{H}}_{low} = \mathbb{H} \upharpoonright_{\mathbb{L}}$	$\mathcal{N}(k, \mathbf{t}, \hat{\mathbf{H}}_{low}) = \delta$
$\forall i \in \{1, \dots, i_n\}. \mathbf{C} \xrightarrow{\Phi_i} \mathbf{C}_i$	$\mathbf{C} = \mathbf{C}_{i_1}, \dots, \mathbf{C}_{i_n}$
$\mathbf{C} \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}$	
CONFIGURATION-DISTRIBUTION	
$\mathbf{C} = \mathbf{C}_1, \dots, \mathbf{C}_k$	
$\forall 1 \leq i \leq k. \mathbf{C}_i \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}'_i$	$\mathbf{C}' = \mathbf{C}'_1, \dots, \mathbf{C}'_k$
$\mathbf{C} \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}'$	

Figure 16. $\lambda_{conc\mathcal{N}}^{sync}$ (All-Paths-In-Parallel): Operational Semantics (for proof of probabilistic noninterference only)

6.1 Security Properties of $\lambda_{conc\mathcal{N}}^{sync}$

The discrete probability distribution permits us to prove the probabilistic noninterference result [38] for $\lambda_{conc\mathcal{N}}^{sync}$ using a technique which aligns better with those used for λ_{seq}^{sync} and $\lambda_{conc\mathcal{D}}^{sync}$. Our proof technique is as follows: execute in parallel all possible paths indicated by the nondeterministic scheduler, and show the final discrete distribution of low values and their timings are independent of high data. To this end we define an all-paths-in-parallel semantics, corresponding the semantics of $\lambda_{conc\mathcal{N}}^{sync}$, in Figure 16, and the corresponding grammar appears in Figure 15. It captures, in parallel, all possible interleavings of a concurrent program under a nondeterministic scheduler. The reduction relation $\rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}}$, for a table of sync timer functions Φ , a set of low integral or boolean heap locations \mathbb{L} , and a nondeterministic scheduler \mathcal{N} , is defined over discrete distributions of configurations \mathbf{C} ; the corresponding n -step reflexive and transitive closure is denoted as $\rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}}^n$.

The following lemma states the soundness of the all-paths-in-parallel semantics relative to the operational semantics of $\lambda_{conc\mathcal{N}}^{sync}$. We write $\mathbf{C} \in \mathbf{C}$ iff either, $\mathbf{C} = \mathbf{C}$, or for some $\overline{\mathbf{C}}_k$, $\mathbf{C} = \overline{\mathbf{C}}_k$ and $\bigvee_{1 \leq i \leq k} \mathbf{C} \in \mathbf{C}_i$.

Lemma 6.1 (Soundness of All-Paths-In-Parallel Semantics). *If $\mathbf{C} \xrightarrow{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}'$, $\mathbf{C} \in \mathbf{C}$ and $\mathbf{C} \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}}^n \mathbf{C}'$ then $\mathbf{C}' \in \mathbf{C}'$.*

Proof. By induction on the derivation of $\mathbf{C} \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}}^n \mathbf{C}'$. \square

$\lambda_{conc\mathcal{N}}^{sync}$ is proved to be probabilistically noninterfering by showing that all-paths-in-parallel executions of two pools of threads, which differ only in high values, are strongly bisimilar. Strong bisimilarity implies isomorphism of the distributions of configurations at all intermediate steps of the all-paths-in-parallel execution, with the term *strong bisimilarity* indicating the lock-stepness of bisimilar executions as before. The bisimulation relation, defined below, inductively entails the bisimilarity of corresponding configurations; it is an addendum to $\lambda_{conc\mathcal{D}}^{sync}$'s bisimulation relation (Definition 5.1).

Definition 6.2 ($\lambda_{conc\mathcal{N}}^{sync}$: Bisimulation Relation).

- (Configurations). $(\mathcal{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathcal{H}_2, \mathbf{C}_2)$ iff $\mathbf{C}_1 = (\mathbb{H}_1, \mathbf{T}_1)_{\mathbf{t}_1}$, $\mathbf{C}_2 = (\mathbb{H}_2, \mathbf{T}_2)_{\mathbf{t}_2}$ and $(\mathcal{H}_1, \mathbb{H}_1, \mathbf{T}_1)_{\mathbf{t}_1} \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathcal{H}_2, \mathbb{H}_2, \mathbf{T}_2)_{\mathbf{t}_2}$.

- (Distributions of Configurations). $(\mathbb{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_2, \mathbf{C}_2)$ iff for some $\overline{\mathbb{H}}'_k, \overline{\mathbb{H}}''_k, \overline{\mathbf{C}}'_k, \overline{\mathbf{C}}''_k$, and $\overline{\mathbf{u}}_k$, $\mathbb{H}_1 = \overline{\mathbb{H}}'_k$, $\mathbb{H}_2 = \overline{\mathbb{H}}''_k$, $\mathbf{C}_1 = \overline{\mathbf{C}}'_k$, $\mathbf{C}_2 = \overline{\mathbf{C}}''_k$, $\mathbf{u} = \overline{\mathbf{u}}_k$, and $\forall 1 \leq i \leq k. (\mathbb{H}'_i, \mathbf{C}'_i) \sim_{\Phi\mathbb{L}\mu_i\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}''_i, \mathbf{C}''_i)$.

Note the above bisimulation relation is scheduler-independent, akin to the one in [38].

The inverse of a pool of μ 's, \mathbf{u}^{-1} is defined as, $\mathbf{u}^{-1} = \mu^{-1}$ if $\mathbf{u} = \mu$, and $\mathbf{u}^{-1} = \mathbf{u}_1^{-1}, \dots, \mathbf{u}_k^{-1}$ if $\mathbf{u} = \overline{\mathbf{u}}_k$. Analogously, the composition of distributions of μ 's is defined as, $\mathbf{u} \circ \mathbf{u}' = \mu \circ \mu'$ if $\mathbf{u} = \mu$ and $\mathbf{u}' = \mu'$, and $\mathbf{u} \circ \mathbf{u}' = \mathbf{u}_1 \circ \mathbf{u}'_1, \dots, \mathbf{u}_k \circ \mathbf{u}'_k$ if $\mathbf{u} = \overline{\mathbf{u}}_k$ and $\mathbf{u}' = \overline{\mathbf{u}'_k}$.

