Quantifying and Preventing Side Channels with Substructural Type Systems

Matt Fredrikson Jan Hoffmann Van Chan Ngo
Carnegie Mellon University

Abstract

Static techniques for deriving upper bounds on the resource consumption of programs have been extensively studied. However, there are applications that require more fine-grained information such as the difference between upper and lower bounds or the guaranty that the resource usage of a program does not differ for certain inputs. This article presents two novel substructural type systems for deriving lower bounds and for proving that a program has constant resource consumption for a class of inputs. The type systems are based on the potential method of amortized analysis to achieve compositionality, precision, and automatic inference using off-the-shelf linear optimization. While classic amortized analysis treats potential as an affine resource, the novel type systems treat potential as a relevant and linear resource, respectively. The soundness of the type systems with respect to an operational cost semantics is verified using the proof assistant Agda. The novel constant-resource and lower bound analyses are applied to quantify and prevent security vulnerabilities that leak secret information through resource consumption, such as side channels. First, implementations of the lower bound and constant-resource type systems in Resource Aware ML are used to automatically verify constant-time implementations of list comparison, encryption and decryption routines, database queries, and other resource-sensitive functionality. Second, the type systems are used to implement a method for automatically turning programs into constant-resource programs using LP solving. The method is static, does not require tracking resources at runtime, and works on most programs for which Resource Aware ML can derive an upper bound. Third, a resource-aware noninterference property is introduced. It relaxes the constant-resource requirement on programs, and requires only that resource usage does not leak information about *secret* inputs. This property is statically verified by combining the linear type system for constant resource consumption with an information flow type system.

1 Introduction

Automatic static analysis of the resource consumption of programs is an active area of research. Motivated by applications in embedded and real-time systems [W$^+$08], finding performance bugs [ODL15], and providing feedback to developers [GMC09], static resource analysis techniques have focused on derivation of worst-case bounds [SZV14, CHK$^+$15, BEF$^+$14, ALM12, AM13, DLR12, AFR15, AAG$^+$09]. One successful technique for automatically finding resource bounds at compile time is automatic amortized resource analysis (AARA). The idea of AARA is to combine the potential method of amortized analysis with existing programming languages techniques to achieve automation. For example, AARA has been integrated into type systems to automatically derive linear [HJ03] and polynomial [HH10, HAH12, HS14] bounds for strict and higher-order [JHLH10, Ano15] functional programs.

The main advantages of AARA are compositionality, efficiency, and precision. It has been shown that the technique can automatically derive bounds for complex real-world programs such as parts of the CompCert C Compiler [Ano15] and the cBench benchmark suite [CHS15]. Precision and efficiency stems from the selection of algebraic structures such as multivariate resource polynomials [HAH12] that can represent a wide range of bounds, as well as the reduction of bound inference to efficient LP solving. AARA is naturally compositional since the potential methods integrates reasoning about size changes and resource consumption. However, *existing AARA techniques are limited to worst-case bounds.*

Novel resource type systems The starting point of this paper is the technical insight that *the potential method of amortized analysis can also be used to derive lower bounds*, as well as to *prove that a program has constant resource consumption for a fixed input size.* In classic AARA the potential is used as an affine resource: it must be available to cover cost but excess potential is simply discarded. We show that if potential is treated as a linear resource, then corresponding type derivations prove that programs have *constant* resource consumption, i.e., resource consumption is independent of the execution path. Intuitively, this amounts to requiring that *all* potential must be used to cover the cost and that excess potential is not wasted. Furthermore, we show that if potential is treated as a relevant resource, then we derive *lower bounds* on the resource usage. Following a similar intuition, this requires that all potential is used, but the available potential does not need to be sufficient to cover the remaining cost.

The two novel type systems that we present enjoy the same advantages as classic AARA for upper bounds. They are naturally compositional, often derive precise results, and allow for fully-automated type inference based on LP solving. Moreover, as in classic AARA, they are parametric in the resource of interest and incorporate user-specified resource metrics that assign a constant cost to each basic operation. The type systems discussed in this paper apply to a simple first-order functional language, and use the linear potential annotations from the original work of Hofmann and Jost [HJ03]. This is sufficient to discuss the main technical points although it limits the systems to linear bounds. However, our implementation builds on Resource Aware ML (RAML) [Ano15], and supports polynomial bounds, user-defined data types, and higher-order functions. We formalized the soundness proof of these type systems, as well as that of classic linear AARA, in the proof assistant Agda. Soundness is proved with respect to an operational cost semantics, and like the type systems themselves, is parametric in the resource of interest.

Side channel mitigation In the second half of the paper, we apply our lower-bound and constant-resource type systems to the problem of preventing and quantifying *side channel vulnerabilities.* Side channel attacks extract sensitive information about a program's state through its use of resources such as time, network, and memory. Several notable instances of this type of attack have demonstrated leakage of cryptographic keys [Koc96, BB03, CHVV03, AP13, GBK11] and private user data [HPN11, AKM+15, FS00, BB07, ZJRR14] through such channels.

Whereas traditional notions of information flow can be described in terms of standard program semantics, a similar treatment of side channels requires incorporating the corresponding resource into the semantics and applying quantitative reasoning. This difficulty has led previous work in the area to treat resource use indirectly, by reasoning about the flow of secret information into branching control flow [ABB+16, RQaPA16, BBC+14, MPSW06] or introducing obfuscation components that mask secret-dependent differences in resource use [AZM10, KD09]. These approaches can limit program expressiveness or lead to unnecessary performance penalties.

In contrast, our approach performs quantitative analysis of resource use directly through the constant-resource type system. We consider an adversary that is able to observe the final resource consumption of a program as specified by a cost semantics, and derive a proof that the attacker's observations will not change as the program's inputs do. Although this observation model does not cover all known side-channel attacks, it applies to a large class of attackers that are not able to make intermediate observations of the program's behavior, such as those that reside over a network. Additionally, we show how one can use derived upper and lower bounds to quantify leakage through resource use, by reasoning about the number of distinct observations an attacker can make.

In general, requiring that a program only ever consumes a constant amount of resources is too restrictive. In most settings, it is sufficient to make sure that the resource usage of a program does not depend on selected parts of the input. To account for this, we present a new information flow type system that incorporates our constant-resource type system to reason about an adversary who can observe and manipulate inputs marked public, but can only make observations on secret inputs through the program's resource behavior and public outputs. Intuitively, the guarantee enforced by this type system, *resource-aware noninterference*, requires that the parts of the program affected by secret inputs can only make constant use of resources.

The main technical contribution in this part is the soundness proof of this type system with respect to the cost semantics. The main conceptional contribution is that the type system allows to freely switch between local and global reasoning. One extreme would be to ignore the information flow of the secret values and prove that the whole program has constant resource consumption. The other extreme would be to ensure that every conditional that branches on a secret value (a critical conditional) uses a constant amount of resources. However, there are constant time programs in which individual conditional are not constant time (see Section 4). As a result, we allow different levels of global and local reasoning and in the type system to ensure that every critical conditional occurs in a constant-resource block.

Finally, we show that our type inference algorithm for the constant-resource type system can be used to automatically turn programs into constant-resource programs. To this end, we introduce a *consume* expression that performs resource padding (e.g., sleep for time). The amount of resource padding that is needed is automatically determined by the LP solver and is parametric in the size of the program variables. This technique is more efficient then existing techniques [citations] since it does not change the worst-case resource behavior of the program. Of course, it would be possible to do such a resource padding to the worst-case behavior dynamically at the end of the run of the program. The advantage of our method is that we do not have to keep track of the actual resource usage at runtime and that we automatically derive a proof (a type derivation) that the modified program has constant resource use without reasoning on a meta level. We implemented this technique in RAML.

Contributions We make the following contributions:

- Two novel AARA type systems that derive lower bounds and prove constant resource use, and an implementation of these systems that extends RAML. We evaluate the implementation on several examples, including encryption routines and data processing programs that were previously studied in the context of timing leaks in differentially-private systems [HPN11].
- A mechanization of the soundness proofs the two new type systems and classic AARA for upper bounds in Agda. To the best our knowledge, this is also the first formalization of the soundness of linear AARA for worst-case bounds.
- An information-flow type system that incorporates our constant-time system to prevent

leakage of selected secrets through resource side channels, and an LP-based method that transforms programs into constant-resource versions.

2 Language-level constant-resource programs

In this section we introduce a language, an operational cost semantics, and the notion of constant-resource functions.

2.1 The language

Syntax To discuss the main ideas of our work, it is sufficient to study a purely functional first-order and monomorphic typed functional language with Booleans, integers, pairs, and list data types, pattern matching and recursive functions as given in Fig. 1. The base types of the language are given as follows.

$$\begin{aligned} \text{zero-order types:} \quad & T ::= \text{unit} \mid \text{bool} \mid \text{int} \mid L(T) \mid T * T \\ \text{first-order types:} \quad & G ::= T \to T \end{aligned}$$

In this grammar we use abstract binding trees [Har12] and in examples we use equivalent

e	$::= () \mid \mathsf{true} \mid \mathsf{false} \mid n \mid x$	$::= () \mid \mathsf{true} \mid \mathsf{false} \mid n \mid x$
	$\mid \text{op}_\diamond(x_1, x_2)$	$\mid x_1 \diamond x_2$
	$\mid \text{app}(f, x)$	$\mid f(x)$
	$\mid \text{let}(x, e_1, x.e_2)$	$\mid \text{let } x = e_1 \text{ in } e_2$
	$\mid \text{if}(x, e_t, e_f)$	$\mid \text{if } x \text{ then } e_t \text{ else } e_f$
	$\mid \text{pair}(x_1, x_2)$	$\mid (x_1, x_2)$
	$\mid \text{match}(x, (x_1, x_2).e)$	$\mid \text{match } x \text{ with } (x_1, x_2) \to e$
	$\mid \text{nil}$	$\mid []$
	$\mid \text{cons}(x_1, x_2)$	$\mid x_1 :: x_2$
	$\mid \text{match}(x, e_1, (x_1, x_2).e_2)$	$\mid \text{match } x \text{ with } \mid [] \to e_1 \mid x_1 :: x_2 \to e_2$
	$\mid \text{share}(x, (x_1, x_2).e)$	$\mid \text{share } (x_1, x_2) = x \text{ in } e$

$\diamond \in \{+, -, *, \mathsf{div}, \mathsf{mod}, =, <>, >, <, <=, >=, \mathsf{and}, \mathsf{or}\}$

Figure 1: Syntax of the language

expressions in OCaml syntax. The expressions are in *let normal form*, meaning that they are formed from variables whenever it is possible. It makes the typing rules and semantics simpler without loosing expressivity. The syntactic form *share* has to be use to introduce multiple occurrences of a variable in an expression.

A value is a boolean constant, an integer value n, the empty list nil, a list of values $[v_1, ..., v_n]$, or a pair of values (v_1, v_2). A type context $\Gamma : \text{VID} \to \mathcal{T}$ is a partial mapping from variable identifiers to data types T. A *signature* $\Sigma : \text{FID} \to \mathcal{G}$ is a partial mapping from function identifiers to first-order types G. The typing rules that define a type judgement $\Sigma; \Gamma \vdash e : T$ are standard.

A program is a tuple containing a signature Σ and a finite set of tuples $(e_g, x^g)_{g \in \text{dom}(\Sigma)}$ where e_g is an expression defining the function's body and x^g is the argument. For any e_g, it holds that $\Sigma; x^g : T_1 \vdash e_g : T_2$ if $\Sigma(g) = T_1 \to T_2$.

Operational cost semantics The operational cost semantics defines the resource consumption of programs. It is instrumented with a non-negative resource counter that is incremented or decremented by a constant at every step of the semantics. The semantics is parametric in the cost that is used at each step and we call a particular set of such cost parameters a *cost model.* The constants can be used to indicate the costs of storing or loading a value in the memory, evaluating a primitive operation, binding of a value in the environment, or branching on a Boolean value.

It is possible to further parameterize some constants to obtain a more precise cost model. For example, the cost of calling a function may vary according to the number of the arguments. In the following, we will show that any suitable values can be used for the constants in the cost model and the soundness of the type system does not rely on any specific values for these constants. In the examples, we use a cost model in which the constants are either 1 for each step or 0 for all steps except for calls to the *tick* where *tick(q)* means that we have resource resource usage $q \in \mathbb{Q}$.

The cost semantics is based on a big-step semantics and is formulated using an *environment* $E : \text{VID} \to Val$ that is a finite mapping from a set of variable identifiers to a set of values. We write $E[x \mapsto v]$ to denote a new environment that extends E by adding a new binding $x \mapsto v$.

Evaluation judgements are of the form $E \vdash^{q}_{q'} e \Downarrow v$ where $q, q' \in \mathbb{Q}^{+}_{0}$. The intuitive meaning is that under the environment E and q available resources, e evaluates to the value v without running out of resources and q' resources are available after the evaluation. The evaluation consumes $\delta = q - q'$ resource units. Fig. 2, Fig. 3, and Fig. 4 represent the typing rules for values, the base typing and the evaluation rules for the language, respectively.

$$\text{(V:UNIT)}\ \frac{v = ()}{\models v : \text{unit}} \qquad \text{(V:BOOL)}\ \frac{v \in \{\mathsf{true}, \mathsf{false}\}}{\models v : \text{bool}} \qquad \text{(V:INT)}\ \frac{v \in \mathbb{Z}}{\models v : \text{int}} \qquad \text{(V:PAIR)}\ \frac{\models v_1 : T_1 \qquad \models v_2 : T_2}{\models (v_1, v_2) : T_1 * T_2} \qquad \text{(V:NIL)}\ \frac{v = \mathsf{nil}}{\models v : L(T)}$$

$$\text{(V:LIST)}\ \frac{\models v_i : T \qquad \forall i = 1, \ldots, n}{\models [v_1, \ldots, v_n] : L(T)}$$

Figure 2: Typing rules: values

2.2 Constant-resource programs

Informally, a program is *constant resource* if it has the same quantitative resource consumption under all environments in which values have the same size.

We write $\models v : T$ to denote that v is a well-formed value of type T. The typing rules for values are standard [HJ03, HAH11, Ano16] and we omit them here. Let Γ be a context that maps variable identifiers to base types, an environment E is *well-formed* w.r.t Γ, denoted $\models E : \Gamma$, if $\forall x \in \text{dom}(\Gamma).\ \models E(x) : \Gamma(x)$.

Below we define the notation of *size equivalence*, written $|v| \approx |u|$, which is a binary relation relating two values v and u of the same type T. We write $[v_1, \ldots, v_n]$ to denote a list of n values. Note that we are only considering the sizes of inductive types here. The size of a value of a base type can be arbitrary depending on the hardware architectures (e.g., an integer value can be 32

$$\frac{}{\Sigma;\emptyset \vdash () : \text{unit}}\;(\text{T:Unit}) \qquad \frac{b \in \{\mathsf{true}, \mathsf{false}\}}{\Sigma;\emptyset \vdash b : \text{bool}}\;(\text{T:Bool}) \qquad \frac{n \in \mathbb{Z}}{\Sigma;\emptyset \vdash n : \text{int}}\;(\text{T:Int}) \qquad \frac{x \in \text{dom}(E)}{\Sigma; x : T \vdash x : T}\;(\text{T:Var})$$

$$\frac{\diamond \in \{\mathsf{and}, \mathsf{or}\}}{\Sigma; x_1 : \text{bool}, x_2 : \text{bool} \vdash \text{op}_\diamond(x_1, x_2) : \text{bool}}\;(\text{T:B-Op}) \qquad \frac{\diamond \in \{=, <>, >, <, <=, >=\}}{\Sigma; x_1 : \text{int}, x_2 : \text{int} \vdash \text{op}_\diamond(x_1, x_2) : \text{bool}}\;(\text{T:IB-Op})$$

$$\frac{\diamond \in \{+, -, *, \mathsf{div}, \mathsf{mod}\}}{\Sigma; x_1 : \text{int}, x_2 : \text{int} \vdash \text{op}_\diamond(x_1, x_2) : \text{int}}\;(\text{T:I-Op}) \qquad \frac{\Sigma(g) = T_1 \to T_2}{\Sigma; x : T_1 \vdash \text{app}(g, x) : T_2}\;(\text{T:Fun})$$

$$\frac{\Sigma;\Gamma_1 \vdash e_1 : T_1 \qquad \Sigma;\Gamma_2, x : T_1 \vdash e_2 : T_2}{\Sigma;\Gamma_1,\Gamma_2 \vdash \text{let}(x, e_1, x.e_2) : T_2}\;(\text{T:Let}) \qquad \frac{\Sigma;\Gamma \vdash e_t : T \qquad \Sigma;\Gamma \vdash e_f : T}{\Sigma;\Gamma, x : \text{bool} \vdash \text{if}(x, e_t, e_f) : T}\;(\text{T:If})$$

$$\frac{}{\Sigma; x_1 : T_1, x_2 : T_2 \vdash \text{pair}(x_1, x_2) : T_1 * T_2}\;(\text{T:Pair}) \qquad \frac{\Sigma;\Gamma, x_1 : T_1, x_2 : T_2 \vdash e : T}{\Sigma;\Gamma, x : T_1 * T_2 \vdash \text{match}(x, (x_1, x_2).e) : T}\;(\text{T:Match-P})$$

$$\frac{T \in \mathcal{T}}{\Sigma;\emptyset \vdash \mathsf{nil} : L(T)}\;(\text{T:Nil}) \qquad \frac{}{\Sigma; x_h : T, x_t : L(T) \vdash \text{cons}(x_h, x_t) : L(T)}\;(\text{T:Cons})$$

$$\frac{\Sigma;\Gamma \vdash e_1 : T_1 \qquad \Sigma;\Gamma, x_h : T, x_t : L(T) \vdash e_2 : T_1}{\Sigma;\Gamma, x : L(T) \vdash \text{match}(x, e_1, (x_h, x_t).e_2) : T_1}\;(\text{T:Match-L}) \qquad \frac{\Sigma;\Gamma, x_1 : T, x_2 : T \vdash e : T_1}{\Sigma;\Gamma, x : T \vdash \text{share}(x, (x_1, x_2).e) : T_1}\;(\text{T:Share})$$

$$\frac{\Sigma;\Gamma \vdash e : T_1}{\Sigma;\Gamma, x : T \vdash e : T_1}\;(\text{T:Weakening})$$

Figure 3: Base typing rules: language

or 64 bits).

$$\frac{T \in \{\text{unit}, \text{bool}, \text{int}\}}{|v| \approx |u|} \qquad \frac{|v_1| \approx |u_1| \qquad |v_2| \approx |u_2|}{|(v_1, v_2)| \approx |(u_1, u_2)|} \qquad \frac{m = n \qquad |v_i| \approx |u_i|}{|[v_1, \ldots v_n]| \approx |[u_1, \ldots u_m]|}$$

Let $X \subseteq \text{dom}(\Gamma)$ be a set of variables and E_1, E_2 be two well-formed environments. Then E_1 and E_2 are *size-equivalent* w.r.t X, denoted $E_1 \approx_X E_2$, when they agree on the sizes of the variables in X, that is, $\forall x \in X. |E_1(x)| \approx |E_2(x)|$.

Note that $\approx_X$ is an equivalence relation, i.e., it is reflexive, symmetric and transitive. Using size-equivalence for environments we now formally define the notation of constant-resource expressions as follows.

Definition 1. *An expression e is constant resource w.r.t $X \subseteq dom(\Gamma)$, written $const_X(e)$, if for all*

$$\frac{}{E \vdash^{q+K^{\mathsf{unit}}}_{q} () \Downarrow ()}\ (\text{E:Unit}) \qquad \frac{b \in \{\mathsf{true}, \mathsf{false}\}}{E \vdash^{q+K^{\mathsf{bool}}}_{q} b \Downarrow b}\ (\text{E:Bool}) \qquad \frac{n \in \mathbb{Z}}{E \vdash^{q+K^{\mathsf{int}}}_{q} n \Downarrow n}\ (\text{E:Int}) \qquad \frac{x \in \mathrm{dom}(E)}{E \vdash^{q+K^{\mathsf{var}}}_{q} x \Downarrow E(x)}\ (\text{E:Var})$$

$$\frac{v = E(x_1) \diamond E(x_2)}{E \vdash^{q+K^{\mathsf{op}}}_{q} \mathrm{op}_\diamond(x_1, x_2) \Downarrow v}\ (\text{E:Bin}) \qquad \frac{\Sigma(g) = T_1 \to T_2 \qquad [y^g \mapsto E(x)] \vdash^{q}_{q'} e_g \Downarrow v}{E \vdash^{q+K^{\mathsf{app}}}_{q'} \mathrm{app}(g, x) \Downarrow v}\ (\text{E:Fun})$$

$$\frac{E \vdash^{q-K^{\mathsf{let}}}_{q_1'} e_1 \Downarrow v_1 \qquad E[x \mapsto v_1] \vdash^{q_1'}_{q'} e_2 \Downarrow v}{E \vdash^{q}_{q'} \mathrm{let}(x, e_1, x.e_2) \Downarrow v}\ (\text{E:Let}) \qquad \frac{E(x) = \mathsf{true} \qquad E \vdash^{q-K^{\mathsf{cond}}}_{q'} e_t \Downarrow v}{E \vdash^{q}_{q'} \mathrm{if}(x, e_t, e_f) \Downarrow v}\ (\text{E:If-True})$$

$$\frac{E(x) = \mathsf{false} \qquad E \vdash^{q-K^{\mathsf{cond}}}_{q'} e_f \Downarrow v}{E \vdash^{q}_{q'} \mathrm{if}(x, e_t, e_f) \Downarrow v}\ (\text{E:If-False}) \qquad \frac{x_1, x_2 \in \mathrm{dom}(E) \qquad v = (E(x_1), E(x_2))}{E \vdash^{q+K^{\mathsf{pair}}}_{q} \mathrm{pair}(x_1, x_2) \Downarrow v}\ (\text{E:Pair})$$

$$\frac{}{E \vdash^{q+K^{\mathsf{nil}}}_{q} \mathsf{nil} \Downarrow \mathsf{nil}}\ (\text{E:Nil}) \qquad \frac{E(x) = (v_1, v_2) \qquad E[x_1 \mapsto v_1, x_2 \mapsto v_2] \vdash^{q-K^{\mathsf{matchP}}}_{q'} e \Downarrow v}{E \vdash^{q}_{q'} \mathrm{match}(x, (x_1, x_2).e) \Downarrow v}\ (\text{E:Match-P})$$

$$\frac{x_h, x_t \in \mathrm{dom}(E) \qquad E(x_h) = v_1 \qquad E(x_t) = [v_2, .., v_n]}{E \vdash^{q+K^{\mathsf{cons}}}_{q} \mathrm{cons}(x_h, x_t) \Downarrow [v_1, ..., v_n]}\ (\text{E:Cons}) \qquad \frac{E(x) = \mathsf{nil} \qquad E \vdash^{q-K^{\mathsf{matchN}}}_{q'} e_1 \Downarrow v}{E \vdash^{q}_{q'} \mathrm{match}(x, e_1, (x_h, x_t).e_2) \Downarrow v}\ (\text{E:Match-N})$$

$$\frac{E(x) = v_1 \qquad E[x_1 \mapsto v_1, x_2 \mapsto v_1] \setminus \{x\} \vdash^{q}_{q'} e \Downarrow v}{E \vdash^{q}_{q'} \mathrm{share}(x, (x_1, x_2).e) \Downarrow v}\ (\text{E:Share})$$

$$\frac{E(x) = [v_1, ..., v_n] \qquad E[x_h \mapsto v_1, x_t \mapsto [v_2, ..., v_n]] \vdash^{q-K^{\mathsf{matchL}}}_{q'} e_2 \Downarrow v}{E \vdash^{q}_{q'} \mathrm{match}(x, e_1, (x_h, x_t).e_2) \Downarrow v}\ (\text{E:Match-L})$$

Figure 4: Evaluation rules: language

well-formed environments E_1 and E_2 such that $E_1 \approx_X E_2$, the following statement holds.

$$\textit{If } E_1 \vdash^{p_1}_{p_1'} e \Downarrow v_1 \textit{ and } E_2 \vdash^{p_2}_{p_2'} e \Downarrow v_2 \textit{ then } p_1 - p_1' = p_2 - p_2'$$

We say that a function $g(x_1, \ldots, x_n) = e_g$ is constant resource w.r.t $X \subseteq \{x_1, \ldots, x_n\}$ if $\mathrm{const}_X(e_g)$. If $Y \subseteq X$ and $E_1 \approx_X E_2$ then $E_1 \approx_Y E_2$. Thus we have the following lemma.

