



# Foundations for Parallel Information Flow Control Runtime Systems

Marco Vassena<sup>1</sup>(✉), Gary Soeller<sup>2</sup>, Peter Amidon<sup>2</sup>, Matthew Chan<sup>3</sup>,  
John Renner<sup>2</sup>, and Deian Stefan<sup>2</sup>(✉)

<sup>1</sup> Chalmers University, Gothenburg, Sweden

`vassena@chalmers.se`

<sup>2</sup> UC San Diego, San Diego, USA

`deian@cs.ucsd.edu`

<sup>3</sup> Awake Security, Sunnyvale, USA

**Abstract.** We present the foundations for a new dynamic information flow control (IFC) parallel runtime system, LIO<sub>PAR</sub>. To our knowledge, LIO<sub>PAR</sub> is the first dynamic language-level IFC system to (1) support deterministic parallel thread execution and (2) eliminate both internal- and external-timing covert channels that exploit the runtime system. Most existing IFC systems are vulnerable to external timing attacks because they are built atop vanilla runtime systems that do not account for security—these runtime systems allocate and reclaim shared resources (e.g., CPU-time and memory) *fairly* between threads at different security levels. While such attacks have largely been ignored—or, at best, mitigated—we demonstrate that extending IFC systems with parallelism leads to the *internalization* of these attacks. Our IFC runtime system design addresses these concerns by hierarchically managing resources—both CPU-time and memory—and making resource allocation and reclamation explicit at the language-level. We prove that LIO<sub>PAR</sub> is secure, i.e., it satisfies progress- and timing-sensitive non-interference, even when exposing clock and heap-statistics APIs.

## 1 Introduction

Language-level dynamic information flow control (IFC) is a promising approach to building secure software systems. With IFC, developers specify application-specific, data-dependent security policies. The language-level IFC system—often implemented as a library or as part of a language runtime system—then enforces these policies automatically, by tracking and restricting the flow of information throughout the application. In doing so, IFC can ensure that different application components—even when buggy or malicious—cannot violate data confidentiality or integrity.

---

This work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation program sponsored by DARPA and by gifts from Cisco and Fujitsu. This work was partly done while Marco Vassena and Matthew Chan were at UCSD.

© The Author(s) 2019

F. Nielson and D. Sands (Eds.): POST 2019, LNCS 11426, pp. 1–28, 2019.

[https://doi.org/10.1007/978-3-030-17138-4\\_1](https://doi.org/10.1007/978-3-030-17138-4_1)

The key to making language-level IFC practical lies in designing real-world programming language features and abstractions without giving up on security. Unfortunately, many practical language features are at odds with security. For example, even exposing language features as simple as `if`-statements can expose users to timing attacks [42, 64]. Researchers have made significant strides towards addressing these challenges—many IFC systems now support real-world features and abstractions safely [10, 15, 20, 34, 43, 50, 51, 54, 55, 59, 60, 62, 67, 68]. To the best of our knowledge, though, no existing language-level dynamic IFC supports parallelism. Yet, many applications rely on parallel thread execution. For example, modern Web applications typically handle user requests in parallel, on multiple CPU cores, taking advantage of modern hardware. Web applications built atop state-of-the-art dynamic IFC Web frameworks (e.g., Jacqueline [67], Hails [12, 13], and LMonad [45]), unfortunately, do not handle user requests in parallel—the language-level IFC systems that underlie them (e.g., Jeeves [68] and LIO [54]) do not support parallel thread execution.

In this paper we show that extending most existing IFC systems—even concurrent IFC systems such as LIO—with parallelism is unsafe. The key insight is that most IFC systems *do not* prevent sensitive computations from affecting public computations; they simply prevent public computations from *observing* such sensitive effects. In the sequential and concurrent setting, such effects are only observable to attackers *external* to the program and thus outside the scope of most IFC systems. However, when computations execute in parallel, they are essentially external to one another and thus do not require an observer external to the system—they can observe such effects internally.

Consider a program consisting of three concurrent threads: two public threads— $p_0$  and  $p_1$ —and a secret thread— $s_0$ . On a single core, language-level IFC can ensure that  $p_0$  and  $p_1$  do not learn anything secret by, for example, disallowing them from observing the return values (or lack thereof) of the secret thread. Systems such as LIO are careful to ensure that public threads cannot learn secrets even indirectly (e.g., via covert channels that abuse the runtime system scheduler). But, secret threads *can* leak information to an external observer that monitors public events (e.g., messages from public threads) by influencing the behavior of the public threads. For example,  $s_0$  can terminate (or not) based on a secret and thus affect the amount of time  $p_0$  and  $p_1$  spend executing on the CPU—if  $s_0$  terminated, the runtime allots the whole CPU to public threads, otherwise it only allots, say, two thirds of the CPU to the public threads; this allows an external attacker to trivially infer the secret (e.g., by measuring the rate of messages written to a public channel). Unfortunately, such *external timing attacks* manifest *internally* to the program when threads execute in parallel, on multiple cores. Suppose, for example, that  $p_0$  and  $s_0$  are co-located on a core and run in parallel to  $p_1$ . By terminating early (or not) based on a secret,  $s_0$  affects the CPU time allotted to  $p_0$ , which can be measured by  $p_1$ . For example,  $p_1$  can count the number of messages sent from  $p_0$  on a public channel—the number of  $p_0$  writes indirectly leaks whether or not  $s_0$  terminated.

We demonstrate that such attacks are feasible by building several proof-of-concept programs that exploit the way the runtime system allocates and reclaims *shared* resources to violate LIO’s security guarantees. Then, we design a new dynamic parallel language-level IFC runtime system called LIO<sub>PAR</sub>, which extends LIO to the parallel setting by changing how *shared* runtime system resources—namely CPU-time and memory—are managed. Ordinary runtime systems (e.g., GHC for LIO) *fairly* balance resources between threads; this means that allocations or reclamations for secret LIO threads directly affect resources available for public LIO threads. In contrast, LIO<sub>PAR</sub> makes resource management *explicit* and *hierarchical*. When allocating new resources on behalf of a thread, the LIO<sub>PAR</sub> runtime does not “fairly” steal resources from all threads. Instead, LIO<sub>PAR</sub> demands that the thread requesting the allocation explicitly gives up a portion of its own resources. Similarly, the runtime does not automatically relinquish the resources of a terminated thread—it requires the parent thread to explicitly reclaim them.

Nevertheless, automatic memory management is an integral component of modern language runtimes—high-level languages (e.g., Haskell and thus LIO) are typically garbage collected, relieving developers from manually reclaiming unused memory. Unfortunately, even if memory is hierarchically partitioned, some garbage collection (GC) algorithms, such as GHC’s stop-the-world GC, may introduce timing covert channels [46]. Inspired by previous work on real-time GCs (e.g., [3, 5, 6, 16, 44, 48]), we equip LIO<sub>PAR</sub> with a per-thread, interruptible garbage collector. This strategy is agnostic to the particular GC algorithm used: our hierarchical runtime system only demands that the GC runs within the memory confines of individual threads and their time budget.

In sum, this paper makes three contributions:

- We observe that several external timing attacks *manifest internally* in the presence of parallelism and demonstrate that LIO, when compiled to run on multiple cores, is vulnerable to such attacks (Sect. 2).
- In response to these attacks, we propose a novel parallel runtime system design that safely manages shared resources by enforcing explicit and *hierarchical* resource allocation and reclamation (Sect. 3). To our knowledge, LIO<sub>PAR</sub> is the first parallel language-level dynamic IFC runtime system to address both internal and external timing attacks that abuse the runtime system scheduler, memory allocator, and GC.
- We formalize the LIO<sub>PAR</sub> hierarchical runtime system (Sect. 4) and prove that it satisfies *progress- and timing-sensitive non-interference* (Sect. 5); we believe that this is the first general purpose dynamic IFC runtime system to provide such strong guarantees in the parallel setting [64].

Neither our attack nor our defense is tied to LIO or GHC—we focus on LIO because it already supports concurrency. We believe that extending any existing language-level IFC system with parallelism will pose the same set of challenges—challenges that can be addressed using explicit and hierarchical resource management.

## 2 Internal Manifestation of External Attacks

In this section we give a brief overview of LIO and discuss the implications of shared, finite runtime system resources on security. We demonstrate several external timing attacks against LIO that abuse two such resources—the thread scheduler and garbage collector—and show how running LIO threads in parallel internalizes these attacks.

### 2.1 Overview of the Concurrent LIO Information Flow Control System

At a high level, the goal of an IFC system is to track and restrict the flow of information according to a security policy—almost always a form of *non-interference* [14]. Informally, this policy ensures *confidentiality*, i.e., secret data should not leak to public entities, and *integrity*, i.e., untrusted data should not affect trusted entities.

To this end, LIO tracks the flow of information at a coarse-granularity, by associating *labels* with threads. Implicitly, the thread label classifies all the values in its scope and reflects the sensitivity of the data that it has inspected. Indeed, LIO “raises” the label of a thread to accommodate for reading yet more sensitive data. For example, when a *public* thread reads *secret* data, its label is raised to *secret*—this reflects the fact that the rest of the thread computation may depend on sensitive data. Accordingly, LIO uses the thread’s *current label* or *program counter label* to restrict its communication. For example, a *secret* thread can only communicate with other *secret* threads.

In LIO, developers can express programs that manipulate data of varying sensitivity—for example programs that handle both *public* and *secret* data—by forking multiple threads, at run-time, as necessary. However, naively implementing concurrency in an IFC setting is dangerous: concurrency can amplify and internalize the *termination covert channel* [1, 58], for example, by allowing public threads to observe whether or not secret threads terminated. Moreover, concurrency often introduces *internal timing covert channels* wherein secret threads leak information by influencing the scheduling behavior of public threads. Both classes of covert channels are high-bandwidth and easy to exploit.

Stefan et al. [54] were careful to ensure that LIO does not expose these termination and timing covert channels *internally*. LIO ensures that even if secret threads terminate early, loop forever, or otherwise influence the runtime system scheduler, they cannot leak information to public threads. But, secret threads *do* affect public threads with those actions and thus expose timing covert channels *externally*—public threads just cannot detect it. In particular, LIO disallows public threads from (1) directly inspecting the return values (and thus timing and termination behavior) of secret threads, without first raising their program counter label, and (2) observing runtime system resource usage (e.g., elapsed time or memory availability) that would indirectly leak secrets.

LIO prevents public threads from measuring CPU-time usage directly—LIO does not expose a clock API—and indirectly—threads are scheduled fairly

in a round-robin fashion [54]. Similarly, LIO prevents threads from measuring memory usage directly—LIO does not expose APIs for querying heap statistics—and indirectly, through garbage collection cycles (e.g., induced by secret threads) [46]—GHC’s stop-the-world GC stops all threads. Like other IFC systems, the security guarantees of LIO are weaker in practice because its formal model does not account for the GC and assumes memory to be infinite [54, 55].

## 2.2 External Timing Attacks to Runtime Systems

Since secret threads can still influence public threads by abusing the scheduler and GC, LIO is vulnerable to *external timing and termination attacks*, i.e., attacks that leak information to external observers. To illustrate this, we craft several LIO programs consisting of two threads: a public thread  $p$  that writes to the external channel observed by the attacker and a secret thread  $s$ , which abuses the runtime to influence the throughput of the public thread. The secret thread can leak in many ways, for example, thread  $s$  can:

1. *fork bomb*, i.e., fork thousands of secret threads that will be interleaved with  $p$  and thus decrease its write throughput;
2. terminate early to relinquish the CPU to  $p$  and thus double its write throughput;
3. exhaust all memory to crash the program, and thus stop  $p$  from further writing to the channel;
4. force a garbage collection which, because of GHC’s stop-the-world GC, will intermittently stop  $p$  from writing to the channel.