Lemma 6.3 ($\lambda_{conc\mathcal{N}}^{sync}$: Properties of Bisimulation Relation).

- (Symmetry). If $(\mathbb{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_2, \mathbf{C}_2)$ then $(\mathbb{H}_2, \mathbf{C}_2) \sim_{\Phi\mathbb{L}\mu^{-1}\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_1, \mathbf{C}_1)$.
- (Reflexivity).
 - (Configurations). If $\mathbf{C} = (\mathbb{H}, \mathbf{T})_{\mathbf{t}}$ and $(\mathcal{H}, \mathbb{H}, \mathbf{T})_{\mathbf{t}} \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathcal{H}, \mathbb{H}, \mathbf{T})_{\mathbf{t}}$ then $(\mathcal{H}, \mathbf{C}) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathcal{H}, \mathbf{C})$.
 - (Distributions of Configurations). If $\mathbb{H} = \overline{\mathbb{H}}_k$, $\mathbf{C} = \overline{\mathbf{C}}_k$, $\mathbf{u} = \overline{\mathbf{u}}_k$, and $\forall i. 1 \leq i \leq k \implies (\mathbb{H}_i, \mathbf{C}_i) \sim_{\Phi\mathbb{L}\mu_i\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_i, \mathbf{C}_i)$, then $(\mathbb{H}, \mathbf{C}) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}, \mathbf{C})$.
- (Transitivity). If $(\mathbb{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_2, \mathbf{C}_2)$ and $(\mathbb{H}_2, \mathbf{C}_2) \sim_{\Phi\mathbb{L}\mu'\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_3, \mathbf{C}_3)$ then for $\mathbf{u}'' = \mathbf{u}' \circ \mathbf{u}$, $(\mathbb{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu''\mathbb{L}}^{\Gamma\Upsilon} (\mathbb{H}_3, \mathbf{C}_3)$.

Proof. 1. (Symmetry). By Definition 6.2 given Lemma 5.2[1b].

2. (Reflexivity). By Definition 6.2.

3. (Transitivity). By Definition 6.2 and Lemma 5.2[3b]. \square

The following single-step strong bisimulation lemma states: given a pair of bisimilar (distributions of) configurations, if the first one steps (all-paths-in-parallel), then the second one also steps (all-paths-in-parallel), and bisimilarity is preserved in the resulting (distributions of) configurations. For the most part the lemma follows directly from Lemma 5.3[2]; the only exception being the isomorphism of the traces in case 2 below, which is a direct consequence of low scheduler behavior.

We write $\mathbb{H}_1 \subseteq \mathbb{H}_2$ if either, for some \mathcal{H}_1 and \mathcal{H}_2 , $\mathbb{H}_1 = \mathcal{H}_1$, $\mathbb{H}_2 = \mathcal{H}_2$ and $\mathcal{H}_1 \subseteq \mathcal{H}_2$, or for some $\overline{\mathbb{H}}_{1k}$ and $\overline{\mathbb{H}}_{2k}$, $\mathbb{H}_1 = \overline{\mathbb{H}}_{1k}$, $\mathbb{H}_2 = \overline{\mathbb{H}}_{2k}$ and $\overline{\mathbb{H}}_{1k} \subseteq \overline{\mathbb{H}}_{2k}$. Analogously we write $\mathbf{u}_1 \subseteq \mathbf{u}_2$ if either, for some μ_1 and μ_2 , $\mathbf{u}_1 = \mu_1$, $\mathbf{u}_2 = \mu_2$ and $\mu_1 \subseteq \mu_2$, or for some $\overline{\mathbf{u}}_{1k}$ and $\overline{\mathbf{u}}_{2k}$, $\mathbf{u}_1 = \overline{\mathbf{u}}_{1k}$, $\mathbf{u}_2 = \overline{\mathbf{u}}_{2k}$ and $\overline{\mathbf{u}}_{1k} \subseteq \overline{\mathbf{u}}_{2k}$.

Lemma 6.4 ($\lambda_{conc\mathcal{N}}^{sync}$: Strong Bisimulation, 1-step).

- (Configurations). If $\mathbf{C}_1 \xrightarrow{\Phi_i} \mathbf{C}'_1$ and $(\mathcal{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\emptyset\Upsilon} (\mathcal{H}_2, \mathbf{C}_2)$ then for some \mathbf{C}'_2 , $\mathbf{C}_2 \xrightarrow{\Phi_i} \mathbf{C}'_2$ holds, and there exists a $\mathcal{H}'_1, \mathcal{H}'_2$ and μ' such that $\mathcal{H}_1 \subseteq \mathcal{H}'_1$, $\mathcal{H}_2 \subseteq \mathcal{H}'_2$, $\mu \subseteq \mu'$ and $(\mathcal{H}'_1, \mathbf{C}'_1) \sim_{\Phi\mathbb{L}\mu'\mathbb{L}}^{\emptyset\Upsilon} (\mathcal{H}'_2, \mathbf{C}'_2)$.
- (Distributions of Configurations). If $\mathbf{C}_1 \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}'_1$ and $(\mathbb{H}_1, \mathbf{C}_1) \sim_{\Phi\mathbb{L}\mu\mathbb{L}}^{\emptyset\Upsilon} (\mathbb{H}_2, \mathbf{C}_2)$ then for some \mathbf{C}'_2 , $\mathbf{C}_2 \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}'_2$ holds, and there exists a $\mathbb{H}'_1, \mathbb{H}'_2$ and \mathbf{u}' such that $\mathbb{H}_1 \subseteq \mathbb{H}'_1$, $\mathbb{H}_2 \subseteq \mathbb{H}'_2$, $\mathbf{u} \subseteq \mathbf{u}'$ and $(\mathbb{H}'_1, \mathbf{C}'_1) \sim_{\Phi\mathbb{L}\mu'\mathbb{L}}^{\emptyset\Upsilon} (\mathbb{H}'_2, \mathbf{C}'_2)$.

Proof. 1. (Configurations). By Definition 6.2[1] and Lemma 5.3[2].

