

Safety Guarantees from Explicit Resource Management

David Aspinall, Patrick Maier, and Ian Stark

Laboratory for Foundations of Computer Science
School of Informatics, The University of Edinburgh, Scotland
{David.Aspinall,Patrick.Maier,Ian.Stark}@ed.ac.uk

Abstract. We present a language and a program analysis that certifies the safe use of flexible resource management idioms, in particular advance reservation or “block booking” of costly resources. This builds on previous work with *resource managers* that carry out runtime safety checks, by showing how to assist these with compile-time checks. We give a small ANF-style language with explicit resource managers, and introduce a type and effect system that captures their runtime behaviour. In this setting, we identify a notion of *dynamic safety* for running code, and show that dynamically safe code may be executed without runtime checks. We show a similar *static safety* property for type-safe code, and prove that static safety implies dynamic safety. The consequence is that typechecked code can be executed without runtime instrumentation, and is guaranteed to make only appropriate use of resources.

1 Introduction

Safe management of resources is a crucial aspect of software correctness. Bad resource management impacts reliability and security. The more expensive a resource or the more complex its usage pattern, the more important is good management. For example, a media player could crash badly, leaving the hardware in a messy state, if its memory management was governed by the overly optimistic assumption that every request for memory will succeed. Malware on a mobile phone can defraud an unaware user by maliciously sending text messages to premium rate numbers, if there is no effective management of network access [12]. On current mobile platforms such as Java MIDP 2.0, management of network access is commonly left to the user, but users can easily be deceived by social engineering attacks.

Unfortunately, current programming languages do not provide special mechanisms for resource management. Therefore, programmers can only hope that their applications are resource safe, or use necessarily imprecise analyses to try to show this. For example, there are type systems that over-approximate (hopefully tightly) the memory requirements of an application [6], and static analyses that over-approximate the number of text messages being sent by an application [7].

These approaches may fail if a dynamic set of resources must be managed, as with *bulk messaging* where the user wants to send a text message to a number of recipients selected from an address book. Because of the cost of sending text messages, the user must authorise each recipient (i. e., their phone number) explicitly. This could happen individually, just before each message is being sent, or collectively, before sending the

```

send_bulk ::
λ let (r) = res_from_nums (nums) in
  let (m) = init () in
  let (m',r') = enable (m,r) in
  let (n) = size (r') in
  if n then let () = consume (m') in
    ret ()
  else let (m'') = send_msgs (msg,nums,m') in
    let (m''') = assertEmpty (m'') in
    let () = consume (m''') in
    ret () :
(msg:str, nums:str[]) → ()

res_from_nums ::
λ let (i) = length (nums) in
  let (r) = empty () in
  let (r') = res_from_nums' (nums,r,i) in
  ret (r') :
(nums:str[]) → (r':res{ })

res_from_nums' ::
λ if i then let (i') = sub (i,1) in
  let (num) = read (nums,i') in
  let (c) = fromstr (num) in
  let (r,c) = single (c,1) in
  let (r') = sum (r, r,c) in
  let (r'') = res_from_nums' (nums,r'',i') in
  ret (r'')
else let (r') = id (r) in
  ret (r') :
(nums:str[], r:res{ }, i:int) → (r':res{ })

send_msgs ::
λ let (i) = length (nums) in
  let (m') = send_msgs' (msg,nums,m,i) in
  ret (m') :
(msg:str, nums:str[], m:mgr) → (m':mgr)

send_msgs' ::
λ if i then let (i') = sub (i,1) in
  let (num) = read (nums,i') in
  let (m'') = send_msg (msg,num,m) in
  let (m''') = send_msgs' (msg,nums,m'',i') in
  ret (m''')
else let (m') = id (m) in
  ret (m') :
(msg:str, nums:str[], m:mgr, i:int) → (m':mgr)

send_msg ::
λ let (c) = fromstr (num) in
  let (r) = single (c,1) in
  let (m',m_r) = split (m,r) in
  let (m_r') = assertAtLeast (m_r,r) in
  let () = prim_send_msg (msg,num) in
  let () = consume (m_r') in
  ret (m') :
(msg:str, num:str, m:mgr) → (m':mgr)

prim_send_msg ::
λ ... :
(msg:str, num:str) → ()

```

Fig. 1. Bulk messaging application

first message. Collective authorisation, or *block booking* of resources, is preferable but requires detailed resource management, keeping track of the (multi-)set of authorised resources – in this case the permitted phone numbers.

In this paper, we present a language-based mechanism that provides programmers with a safe way to control complex resource usage patterns using a notion of *resource manager*. Figure 1 shows the code of a bulk messaging application using resource managers in our intermediate-level functional programming language. The language and functions used will be explained in full detail in Section 2; for now, we just give an outline of operation. The function `send_bulk` calls `send_msgs` to send the message `msg` to the phone numbers stored in the array `nums`. Along with these two arguments `send_msgs` takes a resource manager `m'` which encapsulates the resources that have been authorised (during the call to `enable`) to send the messages. For each phone number in `nums`, `send_msgs` calls the wrapper function `send_msg`, passing along a resource manager. Prior to calling the primitive `send` function `prim_send_msg`, the wrapper checks (using `assertAtLeast`) whether its input manager `m` contains the resource required to send a message to `num`; if the resource is not present, the program will abort with a runtime error, otherwise `send_msg` removes the resource from the manager (using `split`), and returns the modified manager as `m'`.

The bulk messaging application is (dynamically) resource safe by construction, as the resource managers will trap attempts to abuse resources. The resource manager

abstraction works in tandem with a static analysis, so that programs which can be proved resource safe statically can be treated more efficiently at runtime by removing the dynamic accounting code. In Section 3.2, we prove resource safety statically for the bulk messaging application.

Our contribution is two-fold. In Section 2, we develop a functional programming language for coding complex resource idioms, such as block booking resources in the bulk messaging application. The language is essentially a first-order functional language in administrative normal form (ANF) [10] with a novel type system serving two purposes. First, the type system names input and output parameters of functions and avoids shadowing of previously bound names, thus admitting to view functions as relations (expressed by logical formulae) between their input and output parameters. Second, the language includes a special, linear type for resource managers, where linearity serves as a means of introducing stateful objects into an otherwise pure functional language. Resource managers track what resources a program is allowed to use, and the operational semantics causes the program to go wrong (i. e., abort with a runtime error) as soon as it attempts to abuse resources. This induces a notion of *dynamic resource safety*, which holds if a program never attempts to abuse resources. In this case, accounting is not necessary. As our first result, we show that erasing resource managers does not alter the semantics of dynamically resource safe programs.

Decisions about which resources programs may use are typically guided by *resource policies*. From the point of view of a program, a policy is simply an oracle determining what resources to grant; and we abstract this as a non-deterministic operation on resource managers. This covers many concrete policy mechanisms, both static (e. g., Java-style policy files) or dynamic (e. g., user interaction); see [3] for more on the interaction of resource managers and policies.

In Section 3 we present our second contribution, an effect type system for deriving relational approximations of functions. These approximations are expressed as pairs of constraints in a first-order logic, specifying a pre- and postcondition (or rather, state transforming action) of a given function, similar to Hoare type theory [11]; note that the use of logical formulae as effects is the rationale behind choosing a programming language where functions have named input and output parameters. Typability of functions in the effect type system induces a notion of *static resource safety*. As our second result, we prove a soundness theorem stating that static implies dynamic resource safety. As a corollary, we show that resource managers can always be erased from statically resource safe programs. Proofs have been omitted due to lack of space.

2 A Programming Language for Resource Management

We introduce a simple programming language with built-in constructs for handling resource managers. The language is essentially a simply-typed first-order functional language in ANF [10], with the additional features that functions take and return tuples of values, function types name input and output arguments, scoping avoids shadowing, and the type of resource managers enforces a linearity restriction on its values. The first three of these features are related to giving the language a relational appeal: for the purpose of specifying and reasoning logically, functions ought to be viewed as relations

$\langle \text{fundecl} \rangle ::= \langle \text{prodtype} \rangle \rightarrow \langle \text{prodtype} \rangle$	<i>(built-in function)</i>
$\lambda \langle \text{exp} \rangle : \langle \text{prodtype} \rangle \rightarrow \langle \text{prodtype} \rangle$	<i>(λ-abstraction)</i>
$\langle \text{exp} \rangle ::= \text{if } \langle \text{val} \rangle \text{ then } \langle \text{exp} \rangle \text{ else } \langle \text{exp} \rangle$	<i>(conditional)</i>
$\text{let } (\langle \text{var} \rangle, \dots, \langle \text{var} \rangle) = \langle \text{fun} \rangle (\langle \text{val} \rangle, \dots, \langle \text{val} \rangle) \text{ in } \langle \text{exp} \rangle$	<i>(function call)</i>
$\text{ret } (\langle \text{var} \rangle, \dots, \langle \text{var} \rangle)$	<i>(return)</i>
$\langle \text{val} \rangle ::= \langle \text{const} \rangle \mid \langle \text{var} \rangle$	
$\langle \text{prodtype} \rangle ::= (\langle \text{var} \rangle : \langle \text{type} \rangle, \dots, \langle \text{var} \rangle : \langle \text{type} \rangle)$	
$\langle \text{type} \rangle ::= \langle \text{datatype} \rangle \mid \text{mgr}$	
$\langle \text{datatype} \rangle ::= \text{unit} \mid \text{int} \mid \text{str} \mid \text{res} \mid \text{res}\{\} \mid \langle \text{datatype} \rangle []$	

Fig. 2. BNF grammar

between input and output parameters. The fourth feature is a means of introducing state into a functional language.

The choice for such a language has been inspired by Grail [2], another first-order functional language in ANF. Moreover, Appel [1] argues that ANF, the intermediate language used by many compilers for functional languages, and SSA, the intermediate representation used by most compilers for imperative languages, are essentially the same thing. Therefore, our language should capture the essence of first-order programming languages, whether functional or imperative.