These attacks abuse the runtime’s automatic *allocation* and *reclamation* of shared resources, i.e., CPU time and memory. In particular, attack 1 hinges on the runtime *allocating* CPU time for the new secret threads, thus reducing the CPU time allotted to the public thread. Dually, attack 2 relies on it *reclaiming* the CPU time of terminated threads—it reassigns it to public threads. Similarly, attacks 3 and 4 force the runtime to allocate all the available memory and preemptively reassign CPU time to the GC, respectively.

These attacks are not surprising, but, with the exception of the GC-based attack [46], they are novel in the IFC context. Moreover these attacks are not exhaustive—there are other ways to exploit the runtime system—nor optimized—our implementation leaks sensitive data at a rate of roughly 2bits/second<sup>1</sup>. Nevertheless, they are feasible and—because they abuse the runtime—they are effective against language-level external-timing mitigation techniques, including [54, 71]. The attacks are also feasible on other systems—similar attacks that abuse the GC have been demonstrated for both the V8 and JVM runtimes [46].

---

<sup>1</sup> A more assiduous attacker could craft similar attacks that leak at higher bit-rates.

Core $c_0$		Core $c_1$
Secret Thread ( $s_0$ )	Public Thread ( $p_0$ )	Public Thread ( $p_1$ )
<pre> if secret   then terminate   else forever skip </pre>	<pre> forever   (write chan <math>p_0</math>) </pre>	<pre> for [1..n] (write chan <math>p_1</math>) ms &lt;- read chan <math>w_0</math> &lt;- count <math>p_0</math> ms <math>w_1</math> &lt;- count <math>p_1</math> ms return (<math>w_0 &lt; w_1</math>) </pre>

**Fig. 1.** In this attack three threads run in parallel, colluding to leak secret **secret**. The two public threads write to a *public* output channel; the relative number of messages written on the channel by each thread directly leaks the secret (as inferred by  $p_1$ ). To affect the rate that  $p_0$  can write,  $s_0$  conditionally terminates—which will free up time on core  $c_0$  for  $p_0$  to execute.

### 2.3 Internalizing External Timing Attacks

LIO, like almost all IFC systems, considers external timing out of scope for its attacker model. Unfortunately, when we run LIO threads on multiple cores, in parallel, the allocation and reclamation of resources on behalf of secret threads is indirectly observable by public threads. Unsurprisingly, some of the above external timing attacks manifest internally—a thread running on a parallel core acts as an “external” attacker. To demonstrate the feasibility of such attacks, we describe two variants of the aforementioned scheduler-based attacks which leak sensitive information internally to public threads.

Secret threads can leak information by relinquishing CPU time, which the runtime reclaims and *unsafely* redistributes to public threads running on the same core. Our attack program consists of three threads: two public threads— $p_0$  and  $p_1$ —and a secret thread— $s_0$ . Figure 1 shows the pseudo-code for this attack. Note that the threads are secure in isolation, but leak the value of *secret* when executed in parallel, with a round robin scheduler. In particular, threads  $p_0$  and  $s_0$  run concurrently on core  $c_0$  using half of the CPU time each, while  $p_1$  runs in parallel alone on core  $c_1$  using all the CPU time. Both public threads repeatedly write their respective thread IDs to a *public channel*. The secret thread, on the other hand, loops forever or terminates depending on *secret*. Intuitively, when the secret thread terminates, the runtime system redirects its CPU time to  $p_0$ , causing both  $p_1$  and  $p_0$  to write at the same rate. In converse, when the secret thread does not terminate early,  $p_0$  is scheduled in a round-robin fashion with  $s_0$  on the same core and can thus only write half as fast as  $p_1$ . More specifically:

- If **secret** = **true**, thread  $s_0$  terminates and the runtime system assigns all the CPU time of core  $c_0$  to public thread  $p_0$ , which then writes at the same rate as thread  $p_1$  on core  $c_1$ . Then,  $p_0$  writes as many times as  $p_1$ , which then returns **true**.

- If `secret = false`, secret thread  $s_0$  loops and public thread  $p_0$  shares the CPU time on core  $c_0$  with it. Then,  $p_0$  writes messages at roughly half the rate of thread  $p_1$ , which writes more often—it has all the CPU time on  $c_1$ —and thus returns `false`.<sup>2</sup>

Secret LIO threads can also leak information by allocating many secret threads on a core with public threads—this reduces the CPU-time available to the public threads. For example, using the same setting with three threads from before, the secret thread forks a spinning thread on core  $c_1$  by replacing command `terminate` with command `fork (forever skip) c_1` in the code of thread  $s_0$  in Fig. 1. Intuitively, if `secret` is `false`, then  $p_1$  writes more often than  $p_0$  before, otherwise the write rate of  $p_1$  decreases—it shares core  $c_1$  with the child thread of  $s_0$ —and  $p_0$  writes as often as  $p_1$ .

Not all external timing attacks can be internalized, however. In particular, GHC’s approach to reclaiming memory via a stop-the-world GC simultaneously stops all threads on *all* cores, thus the relative write rate of public threads remain constant. Interestingly, though, implementing LIO on runtimes (e.g., Node.js as proposed by Heule et al. [17]) with modern parallel garbage collectors that do not always stop the world would internalize the GC-based external timing attacks. Similarly, abusing GHC’s memory allocation to exhaust all memory crashes all the program threads and, even though it cannot be internalized, it still results in information leakage.

### 3 Secure, Parallel Runtime System

To address the external and internal timing attacks, we propose a new dynamic IFC runtime system design. Fundamentally, today’s runtime systems are vulnerable because they automatically allocate and reclaim resources that are shared across threads of varying sensitivity. However, the automatic allocation and reclamation is not in itself a problem—it is only a problem because the runtime steals (and grants) resources from (and to) differently-labeled threads.

Our runtime system,  $\text{LIO}_{\text{PAR}}$ , explicitly partitions CPU-time and memory among threads—each thread has a fixed CPU-time and memory *budget* or *quota*. This allows resource management decisions to be made locally, for each thread, independent of the other threads in the system. For example, the runtime scheduler of  $\text{LIO}_{\text{PAR}}$  relies on CPU-time partitioning to ensure that threads always run for a fixed amount of time, irrespective of the other threads running on the same core. Similarly, in  $\text{LIO}_{\text{PAR}}$ , the memory allocator and garbage collector rely on memory partitioning to be able to allocate and collect memory on behalf of a thread without being influenced or otherwise influencing other threads in the system. Furthermore, partitioning resources among threads enables fine-grained control of resources:  $\text{LIO}_{\text{PAR}}$  exposes secure primitives to (i) measure resource usage (e.g., time and memory) and (ii) elicit garbage collection cycles.

<sup>2</sup> The attacker needs to empirically find parameter  $n$ , so that  $p_1$  writes roughly twice as much as thread  $p_0$  with half CPU time on core  $c_0$ .

The  $\text{LIO}_{\text{PAR}}$  runtime does not automatically balance resources between threads. Instead,  $\text{LIO}_{\text{PAR}}$  makes resource management explicit at the language level. When forking a new thread, for example,  $\text{LIO}_{\text{PAR}}$  demands that the parent thread give up part of its CPU-time and memory budgets to the children. Indeed,  $\text{LIO}_{\text{PAR}}$  even manages core ownership or *capabilities* that allow threads to fork threads across cores. This approach ensures that allocating new threads does not indirectly leak any information externally or to other threads. Dually, the  $\text{LIO}_{\text{PAR}}$  runtime does not re-purpose unused memory or CPU-time, even when a thread terminates or “dies” abruptly—parent threads must explicitly kill their children when they wish to reclaim their resources.

To ensure that CPU-time and memory can always be reclaimed,  $\text{LIO}_{\text{PAR}}$  allows threads to kill their children at any time. Unsurprisingly, this feature requires restricting the  $\text{LIO}_{\text{PAR}}$  floating-label approach more than that of  $\text{LIO}$ — $\text{LIO}_{\text{PAR}}$  threads cannot raise their current label if they have already forked other threads. As a result, in  $\text{LIO}_{\text{PAR}}$  threads form a *hierarchy*—children threads are always at least as sensitive as their parent—and thus it is secure to expose an API to *allocate* and *reclaim* resources.

**Attacks Revisited.**  $\text{LIO}_{\text{PAR}}$  enforces security against *reclamation-based attacks* because secret threads cannot automatically relinquish their resources. For example, our hierarchical runtime system stops the attack in Fig. 1: even if secret thread  $s_0$  terminates (`secret = true`), the throughput of public thread  $p_0$  remains constant— $\text{LIO}_{\text{PAR}}$  does not reassign the CPU time of  $s_0$  to  $p_0$ , but keeps  $s_0$  spinning until it gets killed. Similarly,  $\text{LIO}_{\text{PAR}}$  protects against *allocation-based attacks* because secret threads cannot steal resources owned by other public threads. For example, the *fork-bomb* variant of the previous attack fails because  $\text{LIO}_{\text{PAR}}$  aborts command `fork` (`forever skip`)  $c_1$ —thread  $s_0$  does not own the core capability  $c_1$ —and thus the throughput of  $p_1$  remains the same. In order to substantiate these claims, we first formalize the design of the *hierarchical* runtime system (Sect. 4) and establish its security guarantees (Sect. 5).

**Trust Model.** This work addresses attacks that exploit runtime system resource management—in particular memory and CPU-time. We do not address attacks that exploit other shared runtime system state (e.g., event loops [63], lazy evaluation [7, 59]), shared operating system state (e.g., file system locks [24], events and I/O [22, 32]), or shared hardware (e.g., caches, buses, pipelines and hardware threads [11, 47]) Though these are valid concerns, they are orthogonal and outside the scope of this paper.

## 4 Hierarchical Calculus

In this section we present the formal semantics of  $\text{LIO}_{\text{PAR}}$ . We model  $\text{LIO}_{\text{PAR}}$  as a security monitor that executes simply typed  $\lambda$ -calculus terms extended with *LIO* security primitives on an abstract machine in the style of Sestoft [53]. The security monitor reduces secure programs and aborts the execution of leaky programs.



Label	$\ell, pc, cl \in \mathcal{L}$	Params.	$\mu ::= (h, cl)$
Cores	$k \in \{1 \dots \kappa\}, K \in \mathcal{P}(\{1 \dots \kappa\})$	Heap	$\Delta \in \text{Var} \rightarrow \text{Term}$
Thread Id	$n \in \mathbb{N}, N \in \mathcal{P}(\mathbb{N})$	Budgets	$h, b \in \mathbb{N}$
Type	$\tau ::= () \mid \tau_1 \rightarrow \tau_2 \mid \text{Bool} \mid \mathcal{L} \mid \text{LIO } \tau \mid \text{Labeled } \tau$ $\mid \text{TId} \mid \text{Core} \mid \mathcal{P}(\{1 \dots \kappa\}) \mid \mathbb{N}$		
Value	$v ::= () \mid \lambda x. t \mid \text{True} \mid \text{False} \mid \ell \mid \text{return } t \mid \text{Labeled } \ell \ t^\circ \mid n \mid k \mid K$		
Term	$t ::= v \mid x \mid t_1 \ t_2 \mid \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \mid t_1 \gg t_2 \mid \text{label } t_1 \ t_2$ $\mid \text{unlabel } t \mid \text{fork } t_1 \ t_2 \ t_3 \ t_4 \ t_5 \mid \text{spawn } t_1 \ t_2 \ t_3 \ t_4 \ t_5 \mid \text{kill } t$ $\mid \text{size} \mid \text{time} \mid \text{wait } t \mid \text{send } t_1 \ t_2 \mid \text{receive}$		
CTerm	$t^\circ ::= t \text{ such that } fv(t) = \emptyset$		
Cont.	$C ::= x \mid \text{then } t_2 \text{ else } t_3 \mid \gg t_2 \mid \text{label } t \mid \text{unlabel} \mid \text{fork } t_1 \ t_2 \ t_3 \ t_4$ $\mid \text{spawn } t_1 \ t_2 \ t_3 \ t_4 \mid \text{kill} \mid \text{send } t$		
Stack	$S ::= [] \mid C : S$		
State	$s ::= (\Delta, pc, N \mid t, S)$		