2. (Distributions of Configurations). By induction on the derivation of $\mathbf{C}_1 \rightsquigarrow_{\Phi\mathbb{L}\mathcal{N}} \mathbf{C}'_1$. Following are the possible semantics rule from Figure 16 applicable at the root of this derivation.

- THREAD-POOL- \mathcal{N} -ALL-PATHS. For some $\mathbf{C}_1, \mathbb{H}_1, \mathbf{T}_1, \mathbf{t}_1$, $\overline{\zeta}_{1k}, \hat{\mathbf{H}}_{low}, \delta, \overline{i}_n$ and $\mathbf{C}_{1i_1}, \dots, \mathbf{C}_{1i_n}$, $\mathbf{C}_1 = \mathbf{C}_1$, $\mathbf{C}_1 = (\mathbb{H}_1, \mathbf{T}_1)_{\mathbf{t}_1}$, $\mathbf{T}_1 = \overline{\zeta}_{1k}$, $\hat{\mathbf{H}}_{low} = \mathbb{H}_1 \upharpoonright_{\mathbb{L}}$, $\mathcal{N}(k, \mathbf{t}_1, \hat{\mathbf{H}}_{low}) = \delta$,

$\delta = i_1, \dots, i_n, \forall i \in i_1, \dots, i_n. C_1 \xrightarrow{\Phi_i} C_{1_i}$ and $C'_1 = C_{1_{i_1}}, \dots, C_{1_{i_n}}$.

By premise we know $(\mathbb{H}_1, C_1) \sim_{\Phi \ell \mathbf{u} \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}_2, C_2)$; then knowing $C_1 = C_1$ by Definition 6.2[2] for some $C_2, \mathcal{H}_1, \mathcal{H}_2$ and $\mu, C_2 = C_2, \mathbb{H}_1 = \mathcal{H}_1, \mathbb{H}_2 = \mathcal{H}_2, \mathbf{u} = \mu$ and $(\mathcal{H}_1, C_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}_2, C_2)$. We know for all $i \in i_1, \dots, i_n, C_1 \xrightarrow{\Phi_i} C_{1_i}$; then by applying Lemma 6.4[1] to each of the aforementioned reductions we get, for all $i \in i_1, \dots, i_n$, there exists a C_{2_i} such that $C_2 \xrightarrow{\Phi_i} C_{2_i}$ and there exists a $\mathcal{H}_{1_i}, \mathcal{H}_{2_i}$ and μ_i such that $\mathcal{H}_1 \subseteq \mathcal{H}_{1_i}, \mathcal{H}_2 \subseteq \mathcal{H}_{2_i}, \mu \subseteq \mu_i$ and $(\mathcal{H}_{1_i}, C_{1_i}) \sim_{\Phi \ell \mu_i \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}_{2_i}, C_{2_i})$. Letting $\mathbb{H}'_1 = \mathcal{H}_{1_{i_1}}, \dots, \mathcal{H}_{1_{i_n}}, \mathbb{H}'_2 = \mathcal{H}_{2_{i_1}}, \dots, \mathcal{H}_{2_{i_n}}$ and $\mathbf{u}' = \mu_{i_1}, \dots, \mu_{i_n}$, by Definition 6.2[2] we get $(\mathbb{H}'_1, C_{1_{i_1}}, \dots, C_{1_{i_n}}) \sim_{\Phi \ell \mathbf{u}' \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_2, C_{2_{i_1}}, \dots, C_{2_{i_n}})$.

By premise we know $(\mathbb{H}_1, C_1) \sim_{\Phi \ell \mathbf{u} \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}_2, C_2)$, that is, $(\mathcal{H}_1, C_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}_2, C_2)$. We also know $C_1 = (H_1, T_1)_{\mathbf{t}_1}$. Then by Definition 6.2[1] for some H_2, T_2 and $\mathbf{t}_2, C_2 = (H_2, T_2)_{\mathbf{t}_2}$ and $(\mathcal{H}_1, H_1, T_1)_{\mathbf{t}_1} \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}_2, H_2, T_2)_{\mathbf{t}_2}$; by Definition 5.1[5,2] we get $(\mathcal{H}_1, H_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset} (\mathcal{H}_2, H_2), \mathbf{t}_1 = \mathbf{t}_2$ and for some $\zeta_{2_k}, T_2 = \zeta_{2_k}$. By Lemma 5.2[5] $H_1 \upharpoonright \mathbb{L} = H_2 \upharpoonright \mathbb{L} = \hat{H}_{low}$. Hence $\mathcal{N}(k, \mathbf{t}_2, \hat{H}_{low}) = \mathcal{N}(k, \mathbf{t}_1, \hat{H}_{low}) = \delta$. Then by $\text{THREAD-POOL-}\mathcal{N}_{\text{ALL-PATHS}}$ for $C'_2 = C_{2_{i_1}}, \dots, C_{2_{i_n}}, C_2 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}} C'_2$, that is, $C_2 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}} C'_2$ holds, the first requisite result.

We know $(\mathbb{H}'_1, C_{1_{i_1}}, \dots, C_{1_{i_n}}) \sim_{\Phi \ell \mathbf{u}' \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_2, C_{2_{i_1}}, \dots, C_{2_{i_n}})$, that is, $(\mathbb{H}'_1, C'_1) \sim_{\Phi \ell \mathbf{u}' \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_2, C'_2)$, the second requisite result.

- (b) **CONFIGURATION-DISTRIBUTION.** For some $\overline{C_{1_k}}$, and $C'_1 = C_{1_1}, \dots, C_{1_k}, \forall 1 \leq i \leq k. C_{1_i} \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}} C'_{1_i}$, and $C'_1 = C'_{1_1}, \dots, C'_{1_k}$.

By premise $(\mathbb{H}_1, C_1) \sim_{\Phi \ell \mathbf{u} \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}_2, C_2)$; then by Definition 6.2[2] for some $\overline{C_{2_k}}, \mathbb{H}_{1_k}, \mathbb{H}_{2_k}$, and $\overline{\mathbf{u}_k}, C_2 = \overline{C_{2_k}}, \mathbb{H}_1 = \mathbb{H}_{1_k}, \mathbb{H}_2 = \mathbb{H}_{2_k}, \mathbf{u} = \overline{\mathbf{u}_k}$, and for all i such that $1 \leq i \leq k, (\mathbb{H}_{1_i}, C_{1_i}) \sim_{\Phi \ell \mathbf{u}_i \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}_{2_i}, C_{2_i})$ holds.

