

## Types for Safe C-Level Programming Part 3: Basic Cyclone-Style Region- Based Memory Management

Dan Grossman  
University of Washington  
26 July 2007

## C-level Quantified Types

- As usual, a type variable hides a type's identity
  - Still usable because multiple in same scope hide the same type
- For code reuse and abstraction
- But so far, if you have a  $\tau^*$  (and  $\tau$  has known size), then you can dereference it
  - If the pointed-to location has been *deallocated*, this is broken ("should get stuck")
  - Cannot happen in a garbage-collected language
- All this type-variable stuff will help us!

26 July 2007

Dan Grossman, 2007 Summer School

2

## Safe Memory Management

- Accessing recycled memory violates safety (*dangling pointers*)
- *Memory leaks* crash programs
- In most safe languages, objects conceptually *live forever*
- Implementations use *garbage collection*
- Cyclone needs *more options*, without sacrificing safety/performance

26 July 2007

Dan Grossman, 2007 Summer School

3

## The Selling Points

- **Sound**: programs never follow dangling pointers
- **Static**: no "has it been deallocated" run-time checks
- **Convenient**: few explicit annotations, often allow address-of-locals
- **Exposed**: users control lifetime/placement of objects
- **Comprehensive**: uniform treatment of stack and heap
- **Scalable**: all analysis intraprocedural


26 July 2007

Dan Grossman, 2007 Summer School

4

## Regions

- a.k.a. zones, arenas, ...
- Every object is in exactly one region
- All objects in a region are deallocated simultaneously (no *free* on an object)
- Allocation via a region *handle*



An old idea with some support in languages (e.g., RC)  
and implementations (e.g., ML Kit)

26 July 2007

Dan Grossman, 2007 Summer School

5

## Cyclone Regions

- **heap region**: one, lives forever, conservatively GC'd
- **stack regions**: correspond to local-declaration blocks:  

```
{int x; int y; s}
```
- **dynamic regions**: lexically scoped lifetime, but growable:  

```
{ region r; s }
```
- allocation: `xnew(r, 3)`, where `r` is a *handle*
- handles are first-class
  - caller decides where, callee decides how much
  - heap's handle: `heap_region`
  - stack region's handle: none

26 July 2007

Dan Grossman, 2007 Summer School

6

## That's the Easy Part

The implementation is *dirt simple* because the type system statically prevents dangling pointers

```
void f() {
  int* x;
  if(1) {
    int y=0;
    x=&y;
  }
  *x;
}

int* g(region_t r) {
  return rnew(r,3);
}

void f() {
  int* x;
  { region r;
    x=g(r);
  }
  *x;
}
```

26 July 2007

Dan Grossman, 2007 Summer School

7

## The Big Restriction

- Annotate all pointer types with a *region name* (a type variable of region kind)
- `int*p` can point only into the region created by the construct that introduces `p`
  - heap introduces  $p_H$
  - `L:...` introduces  $p_L$
  - `{region r; s}` introduces  $p_r$ 
    - `r` has type `region_t<p_r>`

26 July 2007

Dan Grossman, 2007 Summer School

8

## So What?

*Perhaps the scope of type variables suffices*

```
void f() {
  int*p_L x;
  if(1) {
    L: int y=0;
    x=&y;
  }
  *x;
}
```

- type of `x` makes no sense
- good intuition for now
- but simple scoping will *not* suffice in general

26 July 2007

Dan Grossman, 2007 Summer School

9

## Where We Are

- Basic region constructs
- Type system annotates pointers with type variables of region kind
- More expressive: region *polymorphism*
- More expressive: region *subtyping*
- More convenient: avoid explicit annotations
- Revenge of existential types

26 July 2007

Dan Grossman, 2007 Summer School

10

## Region Polymorphism

Apply everything we did for type variables to region names (only it's more important!)

```
void swap(int *p1 x, int *p2 y){
  int tmp = *x;
  *x = *y;
  *y = tmp;
}

int*p sumptr(region_t<p> r, int x, int y){
  return rnew(r) (x+y);
}
```

26 July 2007

Dan Grossman, 2007 Summer School

11

## Polymorphic Recursion

```
void fact(int*p result, int n) {
  L: int x=1;
  if(n > 1) fact<pr>(&x,n-1);
  *result = x*n;
}

int g = 0;

int main() {
  fact<pH>(&g,6);
  return g;
}
```

26 July 2007


Dan Grossman, 2007 Summer School

12

## Type Definitions

```
struct ILst< $\rho_1, \rho_2$ > {
  int* $\rho_1$  hd;
  struct ILst< $\rho_1, \rho_2$ > * $\rho_2$  tl;
};
```

- What if we said `ILst< $\rho_2, \rho_1$ >` instead?
- Moral: when you're well-trained, you can follow your nose



26 July 2007

Dan Grossman, 2007 Summer School

13

## Region Subtyping

If  $p$  points to an `int` in a region with name  $\rho_1$ , is it ever sound to give  $p$  type `int* $\rho_2$` ?