(APP<sub>1</sub>)

$$\frac{|\Delta| < \mu.h \quad \text{fresh}(x)}{(\Delta, pc, N \mid t_1 \ t_2, S) \rightsquigarrow_\mu (\Delta[x \mapsto t_2], pc, N \mid t_1, x : S)}$$

(APP<sub>2</sub>)

$$(\Delta, pc, N \mid \lambda y. t, x : S) \rightsquigarrow_\mu (\Delta, pc, N \mid t[x / y], S)$$

(VAR)

$$\frac{x \mapsto t \in \Delta}{(\Delta, pc, N \mid x, S) \rightsquigarrow_\mu (\Delta, pc, N \mid t, S)}$$

(BIND<sub>1</sub>)

$$(\Delta, pc, N \mid t_1 \gg t_2, S) \rightsquigarrow_\mu (\Delta, pc, N \mid t_1, \gg t_2 : S)$$

(BIND<sub>2</sub>)

$$(\Delta, pc, N \mid \text{return } t_1, \gg t_2 : S) \rightsquigarrow_\mu (\Delta, pc, N \mid t_2 \ t_1, S)$$

(LABEL<sub>1</sub>)

$$(\Delta, pc, N \mid \text{label } t_1 \ t_2, S) \rightsquigarrow_\mu (\Delta, pc, N \mid t_1, \text{label } t_2 : S)$$

(LABEL<sub>2</sub>)

$$\frac{pc \sqsubseteq \ell \sqsubseteq \mu.cl \quad t^\circ = \Delta^*(t)}{(\Delta, pc, N \mid \ell, \text{label } t : S) \rightsquigarrow_\mu (\Delta, pc, N \mid \text{return } (\text{Labeled } \ell \ t^\circ), S)}$$

(UNLABEL<sub>1</sub>)

$$(\Delta, pc, N \mid \text{unlabel } t, S) \rightsquigarrow_\mu (\Delta, pc, N \mid t, \text{unlabel} : S)$$

(UNLABEL<sub>2</sub>)

$$\frac{pc \sqcup \ell \sqsubseteq \mu.cl}{(\Delta, pc, N \mid \text{Labeled } \ell \ t, \text{unlabel} : S) \rightsquigarrow_\mu (\Delta, pc \sqcup \ell, N \mid \text{return } t, S)}$$

**Fig. 2.** Sequential LIO<sub>PAR</sub>.

**Semantics.** The state of the monitor, written  $(\Delta, pc, N \mid t, S)$ , stores the state of a thread under execution and consists of a heap  $\Delta$  that maps variables to terms, the thread’s program counter label  $pc$ , the set  $N$  containing the identifiers of the thread’s children, the term currently under reduction  $t$  and a stack of continuations  $S$ . Figure 2 shows the interesting rules of the sequential small-step operational semantics of the security monitor. The notation  $s \rightsquigarrow_{\mu} s'$  denotes a transition of the machine in state  $s$  that reduces to state  $s'$  in one step with thread parameters  $\mu = (h, cl)$ .<sup>3</sup> Since we are interested in modeling a system with *finite* resources, we parameterize the transition with the maximum heap size  $h \in \mathbb{N}$ . Additionally, the clearance label  $cl$  represents an upper bound over the sensitivity of the thread’s floating counter label  $pc$ . Rule [APP<sub>1</sub>] begins a function application. Since our calculus is call-by-name, the function argument is saved as a *thunk* (i.e., an unevaluated expression) on the heap at fresh location  $x$  and the indirection is pushed on the stack for future lookups.<sup>4</sup> Note that the rule allocates memory on the heap, thus the premise  $|\Delta| < h$  forbids a heap overflow, where the notation  $|\Delta|$  denotes the size of the heap  $\Delta$ , i.e., the number of bindings that it contains.<sup>5</sup> To avoid overflows, a thread can measure the size of its own heap via primitive *size* (Sect. 4.2). If  $t_1$  evaluates to a function, e.g.,  $\lambda y.t$ , rule [APP<sub>2</sub>] starts evaluating the body, in which the bound variable  $y$  is substituted with the heap-allocated argument  $x$ , i.e.,  $t[x/y]$ . When the evaluation of the function body requires the value of the argument, variable  $x$  is looked up in the heap (rule [VAR]). In the next paragraph we present the rules of the basic security primitives. The other sequential rules are available in the extended version of this paper.

**Security Primitives.** A labeled value *Labeled*  $\ell t^\circ$  of type *Labeled*  $\tau$  consists of term  $t$  of type  $\tau$  and a label  $\ell$ , which reflects the sensitivity of the content. The annotation  $t^\circ$  denotes that term  $t$  is *closed* and does not contain any free variable, i.e.,  $fv(t) = \emptyset$ . We restrict the syntax of labeled values with closed terms for security reasons. Intuitively, LIO<sub>PAR</sub> allocates free variables inside a secret labeled values on the heap, which then leaks information to public threads with its size. For example, a public thread could distinguish between two secret values, e.g., *Labeled*  $H x$  with heap  $\Delta = [x \mapsto 42]$ , and *Labeled*  $H 0$  with heap  $\Delta = \emptyset$ , by measuring the size of the heap. To avoid that, labeled values are closed and the size of the heap of a thread at a certain security level, is not affected by data labeled at different security levels. A term of type *LIO*  $\tau$  is a secure computation that performs side effects and returns a result of type  $\tau$ . Secure computations are structured using standard monadic constructs *return*  $t$ , which embeds term  $t$  in the monad, and *bind*, written  $t_1 \gg t_2$ , which sequentially

<sup>3</sup> We use record notation, i.e.,  $\mu.h$  and  $\mu.cl$ , to access the components of  $\mu$ .

<sup>4</sup> The calculus does not feature lazy evaluation. Laziness, because of *sharing*, introduces a covert channel, which has already been considered in previous work [59].

<sup>5</sup> To simplify reasoning, our generic memory model is basic and assumes a uniform size for all the objects stored in the heap. We believe that it is possible to refine our generic model with more accurate memory models (e.g., GHC’s tagless G-machine (STG) [23], the basis for GHC’s runtime [39]), but leave this to future work.

composes two monadic actions, the second of which takes the result of the first as an argument. Rule  $[\text{BIND}_1]$  deconstructs a computation  $t_1 \gg t_2$  into term  $t_1$  to be reduced first and pushes on the stack the continuation  $\gg t_2$  to be invoked after term  $t_1$ .<sup>6</sup> Then, the second rule  $[\text{BIND}_2]$  pops the topmost continuation placed on the stack (i.e.,  $\gg t_2$ ) and evaluates it with the result of the first computation (i.e.,  $t_2 \ t_1$ ), which is considered complete when it evaluates to a monadic value, i.e., to syntactic form *return*  $t_1$ . The runtime monitor secures the interaction between computations and labeled values. In particular, secure computations can construct and inspect labeled values exclusively with monadic primitives *label* and *unlabel* respectively. Rules  $[\text{LABEL}_1]$  and  $[\text{UNLABEL}_1]$  are straightforward and follow the pattern seen in the other rules. Rule  $[\text{LABEL}_2]$  generates a labeled value at security level  $\ell$ , subject to the constraint  $pc \sqsubseteq \ell \sqsubseteq cl$ , which prevents a computation from labeling values below the program counter label  $pc$  or above the clearance label  $cl$ .<sup>7</sup> The rule computes the closure of the content, i.e., closed term  $t^\circ$ , by recursively substituting every free variable in term  $t$  with its value in the heap, written  $\Delta^*(t)$ . Rule  $[\text{UNLABEL}_2]$  extracts the content of a labeled value and taints the program counter label with its label, i.e., it rises it to  $pc \sqcup \ell$ , to reflect the sensitivity of the data that is now in scope. The premise  $pc \sqcup \ell \sqsubseteq cl$  ensures that the program counter label does not float over the clearance  $cl$ . Thus, the run-time monitor prevents the program counter label from floating above the clearance label (i.e.,  $pc \sqsubseteq cl$  always holds).

The calculus also includes concurrent primitives to allocate resources when forking threads (*fork* and *spawn* in Sect. 4.1), reclaim resources and measure resource usage (*kill*, *size*, and *time* in Sect. 4.2), threads synchronization and communication (*wait*, *send* and *receive* in the extended version of this paper).

## 4.1 Core Scheduler

In this section, we extend  $\text{LIO}_{\text{PAR}}$  with concurrency, which enables (i) *interleaved* execution of threads on a single core and (ii) *simultaneous* execution on  $\kappa$  cores. To protect against attacks that exploit the automatic management of shared *finite* resource (e.g., those in Sect. 2.3),  $\text{LIO}_{\text{PAR}}$  maintains a resource budget for each running thread and updates it as threads allocate and reclaim resources. Since  $\kappa$  threads execute at the same time, those changes must be coordinated in order to preserve the consistency of the resource budgets and guarantee *deterministic parallelism*. For this reason, the hierarchical runtime system is split in two components: (i) the *core scheduler*, which executes threads on a single core, ensures that they respect their resource budgets and performs security checks, and (ii) the top-level *parallel scheduler*, which synchronizes the execution on multiple cores and reassigns resources by updating the resource budgets according to the instructions of the core schedulers. We now introduce the core scheduler and describe the top-level parallel scheduler in Sect. 4.3.

<sup>6</sup> Even though the stack size is unbounded in this model, we could account for its memory usage by explicitly allocating it on the heap, in the style of Yang et al. [66].

<sup>7</sup> The labels form a security lattice  $(\mathcal{L}, \sqcup, \sqsubseteq)$ .

Thread Map $T \in TId \rightarrow State$	Global State $\Sigma ::= (T, B, H, \theta, \omega)$
Time Map $B \in TId \rightarrow \mathbb{N}$	Core Queue $Q ::= \langle n^b \rangle \mid \langle Q_1 \mid Q_2 \rangle$
Size Map $H \in TId \rightarrow \mathbb{N}$	Event $e ::= \epsilon \mid \mathbf{fork}(\Delta, n, t, b, h)$
Core Map $\theta \in TId \rightarrow \mathcal{P}(\{1.. \kappa\})$	$\mid \mathbf{spawn}(\Delta, n, t, K)$
Clock $\omega \in \mathbb{N}$	$\mid \mathbf{kill}(n) \mid \mathbf{send}(n, t)$

  

STEP	
$\Sigma.T(n) = s$	$\mu = (\Sigma.H(n), n.cl) \quad s \rightsquigarrow_\mu s'$
$\frac{}{Q[\langle n^{1+b} \rangle] \xrightarrow{(n, s', \epsilon)}_\Sigma Q[\langle n^b \rangle]}$	

  

FORK	
$\Sigma.T(n) = (\Delta, pc, N \mid b_2, \mathbf{fork} \ell_L \ell_H h_2 t : S)$	
$pc \sqsubseteq \ell_L \quad n' \leftarrow \mathbf{fresh}^{TId}(\ell_L, \ell_H, n.k)$	
$s = (\Delta, pc, \{n'\} \cup N \mid \mathbf{return} \, n', S) \quad \Delta' = \{x \mapsto \Delta(x) \mid x \in fv^*(t, \Delta)\}$	
$\Sigma.H(n) = h_1 + h_2 \quad  \Delta  \leq h_1 \quad  \Delta'  \leq h_2$	
$\frac{}{Q[\langle n^{1+b_1+b_2} \rangle] \xrightarrow{(n, s, \mathbf{fork}(\Delta', n', t, b_2, h_2))}_\Sigma Q[\langle \langle n^{b_1} \rangle \mid \langle n^{b_2} \rangle \rangle]}$	

  