By induction hypothesis for all i such that $1 \leq i \leq k$, for some $C'_{2_i}, C_{2_i} \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}} C'_{2_i}$, and there exists a $\mathbb{H}'_{1_i}, \mathbb{H}'_{2_i}$ and \mathbf{u}'_i such that $\mathbb{H}_{1_i} \subseteq \mathbb{H}'_{1_i}, \mathbb{H}_{2_i} \subseteq \mathbb{H}'_{2_i}, \mathbf{u}_i \subseteq \mathbf{u}'_i$ and $(\mathbb{H}'_{1_i}, C'_{1_i}) \sim_{\Phi \ell \mathbf{u}'_i \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_{2_i}, C'_{2_i})$.

By **CONFIGURATION-DISTRIBUTION**, letting $C'_2 = C'_{2_1}, \dots, C'_{2_k}, C_2 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}} C'_2$, the first requisite result. Then by Definition 6.2[2], for $\mathbb{H}'_1 = \mathbb{H}'_{1_1}, \dots, \mathbb{H}'_{1_k}, \mathbb{H}'_2 = \mathbb{H}'_{2_1}, \dots, \mathbb{H}'_{2_k}$ and $\mathbf{u}' = \mathbf{u}'_1, \dots, \mathbf{u}'_k, (\mathbb{H}'_1, C'_1) \sim_{\Phi \ell \mathbf{u}' \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_2, C'_2)$ holds, the second requisite result. \square

The following n -step strong bisimulation lemma then states: given a pair of bisimilar (distributions of) configurations, if the first one takes n steps (all-paths-in-parallel), then the second one also takes n steps (all-paths-in-parallel), and bisimilarity is preserved in the resulting (distributions of) configurations.

Lemma 6.5 ($\lambda_{conc \mathcal{N}}^{sync}$: Strong Bisimulation, n -steps). *If $C_1 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}}^n C'_1$ and $(\mathbb{H}_1, C_1) \sim_{\Phi \ell \mathbf{u} \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}_2, C_2)$ then for some $C'_2, C_2 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}}^n C'_2$ holds, and there exists a $\mathbb{H}'_1, \mathbb{H}'_2$ and \mathbf{u}' such that $\mathbb{H}_1 \subseteq \mathbb{H}'_1, \mathbb{H}_2 \subseteq \mathbb{H}'_2, \mathbf{u} \subseteq \mathbf{u}'$ and $(\mathbb{H}'_1, C'_1) \sim_{\Phi \ell \mathbf{u}' \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_2, C'_2)$.*

Proof. By induction on the derivation of $C_1 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}}^n C'_1$ given Lemma 6.4[2]. \square

The concatenation operator $@$ on $\hat{\Delta}$'s is defined as, $\hat{\Delta}_1 @ \hat{\Delta}_2 = \hat{\Delta}_3$ iff $\hat{\Delta}_1 = \hat{v}_j, \hat{\Delta}_2 = \hat{v}'_k$ and $\hat{\Delta}_3 = \hat{v}_1, \dots, \hat{v}_j, \hat{v}'_1, \dots, \hat{v}'_k$.

Definition 6.6 (Ground to Discrete Distribution of Values).

1. (Configuration). $[C]_i = \hat{v}$ if $C = (H, T)_{\mathbf{t}}$ and $T[i] = (\hat{v}, d)$, for some H, \mathbf{t} and d ; otherwise $[C]_i = \cdot$.
2. (Distribution of Configurations). $[C]_i = \hat{\Delta}$ iff $C = \overline{C_k}, \forall 1 \leq j \leq k. [C_j]_i = \hat{\Delta}_j$, and $\hat{\Delta} = \hat{\Delta}_1 @ \dots @ \hat{\Delta}_k$.

Lemma 6.7 ("Low" Int/Bool Type). *If $(\mathbb{H}_1, C_1) \sim_{\Phi \ell \mathbf{u} \mathbb{L}}^{\Gamma \Upsilon} (\mathbb{H}_2, C_2), \Upsilon[i] = \hat{\tau}$ selevel($\hat{\tau}$) $\leq \ell$ and $[C_1]_i = \hat{\Delta}$ then $[C_2]_i = \hat{\Delta}$.*

Proof. By Definition 6.2, Definition 5.1, Definition 4.3[1], Lemma 4.4[1c] and Definition 6.6. \square

The following defines the convergence of a thread to a distribution of values under a nondeterministic scheduler.

Definition 6.8 (Convergence under a Nondeterministic Scheduler). *A thread i in a configuration C , given a table of sync timer functions Φ , a set of heap locations \mathbb{L} and a nondeterministic scheduler \mathcal{N} , converges to a distribution of values $\hat{\Delta}$ in n steps iff $C \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}}^n C$, for some C , and $[C]_i = \hat{\Delta}$.*

The strong bisimulation lemma then entails the following property of probabilistic noninterference for $\lambda_{conc \mathcal{N}}^{sync}$. It states: if a thread in a concurrent configuration converges to a certain distribution of low integral or boolean values after a certain number of steps, then the same thread in that concurrent configuration, but possibly differing in high values from the first one, converges to the same distribution of values after the same number of steps as the first one.

