


On the Usefulness of Liveness for Garbage Collection and Leak Detection

Martin Hirzel, Amer Diwan, and Antony Hosking
{hirzel,diwan}@cs.colorado.edu hosking@cs.purdue.edu

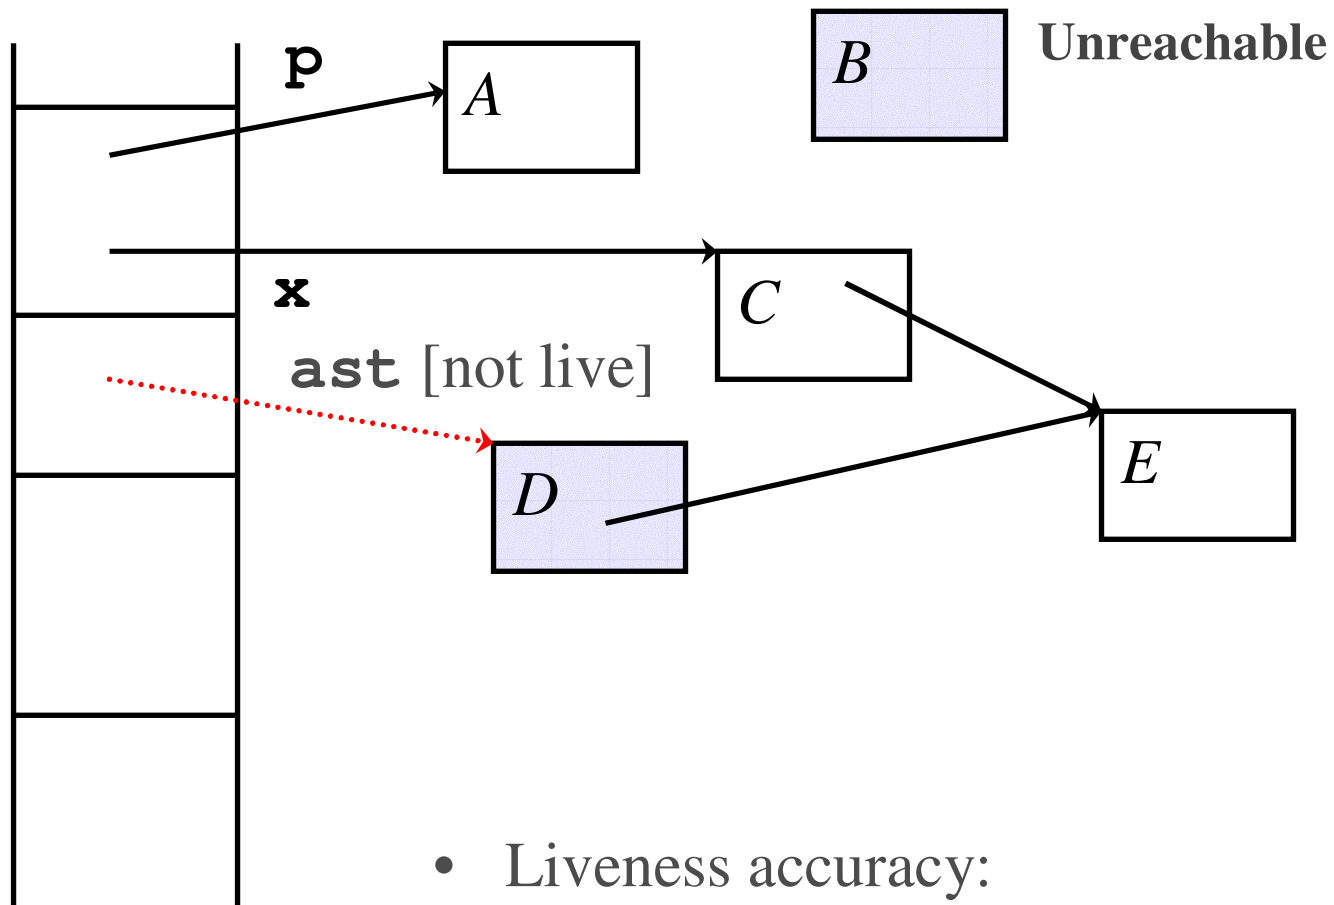
ECOOP June 2001 Budapest, Hungary

What is Liveness?

- A variable is *live* if its value will be used in the future.