- If so, let `int* $\rho_1$  < int* $\rho_2$`
- Region subtyping is the *outlives* relationship  

```
void f() { region  $\rho_1$ ; ... { region  $\rho_2$ ; ... }}
```
- But pointers are still invariant:  
`int* $\rho_1$ * $p$  < int* $\rho_2$ * $p$`  only if  $\rho_1 = \rho_2$
- Still following our nose

26 July 2007

Dan Grossman, 2007 Summer School

14

## Subtyping cont'd

- Thanks to LIFO, a new region is outlived by all others
- The heap outlives everything

```
void f (int b, int* $\rho_1$  p1, int* $\rho_2$  p2) {
  L: int* $\rho_L$  p;
  if (b) p=p1; else p=p2;
  /* ...do something with p... */
}
```

- Moving beyond LIFO restricts subtyping, but the user has more options

26 July 2007

Dan Grossman, 2007 Summer School

15

## Where We Are

- Basic region region constructs
- Type system annotates pointers with type variables of region kind
- More expressive: region *polymorphism*
- More expressive: region *subtyping*
- More convenient: avoid explicit annotations
- Revenge of existential types

26 July 2007

Dan Grossman, 2007 Summer School

16

## Who Wants to Write All That?

- Intraprocedural *inference*
  - determine region annotation based on uses
  - same for polymorphic instantiation
  - based on unification (as usual)
  - so forget all those `L:` things
- Rest is by *defaults*
  - Parameter types get fresh region names (so default is region-polymorphic with no equalities)
  - Everything else (return values, globals, struct fields) gets  $\rho_H$

26 July 2007

Dan Grossman, 2007 Summer School

17

## Examples

```
void fact(int* result, int n) {
  int x = 1;
  if (n > 1) fact(&x, n-1);
  *result = x*n;
}
void g(int* $\rho$  pp, int* $\rho$  p) { *pp = p; }
```

- The callee ends up writing just the equalities the caller needs to know; caller writes nothing
- Same rules for parameters to structs and typedefs
- In porting, "one region annotation per 200 lines"

26 July 2007

Dan Grossman, 2007 Summer School

18

## But Are We Sound?

- Because types can mention only in-scope type variables, it is hard to create a dangling pointer
- But not impossible: an existential can hide type variables
- Without built-in closures/objects, eliminating existential types is a real loss
- With built-in closures/objects, you have the same problem:  $(\text{fn } x \rightarrow (*y) + x) : \text{int} \rightarrow \text{int}$

26 July 2007

Dan Grossman, 2007 Summer School

19

## The Problem

```

struct T { <α>
  int (*f) (α);
  α env;
};

int read(int*p x) { return *x; }

struct T dangle() {
  L: int x = 0;
  struct T ans =
    T(read<pLpL
  return ans;
}
  
```

26 July 2007

Dan Grossman, 2007 Summer School

20

## And The Dereference

```

void bad() {
  let T{<β> .f=fp, .env=ev} = dangle();
  fp(ev);
}
  
```

Strategy:

- Make the system “feel like” the scope-rule except when using existentials
- Make existentials usable (strengthen `struct T`)
- Allow dangling pointers, prohibit dereferencing them

26 July 2007

Dan Grossman, 2007 Summer School

21

## Capabilities and Effects

- Attach a compile-time *capability* (a set of region names) to each program point
- Dereference requires region name in capability
- Region-creation constructs add to the capability, *existential unpacks do not*
- Each function has an *effect* (a set of region names)
  - body checked with effect as capability
  - call-site checks effect (after type instantiation) is a subset of capability

26 July 2007

Dan Grossman, 2007 Summer School

22

## Not Much Has Changed Yet...

If we let the *default effect* be the region names in the prototype (and  $\rho_H$ ), everything *seems fine*

```

void fact(int*p result, int n ;{p}) {
  L: int x = 1;
  if(n > 1) fact<pLpH

