

## Program Equivalence

(slides due to Rahul Sharma)

Prof. Aiken, Barrett & Dill CS 357  
Lecture 13
1

## Motivation

- Verification is specification-limited
  - We need specifications to verification
  - And specifications are hard to come by
- Much research focuses on “well-known” specs
  - Buffer overruns, null dereference, string termination, integer overflows ...

Prof. Aiken, Barrett & Dill CS 357  
Lecture 13
2


## Program Equivalence

- Are two programs equal?
- Arbitrarily hard verification problem
  - Instances vary from easy to impossible
- And a very common “specification”
  - Specifically, for program optimization
  - Two equal programs, one is faster than the other

Prof. Aiken, Barrett & Dill CS 357  
Lecture 13
3

## Performance



- Many systems where code performance matters
  - Compute-bound
  - Repeatedly executed
- Scientific computing
- Graphics
- Low-latency server code
- Encryption/decryption
- ...



Prof. Aiken, Barrett & Dill CS 357  
Lecture 13
4

## More Pain, More Gain



- Programmer Bob requires the best possible code
- Use “best available compiler -O best”
  - E.g, gcc -O3, clang -O3, icc -O3
- Provide unchecked annotations to the compilers
  - E.g, restrict, \_\_builtin
- Optimize by hand

Prof. Aiken, Barrett & Dill CS 357  
Lecture 13
5

## Goal

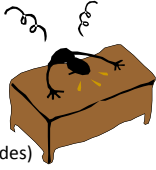
- Minimum pain, maximum gain
- Machine code in 64 bit x86
  - The universal ISA for desktops/servers/laptops
- Automatically generate best possible x86 code

Prof. Aiken, Barrett & Dill CS 357  
Lecture 13
6

## X86 Optimizations are Hard

- x86 is a CISC instruction set
- Many complicated instructions
  - E.g., dpp, popcnt, crc, ...
  - ~2000 instruction variants (~300 opcodes)
  - 3439 page manual
- x86 optimizations requires expert knowledge
  - Experts take hours, days, or even weeks
  - For each new processor



Profs. Aiken, Barrett & Dill CS 357  
Lecture 13 7

## Superoptimization

- Massalin [ASPLOS 87], Bansal and Aiken [ASPLOS 06]
  - Enumerate all possible straight line programs
- STOKE, Schkufza, Sharma, Aiken [ASPLOS 13]
  - Random enumeration instead of exhaustive
- Correctness of straight line code is relatively easy
- But we want to prove correctness of loops ...

Profs. Aiken, Barrett & Dill CS 357  
Lecture 13 8

## Outline

- Equivalence checker for x86
  - Prove correctness of aggressive optimizations
- From checker to optimizer
  - Different application of verification technology
- Use information from context: even better code
- Conclusion

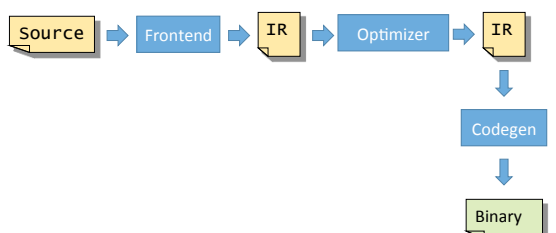
Profs. Aiken, Barrett & Dill CS 357  
Lecture 13 9

## Outline

- **Equivalence checker for x86**
- From checker to optimizer
- Use information from context: even better code
- Conclusion


Profs. Aiken, Barrett & Dill CS 357  
Lecture 13 10

## Compiler structure



Profs. Aiken, Barrett & Dill CS 357  
Lecture 13 11

## Certified Compilation



- Implementation requires a mammoth effort
  - CompCert does not do any loop optimizations
- Maintaining a certified compiler is difficult

Profs. Aiken, Barrett & Dill CS 357  
Lecture 13 12

### Translation Validation

- Instrument the compiler (gcc/llvm)
- Proves correctness of an IR and not the binary
- Validating binaries is harder
  - Lack of structure, static analysis is difficult