2.1 Syntax and Static Semantics

Grammar. Figure 2 shows the grammar of the programming language. The nonterminals $\langle \text{fun} \rangle$, $\langle \text{var} \rangle$ and $\langle \text{const} \rangle$ represent *functions*, *variables* and *constants*, respectively. A *program* Π is a partial function from $\langle \text{fun} \rangle$ to $\langle \text{fundecl} \rangle$, i. e., Π maps functions to function declarations, which are either type declarations for built-in functions or λ -abstractions (with type annotations serving as variable binders). We use the notation $\Pi(f) = [\lambda \dots] \sigma \rightarrow \sigma'$ if we are only interested in the type of f , regardless whether f is built-in or a λ -abstraction. By $\text{dom}(\Pi)$, we denote the domain of Π . We denote the restriction of Π to the built-in functions by Π_0 , i. e., $\Pi(f)$ is a λ -abstraction if and only if $f \in \text{dom}(\Pi) \setminus \text{dom}(\Pi_0)$. We assume that Π_0 declares exactly the functions that are shown in Figure 4.

The grammar of *expressions* $e \in \langle \text{exp} \rangle$ and *values* $v \in \langle \text{val} \rangle$ is quite standard for a first-order functional language in ANF. Throughout, functions operate on tuples of values, which is reflected by the syntax for function call and return. The sets of free and bound (by the let-construct) variables of an expression e , denoted by $\text{free}(e)$ and $\text{bound}(e)$ respectively, are defined in the usual way.

Datatypes $\tau \in \langle \text{datatype} \rangle$ comprise the unit type, integers, strings, resources, multisets of resources, and arrays. A *type* $\tau \in \langle \text{type} \rangle$ is either a datatype or the special type of resource managers, denoted **mgr**. See Section 2.2 for the interpretations of types. A tuple $(x_1 : \tau_1, \dots, x_n : \tau_n) \in \langle \text{prodtype} \rangle$ is a *product type* if the variables x_1, \dots, x_n are pairwise distinct. Product types appear to associate types to variables, but they really associate variables *and* types to positions in tuples. A pair of product

types of the form $(x_1:\tau_1, \dots, x_m:\tau_m) \rightarrow (x'_1:\tau'_1, \dots, x'_n:\tau'_n)$ forms a *function type* if the variable sets $\{x_1, \dots, x_m\}$ and $\{x'_1, \dots, x'_n\}$ are disjoint. We call the product types to the left and right of the arrow *argument type* and *return type*, respectively. As an example consider the type of the function `send_msg` from Figure 1. It states that `send_msg` takes two strings and a resource manager and returns a resource manager, while at the same time binding the names of the formal input parameters `msg`, `num` and `m` and announcing that the formal output parameter will be `m'`.

Static typing. A *type environment* Γ is a functional association list of type declarations of the form $x:\tau$, where x is a variable and τ a type. Being functional implies that whenever Γ contains two type declarations $x:\tau$ and $x:\tau'$ we must have $\tau = \tau'$. Therefore, Γ can be seen as a partial function mapping variables to types. By $\text{dom}(\Gamma)$, we denote the domain of this partial function, and for $x \in \text{dom}(\Gamma)$, we may write $\Gamma(x)$ for the unique type which Γ associates to x . We write type environments as comma-separated lists, the empty list being denoted by \emptyset . The restriction $\Gamma|_X$ of Γ to a set of variables X , is defined in the usual way and induces a partial order \succeq type environments, where $\Gamma' \succeq \Gamma$ iff $\Gamma'|_{\text{dom}(\Gamma)} = \Gamma$.

We call a type environment $\Gamma = x_1:\tau_1, \dots, x_n:\tau_n$ *linear* if the variables x_1, \dots, x_n are pairwise distinct. Note that such a linear type environment Γ may be viewed as a product type $\sigma = (x_1:\tau_1, \dots, x_n:\tau_n)$, and vice versa. Occasionally, we will write $\Pi(f) = [\lambda \dots] \Gamma \rightarrow \Delta$ to emphasise that argument and return types of the function f are to be viewed as linear type environments.

Figure 3 shows the typing rules for the programming language. The judgement $C; \Gamma \vdash v : \tau$ expresses that the value v has type τ in type environment Γ and context C , where a *context* is a set of variables (generally the set of variables occurring in some super-expression of v). Note that (T-const) restricts program constants to the unit value, integers and strings, which are the interpretations of the types `unit`, `int` and `str`, respectively (see Section 2.2). All other types are abstract in the sense that their values can only be accessed through built-in functions.

The judgement $C; \Gamma \vdash_{\Pi} e : \sigma$ means that the expression e has product type σ in type environment Γ , context C and program Π . If the program is understood we may write $C; \Gamma \vdash e : \sigma$. There are three things worth noting about expression typing. First, although the type system is linear, weakening and contraction are available to all types but `mgr`, rendering `mgr` the sole linear type of the language. Second, the side condition of (T-let) ensures that let-bound variables do not shadow any variables in the context (which is generally a superset of the set of variables occurring in the let-expression). Third, the rule (T-ret) matches the variables in the return expression to the variables in the product type, thus enforcing that an expression uniformly uses the same variables to return its results (even though these return variables may be let-bound in different branches of the expression). Note that (T-ret) is the only rule to exploit type information about variables. Finally, the judgement $\Gamma \vdash e : \sigma$ (or $\Gamma \vdash_{\Pi} e : \sigma$ if we want to stress the program Π) means that e has product type σ in a linear type environment Γ .

The judgement $\Pi \vdash f$ states that f is a well-typed λ -abstraction in program Π . Note that the syntax of λ -abstractions does not appear to bind variables, yet it does bind the variables hidden in the argument type. Note also that the restriction on function

Typing of values $C; \Gamma \vdash v : \tau$	
(T-var) $\frac{}{C; x:\tau \vdash x:\tau}$ if $x \in C$	(T-const) $\frac{}{C; \emptyset \vdash d:\tau}$ if $\begin{cases} d \in \tau \wedge \\ \tau \in \{\mathbf{unit}, \mathbf{int}, \mathbf{str}\} \end{cases}$
Typing of expressions $C; \Gamma \vdash e : \sigma$	
(T-weak) $\frac{C; \Gamma \vdash e : \sigma}{C; \Gamma, x:\tau \vdash e : \sigma}$ if $\begin{cases} x \in C \wedge \\ \tau \neq \mathbf{mgr} \end{cases}$	(T-contr) $\frac{C; \Gamma, x:\tau, x:\tau \vdash e : \sigma}{C; \Gamma, x:\tau \vdash e : \sigma}$ if $\tau \neq \mathbf{mgr}$
(T-if) $\frac{C; \Gamma \vdash v : \mathbf{int} \quad C; \Gamma' \vdash e_1 : \sigma \quad C; \Gamma' \vdash e_2 : \sigma}{C; \Gamma, \Gamma' \vdash \mathbf{if } v \mathbf{ then } e_1 \mathbf{ else } e_2 : \sigma}$	(T-xch) $\frac{C; \Gamma, \Gamma' \vdash e : \sigma}{C; \Gamma', \Gamma \vdash e : \sigma}$
(T-ret) $\frac{C; \Gamma_1 \vdash x_1 : \tau_1 \quad \dots \quad C; \Gamma_n \vdash x_n : \tau_n}{C; \Gamma_1, \dots, \Gamma_n \vdash \mathbf{ret } (x_1, \dots, x_n) : (x_1:\tau_1, \dots, x_n:\tau_n)}$	
(T-let) $\frac{\begin{array}{c} \Pi(f) = [\lambda \dots](z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow (z'_1:\tau'_1, \dots, z'_n:\tau'_n) \\ C; \Gamma_1 \vdash v_1 : \tau_1 \quad \dots \quad C; \Gamma_n \vdash v_m : \tau_m \\ C \cup \{x'_1, \dots, x'_n\}; \Gamma', x'_1:\tau'_1, \dots, x'_n:\tau'_n \vdash e' : \sigma'' \end{array}}{C; \Gamma_1, \dots, \Gamma_m, \Gamma' \vdash \mathbf{let } (x'_1, \dots, x'_n) = f(v_1, \dots, v_m) \mathbf{ in } e' : \sigma''}$ if (*)	
where (*) $\begin{cases} x'_1, \dots, x'_n \text{ pairwise distinct } \wedge \\ x'_1, \dots, x'_n \notin C \cup \text{dom}(\Gamma') \end{cases}$	
Typing of expressions $\Gamma \vdash e : \sigma$	Well-typedness of λ-abstractions $\Pi \vdash f$
(T-lin) $\frac{\text{dom}(\Gamma); \Gamma \vdash e : \sigma}{\Gamma \vdash e : \sigma}$ if Γ linear	(T-lam) $\frac{\begin{array}{c} \Pi(f) = \lambda e : (x_1:\tau_1, \dots, x_m:\tau_m) \rightarrow \sigma' \\ x_1:\tau_1, \dots, x_m:\tau_m \vdash e : \sigma' \end{array}}{\Pi \vdash f}$

Fig. 3. Typing rules (for a fixed program Π)

types means that the return variables of the body of a λ -abstraction must be disjoint from its argument variables. Finally, we call a program Π *well-typed* if $\Pi \vdash f$ for all $f \in \text{dom}(\Pi) \setminus \text{dom}(\Pi_0)$.

Lemma 1. *Let e be an expression (referring to an implicit program Π), Γ a type environment and σ a product type.*

1. *If $\Gamma \vdash e : \sigma$ then $\text{free}(e) \subseteq \text{dom}(\Gamma)$ and $\text{bound}(e) \cap \text{dom}(\Gamma) = \emptyset$.*
2. *If $\Gamma \vdash e : \sigma$ and $X \supseteq \text{free}(e)$ then $\Gamma|_X \vdash e : \sigma$.*

2.2 Interpretation of Types and Effects of Built-in Functions

Constraints. To provide a formal semantics for the built-in functions, we introduce a many-sorted first-order language \mathcal{L} with equality. Sorts of \mathcal{L} are the datatypes of the programming language (note that this excludes the type \mathbf{mgr}). Formulae of \mathcal{L} are formed from atomic formulae using the usual Boolean connectives $\neg, \wedge, \vee, \Rightarrow$ and \Leftrightarrow (in decreasing order of precedence), and the quantifiers $\forall x:\tau$ and $\exists x:\tau$, where $x \in \langle \text{var} \rangle$ is a variable and $\tau \in \langle \text{datatype} \rangle$ a sort. Atomic formulae are the Boolean constants \top and \perp , or are constructed from terms using the binary equality predicate \approx (which

is available for all sorts), the binary inequality predicate \leq on sort **int** or the binary inclusion predicate \subseteq on sort **res** $\{\}$. Terms are constructed from variables in $\langle \text{var} \rangle$ and the term constructors, which are introduced below, alongside associating the sorts to specific interpretations.