Lemma 1. *For all e, X, and $Y \subseteq X$, if $const_Y(e)$ then $const_X(e)$.*

```
let rec compare(h,l) = match h with
  | [] → (match l with | [] → Raml.tick 1.0; true
    | y::ys → Raml.tick 1.0; false)
  | x::xs → match l with | [] → Raml.tick 1.0; false
    | y::ys → if (x = y) then
        Raml.tick 5.0; compare(xs,ys)
      else Raml.tick 5.0; false

let rec p_compare(h,l) =
 let rec aux(r,h,l) = match h with
  | [] → (match l with | [] → Raml.tick 1.0; r
    | y::ys → Raml.tick 1.0; false)
  | x::xs → match l with | [] → Raml.tick 1.0; false
    | y::ys → if (x = y) then
        Raml.tick 5.0; aux(r,xs,ys)
      else Raml.tick 5.0; aux(false,xs,ys)
  in aux(true,h,l)
```

Figure 5: The list comparison function *compare* is not constant resource, while the manually padded function *p_compare* is constant resource w.r.t h and l.

Example The function *compare* in Fig. 5 is not constant-resource function w.r.t h and l when the cost model is defined using *tick* annotations. Since the execution cost of the two branches of the conditional dependents on the relation of x and y. The function *p_compare* is a manually padded version with a dummy computation that is constant w.r.t h and l. However, it is not constant w.r.t h. For instance, *p_compare([1;2;3],[0;1;2])* has cost 16 but *p_compare([1;2;1],[0;1])* has cost $12 \neq 16$. If we further pad the nil case with *Raml.tick 5.0; aux false xs []* to make the function to always iterate all of h's nodes, then it is constant w.r.t h.

In Section 5, we will provide a better way to transform a program into constant with our extended expression *consume*. Users insert *consume* expressions into program-under-consideration then our analyzer will infer automatically the amount of resource units needed to spend to make the program constant.

3 Type systems for lower bounds and constant resource usage

In this section we introduce two substructural resource-annotated type systems: The type system for constant resource usage is linear and the one for lower bounds is relevant.

3.1 Background

Amortized analysis The potential method of amortized analysis has been introduced [Tar85] to bound the *worst-case resource usage* of a sequence of data structure operations. The key idea is to incorporate a non-negative potential into the analysis that can be used to pay (costly) operations.

To statically analyze a program with the potential method, a mapping from program points to potentials must be established. One has to show that the potential at every program point suffices to cover the cost of any possible evaluation step and the potential of the next program

```
let rec filter_succ l = match l with
  | [] → Raml.tick 1.0; []
  | x::xs →
    if x > 0 then Raml.tick 8.0; filter_succ xs
    else Raml.tick 3.0; (x+1)::filter_succ xs

let fs_twice l = filter_succ (filter_succ l)
```

Figure 6: Two OCaml functions with linear resource usage. The *worst-case* number of *ticks* executed by *fitler_succ(ℓ)* and *fs_twice(ℓ)* is $8|\ell|+1$ and $11|\ell|+2$ respectively. In the *best-case* the functions execute $3|\ell|+1$ and $6|\ell|+2$ ticks, respectively. The resource consumption is not constant.

point. The initial potential is then an upper bound on the resource usage of the program.

Linear potential for upper bounds To automate amortized analysis, we fix a format of the potential functions and use LP solving to find the optimal coefficients. To infer linear potential functions, inductive data types are annotated with a non-negative rational numbers q [HJ03]. For example, the type $L^q(\mathsf{bool})$ of Boolean lists with potential q defines potential $\Phi([b_1,\ldots,b_n] : L^q(\mathsf{bool})) = q{\cdot}n$. Type rules statically verify that the initial potential is sufficient to cover the operations over the data structure for any possible evaluation of the program.

This idea is best explained by example. Consider the function *filter_succ* below that filters out positive numbers and increments non-positive numbers. As in RAML, we use OCaml syntax and *tick* commands to specify resource usage. If we filter out a number then we have a high cost (8 resource units) since *x* is, e.g., sent to an external device. If *x* is incremented we have a lower cost of 3 resource units. As a result, the worst-case resource consumption of *filter_succ(ℓ)* is $8|\ell|+1$ (where 1 is for the cost that occurs in the nil case of the match). The function *fs_twice(ℓ)* applies *filter_succ* twice, to ℓ and to the result of *filter_succ(ℓ)*. The worst-case behavior appears if no list element is filtered out in the first call and all elements are filtered out in the second call. The worst-case behavior is thus $11|\ell|+2$. These upper bounds can be expressed with the following annotated function types, which can be derived using local type rules in Fig. 7.

$$\begin{aligned} \mathsf{filter\_succ}: &\quad L^8(\text{int}) \xrightarrow{1/0} L^0(\text{int}) \\ \mathsf{fs\_twice}: &\quad L^{11}(\text{int}) \xrightarrow{2/0} L^0(\text{int}) \end{aligned}$$

Intuitively, the first function type states that an initial potential of $8|\ell|+1$ is sufficient to cover the cost of *filter_succ(ℓ)* and there is $0|\ell'|+0$ potential left where ℓ' is the result of the computation. This is just one possible potential annotation of many. The right choice of the potential annotation depends on the use of the function result. For example, for the inner call of *filter_succ* in *fs_twice* we need the following annotation.

$$\mathsf{filter\_succ}: \quad L^{11}(\text{int}) \xrightarrow{2/1} L^8(\text{int})$$

It states that the initial potential of $11|\ell|+2$ is sufficient to cover the cost of *filter_succ(ℓ)* and there is $8|\ell'|+1$ potential left to be assigned to the returned list ℓ'. The potential of the result can then be used with the previous type of *filter_succ* to pay for the cost of the outer call.

$$\mathsf{filter\_succ}: \quad L^p(\text{int}) \xrightarrow{q/q'} L^r(\text{int}) \mid q \geq q'+1 \wedge p \geq 8 \wedge p \geq 3+r$$

We can summarize all possible types of *filter_succ* with a linear constraint system. In the type inference, we generate such a constraint system and solve it with an off-the-shelf LP solver to derive a concrete bound. To obtain tight bounds, we perform a whole-program analysis and minimize the coefficients in the input potential.

Surprisingly, this approach—as well as the new concepts we introduce here—can be extended to polynomial bounds [HAH12], higher-order functions [JHLH10, Ano15], polymorphism [JLH+09], and user-defined inductive types [JLH+09, Ano15].

3.2 Resource annotations

To explain the new ideas of this work, we focus on linear potential annotations for lists in a first-order language. However, the results extend to and have been implement with multivariate polynomial potential, user-defined data types and higher-order functions.

The resource-annotated types are base types in which the inductive data types are annotated with non-negative rational numbers, called *resource annotations*. These numbers represent a potential per element in the list that can be used to "pay" for resource consumption during the evaluation of a program. The annotated data types of the language are given as follows.

$$A ::= \text{unit} \mid \text{bool} \mid \text{int} \mid L^p(A) \mid A * A \quad (\text{for } p \in \mathbb{Q}_0^+)$$

Let $\mathcal{A}$ be the set of resource-annotated data types. A type context, $\Gamma^r : \text{VID} \to \mathcal{A}$, is a partial mapping from variable identifiers to resource-annotated types. The *underlying* base type and base type context denoted by $\widehat{A}$, and $\widehat{\Gamma^r}$ respectively can be obtained by removing the annotations. We extend all definitions such as $|v|$, $\models E : \Gamma$ and $\approx$ for base data types to resource-annotated data types by ignoring the annotations.

We now formally define the notation of *potential* representing how resource is associated with runtime values. The potential of a value v of type A, written $\Phi(v : A)$, is defined by the function $\Phi : \mathit{Val} \to \mathbb{Q}_0^+$ as follows.

$$\begin{aligned}
&\Phi(() : \text{unit}) = \Phi(b : \text{bool}) = \Phi(n : \text{int}) = 0\\
&\Phi((v_1, v_2) : A_1 * A_2) = \Phi(v_1 : A_1) + \Phi(v_2 : A_2)\\
&\Phi([v_1, \cdots, v_n] : L^p(A)) = n\cdot p + \Sigma_{i=1}^n \Phi(v_i : A)
\end{aligned}$$

Example The potential of a list $v = [b_1, \cdots, b_n]$ of type $L^p(\text{bool})$ is $n\cdot p$. Similarly, a list of lists of Booleans values $v = [v_1, \cdots, v_n]$ of type $L^p(L^q(\text{bool}))$, where $v_i = [b_{i1}, \cdots, b_{im_i}]$, has the potential $n\cdot p + (m_1 + \cdots + m_n)\cdot q$.

Let Γ^r be a context and E be a well-formed environment w.r.t Γ^r. The potential of $X \subseteq \text{dom}(\Gamma^r)$ under E is defined as $\Phi_E(X : \Gamma^r) = \Sigma_{x \in X} \Phi(E(x) : \Gamma^r(x))$. The potential of Γ^r is $\Phi_E(\Gamma^r) = \Phi_E(\text{dom}(\Gamma^r) : \Gamma^r)$. Note that if $x \notin X$ then $\Phi_E(X)) = \Phi_{E[x \mapsto v]}(X)$. The following lemma states that the potential is the same under two well-formed size-equivalent environments.

Lemma 2. *If* $E_1 \approx_X E_2$ *then* $\Phi_{E_1}(X : \Gamma^r) = \Phi_{E_2}(X : \Gamma^r)$.

Annotated first-order data types are given as follows, where q and q' are rational numbers.

$$F ::= A_1 \xrightarrow{q/q'} A_2$$

Let $\mathcal{F}$ be the set of the annotated first-order types. A resource-annotated signature $\Sigma^r : \text{FID} \to \wp(\mathcal{F}) \setminus \{\emptyset\}$ is a partial mapping from function identifiers to a non-empty sets of annotated first-order types. That means a function can have different resource annotations depending on the context. The *underlying* base types are denoted by $\widehat{F}$. and the underlying base signature is denoted by $\widehat{\Sigma^r}$ where $\widehat{\Sigma^r}(f) = \widehat{\Sigma^r(f)}$.

$$\text{(A:UNIT)}\quad \frac{}{\Sigma^r;\varnothing \mid\frac{K^{\mathsf{unit}}}{0}\; ():\mathrm{unit}} \qquad \text{(A:BOOL)}\quad \frac{b \in \{\mathsf{true},\mathsf{false}\}}{\Sigma^r;\varnothing \mid\frac{K^{\mathsf{bool}}}{0}\; b:\mathrm{bool}} \qquad \text{(A:INT)}\quad \frac{n \in \mathbb{Z}}{\Sigma^r;\varnothing \mid\frac{K^{\mathsf{int}}}{0}\; n:\mathrm{int}} \qquad \text{(A:VAR)}\quad \frac{}{\Sigma^r;x:A \mid\frac{K^{\mathsf{var}}}{0}\; x:A}$$

$$\text{(A:B-OP)}\quad \frac{\diamond \in \{\mathsf{and},\mathsf{or}\}}{\Sigma^r;x_1:\mathrm{bool},x_2:\mathrm{bool} \mid\frac{K^{\mathsf{op}}}{0}\; \mathrm{op}_\diamond(x_1,x_2):\mathrm{bool}} \qquad \text{(A:IB-OP)}\quad \frac{\diamond \in \{=,<>,>,<,<=,>=\}}{\Sigma^r;x_1:\mathrm{int},x_2:\mathrm{int} \mid\frac{K^{\mathsf{op}}}{0}\; \mathrm{op}_\diamond(x_1,x_2):\mathrm{bool}}$$

$$\text{(A:I-OP)}\quad \frac{\diamond \in \{+,-,*,\mathsf{div},\mathsf{mod}\}}{\Sigma^r;x_1:\mathrm{int},x_2:\mathrm{int} \mid\frac{K^{\mathsf{op}}}{0}\; \mathrm{op}_\diamond(x_1,x_2):\mathrm{int}} \qquad \text{(A:FUN)}\quad \frac{\Sigma^r(f) = A_1 \xrightarrow{q/q'} A_2}{\Sigma^r;x:A_1 \mid\frac{q+K^{\mathsf{app}}}{q'}\; \mathrm{app}(f,x):A_2}$$

$$\text{(A:LET)}\quad \frac{\Sigma^r;\Gamma_1^r \mid\frac{q-K^{\mathsf{let}}}{q_1'}\; e_1:A_1 \qquad \Sigma^r;\Gamma_2^r,x:A_1 \mid\frac{q_1'}{q'}\; e_2:A_2}{\Sigma^r;\Gamma_1^r,\Gamma_2^r \mid\frac{q}{q'}\; \mathrm{let}(x,e_1,x.e_2):A_2}$$

$$\text{(A:IF)}\quad \frac{\Sigma^r;\Gamma^r \mid\frac{q-K^{\mathsf{cond}}}{q'}\; e_t:A \qquad \Sigma^r;\Gamma^r \mid\frac{q-K^{\mathsf{cond}}}{q'}\; e_f:A}{\Sigma^r;\Gamma^r,x:\mathrm{bool} \mid\frac{q}{q'}\; \mathrm{if}(x,e_t,e_f):A}$$

$$\text{(A:PAIR)}\quad \frac{}{\Sigma^r;x_1:A_1,x_2:A_2 \mid\frac{K^{\mathsf{pair}}}{0}\; \mathrm{pair}(x_1,x_2):A_1 * A_2} \qquad \text{(A:MATCH-P)}\quad \frac{\Sigma^r;\Gamma^r,x_1:A_1,x_2:A_2 \mid\frac{q-K^{\mathsf{matchP}}}{q'}\; e:A}{\Sigma^r;\Gamma^r,x:A_1 * A_2 \mid\frac{q}{q'}\; \mathrm{match}(x,(x_1,x_2).e):A}$$

$$\text{(A:NIL)}\quad \frac{A \in \mathcal{A}}{\Sigma^r;\varnothing \mid\frac{K^{\mathsf{nil}}}{0}\; \mathsf{nil}:L^p(A)} \qquad \text{(A:CONS)}\quad \frac{}{\Sigma^r;x_h:A,x_t:L^p(A) \mid\frac{p+K^{\mathsf{cons}}}{0}\; \mathrm{cons}(x_h,x_t):L^p(A)}$$

$$\text{(A:MATCH-L)}\quad \frac{\Sigma^r;\Gamma^r \mid\frac{q-K^{\mathsf{matchN}}}{q'}\; e_1:A_1 \qquad \Sigma^r;\Gamma^r,x_h:A,x_t:L^p(A) \mid\frac{q+p-K^{\mathsf{matchL}}}{q'}\; e_2:A_1}{\Sigma^r;\Gamma^r,x:L^p(A) \mid\frac{q}{q'}\; \mathrm{match}(x,e_1,(x_h,x_t).e_2):A_1}$$

$$\text{(A:SHARE)}\quad \frac{\Sigma^r;\Gamma^r,x_1:A_1,x_2:A_2 \mid\frac{q}{q'}\; e:B \qquad \curlyvee(A \mid A_1,A_2)}{\Sigma^r;\Gamma^r,x:A \mid\frac{q}{q'}\; \mathrm{share}(x,(x_1,x_2).e):B}$$

Figure 7: Common typing rules: upper bounds, constant, and lower bounds

3.3 Type system for constant resource consumption

The typing rules of the constant-resource type system define judgments of the form

$$\Sigma^r;\Gamma^r \mid\frac{q}{q'}\; e:A$$

where e is an expression and $q, q' \in \mathbb{Q}_0^+$. The intended meaning is that in the environment E, $q + \Phi_E(\Gamma^r)$ resource units are sufficient to evaluate e to a value v with type A and there are *exactly* $q' + \Phi(v : A)$ resource units left over.

The typing rules form a *linear* type system. It ensures that every variable is used exactly once by allowing exchange but not weakening or contraction [Wal02]. The rules can be organized into syntax directed and structural rules.

Syntax-directed rules The syntax-directed rules listed in Fig. 7 are shared among all type systems. Rules like A:VAR and A:B-OP for leaf expressions (e.g., variable, binary operations, pairs) have fixed costs as specified by the constants K^x. Note that we require all available potential to be spent. The cost of the function call is represented by the constant K^{app} in the rule A:FUN and the argument carries the potential to pay for the function execution. In the rule A:LET, the cost of binding is represented by the constant K^{let}. The potentials carried by the contexts Γ_1^r and Γ_2^r are passed sequentially through the sub derivations. Note that the contexts are disjoint since our type system is linear. Multiple uses of variables must be introduced through the rule A:SHARE. The rule A:IF is the key rule for ensuring constant resource usage. By using the same context Γ^r for typing both e_t and e_f, we ensure that the conditional expression has the same resource usage in size-equivalent environments independent of the value of the Boolean variable x. The rules for inductive data types are crucial for the interaction of the linear potential annotations with the constant potential. The rule A:CONS shows how constant potential can be associated with a new data structure. The dual is the rule A:MATCH-L, which shows how potential associated with data can be released. It is important that these transitions are made in a linear fashion: potential is neither lost or gained.

Sharing The *share expression* makes multiple uses of a variable explicit. While multiple uses of a variable seem to be in conflict with the linear type discipline, the *sharing relation* $\curlyvee(A \mid A_1, A_2)$ ensures that potential is treated in a linear way. It apportions potential to ensure that the total potential associated with all uses is equal to the potential initially associated with the variable. This relation is only defined for structurally-identical types which differ in at most the resource annotations as follows.

$$\frac{A \in \{\text{unit}, \text{bool}, \text{int}\}}{\curlyvee(A \mid A, A)} \qquad \frac{\curlyvee(A \mid A_1, A_2) \qquad p = p_1 + p_2}{\curlyvee(L^p(A) \mid L^{p_1}(A_1), L^{p_2}(A_2))} \qquad \frac{\curlyvee(A \mid A_1, A_2) \qquad \curlyvee(B \mid B_1, B_2)}{\curlyvee(A * B \mid A_1 * B_1, A_2 * B_2)}$$

Lemma 3. *If* $\curlyvee(A \mid A_1, A_2)$ *then* $\widehat{A} = \widehat{A_1} = \widehat{A_2}$ *and* $\forall v.\ \Phi(v : A) = \Phi(v : A_1) + \Phi(v : A_2)$

Proof. By induction on the definition of the sharing relation. □

Structural rules To allow more programs to be typed we add two structural rules to the type system which can be applied to every expression. These rules are specific to the the constant-resource type system.

(C:WEAKENING)
$$\frac{\Sigma^r; \Gamma^r \mathrel{\substack{q \\ \vdash \\ q'}} e : B \qquad \curlyvee(A \mid A, A)}{\Sigma^r; \Gamma^r, x : A \mathrel{\substack{q \\ \vdash \\ q'}} e : B}$$

(C:RELAX)
$$\frac{\Sigma^r; \Gamma^r \mathrel{\substack{p \\ \vdash \\ p'}} e : A \qquad q \geq p \qquad q - p = q' - p'}{\Sigma^r; \Gamma^r \mathrel{\substack{q \\ \vdash \\ q'}} e : A}$$

The rule C:RELAX reflects the fact that if it is sufficient to evaluate e with p available resource units and there are p' resource units left over then e can be evaluated with $p+c$ resource units and there are exactly $p'+c$ resource left over, where $c \in \mathbb{Q}_0^+$. Rule C:WEAKENING states that an extra variable can be added into the given context if its potential is zero. The condition is enforced by $\curlyvee(A \mid A, A)$ since $\Phi(v:A) = \Phi(v:A) + \Phi(v:A)$ or $\Phi(v:A) = 0$. The rules can be used in branchings such as the conditional or the pattern match to ensure that subexpressions are typed using the same contexts and potential annotations.

Example Consider again the function *p_compare* in Fig. 5 in which the resource consumption is defined using *tick* annotations. The resource usage of *p_compare(h, ℓ)* is constant w.r.t h, that is, it is exactly $5|h|+1$. This can be reflected by the following type.

$$\mathsf{p\_compare}: \quad (L^5(\text{int}), L^0(\text{int})) \xrightarrow{1/0} \text{bool}$$

It can be understood as follows. If the input list h carries 5 potential units per element then it is sufficient to cover the cost of *p_compare(h, ℓ)*, no potential is wasted, and 0 potential is left.