SPAWN	
$\Sigma.T(n) = (\Delta, pc, N \mid k, \mathbf{spawn} \ell_L \ell_H K_1 t : S)$	
$\Sigma.\theta(n) = \{k\} \cup K_1 \cup K_2 \quad pc \sqsubseteq \ell_L \quad n' \leftarrow \mathbf{fresh}^{TId}(\ell_L, \ell_H, k)$	
$s = (\Delta, pc, \{n'\} \cup N \mid \mathbf{return} \, n', S) \quad \Delta' = \{x \mapsto \Delta(x) \mid x \in fv^*(t, \Delta)\}$	
$\frac{}{Q[\langle n^{1+b} \rangle] \xrightarrow{(n, s, \mathbf{spawn}(\Delta', n', t, K_1))}_\Sigma Q[\langle n^b \rangle]}$	

  

STUCK	
$\Sigma.T(n) = s$	$MaxHeapSize(s, \Sigma.H(n)) \vee UnlabelStuck(n, \Sigma.T) \vee$
	$ForkStuck(n, \Sigma.H, \Sigma.T) \vee SpawnStuck(s, \theta(n)) \vee ValueStuck(s) \vee$
	$WaitStuck(n, T) \vee ReceiveStuck(s) \vee KillStuck(s)$
$\frac{}{Q[\langle n^{1+b} \rangle] \xrightarrow{(n, s, \epsilon)}_\Sigma Q[\langle n^b \rangle]}$	

  

CONTEXT SWITCH	
$s_o = ([, \perp, \emptyset \mid \mathbf{return} \, (), []]$	
$\frac{}{Q[\langle n^0 \rangle] \xrightarrow{(o, s_o, \epsilon)}_\Sigma Q[\langle n_1^{\Sigma.B(n_1)} \rangle, \dots, \langle n_{ Q }^{\Sigma.B(n_{ Q })} \rangle]}$	

Fig. 3. Concurrent LIO<sub>PAR</sub>.

**Syntax.** Figure 3 presents the core scheduler, which has access to the global state  $\Sigma = (T, B, H, \theta, \omega)$ , consisting of a thread pool map  $T$ , which maps a thread id to the corresponding thread's current state, the time budget map  $B$ , a memory budget map  $H$ , core capabilities map  $\theta$ , and the global clock  $\omega$ . Using these maps, the core scheduler ensures that thread  $n$ : (i) performs  $B(n)$  uninterrupted steps until the next thread takes over, (ii) does not grow its heap above its maximum heap size  $H(n)$ , and (iii) has exclusive access to the *free* core capabilities  $\theta(n)$ . Furthermore, each thread id  $n$  records the *initial* current label when the thread was created ( $n.pc$ ), its clearance ( $n.cl$ ), and the core where it runs ( $n.k$ ), so that

the runtime system can enforce security. Notice that thread ids are *opaque* to threads—they cannot force them nor access their fields.

**Hierarchical Scheduling.** The core scheduler performs *deterministic* and *hierarchical* scheduling—threads lower in the hierarchy are scheduled first, i.e., parent threads are scheduled before their children. The scheduler manages a core run queue  $Q$ , which is structured as a binary tree with leaves storing thread ids and residual time budgets. The notation  $n^b$  indicates that thread  $n$  can run for  $b$  more steps before the next thread runs. When a new thread is spawned, the scheduler creates a subtree with the parent thread on the left and the child on the right. The scheduler can therefore find the thread with the highest priority by following the left spine of the tree and backtracking to the right if a thread has no residual budget.<sup>8</sup> We write  $Q[\langle n^b \rangle]$  to mean the first thread encountered via this traversal is  $n$  with budget  $b$ . As a result, given the slice  $Q[\langle n^{1+b} \rangle]$ , thread  $n$  is the next thread to run, and  $Q[\langle n^0 \rangle]$  occurs only if *all* threads in the queue have zero residual budget. We overload this notation to represent tree updates: a rule  $Q[\langle n^{1+b} \rangle] \rightarrow Q[\langle n^b \rangle]$  finds the next thread to run in queue  $Q$  and decreases its budget by one.

**Semantics.** Figure 3 formally defines the transition  $Q \xrightarrow{(n,s,e)}_{\Sigma} Q'$ , which represents an execution step of the *core scheduler* that schedules thread  $n$  in core queue  $Q$ , executes it with global state  $\Sigma = (T, B, H, \theta, \omega)$  and updates the queue to  $Q'$ . Additionally, the core scheduler informs the parallel scheduler of the final state  $s$  of the thread and requests on its behalf to update the global state by means of event message  $e$ . In rule [STEP], the scheduler retrieves the next thread in the schedule, i.e.,  $Q[\langle n^{1+b} \rangle]$  and its state in the thread pool from the global state, i.e.,  $\Sigma.T(n) = s$ . Then, it executes the thread for one sequential step with its memory budget and clearance, i.e.,  $s \rightsquigarrow_{\mu} s'$  with  $\mu = (\Sigma.H(n), n.cl)$ , sends the empty event  $\epsilon$  to the parallel scheduler, and decrements the thread's residual budget in the final queue, i.e.,  $Q[\langle n^b \rangle]$ . In rule [FORK], thread  $n$  creates a new thread  $t$  with initial label  $\ell_L$  and clearance  $\ell_H$ , such that  $\ell_L \sqsubseteq \ell_H$  and  $pc \sqsubseteq \ell_L$ . The child thread runs on the same core of the parent thread, i.e.,  $n.k$ , with fresh id  $n'$ , which is then added to the set of children, i.e.,  $\{n'\} \cup N$ . Since parent and child threads do not share memory, the core scheduler must copy the portion of the parent's private heap reachable by the child's thread, i.e.,  $\Delta'$ ; we do this by copying the bindings of the variables that are transitively reachable from  $t$ , i.e.,  $fv^*(t, \Delta)$ , from the parent's heap  $\Delta$ . The parent thread gives  $h_2$  of its memory budget  $\Sigma.H(n)$  to its child. The conditions  $|\Delta| \leq h_1$  and  $|\Delta'| \leq h_2$ , ensure that the heaps do not overflow their new budgets. Similarly, the core scheduler splits the residual time budget of

<sup>8</sup> When implemented, this procedure might introduce a timing channel that leaks the number of threads running on the core. In practice, techniques from real time schedulers can be used to protect against such timing channels. The model of LIO<sub>PAR</sub> does not capture the execution time of the runtime system itself and thus this issue does not arise in the security proofs.

the parent into  $b_1$  and  $b_2$  and informs the parallel scheduler about the new thread and its resources with event **fork**( $\Delta'$ ,  $n'$ ,  $t$ ,  $b_2$ ,  $h_2$ ), and lastly updates the tree  $Q$  by replacing the leaf  $\langle n^{1+b_1+b_2} \rangle$  with the two-leaves tree  $\langle \langle n^{b_1} \rangle | \langle n'^{b_2} \rangle \rangle$ , so that the child thread will be scheduled immediately after the parent has consumed its remaining budget  $b_1$ , as explained above. Rule [SPAWN] is similar to [FORK], but consumes core capability resources instead of time and memory. In this case, the core scheduler checks that the parent thread owns the core where the child is scheduled and the core capabilities assigned to the child, i.e.,  $\theta(n) = \{k\} \cup K_1 \cup K_2$  for some set  $K_2$ , and informs the parallel scheduler with event **spawn**( $\Delta'$ ,  $n'$ ,  $t$ ,  $K_1$ ). Rule [STUCK] performs busy waiting by consuming the time budget of the scheduled thread, when it is *stuck* and cannot make any progress—the premises of the rule enumerate the conditions under which this can occur (see the extended version of this paper for details). Lastly, in rule [CONTEXTSWITCH] all the threads scheduled in the core queue have consumed their time budget, i.e.,  $Q[\langle n^0 \rangle]$  and the core scheduler resets their residual budget using the budget map  $\Sigma.B$ . In the rule, the notation  $Q[\langle n_i^b \rangle]$  selects the  $i$ -th leaf, where  $i \in \{1 \dots |Q|\}$  and  $|Q|$  denotes the number of leaves of tree  $Q$  and symbol  $\circ$  denotes the thread identifier of the core scheduler, which updates a dummy thread that simply spins during a context-switch or whenever the core is unused.

## 4.2 Resource Reclamation and Observations

The calculus presented so far enables threads to manage their time, memory and core capabilities hierarchically, but does not provide any primitive to reclaim their resources. This section rectifies this by introducing (i) a primitive to kill a thread and return its resources back to the owner and (ii) a primitive to elicit a garbage collection cycle and reclaim unused memory. Furthermore, we demonstrate that the runtime system presented in this paper is robust against timing attacks by exposing a timer API allowing threads to access a global clock.<sup>9</sup> Intuitively, it is secure to expose this feature because LIO<sub>PAR</sub> ensures that the time spent executing high threads is fixed in advanced, so timing measurements of low threads remain unaffected. Lastly, since memory is hierarchically partitioned, each thread can securely query the current size of its *private heap*, enabling fine-grained control over the garbage collector.

**Kill.** A parent thread can reclaim the resources given to its child thread  $n'$ , by executing *kill*  $n'$ . If the child thread has itself forked or spawned other threads, they are transitively killed and their resources returned to the parent thread. The concurrent rule [KILL<sub>2</sub>] in Fig. 4 initiates this process, which is completed by the parallel scheduler via event **kill**( $n'$ ). Note that the rule applies only when the thread killed is a *direct* child of the parent thread—that is when the parent’s children set has shape  $\{n'\} \cup N$  for some set  $N$ . Now that threads can unrestrictedly reclaim resources by killing their children, we must revise the primitive

<sup>9</sup> An *external* attacker can take timing measurements using network communications. An attacker equipped with an *internal* clock is equally powerful but simpler to formalize [46].

$$\begin{array}{c}
\text{KILL}_2 \\
\frac{\Sigma.T(n) = (\Delta, pc, \{n'\} \cup N \mid n', kill : S) \quad s = (\Delta, pc, N \mid return (), S)}{Q[\langle n^{1+b} \rangle] \xrightarrow{(n, s, kill(n'))}_{\Sigma} Q[\langle n^b \rangle]} \\
\\
\text{UNLABEL}_2 \\
\frac{pc \sqcup \ell \sqsubseteq \mu.cl \quad \forall n \in N . pc \sqcup \ell \sqsubseteq n.pc}{(\Delta, pc, N \mid Labeled \ell t, unlabel : S) \rightsquigarrow_{\mu} (\Delta, pc \sqcup \ell, N \mid return t, S)} \\
\\
\text{GC} \\
\frac{R = fv^*(t, \Delta) \cup fv^*(S, \Delta) \quad \Delta' = \{x \mapsto \Delta(x) \mid x \in R\}}{\langle \Delta, pc, N \mid gc t, S \rangle \rightsquigarrow_{\mu} \langle \Delta', pc, N \mid t, S \rangle} \\
\\
\text{APP-GC} \\
\frac{|\Delta| \equiv \mu.h}{\langle \Delta, pc, N \mid t_1 t_2, S \rangle \rightsquigarrow_{\mu} \langle \Delta, pc, N \mid gc (t_1 t_2), S \rangle} \\
\\
\text{SIZE} \\
\langle \Delta, pc, N \mid size, S \rangle \rightsquigarrow_{\mu} \langle \Delta, pc, N \mid return |\Delta|, S \rangle \\
\\
\text{TIME} \\
\frac{\Sigma.T(n) = (\Delta, pc, N \mid time, S) \quad s = (\Delta, pc, N \mid return \Sigma.\omega, S)}{Q[\langle n^{1+b} \rangle] \xrightarrow{(n, s, \epsilon)}_{\Sigma} Q[\langle n^b \rangle]}
\end{array}$$

**Fig. 4.** LIO<sub>PAR</sub> with resource reclamation and observation primitives.

*unlabel*, since the naive combination of *kill* and *unlabel* can result in information leakage. This will happen if a public thread forks another public thread, then reads a secret value (raising its label to secret), and based on that decides to kill the child. To close the leak, we modify the rule [UNLABEL<sub>2</sub>] by adding the highlighted premise, causing the primitive *unlabel* to fail whenever the parent thread’s label would float above the *initial* current label of one of its children.