Sort unit is interpreted by the one-element set $\{\star\}$. Its only constant is \star . There are no function symbols.

Sort int is interpreted by the integers with infinity. Constants are the integers plus ∞ .

Function symbols are the usual $- : \mathbf{int} \rightarrow \mathbf{int}$ and $+, \cdot, /, \% : \mathbf{int} \times \mathbf{int} \rightarrow \mathbf{int}$ (where $/$ and $\%$ denote integer division and remainder, respectively).

Sort str is interpreted by the set of strings (over some fixed but unspecified alphabet).

Constants are all strings. The only function symbol is $++ : \mathbf{str} \times \mathbf{str} \rightarrow \mathbf{str}$ (concatenation).

Sort res is interpreted by an arbitrary infinite set (whose elements are termed *resources*). There are no constants, and $\text{fromstr} : \mathbf{str} \rightarrow \mathbf{res}$, an embedding of strings into resources, is the only one function symbol.

Sort res $\{\}$ is interpreted by multisets of resources. It features the constant \emptyset (empty multiset) and the function symbols $\cap, \cup, \uplus : \mathbf{res}\{\} \times \mathbf{res}\{\} \rightarrow \mathbf{res}\{\}$ (intersection, union and sum of multisets, respectively), $|_ : \mathbf{res}\{\} \rightarrow \mathbf{int}$ (size of a multiset), $\text{count} : \mathbf{res}\{\} \times \mathbf{res} \rightarrow \mathbf{int}$ (counting the multiplicity of a resource in a multiset) and $\{- : _ \} : \mathbf{res} \times \mathbf{int} \rightarrow \mathbf{res}\{\}$ (constructing a “singleton” multiset containing a given resource with a given multiplicity and nothing else).

Sort τ \square is interpreted by integer-indexed arrays of elements of sort τ , where an integer-indexed array is a function from an initial segment of the natural numbers to τ . This sort features the constant *null* (array of length 0) and the function symbols $\text{len} : \tau\square \rightarrow \mathbf{int}$ (length of an array), $_[-] : \tau\square \times \mathbf{int} \rightarrow \tau$ (reading at a given index) and $_[- := _] : \tau\square \times \mathbf{int} \times \tau \rightarrow \tau\square$ (updating a given index with a given value). Note that the values of $a[i]$ and $a[i := v]$ are generally unspecified if the index i is out of bounds (i. e., $i < 0$ or $i \geq \text{len}(a)$). As an exception, for $i = \text{len}(a)$, the array $a[i := v]$ properly extends a , i. e., $\text{len}(a[i := v]) = \text{len}(a) + 1$. This models vectors that can grow in size.

Treating the type **mgr** as an alias for the sort **res** $\{\}$, type environments can be seen as associating sorts to variables. Given a type environment Γ and constraint $\phi \in \mathcal{L}$, we write $\Gamma \vdash \phi$ if ϕ is well-sorted w. r. t. Γ ; note that this entails $\text{free}(\phi) \subseteq \text{dom}(\Gamma)$, where $\text{free}(\phi)$ is the set of free variables in ϕ .

Substitutions. A *substitution* μ maps variables $x \in \langle \text{var} \rangle$ to values $\mu(x) \in \langle \text{val} \rangle$ (which are variables again or constants, not arbitrary terms). We denote the domain of a substitution μ by $\text{dom}(\mu)$. Given a type environment Γ , we write $\Gamma\mu$ for the type environment that arises from substituting the variables in Γ according to μ . This is defined recursively: $\emptyset\mu = \emptyset$ and $(\Gamma, x:\tau)\mu$ equals $\Gamma\mu, x:\tau$ if $x \notin \text{dom}(\mu)$, or $\Gamma\mu, \mu(x):\tau$ if $\mu(x) \in \langle \text{var} \rangle$, or $\Gamma\mu$ if $\mu(x) \in \langle \text{const} \rangle$. Note that $\Gamma\mu$ need not be linear even if Γ is. Given a formula ϕ such that $\Gamma \vdash \phi$, we write $\phi\mu$ for the formula obtained by substituting the free variables of ϕ according to μ , avoiding capture. Note that $\Gamma \vdash \phi$ implies $\Gamma\mu \vdash \phi\mu$.

Valuations. Let Γ be a type environment. A Γ -valuation α maps variables $x \in \text{dom}(\Gamma)$ to elements $\alpha(x)$ in the interpretation of the sort $\Gamma(x)$; we call α a *valuation* if we do not care about the particular type environment Γ . We denote the domain of α by $\text{dom}(\alpha)$. Note that $\text{dom}(\alpha) \subseteq \text{dom}(\Gamma)$ but not necessarily $\text{dom}(\alpha) = \text{dom}(\Gamma)$; we call α a *maximal* Γ -valuation if $\text{dom}(\alpha) = \text{dom}(\Gamma)$. Given a Γ -valuation α and a set of variables X , we denote the restriction of α to X by $\alpha|_X$; note that $\text{dom}(\alpha|_X) = \text{dom}(\alpha) \cap X$. Restriction induces a partial order \succeq on Γ -valuations, where $\alpha' \succeq \alpha$ iff $\alpha'|_{\text{dom}(\alpha)} = \alpha$. Given n pairwise distinct variables $x_i \in \text{dom}(\Gamma)$ and corresponding elements d_i in the interpretation of $\Gamma(x_i)$, we write $\alpha\{x_1 \mapsto d_1, \dots, x_n \mapsto d_n\}$ for the Γ -valuation α' that maps the x_i to d_i and all other $x \in \text{dom}(\alpha)$ to $\alpha(x)$. In the special case $\text{dom}(\alpha) = \emptyset$, we may drop α and simply write $\{x_1 \mapsto d_1, \dots, x_n \mapsto d_n\}$.

Entailment. Let $\phi, \psi \in \mathcal{L}$ be constraints such that $\Gamma \vdash \phi$ and $\Gamma \vdash \psi$. Given a Γ -valuation α with $\text{free}(\phi) \subseteq \text{dom}(\alpha)$, we write $\alpha \models \phi$ if α satisfies ϕ . We write $\models \phi$ if $\alpha \models \phi$ for all Γ -valuations α with $\text{free}(\phi) \subseteq \text{dom}(\alpha)$, and we write $\phi \models \psi$ if $\alpha \models \phi$ implies $\alpha \models \psi$ for all Γ -valuations α with $\text{free}(\phi) \cup \text{free}(\psi) \subseteq \text{dom}(\alpha)$. Entailment induces a theory $\mathcal{T} = \{\phi \mid \text{free}(\phi) = \emptyset \wedge \top \models \phi\}$, with respect to which entailment can be reduced to unsatisfiability. Note that unsatisfiability w. r. t. \mathcal{T} is not even semi-decidable as \mathcal{T} contains Peano arithmetic. Thus for reasoning purposes, we will generally approximate \mathcal{T} by weaker theories.

Effects. Let f be a built-in function with $\Pi(f) = \Gamma \rightarrow \Delta$ (viewing argument and return types of f as type environments Γ and Δ , respectively.) An *effect* for f is a pair of constraints ϕ and ψ such that $\Gamma \vdash \phi$ and $\Gamma, \Delta \vdash \psi$. (Note that $\Gamma \rightarrow \Delta$ being a function type implies $\text{dom}(\Gamma) \cap \text{dom}(\Delta) = \emptyset$, hence Γ, Δ is a type environment.) We write $\phi \rightarrow \psi$ to denote such an effect, and we call ϕ its *precondition* and ψ its *action*.

An *effect environment* maps the built-in functions $f \in \text{dom}(\Pi_0)$ to effects for f . Figure 4 displays the effect environment Θ_0 , providing an axiomatic, relational semantics for all $f \in \text{dom}(\Pi_0)$. This semantics ties most built-in functions to corresponding logical operators in a straightforward way; note the non-trivial preconditions for division, reading and writing arrays, and constructing singleton multisets. The effects of functions operating on resource managers warrant some explanation.

init returns an empty manager m' .

enable non-deterministically adds some sub-multiset of r to manager m , returning the result in manager m' ; the complement of the added multiset is returned in r' .

