



# Memory Safety in Systems Languages


Major Area Exam

---

Michael Christensen

June 11, 2018

## Committee:

Ben Hardekopf (  )    Tim Sherwood    Rich Wolski

# What is a System?

Infrastructure software upon which applications are built

Operating Systems

- Process abstraction
- Multiplex physical hardware resources
- Partition and abstract **memory**

Embedded Systems, Compilers, Garbage Collectors, Device Drivers, File Systems

# Systems Languages

- Past systems languages: ALGOL, PL/I, Fortran, BCPL/B, C, Mesa/Cedar, Pascal/Modula-2/Oberon, C++, ...
- C: the de-facto standard
  - Data structure representation control
  - Memory management control
  - Complete mutability via pointers
  - Performant
  - Legacy
- C: the unsafe standard
  - Unchecked array operations
  - Pointers  $\equiv$  arrays
  - Unsafe casts
  - Aliasing
  - **Undefined behavior**

# Memory Safety

- Memory safety error: reads or writes outside the referent's storage
  - Spatial: outside referent's **address bounds**
  - Temporal: outside referent's **lifetime**
- Ideal technique is
  - Efficient and expressive
  - Purely static
  - Precise
  - Automatic
- Memory errors become **type errors**, management happens at **compile-time**


# Spatial Safety

Always access within object's **bounds**

Some approaches:

- Fat Pointers and Shadow Structures
- Referent Objects
- Dependent Types

```
struct { int *arr; int i; } s;  
s.arr = calloc(3, sizeof(int));
```



## Spatial Safety Example

```
1 int find_token(int *data,
2               int *end,
3               int token)
4 {
5     int *p = data;
6     while (p < end) {
7         if (*p == token) break;
8         p++;
9     }
10    return (*p == token);
11 }
```

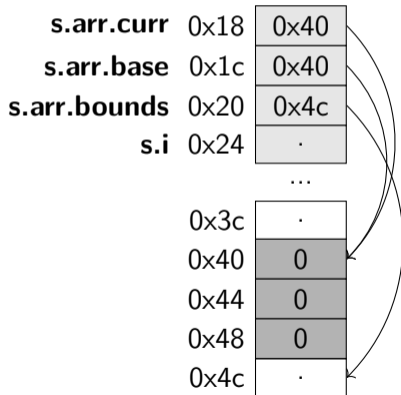
Potential pointer dereference problems:

- Null
- Uninitialized
- Out-of-bounds
- Manufactured

## Fat Pointers

- Added base and bound addresses
- Instrument all pointers and functions
- Insert runtime checks **before dereferences**
- Advantages:
  - Quickly find and retrieve metadata
  - Complete spatial safety
  - No sub-object overflows
- Disadvantages:
  - Breaks binary compatibility
  - Metadata propagation
  - Code bloat, memory usage, runtime overhead
  - Unsafe casts overwriting metadata

```
struct fptr { int *curr;  
             int *base; int *bound; };  
struct { struct fptr arr; int i; } s;  
s.arr.curr = calloc(3, sizeof(int));
```



# Fat Pointer Approaches

SafeC<sup>1</sup>:

- **Safe** pointers have value, base, and size
- Complete spatial safety, if **transparent** storage management and no safe pointer attribute **manipulation**
- 275% space overhead, 2-6x runtime overhead, 0.35-3x code size overhead
- Some static optimization based on still-valid previous checks

Cyclone<sup>2</sup>:

- Annotations for non-array vs array pointers (can specify size)
- Tagged unions and automatic tag injection

---

<sup>1</sup>Austin, Breach, and Sohi, "Efficient Detection of All Pointer and Array Access Errors", 1994.

<sup>2</sup>Jim et al., "Cyclone: A Safe Dialect of C.", 2002.



# Fat Pointer Approaches

## CCured<sup>3</sup>

- Separate pointers on usage (SAFE, SEQ, WILD)
- Whole-program inference to find as many SAFE then SEQ pointers as possible
- Reduce WILD pointers<sup>4</sup> using physical subtyping<sup>5</sup> for upcasts
- Special pointer RTTI carrying runtime type for downcasts

## Fail-Safe C<sup>6</sup>:

- Combines fat pointers w/ fat integers and virtual structure offsets

---

<sup>3</sup>Necula, McPeak, and Weimer, "CCured", 2002.