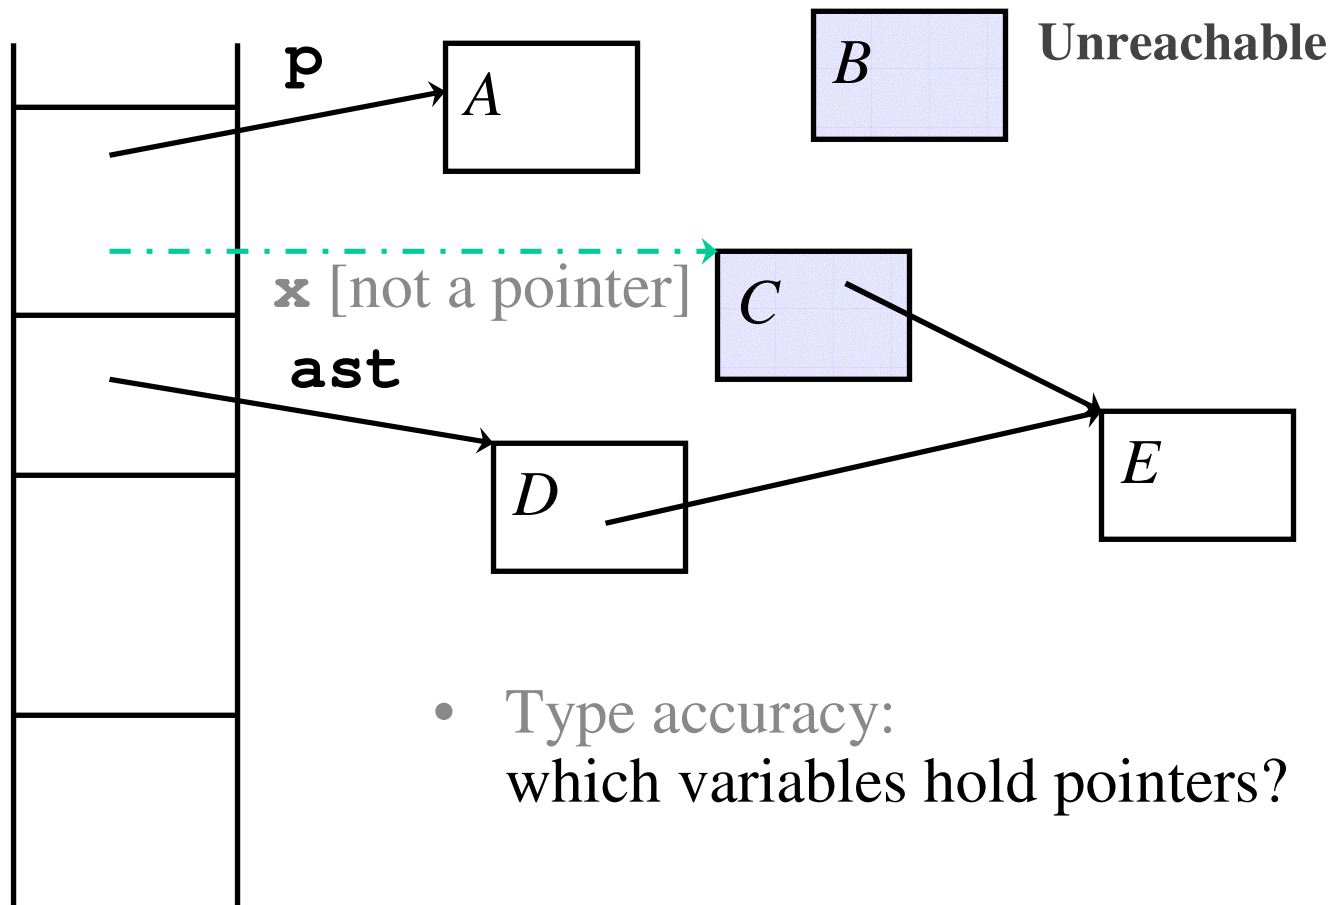
```
...                                     ast  
Tree *ast = parse();  
Cfg *cfg = translate(ast);            
< code that does not use the value  
  of ast >  
...
```


Accuracy and Reachability



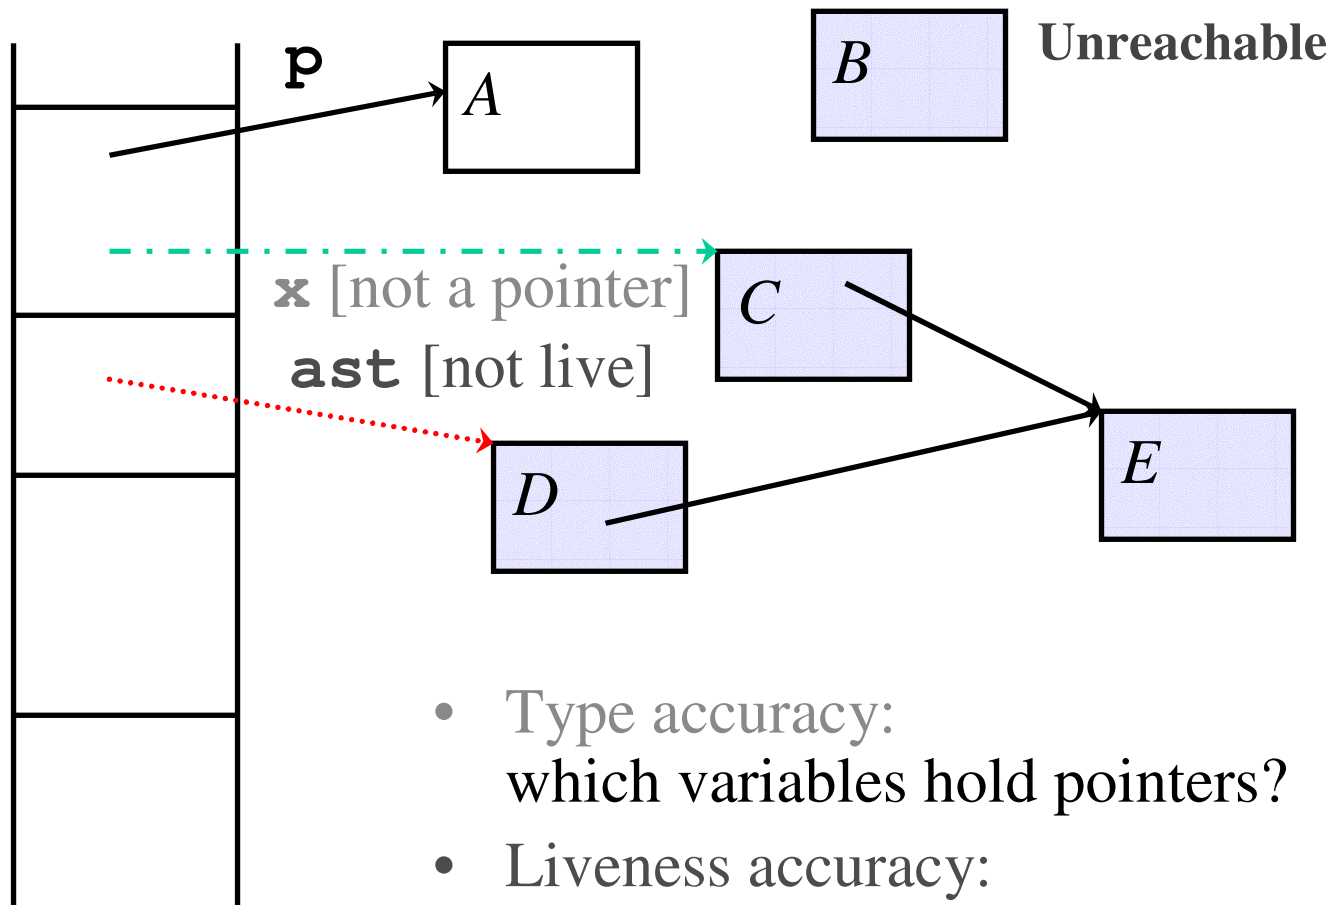
- Liveness accuracy:
which variables are live?