In an implementation [3] the multiset to be added to m would be chosen by some *policy*, perhaps involving security profiles or user input; we use non-determinism to abstractly model such policy mechanisms.

split splits the multiset held by manager m and distributes it to the managers m'_1 and m'_2 such that m'_2 gets the largest possible sub-multiset of r .

join adds the multisets held by managers m_1 and m_2 , returning their sum in m' .

consume is an explicit destructor for manager m and all its resources; the linear type system means that calls to **consume** are necessary even if m is known to be empty.

assertEmpty acts as identity on managers, but subject to the precondition that m is empty; it will be treated specially by the programming language semantics.

f	$\Pi_0(f)$	$\Theta_0(f)$
id_τ	$(x:\tau) \rightarrow (x':\tau)$	$\top \rightarrow x' \approx x$
eq_τ	$(x_1:\tau, x_2:\tau) \rightarrow (i':\text{int})$	$\top \rightarrow i' \approx 1 \wedge x_1 \approx x_2 \vee i' \approx 0 \wedge x_1 \not\approx x_2$
add sub mul div mod leq	$(i_1:\text{int}, i_2:\text{int}) \rightarrow (i':\text{int})$	$\top \rightarrow i' \approx i_1 + i_2$ $\top \rightarrow i' \approx i_1 + (-i_2)$ $\top \rightarrow i' \approx i_1 \cdot i_2$ $i_2 \neq 0 \rightarrow i' \approx i_1 / i_2$ $i_2 \neq 0 \rightarrow i' \approx i_1 \% i_2$ $\top \rightarrow i' \approx 1 \wedge i_1 \leq i_2 \vee i' \approx 0 \wedge i_1 \not\leq i_2$
conc	$(w_1:\text{str}, w_2:\text{str}) \rightarrow (w':\text{str})$	$\top \rightarrow w' \approx w_1 ++ w_2$
fromstr	$(w:\text{str}) \rightarrow (c':\text{res})$	$\top \rightarrow c' \approx \text{fromstr}(w)$
null $_\tau$	$() \rightarrow (a':\tau[])$	$\top \rightarrow a' \approx \text{null}$
length $_\tau$	$(a:\tau[]) \rightarrow (i':\text{int})$	$\top \rightarrow i' \approx \text{len}(a)$
read $_\tau$	$(a:\tau[], i:\text{int}) \rightarrow (x':\tau)$	$0 \leq i \wedge i < \text{len}(a) \rightarrow x' \approx a[i]$
write $_\tau$	$(a:\tau[], i:\text{int}, x:\tau) \rightarrow (a':\tau[])$	$0 \leq i \wedge i \leq \text{len}(a) \rightarrow a' \approx a[i:=x]$
empty	$() \rightarrow (r':\text{res}\{\})$	$\top \rightarrow r' \approx \emptyset$
single	$(c:\text{res}, i:\text{int}) \rightarrow (r':\text{res}\{\})$	$i \geq 0 \rightarrow r' \approx \{c:i\}$
inter union sum	$(r_1:\text{res}\{\}, r_2:\text{res}\{\}) \rightarrow (r':\text{res}\{\})$	$\top \rightarrow r' \approx r_1 \cap r_2$ $\top \rightarrow r' \approx r_1 \cup r_2$ $\top \rightarrow r' \approx r_1 \uplus r_2$
size	$(r:\text{res}\{\}) \rightarrow (i':\text{int})$	$\top \rightarrow i' \approx r $
count	$(r:\text{res}\{\}, c:\text{res}) \rightarrow (i':\text{int})$	$\top \rightarrow i' \approx \text{count}(r, c)$
include	$(r_1:\text{res}\{\}, r_2:\text{res}\{\}) \rightarrow (i':\text{int})$	$\top \rightarrow i' \approx 1 \wedge r_1 \subseteq r_2 \vee i' \approx 0 \wedge r_1 \not\subseteq r_2$
init	$() \rightarrow (m':\text{mgr})$	$\top \rightarrow m' \approx \emptyset$
enable	$(m:\text{mgr}, r:\text{res}\{\}) \rightarrow (m':\text{mgr}, r':\text{res}\{\})$	$\top \rightarrow r' \subseteq r \wedge m \uplus r \approx m' \uplus r'$
split	$(m:\text{mgr}, r:\text{res}\{\}) \rightarrow (m_1':\text{mgr}, m_2':\text{mgr})$	$\top \rightarrow m_2' \approx m \cap r \wedge m \approx m_1' \uplus m_2'$
join	$(m_1:\text{mgr}, m_2:\text{mgr}) \rightarrow (m':\text{mgr})$	$\top \rightarrow m' \approx m_1 \uplus m_2$
consume	$(m:\text{mgr}) \rightarrow ()$	$\top \rightarrow \top$
assertEmpty	$(m:\text{mgr}) \rightarrow (m':\text{mgr})$	$m \approx \emptyset \rightarrow m' \approx m$
assertAtLeast	$(m:\text{mgr}, r:\text{res}\{\}) \rightarrow (m':\text{mgr})$	$r \subseteq m \rightarrow m' \approx m$

Fig. 4. Types and effects of built-in functions. The subscripts τ indicate families of functions indexed by $\tau \in \langle \text{datatype} \rangle$, except for id_τ , which is indexed by $\tau \in \langle \text{type} \rangle$.

assertAtLeast acts as identity on managers, but subject to the precondition that the manager m contains the multiset r ; will be treated specially by the programming language semantics.

To facilitate the presentation of programming language semantics, we capture the logical semantics of effects directly in terms of valuations. Given a built-in function f with $\Pi_0(f) = \Gamma \rightarrow \Delta$ and $\Theta_0(f) = \phi \rightarrow \psi$, we define $\text{Eff}_{\Theta_0}^{\Pi_0}(f)$ to be the set of maximal (Γ, Δ) -valuations such that $\alpha \in \text{Eff}_{\Theta_0}^{\Pi_0}(f)$ if and only if $\alpha \models \phi \wedge \psi$.

2.3 Small-Step Reduction Semantics

We present a stack-based reduction semantics (which is essentially a continuation semantics) for our programming language. We will show that reduction preserves the

resources stored in resource managers, thanks to linearity. Throughout this section, let Π be a fixed well-typed program.

Stacks. We call a tuple $\langle x_1, \dots, x_n | \alpha, e \rangle$ a *frame* if x_1, \dots, x_n is a list of pairwise distinct variables, α is a valuation and e is an expression such that

- $\text{dom}(\alpha) \cap \{x_1, \dots, x_n\} = \emptyset$ and
- $\text{dom}(\alpha) \subseteq \text{free}(e) \subseteq \text{dom}(\alpha) \cup \{x_1, \dots, x_n\}$.

The roles of e (redex) and α (providing values for the free variables of e) should be clear. The x_i are only present if the frame is suspended waiting for a function to return in which case the x_i act as slots for the return values. A *pre-stack* is either $\frac{1}{2}$ or ϵ or $F :: S$, where F is a frame and S is a pre-stack. (Pre-stacks essentially correspond to continuations in an abstract machine interpreting λ -terms in ANF [10].) A *stack* (or Π -stack if we want to emphasise the program Π) is a pre-stack of the form $\frac{1}{2}$ or $\langle \alpha, e \rangle :: S$. We call $\frac{1}{2}$ the *error stack*. A stack of the form $\langle \alpha, \text{ret } (x_1, \dots, x_n) \rangle :: \epsilon$ is called *terminal*. If $F :: S$ is a stack then F is its *top frame*.

Reduction. Figure 5 presents the rules generating the reduction relation \rightsquigarrow_{Π} on stacks. We denote the reflexive-transitive closure of \rightsquigarrow_{Π} by \rightsquigarrow_{Π}^* . As usual Π may be omitted if it is understood. Note that reduction performs an eager garbage collection in that it deallocates unused variables immediately by restricting the valuation α in the post stack to the free variables of the expression e .

Reduction is deterministic, except for calls to the built-in function **enable**.

Proposition 2. *For all stacks S_0 there is at most one stack S_1 such that $S_0 \rightsquigarrow S_1$, unless S_0 is of the form $\langle \alpha, \text{let } (m', r') = \text{enable } (m, r) \text{ in } e \rangle :: S'_0$.*

Typed stacks. Reduction is untyped since type information is not needed at runtime. However, various properties of reduction are best stated if the type of variables is known. Therefore, we annotate stacks with type environments and conservatively extend reduction to typed stacks.

Given a frame $\langle x_1, \dots, x_n | \alpha, e \rangle$, we call $\langle x_1, \dots, x_n | \alpha, e \rangle^{\Gamma}$ a *typed frame* if Γ is a linear type environment such that

- $\text{dom}(\Gamma) = \text{dom}(\alpha) \cup \{x_1, \dots, x_n\}$,
- α is a Γ -valuation, and
- $\Gamma \vdash e : \sigma$ for some product type σ .

A *typed pre-stack* is $\frac{1}{2}$, or ϵ , or $F :: \epsilon$ where F is a typed frame, or $F :: F' :: S'$ where S' is a typed pre-stack and $F = \langle x_1, \dots, x_m | \alpha, e \rangle^{\Gamma}$ and $F' = \langle x'_1, \dots, x'_n | \alpha', e' \rangle^{\Gamma'}$ are typed frames such that $\Gamma \vdash e : (z'_1 : \Gamma'(x'_1), \dots, z'_n : \Gamma'(x'_n))$ for some variables z'_1, \dots, z'_n . A *typed stack* is typed pre-stack of the form $\frac{1}{2}$ or $\langle \alpha, e \rangle^{\Gamma} :: S$. Given a typed frame $F = \langle x_1, \dots, x_n | \alpha, e \rangle^{\Gamma}$, we denote its underlying frame $\langle x_1, \dots, x_n | \alpha, e \rangle$ by F^{\natural} . We extend this notation to typed (pre-)stacks, writing S^{\natural} for the (pre-)stack underlying the typed (pre-)stack S .

The following proposition shows that reduction does not break the invariants maintained by typed stacks.