<sup>4</sup>Necula, Condit, et al., "CCured", 2005.

<sup>5</sup>Siff et al., "Coping with Type Casts in C", 1999.

<sup>6</sup>Oiwa, "Implementation of the Memory-safe Full ANSI-C Compiler", 2009.

## Fat Pointers Preventing Spatial Errors

```
1 int find_token(int *data,
2               int *end,
3               int token)
4 {
5     int *p = data;
6     while (p < end) {
7
8         if (*p == token) break;
9
10        p++;
11    }
12
13    return (*p == token);
14 }
```

```
1 int find_token(int *SEQ data,
2               int *SAFE end,
3               int token)
4 {
5     int *SEQ p = data;
6     while (p.cur < end) {
7         assert(p.base != 0 &&
8                0 <= p.cur &&
9                p.cur < p.bound);
10        if (*p.cur == token) break;
11        p.cur = p.cur + (1 * sizeof(int));
12        p.base = p.base; // optimized out
13        p.bound = p.bound; // " "
14    }
15    ...(repeat lines 7-9)...
16    return (*p.cur == token);
17 }
```

# Pointer-Based – Shadow Structures

MSCC<sup>7</sup>

- *Split* metadata from pointer, preserving layout
- Every value has linked shadow structure mirroring entire data structure
- Transform every function call to take additional metadata parameters
- Wrappers for external functions; cannot detect memory errors

---

<sup>7</sup>Xu, DuVarney, and Sekar, “An Efficient and Backwards-compatible Transformation to Ensure Memory Safety of C Programs”, 2004.

## Shadow Structures Example

```
1
2
3
4
5 int find_token(
6     int *data,
7     int *end,
8     int token)
9 {
10     int *p = data;
11
12     while (p < end) {
13         if (*p == token) break;
14         p++;
15     }
16
17     return (*p == token);
18 }
19 }
```

```
1 struct ptr_info {
2     void *base;
3     unsigned long bound;
4 };
5 int find_token(
6     int *data, ptr_info *data_info,
7     int *end, ptr_info *end_info,
8     int token)
9 {
10     int *p = data;
11     ptr_info p_info = *data_info;
12     while (p < end) {
13         CHECK_SPATIAL(p, sizeof(*p), p_info);
14         if (*p == token) break;
15         p++;
16     }
17     CHECK_SPATIAL(p, sizeof(*p), p_info);
18     return (*p == token);
19 }
```

# Referent Objects

## Objects<sup>8,9</sup>

- Metadata about **objects**, not pointers
- Global database relates every allocated address to corresponding object metadata
- Every pointer to same object shares same metadata
- Bounds check on **pointer arithmetic**
- 2-12x overhead

### Advantages:

- Compatible with uninstrumented code

### Disadvantages:

- Special mechanisms to handle legal OOB pointers
- Splay-tree object lookup overhead
- Incomplete spatial safety: **sub-object overflows**

---

<sup>8</sup>Jones and Kelly, "Backwards-compatible bounds checking for arrays and pointers in C programs", 1997.

<sup>9</sup>Ruwase and Lam, "A Practical Dynamic Buffer Overflow Detector", 2004.

## The Sub-Object Problem

```
1 struct node {char str[3]; void (*func)(); };  
2 struct node *n = (struct node *) malloc(sizeof(node));  
3 char *s = n.str;  
4 strcpy(s, "bad!");
```

- `n` and `s` have the same address  $\Rightarrow$  map to same object in global database
- `strcpy` will see `s`'s size as that of `n`

# Referent Objects Approaches

## SafeCode<sup>10</sup>

- Use automatic pool allocation (APA)<sup>11</sup>
- Use separate, **smaller** data structures to store bounds metadata for **each partition**
- 1.2x overhead

## Baggy Bounds Checking (BBC)<sup>12</sup>

- Compact bounds representation and efficient way to look up object bounds
- Align base addresses to be multiple of padded size
- Replace splay tree with small lookup table
- 0.6x overhead

---

<sup>10</sup>Dhurjati and Adve, "Backwards-compatible Array Bounds Checking for C with Very Low Overhead", 2006.

<sup>11</sup>Lattner and Adve, "Automatic Pool Allocation", 2005.

<sup>12</sup>Akritidis et al., "Baggy Bounds Checking", 2009.

