

Typed Memory Management in a Calculus of Capabilities

C Karl Crary

D
Cornell University

Greg Morrisett
Cornell University

A

A
g
c
p
provide an alternative to garbage collection by making mem-
o
not been clear how to use regions in low-level, type-safe code.
W

C
m
straightforward to compile to a typed assembly language.
Source languages be compiled to our language using
k
times need not be lexically sco6(-) d in our language, yet the

l
y
P T

s
m
p
a

1

A
t
l
i

p
g

region types and lifetimes and how to implement their the-
dis (this calculus was first formalized in 1993 [12]) [24] [35] [66] [68] [50] [31] [9] [32] [5] [6] [12] [3] [10] [2] [13] [15] [84] [19] [40] [25] [13] [6] [20] [16] [15] [10] [27] [10] [19] [24] [15] [11] [10] [2] [76]
operations and in principle could be formally veri ed.

The Tofte-Talpin calculus uses a lexically scoped expres-
sion 0 0 (ang) 280 (d) 18 (e) 11 (lim) 9 (it) -260 (t) 7 (h) 18 (e) -257 (life) 11 (9 (anim) 9 (e) -257 (o) 14 (f) -266 (a) -253 (r) 9 (e) 11 (-)] TJfl -12.982e-1.097 TDfl -0.005e) 1 [(gi) -14 (on)] TJfl /F11 1 Tffl 2.

Both Birkedal *et al.* [4] and Aiken *et al.* [1] observe that the use of memory in many cases. They proposed a series of improvements upon traditional tracing garbage collection in some cases. Although their optimizations are safe, there is no simple proof of correctness that an untrusted client can be safely integrated with the TdWte-Talpn framework and allow arbitrary separation of allocation and deallocation points. Therefore, we have taken a first step towards our goal by constructing a strongly-typed region-based assembly language in which we must ensure that any live memory is eventually deallocated. It will no longer be accessed. Operating systems such as Hydra [41] have solved the access control problem by separating object and requiring that the object is not accessed in the future.

2 Overview of Contributions

In the rest of this paper we describe a strongly typed language called the Capability Calculus. Our language's type system provides an efficient way to check the safety of explicit, arbitrarily ordered region allocation and deallocation instructions using a notion of capability. As in traditional capability systems, our type system keeps track of capabilities. The capabilities in our calculus are a purely static concept and thus their implementation requires no run-time overhead. We have a purely

kinds $::= \text{Type} / \text{Rgn} / \text{Cap}$
constructor vars $::=$
constructors $c ::= j \text{ } j r j C$
types $::= j \text{int} j r \text{ handle } j \&[] : (C_1 \text{ } \dots \text{ } C_n) ! 0 \text{atr} j h_1 \text{ } \dots \text{ } h_i \text{atr}$
regions $r ::= j$
capabilities $C ::= j j f \text{ } r \text{ } g j C_1 \text{ } C_2 j C$
multiplicities $' ::= 1 \text{ } j +$

constructor contexts $::= j \text{ } ; \text{ } : j \text{ } ; \text{ } \text{TD } 0 \text{f} 277 \text{ } T c \text{ } (j) T j \text{ } / F 16 \text{ } 1 \text{ } T f \text{ } 1.124 \text{ } 0 \text{ } T D \text{ } (\text{ }) T j \text{ } / F 10$

memory types $::= f \text{ } \begin{matrix} 1 : 1 ; \dots ; n : n g \\ 1 : 1 ; \dots ; n : n g \end{matrix}$

heap values $v ::= x j i j : ' j \text{ handle } (\text{ }) j v [c]$
 $h ::= \text{fix } f [] (C x_1 : 1 ; \dots ; h x_n) : e j h_1 ; \dots ; h_n v$