13

### Equivalence Checking

Equivalence Checking of Loop Free Code: SMT Solvers

14

### Proof Decomposition

- SMT solvers can prove properties of loop-free code
- We want to prove equivalence of loops
- Decompose proof for loops into subproofs
  - Subproofs about loop free fragment, query SMT solvers
  - Cutpoints: break the loops
  - Invariants: relationships between  $T$  and  $R$  at cutpoints

15

### Proof Decomposition Example

while(i!=0) i--;

16

### Inference

- Given a simulation relation, proofs for loops reduce to subproofs for loop free fragments
  - Query SMT solvers
- Main challenge: infer a simulation relation
  - Infer cutpoints
  - Infer invariants
- We use compilers as black boxes
- Mine relations from data (program executions)

17

### Cutpoints From Data

- Attempt to detect cutpoints
  - Number of times program points are executed

18

### Invariants

- Invariants are restricted to equalities
- Infer invariants from observed data values

**Target  $\mathcal{T}$**

```

movq 8(rsp), rdi
#rdi != 0
movq 8(rsp), rdi
decq rdi
movq rdi, 8(rsp)
retq
    
```

**Data!**

| 8(rsp) | rdi |
|--------|-----|
| 2      | 2   |
| 1      | 1   |
| 0      | 0   |

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 19

### Invariants

- Invariants are restricted to equalities
- Infer invariants from observed data values

**Rewrite  $\mathcal{R}$**

```

movq 8(rsp), r9
#r9 != 0
decq r9
retq
    
```

| 8(rsp) | rdi | r9' |
|--------|-----|-----|
| 2      | 2   | 2   |
| 1      | 1   | 1   |
| 0      | 0   | 0   |

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 20

### Linear algebra

| 8(rsp) | rdi | r9' |
|--------|-----|-----|
| 2      | 2   | 2   |
| 1      | 1   | 1   |
| 0      | 0   | 0   |

- Mine all equalities
- Find all  $w$  s.t.  $Aw=0$
- Nullspace

$I = (8(rsp)=rdi) \wedge (rdi=r9')$

$4eax = edx * 3$   
 $10eax + edx = ecx'$

- $w \downarrow 1 = [-1, 1, 0]$
- $w \downarrow 2 = [0, 1, -1]$

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 21

### Check Simulation Relation

- Query SMT solvers
  - Incorporate counter-examples in relations
- Sound but not complete
  - If checking succeeds then equivalent
  - Can fail to infer a correct simulation relation
  - Learn from counter-examples

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 22

### Results [OOPSLA'13]

| Program     | Run-time (s) |
|-------------|--------------|
| lerner1a    | 12.94        |
| lerner3b    | 53.72        |
| bansal      | 9.89         |
| chomp       | 11.00        |
| fannkuch    | 17.03        |
| knucleotide | 6.56         |
| lists       | 1.40         |
| nsievebits* | 36.44        |
| nsieve*     | 166.31       |
| qsort*      | 140.87       |
| sh1*        | 12888.87     |

- $\mathcal{T}$ : from CompCert
- $\mathcal{R}$ : from GCC
- Prove  $\mathcal{T} \approx \mathcal{R}$

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 23

### Outline

- Equivalence checker for x86
- From checker to optimizer**
- Use information from context: even better code
- Conclusion

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 24

## From Checkers to Optimizers

- We have a checker to prove correctness
  - We can be adventurous with optimizations!
- Bombard the checker with random programs
  - Take the fastest correct program
- STOKE (Schkufza, Sharma, Aiken ASPLOS'13)
  - Start from an initial program (unoptimized, empty, ...)
  - Generate new programs by repeated random changes
  - Output the fastest **correct** program found in time limit

## Random Transformations

### original

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

## Random Transformations

### insert

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
imulq rsi, rdx
...
```

### original

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

## Random Transformations

### insert

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq rsi, rdx
...
```

### original

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

### delete

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

## Random Transformations

### insert

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
imulq rsi, rdx
...
```

### original

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

### delete

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

### instruction

```
...
movl ecx, ecx
shrq 32, rsi
salq 16, rcx
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

## Random Transformations

### insert

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
imulq rsi, rdx
...
```

### original

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

### delete

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

### instruction

```
...
movl ecx, ecx
shrq 32, rsi
salq 16, rcx
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