Theorem 6.9 ($\lambda_{conc \mathcal{N}}^{sync}$: Probabilistic Noninterference). *If $\perp, \{x_k \mapsto \tau_k\}, \mathcal{H} \vdash_{\Phi \ell} (H, T) : \Upsilon, \Upsilon[i] = \hat{\tau}, \mathcal{H} \upharpoonright \mathbb{L} = \hat{\mathcal{H}}, \text{selevel}(\hat{\tau}), \text{selevel}(\hat{\mathcal{H}}) \leq \ell, \perp, \emptyset, \emptyset \vdash_{\Phi \ell}^{high} v_k, v'_k : \tau_k, (H_1, T_1) = (H, T)[v_k/x_k], (H_2, T_2) = (H, T)[v'_k/x_k]$ and the thread i in $(H_1, T_1)_{\epsilon}$ given Φ, \mathbb{L} and a nondeterministic scheduler \mathcal{N} converges to a distribution of values $\hat{\Delta}$ in n steps, then the thread i in $(H_2, T_2)_{\epsilon}$ given Φ, \mathbb{L} and \mathcal{N} converges to the same distribution of values $\hat{\Delta}$ in n steps.*

Proof. By premise for $C_1 = (H_1, T_1)_{\epsilon}$ and some $C'_1, C_1 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}}^n C'_1$ and $[C'_1]_i = \hat{\Delta}$. Let $\mathbb{L} = \{\overline{loc_k}\}$, for some loc_k , and $\mu = \{\overline{loc_k} \mapsto loc_k\}$; by Definition 5.1 and 4.3[1,3] $(\mathcal{H}, H_1, T_1)_{\epsilon} \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}, H_2, T_2)_{\epsilon}$. Then by Definition 6.2[1] for $C_2 = (H_2, T_2)_{\epsilon}, (\mathcal{H}, C_1) \sim_{\Phi \ell \mu \mathbb{L}}^{\emptyset \Upsilon} (\mathcal{H}, C_2)$. By Lemma 6.5 ($\lambda_{conc \mathcal{N}}^{sync}$: Strong Bisimulation, n -steps) for some $C'_2, C_2 \rightsquigarrow_{\Phi \mathbb{L} \mathcal{N}}^n C'_2$ holds, and there exists a $\mathbb{H}'_1, \mathbb{H}'_2$ and \mathbf{u}' such that $\mathcal{H}_1 \subseteq \mathbb{H}'_1, \mathcal{H}_2 \subseteq \mathbb{H}'_2, \mu \subseteq \mathbf{u}'$ and $(\mathbb{H}'_1, C'_1) \sim_{\Phi \ell \mathbf{u}' \mathbb{L}}^{\emptyset \Upsilon} (\mathbb{H}'_2, C'_2)$. Then by Lemma 6.7 $[C'_2]_i = \hat{\Delta}$. That is, the thread i in $(H_2, T_2)_{\epsilon}$ given Φ, \mathbb{L} and \mathcal{N} converges to the distribution of values $\hat{\Delta}$ in n steps, the requisite result. \square

Note the above result holds for all nondeterministic schedulers \mathcal{N} , that is, it is scheduler-independent [38].

7. Related Work

Since the seminal work of Volpano and Smith [44] much work has been done to secure timing channels in programming languages. However, all the efforts focus on static analyses; to our knowledge there is no language-based runtime approach to secure timing channels. Further, most existing works [43, 38, 47, 42, 34, 3] concentrate on closing the internal timing channels, leaving the external

ones vulnerable to attacks. Our dynamic synchronization technique secures both external and internal timing channels.

The program transformation approach of Agat [1] is the only known language-based technique to secure both external and internal timing channels. It equalizes the execution times of high branches by a program transformation which pads each branch with instructions that pretend to execute the other branch as well. Various extensions to Agat’s approach have been proposed – [38] employed a minor variant to secure internal timing channels in the context of multi-threaded programs (using skip instructions and dummy forks instead of the “skip-commands”, `skipAsn` and `skipIf`, used by Agat), [10] augmented Agat’s work towards a bytecode-like language, while [46] adapted it to secure typed assembly languages [27]. Recently a variant based on unification, which generates programs smaller in size, has been proposed [19].

The aforementioned works, however, have limited expressiveness in that they do not appear to generalize to incorporate higher-order functions – the exact target code of a branching statement is assumed to be known statically in these works, whereas higher-order function invocations are dynamic in that the exact code to be invoked is only known at runtime. Furthermore, none except [10] allow looping executions on high guards, that is, executions whose running times may depend on secret data – the execution times of such recursive computations are inherently tied to secret information (the high terminating condition) and it does not seem possible to secure them via static approaches without imposing severe restrictions such as the disallowance of subsequent low outputs [10]. We support both fully higher-order function invocations as well as recursive computations with high terminating conditions.

On the other hand, our dynamic synchronization technique endures an overhead of computing the sync times at runtime, and then handling the setting and expiration of the synchronization timers. In practice a hybrid approach may be best: use Agat’s technique wherever applicable, and use our dynamic synchronization technique to secure cases the former cannot handle.

A recent work by Barthe *et al.* [4] proposed a variant of Agat’s cross-copying technique based on a transaction mechanism. Each branch is wrapped in a transaction and then sequentially composed with the other corresponding branch. The transaction around the original branch is committed, while the cross-copied transaction is aborted. This technique is applicable to object-oriented languages with exceptions and methods calls. However it suffers from the same limitations as Agat’s cross-copying approach, in that neither high-recursive programs nor higher-order function applications (dynamic dispatch in their object-oriented context) are supported. Additionally, unlike Agat’s, this work suffers from a runtime overhead of committing and aborting transactions. Neither it is clear if it is applicable in a concurrent setting.

Molnar *et al.* [26] proposed methods to secure C programs against side-channel attacks. It performs a C-to-C program transformation to ensure both branches of conditional branching statements are executed but only results from the branch that would have been executed in the original program are retained through logical masking. The approach has been formalized for a simple imperative language called IncredibL, with only bounded loops (loops that can be fully unrolled statically), if statements and straight-line assignments. It is unclear if the proposed techniques would generalize to a richer language.

Neilson and Schwartzbach [28] presented a domain-specific imperative programming language called SMCL for Secure Multi-party Computation (SMC). To prevent timing leaks its runtime semantics evaluates both branches of if statements with a secret conditionals in sequence and on copies of the local state. After these executions, the results of both branches are merged such that only those of the valid branch are retained. It does not allow loops on

secret conditionals nor calls to recursive functions guarded by high conditions.

Kobayashi and Shirane [15] proposed a Java virtual machine language based type system for timing-sensitive noninterference. The system relies on the computation of control dependency for identifying sensitive regions, and closes timing channels in sequential code by inserting a delay linear with respect to the normal execution time. It does not support objects or subroutines.