```

26 July 2007

Dan Grossman, 2007 Summer School

23

## But What About Polymorphism?

```

struct Lst<α> {
  α hd;
  struct Lst<α>* tl;
};

struct Lst<β>* map(β f(α ;??),
  struct Lst<α>* p l
  ;??);
  
```

- There's no good answer
- Choosing {} prevents using `map` for lists of non-heap pointers (unless  $\varepsilon$  doesn't dereference them)
- The Tofte/Talpin solution: *effect variables*  
a type variable of kind “set of region names”

26 July 2007

Dan Grossman, 2007 Summer School

24

## Effect-Variable Approach

- Let the default effect be:
  - the region names in the prototype (and  $\rho_H$ )
  - the effect variables in the prototype
  - a fresh effect variable

```
struct Lst< $\beta$ >* map(
     $\beta$  f( $\alpha$  ;  $\epsilon_1$ ),
    struct Lst< $\alpha$ > *p l
    ;  $\epsilon_1 + \epsilon_2 + \{\rho\}$ );
```

26 July 2007

Dan Grossman, 2007 Summer School

25

## It Works

```
struct Lst< $\beta$ >* map(
     $\beta$  f( $\alpha$  ;  $\epsilon_1$ ),
    struct Lst< $\alpha$ > *p l
    ;  $\epsilon_1 + \epsilon_2 + \{\rho\}$ );

int read(int* $\rho$  x ;  $\{\rho\} + \epsilon_1$ ) { return *x; }

void g(;;) {
    L: int x=0;
    struct Lst<int* $\rho_L$ >* $\rho_H$  l =
        new Lst(&x, NULL);
    map< $\alpha$ =int* $\rho_L$   $\beta$ =int  $\rho$ = $\rho_H$   $\epsilon_1$ = $\rho_L$   $\epsilon_2$ ={} >
        (read< $\epsilon_1$ ={}  $\rho$ = $\rho_L$ >, l);
}
```

26 July 2007

Dan Grossman, 2007 Summer School

26

## Not Always Convenient

- With *all default effects*, type-checking will never fail because of effects (!)
- Transparent until there's a function pointer in a struct:

```
struct Set< $\alpha$ ,  $\epsilon$ > {
    struct Lst< $\alpha$ > elts;
    int (*cmp) ( $\alpha$ ,  $\alpha$ ;  $\epsilon$ )
};
```

*Clients must know why  $\epsilon$  is there*

- And then there's the compiler-writer  
*It was time to do something new*

26 July 2007

Dan Grossman, 2007 Summer School

27

## Look Ma, No Effect Variables

- Introduce a type-level operator `regions( $\tau$ )`
- `regions( $\tau$ )` means the set of regions mentioned in  $\tau$ , so it's an effect
- `regions( $\tau$ )` reduces to a normal form:
  - `regions(int) = {}`
  - `regions( $\tau$ * $\rho$ ) = regions( $\tau$ ) +  $\{\rho\}$`
  - `regions( $(\tau_1, \dots, \tau_n) \rightarrow \tau$ ) = regions( $\tau_1$ ) + ... + regions( $\tau_n$ ) + regions( $\tau$ )`
  - `regions( $\alpha$ ) = regions( $\alpha$ )`

26 July 2007

Dan Grossman, 2007 Summer School

28

## Simpler Defaults and Type-Checking

- Let the default effect be:
  - the region names in the prototype (and  $\rho_H$ )
  - regions( $\alpha$ ) for all  $\alpha$  in the prototype

```
struct Lst< $\beta$ >* map(
     $\beta$  f( $\alpha$  ; regions( $\alpha$ ) + regions( $\beta$ )),
    struct Lst< $\alpha$ > *p l
    ; regions( $\alpha$ ) + regions( $\beta$ ) +  $\{\rho\}$ );
```

26 July 2007

Dan Grossman, 2007 Summer School

29

## map Works

```
struct Lst< $\beta$ >* map(
     $\beta$  f( $\alpha$  ; regions( $\alpha$ ) + regions( $\beta$ )),
    struct Lst< $\alpha$ > *p l
    ; regions( $\alpha$ ) + regions( $\beta$ ) +  $\{\rho\}$ );

int read(int * $\rho$  x ;  $\{\rho\}$ ) { return *x; }

void g(;;) {
    L: int x=0;
    struct Lst<int* $\rho_L$ >* $\rho_H$  l =
        new Lst(&x, NULL);
    map< $\alpha$ =int* $\rho_L$   $\beta$ =int  $\rho$ = $\rho_H$ >
        (read< $\rho$ = $\rho_L$ >, l);
}
```