## Referent Objects Example

```
1 int find_token(int *data,
2               int *end,
3               int token)
4 {
5     int *p = data;
6     while (p < end) {
7         if (*p == token) break;
8
9
10
11
12         p++;
13     }
14     return (*p == token);
15 }
```

```
1 int find_token(int *data,
2               int *end,
3               int token)
4 {
5     int *p = data;
6     while (p < end) {
7         if (*p == token) break;
8         int *q = p + 1;
9         int size = 1 << TABLE[p]>>4];
10        int base = p & ~(size - 1);
11        assert(q >= base && q - base < size);
12        p++;
13    }
14    return (*p == token);
15 }
```



# Softbound

## Softbound<sup>13</sup>

- Base and bound metadata for each pointer, stored in disjoint metadata table
- Total spatial safety of **pointer-based** approaches
- Source compatibility, separate compilation of **object-based** approaches
- Runtime bounds checks on each dereference
- Propagate metadata as extra arguments
- Arbitrary casts allowed
- 67% overhead

---

<sup>13</sup>Nagarakatte, Zhao, Milo MK Martin, et al., "SoftBound", 2009.

# Dependent Types

Dependent types are *typed-valued functions*<sup>14</sup>

```
Vector : Nat → Type → Type
  nil : Vector 0 a
cons : Πn:Nat.a → Vector n a → Vector (n+1) a
  (cons 'a' (cons 'b' nil)) : Vector 2 Char
  head : Πn:Nat.Vector (n+1) a → a
  head nil ⇒ Rejected!
```

- Based on type theory work by Martin-Löf<sup>15</sup>
- **Undecidability** of type checking: arbitrary computation to check type equality
- Work on defining equality and restricting forms of index terms

---

<sup>14</sup>Pierce, *Advanced topics in types and programming languages*, 2005.

<sup>15</sup>Martin-Löf, "Constructive mathematics and computer programming", 1984.

## Early Uses of Dependent Types

### Dependent ML<sup>16</sup> and Cayenne<sup>17</sup>

- Reduce static array bound checking to constraint satisfiability
- DML uses *indexed types*: limit indices to linear integer and boolean expressions; compile-time decidable
- Cayenne has *no restrictions* on types: undecidability of arbitrary expression equivalence and thus type checking

### Xanadu<sup>18</sup>

- Imperative environment
- Restrict index expressions in types to integer constraint domain

---

<sup>16</sup>Xi and Pfenning, “Eliminating Array Bound Checking Through Dependent Types”, 1998.

<sup>17</sup>Augustsson, “Cayenne—a Language with Dependent Types”, 1998.

<sup>18</sup>Xi, “Imperative programming with dependent types”, 2000.

# Dependent Types in Imperative Languages

SafeDrive<sup>19</sup> and Deputy<sup>20,21</sup>

- User-added annotations relating pointers to bounds
  - `safe`, `sentinel`, `count(n)`, `bound(lo,hi)`
  - Use constants/variables/field names in immediately enclosing scope
- Three-phase pass over annotated C programs, emits C code
  1. Automatic addition of bounds annotations for pointer types
  2. Flow-insensitive type checking (insert run-time checks; helps decidability)
  3. Flow-sensitive check optimization
- More C Support with dependent union tags , Safe TinyOS<sup>22</sup>

---

<sup>19</sup>Zhou et al., "SafeDrive", 2006.

<sup>20</sup>Condit et al., "Dependent Types for Low-Level Programming", 2007.

<sup>21</sup>Anderson, "Static Analysis of C for Hybrid Type Checking", 2007.

<sup>22</sup>Cooprider et al., "Efficient Memory Safety for TinyOS", 2007.

## Deputy Example for Spatial Safety

```
1 int find_token(int *data,
2               int *end,
3               int token)
4 {
5
6     int *p = data;
7     while (p < end) {
8
9         if (*p == token) break;
10
11         p++;
12     }
13
14     return (*p == token);
15 }
16
17
18 }
```

```
1 int find_token(int * bound(data, end) data,
2               int * sentinel end,
3               int token)
4 {
5     assert(data != NULL);
6     assert(end != NULL);
7     int *p bound(p, end) = data;
8     while (p < end) {
9         assert(p != NULL);
10        assert(p < end);
11        if (*p == token) break;
12        assert(p <= p + 1 <= end);
13        p++;
14    }
15    assert(p != NULL);
16    assert(p < end);
17    return (*p == token);
18 }
```