Soundness That soundness theorem states that if e is well-typed in the resource type system and it evaluates to a value v then the difference between the initial and the final potential is the net resources usage. Moreover, if the potential annotations of the return value and all variables not belonging to a set $X \subseteq \text{dom}(\Gamma^r)$ are zero then e is constant-resource w.r.t X. We write $\text{const}_X(e)$ if $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : A$, $\curlyvee(A \mid A, A)$, and $\forall x \in \text{dom}(\Gamma^r) \setminus X.\ \curlyvee(\Gamma^r(x) \mid \Gamma^r(x), \Gamma^r(x))$.

Theorem 1. *If $\models E : \Gamma^r$, $E \vdash e \Downarrow v$, and $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : A$, then for all $p, r \in \mathbb{Q}_0^+$ such that $p = q + \Phi_E(\Gamma^r) + r$, there exists $p' \in \mathbb{Q}_0^+$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$ and $p' = q' + \Phi(v : A) + r$.*

Proof. The proof is done by induction on the length of the derivation of the evaluation judgment and the typing judgment with lexical order, in which the derivation of the evaluation judgment takes priority over the typing derivation. We need to do induction on the length of both evaluation and typing derivations since on one hand, an induction of only typing derivation would fail for the case of function application, which increases the length of the typing derivation, while the length of the evaluation derivation never increases. On the other hand, if the rule C:WEAKENING is final step in the derivation, then the length of typing derivation decreases, while the length of evaluation derivation is unchanged.

A:SHARE Assume that the typing derivation ends with an application of the rule A:SHARE, thus $\Sigma^r;\Gamma^r, x_1 : A_1, x_2 : A_2 \vdash^{q}_{q'} e : B$ and $\curlyvee(A|A_1, A_2)$.

Let $E_1 = E \setminus \{x\} \cup \{[x_1 \mapsto E(x), x_2 \mapsto E(x)]\}$. Since $\models E : \Gamma^r, x : A$ and following the property of the share relation we have $\models E_1 : \Gamma^r, x_1 : A_1, x_2 : A_2$. By the induction hypothesis for e, it holds that for all $p, r \in \mathbb{Q}_0^+$ such that $p = q + \Phi_{E_1}(\Gamma^r, x_1 : A_1, x_2 : A_2) + r$, there exists $p' \in \mathbb{Q}_0^+$ satisfying $E_1 \vdash^{p}_{p'} e \Downarrow v$ and $p' = q' + \Phi(v : B) + r$.

Because $\Phi(E(x) : A) = \Phi(E_1(x_1) : A_1) + \Phi(E_1(x_2) : A_2)$ and $\Phi_E(\Gamma^r) = \Phi_{E_1}(\Gamma^r) = \Phi_{E\setminus\{x\}}(\Gamma^r)$, thus $p = q + \Phi_E(\Gamma^r, x : A) + r$ and there exists p' satisfying $E \vdash^{p}_{p'} \text{share}(x, (x_1, x_2).e) \Downarrow v$.

C:WEAKENING Suppose that the typing derivation ends with an application of the rule C:WEAKENING. Thus we have $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : B$, in which the data type A satisfies $\curlyvee(A \mid A, A)$.

Since $\models E : \Gamma^r, x : A$, it follows $\models E : \Gamma^r$. By the induction hypothesis for e, it holds that for all $p, r \in \mathbb{Q}_0^+$ such that $p = q + \Phi_E(\Gamma^r) + r$, there exists $p' \in \mathbb{Q}_0^+$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$ and $p' = q' + \Phi(v :$

$B) + r$. By the property of the share relation, $\Phi(a : A) = 0$, then we have $p = q + \Phi_E(\Gamma^r, x : A) + r$, $E \vdash^{p}_{p'} e \Downarrow v$ and $p' = q' + \Phi(v : B) + r$ as required.

C:RELAX Suppose that the typing derivation ends with an application of the rule C:RELAX, thus we have $\Sigma^r; \Gamma^r \vdash^{q_1}_{q_1'} e : A$, $q \geq q_1$, and $q - q_1 = q' - q_1'$.

For all $p, r \in \mathbb{Q}_0^+$ such that $p = q + \Phi_E(\Gamma^r) + r = q_1 + \Phi_E(\Gamma^r) + (q - q_1) + r$, we have $\models E : \Gamma^r$. By the induction hypothesis for e in the premise, there exists $p' \in \mathbb{Q}_0^+$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$ and $p' = q_1' + \Phi(v : A) + (q - q_1) + r = q' + \Phi(v : A) + r$.

A:VAR Assume that e is a variable x. If $\Sigma^r; x : A \vdash^{K^{\mathsf{var}}}_{0} x : A$. Thus for all $p, r \in \mathbb{Q}_0^+$ such that $p = K^{\mathsf{var}} + \Phi(v : A) + r$, there exists $p' = \Phi(v : A) + r$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$.

A:UNIT It is similar to the case A:VAR.

A:BOOL It is similar to the case A:VAR.

A:INT It is similar to the case A:VAR.

A:B-OP Assume that e is an expression of the form $\mathrm{op}_\diamond(x_1, x_2)$, where $\diamond = \{\mathsf{and}, \mathsf{or}\}$. Thus $\Sigma^r; x_1 : \mathrm{bool}, x_2 : \mathrm{bool} \vdash^{K^{\mathsf{op}}}_{0} e : \mathrm{bool}$ and $\models E : \{x_1 : \mathrm{bool}, x_2 : \mathrm{bool}\}$. We have $E \vdash^{K^{\mathsf{op}}}_{0} e \Downarrow v$, thus for all $p, r \in \mathbb{Q}_0^+$ such that $p = K^{\mathsf{op}} + r = K^{\mathsf{op}} + \Phi_E(x_1 : \mathrm{bool}, x_2 : \mathrm{bool}) + r$, there exists $p' = \Phi(v : \mathrm{bool}) + r = r$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$.

A:I-OP It is similar to the case A:B-OP.

A:IB-OP It is similar to the case A:B-OP.

A:CONS If e is of the form $\mathrm{cons}(x_1, x_2)$, then the type derivation ends with an application of the rule A:CONS and the evaluation ends with the application of the rule E:CONS. Thus $\Sigma^r; x_1 : A, x_2 : L^{p_1}(A) \vdash^{p_1 + K^{\mathsf{cons}}}_{0} e : L^{p_1}(A)$ and $\models E : \{x_1 : A, x_2 : L^{p_1}(A)\}$.

We have $E \vdash^{K^{\mathsf{cons}}}_{0} e \Downarrow [v_1, ..., v_n]$, where $E(x_1) = v_1$ and $E(x_2) = [v_2, \cdots, v_n]$. Let $\Gamma^r = x_h : A, x_t : L^{p_1}(A)$, for all $p, r \in \mathbb{Q}_0^+$ such that $p = p_1 + K^{\mathsf{cons}} + \Phi_E(\Gamma^r) + r$, there exists $p' \in \mathbb{Q}_0^+$ satisfying $p' = \Phi([v_1, ..., v_n] : L^{p_1}(A)) + r = \Phi_E(\Gamma^r) + p_1 + r$ and $E \vdash^{p}_{p'} e \Downarrow [v_1, ..., v_n]$.

A:PAIR It is similar to the case A:CONS.

A:NIL It is similar to the case A:CONS.

A:MATCH-P Suppose that the typing derivation $\Sigma^r; \Gamma^r, x : A_1 * A_2 \vdash^{q}_{q'} \mathrm{match}(x, (x_1, x_2).e) : A$ ends with an application of the rule A:MATCH-P. Thus $\Sigma^r; \Gamma^r, x_1 : A_1, x_2 : A_2 \vdash^{q - K^{\mathsf{matchP}}}_{q'} e : A$ and $\models E : \Gamma^r, x : A_1 * A_2$.

Let $E_1 = E[x_1 \mapsto v_1, x_2 \mapsto v_2]$ and $\Gamma_1^r = \Gamma^r, x_1 : A_1, x_2 : A_2$. Since $\models v_1 : A_1$, $\models v_2 : A_2$, and $\models E : \Gamma^r$ it holds that $\models E_1 : \Gamma_1^r$. For all $p, r \in \mathbb{Q}_0^+$ such that $p = q + \Phi_E(\Gamma^r, x : A_1 * A_2) + r$, thus $p - K^{\mathsf{matchP}} = q - K^{\mathsf{matchP}} + \Phi_{E_1}(\Gamma_1^r) + r$, by the induction hypothesis for e, there exists $p' \in \mathbb{Q}_0^+$

satisfying $p' = q' + \Phi(v : A) + r$ and $E_1 \vdash^{p-K^{\mathsf{matchP}}}_{p'} e \Downarrow v$. Hence, by the rule E:MATCH-P, there exists $p' = q' + \Phi(v : A) + r$ satisfying $E \vdash^{p}_{p'} \mathrm{match}(x, (x_1, x_2).e) \Downarrow v$.

A:FUN Assume that e is a function application of the form $\mathrm{app}(f, x)$. Thus $\Sigma^r; x : A_1 \vdash^{q+K^{\mathsf{app}}}_{q'} e : A_2$ and $\Sigma^r(f) = A_1 \xrightarrow{q/q'} A_2$. Because the considering program is well-formed, there exists a well-typed expression e_f under the typing context $\Gamma^r_1 = y^{\hat{f}} : A_1$ and the signature Σ^r, or $\Sigma^r; \Gamma^r_1 \vdash^{q}_{q'} e_f : A_2$.

Let $\Gamma^r = x : A_1$, $E(x) = v_1$ and $E_1 = [y^{\hat{f}} \mapsto v_1]$, since $\models E : \Gamma^r$, it follows that $\models E_1 : \Gamma^r_1$. For all $p, r \in \mathbb{Q}^+_0$ such that $p = q + K^{\mathsf{app}} + \Phi_E(\Gamma^r) + r$, since $\Phi_{E_1}(\Gamma^r_1) = \Phi(E_1(y^{\hat{f}}) : A_1) = \Phi_E(\Gamma^r) = \Phi(E(x) : A_1)$, it holds that $p - K^{\mathsf{app}} = q + \Phi_{E_1}(\Gamma^r_1) + r$. By the induction hypothesis for e_f, there exists $p' \in \mathbb{Q}^+_0$ satisfying $p' = q' + \Phi(v : A_2) + r$ and $E_1 \vdash^{p_1}_{p'_1} e_{\hat{f}} \Downarrow v$. Hence, $E \vdash^{p}_{p'} e \Downarrow v$.

A:IF Suppose that e is an expression of the form $\mathrm{if}(x, e_t, e_f)$. Then one of the rules E:IF-TRUE and E:IF-FALSE has been applied in the evaluation derivation depending on the value of x.

Assume that the variable x is assigned the value true in E, or $E(x) = \mathsf{true}$. The typing rule for e has been derived by an application of the rule A:IF using the premise on the left thus $\Sigma^r; \Gamma^r \vdash^{q-K^{\mathsf{cond}}}_{q'} e_t : A$.

Let $\Gamma^r_1 = \Gamma^r, x : \mathrm{bool}$, since $\models E : \Gamma^r_1$, it follows that $\models E : \Gamma^r$. For all $p, r \in \mathbb{Q}^+_0$ such that $p = q + \Phi_E(\Gamma^r_1) + r$, since $\Phi_E(\Gamma^r) = \Phi_E(\Gamma^r_1)$ thus $p_1 = p - K^{\mathsf{cond}} = q - K^{\mathsf{cond}} + \Phi_E(\Gamma^r) + r$. By the induction hypothesis for e_t, there exists $p'_1 \in \mathbb{Q}^+_0$ satisfying $E \vdash^{p_1}_{p'_1} e_t \Downarrow v$ and $p'_1 = q' + \Phi(v : A)$. Hence, by the rule E:IF-TRUE, there exists $p' = p'_1$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$ and $p' = q' + \Phi(v : A)$. If x is assigned the value false in E then it is similar to the case $E(x) = \mathsf{true}$.

A:MATCH-L It is the same as the case of a conditional expression. The evaluation derivation applies one of the rules E:MATCH-N and E:MATCH-L depending on the value of x.

Assume that x is assigned the value $[v_1, ..., v_n]$ under E, or $E(x) = [v_1, ..., v_n]$. Then, the evaluation derivation ends with an application of the rule E:MATCH-L. Let $E_1 = E[x_h \mapsto v_1, x_t \mapsto [v_2, ..., v_n]]$ and $\Gamma^r_1 = \Gamma^r, x_h : A, x_t : L^{p_1}(A)$, the typing derivation ends with an application of the rule A:MATCH-L, thus $\Sigma^r; \Gamma^r_1 \vdash^{q+p_1-K^{\mathsf{matchL}}}_{q'} e_2 : A_1$.

Since $\models [v_1, ..., v_n] : L^{p_1}(A)$, we have $\models v_i : A, \forall i = 1, ..., n$. Hence, it holds that $\models v_1 : A$ and $\models [v_2, ..., v_n] : L^{p_1}(A)$. Finally, we have $\models E_1 : \Gamma^r_1$ (since $\models E : \Gamma^r$ implies $\models E_1 : \Gamma^r$).

For all $p, r \in \mathbb{Q}^+_0$ such that $p = q + \Phi_E(\Gamma^r, x : L^{p_1}(A)) + r$, because $\Phi_E(\Gamma^r, x : L^{p_1}(A)) = \Phi_E(\Gamma^r) + n.p_1 + \Sigma^n_{i=1}\Phi(v_i : A)$, $\Phi_{E_1}(\Gamma^r_1) = \Phi_{E_1}(\Gamma^r) + (n-1).p_1 + \Sigma^n_{i=1}\Phi(v_i : A)$ and $\Phi_{E_1}(\Gamma^r) = \Phi_E(\Gamma^r)$, thus we have $\Phi_{E_1}(\Gamma^r_1) = \Phi_E(\Gamma^r, x : L^{p_1}(A)) - p_1$. Thus $p_2 = p - K^{\mathsf{matchL}} = q + p_1 - K^{\mathsf{matchL}} + \Phi_{E_1}(\Gamma^r_1) + r$. By the induction hypothesis for e_2, there exists $p'_2 \in \mathbb{Q}^+_0$ satisfying $E_1 \vdash^{p_2}_{p'_2} e_2 \Downarrow v$ and $p'_2 = q' + \Phi(v : A_1)$. Hence, there exists $p' = p'_2$ such that $E \vdash^{p}_{p'} e \Downarrow v$. If $E(x) = \mathsf{nil}$ then it is similar to the case A:MATCH-P.

A:LET Assume that e is an expression of the form $\mathrm{let}(x, e_1, x.e_2)$. Hence, the evaluation derivation ends with an application of the rule E:LET. Let $E_1 = E[x \mapsto v_1]$ and $\Gamma^r = \Gamma^r_1, \Gamma^r_2$. The typing derivation ends with an application of the rule A:LET, thus $\Sigma^r; \Gamma^r_1 \vdash^{q-K^{\mathsf{let}}}_{q'_1} e_1 : A_1$ and $\Sigma^r; \Gamma^r_2, x : A_1 \vdash^{q'_1}_{q'} e_2 : A_2$.

For all $p, r \in \mathbb{Q}_0^+$ such that $p = q+\Phi_E(\Gamma^r)+r$, thus $p_1 = p-K^{\text{let}} = q-K^{\text{let}}+\Phi_E(\Gamma_1^r)+\Phi_E(\Gamma_2^r)+r$. Since $\models E:\Gamma^r$, we have $\models E:\Gamma_1^r$. By the induction hypothesis for e_1, there exists $p_1' \in \mathbb{Q}_0^+$ satisfying $E \vdash^{p_1}_{p_1'} e_1 \Downarrow v_1$ and $p_1' = q_1' + \Phi(v_1 : A_1) + \Phi_E(\Gamma_2^r) + r$.

We have $\models E:\Gamma_2^r$, thus $\models E_1 : \Gamma_2^r, x : A_1$. Again by the induction hypothesis for e_2, with $p_2 = p-K^{\text{let}}-(p_1-p_1') = p_1' = q_1'+\Phi_{E_1}(\Gamma_2^r, x:A_1)+r$, there exists $p_2' \in \mathbb{Q}_0^+$ satisfying $E_1 \vdash^{p_2}_{p_2'} e_2 \Downarrow v$ and $p_2' = q' + \Phi(v : A_2) + r$. Hence, by the rule E:LET, there exists $p' = p_2'$ satisfying $E \vdash^{p}_{p'} e \Downarrow v$ and $p' = q' + \Phi(v : A_2)$. □

Theorem 2. *If* $\models E:\Gamma^r$, $E \vdash e \Downarrow v$, $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : A$, $\curlyvee(A \mid A, A)$*, and* $\forall x \in dom(\Gamma^r) \setminus X. \curlyvee(\Gamma^r(x) \mid \Gamma^r(x), \Gamma^r(x))$ *then e is constant resource w.r.t* $X \subseteq dom(\Gamma^r)$.

Proof. First, we prove that if $E \vdash^{p}_{p'} e \Downarrow v$ then $p-p' = q+\Phi_E(\Gamma^r)-(q'+\Phi(v:A))$. Suppose $p-p' \neq q+\Phi_E(\Gamma^r)-(q'+\Phi(v:A))$, there exists always some $r_1, r_2 \in \mathbb{Q}_0^+$ such that $p+r_1 = q+\Phi_E(\Gamma^r)+r_2$. Since $E \vdash^{p}_{p'} e \Downarrow v$, we have $E \vdash^{p+r_1}_{p'+r_1} e \Downarrow v$. By Theorem 1, $p'+r_1 = q'+\Phi(v : A)+r_2$, thus the assumption is contradictory.

Consider any E_1 and E_2 such that $E_1 \approx_X E_2$, hence $E_1 \vdash e \Downarrow v_1$ and $E_2 \vdash e \Downarrow v_2$. For all $p_1, p_1' \in \mathbb{Q}_0^+$ such that $E_1 \vdash^{p_1}_{p_1'} e \Downarrow v_1$, we have $p_1 - p_1' = q+\Phi_{E_1}(\Gamma^r)-(q'+\Phi(v_1:A))$. Similarly, for all $p_2, p_2' \in \mathbb{Q}_0^+$ such that $E_2 \vdash^{p_2}_{p_2'} e \Downarrow v_2$, $p_2 - p_2' = q+\Phi_{E_2}(\Gamma^r)-(q'+\Phi(v_2:A))$. Since $\Phi_{E_1}(X) = \Phi_{E_2}(X)$ by Lemma 2, $\forall x \in \text{dom}(\Gamma^r) \setminus X. \Phi(E_i(x) : \Gamma^r(x)) = 0$, and $\Phi(v_i : A) = 0$, $i = 1,2$. Thus $p_1 - p_1' = p_2 - p_2'$. □

3.4 Type system for upper bounds

If we treat potential as an *affine* resource then we arrive that the original amortized analysis for upper bounds [HJ03]. To this end, we allow unrestricted weakening and a relax rule in which we can waste potential.

$$\text{(U:RELAX)}\ \frac{\Sigma^r;\Gamma^r \vdash^{p}_{p'} e:A \qquad q \geq p \qquad q-p \geq q'-p'}{\Sigma^r;\Gamma^r \vdash^{q}_{q'} e:A} \qquad \text{(U:WEAKENING)}\ \frac{\Sigma^r;\Gamma^r \vdash^{q}_{q'} e:B}{\Sigma^r;\Gamma^r, x:A \vdash^{q}_{q'} e:B}$$

$$\text{(U:SUBTYPE)}\ \frac{\Sigma^r;\Gamma^r \vdash^{q}_{q'} e:A \qquad B <: A}{\Sigma^r;\Gamma^r \vdash^{q}_{q'} e:B} \qquad \text{(U:SUPERTYPE)}\ \frac{\Sigma^r;\Gamma^r, x:B \vdash^{q}_{q'} e:C \qquad B <: A}{\Sigma^r;\Gamma^r, x:A \vdash^{q}_{q'} e:C}$$

Additionally, we can use subtyping to waste linear potential [HJ03]. (See the converse definition for subtyping for lower bounds below.) Similarly to Theorem 1, we can prove the following theorem.

Theorem 3. *If* $\models E:\Gamma^r$, $E \vdash e \Downarrow v$*, and* $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : A$*, then for all* $p, r \in \mathbb{Q}_0^+$ *such that* $p \geq q+\Phi_E(\Gamma^r)+r$*, there exists* $p' \in \mathbb{Q}_0^+$ *satisfying* $E \vdash^{p}_{p'} e \Downarrow v$ *and* $p' \geq q'+\Phi(v:A)+r$.

3.5 Type system for lower bounds

The type judgements for lower bounds have the same form and data types as the type judgements for constant resource usage and upper bounds. However, the intended meaning of the judgment $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e:A$ is the following. Under given environment E, less than $q+\Phi_E(\Gamma)$ resource units are not sufficient to evaluate e to a value v so that more than $q' + \Phi(v : A)$ resource units are left over.

The syntax-directed typing rules are the same as the rules in constant-resource type system as given in Fig. 7. In addition, we have the structural rules in Fig. 8. The rule L:RELAX is dual to U:RELAX. In L:RELAX, potential is treated as a *relevant* resource: We are not allowed to waste potential but we can create potential out of the blue if we ensure that we either use it or pass it to the result. The same idea is formalized for the linear potential with the subtyping rules L:SUBTYPE and L:SUPERTYPE. The subtyping relation is defined as follows.

$$\frac{A \in \{\text{unit}, \text{bool}, \text{int}\}}{A <: A} \qquad \frac{A_1 <: A_2 \quad p_1 \le p_2}{L^{p_1}(A_1) <: L^{p_2}(A_2)} \qquad \frac{A_1 <: A_2 \quad B_1 <: B_2}{A_1 * A_2 <: B_1 * B_2}$$

It holds that if $A <: B$ then $\widehat{A} = \widehat{B}$ and $\Phi(v : A) \le \Phi(v : B)$. Suppose that it is not sufficient to evaluate e with p available resource units to get p' resource units left over. L:SUBTYPE reflects the fact that we also cannot evaluate e with p resources get more than p' resource units after the evaluation. L:SUPERTYPE says that we also cannot evaluate e with less than p and get p' resource units afterwards. The rule A:RELAX reflects the fact that it is impossible to evaluate e with $p + c$ and gets more than $p' + c$ where $c \in \mathbb{Q}_0^+$.