 $1 \text{ } p \text{ } v_2 \text{ } j \text{ } x = j \text{ } \text{De } (\text{at}) j \text{fl } / \text{Per } 4 \text{ } T \text{fl } 0.83 \text{ } 0 \text{ } T D \text{fl } 0.01 \text{ } T c \text{fl } (\text{at}) T j \text{fl } / F 10 \text{ } 1 \text{ } T \text{fl } 1$

memory regions $R ::= f \text{ } \begin{matrix} 1 \text{ } \text{ } h \\ 1 \text{ } \text{ } R_1 ; \dots ; n \text{ } \text{ } R_n g \end{matrix}$

 $M = f \text{ } \begin{matrix} 1 \text{ } \text{ } R_1 ; \dots ; n \text{ } \text{ } R_n g \end{matrix}$

 $P = (M$

which states that (when memory has type Σ , free constructor variables have kinds given by \mathcal{K} and free value variables have types given by \mathcal{T}) it is legal to execute the term e , *provided that the capability C is held*. A related typing judgement is

$$\vdash e : \tau ; \Sigma ; \mathcal{K} ; C \vdash d : \tau^0 ; \Sigma^0 ; C^0$$

which states that if the capability C

Uniqueness

We solve this problem by using bounded quantification to relate τ_1 , τ_2 and τ . Suppose h has type:

$$\mathcal{B}[\tau_1 : \text{Rgn}; \tau_2 : \text{Rgn}; f : \tau_1^+ ; \tau_2^+] \rightarrow \tau \text{ at } r^{\infty}$$

If we hold capability $f r^1 g$, we may call h by instantiating τ_1 and τ_2 with r and instantiating f with $f r$. This instantiation is permissible because $f r^1 g \rightarrow f r^+ ; r g$. As with g , the continuation will possess the capability $r g$, allowing it

$$\tau_1 \text{ and } \tau_2, \text{ since } \tau_1^+ \rightarrow \tau_2^+ g.$$

Bounded quantification solves the problem by revealing some information about a capability τ , while still requiring the function to be parametric over τ . Hence, when the function calls its continuation we regain the stronger capability (to r), although that capability was temporarily hidden in order to duplicate τ . Generally, bounded quantification allows us to hide some privileges when calling a function, and regain those privileges in its continuation. The most important properties of the Capability Calculus are type soundness and complete collection. Each can be proven from the formal semantics in Appendix A. Thus, we support statically checkable attenuation and amplification of capabilities.

) then $(M^0; e^0)$ is not stuck.

The proof of soundness is straightforward, making use

level operations such as the atomic allocation and initialization of data. In the companion technical report [5], we show that the capability constructs interact benignly with the process of type-preserving compilation described by Morrisett *et al.* [24] and we use the techniques described in this paper to modify their typed assembly language to allow explicit deallocation but with a few changes we believe our capability apparatus may be used in a variety of other settings as well. One potential application involves communication across the user-kernel address space boundary in traditional operating systems. Typically, in such systems, when data in user space is presented to the kernel, the kernel must copy that data to ensure its integrity is preserved. However, if a user process has a unique capability for a region to the kernel, the kernel does not have to copy that region's data.

Kernel may alias other regions. Second, after we dynamically check which of the capabilities are valid, we can use the typecase mechanism developed for the TIL compiler [13] to produce type correct code in the capability Calculus. Gay and Aiken [9] have developed a safe region im-

A connection can also be drawn between capabilities and monadic type systems. Work relating effects to monads [21, 28, 18, 8] has viewed effectful functions as pure functions that return transformers. This might be called

but may alias other regions. Second, after we dynamically check which of the capabilities are valid, we can use the typecase mechanism developed for the TIL compiler [13] to produce type correct code in the capability Calculus. Gay and Aiken [9] have developed a safe region im-

to produce type correct code in the capability Calculus.

Gay and Aiken [9] have developed a safe region im-

6 Acknowledgements

We would like to thank Lars Birkedal, Martin Elsman, Dan Grossman, Chris Hawblitzel, Fred Smith, Stephanie Weirich, Steve Zdancewic, and the anonymous reviewers for their comments and suggestions.

References

- [1] Alexander Aiken, Manuel F.