## Abstract Syntax, For Your Consideration

$x, y \in \text{Variables}$     $op \in \text{Binary ops}$     $n \in \text{Integers}$     $comp \in \text{Comparison Ops}$

*Ctors*    $C ::= \text{int} \mid \text{ref} \mid \text{array}$

*Types*    $\tau ::= C \mid \tau_1 \tau_2 \mid \tau e$

*L-exprs*    $l ::= x \mid *e$

*Exprs*    $e ::= n \mid l \mid e_1 \text{ op } e_2$

*Cmds*    $c ::= l := e \mid \text{assert}(\gamma) \mid c_1; c_2 \mid \dots$

*Preds*    $\gamma ::= e_1 \text{ comp } e_2 \mid \text{true} \mid \gamma_1 \wedge \gamma_2$

## Typing Rules, For Your Consideration

Local Expressions:  $\Gamma \vdash_L e : \tau$

$$\frac{\Gamma(x) = \tau}{\Gamma \vdash_L x : \tau} \text{ (LOCAL NAME)} \quad \frac{}{\Gamma \vdash_L n : \text{int}} \text{ (LOCAL NUM)}$$

Non-local Expressions:  $\Gamma \vdash e : \tau \Rightarrow \gamma$

$$\frac{\Gamma \vdash e : \text{ref } \tau \Rightarrow \gamma}{\Gamma \vdash *e : \tau \Rightarrow \gamma} \text{ (DEREF)}$$

Commands:  $\Gamma \vdash c \Rightarrow c'$

$$\frac{x \in \text{Dom}(\Gamma) \quad \forall (y : \tau_y) \in \Gamma, \Gamma \vdash y[e/x] : \tau_y[e/x] \Rightarrow \gamma_y}{\Gamma \vdash x := e \Rightarrow \text{assert}(\bigwedge_{y \in \text{Dom}(\Gamma)} \gamma_y); x := e} \text{ (VAR WRITE)}$$

# The Interesting Rules

Dereferencing

$$\frac{\Gamma \vdash e : \text{array } \tau \ e_{len} \Rightarrow \gamma_e}{\Gamma \vdash *e; \tau \Rightarrow \gamma_e \wedge (0 < e_{len})} \text{ (ARRAY Deref)}$$

Arithmetic

$$\frac{\Gamma \vdash e : \text{array } \tau \ e_{len} \Rightarrow \gamma_e \quad \Gamma \vdash e' : \text{int} \Rightarrow \gamma_{e'}}{\Gamma \vdash e + e' : \text{array } \tau \ (e_{len} - e') \Rightarrow \gamma_e \wedge \gamma'_{e'} \wedge (0 \leq e' \leq e_{len})} \text{ (ARRAY ARITH)}$$



# Dependent Types in Imperative Languages

## Týr<sup>23</sup>

- Augments LLVM IR with dependent pointer types
- Uses programmer annotations insert run-time bounds checks
- LLVM optimizations remove always-true checks; error if always-false

## Checked C<sup>24</sup>

- Extend C with two *checked pointer types*: `_Ptr<T>` and `_Array_ptr<T>`
- Associated bounds expressions indicating where bounds are stored
- Isolate (un)safe code with *checked code regions*

## Low\*<sup>25</sup>

- DSL for verified, efficient low-level programming in F\*
- Write F\* syntax against library modelling lower-level view of C memory

---

<sup>23</sup>De Araújo, Moreira, and Machado, "Týr", 2016.

<sup>24</sup>Ruef et al., "Checked C for Safety, Gradually", 2017.

<sup>25</sup>Protzenko et al., "Verified Low-level Programming Embedded in F\*", 2017.

# Quick Spatial Recap

## Spatial Safety

- Arrays and pointers
- Fat pointers
- Referent objects
- Dependent types

# Temporal Safety

Prevent accessing object that has been **previously deallocated**

- Capabilities and locks
- Effects and regions
- Linear types and ownership

# Temporal Safety Example

```
1 int attach(struct sock *sk) {
2     if (sk->bad) {
3         free(sk); return 1;
4     }
5     return 0;
6 }
7 void mq_notify(sigevent *n) {
8     struct sock_t *sock;
9     while (n->try) {
10        sock = malloc_sock(n->info);
11        if (attach(sock)){
12            //sock = NULL;
13            break;
14        }
15    }
16    if (sock) free(sock);
17 }
```