Example Consider again the functions *filter_succ* and *fs_twice* given in Fig. 6 in which the resource consumption is defined using *tick* annotations. The best-case resource usage of *filter_succ(ℓ)* is $3|\ell| + 1$ and best-case resource usage of *fs_twice(ℓ)* is $6|\ell| + 2$. This can be reflected by the following function types for lower bounds.

$$\begin{aligned} \mathsf{filter\_succ} :&\quad L^3(\text{int}) \xrightarrow{1/0} L^0(\text{int}) \\ \mathsf{fs\_twice} :&\quad L^6(\text{int}) \xrightarrow{2/0} L^0(\text{int}) \end{aligned}$$

To derive the lower bound for *fs_twice*, we need the same compositional reasoning as for the derivation of the upper bound. For the inner call of *filter_succ* we use the type

$$\mathsf{filter\_succ} : \quad L^6(\text{int}) \xrightarrow{2/1} L^3(\text{int}) .$$

It can be understood as follows. If the input list carries 6 potential units per element then, for each element, we can either use all 6 (*if* case) or we can use 3 and assign 3 to the output (*else* case).

The type system for lower bounds is a *relevant* type system [Wal02]. That means every variable is used at least once by allowing *exchange* and *contraction properties*, but not *weakening*. However, we as in the constant-time type system we allow a restricted from of weakening if the potential annotations are zero using the rule L:WEAKENING. The following lemma states formally the contraction property which is derived in Fig. 9.

Lemma 4. *If* $\Sigma^r;\Gamma^r, x_1 : A, x_2 : A \,\Big|\frac{q}{q'}\; e : B$ *then* $\Sigma^r;\Gamma^r, x : A \,\Big|\frac{q}{q'}\; \mathit{share}(x,(x_1,x_2).e) : B$

Proof. The proof is done by applying the share and supertype rules. Given an annotated data type A, there exists always data types A_1 and A_2 such that $\curlyvee(A \mid A_1, A_2)$, $A_1 <: A$, and $A_2 <: A$. The typing derivation is given in Fig. 9. □

The following theorems establish the soundness of the analysis. Theorem 5 is proved by induction and Theorem 4 follows by contradiction.

(L:RELAX)
$$\frac{\Sigma;\Gamma \mid\!\frac{p}{p'}\; e : A \qquad q \geq p \qquad q-p \leq q'-p'}{\Sigma;\Gamma \mid\!\frac{q}{q'}\; e : A}$$

(L:SUBTYPE)
$$\frac{\Sigma;\Gamma \mid\!\frac{q}{q'}\; e : A \qquad A <: B}{\Sigma;\Gamma \mid\!\frac{q}{q'}\; e : B}$$

(L:WEAKENING)
$$\frac{\Sigma^r;\Gamma^r \mid\!\frac{q}{q'}\; e : B \qquad \curlyvee(A \mid A, A)}{\Sigma^r;\Gamma^r, x : A \mid\!\frac{q}{q'}\; e : B}$$

(L:SUPERTYPE)
$$\frac{\Sigma;\Gamma, x : B \mid\!\frac{q}{q'}\; e : C \qquad A <: B}{\Sigma;\Gamma, x : A \mid\!\frac{q}{q'}\; e : C}$$

Figure 8: Structural rules for lower bounds.

(L:CONTRACTION)
$$\cfrac{\cfrac{\cfrac{\Sigma;\Gamma, x_1 : A, x_2 : A \mid\!\frac{q}{q'}\; e : B \qquad A_2 <: A}{\Sigma;\Gamma, x_1 : A, x_2 : A_2 \mid\!\frac{q}{q'}\; e : B \qquad A_1 <: A}}{\Sigma;\Gamma, x_1 : A_1, x_2 : A_2 \mid\!\frac{q}{q'}\; e : B \qquad \curlyvee(A \mid A_1, A_2)}}{\Sigma;\Gamma, x : A \mid\!\frac{q}{q'}\; \mathrm{share}(x, (x_1, x_2).e) : B}$$

Figure 9: Derivation of the contraction rule for lower-bounds.

Theorem 4. *Let $\models E : \Gamma^r$, $E \vdash e \Downarrow v$, and $\Sigma^r;\Gamma^r \mid\!\frac{q}{q'}\; e : A$. Then for all $p, r \in \mathbb{Q}_0^+$ such that $p < q + \Phi_E(\Gamma^r) + r$, there exists no $p' \in \mathbb{Q}_0^+$ satisfying $E \mid\!\frac{p}{p'}\; e \Downarrow v$ and $p' \geq q' + \Phi(v : A) + r$.*

Proof. The proof is relied on Theorem 5. For all $p, r \in \mathbb{Q}_0^+$ such that $p < q + \Phi_E(\Gamma^r) + r$, assume that there exists some $p' \in \mathbb{Q}_0^+$ such that $E \mid\!\frac{p}{p'}\; e \Downarrow v$ and $p' \geq q' + \Phi(v : A) + r$. Thus we have $p - p' < q + \Phi_E(\Gamma^r) - (q' + \Phi(v : A))$.

On the other hand, it holds that $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : A)) \leq p - p'$. The assumption is contradictory. □

Theorem 5. *Let $\models E : \Gamma^r$, $E \vdash e \Downarrow v$, and $\Sigma^r;\Gamma^r \mid\!\frac{q}{q'}\; e : A$. Then for all $p, p' \in \mathbb{Q}_0^+$ such that $E \mid\!\frac{p}{p'}\; e \Downarrow v$ we have $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : A)) \leq p - p'$.*

Proof. The proof is done by induction on the length of the derivation of the evaluation judgment $E \mid\!\frac{p}{p'}\; e \Downarrow v$ and the typing judgment $\Sigma;\Gamma \mid\!\frac{q}{q'}\; e : A$ with lexical order, in which the derivation of the evaluation judgment takes priority over the typing derivation. We need to do induction on the length of both evaluation and typing derivations since on one hand, an induction of only typing derivation would fail for the case of function application, which increases the length of the typing derivation, while the length of the evaluation derivation never increases. On the other hand, if the rules L:WEAKENING and A:SHARE are final step in the derivation, then the length of typing derivation decreases, while the length of evaluation derivation is unchanged.

A:SHARE Assume that the typing derivation ends with an application of the rule A:SHARE, thus $\Sigma^r;\Gamma^r, x_1 : A_1, x_2 : A_2 \mid\!\frac{q}{q'}\; e : B$ and $\curlyvee(A|A_1, A_2)$. Let $E_1 = E \setminus \{x\} \cup \{[x_1 \mapsto E(x), x_2 \mapsto E(x)]\}$. Since $\models E : \Gamma^r, x : A$ and following the property of the share relation we have $\models E_1 : \Gamma^r, x_1 : A_1, x_2 : A_2$.

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \vdash_{p'}^{p} \mathsf{share}(x,(x_1,x_2).e \Downarrow v$, by the rule E:SHARE we have $E_1 \vdash_{p'}^{p} e \Downarrow v$. Hence, by the induction hypothesis for e in the premise, it holds that $q + \Phi_{E_1}(\Gamma^r, x_1 : A_1, x_2 : A_2) - (q' + \Phi(v : B)) \leq p - p'$.

Because $\Phi(E(x) : A) = \Phi(E_1(x_1) : A_1) + \Phi(E_1(x_2) : A_2)$ and $\Phi_E(\Gamma^r) = \Phi_{E_1}(\Gamma^r) = \Phi_{E\setminus\{x\}}(\Gamma^r)$, we have $q + \Phi_E(\Gamma^r, x : A) - (q' + \Phi(v : B)) \leq p - p'$.

L:WEAKENING Suppose that the typing derivation $\Sigma^r; \Gamma^r, x : A \vdash_{q'}^{q} e : B$ ends with an application of the rule L:WEAKENING. Thus we have $\Sigma^r; \Gamma^r \vdash_{q'}^{q} e : B$, in which the data type A satisfies $\curlyvee(A \mid A, A)$. Since $\models E : \Gamma^r, x : A$, it follows that $\models E : \Gamma^r$.

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \vdash_{p'}^{p} e \Downarrow v$, by the induction hypothesis for e in the premise, it holds that $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : B)) \leq p - p'$. By the property of the share relation, $\Phi(a : A) = 0$, hence we have $q + \Phi_E(\Gamma^r, x : A) - (q' + \Phi(v : B)) \leq p - p'$.

L:RELAX Suppose that the typing derivation ends with an application of the rule L:RELAX, thus we have $\Sigma^r; \Gamma^r \vdash_{q_1'}^{q_1} e : A$, $q \geq q_1$, and $q - q_1 \leq q' - q_1'$.

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \vdash_{p'}^{p} e \Downarrow v$, we have $\models E : \Gamma^r$, hence by the induction hypothesis for e in the premise, it holds that $q_1 + \Phi_E(\Gamma^r) - (q_1' + \Phi(v : A)) \leq p - p'$. We have $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : A)) = q_1 + \Phi_E(\Gamma^r) - (q_1' + \Phi(v : A)) + ((q - q_1) - (q' - q_1'))$. Since $q - q_1 \leq q' - q_1'$, it holds that $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : A)) \leq q_1 + \Phi_E(\Gamma^r) - (q_1' + \Phi(v : A)) \leq p - p'$.

A:VAR Assume that e is a variable x. If $\Sigma^r; x : A \vdash_{0}^{K^{\mathsf{var}}} x : A$. Thus for all $p, p' \in \mathbb{Q}_0^+$ such that $E \vdash_{p'}^{p} e \Downarrow v$, we have $p = p' + K^{\mathsf{var}}$, hence $K^{\mathsf{var}} + \Phi(E(x) : A) - \Phi(v : A) \leq p - p' = K^{\mathsf{var}}$.

A:UNIT It is similar to the case A:VAR.

A:BOOL It is similar to the case A:VAR.

A:INT It is similar to the case A:VAR.

A:B-OP Assume that e is an expression of the form $\mathsf{op}_\diamond(x_1, x_2)$, where $\diamond = \{\mathsf{and}, \mathsf{or}\}$. Thus $\Sigma^r; x_1 : \mathrm{bool}, x_2 : \mathrm{bool} \vdash_{0}^{K^{\mathsf{op}}} e : \mathrm{bool}$ and $\models E : \{x_1 : \mathrm{bool}, x_2 : \mathrm{bool}\}$.

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \vdash_{p'}^{p} e \Downarrow v$, we have $K^{\mathsf{op}} + \Phi_E(x_1 : \mathrm{bool}, x_2 : \mathrm{bool}) - \Phi(v : \mathrm{bool}) = K^{\mathsf{op}} \leq p - p' = K^{\mathsf{op}}$.

A:IB-OP It is similar to the case A:B-OP.

A:I-OP It is similar to the case A:B-OP.

A:CONS If e is of the form $\mathrm{cons}(x_1, x_2)$, then the typing derivation ends with an application of the rule A:CONS and the evaluation derivation ends with the application of the rule E:CONS. Thus $\Sigma^r; x_1 : A, x_2 : L^{p_1}(A) \vdash_{0}^{p_1 + K^{\mathsf{cons}}} e : L^{p_1}(A)$ and $\models E : \{x_1 : A, x_2 : L^{p_1}(A)\}$.

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \vdash_{p'}^{p} e \Downarrow v$, we have $p - p' = K^{\mathsf{cons}}$, $E(x_1) = v_1$ and $E(x_2) = [v_2, \cdots, v_n]$. Let $\Gamma^r = x_h : A, x_t : L^{p_1}(A)$, it holds that $p_1 + K^{\mathsf{cons}} + \Phi_E(\Gamma^r) - (\Phi([v_1, ..., v_n] : L^{p_1}(A))) = K^{\mathsf{cons}} \leq p - p'$.

A:PAIR It is similar to the case A:CONS.

A:NIL It is similar to the case A:CONS.

A:MATCH-P Suppose that the typing derivation $\Sigma^r;\Gamma^r, x : A_1 * A_2 \vdash^{q}_{q'} \mathrm{match}(x,(x_1,x_2).e) : A$ ends with an application of the rule A:MATCH-P. Thus $\Sigma^r;\Gamma^r, x_1 : A_1, x_2 : A_2 \vdash^{q-K^{\mathsf{matchP}}}_{q'} e : A$ and $\models E : \Gamma, x : A_1 * A_2$.

Let $E_1 = E[x_1 \mapsto v_1, x_2 \mapsto v_2]$ and $\Gamma^r_1 = \Gamma^r, x_1 : A_1, x_2 : A_2$, since $\models v_1 : A_1$, $\models v_2 : A_2$, and $\models E : \Gamma^r$ it holds that $\models E_1 : \Gamma^r_1$.

For all $p, p' \in \mathbb{Q}^+_0$ such that $E \vdash^{p}_{p'} e \Downarrow v$, by the rule E:MATCH-P we have $E_1 \vdash^{p-K^{\mathsf{matchP}}}_{p'} e \Downarrow v$. Hence, by the induction hypothesis for e in the premise, it holds that $q - K^{\mathsf{matchP}} + \Phi_{E_1}(\Gamma^r_1) - (q' + \Phi(v : A)) \leq p - K^{\mathsf{matchP}} - p'$.

Since $\Phi_E(\Gamma^r, x : A_1 * A_2) = \Phi_{E_1}(\Gamma^r_1)$, it follows that $q + \Phi_E(\Gamma^r, x : A_1 * A_2) - (q' + \Phi(v : A)) \leq p - p'$.

A:FUN Assume that e is a function application of the form $\mathrm{app}(f, x)$. Thus $\Sigma^r; x : A_1 \vdash^{q+K^{\mathsf{app}}}_{q'} e : A_2$ and $\Sigma^r(f) = A_1 \xrightarrow{q/q'} A_2$. Because the considering program is well-formed, there exists a well-typed expression e_f under the typing context $\Gamma^r_1 = y^{\hat{f}} : A_1$ and the signature Σ^r, or $\Sigma^r;\Gamma^r_1 \vdash^{q}_{q'} e_f : A_2$.

Let $\Gamma^r = x : A_1$, $E(x) = v_1$ and $E_1 = [y^{\hat{f}} \mapsto v_1]$, since $\models E : \Gamma^r$, it follows that $\models E_1 : \Gamma^r_1$. For all $p, p' \in \mathbb{Q}^+_0$ such that $E \vdash^{p}_{p'} e \Downarrow v$, we have $E_1 \vdash^{p-K^{\mathsf{app}}}_{p'} e_{\hat{f}} \Downarrow v$. Hence, by the induction hypothesis for e_f, it holds that $q + \Phi_{E_1}(\Gamma^r_1) - (q' + \Phi(v : A_2)) \leq p - K^{\mathsf{app}} - p'$.

Since $\Phi_{E_1}(\Gamma^r_1) = \Phi(E_1(y^{\hat{f}}) : A_1) = \Phi_E(\Gamma^r) = \Phi(E(x) : A_1)$, it follows that $q + K^{\mathsf{app}} + \Phi(E(x) : A_1) - (q' + \Phi(v : A_2)) \leq p - p'$.

A:IF Suppose that e is an expression of the form $\mathrm{if}(x, e_t, e_f)$. Then one of the rules E:IF-TRUE and E:IF-FALSE has been applied in the evaluation derivation depending on the value of x.

Assume that the variable x is assigned the value $\mathtt{true}$ in E, or $E(x) = \mathtt{true}$. The typing rule for e has been derived by an application of the rule A:IF using the premise on the left thus $\Sigma^r;\Gamma^r \vdash^{q-K^{\mathsf{cond}}}_{q'} e_t : A$. Let $\Gamma^r_1 = \Gamma^r, x : \mathrm{bool}$, since $\models E : \Gamma^r_1$, it follows that $\models E : \Gamma^r$.

For all $p, p' \in \mathbb{Q}^+_0$ such that $E \vdash^{p}_{p'} e \Downarrow v$, by the rule E:IF-TRUE we have $E \vdash^{p-K^{\mathsf{cond}}}_{p'} e_t \Downarrow v$. Hence, by the induction hypothesis for e_t, it holds that $q - K^{\mathsf{cond}} + \Phi_E(\Gamma^r) - (q' + \Phi(v : A)) \leq p - K^{\mathsf{cond}} - p'$.

Because $\Phi_E(\Gamma^r) = \Phi_E(\Gamma^r_1)$, it follows $q + \Phi_E(\Gamma^r_1) - (q' + \Phi(v : A)) \leq p - p'$. If $E(x) = \mathsf{false}$ then the proof is similar.

A:MATCH-L It is the same as the case of a conditional expression. The evaluation derivation applies one of the rules E:MATCH-N and E:MATCH-L depending on the value of x.

Assume that x is assigned the value $[v_1, ..., v_n]$ under E, or $E(x) = [v_1, ..., v_n]$. Then, the evaluation derivation ends with an application of the rule E:MATCH-L. Let $E_1 = E[x_h \mapsto v_1, x_t \mapsto [v_2, ..., v_n]]$ and $\Gamma^r_1 = \Gamma^r, x_h : A, x_t : L^{p_1}(A)$, the typing derivation ends with an application of the rule A:MATCH-L, thus $\Sigma^r;\Gamma^r_1 \vdash^{q+p_1-K^{\mathsf{matchL}}}_{q'} e_2 : A_1$.

Since $\models [v_1,...,v_n] : L^{p_1}(A)$, we have $\models v_i : A, \forall i = 1,...,n$. Hence, it holds that $\models v_1 : A$ and $\models [v_2,...,v_n] : L^{p_1}(A)$. Finally, we have $\models E_1 : \Gamma_1^r$ (since $\models E : \Gamma^r$ implies $\models E_1 : \Gamma^r$).

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \mid\frac{p}{p'}\ e \Downarrow v$, by the rule E:MATCH-L we have $E_1 \mid\frac{p - K^{\mathsf{matchL}}}{p'}\ e_2 \Downarrow v$. By the induction hypothesis for e_2, it holds that $q + p_1 - K^{\mathsf{matchL}} + \Phi_{E_1}(\Gamma_1^r) - (q' + \Phi(v : A_1)) \leq p - K^{\mathsf{matchL}} - p'$.

Because $\Phi_E(\Gamma^r, x : L^p(A)) = \Phi_E(\Gamma^r) + n.p_1 + \Sigma_{i=1}^n \Phi(v_i : A)$, $\Phi_{E_1}(\Gamma_1^r) = \Phi_{E_1}(\Gamma^r) + (n-1).p_1 + \Sigma_{i=1}^n \Phi(v_i : A)$ and $\Phi_{E_1}(\Gamma^r) = \Phi_E(\Gamma^r)$, thus we have $\Phi_{E_1}(\Gamma_1^r) = \Phi_E(\Gamma^r, x : L^{p_1}(A)) - p_1$. Therefore, $q + \Phi_E(\Gamma^r, x : L^{p_1}(A)) - (q' + \Phi(v : A_1)) \leq p - p'$. If $E(x) = \mathsf{nil}$ then it is similar to the case A:MATCH-P.

A:LET Assume that e is an expression of the form $\mathrm{let}(x, e_1, x.e_2)$. Hence, the evaluation derivation ends with an application of the rule E:LET. Let $E_1 = E[x \mapsto v_1]$ and $\Gamma^r = \Gamma_1^r, \Gamma_2^r$. The typing derivation ends with an application of the rule A:LET, thus $\Sigma^r; \Gamma_1^r \mid\frac{q - K^{\mathsf{let}}}{q_1'}\ e_1 : A_1$ and $\Sigma^r; \Gamma_2^r, x : A_1 \mid\frac{q_1'}{q'}\ e_2 : A_2$.

For all $p, p' \in \mathbb{Q}_0^+$ such that $E \mid\frac{p}{p'}\ e \Downarrow v$, by the rule E:LET we have $E \mid\frac{p - K^{\mathsf{let}}}{p_1'}\ e_1 \Downarrow v_1$ and $E_1 \mid\frac{p_1'}{p'}\ e_2 \Downarrow v$. Since $\models E : \Gamma^r$, we have $\models E : \Gamma_1^r$. By the induction hypothesis for e_1, it holds that $q - K^{\mathsf{let}} + \Phi_E(\Gamma_1^r) - (q_1' + \Phi(v_1 : A_1)) \leq p - K^{\mathsf{let}} - p_1'$.

We have $\models E : \Gamma_2^r$, thus $\models E_1 : \Gamma_2^r, x : A_1$. Again by the induction hypothesis for e_2, we derive that $q_1' + \Phi_{E_1}(\Gamma_2^r, x : A_1) - (q' + \Phi(v : A_2)) \leq p_1' - p'$.

Sum two in-equations above, it follows $q + \Phi_E(\Gamma_1^r) - \Phi(v_1 : A_1) + \Phi_{E_1}(\Gamma_2^r, x : A_1) - (q' + \Phi(v : A_2)) = q + \Phi_E(\Gamma_1^r, \Gamma_2^r) - (q' + \Phi(v : A_2)) \leq p - p'$.

L:SUBTYPE Assume that the typing derivation ends with an application of the rule L:SUBTYPE, thus $\Sigma^r; \Gamma^r \mid\frac{q}{q'}\ e : A$ and $A <: B$.

By the induction hypothesis for e in the premise, for all $p, p' \in \mathbb{Q}_0^+$ such that $E \mid\frac{p}{p'}\ e \Downarrow v$ it holds that $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : A)) \leq p - p'$.

Because $\Phi(E(x) : A) \leq \Phi(E(x) : B)$ we have $q + \Phi_E(\Gamma^r) - (q' + \Phi(v : B)) \leq p - p'$.

L:SUPERTYPE Assume that the typing derivation ends with an application of the rule L:SUPERTYPE, thus $\Sigma^r; \Gamma^r, x : B \mid\frac{q}{q'}\ e : C$ and $A <: B$. Since $\models E : \Gamma^r, x : A$ and following the property of the subtyping relation we have $\models E : \Gamma^r, x : B$.

By the induction hypothesis for e in the premise, for all $p, p' \in \mathbb{Q}_0^+$ such that $E \mid\frac{p}{p'}\ e \Downarrow v$ it holds that $q + \Phi_E(\Gamma^r, x : B) - (q' + \Phi(v : C)) \leq p - p'$.