**Garbage Collection.** Rule [GC] extends LIO<sub>PAR</sub> with a *time-sensitive hierarchical* garbage collector via the primitive *gc t*. The rule elicits a garbage collection cycle which drops entries that are no longer needed from the heap, and then evaluates *t*. The sub-heap  $\Delta'$  includes the portion of the current heap that is (transitively) *reachable* from the free variables in scope (i.e. those present in the term,  $fv^*(t, \Delta)$  or on the stack  $fv^*(S, \Delta)$ ). After collection, the thread resumes and evaluates term *t* under compacted private heap  $\Delta'$ .<sup>10</sup> In rule [APP-GC], a collection is *automatically* triggered when the thread’s next memory allocation would overflow the heap.

<sup>10</sup> In practice a garbage collection cycle takes time that is proportional to the size of the memory used by the thread. That does not hinder security as long as the garbage collector runs on the thread’s time budget.

**Resource Observations.** All threads in the system share a global fine-grained clock  $\omega$ , which is incremented by the parallel scheduler at each cycle (see below). Rule [TIME] gives all threads unrestricted access to the clock via monadic primitive *time*.

### 4.3 Parallel Scheduler

This section extends LIO<sub>PAR</sub> with *deterministic parallelism*, which allows to execute  $\kappa$  threads simultaneously on as many cores. To this end, we introduce the top-level parallel scheduler, which coordinates simultaneous changes to the global state by updating the resource budgets of the threads in response core events (e.g., fork, spawn, and kill) and ticks the global clock.

Queue Map $\Phi \in \{1 \dots \kappa\} \rightarrow \text{Queue}$	Configuration $c ::= \langle T, B, H, \theta, \omega, \Phi \rangle$
<p>PARALLEL</p> $\frac{\forall i \in \{1 \dots \kappa\}. \Phi(i) \xrightarrow{(n_i, s_i, e_i)}_{\Sigma} Q_i \quad T' = \Sigma.T[n_i \mapsto s_i] \quad \Phi' = \Phi[i \mapsto Q_i] \quad c = \langle T', B, H, \theta, \Sigma.\omega + 1, \Phi' \rangle \quad \langle \Sigma', \Phi'' \rangle = \llbracket \text{sort}[(n_1, e_1), \dots, (n_\kappa, e_\kappa)] \rrbracket^c}{\langle \Sigma, \Phi \rangle \hookrightarrow \langle \Sigma', \Phi'' \rangle}$	
$next(\_, \epsilon, c) = c$ $next(n_1, \text{fork}(\Delta, n_2, t, b, h), c)$ $= \langle T', B', H', \theta'[n_2 \mapsto h], \omega', \Phi' \rangle$ <b>where</b> $\langle T, B, H, \theta, \omega, \Phi \rangle = c$ $s = (\Delta, n_2.pc, \emptyset \mid t, [])$ $T' = T[n_2 \mapsto s]$ $B' = B[n_1 \mapsto B(n_1) - b]$ $H' = H[n_1 \mapsto H(n_1) - h]$ $\theta' = \theta[n_2 \mapsto \emptyset]$	$next(n_1, \text{spawn}(\Delta, n_2, t, K), c)$ $= \langle T', B', H', \theta'[n_2 \mapsto K], \omega, \Phi' \rangle$ <b>where</b> $\langle T, B, H, \theta, \omega, \Phi \rangle = c$ $s = (\Delta, n_2.pc, \emptyset \mid t, [])$ $T' = T[n_2 \mapsto s]$ $B' = B[n_2 \mapsto B_0]$ $H' = H[n_2 \mapsto H_0]$ $\theta' = \theta[n_1 \mapsto \theta(n_1) \setminus \{n_2.k\} \cup K]$ $\Phi' = \Phi[n_2.k \mapsto \langle n_2^{B_0} \rangle]$
$next(n, \text{kill}(n'), \langle T, B, H, \theta, \omega, \Phi \rangle)$ $\mid n \notin \text{Dom}(T) = \langle T, B, H, \theta, \omega, \Phi \rangle$ $\mid n \in \text{Dom}(T) = \langle T \setminus N, B' \setminus N, H' \setminus N, \theta' \setminus N, \omega, \Phi' \rangle$ <b>where</b> $N = \llbracket \{n'\} \rrbracket^T$ $B' = B[n \mapsto B(n) + \sum_{i \in N, i.k=n.k} B(i)]$ $H' = H[n \mapsto H(n) + \sum_{i \in N, i.k=n.k} H(i)]$ $\theta' = \theta[n \mapsto \theta(n) \cup \bigcup_{i \in N} \theta(i) \cup \{i.k \mid i \in N, i.k \neq n.k\}]$ $\Phi' = \lambda k. \Phi[k \mapsto \Phi(k) \setminus N]$	

**Fig. 5.** Top-level parallel scheduler.

**Semantics.** Figure 5 formalizes the operational semantics of the parallel scheduler, which reduces a configuration  $c = \langle \Sigma, \Phi \rangle$  consisting of global state  $\Sigma$  and



core map  $\Phi$  mapping each core to its run queue, to configuration  $c'$  in one step, written  $c \hookrightarrow c'$ , through rule [PARALLEL] only. The rule executes the threads scheduled on each of the  $\kappa$  cores, which all step at once according to the concurrent semantics presented in Sects. 4.1–4.2, with the same current global state  $\Sigma$ . Since the execution of each thread can change  $\Sigma$  *concurrently*, the top-level parallel scheduler reconciles those actions by updating  $\Sigma$  *sequentially* and *deterministically*.<sup>11</sup> First, the scheduler updates the thread pool map  $T$  and core map  $\Phi$  with the final state obtained by running each thread in isolation, i.e.,  $T' = \Sigma.T[n_i \mapsto s_i]$  and  $\Phi' = \Phi[i \mapsto Q_i]$  for  $i \in \{1 \dots \kappa\}$ . Then, it collects all concurrent events generated by the  $\kappa$  threads together with their thread id, sorts the events according to type, i.e.,  $\text{sort}[(n_1, e_1), \dots, (n_\kappa, e_\kappa)]$ , and computes the updated configuration by processing the events in sequence.<sup>12</sup> In particular, new threads are created first (event **spawn**( $\cdot$ ) and **fork**( $\cdot$ )), and then killed (event **kill**( $\cdot$ ))—the ordering between events of the same type is arbitrary and assumed to be fixed. Trivial events ( $\epsilon$ ) do not affect the configuration and thus their ordering is irrelevant. The function  $\langle\langle es \rangle\rangle^c$  computes a final configuration by processing a list of events in order, accumulating configuration updates ( $\text{next}(\cdot)$  updates the current configuration by one event-step):  $\langle\langle (n, e) : es \rangle\rangle^c = \langle\langle es \rangle\rangle^{\text{next}(n, e, c)}$ . When no more events need processing, the configuration is returned  $\langle\langle [] \rangle\rangle^c = c$ .

**Event Processing.** Figure 5 defines function  $\text{next}(n, e, c)$ , which takes a thread identifier  $n$ , the event  $e$  that thread  $n$  generated, the current configuration and outputs the configuration obtained by performing the thread’s action. The empty event  $\epsilon$  is trivial and leaves the state unchanged. Event  $(n_1, \text{fork}(\Delta, n_2, t, b, h))$  indicates that thread  $n_1$  forks thread  $t$  with identifier  $n_2$ , sub-heap  $\Delta$ , time budget  $b$  and maximum heap size  $h$ . The scheduler deducts these resources from the parent’s budgets, i.e.,  $B' = B[n_1 \mapsto B(n_1) - b]$  and  $H' = H[n_1 \mapsto H(n_1) - h]$  and assigns them to the child, i.e.,  $B'[n_2 \mapsto b]$  and  $H'[n_2 \mapsto h]$ .<sup>13</sup> The new child shares the core with the parent—it has no core capabilities i.e.,  $\theta' = \theta[n_2 \mapsto \emptyset]$ —and so the core map is left unchanged. Lastly, the scheduler adds the child to the thread pool and initializes its state, i.e.,  $T[n_2 \mapsto (\Delta, n_2.\ell_L, \emptyset \mid t, [])]$ . The scheduler handles event  $(n_1, \text{spawn}(\Delta, n_2, t, K))$  similarly. The new thread  $t$  gets scheduled on core  $n_2.k$ , i.e.,  $\Phi[n_2.k \mapsto \langle n_2^{B_0} \rangle]$ , where the thread takes all the time and memory resources of the core, i.e.,  $B[n_2 \mapsto B_0]$  and  $H[n_2 \mapsto H_0]$ , and extra core capabilities  $K$ , i.e.,  $\theta'[n_2 \mapsto K]$ . For simplicity, we assume that all cores execute  $B_0$  steps per-cycle and feature a memory of size  $H_0$ . Event  $(n, \text{kill}(n'))$  informs the scheduler that thread  $n$  wishes to kill thread  $n'$ . The scheduler leaves the global state unchanged if the parent thread has already been killed by the time this event is handled, i.e., when the guard  $n \notin \text{Dom}(T)$  is true—the resources of the child  $n'$  will have been reclaimed by another ancestor.

<sup>11</sup> Non-deterministic updates would make the model vulnerable to refinement attacks [40].

<sup>12</sup> Since the clock only needs to be incremented, we could have left it out from the configuration  $c = \langle T', B, H, \theta, \Sigma.\omega + 1, \Phi' \rangle$ ; function  $\langle\langle es \rangle\rangle^c$  does not use nor change its value.

<sup>13</sup> Notice that  $|\Delta| < h$  by rule [FORK].

Otherwise, the scheduler collects the identifiers of the descendants of  $n'$  that are *alive* ( $N = \llbracket \{n'\} \rrbracket^T$ )—they must be killed (and reclaimed) *transitively*. The set  $N$  is computed recursively by  $\llbracket N \rrbracket^T$ , using the thread pool  $T$ , i.e.,  $\llbracket \emptyset \rrbracket^T = \emptyset$ ,  $\llbracket \{n\} \rrbracket^T = \{n\} \cup \llbracket T(n).N \rrbracket^T$  and  $\llbracket N_1 \cup N_2 \rrbracket^T = \llbracket N_1 \rrbracket^T \cup \llbracket N_2 \rrbracket^T$ . The scheduler then increases the time and memory budget of the parent with the sum of the budget of all its descendants scheduled on the *same* core, i.e.,  $\sum_{i \in N, i.k=n.k} B(i)$  (resp.  $\sum_{i \in N, i.k=n.k} H(i)$ )—descendants running on other cores do not share those resources. The scheduler reassigns to the parent thread their core capabilities, which are split between capabilities explicitly assigned but not in use, i.e.,  $\bigcup_{i \in N} \theta(i)$  and core capabilities assigned and in use by running threads, i.e.,  $\{i.k \mid i \in N, i.k \neq n.k\}$ . Lastly, the scheduler removes the killed threads from each core, written  $\Phi(i) \setminus N$ , by pruning the leaves containing killed threads and reassigning their leftover time budget to their parent, see the extended version of this paper for details.