(R-ret)	$\frac{\alpha'' = \alpha' \{x'_1 \mapsto \alpha(x_1), \dots, x'_n \mapsto \alpha(x_n)\}}{\langle \alpha, \mathbf{ret} (x_1, \dots, x_n) :: \langle x'_1, \dots, x'_n \alpha', e' \rangle :: S \rightsquigarrow \langle \alpha'' _{\text{free}(e')}, e' \rangle :: S}$
(R-let ₁ [†])	$\frac{\begin{array}{l} \Pi(f) = \lambda e : (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow \sigma' \\ \alpha' = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \end{array}}{\langle \alpha, \mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} \mathbf{ret} (x'_1, \dots, x'_n) :: S \rightsquigarrow \langle \alpha' _{\text{free}(e)}, e \rangle :: S}$
(R-let ₁)	$\frac{\begin{array}{l} \Pi(f) = \lambda e : (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow \sigma' \quad e' \neq \mathbf{ret} (x'_1, \dots, x'_n) \\ \alpha' = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \end{array}}{\begin{array}{l} \langle \alpha, \mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e' \rangle :: S \\ \rightsquigarrow \langle \alpha' _{\text{free}(e)}, e \rangle :: \langle x'_1, \dots, x'_n \alpha' _{\text{free}(e')}, e' \rangle :: S \end{array}}$
(R-let ₂)	$\frac{\begin{array}{l} \Pi_0(f) = (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow (z'_1:\tau'_1, \dots, z'_n:\tau'_n) \\ \alpha_f = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \quad \alpha'_f \in \text{Eff}_{\Theta_0}^{\Pi_0}(f) \quad \alpha'_f \succeq \alpha_f \\ \alpha' = \alpha \{x'_1 \mapsto \alpha'_f(z'_1), \dots, x'_n \mapsto \alpha'_f(z'_n)\} \end{array}}{\langle \alpha, \mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e' \rangle :: S \rightsquigarrow \langle \alpha' _{\text{free}(e')}, e' \rangle :: S}$
(R-let ₂ [‡])	$\frac{\begin{array}{l} \Pi_0(f) = (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow \sigma' \quad f \in \{\mathbf{assertEmpty}, \mathbf{assertAtLeast}\} \\ \alpha_f = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \quad \forall \alpha'_f \in \text{Eff}_{\Theta_0}^{\Pi_0}(f) : \alpha'_f \not\preceq \alpha_f \end{array}}{\langle \alpha, \mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e' \rangle :: S \rightsquigarrow \zeta}$
(R-if ₁)	$\frac{\alpha(v) \neq 0}{\langle \alpha, \mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2 \rangle :: S \rightsquigarrow \langle \alpha _{\text{free}(e_1)}, e_1 \rangle :: S}$
(R-if ₂)	$\frac{\alpha(v) = 0}{\langle \alpha, \mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2 \rangle :: S \rightsquigarrow \langle \alpha _{\text{free}(e_2)}, e_2 \rangle :: S}$

Fig. 5. Small-step reduction relation \rightsquigarrow (for a fixed program Π). Application of valuations α extends to values $v \in \langle \text{val} \rangle$ in the natural way, i. e., $\alpha(v) = v$ if v is a constant.

Proposition 3. *Let \hat{S}_0 be a typed stack and S_1 a stack. If $\hat{S}_0^{\text{h}} \rightsquigarrow S_1$ then there is a typed stack \hat{S}_1 such that $\hat{S}_1^{\text{h}} = S_1$.*

The proposition justifies the view of reduction on typed stacks as a conservative extension of the reduction relation defined in Figure 5, where reduction on typed stacks is defined by $\hat{S}_0 \rightsquigarrow_{\Pi} \hat{S}_1$ if and only if $\hat{S}_0^{\text{h}} \rightsquigarrow_{\Pi} \hat{S}_1^{\text{h}}$; as usual Π may be omitted if it is understood.

We call a stack S_0 *stuck* if there is no stack S_1 such that $S_0 \rightsquigarrow S_1$, and S_0 is neither terminal nor the error stack. Our next result shows that reduction on typed stacks will get stuck only at calls to built-in functions (other than **assertEmpty** and **assertAtLeast**), and only if the preconditions of these calls fail. As the effects listed in Figure 4 reveal, reduction will get stuck only upon attempts to divide by 0, access arrays out of bounds or construct singleton multisets with negative multiplicity.

Proposition 4. *Let \hat{S} be a typed stack. If \hat{S}^{h} is stuck then it is of the form*

$$\langle \alpha, \mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e' \rangle :: S',$$

$f \in \text{dom}(\Pi_0) \setminus \{\mathbf{assertEmpty}, \mathbf{assertAtLeast}\}$, and there is no $\alpha'_f \in \text{Eff}_{\Theta_0}^{\Pi_0}(f)$ such that $\alpha'_f \succeq \alpha_f$, where α_f is defined as in rule (R-let₂).

Preservation of resources. Given a typed frame $F = \langle x_1, \dots, x_n | \alpha, e \rangle^F$, we define the multiset $\text{res}(F)$ of resources in F by $\text{res}(F) = \bigsqcup \{\alpha(x) \mid x \in \text{dom}(\alpha), \Gamma(x) = \mathbf{mgr}\}$. We extend res to typed non-error stacks by defining $\text{res}(\epsilon) = \emptyset$ and $\text{res}(F :: S) = \text{res}(F) \uplus \text{res}(S)$. Proposition 5 states *resource preservation*: The sum of all resources in the system remains unchanged by reduction, unless the built-in functions **enable** and **consume** are called. The former admits increasing (but not decreasing) the resources, whereas the latter behaves the other way round. Obviously, resource preservation depends on the linearity restriction on type **mgr**, otherwise resources could be duplicated by re-using managers.

Proposition 5. *Let S_0 and S_1 be typed stacks such that $S_0 \rightsquigarrow S_1 \neq \downarrow$.*

1. *If S_0 is of the form $\langle | \alpha, \mathbf{let} (m', r') = \mathbf{enable} (m, r) \mathbf{in} e \rangle^F :: S'_0$ then $\text{res}(S_0) \subseteq \text{res}(S_1)$.*
2. *If S_0 is of the form $\langle | \alpha, \mathbf{let} () = \mathbf{consume} (m) \mathbf{in} e \rangle^F :: S'_0$ then $\text{res}(S_0) \supseteq \text{res}(S_1)$.*
3. *In all other cases, $\text{res}(S_0) = \text{res}(S_1)$.*

2.4 Erasing Resource Managers

According to the reduction semantics, a call to **assertEmpty** or **assertAtLeast** either does nothing¹ or goes wrong, and calling one of these two tests is the only way to go wrong. Hence, if we know that a program cannot go wrong (and Section 3 will present a type system for proving just that) then we can erase all calls to these built-ins (or rather, replace them by true no-ops) and obtain an equivalent program.

In fact, we can do more than that. Once the assertion built-ins are gone, it is even possible to remove the resource managers themselves. By the design of the programming language (in particular, the choice of built-in operations on resource managers) the contents of resource managers cannot influence the values of variables of any other type. Informally, this justifies replacing the resource managers themselves by variables of type **unit** whenever we know that a program cannot go wrong. Erasing resource managers also means that the built-in functions acting on managers can be replaced by simpler ones on **unit**: all of which are no-ops, except for **enable** itself.² The remainder of the section formalises this intuition.

Figure 6 shows the necessary program transformations to erase resource managers. Most fundamentally, erasure maps the manager type **mgr** to the unit type **unit**. Erasure on types determines erasure on product types, type environments, programs and valuations (where erasure uniformly maps the values of **mgr**-variables to \star , the only value of type **unit**), which in turn determines erasure on typed stacks. As outlined

¹ Due to the linearity restriction on resource managers these functions must copy the input manager to an output manager; a true no-op would violate resource preservation.

² We do keep the calls in place, so that erasure preserves the structure of programs; this simplifies reasoning, and does not preclude optimising away no-op calls at a later stage.

<p>Erasure τ° of types τ</p> $\tau^\circ = \mathbf{unit} \quad \text{if } \tau = \mathbf{mgr}$ $\tau^\circ = \tau \quad \text{otherwise}$	<p>Erasure Γ° of type environments Γ</p> $\emptyset^\circ = \emptyset$ $(\Gamma, x:\tau)^\circ = \Gamma^\circ, x:\tau^\circ$
<p>Erasure σ° of product types σ</p> $(x_1:\tau_1, \dots, x_n:\tau_n)^\circ = (x_1:\tau_1^\circ, \dots, x_n:\tau_n^\circ)$	
<p>Erasure Π° of programs Π</p> $\text{dom}(\Pi^\circ) = \text{dom}(\Pi)$ $\Pi^\circ(f) = \lambda e : \sigma^\circ \rightarrow \sigma'^\circ \quad \text{if } \Pi(f) = \lambda e : \sigma \rightarrow \sigma'$ $\Pi^\circ(f) = \sigma^\circ \rightarrow \sigma'^\circ \quad \text{if } \Pi(f) = \sigma \rightarrow \sigma'$	
<p>Erasure Θ_0° of effect environment Θ_0</p> $\text{dom}(\Theta_0^\circ) = \text{dom}(\Theta_0)$ $\Theta_0^\circ(\mathbf{enable}) = \top \rightarrow r' \subseteq r$ $\Theta_0^\circ(f) = \top \rightarrow \top \quad \text{if } \left\{ \begin{array}{l} f \in \{\mathbf{init}, \mathbf{split}, \mathbf{join}, \mathbf{consume}\} \cup \\ \quad \{\mathbf{assertEmpty}, \mathbf{assertAtLeast}\} \end{array} \right.$ $\Theta_0^\circ(f) = \Theta_0(f) \quad \text{otherwise}$	
<p>Erasure α° of Γ-valuations α</p> $\text{dom}(\alpha^\circ) = \text{dom}(\alpha)$ $\alpha^\circ(x) = \star \quad \text{if } \Gamma(x) = \mathbf{mgr}$ $\alpha^\circ(x) = \alpha(x) \quad \text{otherwise}$	
<p>Erasure S° of typed stacks S</p> $\frac{1}{2}^\circ = \frac{1}{2} \quad \epsilon^\circ = \epsilon \quad (\langle x_1, \dots, x_n \alpha, e \rangle^\Gamma :: S)^\circ = \langle x_1, \dots, x_n \alpha^\circ, e \rangle^{\Gamma^\circ} :: S^\circ$	

Fig. 6. Erasure of resource managers

above, erasure on effect environments trivialises the effect of resource manager built-ins, except **enable**, and preserves the effects of all built-ins not operating on managers. The effect of **enable** after erasure is to non-deterministically choose a sub-multiset of r and return its complement in r' . This reflects the fact that calls to **enable** provide points of interaction for the policy (e.g., the user) to decide how many resources the system is granted. Erasing resource managers does not mean that policy decisions are fixed, it just removes the managers' book keeping about those decisions.

Lemma 6. *Let Π be a well-typed program and S a typed Π -stack. Then Π° is a well-typed program and S° a typed Π° -stack.*

Erasure makes trivial the effects of **assertEmpty** and **assertAtLeast**, and in particular, replaces their precondition by \top . Thus a program cannot go wrong after erasure, as rule (R-let $_{\frac{1}{2}}$) will never apply.