Because $\Phi(E(x) : A) \leq \Phi(E(x) : B)$ we have $q + \Phi_E(\Gamma^r, x : A) - (q' + \Phi(v : C)) \leq p - p'$. □

We can see that the relax rules are consistent among these type systems in sense of satisfying the following.

$$(q \geq p \wedge q - p \leq q' - p') \wedge (q \geq p \wedge q - p \geq q' - p') \Leftrightarrow (q \geq p \wedge q - p = q' - p')$$

That means the constraints for upper bounds and lower bounds imply the constraints for constant resource and vice versa.

3.6 Mechanization

We mechanized the soundness proofs for both the two new type systems as well as the classic AARA type system using the proof assistant Agda. The development is roughly 4000 lines of code, which includes the rules of the three types systems, the operational cost semantics, a proof of type preservation for each type system, and the soundness theorems for each type system.

One notable difference is our implementation of the typing contexts. In Agda our contexts are implemented as lists of pairs of variables and their types. Moreover, in our typing rules whenever a variable is added to the context we require a proof that the variable is fresh with respect to the existing context. This requirement is important as it allows us to preserve the invariant that the context is well formed with respect to the environment as we induct over typing and evaluation judgements in our soundness proofs. Furthermore, as our typing contexts are ordered lists we added an *exchange* rule to our typing rules.

Another important detail is in the implementation of potential. Potential $\Phi(v : A)$ for a value only is defined for well formed inputs. Inputs such as $\Phi(\mathsf{nil} : \text{bool})$ are not defined. Agda is total language and as such prohibits users from implementing partial functions. Thus we require in our Agda implementation that when calculating the potential of a value of a given type the user provide a derivation that the value is well formed with respect to that type. Similarly when calculating the potential of a context, $\Phi_E(\Gamma^r)$, with respect to an environment we require that the user provide a derivation that the context is well formed with respect to that environment.

Lastly, whereas the type systems and proofs presented here used positive rational numbers, in the Agda implementation we use natural numbers. This deviation was simply due to the lacking support for rationals in the Agda standard library. By replacing a number of trivial lemmas, mostly related to associativity and commutativity, the proofs and embeddings could be transformed to use rational numbers instead.

A large bulk of the work for the soundness proofs was algebraic manipulation of simply equalities and inequalities. While Agda has some nice support for equality reasoning and for automatically generating proofs for simple goals, these tools fall short and made certain parts of the verification process tedious. Agda, unlike Coq does not have a separate tacit language and a result a number of these proofs which are obvious algebraic manipulations to humans are long and tedious in Agda.

4 A resource-aware security type system

In this section we introduce a new type system that enforces *resource-aware noninterference* to prevent the leakage of information in *high-security* variables through *low-security* channels. In addition to preventing leakage over the usual input/output information flow channels, our system incorporates the constant-resource type system discussed in Section 3 to ensure that leakage does not occur over resource side channels.

The notion of security addressed by our type system considers an attacker who wishes to learn information about certain inputs containing sensitive data by making observations of the program's public outputs and resource usage. We assume an attacker who is able to control the value of any variable she is capable of observing, and thus influence the program's behavior and resource consumption. However, in our model the attacker can only observe the program's total resource usage upon termination, and so cannot distinguish between intermediate states or between terminating and non-terminating executions.

4.1 Security types

To distinguish parts of the program under the attacker's control from those that remain secret, we annotate types with labels ranging over a lattice $(\mathcal{L}, \sqsubseteq, \sqcup, \bot)$. The elements of $\mathcal{L}$ correspond to security levels partially-ordered by $\sqsubseteq$ with a unique bottom element $\bot$. The corresponding basic security types take the form:

$$k \in \mathcal{L} \qquad S ::= (\text{unit}, k) \mid (\text{bool}, k) \mid (\text{int}, k) \mid (L(S), k) \mid S * S$$

A security context Γ^s is a partial mapping from variable identifiers and the program counter pc to security types. The context assigns a type (unit, k) to pc to track information that may propagate through control flow as a result of branching statements. The security type for lists contains a label $L(S)$ for the elements, as well as a label k for the list's length.

As in other information flow type systems, the partial order $k \sqsubseteq k'$ indicates that the class k' is at least as restrictive as k, i.e., k is allowed to flow to k'. We assume a non-trivial security lattice that contains at least two labels: ℓ (low security) and h (high security), with $\ell \sqsubseteq h$. Following the convention defined in FlowCaml [Sim03], we also make use of a *guard relation* $k \lhd S$ which denotes that all of the labels appearing in S are at least as restrictive as k. This is given in Figure 10 along with its dual notion $S \blacktriangleleft k$, called the *collecting relation*, and the standard subtyping relation $S_1 \leq S_2$.

To refer to sets of variables by security class, we write $[\Gamma^s]_{\blacktriangleleft k}$ to denote the set of variable identifiers x in the domain of Γ^s such that $\Gamma^s(x) \blacktriangleleft k$, and define ${}_{k\lhd}[\Gamma^s]$ similarly. These gives us the set of variables upper- and lower-bounded by k, respectively. Conversely, we define $[\Gamma^s]_{\not\blacktriangleleft k} = \{x \in \text{dom}(\Gamma^s) : \Gamma^s(x) \not\blacktriangleleft k\}$, the set of variables more restrictive than k. To refer to the set of variables strictly bounded below by k_1 and above by k_2, we write ${}_{k_1\lhd}[\Gamma^s]_{\blacktriangleleft k_2}$. Given two well-formed environments E_1 and E_2, we say that they are *k-equivalent* with respect to Γ^s, written $E_1 \equiv_k E_2$, if they agree on all variables with label at most k:

$$E_1 \equiv_k E_2 \Leftrightarrow \forall x \in [\Gamma^s]_{\blacktriangleleft k}. E_1(x) = E_2(x)$$

This relation captures the attacker's *observational equivalence* between the two environments.

$$\frac{k \sqsubseteq k' \quad T \in \text{Atoms}}{k \lhd (T, k')} \qquad \frac{k \sqsubseteq k' \quad k \lhd S}{k \lhd (L(S), k')} \qquad \frac{k \lhd S_1 \quad k \lhd S_2}{k \lhd S_1 * S_2}$$

$$\frac{k' \sqsubseteq k \quad T \in \text{Atoms}}{(T, k') \blacktriangleleft k} \qquad \frac{k' \sqsubseteq k \quad S \blacktriangleleft k}{(L(S), k') \blacktriangleleft k} \qquad \frac{S_1 \blacktriangleleft k \quad S_2 \blacktriangleleft k}{S_1 * S_2 \blacktriangleleft k}$$

$$\frac{k \sqsubseteq k' \quad T \in \text{Atoms}}{(\text{T}, k) \leq (\text{T}, k')} \qquad \frac{k \sqsubseteq k' \quad S \leq S'}{(L(S), k) \leq (L(S'), k')} \qquad \frac{S_1 \leq S_1' \quad S_2 \leq S_2'}{S_1 * S_2 \leq S_1' * S_2'}$$

Figure 10: Guards, collecting security labels, and subtyping (Atoms = {unit, int, bool})

The first-order security types take the form:

$$\text{pc} \in \mathcal{L} \qquad F^s ::= S_1 \xrightarrow{\text{pc/const}} S_2 \mid S_1 \xrightarrow{\text{pc}} S_2$$

$$\frac{}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} () : (\text{unit}, \text{pc})}\ (\text{SR:Unit}) \qquad \frac{b \in \{\text{true}, \text{false}\}}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} b : (\text{bool}, \text{pc})}\ (\text{SR:Bool}) \qquad \frac{n \in \mathbb{Z}}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} n : (\text{int}, \text{pc})}\ (\text{SR:Int})$$

$$\frac{x : S \in \Gamma^s \qquad \text{pc} \lhd S}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} x : S}\ (\text{SR:Var})$$

$$\frac{x_1 : (\text{bool}, k_{x_1}) \in \Gamma^s \qquad x_2 : (\text{bool}, k_{x_2}) \in \Gamma^s \qquad \text{pc} \sqsubseteq k_{x_1} \sqcup k_{x_2} \qquad \diamond \in \{\text{and}, \text{or}\}}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} \text{op}_\diamond(x_1, x_2) : (\text{bool}, k_{x_1} \sqcup k_{x_2})}\ (\text{SR:B-Op})$$

$$\frac{\text{pc}; \Sigma^s; \Gamma^s \vdash e : S \qquad \text{const}(e)}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} e : S}\ (\text{SR:C-Gen}) \qquad \frac{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} e : S}{\text{pc}; \Sigma^s; \Gamma^s \vdash e : S}\ (\text{SR:Gen})$$

$$\frac{x_1 : (\text{int}, k_{x_1}) \in \Gamma^s \qquad x_2 : (\text{int}, k_{x_2}) \in \Gamma^s \qquad \text{pc} \sqsubseteq k_{x_1} \sqcup k_{x_2} \qquad \diamond \in \{=, <>, >, <, <=, >=\}}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} \text{op}_\diamond(x_1, x_2) : (\text{bool}, k_{x_1} \sqcup k_{x_2})}\ (\text{SR:IB-Op})$$

$$\frac{x_1 : S_1 \in \Gamma^s \qquad x_2 : S_2 \in \Gamma^s \qquad \text{pc} \lhd S_1 * S_2}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} \text{pair}(x_1, x_2) : S_1 * S_2}\ (\text{SR:Pair})$$

$$\frac{x_1 : (\text{int}, k_{x_1}) \in \Gamma^s \qquad x_2 : (\text{int}, k_{x_2}) \in \Gamma^s \qquad \text{pc} \sqsubseteq k_{x_1} \sqcup k_{x_2} \qquad \diamond \in \{+, -, *, \text{div}, \text{mod}\}}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} \text{op}_\diamond(x_1, x_2) : (\text{int}, k_{x_1} \sqcup k_{x_2})}\ (\text{SR:I-Op})$$

$$\frac{x_h : S \in \Gamma^s \qquad x_t : (L(S), k_x) \in \Gamma^s \qquad \text{pc} \lhd (L(S), k_x)}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} \text{cons}(x_h, x_t) : (L(S), k_x)}\ (\text{SR:Cons})$$

$$\frac{x : S_1 \in \Gamma^s \qquad \Sigma^s(f) = S_1 \xrightarrow{\text{pc}'} S_2 \qquad \text{pc} \sqsubseteq \text{pc}'}{\text{pc}; \Sigma^s; \Gamma^s \vdash \text{app}(f, x) : S_2}\ (\text{SR:Fun})$$

$$\frac{x : S_1 \in \Gamma^s \qquad \Sigma^s(f) = S_1 \xrightarrow{\text{pc}'} S_2 \qquad \text{pc} \sqsubseteq \text{pc}' \qquad S_1 \blacktriangleleft k_1}{\text{pc}; \Sigma^s; \Gamma^s \vdash^{\text{const}} \text{app}(f, x) : S_2}\ (\text{SR:L-Arg}) \qquad \frac{\text{pc}; \Sigma^s; \Gamma^s \vdash e : S \qquad S \le S'}{\text{pc}; \Sigma^s; \Gamma^s \vdash e : S'}\ (\text{SR:SubTyping})$$

Figure 11: Typing rules: security (1 of 2)

The annotation pc indicates the security level of the program counter, i.e., a lower-bound on the label of any observer who is allowed to learn that a given function has been invoked. The const annotation denotes that the function body respects *resource-aware noninterference* defined in

(SR:C-FUN)

$$\frac{x : S_1 \in \Gamma^s \qquad \Sigma^s(f) = S_1 \xrightarrow{\mathrm{pc}'/\mathrm{const}} S_2 \qquad \mathrm{pc} \sqsubseteq \mathrm{pc}'}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} \mathrm{app}(f, x) : S_2}$$

(SR:NIL)

$$\frac{S \in \mathcal{S} \qquad \mathrm{pc} \lhd S}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} \mathsf{nil} : (L(S), \mathrm{pc})}$$

(SR:C-SUBTYPING)

$$\frac{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} e : S \qquad S \leq S'}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} e : S'}$$

(SR:LET)

$$\frac{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash e_1 : S_1 \qquad \mathrm{pc}; \Sigma^s; \Gamma^s, x : S_1 \vdash e_2 : S_2}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash \mathrm{let}(x, e_1, x.e_2) : S_2}$$

(SR:L-LET)

$$\frac{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} e_1 : S_1 \qquad \mathrm{pc}; \Sigma^s; \Gamma^s, x : S_1 \vdash^{\mathrm{const}} e_2 : S_2 \qquad S_1 \blacktriangleleft k_1}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} \mathrm{let}(x, e_1, x.e_2) : S_2}$$

(SR:IF)

$$\frac{x : (\mathrm{bool}, k_x) \in \Gamma^s \qquad \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s \vdash e_t : S \qquad \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s \vdash e_f : S \qquad \mathrm{pc} \sqcup k_x \lhd S}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash \mathrm{if}(x, e_t, e_f) : S}$$

(SR:L-IF)

$$\frac{\begin{array}{c} x : (\mathrm{bool}, k_x) \in \Gamma^s \\ \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} e_t : S \qquad \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} e_f : S \qquad \mathrm{pc} \sqcup k_x \lhd S \qquad k_x \sqsubseteq k_1 \end{array}}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} \mathrm{if}(x, e_t, e_f) : S}$$

(SR:MATCH-P)

$$\frac{x : S_1 * S_2 \in \Gamma^s \qquad \mathrm{pc}; \Sigma^s; \Gamma^s, x_1 : S_1, x_2 : S_2 \vdash e : S}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash \mathrm{match}(x, (x_1, x_2).e) : S}$$

(SR:C-MATCH-P)

$$\frac{x : S_1 * S_2 \in \Gamma^s \qquad \mathrm{pc}; \Sigma^s; \Gamma^s, x_1 : S_1, x_2 : S_2 \vdash^{\mathrm{const}} e : S}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} \mathrm{match}(x, (x_1, x_2).e) : S}$$

(SR:MATCH-L)

$$\frac{\begin{array}{c} x : (L(S), k_x) \in \Gamma^s \\ \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s \vdash e_1 : S_1 \qquad \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s, x_h : S, x_t : (L(S), k_x) \vdash e_2 : S_1 \qquad \mathrm{pc} \sqcup k_x \lhd S_1 \end{array}}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash \mathrm{match}(x, e_1, (x_h, x_t).e_2) : S_1}$$

(SR:C-MATCH-L)

$$\frac{\begin{array}{c} x : (L(S), k_x) \in \Gamma^s \\ \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} e_1 : S_1 \qquad \mathrm{pc} \sqcup k_x; \Sigma^s; \Gamma^s, x_h : S, x_t : (L(S), k_x) \vdash^{\mathrm{const}} e_2 : S_1 \qquad \mathrm{pc} \sqcup k_x \lhd S_1 \end{array}}{\mathrm{pc}; \Sigma^s; \Gamma^s \vdash^{\mathrm{const}} \mathrm{match}(x, e_1, (x_h, x_t).e_2) : S_1}$$

Figure 12: Typing rules: security (2 of 2)

the following section. A *security signature* $\Sigma^s : \mathrm{FID} \to \wp(\mathcal{F}^s) \setminus \{\emptyset\}$ is a finite partial mapping from a set of function identifiers to a *non-empty sets* of first-order security types.

4.2 Resource-aware noninterference

We consider an adversary associated with label $k_1 \in \mathcal{L}$, who can observe and control variables in $[\Gamma^s]_{\blacktriangleleft k_1}$. Intuitively, we say that a program P satisfies resource-aware noninterference at level (k_1, k_2) with respect to Γ^s, where $k_1 \sqsubseteq k_2$, if *1)* the behavior of P does not leak any information about the contents of variables more sensitive than k_1, and *2)* does not leak any information about the contents *or sizes* of variables more sensitive than k_2. The definition follows.

Definition 2. *Let E_1 and E_2 be two well-formed environments and Γ^s be a security context sharing their domain. An expression e satisfies resource-aware noninterference at level (k_1, k_2) for $k_1 \sqsubseteq k_2$, if whenever E_1 and E_2 are:*

1. *observationally-equivalent at k_1: $E_1 \equiv_{k_1} E_2$,*
2. *size-equivalent with respect to ${}_{k_1 \triangleleft}[\Gamma^s]_{\blacktriangleleft k_2}$: $E_1 \approx_{{}_{k_1 \triangleleft}[\Gamma^s]_{\blacktriangleleft k_2}} E_2$*

then the following holds:

$$E_1 \,\Big|\!\frac{p_1}{p_1'}\; e \Downarrow v_1, E_2 \,\Big|\!\frac{p_2}{p_2'}\; e \Downarrow v_2 \implies v_1 = v_2 \wedge p_1 - p_1' = p_2 - p_2'$$

The final condition in Defintion 2 ensures two properties. First, requiring that $v_1 = v_2$ provides noninterference [GM82], given that E_1 and E_2 are observationally-equivalent. Second, the requirement $p_1 - p_1' = p_2 - p_2'$ ensures that the program's resource consumption will remain constant with respect to changes in variables from the set $[\Gamma^s]_{\not\blacktriangleleft k_1}$. This establishes noninterference with respect to the program's final resource consumption, and thus prevents the leakage of secret information through resource side-channels.

Before moving on, we point out an important subtlety in this definition. We require that all variables in ${}_{k_1 \triangleleft}[\Gamma^s]_{\blacktriangleleft k_2}$ begin with equivalent sizes in E_1 and E_2, but not those in ${}_{k_2 \triangleleft}[\Gamma^s]$. By fixing this quantity in the initial environments, we assume that an attacker is able to control and observe it, so it is not protected by the definition. This effectively establishes three classes of variables, i.e., those whose size and content are observable to the k_1-adversary, those whose size (but *not* content) is observable, and those whose size and content remain secret. In the remainder of the text, we will simplify the technical development by assuming that the third and most-restrictive class is empty, and that all of the secret variables reside in ${}_{k_1 \triangleleft}[\Gamma^s]_{\blacktriangleleft k_2}$.

Assumptions and limitations. The definition of resource-aware noninterference given in Definition 2 assumes an adversary whose observations of resource consumption match the cost semantics given in Section 3. Depending on how the costs are parameterized, this may not match the reality of actual resource use in a physical environment on modern hardware. For example, if the processor's instruction cache is not accounted for then this may introduce an exploitable discrepancy between the guarantees provided by the type system and the real-world attacker's observations [BB05, GBK11, OST06]. In this work, we use a cost semantics that is conceptually straightforward, and leave as future work the development of more precise models (such as the one described in by Zhang et al. [ZAM12]) that are faithful to the subtleties of hardware platforms.

4.3 Proving resource-aware noninterference

There are two extreme ways of proving resource-aware noninterference. Assume we already have established classic noninterference by using an information-flow type system. The first way is to additionally prove constant resource usage *globally* by forgetting the security labels and showing that the program has constant resource usage. This is a sound approach but it requires us to reason about parts of the programs that are not affected by secret data. It would therefore result the rejection of programs that have the resource-aware noninterference property but are not constant resource. The second way is to prove constant resource usage *locally* by ensuring that every conditional that branches on secret values is constant time. However, this local approach is problematic because it is not compositional. Consider the following examples in which *rev* is the standard reverse function.

```
let f1(b,x) =
  let z = if b then x else [] in rev z

let f2(b,x,y) =
  let z = if b then let _ = rev y in x
          else let _ = rev x in y
  in rev z
```

If we assume a cost model in which we count the number of function calls then the cost of *rev(x)* is $|x|$. So *rev* is constant resource w.r.t its argument. Moreover, the expression *if b then x else []* is constant resource. However, *f1* is not constant resource. In contrast, the conditional in the function *f2* is not constant resource. However, *f2* is a constant resource function. The function *f2* can be automatically analyzed with the constant-resource type system from Section 3 while *f1* is correctly rejected.

The idea of our type system for resource-aware noninterference is to allow both global and local reasoning about resource consumption as well as arbitrary intermediate levels. *We ensure that every expression that is typed in a high security context is part of a constant resource expression.* In this way, we get the benefits of local reasoning without loosing compositionality.

4.4 Typing rules and soundness

We combine our type system for constant resource usage with a standard information flow type system which based on FlowCaml [PS02]. The interface between the two type systems is relatively light and the idea is applicable to other methods for proving constant resource use as well as other security type systems.

In the type judgement, an expression is typed under a type context Γ^s and a label pc. The pc label can be considered an upper bound on the security labels of all values that affect the control flow of the expression and a lower bound on the labels of the function's effects [PS02]. As mentioned earlier, we will simplify the technical development by assuming that the third and most-restrictive class is empty, and that all of the secret variables reside in ${}_{k_1\triangleleft}[\Gamma^s]_{\blacktriangleleft k_2}$, that is, the typing rules here guarantee that well-typed expressions provably satisfy the resource-aware noninterference property w.r.t. changes in variables from the set $[\Gamma^s]_{\not\blacktriangleleft k_1}$.

We define two type judgements of the following form, in which we write $\mathrm{const}(e)$ if there exists Γ^r and Σ^r such that $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : A, \curlyvee(A \mid A, A)$, and $\forall x \in [\Gamma^s]_{\not\blacktriangleleft k_1}. \curlyvee(\Gamma^r(x) \mid \Gamma^r(x), \Gamma^r(x))$.

$$\mathrm{pc};\Sigma^s;\Gamma^s \vdash^{\mathrm{const}} e : S \quad \text{and} \quad \mathrm{pc};\Sigma^s;\Gamma^s \vdash e : S\,.$$

The judgement with the const annotation states that under a security configuration given by Γ^s and the label pc, e has type S and it satisfies resource-ware noninterference w.r.t. changes in variables from the set $[\Gamma^s]_{\not\blacktriangleleft k_1}$. The second judgement indicates that e satisfies the noninterference property but does not make any guarantees about resource-based side channels. The typing rules are given in Fig. 11 and Fig. 12. We implicitly assume that the security types and the resource-annotated counterparts have the same base types.

Note that the standard information flow typing rules [HR98, PS02] can be obtained by removing the const annotation from all judgements. Consider for instance the rule SR:IF for conditional expressions. By executing the true or false branches, an adversary could gain information about the conditional value whose security label is k_x. Therefore the conditional expression must be type-checked under a security assumption at least as restrictive as pc and k_x. This is a standard requirement in any information flow type system. In the following we will focus on explaining how the rules restrict the observable resource usage instead of these classic noninterference aspects.