Potential pointer dereference problems:

- Double frees
- Dangling pointers

A Real Bug

- Linux Kernel in `ipc/mqueue.c`
- July 2017
- [https://bugzilla.redhat.com/show\\_bug.cgi?id=1470659](https://bugzilla.redhat.com/show_bug.cgi?id=1470659)

Goals:

- Good: Detecting use-after-free
- Better: Eliminating free entirely

# A Comment on Garbage Collection

## Garbage collection

- Relinquish control of object location and layout to runtime
- Complete temporal safety, but...
  - Non-zero overhead
  - Drag
  - Loss of real-time guarantees/predictability
  - Reduced reference locality, increased page fault/cache miss rates
- Some spatial approaches (e.g. Fail-Safe C, CCured) use Boehm-Demers-Weister<sup>26</sup>

---

<sup>26</sup>Boehm and Weiser, "Garbage Collection in an Uncooperative Environment", 1988.

# Capabilities and Locks

- SafeC, MSCC
  - Unique capability associated with each memory block
  - Stored in capability store, marked invalid on free
  - Check if pointer's capability copy is still valid on dereference
- CETS<sup>27,28</sup>
  - Each allocation has unique (never reused) key and lock address
  - Freeing allocated object changes value at lock location, so key and lock value don't match
- Memsafe<sup>29</sup>
  - Set bounds of deallocated pointer to invalid value

---

<sup>27</sup>Nagarakatte, Zhao, Milo M.K. Martin, et al., "CETS", 2010.

<sup>28</sup>Nagarakatte, M. M. K. Martin, and Zdancewic, "Everything You Want to Know About Pointer-Based Checking", 2015.

<sup>29</sup>Simpson and Barua, "MemSafe", 2013.

# Lock Example

```
1 int attach(struct sock *sk,
2     key_t sk_key,
3     lock_t *sk_lock_addr) {
4     if (sk_key != *sk_lock_addr)
5         abort();
6     if (sk->bad) {
7         if (Freeable_ptrs_map.lookup(sk_key) != sk)
8             abort();
9         free(sk);
10        *sk_lock_addr = INVALID_KEY;
11        deallocate_lock(sk_lock_addr);
12        return 1;
13    }
14    return 0;
15 }
```

```
1 void mq_notify(sigevent *n) {
2     struct sock_t *sock;
3     key_t sock_key;
4     lock_t *sock_lock_addr;
5     while (n->try) {
6         sock = malloc_sock(n->info);
7         sock_key = Next_key++;
8         sock_lock_addr = allocate_lock();
9         *(sock_lock_addr) = sock_key;
10        Freeable_ptrs_map.insert(sock_key, sock);
11        if (attach(sock)){
12            break;
13        }
14    }
15    if (sock) {
16        if (Freeable_ptrs_map.lookup(sock_key) != sock)
17            abort();
18        free(sock);
19        *(sock_lock_addr) = INVALID_KEY;
20        deallocate_lock(sock_lock_addr);
21    }
22 }
```

# Effects

Effect types describe the *effects* of the computation leading to a value<sup>30</sup>

- Opening a file
- Modifying an object

---

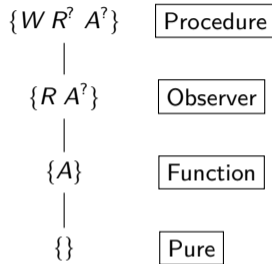
<sup>30</sup>Pierce, *Advanced topics in types and programming languages*, 2005.



# Effects

Fluent Languages,<sup>31</sup> MFX<sup>32</sup>

- Mix functional and imperative languages
- Every expression has a **effect class**, restricting which subroutines or sublanguage features it may use
- Identify **referentially transparent** expressions via a side effect lattice
- Effect masking: inference to delimit regions of memory and their lifetimes



E.g.

- update is Procedure
- nth is Observer
- arrayCreate is Function
- length is Pure

---

<sup>31</sup>David K. Gifford and John M. Lucassen, "Integrating functional and imperative programming", 1986.

<sup>32</sup>J. M. Lucassen and D. K. Gifford, "Polymorphic Effect Systems", 1988.