Proposition 7. *Let Π be a well-typed program and S a Π° -stack S . Then $S \not\rightsquigarrow_{\Pi^\circ}^* \frac{1}{2}$.*

The next result states that the small-step reduction relation \rightsquigarrow_{Π} of a program Π is almost bisimulation equivalent to the reduction relation $\rightsquigarrow_{\Pi^\circ}$ of its erasure. In fact,

it shows that the relation $R = \{\langle S, S^\circ \rangle \mid S \text{ is a } \Pi\text{-stack}\}$ would be a bisimulation if \rightsquigarrow_Π could not reduce stacks to the error stack \perp . Put differently, if Π cannot go wrong then \rightsquigarrow_Π and $\rightsquigarrow_{\Pi^\circ}$ are bisimulation equivalent. The proof of this theorem is by case analysis on the reduction relation \rightsquigarrow_Π of the unerased program. As a corollary, we get that reachability in the erased program is essentially the same as reachability in the unerased one, provided that the unerased program cannot go wrong.

Theorem 8. *Let Π be a well-typed program and \hat{S}_0 a typed Π -stack with $\hat{S}_0 \not\rightsquigarrow_\Pi \perp$.*

1. *For all typed Π -stacks \hat{S}_1 , if $\hat{S}_0 \rightsquigarrow_\Pi \hat{S}_1$ then $\hat{S}_0^\circ \rightsquigarrow_{\Pi^\circ} \hat{S}_1^\circ$.*
2. *For all typed Π° -stacks S_1 , if $\hat{S}_0^\circ \rightsquigarrow_{\Pi^\circ} S_1$ then there is a typed Π -stack \hat{S}_1 such that $\hat{S}_0 \rightsquigarrow_\Pi \hat{S}_1$ and $\hat{S}_1^\circ = S_1$.*

Corollary 9. *Let Π be a well-typed program and S_0 a typed Π -stack. If $S_0 \not\rightsquigarrow_\Pi^* \perp$ then $\{S^\circ \mid S_0 \rightsquigarrow_\Pi^* S\} = \{S \mid S_0^\circ \rightsquigarrow_{\Pi^\circ}^* S\}$.*

What distinguishes erasure of resource managers from other erasure results (e.g., type erasure during compilation, Java generics erasure) is that here, erasure does not completely remove a language construct. Instead, it removes the book keeping but retains the semantically important bit that deals with dynamic policy decisions.

2.5 Big-Step Relational Semantics

The reduction semantics presented in Section 2.3 is good for showing preservation properties, like the preservation of resources. However, it does not easily yield a relational view on functions, relating input and output parameters. This is achieved by a relational semantics, which we will prove equivalent to the reduction semantics. Contrary to the reduction semantics, which was originally untyped and had type environments added conservatively, the relational semantics will be typed from the start. (Types do not hurt here, as the relational semantics is not geared towards execution.)

Throughout this section, we assume that Π is a well-typed program. A *state* β is either the error state \perp or a normal state $\langle \Gamma; \alpha \rangle$, where Γ is a linear type environment and α a maximal Γ -valuation. Given an expression e , a normal state $\langle \Gamma; \alpha \rangle$ and a state β' , we define the judgement $e, \langle \Gamma; \alpha \rangle \Downarrow_\Pi \beta'$ (or $e, \langle \Gamma; \alpha \rangle \Downarrow \beta'$ if Π is understood) by the rules in Figure 7 if $\text{dom}(\Gamma) \cap \text{bound}(e) = \emptyset$ and there are Γ_e and σ such that $\Gamma \succeq \Gamma_e$ and $\Gamma_e \vdash e : \sigma$. The intended meaning of $e, \langle \Gamma; \alpha \rangle \Downarrow \beta'$ is that evaluating expression e in state $\langle \Gamma; \alpha \rangle$ may terminate and result in state β' .

The reduction semantics deallocates variables once they become unused (an eager garbage collection, so to say), which is essential for the linear variables as otherwise resource preservation would not hold. However, the intermediate values of variables are thus lost. In contrast, the relational semantics names and records all intermediate values, even the linear ones, as $e, \langle \Gamma; \alpha \rangle \Downarrow \langle \Gamma'; \alpha' \rangle$ implies $\Gamma' \succeq \Gamma$ and $\alpha' \succeq \alpha$.

By definition, violations of resource safety manifest themselves in reductions ending in the error stack, and hence reductions which diverge or get stuck cannot violate resource safety. Therefore, resource safety is not affected by the fact that the relational semantics ignores such reductions. Under this proviso, Proposition 10 shows the equivalence of reduction and relational semantics.

Evaluation of expressions $e, \langle \Gamma; \alpha \rangle \Downarrow \beta'$	
(E-ret)	$\frac{}{\mathbf{ret} (x_1, \dots, x_n), \langle \Gamma; \alpha \rangle \Downarrow \langle \Gamma; \alpha \rangle}$
(E-let ₁)	$\frac{\begin{array}{l} \Pi(f) = \lambda e : (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow (z'_1:\tau'_1, \dots, z'_n:\tau'_n) \quad \Gamma_f = z_1:\tau_1, \dots, z_m:\tau_m \\ \alpha_f = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \quad e, \langle \Gamma_f; \alpha_f \rangle \Downarrow \langle \Gamma'_f; \alpha'_f \rangle \\ \Gamma' = \Gamma, x'_1:\tau'_1, \dots, x'_n:\tau'_n \quad \alpha' = \alpha\{x'_1 \mapsto \alpha'_f(z'_1), \dots, x'_n \mapsto \alpha'_f(z'_n)\} \\ e', \langle \Gamma'; \alpha' \rangle \Downarrow \beta'' \end{array}}{\mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e', \langle \Gamma; \alpha \rangle \Downarrow \beta''}$
(E-let ₁ [‡])	$\frac{\begin{array}{l} \Pi(f) = \lambda e : (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow \sigma' \quad \Gamma_f = z_1:\tau_1, \dots, z_m:\tau_m \\ \alpha_f = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \quad e, \langle \Gamma_f; \alpha_f \rangle \Downarrow \ddagger \end{array}}{\mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e', \langle \Gamma; \alpha \rangle \Downarrow \ddagger}$
(E-let ₂)	$\frac{\begin{array}{l} \Pi(f) = (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow (z'_1:\tau'_1, \dots, z'_n:\tau'_n) \\ \alpha_f = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \quad \alpha'_f \in \mathit{Eff}_{\emptyset_0}^{\Pi_0}(f) \quad \alpha'_f \succeq \alpha_f \\ \Gamma' = \Gamma, x'_1:\tau'_1, \dots, x'_n:\tau'_n \quad \alpha' = \alpha\{x'_1 \mapsto \alpha'_f(z'_1), \dots, x'_n \mapsto \alpha'_f(z'_n)\} \\ e', \langle \Gamma'; \alpha' \rangle \Downarrow \beta'' \end{array}}{\mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e', \langle \Gamma; \alpha \rangle \Downarrow \beta''}$
(E-let ₂ [‡])	$\frac{\begin{array}{l} \Pi(f) = \lambda e : (z_1:\tau_1, \dots, z_m:\tau_m) \rightarrow \sigma' \quad f \in \{\mathbf{assertEmpty}, \mathbf{assertAtLeast}\} \\ \alpha_f = \{z_1 \mapsto \alpha(v_1), \dots, z_m \mapsto \alpha(v_m)\} \quad \forall \alpha'_f \in \mathit{Eff}_{\emptyset_0}^{\Pi_0}(f) : \alpha'_f \not\preceq \alpha_f \end{array}}{\mathbf{let} (x'_1, \dots, x'_n) = f (v_1, \dots, v_m) \mathbf{in} e', \langle \Gamma; \alpha \rangle \Downarrow \ddagger}$
(E-if ₁)	$\frac{e_1, \langle \Gamma; \alpha \rangle \Downarrow \beta'}{\mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2, \langle \Gamma; \alpha \rangle \Downarrow \beta'} \text{ if } \alpha(v) \neq 0$
(E-if ₂)	$\frac{e_2, \langle \Gamma; \alpha \rangle \Downarrow \beta'}{\mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2, \langle \Gamma; \alpha \rangle \Downarrow \beta'} \text{ if } \alpha(v) = 0$

Fig. 7. Big-step evaluation relation (for a fixed program Π)

Proposition 10. *Let $\langle \Gamma; \alpha \rangle$ and $\langle \Gamma'; \alpha' \rangle$ be states. Let e be an expression such that $\text{dom}(\Gamma) = \text{free}(e)$ and $\Gamma \vdash e : \sigma$ for some product type σ . Then*

1. $e, \langle \Gamma; \alpha \rangle \Downarrow \ddagger$ if and only if $\langle \alpha, e \rangle^\Gamma :: \epsilon \rightsquigarrow^* \ddagger$, and
2. $e, \langle \Gamma; \alpha \rangle \Downarrow \langle \Gamma'; \alpha' \rangle$ if and only if there is a typed stack $\langle \alpha'', \mathbf{ret} (x_1, \dots, x_n) \rangle^{\Gamma''} :: \epsilon$ such that $\langle \alpha, e \rangle^\Gamma :: \epsilon \rightsquigarrow^* \langle \alpha'', \mathbf{ret} (x_1, \dots, x_n) \rangle^{\Gamma''} :: \epsilon$ and $\Gamma' \succeq \Gamma''$ and $\alpha' \succeq \alpha''$.

3 Effect Type System

In this section, we will develop a type system to statically guarantee dynamic resource safety, i. e., the absence of reductions to the error stack \ddagger . We will do so by annotating functions with effects and then extending the notion of effect to a judgement on expressions, which we will define by a simple set of typing rules.