*ACM SIG-
PLAN Conference on Programming Language Design
and Implementation*

- [27] George Necula and Peter Lee. The design and implementation of a certifying compiler. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 333 { 344, Montreal, June 1998.
- [28] Simon L. Peyton Jones and Philip Wadler. Imperative functional programming. In *Twentieth ACM Symposium on Principles of Programming Languages*, Charleston, South Carolina, January 1993.
- [29] John C. Reynolds. Definitional interpreters for higher-order programming languages. In *Conference Record of the 25th Annual ACM Conference*, pages 717{740, Boston, August 1972.
- [30] John C. Reynolds. Syntactic control of interference. In *Fifth ACM Symposium on Principles of Programming Languages*, pages 39{46, Tucson, Arizona, 1978.
- [31] John C. Reynolds. Syntactic control of interference in analysis problems in languages with destructive updating. *ACM Transactions on Programming Languages and Systems*, 20(1):1{50, January 1996.
- [33] J.-P. Talpin and P. Jouvelot. Polymorphic type, region, and effect inference. *Journal of Functional Programming*, 2(1):168-197, 1992.
- [38] Philip Wadler. A taste of linear logic. In *Mathematical Foundations of Computer Science*, volume 711 of *LNCS*, Gdansk, Poland, August 1993. Springer-Verlag.
- [39] Robert Wahbe, Steven Lucco, Thomas Anderson, and Susan Graham. Efficient software-based fault isolation. In *Fourteenth ACM Symposium on Operating Systems Principles*, pages 203{216, Asheville, December 1993.
- [40] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38{94, 1994.
- [41] W. A. Wulf, R. Levin, and S. P. Harbison. *Hydra/C.mmp: An Experimental Computer System*. McGraw-Hill, New York, NY, 1981.

$(M; e) \not\models P$	
if $e =$	then $P =$
$\text{let } x = v \text{ in } e'$	$(M; e'[v=x])$
$\text{let } x = i p j \text{ in } e'$	$(M; e')$
$\text{let } x = \text{hat}(\text{handle}($	$($
$\text{and } \neq \text{Dom}(M)$	where \neq
$\text{let } x = i(:) \text{ in } e'$	$(M; e' = x])$
	$(Mf$
	where $\neq M$ and $\neq e$
$\text{let freern}(\$	$(Mn ; e)$
$\text{if } 0 \text{ then } e \text{ else } e$	$(M; e)$
$\text{if } 0 \text{ then } e \text{ else } e$	

aratin d is well-fdC-4(9)-9(m)-9(e)-7(d)-321(a)-4(nd)-375(pr)-9(d)-30(duc)-7(e)-7(s)-38

Expression e is well-fdrmed.

θ

θ

θ

$(\theta \in \text{Dom}(\theta))$

θ

$C : \text{Cap}$

θ

θ

θ

; ; ' hat r:

$$\frac{\begin{array}{c} \text{; } ; \text{' } Tm6 : Type \quad (\text{for } 1 \leq i \leq n) \\ \text{; } ; \text{' } fix f[\text{' } TC : x_1.T06 :: \dots x_n.Tn6 \end{array}}{\text{; } ; \text{' } Tm6 : Type \quad (\text{for } 1 \leq i \leq n)} \quad \text{; } ; \text{' } Tm6 : Type \quad (\text{for } 1 \leq i \leq n)$$
$$\frac{\begin{array}{c} ; \quad ; \quad ' \quad v_i : \text{T064} \text{ (for } 1 \leq i \leq n) \quad ' \quad r : \text{Rgn} \\ ; \quad ; \quad ' \quad \text{hat } r : \text{T064} \end{array}}{\begin{array}{c} ; \quad ; \quad ' \quad h \, v_1 : \dots : h \, v_n : \text{T064} \quad ' \quad i : \text{Rgn} \\ ; \quad ; \quad ' \quad \text{hat } i : \text{T064} \end{array}} \quad \text{Type}$$

```

; ; ' v : T064

```

• •

/ /

' X