# Effects

## Type and Effect Systems<sup>33</sup>

- Extend the simply-typed lambda calculus with annotations
- Various forms of analyses, incl. Side Effect Analysis and **Region** Inference

Exprs  $e ::= c \mid x \mid \text{fn}_{\pi} x \Rightarrow e \mid e_1 e_2 \mid \text{new}_{\pi} x := e_1 \text{ in } e_2 \mid !x \mid x := e$

Types  $\tau ::= \text{int} \mid \text{bool} \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \xrightarrow{\phi} \tau_2 \mid \tau \text{ ref} \mid \tau \text{ ref } \varrho$

Effect **Region**  $\phi ::= \{\!|\pi|\!\} \mid \{\pi :=\} \mid \{\text{new } \pi\} \mid \phi_1 \cup \phi_2 \mid \emptyset$

PrgPts  $\varrho ::= \{\pi\} \mid \varrho_1 \cup \varrho_2 \mid \emptyset$

$$\Gamma \vdash e : \tau \ \& \ \phi$$

---

<sup>33</sup>F. Nielson and H. R. Nielson, "Type and Effect Systems", 1999.

# Region-Based Memory Management

## Regions<sup>34</sup>

- Divide heap into stack of sub-heaps (i.e. regions)
- Regions grow on individual allocation; **entire** region deallocated
- Explicit region polymorphism
- Sound type system **guarantees safety** of deallocations
- Region inference identifies:
  - Points where entire regions are allocated and deallocated
  - Into which region values should go
- Unreasonable object lifetimes due to LIFO ordering of region lifetimes

---

<sup>34</sup>Tofte and Talpin, "Region-Based Memory Management", 1997.

# Early Use of Regions

## Capability Calculus<sup>35</sup>

- Arbitrarily-ordered region allocation/deallocation, via capability tracking
- Capability: set of regions presently valid to access

Also see:

- RC<sup>36</sup>
- Reaps<sup>37</sup>
- Control-C<sup>38</sup> and Type Homogeneity<sup>39</sup>

---

<sup>35</sup>Crary, Walker, and Morrisett, "Typed Memory Management in a Calculus of Capabilities", 1999.

<sup>36</sup>Gay and Aiken, "Language Support for Regions", 2001.

<sup>37</sup>Berger, Zorn, and McKinley, "OOPSLA 2002", 2002.

<sup>38</sup>Kowshik, Dhurjati, and Adve, "Ensuring Code Safety Without Runtime Checks for Real-time Control Systems", 2002.

<sup>39</sup>Dhurjati, Kowshik, et al., "Memory safety without runtime checks or garbage collection", 2003.

# Regions in Cyclone

Cyclone,<sup>40</sup> again!

- Three region types:
  - single **heap** region
  - **stack** regions
  - **dynamic** regions
- Lifetime subtyping: region A <: region B  $\Leftrightarrow$  region A **outlives** region B
- Sane defaults by inferring region annotations on pointer types

---

<sup>40</sup>Grossman et al., "Region-based Memory Management in Cyclone", 2002.

## Regions in Cyclone

- Pointer can escape scope of their regions
- Type system tracks **capability** (live regions set)
- Check needed **liveness** on pointer dereference
- Function **effect**: set of regions it might access
- Calculate function's effect from prototype alone

```
1 struct Set< $\alpha$ ,  $\rho$ ,  $\epsilon$ >
2 {
3     list_t< $\alpha$ ,  $\rho$ > elts;
4     int (*cmp)( $\alpha$ , $\alpha$ ;  $\epsilon$ );
5 }
```

```
1 struct Set< $\alpha$ ,  $\rho$ >
2 {
3     list_t< $\alpha$ ,  $\rho$ > elts;
4     int (*cmp)( $\alpha$ , $\alpha$ ; regions_of( $\alpha$ ));
5 }
```

## Some Cyclone Abstract Syntax

kinds	$\kappa$	$::= T \mid R$
type and region vars	$\alpha, \rho$	
region sets	$\epsilon$	$::= \emptyset \mid \alpha \mid \epsilon_1 \cup \epsilon_2$
region constraints	$\gamma$	$::= \emptyset \mid \gamma, \epsilon <: \rho$
constructors	$\tau$	$::= \alpha \mid \text{int} \mid \tau_1 \xrightarrow{\epsilon} \tau_2 \mid \tau @ \rho \mid \text{handle}(\rho) \mid \forall \alpha : \kappa \triangleright \gamma. \tau \mid \dots$
expressions	$e$	$::= x_\rho \mid v \mid e \langle \tau \rangle \mid * e \mid \text{new}(e_1)e_2 \mid e_1(e_2) \mid \&e \mid \dots$
functions	$f$	$::= \rho : (\tau_1 x_\rho) \xrightarrow{\epsilon} \tau_2 = \{s\} \mid \Lambda \alpha : \kappa \triangleright \gamma. f$
statements	$s$	$::= e \mid s_1; s_2 \mid \text{if } (e) s_1 \text{ else } s_2 \mid \rho : \{ \tau x_\rho = e; s \} \mid \text{region} \langle \rho \rangle x_\rho s \mid \dots$
		...