Accuracy and Reachability



- Type accuracy:
which variables hold pointers?

Accuracy and Reachability

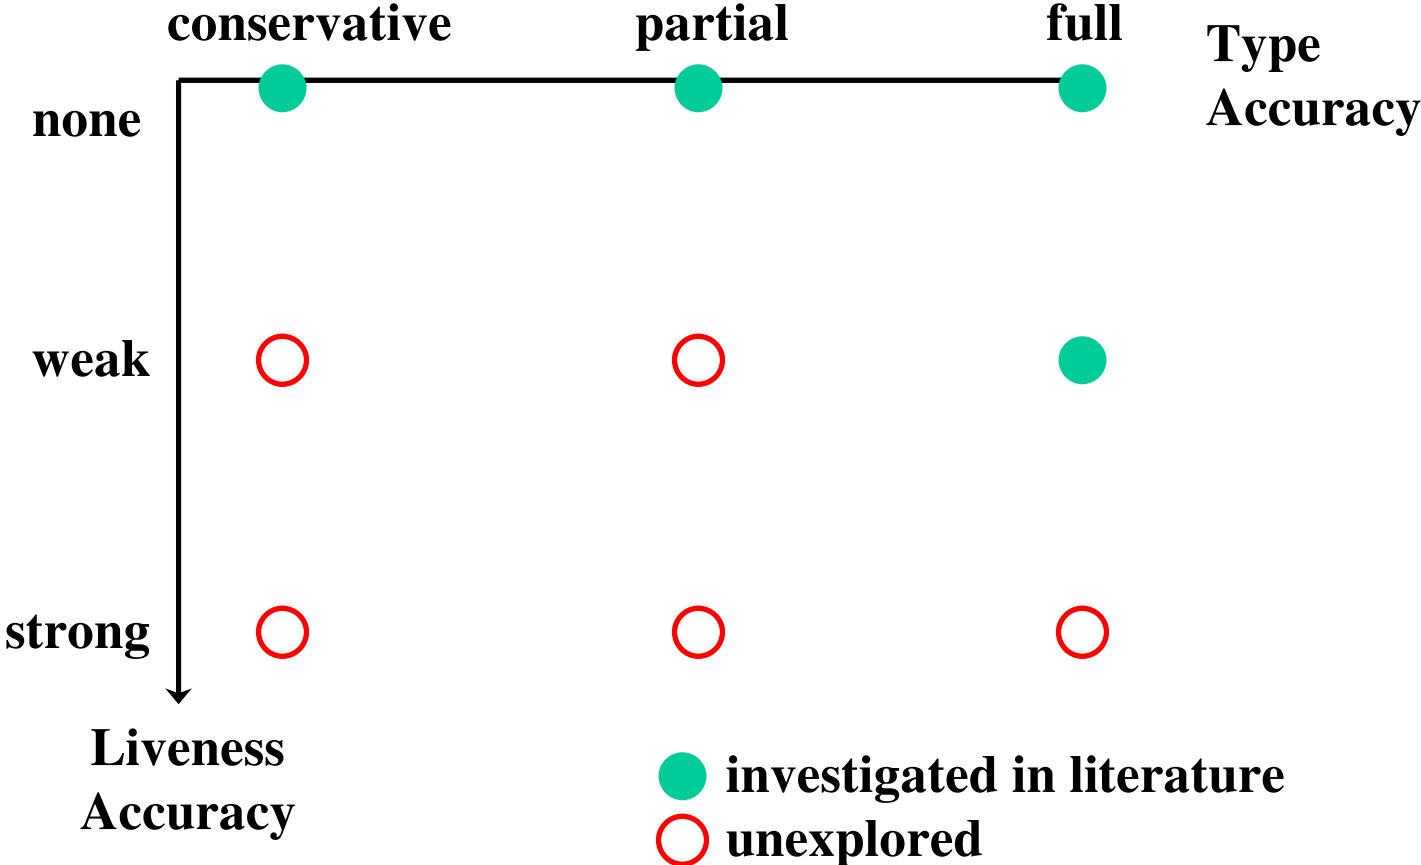


- Type accuracy:
which variables hold pointers?
- Liveness accuracy:
which variables are live?