The most interesting rules are SR:C-GEN and the rules for and let expressions and conditionals, which block leakage over resource usage when branching on high security data. SR:C-GEN allows us to *globally* reason about constant resource usage for an arbitrary subexpression that has the noninterference property. For example, we can apply SR:IF, the standard rule for conditionals, first and then SR:C-GEN to prove that the expression is constant resource. Alternatively, we can use rules such as SR:L-IF and SR:L-LET to *locally* reason about resource use.

The rule SR:L-LET reflects the fact that if both e_1 and e_2 have the resource-aware noninterference property and the size of x only depends on low security data then $\text{let}(x, e_1, x.e_2)$ has the resource-aware noninterference property. The reasoning is similar for rule SR:L-IF where we require that the variable x does not depend on high security data.

Leaf expressions such as $\text{op}_\diamond(x_1, x_2)$ and $\text{cons}(x_h, x_t)$ have constant resource usage. Thus their judgments are always associated with the qualifier const as shown in the rule SR:B-OP. The rule SR:C-FUN states that if a function's body has the resource-aware noninterference property then the function application has the resource-aware noninterference property too. If the argument's label is low security data, bounded below by k_1, then the function application has the resource-aware noninterference property since the value of the argument is always the same under any k-equivalent environments. It is reflected by rule SR:L-ARG.

Example Recall functions *compare* and *p_compare* in Fig. 5. Suppose the content of the first list is secret and the length is public. Thus it has type $(L(\text{int}, h), \ell)$. While the second list controlled by adversaries is public, hence it has type $(L(\text{int}, \ell), \ell)$. Assume that the pc label is ℓ and $[\Gamma^s]_{\not\nearrow k_1} = [\Gamma^s]_{\not\nearrow \ell}$. The return value's label depends on the content of the first list elements whose label is h. Thus it must be assigned the label h to make the functions well-typed.

$$\begin{aligned} \text{compare}: &\quad ((L(\text{int}, h), \ell), (L(\text{int}, \ell), \ell)) \xrightarrow{\ell} (\text{bool}, h) \\ \text{p\_compare}: &\quad ((L(\text{int}, h), \ell), (L(\text{int}, \ell), \ell)) \xrightarrow{\ell/\text{const}} (\text{bool}, h) \end{aligned}$$

Here, both functions satisfy the noninterference property at security label ℓ. However, only *p_compare* is resource-aware noninterference function w.r.t $[\Gamma^s]_{\not\nearrow \ell}$, or the secret list.

Consider the following function *cond_rev* in which *rev* is the standard reverse function.

```
let cond_rev(l1,l2,b1,b2) =
  if b1 then let r =
    if b2 then rev l1; l2 else rev l2; l1
    in rev r; ()
  else ()
```

Assume that l_1, l_2, b_1 and b_2 have types $(L(\text{int}, h), \ell)$, $(L(\text{int}, h), \ell)$, (bool, ℓ), and (bool, h), respectively. Given the *rev* function is constant w.r.t the argument, the inner *if* is not resource-aware noninterference. However, the *let* expression is resource-aware noninterference w.r.t $[\Gamma^s]_{\blacktriangleleft \ell} = \{l_1, l_2, b_2\}$ by applying the rule SR:C-GEN. Finally, the outer *if* branching on low security data and its branches of are resource-aware noninterference, has resource-aware noninterference property w.r.t $\{l_1, l_2, b_2\}$ at level ℓ by the rule SR:L-IF. We obtain the following inferred type.

$$\text{cond\_rev} : \quad ((L(\text{int}, h), \ell), (L(\text{int}, h), \ell), (\text{bool}, \ell), (\text{bool}, h)) \xrightarrow{\ell/\text{const}} (\text{unit}, \ell)$$

We now prove the soundness of the type system w.r.t the definition of resource-aware noninterference. The soundness theorem states that if e is well-typed expression with the const annotation then it is resource-aware noninterference expression at level k_1.

The following two lemmas are needed in the soundness proof. The first lemma states that the type system satisfies the standard *simple security* property [VSI96] and the second shows that the type system prove classic noninterference.

Lemma 5. *Let* $pc; \Sigma^s; \Gamma^s \vdash e : S$ *or* $pc; \Sigma^s; \Gamma^s \vdash^{const} e : S$. *For all variables x in e, if $S \blacktriangleleft k_1$ then* $\Gamma^s(x) \blacktriangleleft k_1$.

Proof. By induction on the structure of the typing derivation.

SR:UNIT There is no variable thus it follows immediately.

SR:BOOL It is similar to the case SR:UNIT.

SR:INT It is similar to the case SR:UNIT.

SR:VAR Since $\Gamma^s(x) = S$, if $S \blacktriangleleft k_1$ then $\Gamma^s(x) \blacktriangleleft k_1$.

SR:B-OP If $(\text{bool}, k_{x_1} \sqcup k_{x_2}) \blacktriangleleft k_1$ then $\Gamma^s(x_1) = (\text{bool}, k_{x_1}) \blacktriangleleft k_1$ and $\Gamma^s(x_2) = (\text{bool}, k_{x_2}) \blacktriangleleft k_1$.

SR:IB-OP It is similar to the case SR:B-OP.

SR:I-OP It is similar to the case SR:B-OP.

SR:GEN - SR:C-GEN By induction for e in the premise, it follows.

SR:FUN Because e is well-formed program, there exists a well-typed expression e_f such that pc'; $\Sigma^s; \Gamma^s \vdash e_f : S_2$. By induction for e_f, for all variables x in e, if $S \blacktriangleleft k_1$ then $\Gamma^s(x) \blacktriangleleft k_1$. It is similar for SR:L-ARG and SR:C-FUN.

S:LET If $S_2 \blacktriangleleft k_1$ then by induction for e_2, $S_1 \blacktriangleleft k_1$. Thus for all variable x in e, it is a variable in e_1 or e_2. By induction for e_1 and e_2, it follows. It is similar for SR:L-LET.

SR:IF If $S \blacktriangleleft k_1$ then by the hypothesis $(\text{bool}, k_x) \blacktriangleleft k_1$. For all variable y in e, it is a variable in e_t or e_f. By induction for e_t and e_f, it follows. It is similar for SR:L-IF.

SR:PAIR If $S_1 * S_2 \blacktriangleleft k_1$ then $\Gamma^s(x_1) = S_1 \blacktriangleleft k_1$ and $\Gamma^s(x_2) = S_2 \blacktriangleleft k_1$.

SR:MATCH-P If $S \blacktriangleleft k_1$ then by induction for e, $\Gamma^s(x_1) \blacktriangleleft k_1$ and $\Gamma^s(x_2) \blacktriangleleft k_1$. Thus $\Gamma^s(x) \blacktriangleleft k_1$. For all other variables y in e, again by induction for e, if $S \blacktriangleleft k_1$ then $\Gamma^s(y) \blacktriangleleft k_1$. It is similar for SR:C-MATCH-P.

SR:NIL It is similar to the case SR:UNIT.

SR:CONS If $(L(S), k_x) \blacktriangleleft k_1$ then $\Gamma^s(x_h) = S \blacktriangleleft k_1$ and $\Gamma^s(x_t) = (L(S), k_x) \blacktriangleleft k_1$.

SR:MATCH-L If $S_1 \blacktriangleleft k_1$ then by induction for e_2, $\Gamma^s(x_h) = S \blacktriangleleft k_1$ and $\Gamma^s(x_t) = (L(S), k_x) \blacktriangleleft k_1$. Thus $\Gamma^s(x) \blacktriangleleft k_1$. For all other variables y in e, y is a variable in e_1 or e_2. Again by induction for e_1 and e_2, if $S_1 \blacktriangleleft k_1$ then $\Gamma^s(y) \blacktriangleleft k_1$.

SR:SUBTYPING By the subtyping relation, if $S' \blacktriangleleft k_1$ then $S \blacktriangleleft k_1$. Thus by induction for e in the premise, for all variables x in e, if $S \blacktriangleleft k_1$ then $\Gamma^s(x) \blacktriangleleft k_1$. It is similar for SR:C-SUBTYPING. □

Lemma 6. *Let* $pc;\Sigma^s;\Gamma^s \vdash e : S$ *or* $pc;\Sigma^s;\Gamma^s \vdash^{const} e : S$, $E_1 \vdash e \Downarrow v_1$, $E_2 \vdash e \Downarrow v_2$, *and* $E_1 \equiv_{k_1} E_2$. *Then* $v_1 = v_2$ *if* $S \blacktriangleleft k_1$.

Proof. The proof is done by induction on the structure of the evaluation derivation and the typing derivation.

SR:UNIT Suppose the evaluation derivation of e ends with an application of the rule E:UNIT, thus $E_1 \vdash e \Downarrow ()$ and $E_2 \vdash e \Downarrow ()$. Hence, it follows.

SR:BOOL It is similar to the case SR:UNIT.

SR:INT It is similar to the case SR:UNIT.

SR:VAR Suppose the evaluation derivation ends with an application of the rule E:VAR, thus $E_1(x) = v_1$ and $E_2(x) = v_2$. The typing derivation ends with an application of the rule SR:VAR, thus $\Gamma^s(x) = S$. If $S \blacktriangleleft k_1$, by the hypothesis $E_1(x) = E_2(x)$ since $x \in \text{dom}(E_i), i = \{1, 2\}$.

SR:B-OP Suppose the evaluation derivation ends with an application of the rule E:BIN, thus $E_1(x_1) \diamond E_1(x_2) = v_1$ and $E_2(x_1) \diamond E_2(x_2) = v_2$. The typing derivation ends with an application of the rules SR:B-OP or SR:GEN. We have $k_{x_1} \lhd S$ and $k_{x_2} \lhd S$. If $S \blacktriangleleft k_1$ then $k_{x_1} \sqsubseteq k_1$ and $k_{x_2} \sqsubseteq k_1$. By the hypothesis, we have $E_1(x_1) = E_2(x_1)$ and $E_1(x_2) = E_2(x_2)$, thus $v_1 = v_2$.

SR:IB-OP It is similar to the case SR:B-OP.

SR:I-OP It is similar to the case SR:B-OP.

SR:GEN-SR:C-GEN By induction for e in the premise, it follows that if $S \blacktriangleleft k_1$ then $v_1 = v_2$.

SR:FUN Suppose the evaluation derivation ends with an application of the rule E:FUN, thus $\Sigma(g) = T_1 \to T_2$ and $[y^g \mapsto E_i(x)] \vdash e_g \Downarrow v_i$ for $i = \{1,2\}$. The typing derivation ends with an application of the following rules.

- Case SR:FUN. Because e is well-formed program, there exists a well-typed expression e_f such that pc';Σ^s;$\Gamma^s \vdash e_f : S_2$ and $e_{\hat{f}} = e_g$. By induction for e_f, if $S_2 \blacktriangleleft k_1$ then $v_1 = v_2$.
- Case SR:L-ARG. It is similar to the case SR:FUN.
- Case SR:C-FUN. It is similar to the case SR:FUN.
- Case SR:GEN and SR:C-GEN. It follows.

S:LET Suppose the evaluation derivation ends with an application of the rule E:LET, thus $E_i \vdash e_1 \Downarrow v_1^i$ and $E_i[x \mapsto v_1^i] \vdash e_2 \Downarrow v_i$ for $i = \{1,2\}$. The typing derivation ends with an application of the following rules.

- Case SR:L-LET. If $S_2 \blacktriangleleft k_1$, by the simple security lemma, it holds that $S_1 \blacktriangleleft k_1$. By induction for e_1, we have $v_1^1 = v_1^2$, so $E_1[x \mapsto v_1^1] \equiv_k E_2[x \mapsto v_1^2]$. Again by induction for e_2, we have $v_1 = v_2$.
- Case SR:LET. It is similar to the case SR:L-LET.
- Case SR:GEN and SR:C-GEN. It follows.

SR:IF Suppose e is of the form if(x, e_t, e_f), the evaluation derivation ends with an application of the rule E:IF-TRUE or the rule E:IF-FALSE. The typing derivation ends with an application of the following rules.

- Case SR:L-IF. By the hypothesis we have $k_x \sqsubseteq k_1$, thus $E_1(x) = E_2(x)$. Assume that $E_1(x) = \mathsf{true}$, then $E_1 \vdash e_t \Downarrow v_1$ and $E_2 \vdash e_t \Downarrow v_2$. By induction for e_t we have $v_1 = v_2$ if $S \blacktriangleleft k_1$. It is similar for $E_1(x) = \mathsf{false}$.
- Case SR:IF. If $k_x \sqsubseteq k_1$ the proof is similar to the case SR:L-IF. Otherwise, $k_x \not\sqsubseteq k_1$, thus by the simple security lemma we have $S \not\blacktriangleleft k_1$.
- Case SR:GEN and SR:C-GEN. It follows.

SR:PAIR Suppose the evaluation derivation ends with an application of the rule E:PAIR, thus $(E_i(x_1), E_i(x_2)) = v_i$ for $i = \{1,2\}$. The typing derivation ends with an application of the rules SR:PAIR or SR:GEN.

If $S_1 * S_2 \blacktriangleleft k$, then by the simple security lemma we have $S_1 \blacktriangleleft k_1$ and $S_2 \blacktriangleleft k_1$. Hence it follows $v_1 = v_2$.

SR:MATCH-P Suppose the evaluation derivation ends with an application of the rule E:MATCH-P, thus $E_i(x) = (v_1^i, v_2^i)$ and $E_i[x_1 \mapsto v_1^i, x_2 \mapsto v_2^i] \vdash e \Downarrow v_i$ for $i = \{1,2\}$. The typing derivation ends with an application of the following rules.

- Case SR:MATCH-P. If $S \blacktriangleleft k_1$, then by the simple security lemma we have $S_1 * S_2 \blacktriangleleft k_1$. By the hypothesis, $E_1(x) = E_2(x)$, thus $v_1^1 = v_1^2$ and $v_2^1 = v_2^2$. Hence, $E_1[x_1 \mapsto v_1^1, x_2 \mapsto v_2^1] \equiv_k E_2[x_1 \mapsto v_1^2, x_2 \mapsto v_2^2]$, by induction for e in the premise, it holds that $v_1 = v_2$.

- Case SR:C-MATCH-P. It is similar to the case SR:MATCH-P.
- Case SR:GEN and SR:C-GEN. It follows.

SR:NIL It is similar to the case SR:UNIT.

SR:CONS Suppose the evaluation derivation ends with an application of the rule E:CONS, thus $E_i(x_h) = v_1^i$ and $E_i(x_t) = [v_2^i, \cdots, v_n^i]$ for $i = \{1,2\}$. The typing derivation ends with an application of the rules SR:CONS or SR:GEN. If $(L(S), k_x) \blacktriangleleft k_1$ then by the hypothesis we have $v_1^1 = v_1^2$ and $[v_2^1, \cdots, v_n^1] = [v_2^2, \cdots, v_n^2]$. Thus $E_1(\mathsf{cons}(x_h, x_t)) = E_2(\mathsf{cons}(x_h, x_t))$.

SR:MATCH-L Suppose e is of the form match$(x, e_1, (x_h, x_t).e_2)$, the evaluation derivation ends with an application of the rule E:MATCH-N or the rule E:MATCH-L. The typing derivation ends with an application of the following rules.

- Case SR:MATCH-L. If $S_1 \blacktriangleleft k_1$, then by the simple security lemma we have $(L(S), k_x) \blacktriangleleft k_1$. By the hypothesis we have $E_1(x) = E_2(x)$. Assume that $E_1(x) = E_2(x) = [v_1, \cdots, v_n]$, by the rule E:MATCH-L we have $E_i[x_h \mapsto v_1, x_t \mapsto [v_2, ..., v_n]] \vdash e_2 \Downarrow v_i$ for $i = \{1,2\}$. Since $E_1[x_h \mapsto v_1, x_t \mapsto [v_2, ..., v_n]] \equiv_k E_2[x_h \mapsto v_1, x_t \mapsto [v_2, ..., v_n]]$, by induction for e_2, it holds that $v_1 = v_2$ if $S_1 \blacktriangleleft k_1$. It is similar for $E_1(x) = E_2(x) = \mathsf{nil}$.
- Case SR:C-MATCH-L.It is similar to the case SR:MATCH-L.
- Case SR:GEN and SR:C-GEN. It follows.

SR:SUBTYPING Suppose the typing derivation ends with the rule SR:SUBTYPING. If $S' \blacktriangleleft k_1$ then $S \blacktriangleleft k_1$. Thus by induction for e in the premise it follows. It is similar for SR:C-SUBTYPING. □

Theorem 6. *If* $\models E : \Gamma^s$, $E \vdash e \Downarrow v$, *and* $pc; \Sigma^s; \Gamma^s \underset{}{\overset{const}{\vdash}} e : S$ *then e is resource-aware noninterference expression at level* k_1.

Proof. The proof is done by induction on the structure of the typing derivation and the evaluation derivation. Let X be the set of variables $[\Gamma^s]_{\blacktriangleleft k_1}$. For all environments E_1, E_2 such that $E_1 \approx_X E_2$ and $E_1 \equiv_{k_1} E_2$, if $E_1 \frac{p_1}{p_1'} e \Downarrow v_1$ and $E_2 \frac{p_2}{p_2'} e \Downarrow v_2$. We then show that $p_1 - p_1' = p_2 - p_2'$ and $v_1 = v_2$ if $S \blacktriangleleft k_1$. By Lemma 6, e satisfies the noninterference property at security label k_1. Thus we need to prove that $p_1 - p_1' = p_2 - p_2'$.

SR:UNIT Suppose the evaluation derivation of e ends with an application of the rule E:UNIT, thus $p_1 - p_1' = p_2 - p_2' = K^{\mathsf{unit}}$.

SR:BOOL It is similar to the case SR:UNIT.

SR:INT It is similar to the case SR:UNIT.

SR:VAR It is similar to the case SR:UNIT.

SR:B-OP Suppose the evaluation derivation ends with an application of the rule E:BIN, thus $E_1 \vdash^{p'_1+K^{\mathsf{op}}}_{p'_1} e \Downarrow v_1$ and $E_1 \vdash^{p'_2+K^{\mathsf{op}}}_{p'_2} e \Downarrow v_1$. We have $p_1 - p'_1 = p_2 - p'_2 = K^{\mathsf{op}}$.

SR:IB-OP It is similar to the case SR:B-OP.

SR:I-OP It is similar to the case SR:B-OP.

SR:C-GEN By the hypothesis we have const(e), thus it holds that $\Sigma^r;\Gamma^r \vdash^{q}_{q'} e : A$ and $\curlyvee(A \mid A, A)$. By the constant-resource theorem, for all $p_1, p'_1, p_2, p'_2 \in \mathbb{Q}^+_0$ such that $E_1 \vdash^{p_1}_{p'_1} e \Downarrow v_1$ and $E_2 \vdash^{p_2}_{p'_2} e \Downarrow v_2$, we have $p_1 - p'_1 = q + \Phi_{E_1}(\Gamma^r) - (q' + \Phi(v_1 : A))$ and $p_2 - p'_2 = q + \Phi_{E_2}(\Gamma^r) - (q' + \Phi(v_2 : A))$.

Since $E_1 \approx_X E_2$, $\Phi_{E_1}(X) = \Phi_{E_2}(X)$. For all $y \notin X$, $E_1(y) = E_2(y)$ since $E_1 \equiv_{k_1} E_2$, thus $\Phi(E_1(y)) = \Phi(E_2(y))$. Hence, $\Phi_{E_1}(\Gamma^r) = \Phi_{E_2}(\Gamma^r)$, it follows $p_1 - p'_1 = p_2 - p'_2$.

SR:FUN Suppose e is of the form app(f, x), thus the typing derivation ends with an application of either the rule SR:L-ARG, SR:C-FUN, or SR:C-GEN.

- Case SR:L-ARG. By the hypothesis we have $E_1(x) = E_2(x)$, it follows $p_1 - p'_1 = p_2 - p'_2$.
- Case SR:C-FUN. Because e is well-formed, there exists a well-typed expression e_f such that pc';$\Sigma^s;\Gamma^s \vdash^{\mathrm{const}} e_f : S_2$. By induction for e_f which is resource-aware noninterference w.r.t X, $p_1 - K^{\mathsf{app}} - p'_1 = p_2 - K^{\mathsf{app}} - p'_2$, it follows.
- Case SR:C-GEN. By the case SR:C-GEN it follows.

SR:LET Suppose e is of the form let$(x, e_1, x.e_2)$, thus the typing derivation ends with an application of either the rule SR:L-LET or SR:C-GEN.

- Case SR:L-LET. Suppose the evaluations $E_1 \vdash^{p_1 - K^{\mathsf{let}}}_{p'} e_1 \Downarrow v^1_1$, $E_2 \vdash^{p_2 - K^{\mathsf{let}}}_{p"} e_1 \Downarrow v^2_1$, $E_1[x \mapsto v^1_1] \vdash^{p'}_{p'_1} e_2 \Downarrow v_1$, and $E_2[x \mapsto v^2_1] \vdash^{p"}_{p'_2} e_2 \Downarrow v_2$. By induction for e_1 that is resource-aware noninterference w.r.t X, $p_1 - K^{\mathsf{let}} - p' = p_2 - K^{\mathsf{let}} - p"$. By the hypothesis $v^1_1 = v^2_1$. Thus $E_1[x \mapsto v^1_1] \approx_X E_2[x \mapsto v^2_1]$ and $E_1[x \mapsto v^1_1] \equiv_{k_1} E_2[x \mapsto v^2_1]$, by induction for e_2 that is resource-aware noninterference w.r.t X, we have $p' - p'_1 = p" - p'_2$. Hence, $p_1 - p'_1 = p_2 - p'_2$.
- Case SR:C-GEN. By the case SR:C-GEN it follows.

SR:IF Suppose e is of the form if(x, e_t, e_f), thus the typing derivation ends with an application of either the rule SR:L-IF or SR:C-GEN.