## 5 Security Guarantees

In this section we show that  $\text{LIO}_{\text{PAR}}$  satisfies a strong security condition that ensures timing-agreement of threads and rules out timing covert channels. In Sect. 5.1, we describe our proof technique based on *term erasure*, which has been used to verify security guarantees of functional programming languages [30], IFC libraries [8, 17, 54, 56, 61], and an IFC runtime system [59]. In Sect. 5.2, we formally prove security, i.e., *progress- and timing-sensitive non-interference*, a strong form of non-interference [14], inspired by Volpano and Smith [64]—to our knowledge, it is considered here for the first time in the context of parallel runtime systems. Works that do not address external timing channels [59, 62] normally prove *progress-sensitive* non-interference, wherein the number of execution steps of a program may differ in two runs based on a secret. This condition is insufficient in the parallel setting: both public and secret threads may step simultaneously on different cores and any difference in the number of execution steps would introduce *external* and *internal* timing attacks. Similar to previous works on secure multi-threaded systems [36, 52], we establish a *strong* low-bisimulation property of the parallel scheduler, which guarantees that attacker-indistinguishable configurations execute in lock-step and remain indistinguishable. Theorem 1 and Corollary 1 use this property to ensure that any two related parallel programs execute in exactly the same number of steps.

### 5.1 Erasure Function

The term erasure technique relies on an *erasure function*, written  $\varepsilon_L(\cdot)$ , which rewrites secret data above the attacker’s level  $L$  to special term  $\bullet$ , in all the syntactic categories: values, terms, heaps, stacks, global states and configurations.<sup>14</sup> Once the erasure function is defined, the core of the proof technique

<sup>14</sup> For ease of exposition, we use the two-point lattices  $\{L, H\}$ , where  $H \not\sqsubseteq L$  is the only disallowed flow. Neither our proofs nor our model rely on this particular lattice.

consists of proving an essential *commutativity* relationship between the erasure function and reduction steps: given a step  $c \hookrightarrow c'$ , there must exist a reduction that *simulates* the original reduction between the erased configurations, i.e.,  $\varepsilon_L(c) \hookrightarrow \varepsilon_L(c')$ . Intuitively, if the configuration  $c$  leaked secret data while stepping to  $c'$ , that data would be classified as public in  $c'$  and thus would remain in  $\varepsilon_L(c')$ —but such secret data would be erased by  $\varepsilon_L(c)$  and the property would not hold. The erasure function leaves ground values, e.g.,  $()$ , unchanged and on most terms it acts homomorphically, e.g.,  $\varepsilon_L(t_1 \ t_2) = \varepsilon_L(t_1) \ \varepsilon_L(t_2)$ . The interesting cases are for labeled values, thread configurations, and resource maps. The erasure function removes the content of secret labeled values, i.e.,  $\varepsilon_L(\text{Labeled } H \ t^\circ) = \text{Labeled } H \ \bullet$ , and erases the content recursively otherwise, i.e.,  $\varepsilon_L(\text{Labeled } L \ t^\circ) = \text{Labeled } L \ \varepsilon_L(t)^\circ$ . The state of a thread is erased per-component, homomorphically if the program counter label is public, i.e.,  $\varepsilon_L(\Delta, L, N, | \ t, S) = (\varepsilon_L(\Delta), L, N \mid \varepsilon_L(t), \varepsilon_L(S))$ , and in full otherwise, i.e.,  $\varepsilon_L(\Delta, H, N, | \ t, S) = (\bullet, \bullet, \bullet \mid \bullet, \bullet)$ .

**Resource Erasure.** Since  $\text{LIO}_{\text{PAR}}$  manages resources explicitly, the simulation property above requires to define the erasure function for resources as well. The erasure function should *preserve* information about the resources (e.g., time, memory, and core capabilities) of *public threads*, since the attacker can explicitly assign resources (e.g., with *fork* and *swap*) and measure them (e.g., with *size*). But what about the resources of secret threads? One might think that such information is secret and thus it should be erased—intuitively, a thread might decide to assign, say, half of its time budget to its secret child depending on secret information. However, public threads can also assign (public) resources to a secret thread when forking: even though these resources currently belong to the secret child, they are *temporary*—the public parent might reclaim them later. Thus, we cannot associate the sensitivity of the resources of a thread with its program counter label when resources are managed *hierarchically*, as in  $\text{LIO}_{\text{PAR}}$ . Instead, we associate the security level of the resources of a secret thread with the sensitivity of its parent: the resources of a secret thread are *public* information whenever the program counter label of the parent is public and *secret* information otherwise. Furthermore, since resource reclamation is transitive, the erasure function cannot discard secret resources, but must rather redistribute them to the hierarchically closest set of public resources, as when *killing* them.

**Time Budget.** First, we project the identifiers of *public* threads from the thread pool  $T : \text{Dom}_L(T) = \{n_L \mid n \in \text{Dom}(T) \wedge T(n).pc \equiv L\}$ , where notation  $n_L$  indicates that the program counter label of thread  $n$  is public. Then, the set  $P = \bigcup_{n \in \text{Dom}_L(T)} \{n\} \cup T(n).N$  contains the identifiers of all the public threads and their immediate children.<sup>15</sup> The resources of threads  $n \in P$  are public information. However, the program counter label of a thread  $n \in P$  is not necessarily public, as explained previously. Hence  $P$  can be disjointly partitioned

<sup>15</sup> The id of the spinning thread on each free core is also public, i.e.,  $o_k \in P$  for  $k \in \{1.. \kappa\}$ .

by program counter label:  $P = P_L \cup P_H$ , where  $P_L = \{n_L \mid n \in P\}$  and  $P_H = \{n_H \mid n \in P\}$ . Erasure of the budget map then proceeds on this partition, leaving the budget of the public threads untouched, and summing the budget of their secret children threads to the budgets of their descendants, which are instead omitted. In symbols,  $\varepsilon_L(B) = B_L \cup B_H$ , where  $B_L = \{n_L \mapsto B(n_L) \mid n_L \in P_L\}$  and  $B_H = \{n_H \mapsto B(n_H) + \sum_{i \in \llbracket \{n_H\} \rrbracket^T} B(i) \mid n_H \in P_H\}$ .

**Queue Erasure.** The erasure of core queues follows the same intuition, preserving public and secret threads  $n \in P$  and trimming all other secret threads  $n_H \notin P$ . Since queues annotate thread ids with their residual time budgets, the erasure function must reassign the budgets of all *secret* threads  $n'_H \notin P$  to their closest ancestor  $n \in P$  on the same core. The ancestor  $n \in P$  could be either (i) another *secret* thread on the same core, i.e.,  $n_H \in P$ , or, (ii) the spinning thread of that core,  $\circ \in P$  if there is no other thread  $n \in P$  on that core—the difference between these two cases lies on whether the original thread  $n'$  was *forked* or *spawned* on that core. More formally, if the queue contains no thread  $n \in P$ , then the function replaces the queue altogether with the spinning thread and returns the residual budgets of the threads to it, i.e.,  $\varepsilon_L(Q) = \langle \circ^B \rangle$  if  $n_i \notin P$  and  $B = \sum b_i$ , for each leaf  $Q[\langle n_i^{b_i} \rangle]$  where  $i \in \{1 \dots |Q|\}$ . Otherwise, the core contains at least a thread  $n_H \in P$  and the erasure function returns the residual time budget of its secret descendants, i.e.,  $\varepsilon_L(Q) = Q \downarrow_L$  by combining the effects of the following mutually recursive functions:

$$\begin{aligned} \langle n^b \rangle \downarrow_L &= \langle n^b \rangle & \langle n_{1H}^{b_1} \rangle \curlyvee \langle n_{2H}^{b_2} \rangle &= \langle n_{1H}^{b_1+b_2} \rangle \\ \langle Q_1, Q_2 \rangle \downarrow_L &= (\langle Q_1 \rangle \downarrow_L) \curlyvee (\langle Q_2 \rangle \downarrow_L) & Q_1 \curlyvee Q_2 &= \langle Q_1, Q_2 \rangle \end{aligned}$$

The interesting case is  $\langle n_{1H}^{b_1} \rangle \curlyvee \langle n_{2H}^{b_2} \rangle$ , which reassigns the budget of the child (the right leaf  $\langle n_{2H}^{b_2} \rangle$ ) to the parent (the left leaf  $\langle n_{1H}^{b_1} \rangle$ ), by rewriting the subtree into  $\langle n_{1H}^{b_1+b_2} \rangle$ .

## 5.2 Timing-Sensitive Non-interference

The proof of progress- and timing-sensitive non-interference relies on two fundamental properties, i.e., *determinacy* and *simulation* of parallel reductions. Determinacy requires that the reduction relation is deterministic.

**Proposition 1 (Determinism).** *If  $c_1 \hookrightarrow c_2$  and  $c_1 \hookrightarrow c_3$  then  $c_2 \equiv c_3$ .*

The equivalence in the statement denotes alpha-equivalence, i.e., up to the choice of variable names. We now show that the parallel scheduler preserves  $L$ -equivalence of parallel configurations.

**Definition 1 ( $L$ -equivalence).** *Two configurations  $c_1$  and  $c_2$  are indistinguishable from an attacker at security level  $L$ , written  $c_1 \approx_L c_2$ , if and only if  $\varepsilon_L(c_1) \equiv \varepsilon_L(c_2)$ .*

**Proposition 2 (Parallel simulation).** *If  $c \hookrightarrow c'$ , then  $\varepsilon_L(c) \hookrightarrow \varepsilon_L(c')$ .*

By combining *determinism* (Proposition 1) and *parallel simulation* (Proposition 2), we prove *progress-insensitive non-interference*, which assumes progress of both configurations.

**Proposition 3 (Progress-insensitive non-interference).** *If  $c_1 \hookrightarrow c'_1$ ,  $c_2 \hookrightarrow c'_2$  and  $c_1 \approx_L c_2$ , then  $c'_1 \approx_L c'_2$ .*

In order to lift this result to be progress-sensitive, we first prove *timing-sensitive progress*. Intuitively, if a *valid* configuration steps then any low equivalent parallel configuration also steps.<sup>16</sup>

**Proposition 4 (Timing-sensitive progress).** *Given a valid configuration  $c_1$  and a parallel reduction step  $c_1 \hookrightarrow c'_1$  and  $c_1 \approx_L c_2$ , then there exists  $c'_2$ , such that  $c_2 \hookrightarrow c'_2$ .*

Using progress-insensitive non-interference, i.e., Proposition 3 and timing-sensitive progress, i.e., Proposition 4 in combination, we obtain a *strong L*-bisimulation property between configurations and prove *progress- and timing-sensitive non-interference*.

**Theorem 1 (Progress- and timing-sensitive non-interference).** *For all valid configurations  $c_1$  and  $c_2$ , if  $c_1 \hookrightarrow c'_1$  and  $c_1 \approx_L c_2$ , then there exists a configuration  $c'_2$ , such that  $c_2 \hookrightarrow c'_2$  and  $c'_1 \approx_L c'_2$ .*

The following corollary instantiates the non-interference security theorem from above for a given  $\text{LIO}_{\text{PAR}}$  parallel program, that explicitly rules out leaks via timing channels. In the following, the notation  $\hookrightarrow_u$  denotes  $u$  reduction steps of the parallel scheduler.

**Corollary 1.** *Given a well-typed  $\text{LIO}_{\text{PAR}}$  program  $t$  of type Labeled  $\tau_1 \rightarrow \text{LIO } \tau_2$  and two closed secrets  $t_1^\circ, t_2^\circ :: \tau_1$ , let  $s_i = ([, L, \emptyset, | t \text{ (Labeled } H \text{ } t_i^\circ), [ ])$ ,  $c_i = (T_i, B, H, \theta, 0, \Phi_i)$ , where  $T_i = [n_L \mapsto s_i, \circ_j \mapsto s_\circ]$ ,  $B = [n_L \mapsto B_0, \circ_j \mapsto 0]$ ,  $H = [n_L \mapsto H_0, \circ_j \mapsto H_0]$ ,  $\theta = [n_L \mapsto \{2.. \kappa\}, \circ_j \mapsto \emptyset]$ ,  $\Phi_i = [1 \mapsto \langle s_i \rangle, 2 \mapsto \langle \circ_2 \rangle, \dots, \kappa \mapsto \langle \circ_\kappa \rangle]$ , for  $i \in \{1, 2\}$ ,  $j \in \{1.. \kappa\}$  and thread identifier  $n_L$  such that  $n.k = 1$  and  $n.cl = H$ . If  $c_1 \hookrightarrow_u c'_1$ , then there exists configuration  $c'_2$ , such that  $c_2 \hookrightarrow_u c'_2$  and  $c'_1 \approx_L c'_2$ .*

To conclude, we show that the *timing-sensitive* security guarantees of  $\text{LIO}_{\text{PAR}}$  extend to concurrent *single-core* programs by instantiating Corollary 1 with  $\kappa = 1$ .