### opcode

```
...
movl ecx, ecx
shrq 32, rsi
andl 0xffffffff, r9d
movq rcx, rax
subl edx, edx
imulq r9, rax
...
```

### Random Transformations

**insert**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
imulq rsi, rdx
...
```

**original**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

**opcode**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

**delete**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

**instruction**

```
...
movl ecx, ecx
shrq 32, rsi
salq 16, rcx
movl edx, edx
imulq r9, rax
...
```

**operand**

```
...
movl ecx, ecx
shrq 32, rcx
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 31

### Random Transformations

**insert**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
imulq rsi, rdx
...
```

**original**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

**opcode**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

**delete**

```
...
movl ecx, ecx
shrq 32, rsi
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

**instruction**

```
...
movl ecx, ecx
shrq 32, rsi
salq 16, rcx
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

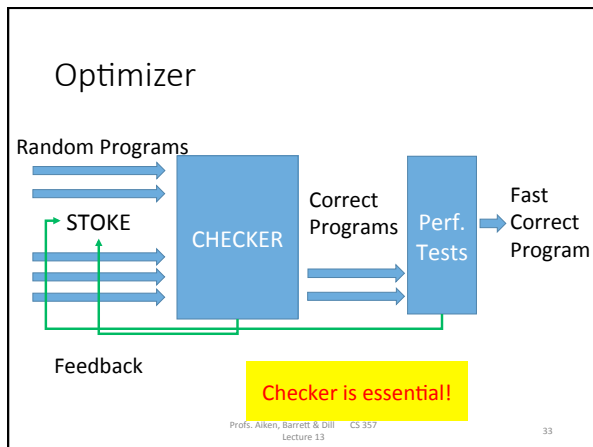
**swap**

```
...
movl ecx, ecx
movl edx, edx
shrq 32, rsi
imulq r9, rax
...
```

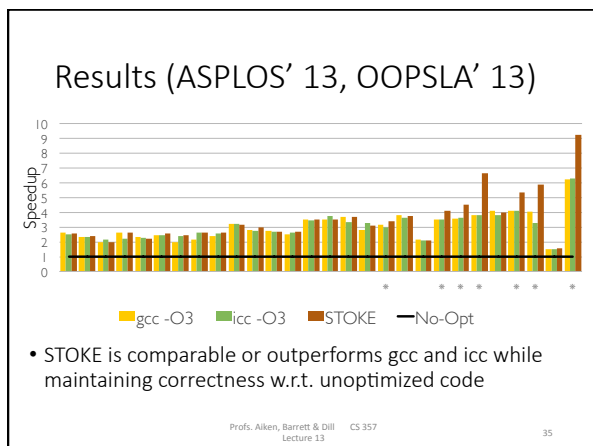
**operand**