On a different note, Siveroni *et al.* [40, 31] apply a padding technique, similar to but more efficient than Agat’s, to transform two processes to have indistinguishable timing behavior. The work is not programming language specific, but based probabilistic transition systems (PTS) [12]. In principle it is adaptable to any probabilistic language whose semantics can be expressed in terms of a PTS model, as for example our $\lambda_{conc_{\mathcal{N}}}^{sync}$ language.

8. Towards a Realistic System

Language Expressiveness The λ_{seq}^{sync} language does not support exceptions or interactive inputs and outputs. Existing works [7, 4] have studied exceptions in the context of timing channels and have shown disallowing exceptions thrown in high contexts from escaping to low contexts prevents timing leaks. λ_{seq}^{sync} can be extended to incorporate exceptions using the same restriction – the synchronization construct delineates the enclosed high context from the enclosing low context, so exceptions must be disallowed from escaping the synchronization construct; exceptions not handled inside the synchronization context must be ignored. Interactive inputs and outputs can be incorporated using a similar philosophy; low inputs and outputs must both be disallowed in high contexts – requests for low inputs in high contexts must be fulfilled with dummy values instead, and low outputs in high contexts must be omitted, as in [21].

The $\lambda_{conc_{\mathcal{D}}}^{sync}$ and $\lambda_{conc_{\mathcal{N}}}^{sync}$ languages do not support dynamic thread creation. Existing works [38, 34] have studied dynamic thread creation in the context of timing channels – [38] proposed cross-copying the fork instructions in high branching statements as dummy forks, along the lines of Agat’s technique, to secure timing channels; while [34] suggested partitioning the thread pool into low and high parts, such that threads created in high contexts are “hidden” in the high part, while the low part contains threads generated in low contexts; the timings of the high threads are hidden, via the scheduler, from those of the low threads. We believe $\lambda_{conc_{\mathcal{D}}}^{sync}$ and $\lambda_{conc_{\mathcal{N}}}^{sync}$ can be extended to incorporate forking following the thread pool partitioning technique [34].

Sabelfeld [35] studied the effects of synchronization primitives such as semaphores on timing channels, and concluded all synchronizations on high semaphores or in high contexts are susceptible to information leaks; $\lambda_{conc_{\mathcal{D}}}^{sync}$ and $\lambda_{conc_{\mathcal{N}}}^{sync}$ can be extended to incorporate synchronization along the same lines. Further we believe our techniques are adaptable to distributed programs as well along the lines of [23, 36].

Our formal model does not support declassification of secure data, a desirable feature in practice. One interesting aspect of our runtime technique is the parameterization of the sync times on low values; when coupled with declassification this parameterization can provide a mechanism to declassify *just* the timing channels. For example, if one or more parameters to a sync timer function (the \bar{i}_k of IF-SYNC and APP-SYNC rules in Figure 8) is declassified, only the associated timing channel has in effect been allowed to disseminate high information, and if there are no attackers observing that timing channel the high data remains secret.

Implementation of Sync Timers Interval timers provided by modern operating systems can be used to realize our dynamic synchronization technique. On Unix the `setitimer` function, included

in `<sys/time.h>`, provides a mechanism for a process to interrupt itself in the future by setting a timer; the process receives a signal when the timer expires. `ITIMER_VIRTUAL` is an interval timer which counts the actual execution time of a process (its virtual time), and sends a `sigvtalrm` signal to the process when it expires. The virtual timer can be used to directly realize our dynamic synchronization technique: a compiler or a preprocessor can instrument the source program such that just before executing the code to be synchronized, `ITIMER_VIRTUAL` is set to the sync time indicated by the associated sync timer function. The `sigvtalrm` handler can then realize the `SYNC-DONE` and `SYNC-ABORT` rules of Figure 8. To be more precise, consider the following λ_{seq}^{sync} program fragment,

$$\text{let } y = (\text{if}_{p, \overline{v_k}}^v p \text{ then } p_1 \text{ else } p_2) \text{ in } \dots \quad (5)$$

where the guard p is typed high, implying the computation of branching statement p must be synchronized. The following Caml [22] program fragment implements our dynamic synchronization technique for (5):

```
let f = fun _ -> failwith "timeout" in
Sys.set_signal (Sys.sigvtalrm) (Sys.Signal_handle f);
...
let h = p in let x = ref v in let sync_time =  $\Phi(p)(\overline{v_k})$  in
try
  Unix.setitimer Unix.ITIMER_VIRTUAL sync_time;
  x := if h then p1 else p2;
  while true do () done (* analogue of SYNC-PAD *)
with
  Failure "timeout" -> ();
let y = !x in ...
```

Note we have empirically validated the above program fragment in OCaml 3.10.1 on Linux. Caml libraries provide the `Unix.setitimer` system call which sets the virtual timer `ITIMER_VIRTUAL` to `sync_time`, and the `Sys.set_signal` function which installs a signal handler f for the signal `sigvtalrm` generated when `ITIMER_VIRTUAL` expires. The nonterminating loop ‘while true do () done’ captures the semantics of the `SYNC-PAD` rule in Figure 8 – if the computation of the preceding branching statement terminates before the timer expires the while-loop pads out the residual sync time. Observe the signal handler f for `sigvtalrm` simply throws the exception `Failure "timeout"` when the sync timer expires; the exception is caught and discarded by the `try-catch` statement, and the control is then transferred to the next statement ‘let y = !x in ...’. Also notice the computed value of the branching statement is assigned to the same heap location, referenced by the variable x , as the default placeholder value v ; in effect the value of $!x$ flowing into y is the value of the branch computation if the sync timer expired after the assignment ‘ $x := \text{if } h \text{ then } p_1 \text{ else } p_2$ ’ had happened; otherwise the default placeholder value flows into y .

Note the transformation illustrated above, from program fragment (5) to the program fragment (6), is generic and applicable to function application statements as well; further it can be done automatically by a preprocessor, and to target languages other than Caml (in C, for example, `siglongjmp()` can be used to perform the nonlocal transfer of control via the signal handler, as opposed to an exception in the above illustration). Also the above implementation denotes only one of the possibly many ways to realize our dynamic synchronization technique; the above method has the advantage of using mechanisms readily available on existing systems. It may be possible to devise more efficient techniques, say for example without the overhead of systems calls and/or context-switching; the exploration of such mechanisms is a direction for future investigation.