## 6 Limitations

**Implementation.** Implementing  $\text{LIO}_{\text{PAR}}$  is a serious undertaking that requires a major redesign of GHC's runtime system. Conventional runtime systems freely

<sup>16</sup> A configuration is valid if satisfies several basic properties, e.g., it does not contain special term  $\bullet$ . See the extended version of this paper for details.

share resources among threads to boost performance and guarantee fairness. For instance, in GHC, threads share heap objects to save memory space and execution time (when evaluating expressions). In contrast,  $\text{LIO}_{\text{PAR}}$  strictly partitions resources to enforce security—threads at different security labels cannot share heap objects. As a result, the GHC memory allocator must be adapted to isolate threads’ private heap, so that allocation and collection can occur independently and in parallel. Similarly, the GHC “fair” round robin scheduler must be heavily modified to keep track of and manage threads’ time budget, to preemptively perform a context switch when their time slice is up.

**Programming Model.** Since resource management is explicit, building applications atop  $\text{LIO}_{\text{PAR}}$  introduces new challenges—the programmer must explicitly choose resource bounds for each thread. If done poorly, threads can spend excessive amounts of time sitting idle when given too much CPU time, or garbage collecting when not given enough heap space. The problem of tuning resource allocation parameters is not unique to  $\text{LIO}_{\text{PAR}}$ —Yang and Mazières’ [66] propose to use GHC profiling mechanisms to determine heap size while the real-time garbage collector by Henriksson [16] required the programmer to specify the worst case execution time, period, and worst-case allocation of each high-priority thread. Das and Hoffmann [9] demonstrate a more automatic approach—they apply machine learning techniques to statically determine upper bounds on execution time and heap usage of OCaml programs. Similar techniques could be applied to  $\text{LIO}_{\text{PAR}}$  in order to determine the most efficient resource partitions. Moreover, this challenge is not unique to real-time systems or  $\text{LIO}_{\text{PAR}}$ ; choosing privacy parameters in differential privacy, for example, shares many similarities [21, 29].

The  $\text{LIO}_{\text{PAR}}$  programming model is also likely easier to use in certain application domains—e.g., web applications where the tail latency of a route can inform the thread bounds, or embedded systems where similar latency requirements are the norm. Nevertheless, in order to simplify programming with  $\text{LIO}_{\text{PAR}}$ , we intend to introduce privileges (and thus declassification) similar to LIO [12, 56] or COWL [57].

Coarse-grained, floating-label systems such as LIO and  $\text{LIO}_{\text{PAR}}$  can suffer *label creep*, wherein the current computation gets tainted to a point where it cannot perform any useful writes [55]. Sequential LIO [56] addresses label creep through a primitive, `toLabeled`, which executes a computation (that may raise the current label) in a separate context and restores the current label upon its termination. Similar to concurrent LIO [54],  $\text{LIO}_{\text{PAR}}$  relies on `fork` to address label creep and not `toLabeled`—the latter exposes the termination covert-channel [54]. Even though  $\text{LIO}_{\text{PAR}}$  has a more restricted floating-label semantics than concurrent LIO,  $\text{LIO}_{\text{PAR}}$  also supports parallel execution, garbage collection, and new APIs for getting heap statistics, counting elapsed time, and killing threads.

## 7 Related Work

There is substantial work on language-level IFC systems [10, 15, 20, 34, 43, 50, 51, 54, 55, 67, 68]. Our work builds on these efforts in several ways. Firstly, LIO<sub>PAR</sub> extends the concurrent LIO IFC system [54] with parallelism—to our knowledge, this is the first *dynamic* IFC system to support parallelism and address the internalization of external timing channels. Previous static IFC systems implicitly allow for parallelism, e.g., Muller and Chong’s [41], several works on IFC  $\pi$ -calculi [18, 19, 25], and Rafnsson et al. [49] recent foundations for composable timing-sensitive interactive systems. These efforts, however, do not model runtime system resource management. Volpano and Smith [64] enforce a timing agreement condition, similar to ours, but for a static concurrent IFC system. Mantel et al. [37] and Li et al. [31] prove non-interference for static, concurrent systems, using rely-guarantee reasoning.

Unlike most of these previous efforts, our hierarchical runtime system also eliminates classes of resource-based external timing channels, such as memory exhaustion and garbage collection. Pedersen and Askarov [46], however, were the first to identify automatic memory management to be a source of covert channels for IFC systems and demonstrate the feasibility of attacks against both V8 and the JVM. They propose a sequential static IFC language with labeled-partitioned memory and a label-aware timing-sensitive garbage collector, which is vulnerable to *external timing* attacks and satisfies only *termination-insensitive* non-interference.

Previous work on language-based systems—namely [35, 66]—identify memory retention and memory exhaustion as a source of denial-of-service (DOS) attacks. Memory retention and exhaustion can also be used as covert channels. In addressing those covert channels, LIO<sub>PAR</sub> also addresses the DOS attacks outlined by these efforts. Indeed, our work generalizes Yang and Mazières’ [66] region-based allocation framework with region-based garbage collection and hierarchical scheduling.

Our LIO<sub>PAR</sub> design also borrows ideas from the secure operating system community. Our explicit hierarchical memory management is conceptually similar to HiStar’s container abstraction [69]. In HiStar, containers—subject to quotas, i.e., space limits—are used to hierarchically allocate and deallocate objects. LIO<sub>PAR</sub> adopts this idea at the language-level and automates the allocation and reclamation. Moreover, we hierarchically partition CPU-time; Zeldovich et al. [69], however, did observe that their container abstraction can be repurposed to enforce CPU quotas. Deterland [65] splits time into ticks to address internal timing channels and mitigate external timing ones. Deterland builds on Determinator [4], an OS that executes parallel applications deterministically and efficiently. LIO<sub>PAR</sub> adopts many ideas from these systems—both the deterministic parallelism and ticks (semantic steps)—to the language-level. Deterministic parallelism at the language-level has also been explored previous to this work [27, 28, 38], but, different from these efforts, LIO<sub>PAR</sub> also hierarchically manages resources to eliminate classes of external timing channels.



Fabric [33,34] and DStar [70] are distributed IFC systems. Though we believe that our techniques would scale beyond multi-core systems (e.g., to data centers), LIO<sub>PAR</sub> will likely not easily scale to large distributed systems like Fabric and DStar. Different from Fabric and DStar, however, LIO<sub>PAR</sub> addresses both internal and external timing channels that result from running code in parallel.

Our hierarchical resource management approach is not unique—other countermeasures to external timing channels have been studied. Hu [22], for example, mitigates both timing channels in the VAX/VMM system [32] using “fuzzy time”—an idea recently adopted to browsers [26]. Askarov et al.’s [2] mitigate external timing channels using predicative black-box mitigation, which delays events and thus bound information leakage. Rather than using noise as in the fuzzy time technique, however, they predict the schedule of future events. Some of these approaches have also been adopted at the language-level [46,54,71]. We find these techniques largely orthogonal: they can be used alongside our techniques to mitigate timing channels we do not eliminate.

Real-time systems—when developed with garbage collected languages [3,5,6,16]—face similar challenges as this work. Blelloch and Cheng [6] describe a real-time garbage collector (RTGC) for multi-core programs with *provable* resource bounds—LIO<sub>PAR</sub> *enforces* resource bounds instead. A more recent RTGC created by Auerbach et al. [3] describes a technique to “tax” threads into contributing to garbage collection as they utilize more resources. Henricksson [16] describes a RTGC capable of enforcing hard and soft deadlines, once given upper bounds on space and time resources used by threads. Most similarly to LIO<sub>PAR</sub>, Pizlo et al. [48] implement a hierarchical RTGC algorithm that independently collects partitioned heaps.

## 8 Conclusion

Language-based IFC systems built atop off-the-shelf runtime systems are vulnerable to resource-based external-timing attacks. When these systems are extended with thread parallelism these attacks become yet more vicious—they can be carried out internally. We presented LIO<sub>PAR</sub>, the design of the first dynamic IFC hierarchical runtime system that supports deterministic parallelism and eliminates both resource-based internal- and external-timing covert channels. To our knowledge, LIO<sub>PAR</sub> is the first parallel system to satisfy progress- and timing-sensitive non-interference.

## References

1. Askarov, A., Hunt, S., Sabelfeld, A., Sands, D.: Termination-insensitive noninterference leaks more than just a bit. In: Jajodia, S., Lopez, J. (eds.) ESORICS 2008. LNCS, vol. 5283, pp. 333–348. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-88313-5\\_22](https://doi.org/10.1007/978-3-540-88313-5_22)
2. Askarov, A., Zhang, D., Myers, A.C.: Predictive black-box mitigation of timing channels. In: Proceedings of the 17th ACM Conference on Computer and Communications Security, pp. 297–307. ACM (2010)



3. Auerbach, J., et al.: Tax-and-spend: democratic scheduling for real-time garbage collection. In: Proceedings of the 8th ACM International Conference on Embedded Software, pp. 245–254. ACM (2008)
4. Aviram, A., Weng, S.-C., Hu, S., Ford, B.: Efficient system-enforced deterministic parallelism. *Commun. ACM* **55**(5), 111–119 (2012)
5. Baker Jr., H.G.: List processing in real time on a serial computer. *Commun. ACM* **21**(4), 280–294 (1978)
6. Blueloch, G.E., Cheng, P.: On bounding time and space for multiprocessor garbage collection. *ACM SIGPLAN Not.* **34**, 104–117 (1999)
7. Buiras, P., Russo, A.: Lazy programs leak secrets. In: Riis Nielson, H., Gollmann, D. (eds.) *NordSec 2013. LNCS*, vol. 8208, pp. 116–122. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-41488-6\\_8](https://doi.org/10.1007/978-3-642-41488-6_8)
8. Buiras, P., Vytiniotis, D., Russo, A.: HLIO: mixing static and dynamic typing for information-flow control in Haskell. In: *ACM SIGPLAN International Conference on Functional Programming. ACM* (2015)
9. Das, A., Hoffmann, J.: ML for ML: learning cost semantics by experiment. In: Legay, A., Margaria, T. (eds.) *TACAS 2017. LNCS*, vol. 10205, pp. 190–207. Springer, Heidelberg (2017). [https://doi.org/10.1007/978-3-662-54577-5\\_11](https://doi.org/10.1007/978-3-662-54577-5_11)
10. Fernandes, E., Paupore, J., Rahmati, A., Simionato, D., Conti, M., Prakash, A.: FlowFence: practical data protection for emerging IoT application frameworks. In: *USENIX Security Symposium*, pp. 531–548 (2016)
11. Ge, Q., Yarom, Y., Cock, D., Heiser, G.: A survey of microarchitectural timing attacks and countermeasures on contemporary hardware. *J. Cryptographic Eng.* **8**, 1–27 (2016)
12. Giffin, D.B., et al.: Hails: protecting data privacy in untrusted web applications. *J. Comput. Secur.* **25**(4–5), 427–461 (2017)
13. Giffin, D.B., et al.: Hails: protecting data privacy in untrusted web applications. In: *Proceedings of the Symposium on Operating Systems Design and Implementation. USENIX* (2012)
14. Goguen, J.A., Meseguer, J.: Unwinding and inference control, pp. 75–86, April 1984
15. Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A.: JSFlow: tracking information flow in JavaScript and its APIs. In: *Proceedings of the 29th Annual ACM Symposium on Applied Computing*, pp. 1663–1671. ACM (2014)
16. Henriksson, R.: Scheduling garbage collection in embedded systems. Ph.D. thesis, Department of Computer Science (1998)
17. Heule, S., Stefan, D., Yang, E.Z., Mitchell, J.C., Russo, A.: IFC inside: retrofitting languages with dynamic information flow control. In: Focardi, R., Myers, A. (eds.) *POST 2015. LNCS*, vol. 9036, pp. 11–31. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46666-7\\_2](https://doi.org/10.1007/978-3-662-46666-7_2)
18. Honda, K., Vasconcelos, V., Yoshida, N.: Secure information flow as typed process behaviour. In: Smolka, G. (ed.) *ESOP 2000. LNCS*, vol. 1782, pp. 180–199. Springer, Heidelberg (2000). [https://doi.org/10.1007/3-540-46425-5\\_12](https://doi.org/10.1007/3-540-46425-5_12)
19. Honda, K., Yoshida, N.: A uniform type structure for secure information flow. *ACM Trans. Program. Lang. Syst. (TOPLAS)* **29**(6), 31 (2007)
20. Hritcu, C., Greenberg, M., Karel, B., Pierce, B.C., Morrisett, G.: All your IFCException are belong to us. In: *2013 IEEE Symposium on Security and Privacy (SP)*, pp. 3–17. IEEE (2013)
21. Hsu, J., et al.: Differential privacy: an economic method for choosing epsilon. In: *Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium, CSF 2014*, pp. 398–410. IEEE Computer Society, Washington, DC (2014)