```
...
movl ecx, ecx
shrq 32, rcx
andi 0xffffffff, r9d
movq rcx, rax
movl edx, edx
imulq r9, rax
...
```

Prof. Aiken, Barrett & Dill CS 357 Lecture 13 32



- ### Benchmarks
- **Synthesis Kernels:** 25 loop-free kernels taken from A Hacker's Delight (Gulwani et al. '11)
  - **SSL:** Montgomery multiplication kernel
  - **Heap Modifying:** Linked list traversal
  - **Linear Algebra:** SAXPY from BLAS
- Prof. Aiken, Barrett & Dill CS 357 Lecture 13 34



- ### Observation
- Speedups are good: 70% over production compiler
  - However, x86 experts produce much better code
  - Why can't we get 2X, 3X speedups over gcc -O3?
  - Checker rejects many "good" programs
    - Work perfectly with the application
    - Fail on some weird corner case that cannot arise
- Prof. Aiken, Barrett & Dill CS 357 Lecture 13 36

### Outline

- Equivalence checker for x86
- From checker to optimizer
- Use information from context: even better code
- Conclusion

Profs. Aiken, Barrett & Dill CS 357 Lecture 13 37

### Example

|   |  |
|---|--|
| <ul style="list-style-type: none"> <li>• Target</li> </ul> <pre>foo(int32_t* a,     int32_t* b) *a++=*b++; *a++=*b++; *a++=*b++; *a++=*b++; return;</pre> | <ul style="list-style-type: none"> <li>• Expert Rewrite</li> </ul> <pre>MOVAPD xmm0, [b] MOVAPD [a], xmm0 RET</pre> <ul style="list-style-type: none"> <li>• Incorrect, if a and b overlap</li> <li>• Segfault if a and b are not 16 byte aligned</li> </ul> |
|---|--|

Profs. Aiken, Barrett & Dill CS 357 Lecture 13 38

### Conditions

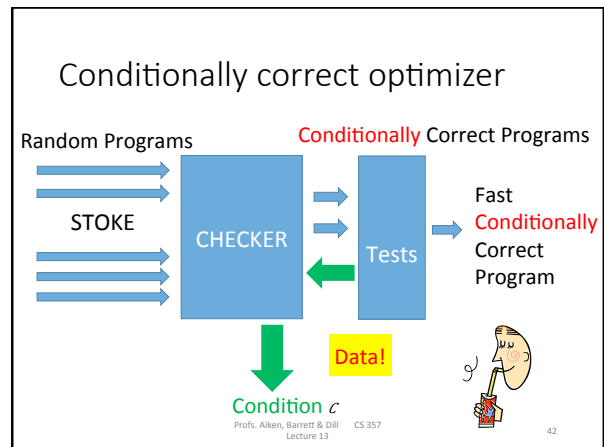
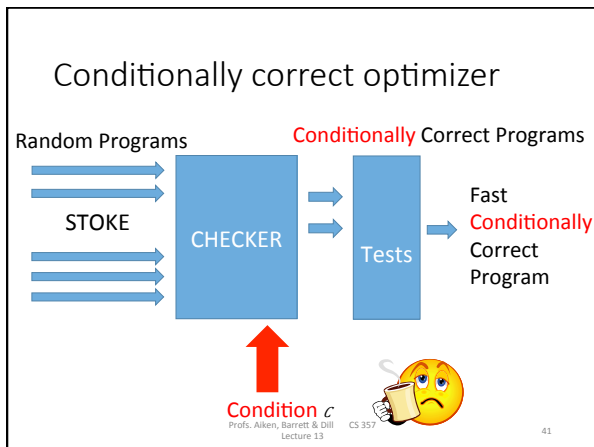
- Cannot use the fast rewrite in an arbitrary context
- Developers know conditions on the context
  - Used in adding unchecked annotations, writing assembly
- Unconditionally correct optimizations are effective
  - But fall short of the fastest code desired
- For performance, we need *conditional equivalence*

Profs. Aiken, Barrett & Dill CS 357 Lecture 13 39

### Conditional Correctness

|   |  |
|---|--|
| <ul style="list-style-type: none"> <li>• Target</li> </ul> <pre>foo(int32_t* a,     int32_t* b) *a++=*b++; *a++=*b++; *a++=*b++; *a++=*b++; return;</pre> | <ul style="list-style-type: none"> <li>• Rewrite</li> </ul> <pre>MOVAPD xmm0, [b] MOVAPD [a], xmm0 RET</pre> <ul style="list-style-type: none"> <li>• Conditionally correct                     <ul style="list-style-type: none"> <li>• restrict(a), restrict(b)</li> <li>• a, b are 16 byte aligned</li> </ul> </li> </ul> |
|---|--|


Profs. Aiken, Barrett & Dill CS 357 Lecture 13 40



## Condition Inference


- Tests encode the conditions on contexts implicitly
- Learn facts about tests
  - Keep adding facts until the proof succeeds
- Aliasing: two memory dereferences do not alias
- Alignment: an address is x-byte aligned
- Equalities, inequalities, floating point specific

## Conditions




1. Analyze A to verify C
2. Manually verify C
3. Runtime check

## Conditional GCC




Annotations can help a compiler significantly

## Conditional STOKE



STOKE with condition inference is comparable or outperforms gcc with or without annotations


## Runtime Checks



Overhead of condition checking can be bearable

## Aliasing

- Restricting aliasing results in simpler queries
  - For equivalence,  $p=q \vee p \neq q$
  - Byte addressable memory, vector instructions





## Conclusion

- Proving the correctness of optimizations is hard
  - Facilitated by data-driven invariant inference
- Better checkers lead to better optimizers
- Context insensitive optimizations fall short
  - Leverage context by inferring conditions from data
- Generate fast and provably correct code