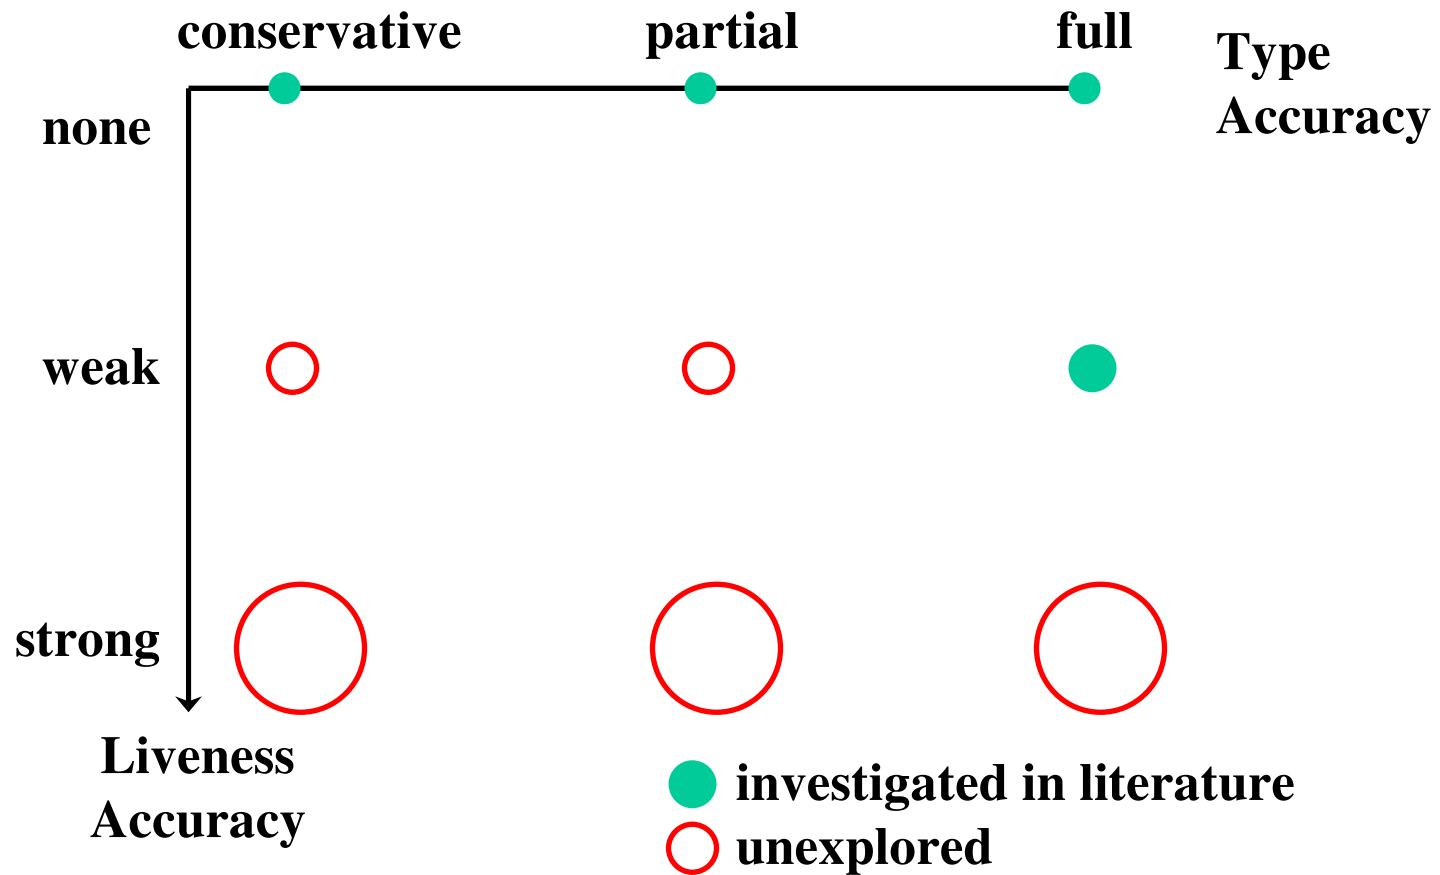
The Questions

- Can liveness accuracy benefit reachability traversals?
 - Useful for garbage collection?
 - Useful for leak detection?
- What kind of analysis is necessary?