- Case SR:L-IF. By the hypothesis we have $E_1(x) = E_2(x)$. Assume that $E_1(x) = E_2(x) = \mathsf{true}$, by the evaluation rule E:IF-TRUE, $E_1 \vdash^{p_1 - K^{\mathsf{cond}}}_{p'_1} e_t \Downarrow v_1$ and $E_2 \vdash^{p_2 - K^{\mathsf{cond}}}_{p'_2} e_t \Downarrow v_2$. By induction for e_t that is resource-aware noninterference w.r.t X, we have $p_1 - p'_1 = p_2 - p'_2$. It is similar for $E_1(x) = E_2(x) = \mathsf{false}$.
- Case SR:C-GEN. Since $E_1 \approx_X E_2$ w.r.t Γ^s, we have $E_1 \approx_X E_2$ w.r.t Γ^r. By the hypothesis we have const(e). Thus by the soundness theorem of constant resource type system, it follows $p_1 - p'_1 = p_2 - p'_2$.

SR:PAIR It is similar to the case SR:B-OP.

SR:MATCH-P Suppose e is of the form match$(x, (x_1, x_2).e)$, thus the typing derivation ends with an application of either the rule SR:C-MATCH-P or SR:C-GEN.

- Case SR:C-MATCH-P. Let $E'_1 = E_1[x_1 \mapsto v^1_1, x_2 \mapsto v^1_2]$ and $E'_2 = E_2[x_1 \mapsto v^2_1, x_2 \mapsto v^2_2]$. If $x \in X$ then $|E_1(x)| \approx |E_2(x)|$. Thus $|E'_1(x_1)| \approx |E'_2(x_1)|$ and $|E'_1(x_2)| \approx |E'_2(x_2)|$. Hence, $E'_1 \approx_{X \cup \{x_1, x_2\}} E'_2$, by induction for e in the premise which is resource-aware noninterference w.r.t $X \cup \{x_1, x_2\}$, $p_1 - K^{\mathsf{matchP}} - p'_1 = p_2 - K^{\mathsf{matchP}} - p'_2$, it follows. If $x \notin X$ then $E_1(x) = E_2(x)$, it is similar.
- Case SR:C-GEN. By the case SR:C-GEN it follows.

SR:NIL It is similar to the case SR:UNIT.

SR:CONS It is similar to the case SR:B-OP.

SR:MATCH-L Suppose e is of the form match$(x, e_1, (x_h, x_t).e_2)$, thus the typing derivation ends with an application of either the rule SR:C-MATCH-L or SR:C-GEN.

- Case SR:C-MATCH-L. Let $E'_1 = E_1[x_h \mapsto v^1_1, x_t \mapsto v^1_2]$ and $E'_2 = E_2[x_h \mapsto v^2_1, x_t \mapsto v^2_2]$. If $x \in X$ then $|E_1(x)| \approx |E_2(x)|$. Suppose $E_1(x)$ and $E_2(x)$ are different from nil, $|E'_1(x_h)| \approx |E'_2(x_h)|$ and $|E'_1(x_t)| \approx |E'_2(x_t)|$. Hence, $E'_1 \approx_{X \cup \{x_t, x_h\}} E'_2$, by induction for e_2 which is resource-aware noninterference w.r.t $X \cup \{x_t, x_h\}$, we have $p_1 - K^{\mathsf{matchL}} - p'_1 = p_2 - K^{\mathsf{matchL}} - p'_2$, thus $p_1 - p'_1 = p_2 - p'_2$. If $E_1(x) = E_2(x) = \mathsf{nil}$ then by induction for e_1 that is resource-aware noninterference w.r.t X, it follows. If $x \notin X$ then $E_1(x) = E_2(x)$, it is similar.
- Case SR:C-GEN. By the case SR:C-GEN it follows.

SR:SUBTYPING The typing derivation ends with an application of either the rule SR:C-SHARE or SR:C-GEN.

- Case SR:C-SUBTYPING. By induction for e in the premise, $p_1 - p'_1 = p_2 - p'_2$.
- Case SR:C-GEN. By the case SR:C-GEN it follows.

□

5 Quantifying and transforming out leakages

We present techniques to quantify the amount of information leakage through resource usage and transform leaky programs into constant resource programs. The quantification relies on the lower and upper bounds inferred by our resource type systems. The transformation pads the programs with dummy computations so that the evaluations consume the same amount of resource usage and the outputs are identical with the original programs. In the current implementation, these dummy computations are added into programs by users and the padding parameters are automatically added by our analyzer to obtain the optimal values. It would be straightforward to make the process fully automatic but the interactive flavor of our approach helps to get a better understanding of the system.

5.1 Quantification

Recall from Section 4 that we assume an adversary at level k_1 who is always able to observe *1)* the values of variables in $[\Gamma^s]_{\blacktriangleleft k_1}$, and *2)* the final resource consumption of the program. For many programs, it may be the case that changes to the secret variables $[\Gamma^s]_{\not\blacktriangleleft k_1}$ effect observable differences in the program's final resource consumption, but only allow the attacker to learn partial information about the corresponding secrets. In this section, we show that the upper and lower-bound information provided by our type systems allow us to derive bounds on the amount of partial information that is leaked.

To quantify the amount of leaked information, we measure the number of distinct environments that the attacker could deduce as having produced a given resource consumption observation. However, becuase there may be an unbounded number of such environments, we parameterize this quantity on the size of the values contained in each environment. Let $\mathbf{E}^N$ denote the space of environments with values of size characterized by N. Given an environment E and expression e, define $U(E,e) = p_\delta$ such that $E \vdash^{p}_{p'} e \Downarrow v$ and $p_\delta = p - p'$. Then for an expression e and resource observation p_δ, we define the set $R_N(e, p_\delta)$ which captures the attacker's uncertainty about the environment which produced p_δ:

$$R_N(e, p_\delta) = \{E' \in \mathbf{E}^N : U(E,e) = p_\delta\}$$

Notice that when $|R_N(e,p)| = 1$, the attacker can deduce exactly which environment was used, whereas when this quantity is large little additional information is learned from p_δ. This gives us a natural definition of leakage, which is obtained by aggregating the inverse of the cardinality of R_N over the possible initial environments of e:

$$C_N(e) = \left(\sum_{E \in \mathbf{E}^N} \frac{1}{|R_N(e, U(E,e))|}\right) - 1$$

$C_N(e)$ corresponds to our intuition about leakage. When e leaks no information through resource consumption, then each term in the summation will be $1/|\mathbf{E}^{sizes}|$ giving $C_N(e) = 0$, whereas if e leaks perfect information about its starting environment then each term will be 1, leading to $C_N(e) = |\mathbf{E}^N| - 1$.

Theorem 7. *Let P^e_N be the complete set of resource observations producible by expression e under environments of size N, i.e.,*

$$P^e_N = \{p : \exists E \in \mathbf{E}^N . U(E,e) = p\}$$

Then $|P^e_N| = C_N(e) + 1$.

Proof. Observe that P^e_N partitions $\mathbf{E}^N$ into sets defined by $R_N(e,p)$ for each $p \in P^e_N$. This relationship allows us to swap the index of the summation from the definition of $C_N(e)$, obtaining an equivalent sum over P^e_N:

$$\begin{aligned} C_N(e) + 1 &= \sum_{E \in \mathbf{E}^N} \frac{1}{|R_N(e, U(E,e))|} = \sum_{\{p : \exists E \in \mathbf{E}^N . U(E,e) = p\}} 1 \\ &= |P^e_N| \end{aligned}$$

□

Lemma 7. *Let $l_e(N)$ and $u_e(N)$ be lower and upper-bounds on the resource consumption of e for inputs of size N. If $U(E,e) \in \mathbb{Z}$ for all environments E, then $C_N(e) \le u_e(N) - l_e(N)$.*

Lemma 8. *Assume that environments are sampled uniform-randomly from* $\mathbf{E}^N$*. Then the Shannon entropy of* P_N^e *is given by* $C_N(e)$*:* $H(P_N^e) \leq \log_2(C_N(e)+1)$

Lemma 7 leverages Theorem 7 to derive an upper-bound on leakage from upper and lower-bounds on resource usage. This result only holds when the possible resource observations of e are integral, as this ensures that the interval $[l_e(N), u_e(N)] \supseteq P_N^e$ is finite. Lemma 8 relates $C_N(e)$ to Shannon entropy, which is commonly used to characterize information leakage [ZAM12, KMO12, KB07].

5.2 Transformation

To transform programs into constant resource programs we extend the type system for constant resource use from Section 3. Recall that the type system treats potential in a linear fashion to ensure that potential is not wasted. We will now add *sinks* for potential which will be able to absorb excess potential. At runtime the the sinks will consume the exact amount of resources that have been statically-absorbed to ensure that potential is still treated in a linear way. The advantage of this approach is that the worst-case resource consumption is often not affected by the transformation. Additionally, we do not need to keep track of resource usage at runtime to pad the resource usage at the sinks, because the amount of resource that must be discarded is statically-determined by the type system. Finally, we automatically obtain a type derivation that serves as a proof that the transformation is constant-resource.

More precisely, the sinks are represented by the syntactic form: $\text{consume}_{(\text{A},p)}(x)$. Here, A is a resource-annotated type and $p \in \mathbb{Q}_{\geq 0}$ is a non-negative rational number. The idea is that A and p define the resource consumption of the expression. In the implementation, the user only has to write $\text{consume}(x)$, and the annotations are added via automatic syntax elaboration during the resource type inference.

Let E be a well-formed environment w.r.t Γ^r. For every $x \in \text{dom}(\Gamma)$ with $\Gamma^r(x) = A$, the expression $\text{consume}_{(\text{A},p)}(x)$ consumes $\Phi(E(x) : A) + p$ resource units and evaluate to $()$. The evaluation and typing rules for sinks are:

(T:CONSUME)
$$\frac{}{\Sigma; x : T_x \vdash \text{consume}_{(A,p)}(x) : \text{unit}}$$

(A:CONSUME)
$$\frac{}{\Sigma^r; x{:}A \mid\!\frac{p}{0}\; \text{consume}_{(A,p)}(x) : \text{unit}}$$

(E:CONSUME)
$$\frac{q = q' + \Phi(E(x) : A) + p}{E \mid\!\frac{q}{q'}\; \text{consume}_{(\text{A},p)}(x) \Downarrow ()}$$

(SR:CONSUME)
$$\frac{x : S_x \in \Gamma^s \qquad \text{pc} \sqsubseteq k \qquad S_x \blacktriangleleft k}{\text{pc}; \Sigma^s; \Gamma^s \vdash \text{consume}_{(A,p)}(x) : (\text{unit}, k)}$$

The extension of the proof of Theorem 1 to consume expressions is straightforward.

Adding consume expressions Let e_i be a subexpression of e and let e_i' be the expression $\text{let}(z, \text{consume}(x_1, \cdots, x_n), z.e_i)$ for some variables x_i. Let e' be the expression obtained from e by replacing e_i with e_i'. We write $e \hookrightarrow e'$ for such a transformation. Note that additional *share* and *let* expressions have to be added to convert e_i' into share-let normal form. The following lemma states that the transformed expression preserves the base type and evaluated value.

Lemma 9. *If* $\Sigma; \Gamma \vdash e : T$*,* $E \vdash e \Downarrow v$*, and* $e \hookrightarrow e'$ *then* $\Sigma; \Gamma \vdash e' : T$ *and* $E \vdash e' \Downarrow v$*.*

Proof. The proof is done by showing that if $\Sigma; \Gamma \vdash e_i : T_i$ and $E \vdash e_i \Downarrow v_i$ then $\Sigma; \Gamma \vdash e_i' : T_i$ and $E \vdash e_i' \Downarrow v_i$. It is easy to prove based on the form the sub-expression e_i. For instance, we show the case of conditional expression. Suppose e_i of the form $\text{if}(x, e_t, e_f)$, thus $e'_i :=$ $\text{let}(z, \text{consume}(x_1, \cdots, x_n), z.\text{if}(x, e_t, e_f))$. By the weakening rule, $\Sigma; \Gamma, z : \text{unit} \vdash \text{if}(x, e_t, e_f) : T_i$. Then by the typing rules for *let* and *share* expressions, $\Sigma; \Gamma \vdash e'_i : T_i$. Since $E \vdash e_i \Downarrow v_i$, thus $E[z \mapsto ()] \vdash e_i \Downarrow v_i$. By the evaluation rules for *let* and *share* expressions, $E \vdash e'_i \Downarrow v_i$. □

To transform an expression e into a constant resource expressions we perform multiple transformations $e \hookrightarrow e'$ which do not affect the type and semantics of e. This can be done automatically but in our implementation it works in an interactive fashion, meaning that users are responsible for the locations where consume expressions are put. The analyzer will infer the annotations A and constants p of the given consume expressions during type inference. If the inference is successful then we have $\text{const}(e')$ for the transformed program e'.

Example Recall the function *compare* form Fig. 5. To turn *compare* into a constant resource function. We insert *consume* expressions as shown below. Users can insert many *consume* expressions and the analyzer will determine which *consume* the are actually needed.

```
let rec c_compare h l = match h with
  | [] → (match l with | [] → Raml.tick 1.0; true
    | y::ys → Raml.tick 1.0; false)
  | x::xs → match l with
    | [] → Raml.tick 1.0; Raml.consume xs; false
    | y::ys → if (x = y) then
        Raml.tick 5.0; c_compare xs ys
      else Raml.tick 5.0; Raml.consume xs; false
```

We automatically obtain the following typing of the transformed function and the consume expressions:

$$\begin{array}{rll} \text{c\_compare}: & (L^5(\text{int}), L^0(\text{int})) \xrightarrow{1/0} \text{bool} & \\ \text{consume}: & L^5(\text{int}) \xrightarrow{5/0} \text{unit} & \text{(at line 5)} \\ \text{consume}: & L^5(\text{int}) \xrightarrow{1/0} \text{unit} & \text{(at line 8)} \end{array}$$

The worst-case resource consumption of the unmodified function *c_compare h l* is $1 + 5|h|$. Thus the consumption of the first *consume* must be $5 + 5(|h| - 1 - |\ell|)$ when h is longer than l. Otherwise, the consumption is zero. The second one consumes $1 + 5(|h_1| - 1)$, where h_1 is the sub-list of h from the first node which is different from the corresponding node in l.

6 Implementation and Evaluation

Type Inference The type inference for the type systems for constant resource and lower bounds are implemented in RAML [Ano15]. RAML is integrated in Inria's OCaml compiler and supports polynomial bounds, user-defined inductive types, higher-order functions, polymorphism, and arrays and references. All features are implemented for the new type systems and are basically orthogonal to the new ideas that we explained in the simplified setting of this article. The implementation is publicly available as source code and in an easy-to-use web interface [Hof16].

The type inference is technically similar to the inference of upper bounds [HJ03]. We first integrate the structural rules of the respective type system in the syntax directed rules. For example, weakening and relaxation is applied at branching points such as conditional and pattern matching. We then compute a type derivation in which all resource annotations are replace by (yet unknown) variables. For each type rule we produce a set of linear constraints

Constant Function	LOC	Metric	Resource Usage	Time
cond_rev : $(L(\text{int}), L(\text{int}), \text{bool}) \to \text{unit}$	20	steps	$13n{+}13x{+}35$	0.03s
trunc_rev : $(L(\text{int}), \text{int}) \to L(\text{int})$	28	func calls	$1n$	0.06s
ipquery : $L(\text{logline}) \to (L(\text{int}), L(\text{int}))$	86	steps	$86n{+}99$	0.86s
kmeans : $L(\text{float}, \text{float}) \to L(\text{float}, \text{float})$	170	steps	$1246n{+}3784$	8.18s
tea_enc : $(L(\text{int}), L(\text{int}), \text{nat}) \to L(\text{int})$	306	ticks	$128n^2z{+}32nxz{+}1184nz{+}96n{+}128z{+}96$	13.73s
tea_dec : $(L(\text{int}), L(\text{int}), \text{nat}) \to L(\text{int})$	306	ticks	$128n^2z{+}32nxz{+}1184nz{+}96n{+}96z{+}96$	14.34s

Table 1: The computed resource usage in case of constant function, the computed lower and upper bounds, and the run-time of the analysis in seconds. Note that constant resource usage, lower and upper bounds are the same when a function is constant. In the computed resource usage n is the size of the first argument, $m = \max_{1\leq i\leq n} m_i$ where m_i are the sizes the first argument's elements, x is the size of the second argument, $y = \max_{1\leq i\leq n} y_i$ where y_i are the sizes the second argument's elements, and z is the value of the third argument.

that specify the properties of valid annotations. These linear constraints are then solved by the LP solver CLP to obtain a type derivation in which the annotations are rational numbers.

An interesting challenge lies in finding a solution for the linear constraints that leads to the best bound for a given function. For upper bounds, we simply disregard the potential of the result type and provide an objective function that minimizes the annotations of the arguments. The same strategy works the constant-time type systems. An interesting property is that the solution to the linear program is unique if we require that the potential of the result type is zero. To obtain the optimal lower bound we want to maximize the potential of the arguments and minimize the potential of the result. We currently simply maximize the potential of the arguments while requiring the potential of the result to be zero. Another approach would be to first minimize the output potential and then maximize the input potential.

Resource-aware noninterference We are currently integrating our constant-time type system with FlowCaml [Sim03]. The combined inference is based on the typing rules in Fig. 11. It is possible to derive a set of type inference rules in the same way as for FlowCaml [SMZ99, PS02]. One of the challenges in the integration is interfacing FlowCaml's type inference with our constant-time type system in rule SR:C-GEN. In the implementation, we intend for each application of SR:C-GEN to generate an intermediate representation of the expression in RAML for the expression under consideration, in which all types are annotated with fresh resource annotations along with the set of variables X. The expression is marked with the qualifier const if a RAML can prove that it is constant time. The type inference algorithm always tries to apply the syntax-directed rules first before using SR:C-GEN.

Evaluation Table 1 and shows the verification of constant resource usage, while Table 2 shows the computation of lower and upper bounds for number of functions with different size in terms of number of lines of code (LOC). The cost models are specified by several different cost metric, i.e., number of evaluation steps, number of multiplication operations. Note that the computed upper bounds are also the resource usages of functions which are padded using *consume* expressions. The experiments were run on machine with Intel Core i5 2.4 GHz processor and 8GB RAM under the OS X 10.11.5. The run-time of the analysis of varies from 0.02 to 14.34 seconds depending on the function code's complexity. The example programs that we analyzed consist

Function	LOC	Metric	Lower Bound	Time	Upper Bound	Time
compare : $(L(\text{int}), L(\text{int})) \rightarrow \text{bool}$	60	steps	7	0.05s	$16n+7$	0.09s
find : $(L(\text{int}), \text{int}) \rightarrow \text{bool}$	40	steps	5	0.04s	$14n+5$	0.02s
rsa : $(L(\text{bool}), \text{int}, \text{int}) \rightarrow \text{int}$	42	multiplications	$1n$	0.07s	$2n$	0.05s
filter : $L(\text{int}) \rightarrow L(\text{int})$	30	steps	$13n+5$	0.05s	$20n+5$	0.04s
isortlist : $L(L(\text{int})) \rightarrow L(L(\text{int}))$	60	steps	$21n+5$	0.13s	$12n^2+9n$ $+10n^2m$ $-10nm+5$	0.43s
bfs_tree : $(\text{btree}, \text{int}) \rightarrow \text{btree option}$	116	steps	15	0.30s	$92n+24$	0.32s

Table 2: The computed resource usage in case of constant function, the computed lower and upper bounds, and the run-time of the analysis in seconds. Note that constant resource usage, lower and upper bounds are the same when a function is constant. In the computed resource usage n is the size of the first argument, $m = \max_{1\leq i\leq n} m_i$ where m_i are the sizes the first argument's elements.

of commonly-used primitives (cond_rev, trunc_rev, compare, find, filter), functions related to cryptography (tea_enc, tea_dec, rsa), and examples taken from Haeberlan et al. [HPN11] (ipquery, kmeans). The encryption functions tea_enc and tea_dec correspond to the encryption and decryption routines of the Corrected Block Tiny Encryption Algorithm [Yar10], a block cipher presented by Needham and Wheeler in an unpublished technical report in 1998. Our implementation correctly identifies these operations as constant-time in the number of primitive operations performed. We applied this cost model for these examples due to the presence of bitwise operations in the original algorithm, which are not currently supported in RAML. In order to derive a more meaningful bound, we implemented bitwise operations in the example source and counted them as single operations.

The two examples taken from Haeberlen et al. [HPN11] were originally created in a study of timing attacks in differentially-private data processing systems. ipquery applies pattern matching to a database derived from Apache server logs, counting the number of matches and non-matches. kmeans implements the k-means clustering algorithm [Mac67], which partitions a set of geometric points into k clusters that minimize the total inter-cluster distance between points. Haeberlen et al. demonstrated that when a query applied to a dataset introduces attacker-observable timing variations, then the privacy guarantees provided by differential privacy are negated. To address this, they proposed a mitigation approach that enforces constant-time behavior by aborting or padding the query's runtime. Our implementation is able to determine that these queries were constant-time to begin with, and thus did not need black-box mitigation.

7 Related work

Resource bounds Our work builds on past research on automatic amortized resource analysis (AARA). AARA has been introduced by Hofmann and Jost for a strict first-order functional language with built-in data types to derive linear heap-memory bounds [HJ03]. It has then been extended to polynomial bounds [HH10, HAH12, HS14] for strict and higher-order [JHLH10, Ano15] functions. AARA has also been used to derive linear bounds for lazy functional programs [SVF+12, VJFH15] and object-oriented programs [HJ06, HR13]. In another line of work, the technique has been integrated into separation logic [Atk10] to derive bounds that depend on mutable data structures, and into Hoare logic to derive linear bounds that depend on integers [CHRS14, CHS15]. The potential method of amortized analysis has also been used to manually verify the complexity of algorithms and data structures using proof assistants [Nip15, CP15].

As discussed in the introduction, AARA has been successfully extended to other resources and language features [JHLH10, Cam09, SVF$^+$12, VJFH15, HJ06, HR13, Atk10] and to polynomial bounds [HH10, HAH11, HAH12, HM14, HM15]. Amortized analysis has also been used to verify bounds on algorithms and data structures with proof assistants [Nip15, CP15]. In contrast to our work, these techniques can only derive upper bounds and prove constant resource consumption. This focus on upper bounds is shared with automatic resource analysis techniques that based on sized types [VH03, Vas08], linear dependent types [LG11, LP13], and other type systems [CW00, Dan08, ÇGA15]. Similarly, semiautomatic analyses [Gro01, Ben04, DLR12, ALM12] focus on upper bounds too.