The resolution of interval timers provided by the `setitimer` function is in microseconds; if deemed too coarse, the `POSIX.1b` high resolution timers (HRTs) provide higher resolution timers in

nanoseconds, and can be used instead. However a resolution in the order of nanoseconds is unlikely to impart better efficiency given the overhead of setting up the interval timers along with the subsequent handling of the timer interrupts, both of which are system calls involving context-switching, is nonetheless likely to be on the order of microseconds on modern machines. In fact, given the overhead of runtime synchronization it may be too expensive to individually synchronize each branching statement of a high guard, especially if it is in a tight loop. Coarser granularities for synchronization could be used to amortize the associated overhead; for example, it may be more efficient to synchronize the execution of an entire loop as opposed to each individual branching statement inside it each time around. The exploration of such optimizations is an important direction for future research.

In typical multiprocessing environments there is likely to be some delay in the arrival of the timer interrupt after the corresponding timer has expired, depending on the system load, CPU pipelining, and other factors. However such delays are likely to be “random”, that is, likely governed by factors extraneous to the process under synchronization, thus independent of the values, both high and low, local to the process. Such potential delays are, therefore, unlikely to leak information; however a closer scrutiny is desirable to confirm the lack of useful information in such delays.

Securing Resource Channels Time is a form of resource. Recall our philosophy behind securing timing channels is to fix the execution times of high computations based on low values, and if a high computation does not finish in the allotted time then it is aborted and replaced with a placeholder value. A similar philosophy can be applied to other forms of resource channels as well.

Memory usage is a form of resource channel – the memory consumption pattern of a program may depend on secret data, and that could be used by the attacker to gain classified information. For example, one branch of a high branching statement may allocate a lot more memory than the other branch; an attacker could then gain information on the secret guard by tracking the memory usage of that program. One potential approach to securing such memory usage channels is to preallocate a fixed amount of memory for use by the high portion of a program computation; if the high computation exhausts the preallocated memory then the bluffing technique takes over, and all further requests for memory by the high computation are satisfied by recycling memory from the preallocated space. Analogous to sync times, the amount of preallocated high memory may be parameterized over low values.

The technique sketched above for securing memory usage channels can be made robust in presence of garbage collection as well – the high memory is clearly demarcated, hence the garbage collector can be restricted to spend a fixed amount of time to clean the high memory regardless of the amount of garbage in it; as above, the bluffing technique provides a fallback in case the high memory runs out of space or the collection time-limit is exceeded.

Other forms of resource channel such as those associated with utilization of disc space, network sockets, etc. can be secured in a similar manner: by fixing the usage of high computation. The bluffing technique is critical to the soundness of this approach, providing a fallback when such resource fixtures are found to be inadequate.

Guarding against Cache Attacks Consider the following program, adapted from [2],

$$\begin{aligned} z &:= x; \\ \text{if } (h == 1) &\text{ then } z := x \text{ else } z := y; \\ \text{output}^{\text{low}} &(\text{"done"}); \end{aligned} \quad (7)$$

where h holds a single high bit. Assuming neither x nor y is in the cpu cache before the execution of this code, the running time of the branching statement is likely to reveal the value of

h – the value of x will likely be cached after the first assignment $z := x$; hence the later assignment $z := x$ in the then-branch will probably run a bit faster than the assignment $z := y$ in the else-branch. This variance in the running time can be used by an attacker to gain information on h . Such information leaks induced by variance in cache behavior are referred to as *cache leaks*, and the corresponding attacks are called *cache attacks*. The capacity of covert timing channels implemented through cache leaks is not likely to be very high but it is certainly high enough to warrant attention, as shown in [2] via a simple experiment; since then many attacks exploiting the timing variability due to cache effects have been demonstrated [30, 29].

Our dynamic synchronization technique equalizes the running times of corresponding branches regardless of the state of the cache; hence, it would be robust against cache attacks for program (7). However, our technique does *not* synchronize the caches at the end of different executions of a high branching statement. Consider the following variant of the above program,

```

if ( $h == 1$ ) then  $z := x$  else  $z := y$ ;
 $z := x$ ;
outputlow("done");

```

(8)

Again assuming neither x nor y is in the cache before the execution of (8), if the then-branch is taken, x is likely to be in the cache, and y likely not, at the end the branching statement; the opposite holds if the else-branch were taken. Consequently, the subsequent assignment $z := x$ may run a bit faster in the case of then-branch compared to the else-branch, revealing the value of h via the timing of the low output “done”.

One simple mitigation strategy would be to flush the entire cache at the end of each synchronized run of a branch, wiping out any possible correlation between the high guard and cached values. The cache can be flushed by reading enough garbage into it [2].

9. Conclusion

We have demonstrated how timing channels can be provably eliminated in both sequential and concurrent programs, in presence of deterministic as well as nondeterministic schedulers. Our techniques advance the current state of the art by showing how timing channels can be eliminated in programs with higher-order functions, with recursive computations under high guards, and in the presence of libraries for which the source code is not available. We have illustrated how our dynamic synchronization technique can be realized on standard computing platforms using interval timers. Furthermore, we believe the technique of bluff computation holds promise in securing resource channels as well. Our solution is general and can be incorporated into arbitrary software products.

The paper makes several technical contributions. Our strong bisimulation proof technique is a general technique for proving noninterference in presence of timing channels. Our method of employing discrete probability distributions, and the associated all-paths-in-parallel operational semantics, allows for a more direct proof of probabilistic noninterference in concurrent programs. The notion of bluff computation generalizes the lenient execution model of [7] to be more broadly applicable.