Motivation



Results in a Nutshell



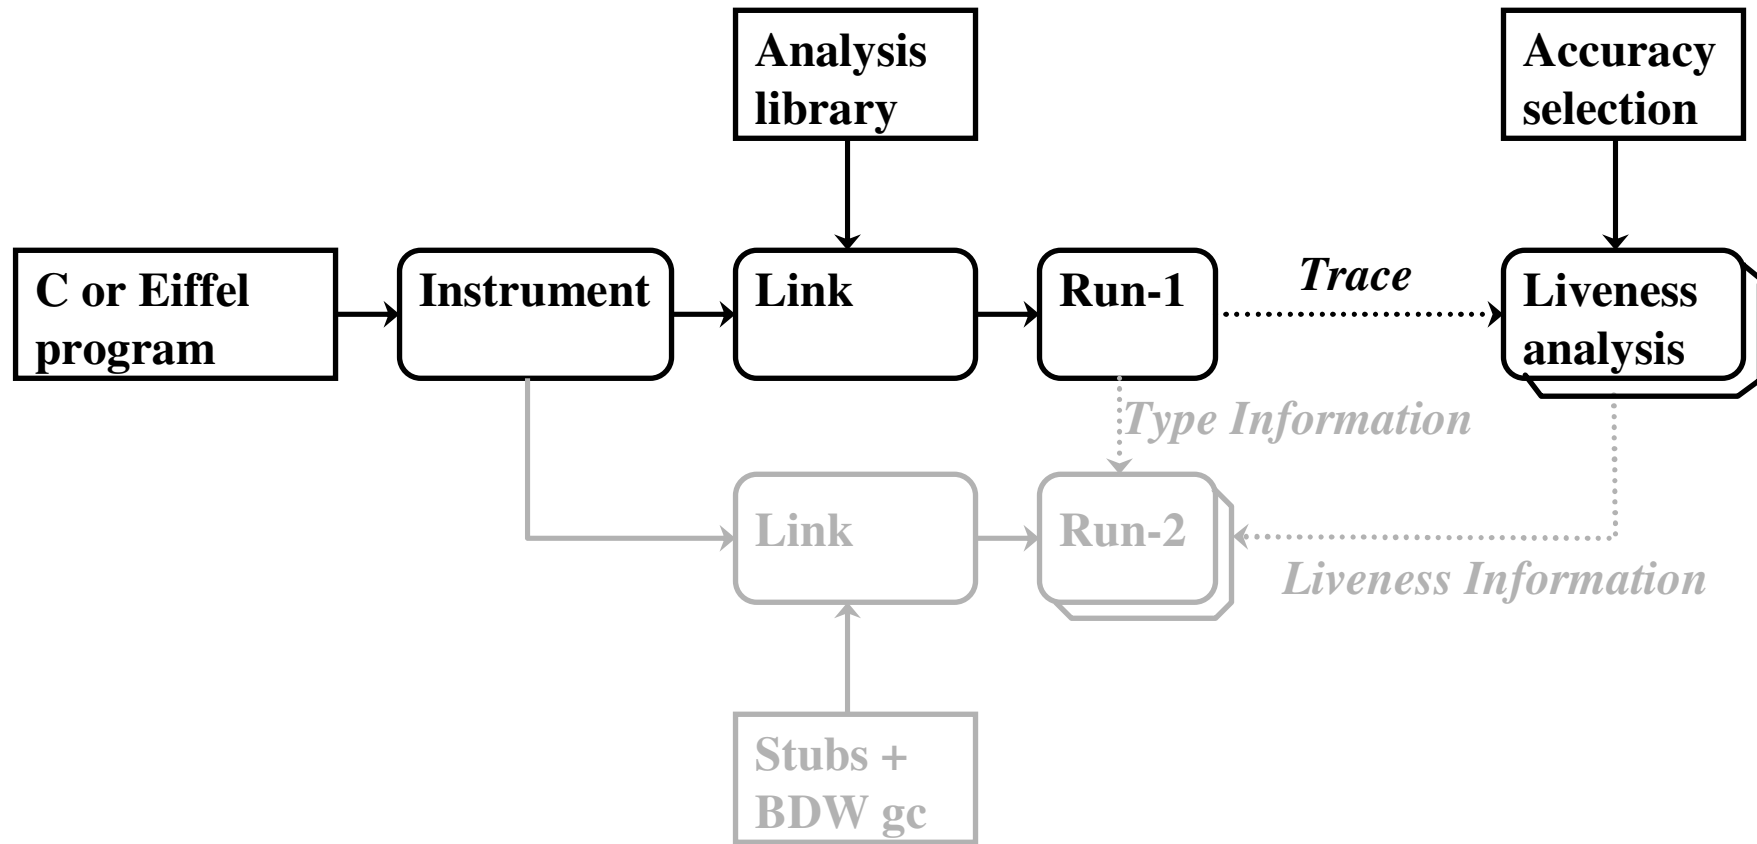
Outline

Introduction	<ul style="list-style-type: none">• Motivation• Preview of results
Methodology	<ul style="list-style-type: none">• Obtaining liveness information• Metrics
Results	<ul style="list-style-type: none">• Reachable heap given various accuracy schemes
Conclusion	<ul style="list-style-type: none">• Related work• Summary