26 July 2007

Dan Grossman, 2007 Summer School

30

## Function-Pointers Work

- With all default effects *and no existentials*, type-checking still won't fail due to effects
- And we fixed the struct problem:

```
struct Set<α> {
  struct Lst<α> elts;
  int (*cmp) (α, α ; regions(α))
};
```

26 July 2007

Dan Grossman, 2007 Summer School

31

## Now Where Were We?

- Existential types allowed dangling pointers, so we added effects
- The effect of polymorphic functions wasn't clear; we explored two solutions
  - effect variables (previous work)
    - simpler
    - better interaction with structs
- Now back to existential types
  - effect variables (already enough)
  - regions(τ) (need one more addition)

26 July 2007

Dan Grossman, 2007 Summer School


32

## Effect-Variable Solution

```
struct T<ε>{ <α>
  int (*f) (α ; ε);
  α env;
};
```

```
int read(int* p x; {p}) { return *x; }
```

```
struct T<{pL}> dangle() {
  L: int x = 0;
  struct T ans =
    T(read<pL>, &x) ; //int*pL ret addr
  return ans;
}
```



26 July 2007

Dan Grossman, 2007 Summer School


33

## Cyclone Solution, Take 1

```
struct T { <α>
  int (*f) (α ; regions(α));
  α env;
};
```

```
int read(int* p x; {p}) { return *x; }
```

```
struct T dangle() {
  L: int x = 0;
  struct T ans =
    T(read<pL>, &x) ; //int*pL ret addr
  return ans;
}
```



26 July 2007

Dan Grossman, 2007 Summer School

34

## Allowed, But Useless!

```
void bad() {
  let T{<β> .f=fp, .env=ev} = dangle();
  fp(ev); // need regions(β)
}
```

- We need some way to "leak" the capability needed to call the function, preferably without an effect variable
- The addition: a *region bound*

26 July 2007

Dan Grossman, 2007 Summer School


35

## Cyclone Solution, Take 2

```
struct T<ρB> { <α> α > ρB
  int (*f) (α ; regions(α));
  α env;
};
```

```
int read(int* p x; {p}) { return *x; }
```

```
struct T<ρL> dangle() {
  L: int x = 0;
  struct T<ρL> ans =
    T(read<ρL>, &x) ; //int*pL
  return ans;
}
```



26 July 2007

Dan Grossman, 2007 Summer School

36

## Not Always Useless

```
struct T<PB> { <α> α > PB
  int (*f) (α ; regions(α));
  α env;
};

struct T<P> no_dangle(region_t<P> r ; {P});

void no_bad(region_t<P> r ; {P}) {
  let T{<β> .f=fp, .env=ev} = no_dangle(r);
  fp(ev); // have P and P ⇒ regions(β)
}
```

*"Reduces effect to a single region"*

26 July 2007

Dan Grossman, 2007 Summer School

37

## Effects Summary

- Without existentials (closures, objects), simple region annotations sufficed
- With hidden types, we need effects
- With effects and polymorphism, we need abstract sets of region names
  - effect variables worked but were complicated and made function pointers in structs clumsy
  - regions(α) and region bounds were our technical contributions

26 July 2007

Dan Grossman, 2007 Summer School

38

## We Proved It

- 40 pages of formalization and proof
- Heap organized into a stack of regions at run-time
- Quantified types can introduce region bounds of the form  $\epsilon > P$
- "Outlives" subtyping with subsumption rule
- Type Safety proof shows
  - no dangling-pointer dereference
  - all regions are deallocated ("no leaks")
- Difficulties
  - type substitution and regions(α)
  - proving LIFO preserved

26 July 2007

Dan Grossman, 2007 Summer School

39

## Scaling it up (another 3 years)

Region types and effects form the core of Cyclone's type system for memory management

- Defaults are *crucial* for hiding most of it most of the time!
- But LIFO is too restrictive; need more options
  - "Dynamic regions" can be deallocated whenever
  - Statically prevent deallocation while "using"
  - Check for deallocation before "using"
  - Combine with *unique pointers* to avoid leaking the space needed to do the check
  - See SCP05/ISMM04 papers (after PLDI02 paper)

26 July 2007

Dan Grossman, 2007 Summer School

40

## Conclusion

- Making an efficient, safe, convenient C is a lot of work
- Combine cutting-edge language theory with careful engineering and user-interaction
- Must get the common case right
- Formal models take a lot of taste to make as simple as possible and no simpler
  - They don't all have to look like ML or TAL

26 July 2007

Dan Grossman, 2007 Summer School

41