References

- [1] J. Agat. Transforming out timing leaks. In *POPL*, 2000. 1, 3, 4, 22
- [2] J. Agat and D. Sands. On confidentiality and algorithms. In *SP*, 2001. 1, 4, 23, 24
- [3] G. Barthe, T. Rezk, A. Russo, and A. Sabelfeld. Security of multithreaded programs by compilation. In *ESORICS*, 2007. 1, 21
- [4] G. Barthe, T. Rezk, and M. Warnier. Preventing timing leaks through transactional branching instructions. In *QAPL*, 2005. 22
- [5] G. Boudol and I. Castellani. Noninterference for concurrent programs. In *ICALP*, 2001. 1, 4
- [6] D. Brumley and D. Boneh. Remote timing attacks are practical. In *USENIX Security Symposium*, 2003. 1
- [7] Z. Deng and G. Smith. Lenient array operations for practical secure information flow. In *CSFW*, 2004. 3, 22, 24
- [8] D. E. R. Denning. *Cryptography and Data Security*. 1982. 6
- [9] J. A. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and Privacy*, 1982. 1, 10
- [10] D. Hedin and D. Sands. Timing aware information flow security for a javacard-like bytecode. In *BYTECODE*, 2005. 1, 22
- [11] W.-M. Hu. Reducing timing channels with fuzzy time. In *IEEE Symposium Research in Security and Privacy*, 1991. 1
- [12] B. Jonsson, K. Larsen, and W. Yi. Probabilistic extensions of process algebras. In *Handbook of Process Algebra*, 2001. 22
- [13] M. H. Kang and I. S. Moskowitz. A pump for rapid, reliable, secure communication. In *CCS*, 1993. 1
- [14] J. Kelsey, B. Schneier, D. Wagner, and C. Hall. Side channel cryptanalysis of product ciphers. In *ESORICS*, 1998. 1
- [15] N. Kobayashi and K. Shirane. Type-based information flow analysis for a low-level language. In *APLAS*, 2002. 22
- [16] P. Kocher, J. Jaffe, and B. Jun. Using unpredictable information to minimize leakage from smartcards and other cryptosystems. In *USA patent, International Publication number WO 99/63696*, 1999. 1
- [17] P. C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *CRYPTO*, 1996. 1
- [18] B. Köpf and D. Basin. Timing-sensitive information flow analysis for synchronous systems. In *ESORICS*, 2006. 1
- [19] B. Köpf and H. Mantel. Eliminating implicit information leaks by transformational typing and unification. In *FAST*, 2005. 22
- [20] B. W. Lampson. A note on the confinement problem. *Communications of the ACM*, 16(10):613–615, 1973. 1
- [21] G. Le Guernic, A. Banerjee, T. Jensen, and D. A. Schmidt. Automata-based confidentiality monitoring. In *ASIAN*, 2006. 22
- [22] X. Leroy, D. Doligez, J. Garrigue, D. Rémy, and J. Vouillon. The Caml Language. <http://caml.inria.fr/>. 2, 23
- [23] H. Mantel and A. Sabelfeld. A unifying approach to the security of distributed and multi-threaded programs. *J. Comput. Secur.*, 2003. 22
- [24] D. May, H. L. Muller, and N. P. Smart. Non-deterministic processors. In *ACISP '01: Proceedings of the 6th Australasian Conference on Information Security and Privacy*, pages 115–129, London, UK, 2001. Springer-Verlag. 1
- [25] D. May, H. L. Muller, and N. P. Smart. Random register renaming to foil dpa. In *CHES '01: Proceedings of the Third International Workshop on Cryptographic Hardware and Embedded Systems*, 2001. 1
- [26] D. Molnar, M. Piotrowski, D. Schultz, and D. Wagner. The program counter security model: Automatic detection and removal of control-flow side channel attacks. In *ICISC*, 2005. 22
- [27] G. Morrisett, D. Walker, K. Crary, and N. Glew. From system f to typed assembly language. In *POPL '98: Symposium on Principles of programming languages*, 1998. 22
- [28] J. D. Nielsen and M. I. Schwartzbach. A domain-specific programming language for secure multiparty computation. In *PLAS*, 2007. 22
- [29] D. Osvik, A. Shamir, and E. Tromer. Cache attacks and countermeasures: the case of aes, 2005. 24
- [30] D. Page. Defending against cache based side-channel attacks. *Information Security Technical Report*, 2003. 24

- [31] A. D. Pierro, C. Hankin, I. Siveroni, and H. Wiklicky. Tempus fugit: How to plug it. *The Journal of Logic and Algebraic Programming* 72 (2007) 173190, 2007. 22
- [32] F. Pottier and V. Simonet. Information flow inference for ML. *TOPLAS*, 2003. 6, 7
- [33] R. L. Rivest, A. Shamir, and L. M. Adelman. A method for obtaining digital signatures and public-key cryptosystems. In *Communications of the ACM*, 1978. 2
- [34] A. Russo and A. Sabelfeld. Securing interaction between threads and the scheduler. In *CSFW*, 2006. 1, 17, 21, 22
- [35] A. Sabelfeld. The impact of synchronisation on secure information flow in concurrent programs. In *PSI '02: Perspectives of System Informatics*, 2001. 22
- [36] A. Sabelfeld and H. Mantel. Static confidentiality enforcement for distributed programs. In *SAS'02: Symposium on Static Analysis*, 2002. 22
- [37] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 2003. 1
- [38] A. Sabelfeld and D. Sands. Probabilistic noninterference for multi-threaded programs. In *CSFW*, 2000. 1, 17, 19, 20, 21, 22
- [39] P. Shroff, S. F. Smith, and M. Thober. Dynamic dependency monitoring to secure information flow. In *CSF*, 2007. 6
- [40] I. Siveroni. Filling out the gaps: A padding algorithm for transforming out timing leaks. *Proceedings of the Third Workshop on Quantitative Aspects of Programming Languages (QAPL 2005)*, 2005. 22
- [41] G. Smith. A new type system for secure information flow. In *CSFW*, 2001. 1, 4
- [42] G. Smith. Improved typings for probabilistic noninterference in a multi-threaded language. *Journal of Computer Security*, 2006. 21
- [43] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *POPL*, 1998. 1, 21
- [44] Volpano and Smith. Eliminating covert flows with minimum typings. In *CSFW*, 1997. 21
- [45] R. Wilhelm et al. The worst-case execution time problem – overview of methods and survey of tools. *TECS*, 2007. 3
- [46] D. Yu and N. Islam. A typed assembly language for confidentiality. In *APLAS*, 2007. 22
- [47] S. Zdancewic and A. Myers. Observational determinism for concurrent program security. *CSFW*, 2003. 1, 4, 21
- [48] Y. Zhou and D. Feng. Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing. In *Cryptology ePrint Archive, Report 2005/388*, 2005. 1