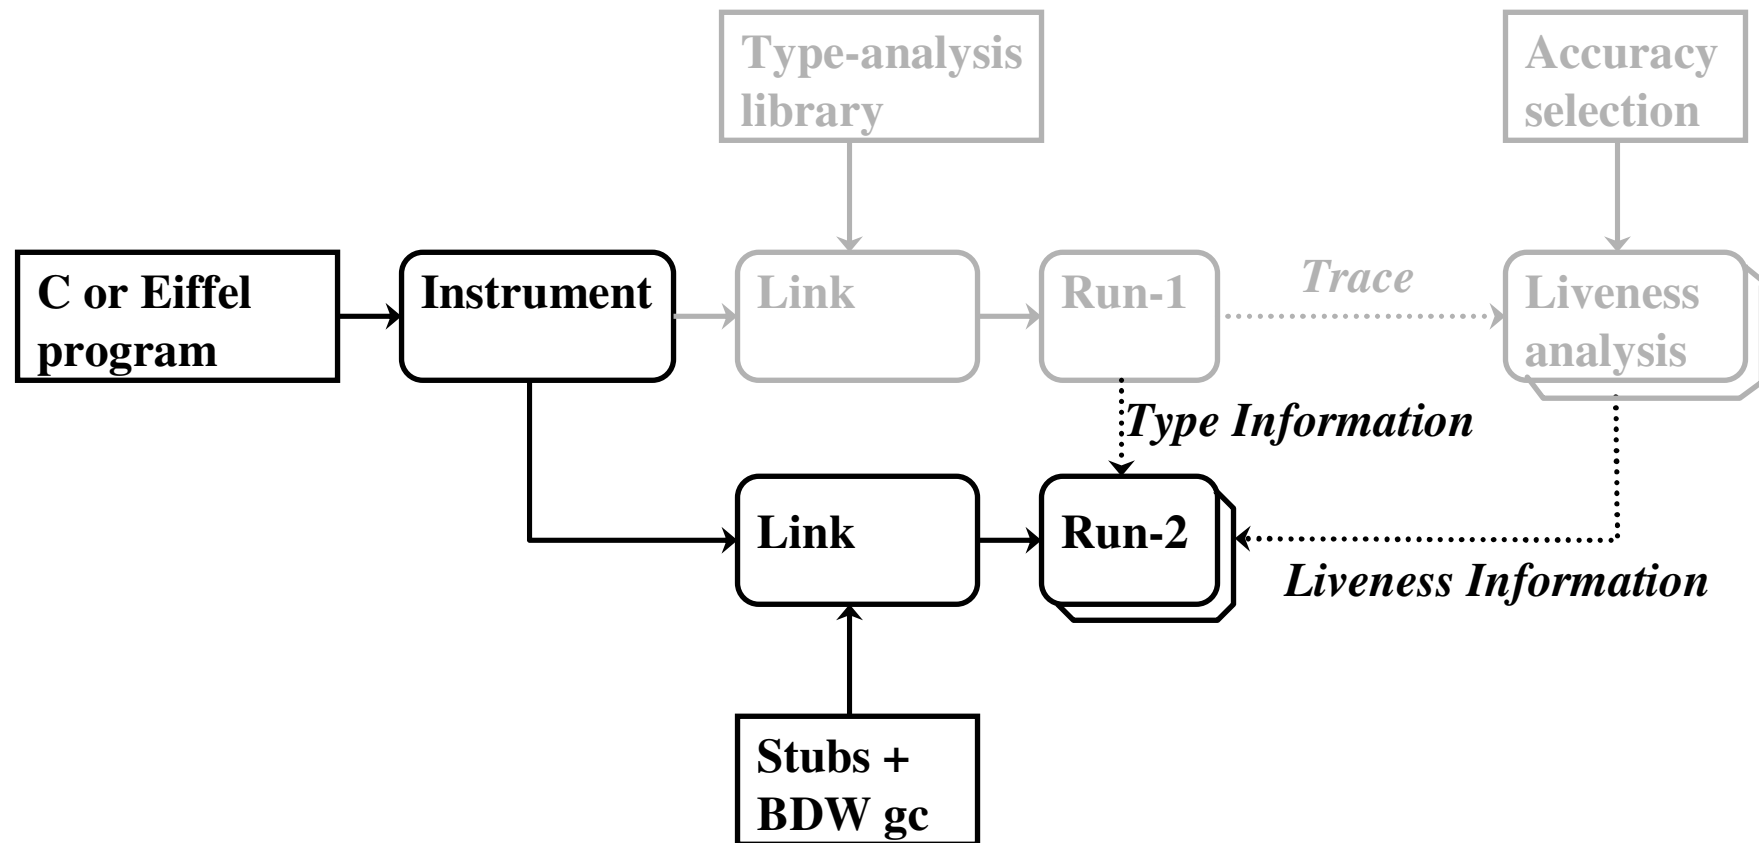
Liveness Approaches

- Static analysis
 - Compiler-analysis of source code
 - Disadvantage: difficult, cost + benefits unclear
- Dynamic analysis
 - Trace-based analysis
 - Disadvantage: two runs needed, limit study

Infrastructure for Experiments



Infrastructure for Experiments



Generating the Trace

Original Code	Instrumented Code	Trace
<pre>x = malloc(4*N); i = 0; while(i < N){ y = x + 4*i; *y = i; i++; }</pre>	<pre>x = malloc(4*N); note_allocation(); note_assign(&x); i = 0; note_assign(&i); while(i < N){ y = x + 4*i; note_assign(&y, &x, &i); *y = i; note_use(&y); note_assign(y, &i); i++; note_assign(&i, &i); }</pre>	<pre>... allocation(91) assign(x) assign(i) assign(y, x, i) use(y) assign(91.0, i) assign(i, i) assign(y, x, i) use(y) assign(91.4, i) ...</pre>