3.1 Effect Type System

We extend the notion of effect $\phi \rightarrow \psi$ from built-in functions to λ -abstractions. To be precise, $\phi \rightarrow \psi$ is an *effect* for f if $\Gamma \vdash \phi$ and $\Gamma, \Delta \vdash \psi$, where $\Pi(f) = [\lambda \dots] \Gamma \rightarrow \Delta$, regardless of whether f is built-in or a λ -abstraction. In line with this extension, an *effect environment* Θ maps all functions $f \in \text{dom}(\Pi)$ to effects $\Theta(f)$ for f .

In order to derive the effects of λ -abstractions, we generalise effects to effect types for expressions and develop a type system for inductively constructing such effect types. Effects relate input and output parameters of functions by logical formulae. Likewise, effect types shall relate input and output parameters of expressions. Here, the input parameters of an expression are its free variables; the output parameters are those variables that are not free yet but will become free during reduction, i. e., the (let-)bound variables. Formally, an *effect type* $\Gamma; \phi \rightarrow \Delta; \psi$ is a pair of constraints ϕ and ψ together with a pair of type environments Γ and Δ such that $\text{dom}(\Gamma) \cap \text{dom}(\Delta) = \emptyset$ and $\Gamma \vdash \phi$ and $\Gamma, \Delta \vdash \psi$. We call ϕ and ψ *precondition* and *action*, and Γ and Δ *input* and *output (parameters)*, respectively. Given an expression e , we say that an effect type $\Gamma; \phi \rightarrow \Delta; \psi$ is an *effect type for* e if $\text{dom}(\Gamma) \cap \text{bound}(e) = \emptyset$.

We say that an effect type $\Gamma; \phi \rightarrow \Delta; \psi$ is *stronger than* an effect type $\Gamma'; \phi' \rightarrow \Delta'; \psi'$, denoted by $\Gamma; \phi \rightarrow \Delta; \psi \supseteq \Gamma'; \phi' \rightarrow \Delta'; \psi'$, if $\phi' \models \phi$ and $(\phi' \wedge \psi) \models \psi'$, i. e., the stronger effect type $\Gamma; \phi \rightarrow \Delta; \psi$ has a weaker precondition but stronger action. The stronger-than relation \supseteq is a quasi-order, i. e., reflexive and transitive, and induces an equivalence relation on effect types, the *as-strong-as* relation, which we denote by \equiv . Note that for every effect type $\Gamma; \phi \rightarrow \Delta; \psi$ is as strong as an effect type $\Gamma'; \phi \rightarrow \Delta'; \psi$ with linear type environments Γ' and Δ' .

Figure 8 presents the typing rules for deriving effect types. There, the judgement $\Theta \vdash_{\Pi} e : \Gamma; \phi \rightarrow \Delta; \psi$ states that expression e has effect type $\Gamma; \phi \rightarrow \Delta; \psi$ in the context of program Π and effect environment Θ . If Π is understood, we may omit it and write $\Theta \vdash e : \Gamma; \phi \rightarrow \Delta; \psi$ instead. The judgement $\Pi, \Theta \vdash f$ means that the effect type ascribed to a λ -abstraction f by Θ and Π is consistent with the effect type derived for the body of f . We say that Θ is an *admissible* effect environment for a program Π if $\Pi, \Theta \vdash f$ for all λ -abstractions $f \in \text{dom}(\Pi) \setminus \text{dom}(\Pi_0)$.

Lemma 11. *Let e be an expression, Θ an effect environment (referring to an implicit program Π) and $\Gamma; \phi \rightarrow \Delta; \psi$ an effect type. If $\Theta \vdash e : \Gamma; \phi \rightarrow \Delta; \psi$ then $\Gamma; \phi \rightarrow \Delta; \psi$ is an effect type for e .*

Theorem 12 states soundness of effect typing w. r. t. the big-step relational semantics. The proof is by double induction on the derivation of relational semantics judgements over the derivation of effect type judgements. As a corollary, we get that reduction starting from a state that satisfies the precondition can't go wrong, hence resource managers can be erased. In fact, the untyped reductions in the erased program match exactly the typed reductions in the original program.

Theorem 12. *Let Θ be an admissible effect environment for a well-typed program Π . Let e be an expression and $\Gamma; \phi \rightarrow \Delta; \psi$ an effect type such that $\Theta \vdash e : \Gamma; \phi \rightarrow \Delta; \psi$. Let $\langle \Gamma; \alpha \rangle$ and β' be states such that $e, \langle \Gamma; \alpha \rangle \Downarrow \beta'$ (which implies $\Gamma_e \vdash e : \sigma$ for some Γ_e, σ). If $\alpha \models \phi$ then $\beta' = \langle \Gamma'; \alpha' \rangle$ for some Γ' and α' such that $\alpha' \models \phi \wedge \psi$. (In particular, if $\alpha \models \phi$ then $\beta' \neq \perp$.)*

Typing of expression effects $\Theta \vdash e : \Gamma; \phi \rightarrow \Delta; \psi$	
(ET-weak)	$\frac{\Theta \vdash e : \Gamma; \phi \rightarrow \Delta; \psi}{\Theta \vdash e : \Gamma'; \phi' \rightarrow \Delta'; \psi'} \text{ if } \begin{cases} \text{dom}(\Gamma') \cap \text{bound}(e) = \emptyset \wedge \\ \Gamma; \phi \rightarrow \Delta; \psi \supseteq \Gamma'; \phi' \rightarrow \Delta'; \psi' \end{cases}$
(ET-ret)	$\frac{}{\Theta \vdash \mathbf{ret} (x_1, \dots, x_n) : \emptyset; \top \rightarrow \emptyset; \top}$
(ET-if)	$\frac{\Theta \vdash e_1 : \Gamma; v \not\approx 0 \wedge \phi \rightarrow \Delta; \psi \quad \Theta \vdash e_2 : \Gamma; v \approx 0 \wedge \phi \rightarrow \Delta; \psi}{\Theta \vdash \mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2 : \Gamma; \phi \rightarrow \Delta; \psi}$
(ET-let)	$\frac{\begin{array}{l} \Pi(f) = [\lambda \dots] \Gamma \rightarrow \Delta \quad \Gamma = z_1:\tau_1, \dots, z_m:\tau_m \quad \Delta = z'_1:\tau'_1, \dots, z'_n:\tau'_n \\ \Theta(f) = \phi \rightarrow \psi \quad \mu = \{z_1 \mapsto v_1, \dots, z_m \mapsto v_m, z'_1 \mapsto x'_1, \dots, z'_n \mapsto x'_n\} \\ \Theta \vdash e' : \Gamma', \Delta'; \phi' \wedge \psi' \rightarrow \Delta''; \psi'' \end{array}}{\Theta \vdash \mathbf{let} (x'_1, \dots, x'_n) = f(v_1, \dots, v_m) \mathbf{in} e' : \Gamma'; \phi' \rightarrow \Delta', \Delta''; \psi' \wedge \psi''} \text{ if } (*)$ <div style="text-align: right; margin-right: 20px;">where (*) $\begin{cases} \text{dom}(\Gamma') \cap \{x'_1, \dots, x'_n\} = \emptyset \wedge \\ \Gamma\mu; \phi\mu \rightarrow \Delta\mu; \psi\mu \supseteq \Gamma'; \phi' \rightarrow \Delta'; \psi' \end{cases}$</div>
Well-typedness of λ-abstraction effects $\Pi, \Theta \vdash f$	
(ET-lam)	$\frac{\Pi(f) = \lambda e : \Gamma \rightarrow \Delta \quad \Theta(f) = \phi \rightarrow \psi \quad \Theta \vdash e : \Gamma; \phi \rightarrow \Delta; \psi}{\Pi, \Theta \vdash f}$

Fig. 8. Typing rules for effect types (for a fixed program Π)

Corollary 13. *Let Θ be an admissible effect environment for a well-typed program Π . Let e be an expression and $\Gamma; \phi \rightarrow \Delta; \psi$ an effect type such that $\Theta \vdash_{\Pi} e : \Gamma; \phi \rightarrow \Delta; \psi$. Let α be a maximal Γ -valuation, and let $\hat{S}_0 = \langle |\alpha|_{\text{free}(e)}, e \rangle^{\Gamma}_{\text{free}(e)} :: \epsilon$ be a typed Π -stack (which implies $\Gamma|_{\text{free}(e)} \vdash_{\Pi} e : \sigma$ for some σ). If $\alpha \models \phi$ then*

1. $\hat{S}_0 \not\rightsquigarrow_{\Pi}^* \hat{\zeta}$ and
2. for all (untyped) Π° -stacks S , $\hat{S}_0^{\circ\ddagger} \rightsquigarrow_{\Pi^{\circ}}^* S$ if and only if there is a typed Π -stack \hat{S} such that $\hat{S}_0 \rightsquigarrow_{\Pi}^* \hat{S}$ and $\hat{S}^{\circ\ddagger} = S$. (In particular, $\hat{S}_0^{\circ\ddagger} \not\rightsquigarrow_{\Pi^{\circ}}^* \hat{\zeta}$.)

3.2 Example: Bulk Messaging Application

To illustrate the use of the effect type system, we revisit the example from Figure 1. The interesting bits of code are in the functions `send_bulk` and `send_msg`.