## Example Cyclone Judgments

$\Delta; \Gamma; \gamma; \epsilon; \tau \vdash_{stmt} S$

$\Delta; \Gamma; \gamma; \epsilon \vdash e : \tau$

$\Gamma \vdash \epsilon \Rightarrow \rho$

$\Delta$  : in-scope type/region vars

$\Gamma$  : mapping of in-scope vars to types

$\gamma$  : constraints relating lifetimes

$\epsilon$  : capability, i.e. live regions

$$\frac{\gamma \vdash \epsilon \Rightarrow \rho}{\Delta; \Gamma; \gamma; \epsilon \vdash x_\rho : \Gamma(x_\rho)} \text{ (VAR)} \quad \frac{\Delta; \Gamma; \gamma; \epsilon \vdash e : \tau * \rho \quad \gamma \vdash \epsilon \Rightarrow \rho}{\Delta; \Gamma; \gamma; \epsilon \vdash *e : \tau} \text{ (DEREF)}$$

$$\frac{\Delta; \Gamma; \gamma; \epsilon \vdash e_1 : \tau_2 \xrightarrow{\epsilon_1} \tau \quad \Delta; \Gamma; \gamma; \epsilon \vdash e_2 : \tau_2 \quad \gamma \vdash \epsilon \Rightarrow \epsilon_1}{\Delta; \Gamma; \gamma; \epsilon \vdash e_1(e_2) : \tau} \text{ (CALL)}$$

$$\frac{\Delta; \Gamma; \gamma; \epsilon \vdash e : \forall \alpha : \kappa \triangleright \gamma_1. \tau_2 \quad \Delta \vdash \tau : \kappa \quad \gamma \vdash \gamma_1[\tau_1/\alpha]}{\Delta; \Gamma; \gamma; \epsilon \vdash e\langle \tau_1 \rangle : \tau_2[\tau_1/\alpha]} \text{ (TYPE-INST)}$$



# Cyclone Results

Soundness Theorem<sup>41</sup>:

- The program cannot get stuck from type errors or dangling-pointer dereferences.
- The terminating program deallocates all regions it allocates

Benchmarks

- 86 lines of region annotation-related changes across 18,000 lines (6%)
- Eliminate heap allocation entirely for web-server
- Near-zero overhead (from garbage collection and bounds checking)

---

<sup>41</sup>Grossman et al., *Formal Type Soundness for Cyclone's Region System*, 2001.

# Linear Types

Linear types<sup>42,43</sup> ensure that every variable is used exactly *once*.

- The **world** is a non-duplicatable resource
- Track proper modification of world
- Efficiency: safe to destructively update an array
- **Memory management: can immediately collect used values**

---

<sup>42</sup>Girard, "Linear logic", 1987.

<sup>43</sup>Wadler, "Linear types can change the world", 1990.

# Applied Linear Types

Clay<sup>44</sup>

- Type-theoretic basis for giving type-safe code more control over memory
- Singleton types to type check loads, coercion functions to modify values' type safely

PACLANG<sup>45</sup>

- Program network processors for handling packets
- Unique ownership property: each packet in heap is referenced by exactly one thread
- Allow mutable aliasing within the same thread
- Operations for a functions to 1) take ownership or 2) create local aliases

---

<sup>44</sup>Hawblitzel et al., "Low-Level Linear Memory Management", 2004.

<sup>45</sup>Ennals, Sharp, and Mycroft, "Linear Types for Packet Processing", 2004.