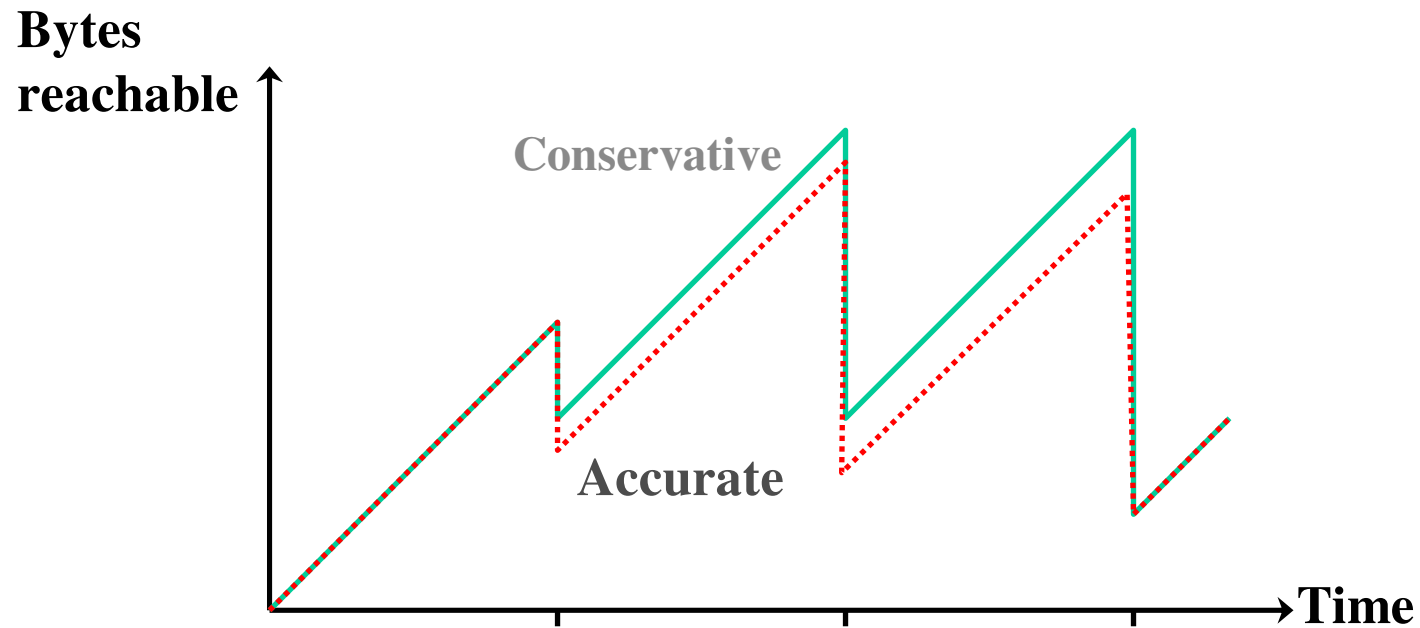
Analyzing the Trace

Trace	Simulated Liveness State	Resulting Information
...		
<i>allocation(91)</i>		
<i>assign(x)</i>		
<i>assign(i)</i>		
<i>assign(y, x, i)</i>		T1: { <i>x, i</i> }
<i>use(y)</i>		
<i>assign(91.0, i)</i>		
<i>assign(i, i)</i>		T1: { <i>x, i</i> }
<i>assign(y, x, i)</i>		
<i>use(y)</i>		
<i>assign(91.4, i)</i>		T2: {}
...		

The diagram illustrates the simulated liveness state for three variables: *x*, *y*, and *i*. The trace consists of several operations: *allocation(91)*, *assign(x)*, *assign(i)*, *assign(y, x, i)*, *use(y)*, *assign(91.0, i)*, *assign(i, i)*, *assign(y, x, i)*, *use(y)*, and *assign(91.4, i)*. The simulated liveness state is shown as green bars for each variable. The first red horizontal line separates the first segment (T1: {*x, i*}) from the second segment. The second red horizontal line separates the second segment (T1: {*x, i*}) from the third segment (T2: {}). The third red horizontal line separates the third segment (T2: {}) from the final segment. The liveness state for *x* and *i* is active during the first two segments, while *y* is active during both segments. In the third segment, *x*, *y*, and *i* are not live.

“A variable is *live* if its value will be used in the future.”

Usefulness Metric for Accuracy



Reachability traversal:	1	2	3
Reachability reduction:	10%	20%	0%

⇒ **The accurate scheme reduced reachability by 10% on average.**

Outline


Introduction	<ul style="list-style-type: none">• Motivation• Preview of results
Methodology	<ul style="list-style-type: none">• Obtaining liveness information• Metrics
Results	<ul style="list-style-type: none">• Reachable heap given various levels of accuracy
Conclusion	<ul style="list-style-type: none">• Related work• Summary

Benchmarks

Name	Language	Lines of Code	Total allocation [Bytes]	Author/Source
Programs written with GC in mind:				
gctest3	C	85	2 200 004	Bartlett
gctest	C	196	1 123 180	Bartlett
bshift	Eiffel	350	28 700	Hirzel
erbt	Eiffel	927	222 300	Durian
ebignum	Eiffel	3 137	109 548	Hillion
li	C	7 597	9 030 872	Spec95
gegrep	Eiffel	17 185	106 392	Bezault
Programs with explicit deallocation:				
anagram	C	647	259 512	Austin
ks	C	782	7 920	Austin
ft	C	2 156	166 832	Austin
yacr2	C	3 979	41 380	Austin
bc	C	7 308	12 382 400	Austin
gzip	C	8 163	14 180	GNU
jpeg	C	31 211	148 664	Spec95