The function `send_bulk` first builds up a multiset of resources r by converting the strings representing phone numbers in `nums` into resources. Next it attempts to authorise the use of all resources by having `enable` add r to an empty resource manager m . If this fails, i. e., the multiset r' returned by `enable` is of non-zero size, `send_bulk` terminates (after destroying m' and whatever resources it holds).³ If authorising all resources succeeds, `send_bulk` calls `send_msgs` to actually send the messages while checking

³ A more sophisticated version of the application could deal more gracefully with `enable` granting only part of the requested resources. This would require more complex code to inspect the multisets r and r' (but not the resource manager m').

f	$\Theta(f)$
<code>send_bulk</code>	$\top \rightarrow \top$
<code>res_from_nums</code>	$\top \rightarrow r \approx \text{bagof}(\text{map}_{\text{fromstr}}(\text{nums}))$
<code>res_from_nums'</code>	$0 \leq i \leq \text{len}(\text{nums}) \wedge r' \approx \text{bagof}(\text{map}_{\text{fromstr}}(\text{subarray}(\text{nums}, i, \text{len}(\text{nums}))))$ $\rightarrow r \approx \text{bagof}(\text{map}_{\text{fromstr}}(\text{nums}))$
<code>send_msgs</code>	$\text{bagof}(\text{map}_{\text{fromstr}}(\text{nums})) \subseteq m \rightarrow m \approx m' \uplus \text{bagof}(\text{map}_{\text{fromstr}}(\text{nums}))$
<code>send_msgs'</code>	$0 \leq i \leq \text{len}(\text{nums}) \wedge \text{bagof}(\text{map}_{\text{fromstr}}(\text{subarray}(\text{nums}, 0, i))) \subseteq m$ $\rightarrow m \approx m' \uplus \text{bagof}(\text{map}_{\text{fromstr}}(\text{subarray}(\text{nums}, 0, i)))$
<code>send_msg</code>	$\text{count}(m, \text{fromstr}(\text{num})) \geq 1 \rightarrow m \approx m' \uplus \{\text{fromstr}(\text{num}):1\}$
<code>prim_send_msg</code>	$\top \rightarrow \top$

$\forall a : \text{len}(\text{map}_{\text{fromstr}}(a)) \approx \text{len}(a)$
$\forall a \forall i : 0 \leq i < \text{len}(a) \Rightarrow \text{map}_{\text{fromstr}}(a)[i] \approx \text{fromstr}(a[i])$
$\forall a \forall j \forall k : 0 \leq j < k \leq \text{len}(a) \Rightarrow \text{len}(\text{subarray}(a, j, k)) = k - j$
$\forall a \forall j \forall k \forall i : 0 \leq j < k \leq \text{len}(a) \wedge 0 \leq i < \text{len}(\text{subarray}(a, j, k)) \Rightarrow \text{subarray}(a, j, k)[i] = a[j + i]$
$\forall a : \text{bagof}(a) \approx \text{len}(a)$
$\forall a : \text{len}(a) \approx 1 \Rightarrow \text{bagof}(a) \approx \{a[0]:1\}$
$\forall a \forall k : 0 \leq k \leq \text{len}(a) \Rightarrow \text{bagof}(a) \approx \text{bagof}(\text{subarray}(a, 0, k)) \uplus \text{bagof}(\text{subarray}(a, k, \text{len}(a)))$

Fig. 9. Bulk messaging application: admissible effect environment Θ and axiomatisation of theory extension; for the sake of readability sort information is suppressed in the axioms

that the manager m' contains the required resources. After that, `send_bulk` checks that `send_msgs` has used up all resources by asserting that the returned manager m' is empty; failing this assertion will trigger a runtime error. Finally, `send_bulk` explicitly destroys the empty manager m' and terminates.

The function `send_msg` sends one message, checking whether the resource manager m holds the resource required. It does so by converting the string `num` into a singleton multiset of resources r . Then it splits the manager m into m' and m_r , so that m_r contains at most the resources in r . Next, `send_msg` asserts that m_r contains at least r ; failing this assertion will trigger a runtime error. Succeeding the assertion, `send_msg` calls the primitive send function, destroys the now used resource by consuming m_r , and returns the remaining resources in the manager m' .

The bulk messaging example is statically resource safe, as witnessed by the admissible effect environment displayed in Figure 9. Of particular interest is the effect $\top \rightarrow \top$ ascribed to the main function `send_bulk`. This least informative effect expresses nothing about the function itself but implies the absence of runtime errors via Corollary 13.

The effects require an extension of the theory \mathcal{T} (see Section 2.2) by three new functions, axiomatised in Figure 9. The function `map` maps an array of strings to an array of resources, `subarray` takes an array and cuts out the sub-array between two given indices, and `bagof` converts an array of resources to a multiset (containing the same elements with the same multiplicity). Note that the axiomatisation of `bagof` is not complete⁴ but sufficient for our purposes.

Effect type checking, e. g., for checking admissibility of the effect environment Θ from Figure 9, requires checking the side condition of the weakening rule (ET-weak),

⁴ A complete axiomatisation of `bagof` is possible in the full first-order theory of multisets and arrays but it is much more complicated and unusable in practise.

which involves checking logical entailment w. r. t. to an extension of the theory \mathcal{T} . Due to the high undecidability of \mathcal{T} , we actually check entailment w. r. t. (an extension of) an approximation of \mathcal{T} ; in particular, we approximate multiplication and division by uninterpreted functions. For the bulk messaging example, we used an SMT solver [4] that can handle linear integer arithmetic and arrays. We added axioms for multisets and the axioms in Figure 9. Due to an incomplete quantifier instantiation heuristic, we had to instantiate a number of these axioms by hand, yet eventually, the solver was able to prove all the entailments required by the weakening rules.

Even though arising from a single example, we believe that the extension of the theories of multisets and arrays with the functions *subarray* and *bagof* is quite generic and could prove useful in many cases.

4 Conclusion

We have presented a programming language with support for complex resource management, close to the standard SSA/ANF forms of compiler intermediate languages [1]. By construction, programs are *dynamically resource safe* in that any attempts to abuse resources are trapped. We have extended the language with an effect type system which guarantees the for well-typed programs no such attempts occur: we have *static resource safety*. In addition, for such programs the bookkeeping required by dynamic resource management can be erased.

Related Work. Many tools and methods have been proposed to assist with resource management at runtime, e.g., in Java, the JRes [9] and J-Seal [8] frameworks. Generally, these aim to enable programs to react to fluctuations of resources caused by an unpredictable environment. Our aim, however is to track the flow of resources through the program, where the environment can influence the availability of resources only at well-understood points of interaction with the program and with clear availability policies. This offers the chance for more precise resource control whose behaviour can be predicted statically.

This paper builds on previous work [3] with a Java library implementing resource managers and focusing on the dynamic aspects of resource management policies. This Java library supports essentially the same operations on resource managers as our functional language, except that state is realised by destructive updates instead of linear types. While [3] does not provide a static analysis to prove static resource safety, it does outline how dynamic accounting could be erased if static resource safety were provable. Our work here shows one way to do just that.

Our approach is in line with a general trend of providing the programmer with language-based mechanisms for security and additional static analyses (often using type systems) which use these mechanisms. This combination provides a desirable graceful degradation: if static analysis succeeds in proving certain properties, then the program may be optimised without affecting security. Yet, even if the analyses fail the language based mechanisms will enforce the security properties at runtime.

The context of our work is the MOBIUS project [5] on proof-carrying code (PCC) for mobile devices. Our effect type system is very simple and in principle well-suited for a PCC setting where checkers themselves are resource bounded. However, the weakening

rule relies on checking logical entailment in a first-order theory, which is undecidable in general. Therefore, a certificate for PCC need not only provide a type derivation tree but also proofs (in some proof system) for the entailment checks in the weakening rule. The development of a suitable such proof system is a topic for further research, as is the investigation of decidable fragments of relevant first-order theories.

Acknowledgements. This work was funded in part by the Sixth Framework programme of the European Community under the MOBIUS project FP6-015905. This paper reflects only the authors' views and the European Community is not liable for any use that may be made of the information contained therein. Ian Stark was also supported by an Advanced Research Fellowship from the UK Engineering and Physical Sciences Research Council, EPSRC project GR/R76950/01.

References

- [1] Appel, A.W.: SSA is functional programming. SIGPLAN Notices 33(4), 17–20 (1998)
- [2] Aspinall, D., Beringer, L., Hofmann, M., Loidl, H.-W., Momigliano, A.: A program logic for resources. Theoret. Comput. Sci. 389(3), 411–445 (2007)
- [3] Aspinall, D., Maier, P., Stark, I.: Monitoring external resources in Java MIDP. Electron. Notes Theor. Comput. Sci. 197, 17–30 (2008)
- [4] Barrett, C., de Moura, L., Stump, A.: Design and results of the 2nd annual satisfiability modulo theories competition. Form. Meth. Syst. Des. 31(3), 221–239 (2007)
- [5] Barthe, G., Beringer, L., Crégut, P., Grégoire, B., Hofmann, M., Müller, P., Poll, E., Puebla, G., Stark, I., Vétillard, E.: MOBIUS: Mobility, ubiquity, security. In: Montanari, U., Sannella, D., Bruni, R. (eds.) TGC 2006. LNCS, vol. 4661, pp. 10–29. Springer, Heidelberg (2007)
- [6] Beringer, L., Hofmann, M., Momigliano, A., Shkaravska, O.: Automatic certification of heap consumption. In: Baader, F., Voronkov, A. (eds.) LPAR 2004. LNCS, vol. 3452, pp. 347–362. Springer, Heidelberg (2005)
- [7] Besson, F., Dufay, G., Jensen, T.P.: A formal model of access control for mobile interactive devices. In: Gollmann, D., Meier, J., Sabelfeld, A. (eds.) ESORICS 2006. LNCS, vol. 4189, pp. 110–126. Springer, Heidelberg (2006)
- [8] Binder, W., Hulaas, J., Villazón, A.: Portable resource control in Java. In: Proc. OOPSLA 2001, pp. 139–155. ACM, New York (2001)
- [9] Czajkowski, G., von Eicken, T.: JRes: A resource accounting interface for Java. In: Proc. OOPSLA 1998, pp. 21–35. ACM, New York (1998)
- [10] Flanagan, C., Sabry, A., Duba, B.F., Felleisen, M.: The essence of compiling with continuations. In: Proc. PLDI 1993, pp. 237–247. ACM, New York (1993)
- [11] Nanevski, A., Ahmed, A., Morrisett, G., Birkedal, L.: Abstract predicates and mutable ADTs in Hoare type theory. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 189–204. Springer, Heidelberg (2007)
- [12] Unknown: Redbrowser. A, J2ME trojan. Identified in February 2006 as Redbrowser. A (F-Secure), J2ME/Redbrowser.a (McAfee), Trojan. Redbrowser. A (Symantec), Trojan-SMS.J2ME.Redbrowser.a (Kaspersky Lab)