# COGENT

COGENT<sup>46, 47</sup>

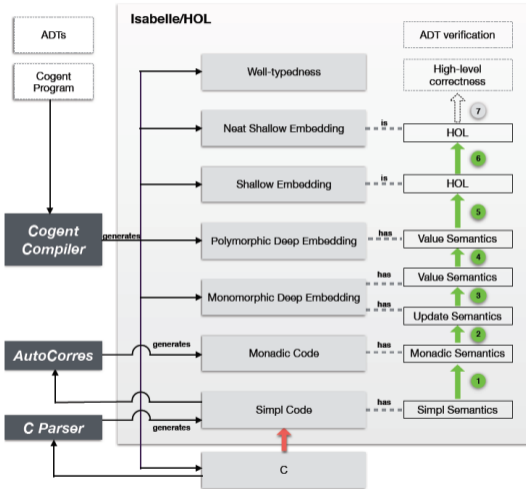
- Pure, polymorphic language with linear types for writing low-level systems code
- Efficient machine code with in-place updates
- Linear types:
  - Ensure safe handling of heap-allocated objects
  - Equational functional semantics via mutable state/imperative effects
- Reason with **interactive theorem prover**
- Missing functionality can be implemented in C, manually verified
- No trusted compiler, runtime, or garbage collector needed

---

<sup>46</sup>Amani et al., "Cogent", 2016.

<sup>47</sup>O'Connor et al., "COGENT", 2016.

# COGENT



## Person-Months of Work

- Proof Framework: 33.5
- Compiler: 10
- Proofs: 18

## Lines (kLOC)

- Isabelle theorems: 17
- Compiler: 9.5
- ext2 Filesystem: 6.5 (Isabelle/HOL: 76.7)

## Other Linear Types

### Quasi-linear types<sup>48</sup>

- Distinguish consumed values from those that may be returned
- Use  $\kappa$  to control how often a variable of type  $\tau^\kappa$  is used (many times locally)

### Vault<sup>49</sup>

- Keys associate static capabilities with run-time resources
- Annotate functions with effect clause (pre- and post-conditions on held-key set)
- Windows 2000 locking errors, IRP ownership model

### Ordered types for memory layout<sup>50</sup>

- Variables must be used in order  $\Rightarrow$  memory locations
- *Orderly lambda calculus* for size-preserving memory operations

---

<sup>48</sup>Kobayashi, "Quasi-linear Types", 1999.

<sup>49</sup>DeLine and Fähndrich, "Enforcing High-level Protocols in Low-level Software", 2001.

<sup>50</sup>Petersen et al., "A Type Theory for Memory Allocation and Data Layout", 2003.

# Ownership

Types can represent **ownership** and prevent *aliasing* and *mutation* on the same location.

## LCL<sup>51</sup>

- **owned** annotation to denote reference with obligation to release storage
- **dependent** annotation for sharing; user ensures lifetimes contained properly

## Ownership Types<sup>52</sup>

- Object's definition includes **unique** object context that owns it

## Singularity<sup>53</sup>

- Type system tracks resources, passes ownership of arguments to callee

---

<sup>51</sup>Evans, "Static Detection of Dynamic Memory Errors", 1996.

<sup>52</sup>Clarke, Potter, and Noble, "Ownership Types for Flexible Alias Protection", 1998.

<sup>53</sup>Fähndrich et al., "Language Support for Fast and Reliable Message-based Communication in Singularity OS", 2006.

- Ownership and lifetimes
- Type system enforces that objects have unique **owners**
- Objects may be borrowed for no longer than owner
- Object automatically **deallocated** when owner leaves scope
- Unsafe sections for mutating raw pointers and aliased state
- Regions  $\approx$  lifetimes (region capabilities  $\approx$  lifetime tokens)

```
1 let (snd, rcv) = channel();
2 join(move || {
3     let mut v = Vec::new();
4     v.push(0);
5     snd.send(v);
6     v.push(1);
7 },
8 move || {
9     let v = rcv.recv().unwrap();
10    println!("Received: ␣{:?}", v);
11 });
```

---

<sup>54</sup>Matsakis and Klock, "The Rust Language", 2014.

<sup>55</sup>Levy et al., "Ownership is Theft", 2015.

<sup>56</sup>Jung et al., "RustBelt", 2017.



## Quick Temporal Recap

- Capabilities and pointer-based metadata
- Effects and regions
- Linear types
- Ownership and borrowing

## Conclusion

- Memory errors  $\equiv$  *type errors*
- *Static* memory management
- Isolate unsafe world
- Be reasonable and optimistic

Thanks

Thanks!