Usefulness of Liveness

Reachability reduction with strongest liveness

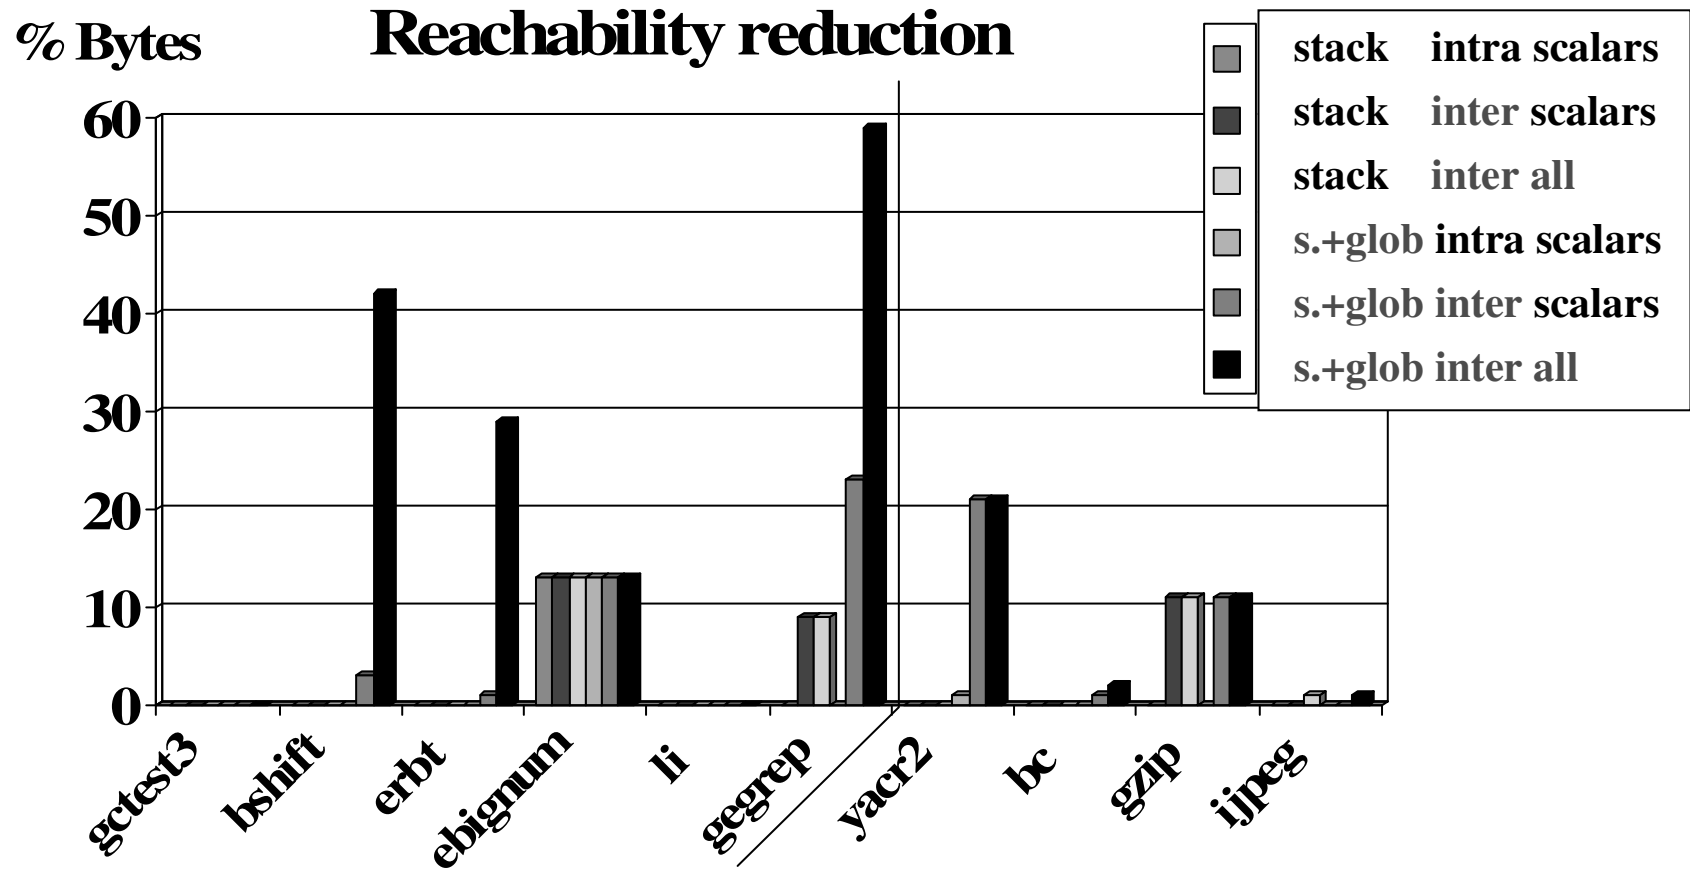


Usefulness of Liveness

Reachability reduction with strongest liveness




Different Levels of Liveness



Type versus Liveness Accuracy

Reachability reduction

% Bytes



Validation

- Comparing liveness information found in different runs
 - For how many locations did the obtained liveness information differ?

Benchmark	Stack % different	Global % different
gegrep	0.7	0.0
yacr2	2.7	0.0
gzip	1.3	2.2

Outline

Introduction	<ul style="list-style-type: none">• Motivation• Preview of results
Methodology	<ul style="list-style-type: none">• Obtaining liveness information• Metrics
Results	<ul style="list-style-type: none">• Reachable heap given various levels of accuracy
Conclusion	<ul style="list-style-type: none">• Related work• Summary

Related Work

- Evaluating Accuracy
 - Hirzel, Diwan: On the type accuracy of garbage collection. ISMM 2000.
 - Shaham, Kolodner, Sagiv: On the effectiveness of GC in Java. ISMM 2000.
- Implementing Accuracy
 - [Bartlett1988] [DiwanMossHudson1992]
[SmithMorrisett1998] [Zorn1993]
[AgesenDetlefsMoss1998] ...

Summary

- Liveness accuracy can be very useful for reachability traversals.
- Strong analyses are necessary to reach significantly fewer Bytes.
- Type accuracy was not very useful in these experiments.