22. Hu, W.-M.: Reducing timing channels with fuzzy time. *J. Comput. Secur.* **1**(3–4), 233–254 (1992)
23. Jones, S.L.P.: Implementing lazy functional languages on stock hardware: the spineless tagless G-machine. *J. Funct. Program.* **2**, 127–202 (1992)
24. Kemmerer, R.A.: Shared resource matrix methodology: an approach to identifying storage and timing channels. *ACM Trans. Comput. Syst. (TOCS)* **1**(3), 256–277 (1983)
25. Kobayashi, N.: Type-based information flow analysis for the  $\pi$ -calculus. *Acta Informatica* **42**(4–5), 291–347 (2005)
26. Kohlbrenner, D., Shacham, H.: Trusted browsers for uncertain times. In: *USENIX Security Symposium*, pp. 463–480 (2016)
27. Kuper, L., Newton, R.R.: LVars: lattice-based data structures for deterministic parallelism. In: *Proceedings of the 2nd ACM SIGPLAN Workshop on Functional High-Performance Computing*, pp. 71–84. ACM (2013)
28. Kuper, L., Todd, A., Tobin-Hochstadt, S., Newton, R.R.: Taming the parallel effect zoo: extensible deterministic parallelism with LVish. *ACM SIGPLAN Not.* **49**(6), 2–14 (2014)
29. Lee, J., Clifton, C.: How much is enough? Choosing  $\epsilon$  for differential privacy. In: Lai, X., Zhou, J., Li, H. (eds.) *ISC 2011*. LNCS, vol. 7001, pp. 325–340. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-24861-0\\_22](https://doi.org/10.1007/978-3-642-24861-0_22)
30. Li, P., Zdancewic, S.: Arrows for secure information flow. *Theoret. Comput. Sci.* **411**(19), 1974–1994 (2010)
31. Li, X., Mantel, H., Tasch, M.: Taming message-passing communication in compositional reasoning about confidentiality. In: Chang, B.-Y.E. (ed.) *APLAS 2017*. LNCS, vol. 10695, pp. 45–66. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-71237-6\\_3](https://doi.org/10.1007/978-3-319-71237-6_3)
32. Lipner, S., Jaeger, T., Zurko, M.E.: Lessons from VAX/SVS for high-assurance VM systems. *IEEE Secur. Priv.* **10**(6), 26–35 (2012)
33. Liu, J., Arden, O., George, M.D., Myers, A.C.: Fabric: building open distributed systems securely by construction. *J. Comput. Secur.* **25**(4–5), 367–426 (2017)
34. Liu, J., George, M.D., Vikram, K., Qi, X., Waye, L., Myers, A.C.: Fabric: a platform for secure distributed computation and storage. In: *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*. ACM (2009)
35. Liu, J., Myers, A.C.: Defining and enforcing referential security. In: Abadi, M., Kremer, S. (eds.) *POST 2014*. LNCS, vol. 8414, pp. 199–219. Springer, Heidelberg (2014). [https://doi.org/10.1007/978-3-642-54792-8\\_11](https://doi.org/10.1007/978-3-642-54792-8_11)
36. Mantel, H., Sabelfeld, A.: A unifying approach to the security of distributed and multi-threaded programs. *J. Comput. Secur.* **11**(4), 615–676 (2003)
37. Mantel, H., Sands, D., Sudbrock, H.: Assumptions and guarantees for compositional noninterference. In: *2011 IEEE 24th Computer Security Foundations Symposium*, pp. 218–232, June 2011
38. Marlow, S., Newton, R., Peyton Jones, S.: A monad for deterministic parallelism. *ACM SIGPLAN Not.* **46**(12), 71–82 (2012)
39. Marlow, S., Peyton Jones, S.: Making a fast curry: push/enter vs. eval/apply for higher-order languages. *J. Funct. Program.* **16**(4–5), 415–449 (2006)
40. McCullough, D.: Specifications for multi-level security and a hook-up. In: *1987 IEEE Symposium on Security and Privacy (SP)*, p. 161, April 1987
41. Muller, S., Chong, S.: Towards a practical secure concurrent language. In: *Proceedings of the 25th Annual ACM SIGPLAN Conference on Object-Oriented Programming Languages, Systems, Languages, and Applications*, pp. 57–74. ACM Press, New York, October 2012

42. Myers, A.C., Zheng, L., Zdancewic, S., Chong, S., Nystrom, N.: Jif 3.0: Java information flow, July 2006
43. Nadkarni, A., Andow, B., Enck, W., Jha, S.: Practical DIFC enforcement on android. In: USENIX Security Symposium, pp. 1119–1136 (2016)
44. North, S.C., Reppy, J.H.: Concurrent garbage collection on stock hardware. In: Kahn, G. (ed.) FPCA 1987. LNCS, vol. 274, pp. 113–133. Springer, Heidelberg (1987). [https://doi.org/10.1007/3-540-18317-5\\_8](https://doi.org/10.1007/3-540-18317-5_8)
45. Parker, J.L.: LMonad: information flow control for Haskell web applications. Ph.D. thesis, University of Maryland, College Park (2014)
46. Pedersen, M.V., Askarov, A.: From trash to treasure: timing-sensitive garbage collection. In: Proceedings of the 38th IEEE Symposium on Security and Privacy. IEEE (2017)
47. Percival, C.: Cache missing for fun and profit (2005)
48. Pizlo, F., Hosking, A.L., Vitek, J.: Hierarchical real-time garbage collection. In: Proceedings of the 2007 ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems, LCTES 2007, pp. 123–133. ACM, New York (2007)
49. Rafnsson, W., Jia, L., Bauer, L.: Timing-sensitive noninterference through composition. In: Maffei, M., Ryan, M. (eds.) POST 2017. LNCS, vol. 10204, pp. 3–25. Springer, Heidelberg (2017). [https://doi.org/10.1007/978-3-662-54455-6\\_1](https://doi.org/10.1007/978-3-662-54455-6_1)
50. Roy, I., Porter, D.E., Bond, M.D., McKinley, K.S., Witchel, E.: Laminar: practical fine-grained decentralized information flow control, vol. 44. ACM (2009)
51. Russo, A.: Functional pearl: two can keep a secret, if one of them uses Haskell. ACM SIGPLAN Not. **50**, 280–288 (2015)
52. Sabelfeld, A., Sands, D.: Probabilistic noninterference for multi-threaded programs. In: Proceedings of the 13th IEEE Workshop on Computer Security Foundations, CSFW 2000, p. 200. IEEE Computer Society, Washington, DC (2000)
53. Sestoft, P.: Deriving a lazy abstract machine. J. Funct. Program. **7**(3), 231–264 (1997)
54. Stefan, D., Russo, A., Buiras, P., Levy, A., Mitchell, J.C., Mazières, D.: Addressing covert termination and timing channels in concurrent information flow systems. In: International Conference on Functional Programming (ICFP). ACM SIGPLAN, September 2012
55. Stefan, D., Russo, A., Mazières, D., Mitchell, J.C.: Flexible dynamic information flow control in the presence of exceptions. J. Funct. Program. **27** (2017)
56. Stefan, D., Russo, A., Mitchell, J.C., Mazières, D.: Flexible dynamic information flow control in Haskell. In: Haskell Symposium. ACM SIGPLAN, September 2011
57. Stefan, D., et al.: Protecting users by confining JavaScript with COWL. In: USENIX Symposium on Operating Systems Design and Implementation. USENIX Association (2014)
58. Tsai, T.-C., Russo, A., Hughes, J.: A library for secure multi-threaded information flow in Haskell. In: 20th IEEE Computer Security Foundations Symposium, CSF 2007, pp. 187–202. IEEE (2007)
59. Vassena, M., Breitner, J., Russo, A.: Securing concurrent lazy programs against information leakage. In: 30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, 21–25 August 2017, pp. 37–52 (2017)
60. Vassena, M., Buiras, P., Waye, L., Russo, A.: Flexible manipulation of labeled values for information-flow control libraries. In: Askoxylakis, I., Ioannidis, S., Katsikas, S., Meadows, C. (eds.) ESORICS 2016. LNCS, vol. 9878, pp. 538–557. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-45744-4\\_27](https://doi.org/10.1007/978-3-319-45744-4_27)

61. Vassena, M., Russo, A.: On formalizing information-flow control libraries. In: Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security, PLAS 2016, pp. 15–28. ACM, New York (2016)
62. Vassena, M., Russo, A., Buiras, P., Waye, L.: MAC a verified static information-flow control library. *J. Log. Algebraic Methods Program.* (2017)
63. Vila, P., Köpf, B.: Loophole: timing attacks on shared event loops in chrome. In: USENIX Security Symposium (2017)
64. Volpano, D., Smith, G.: Eliminating covert flows with minimum typings. In: Proceedings of the 10th IEEE Workshop on Computer Security Foundations, CSFW 1997, p. 156. IEEE Computer Society, Washington, DC (1997)
65. Wu, W., Zhai, E., Wolinsky, D.I., Ford, B., Gu, L., Jackowitz, D.: Warding off timing attacks in Deterland. In: Conference on Timely Results in Operating Systems, Monterey, CS, US (2015)
66. Yang, E.Z., Mazières, D.: Dynamic space limits for Haskell. In: Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2014, pp. 588–598. ACM, New York (2014)
67. Yang, J., Hance, T., Austin, T.H., Solar-Lezama, A., Flanagan, C., Chong, S.: Precise, dynamic information flow for database-backed applications. *ACM SIGPLAN Not.* **51**, 631–647 (2016)
68. Yang, J., Yessenov, K., Solar-Lezama, A.: A language for automatically enforcing privacy policies. *ACM SIGPLAN Not.* **47**, 85–96 (2012)
69. Zeldovich, N., Boyd-Wickizer, S., Kohler, E., Mazières, D.: Making information flow explicit in HiStar. In: Proceedings of the 7th Symposium on Operating Systems Design and Implementation, pp. 263–278. USENIX Association (2006)
70. Zeldovich, N., Boyd-Wickizer, S., Mazieres, D.: Securing distributed systems with information flow control. *NSDI* **8**, 293–308 (2008)
71. Zhang, D., Askarov, A., Myers, A.C.: Predictive mitigation of timing channels in interactive systems. In: Proceedings of the 18th ACM Conference on Computer and Communications Security, pp. 563–574. ACM (2011)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