Automatic resource bound analysis is also actively studied for imperative languages using recurrence relations [ABG12, FH14, AFR15] and abstract interpretation [GMC09, BHHK10, ZSGV11, SZV14, CHK$^+$15]. While these techniques focus on worst-case bounds, it is possible to use similar techniques for deriving lower bounds [AGM13]. The advantage of our method is that it is compositional, deal well with amortization effects, and works for language features such as user-defined data types and higher-order functions. Another approach to (worst-case) bound analysis is based on techniques from term rewriting [AM13, NEG13, BEF$^+$14], which mainly focus on upper bounds. One line of work [FNH$^+$16] derives lower bounds on the *worst-case* behavior of programs which is different from our lower bounds on the best-case behavior.

Side channels Analyzing and mitigating potential sources of side channel leakage is an increasingly well-studied area. Several groups have proposed using type systems or other program analyses to transform programs into constant-time versions by padding out branches and loops with “dummy” commands [Aga00, HS05, CVBS09, ZAM12, BRW06, MPSW06]. Because these systems do not account for timing explicitly, as is the case for our work, this approach will in nearly all cases introduce an unnecessary performance penalty. The most recent of these systems by Zhang et al. [ZAM12] describes an approach for mitigating side channels using a combination of security types, hardware assistance, and predictive mitigation [ZAM11]. Unlike the type system given in Section 4, theirs does not guarantee that information is not leaked through timing. Rather, they show that the amount of this leakage is bounded by the variation of the mitigation commands.

Köpf and Basin [KB07] presented an information-theoretic model for adaptive side channel attacks that occur over multiple runs of a program, as well as an automated analysis for measuring the corresponding leakage. Because their analysis is doubly-exponential in the number of steps taken by the attacker, they describe an approximate version based on a greedy heuristic. Mardziel et al. later generalized this model to probabilistic systems [MAHC14], secrets that change over time, and wait-adaptive adversaries. Pasareanu et al. [PPM16] proposed a symbolic approach for the multi-run setting based on MaxSAT and model counting. Doychev et al. [DFK$^+$13] and Köpf et al. [KMO12] consider cache side channels, and present analyses that over-approximate leakage using model-counting techniques. While these analyses are sometimes able to derive useful bounds on the leakage produced by binaries on real hardware, they do not incorporate security labels to distinguish between different sources, and were not applied to verifying constant-time behavior.

FlowTracker [RQaPA16] and ct-verif [ABB$^+$16] are both constant-time analyses built on top of LLVM which reason about timing and other side-channel behavior indirectly through control and address-dependence on secret inputs. VirtualCert [BBC$^+$14] instruments CompCert with a constant-time analysis based on similar reasoning about control and address-dependence. These approaches are intended for code that has been written in “constant-time style”, and thus impose effective restrictions on the expressiveness of the programs that they will work on. Because our approach reasons about resources explicitly, it imposes no a priori restrictions on program expressiveness.

Information flow A long line of prior work looks at preventing information flows using type systems. Sabelfeld and Myers [SM03] present an excellent overview of much of the early work in this area. The work most closely related to our security type system is FlowCaml [PS02], which provides a type system that enforces noninterference for a core of ML with references, exceptions, and let-polymorphism. The portion of our type system that applies to traditional noninterference coincides with the rules used in FlowCaml. However, the rules in our type system are not only designed to track flows of information, but they are also used to incorporate the information flow and resource usage behavior such as the rules SR:L-IF and SR:L-LET. Moreover, our type system constructs a flexible interface between FlowCaml and the constant resource type system for reasoning about resource consumption, meaning that the rules can be easily adapted to integrate into any information flow type system.

The primary difference between our work and the prior work on information flow type systems is best summarized in terms of our attacker model. Whereas prior work assumes an attacker that can manipulate low-security inputs and observe low-security outputs, our type system enhances this attacker by granting the ability to observe the program’s final resource consumption. This broadens the relevant class of attacks to include resource side channels, which we prevent by extending a traditional information flow type system with explicit reasoning about the resource behavior of the program using AARA.

8 Conclusion

We have introduced new substructural type systems for automatically deriving lower bounds and proving constant resource usage. The evaluation with the implementation in RAML shows that the technique extends beyond the core language that we study in this paper and works for realistic example programs. We have shown how the new type systems can interact with information-flow type systems to prove resource-aware noninterference. Moreover, the type system for constant resource can be used to automatically remove side-channel vulnerabilities from programs.

There are many interesting connection between security and (automatic) quantitative resource analysis that we plan to study in the future. Two concrete projects that we already started are the integration of the type systems for upper and lower bounds with information flow type systems to precisely quantify the resource-based information leakage at certain security levels. Another direction is to more precisely characterize the amount of information that can be obtained about secretes by making one particular resource-usage observation.

References

[AAG+09] Elvira Albert, Puri Arenas, Samir Genaim, Miguel Gómez-Zamalloa, German Puebla, D. Ramírez, G. Román, and Damiano Zanardini. Termination and Cost Analysis with COSTA and its User Interfaces. *Electr. Notes Theor. Comput. Sci.*, 258(1):109–121, 2009.

[ABB+16] Jose Carlos Bacelar Almeida, Manuel Barbosa, Gilles Barthe, François Dupressoir, and Michael Emmi. Verifying constant-time implementations. In *25th USENIX Security Symposium (to appear)*, Austin, TX, August 2016.

[ABG12] Diego Esteban Alonso-Blas and Samir Genaim. On the limits of the classical approach to cost analysis. In *19th Int. Static Analysis Symp. (SAS'12)*, 2012.

[AFR15] Elvira Albert, Jesús Correas Fernández, and Guillermo Román-Díez. Non-cumulative Resource Analysis. In *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, (TACAS'15)*, 2015.

[Aga00] Johan Agat. Transforming out timing leaks. In *Proceedings of the 27th ACM Symposium on Principles of Programming Languages*, 2000.

[AGM13] Elvira Albert, Samir Genaim, and Abu Naser Masud. On the Inference of Resource Usage Upper and Lower Bounds. *ACM Transactions on Computational Logic*, 14(3), 2013.

[AKM+15] Marc Andrysco, David Kohlbrenner, Keaton Mowery, Ranjit Jhala, Sorin Lerner, and Hovav Shacham. On subnormal floating point and abnormal timing. In *2015 IEEE Symposium on Security and Privacy*, 2015.

[ALM12] Martin Avanzini, Ugo Dal Lago, and Georg Moser. Analysing the Complexity of Functional Programs: Higher-Order Meets First-Order. In *29th Int. Conf. on Functional Programming (ICFP'15)*, 2012.

[AM13] Martin Avanzini and Georg Moser. A Combination Framework for Complexity. In *24th International Conference on Rewriting Techniques and Applications (RTA'13)*, 2013.

[Ano15] Anonymous. Towards Automatic Resource Bound Analysis for OCaml. Technical Report TR-??-???, Removed for double blind reviewing, 2015.

[Ano16] Anonymous. Quantifying and Preventing Side Channels with Substructural Type Systems. Technical Report TR-??-???, Removed for double blind reviewing. Submitted as anonymous auxiliary material, 2016.

[AP13] Nadhem J. AlFardan and Kenneth G. Paterson. Lucky thirteen: Breaking the TLS and DTLS record protocols. In *2013 IEEE Symposium on Security and Privacy*, 2013.

[Atk10] Robert Atkey. Amortised Resource Analysis with Separation Logic. In *19th Euro. Symp. on Prog. (ESOP'10)*, 2010.

[AZM10] Aslan Askarov, Danfeng Zhang, and Andrew C. Myers. Predictive black-box mitigation of timing channels. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, 2010.

[BB03] David Brumley and Dan Boneh. Remote timing attacks are practical. In *Proceedings of the 12th Annual USENIX Security Symposium*, Berkeley, CA, USA, 2003. USENIX Association.

[BB05] D. Brumley and D. Boneh. Remote timing attacks are practical. In *Computer Networks*, 2005.

[BB07] Andrew Bortz and Dan Boneh. Exposing private information by timing web applications. In *Proceedings of the 16th International Conference on World Wide Web*, 2007.

[BBC+14] Gilles Barthe, Gustavo Betarte, Juan Campo, Carlos Luna, and David Pichardie. System-level non-interference for constant-time cryptography. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security*, 2014.

[BEF+14] Marc Brockschmidt, Fabian Emmes, Stephan Falke, Carsten Fuhs, and Jürgen Giesl. Alternating Runtime and Size Complexity Analysis of Integer Programs. In *20th Int. Conf. on Tools and Alg. for the Constr. and Anal. of Systems (TACAS'14)*, 2014.

[Ben04] Ralph Benzinger. Automated Higher-Order Complexity Analysis. *Theor. Comput. Sci.*, 318(1-2):79–103, 2004.

[BHHK10] Régis Blanc, Thomas A. Henzinger, Thibaud Hottelier, and Laura Kovács. ABC: Algebraic Bound Computation for Loops. In *Logic for Prog., AI., and Reasoning - 16th Int. Conf. (LPAR'10)*, 2010.

[BRW06] Gilles Barthe, Tamara Rezk, and Martijn Warnier. Preventing timing leaks through transactional branching instructions. *Electron. Notes Theor. Comput. Sci.*, 153(2), May 2006.

[Cam09] Brian Campbell. Amortised Memory Analysis using the Depth of Data Structures. In *18th Euro. Symp. on Prog. (ESOP'09)*, 2009.

[ÇGA15] Ezgi Çiçek, Deepak Garg, and Umut A. Acar. Refinement Types for Incremental Computational Complexity. In *24th European Symposium on Programming (ESOP'15)*, 2015.

[CHK+15] Pavol Cerný, Thomas A. Henzinger, Laura Kovács, Arjun Radhakrishna, and Jakob Zwirchmayr. Segment Abstraction for Worst-Case Execution Time Analysis. In *24th European Symposium on Programming (ESOP'15)*, 2015.

[CHRS14] Quentin Carbonneaux, Jan Hoffmann, Tahina Ramananandro, and Zhong Shao. End-to-End Verification of Stack-Space Bounds for C Programs. In *Conf. on Prog. Lang. Design and Impl. (PLDI'14)*, page 30, 2014.

[CHS15] Quentin Carbonneaux, Jan Hoffmann, and Zhong Shao. Compositional Certified Resource Bounds. In *36th Conf. on Prog. Lang. Design and Impl. (PLDI'15)*, 2015.

[CHVV03] Brice Canvel, Alain Hiltgen, Serge Vaudenay, and Martin Vuagnoux. *Password Interception in a SSL/TLS Channel.* Berlin, Heidelberg, 2003.

[CP15] Arthur Charguéraud and François Pottier. Machine-Checked Verification of the Correctness and Amortized Complexity of an Efficient Union-Find Implementation. In *Interactive Theorem Proving - 6th International Conference (ITP'15)*, 2015.

[CVBS09] Bart Coppens, Ingrid Verbauwhede, Koen De Bosschere, and Bjorn De Sutter. Practical mitigations for timing-based side-channel attacks on modern x86 processors. In *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy*, 2009.

[CW00] Karl Crary and Stephanie Weirich. Resource Bound Certification. In *27th ACM Symp. on Principles of Prog. Langs. (POPL'00)*, pages 184–198, 2000.

[Dan08] Nils Anders Danielsson. Lightweight Semiformal Time Complexity Analysis for Purely Functional Data Structures. In *35th ACM Symp. on Principles Prog. Langs. (POPL'08)*, 2008.

[DFK+13] Goran Doychev, Dominik Feld, Boris Kopf, Laurent Mauborgne, and Jan Reineke. Cacheaudit: A tool for the static analysis of cache side channels. In *Proceedings of the 22nd USENIX Security Symposium*, Washington, D.C., 2013. USENIX.

[DLR12] Norman Danner, Daniel R. Licata, and Ramyaa Ramyaa. Denotational Cost Semantics for Functional Languages with Inductive Types. In *29th Int. Conf. on Functional Programming (ICFP'15)*, 2012.

[FH14] Antonio Flores-Montoya and Reiner Hähnle. Resource Analysis of Complex Programs with Cost Equations. In *Programming Languages and Systems - 12th Asian Symposiu (APLAS'14)*, 2014.

[FNH+16] Florian Frohn, M. Naaf, Jera Hensel, Marc Brockschmidt, and Jürgen Giesl. Lower Runtime Bounds for Integer Programs. In *Automated Reasoning - 8th International Joint Conference (IJCAR'16)*, 2016.

[FS00] Edward W. Felten and Michael A. Schneider. Timing attacks on web privacy. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, 2000.

[GBK11] David Gullasch, Endre Bangerter, and Stephan Krenn. Cache games – bringing access-based cache attacks on aes to practice. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, 2011.

[GM82] Joseph Goguen and José Meseguer. Security Policies and Security Models. In *In Proceedings of the 1982 IEEE Symposium on Security and Privacy*, 1982.

[GMC09] Sumit Gulwani, Krishna K. Mehra, and Trishul M. Chilimbi. SPEED: Precise and Efficient Static Estimation of Program Computational Complexity. In *36th ACM Symp. on Principles of Prog. Langs. (POPL'09)*, 2009.

[Gro01] Bernd Grobauer. Cost Recurrences for DML Programs. In *6th Int. Conf. on Funct. Prog. (ICFP'01)*, pages 253–264, 2001.

[HAH11] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate Amortized Resource Analysis. In *38th ACM Symp. on Principles of Prog. Langs. (POPL'11)*, 2011.

[HAH12] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate Amortized Resource Analysis. *ACM Trans. Program. Lang. Syst.*, 2012.

[Har12] Robert Harper. *Practical Foundations for Programming Languages.* Cambridge University Press, 2012.

[HH10] Jan Hoffmann and Martin Hofmann. Amortized Resource Analysis with Polynomial Potential. In *19th Euro. Symp. on Prog. (ESOP'10)*, 2010.

[HJ03] Martin Hofmann and Steffen Jost. Static Prediction of Heap Space Usage for First-Order Functional Programs. In *30th ACM Symp. on Principles of Prog. Langs. (POPL'03)*, 2003.

[HJ06] Martin Hofmann and Steffen Jost. Type-Based Amortised Heap-Space Analysis. In *15th Euro. Symp. on Prog. (ESOP'06)*, 2006.

[HM14] Martin Hofmann and Georg Moser. Amortised Resource Analysis and Typed Polynomial Interpretations. In *Rewriting and Typed Lambda Calculi (RTA-TLCA;14)*, 2014.

[HM15] Martin Hofmann and Georg Moser. Multivariate Amortised Resource Analysis for Term Rewrite Systems. In *13th International Conference on Typed Lambda Calculi and Applications (TLCA'15)*, 2015.

[Hof16] Jan Hoffmann. RAML Web Site, 2016.

[HPN11] Andreas Haeberlen, Benjamin C. Pierce, and Arjun Narayan. Differential privacy under fire. In *Proceedings of the 20th Annual USENIX Security Symposium*, 2011.

[HR98] Nevin Heintze and Jon G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *In Conference Record of the 25th ACM Symposium on Principles of Programming Languages*, 1998.

[HR13] Martin Hofmann and Dulma Rodriguez. Automatic Type Inference for Amortised Heap-Space Analysis. In *22nd Euro. Symp. on Prog. (ESOP'13)*, 2013.

[HS05] Daniel Hedin and David Sands. Timing aware information flow security for a javacard-like bytecode. *Electron. Notes Theor. Comput. Sci.*, 141(1), December 2005.

[HS14] Jan Hoffmann and Zhong Shao. Type-Based Amortized Resource Analysis with Integers and Arrays. In *12th International Symposium on Functional and Logic Programming (FLOPS'14)*, 2014.

[JHLH10] Steffen Jost, Kevin Hammond, Hans-Wolfgang Loidl, and Martin Hofmann. Static Determination of Quantitative Resource Usage for Higher-Order Programs. In *37th ACM Symp. on Principles of Prog. Langs. (POPL'10)*, 2010.

[JLH$^+$09] Steffen Jost, Hans-Wolfgang Loidl, Kevin Hammond, Norman Scaife, and Martin Hofmann. Carbon Credits for Resource-Bounded Computations using Amortised Analysis. In *16th Symp. on Form. Meth. (FM'09)*, 2009.

[KB07] Boris Köpf and David Basin. An information-theoretic model for adaptive side-channel attacks. In *Proceedings of the 14th ACM Conference on Computer and Communications Security*, 2007.

[KD09] B. Köpf and M. Dürmuth. A provably secure and efficient countermeasure against timing attacks. In *Proceedings of the 22nd Annual IEEE Computer Security Foundations Symposium*, July 2009.

[KMO12] Boris Köpf, Laurent Mauborgne, and Martín Ochoa. Automatic quantification of cache side-channels. In *Proceedings of the 24th International Conference on Computer Aided Verification*, 2012.

[Koc96] Paul C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *16th Annual International Cryptology Conference (CRYPTO'96)*, 1996.

[LG11] Ugo Dal Lago and Marco Gaboardi. Linear Dependent Types and Relative Completeness. In *26th IEEE Symp. on Logic in Computer Science (LICS'11)*, 2011.

[LP13] Ugo Dal Lago and Barbara Petit. The Geometry of Types. In *40th ACM Symp. on Principles Prog. Langs. (POPL'13)*, 2013.

[Mac67] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, Calif., 1967. University of California Press.

[MAHC14] Piotr Mardziel, Mario S. Alvim, Michael Hicks, and Michael R. Clarkson. Quantifying information flow for dynamic secrets. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, 2014.

[MPSW06] David Molnar, Matt Piotrowski, David Schultz, and David Wagner. The program counter security model: Automatic detection and removal of control-flow side channel attacks. In *Proceedings of the 8th International Conference on Information Security and Cryptology*, 2006.

[NEG13] Lars Noschinski, Fabian Emmes, and Jürgen Giesl. Analyzing Innermost Runtime Complexity of Term Rewriting by Dependency Pairs. *J. Autom. Reasoning*, 51(1):27–56, 2013.

[Nip15] Tobias Nipkow. Amortized Complexity Verified. In *Interactive Theorem Proving - 6th International Conference (ITP'15)*, 2015.

[ODL15] Oswaldo Olivo, Isil Dillig, and Calvin Lin. Static Detection of Asymptotic Performance Bugs in Collection Traversals. In *Conference on Programming Language Design and Implementation (PLDI'15)*, 2015.

[OST06] D. Osvik, A. Shamir, and E. Tromer. Cache attacks and countermeasures: the case of AES. In *Topics in Cryptography*, 2006.

[PPM16] Corina Pasareanu, Quoc-San Phan, and Pasquale Malacaria. Multi-run side-channel analysis using symbolic execution and max-smt. In *Proceedings of the 29th IEEE Computer Security Foundations Symposium*, 2016.

[PS02] Francois Pottier and Vicent Simonet. Information Flow Inference for ML. In *In Proceedings of the 29th ACM Symposium on Principles Of Programming Languages*, 2002.

[RQaPA16] Bruno Rodrigues, Fernando Magno Quintão Pereira, and Diego F. Aranha. Sparse representation of implicit flows with applications to side-channel detection. In *Proceedings of the 25th International Conference on Compiler Construction*, 2016.

[Sim03] Vincent Simonet. Flow Caml. `http://cristal.inria.fr/~simonet/soft/flowcaml/`, 2003.

[SM03] Andrei Sabelfeld and Andrew C. Myers. Language-Based Information Flow Security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.

[SMZ99] Martin Sulzmann, Martin Müller, and Christoph Zenger. Hindley/Milner Style Type Systems in Constraint Form. In *Research Report ACRC–99–009, University of South Australia, School of Computer and Information Science*, 1999.

[SVF+12] Hugo R. Simões, Pedro B. Vasconcelos, Mário Florido, Steffen Jost, and Kevin Hammond. Automatic Amortised Analysis of Dynamic Memory Allocation for Lazy Functional Programs. In *17th Int. Conf. on Funct. Prog. (ICFP'12)*, 2012.

[SZV14] Moritz Sinn, Florian Zuleger, and Helmut Veith. A Simple and Scalable Approach to Bound Analysis and Amortized Complexity Analysis. In *Computer Aided Verification - 26th Int. Conf. (CAV'14)*, 2014.

[Tar85] Robert Endre Tarjan. Amortized Computational Complexity. *SIAM J. Algebraic Discrete Methods*, 6(2):306–318, 1985.

[Vas08] Pedro Vasconcelos. *Space Cost Analysis Using Sized Types.* PhD thesis, School of Computer Science, University of St Andrews, 2008.

[VH03] Pedro B. Vasconcelos and Kevin Hammond. Inferring Costs for Recursive, Polymorphic and Higher-Order Functional Programs. In *Int. Workshop on Impl. of Funct. Langs. (IFL'03)*, 2003.

[VJFH15] Pedro B. Vasconcelos, Steffen Jost, Mário Florido, and Kevin Hammond. Type-Based Allocation Analysis for Co-recursion in Lazy Functional Languages. In *24th European Symposium on Programming (ESOP'15)*, 2015.

[VSI96] Dennis Volpano, Geoffrey Smith, and Cynthia Irvine. A sound type system for secure flow analysis. In *Journal of Computer Security*, pages 167–187, 1996.

[W+08] Reinhard Wilhelm et al. The Worst-Case Execution-Time Problem — Overview of Methods and Survey of Tools. *ACM Trans. Embedded Comput. Syst.*, 7(3), 2008.

[Wal02] David Walker. Substructural Type Systems. *Advanced Topics in Types and Programming Languages. MIT Press*, pages 3–43, 2002.

[Yar10] Elias Yarrkov. Cryptanalysis of XXTEA, 2010.

[ZAM11] Danfeng Zhang, Aslan Askarov, and Andrew C. Myers. Predictive mitigation of timing channels in interactive systems. In *Proceedings of the 18th ACM Conference on Computer and Communications Security*, 2011.

[ZAM12] Danfeng Zhang, Aslan Askarov, and Andrew C. Myers. Language-Based Control and Mitigation of Timing Channels. In *In Proceedings of the SIGPLAN Conference on Programming Language Design and Implementation*, 2012.

[ZJRR14] Yinqian Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. Cross-tenant side-channel attacks in paas clouds. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security*, 2014.

[ZSGV11] Florian Zuleger, Moritz Sinn, Sumit Gulwani, and Helmut Veith. Bound Analysis of Imperative Programs with the Size-change Abstraction. In *18th Int. Static Analysis Symp. (SAS'11)*, 2011.