

Coz: Finding Code that Counts with Casual Profiling

Charlie Curtsinger, Emery D. Berger

Grinell College
University of Massachusetts Amherst

Presented at SOSP 2015

Coz: Finding Code that Counts with ~~Causal~~ Casual Profiling

Charlie Curtsinger, Emery D. Berger

Grinell College
University of Massachusetts Amherst

Presented at SOSP 2015

Motivation

Conventional Profilers:

- don't tell you what to optimize
- don't tell you about gains
- are not really made for concurrency

example.cpp

```
1 void a() { // ~6.7 seconds
2     for(volatile size_t x=0; x<2000000000; x++) {}
3 }
4 void b() { // ~6.4 seconds
5     for(volatile size_t y=0; y<1900000000; y++) {}
6 }
7 int main() {
8     // Spawn both threads and wait for them.
9     thread a_thread(a), b_thread(b);
10    a_thread.join(); b_thread.join();
11 }
```

Motivation

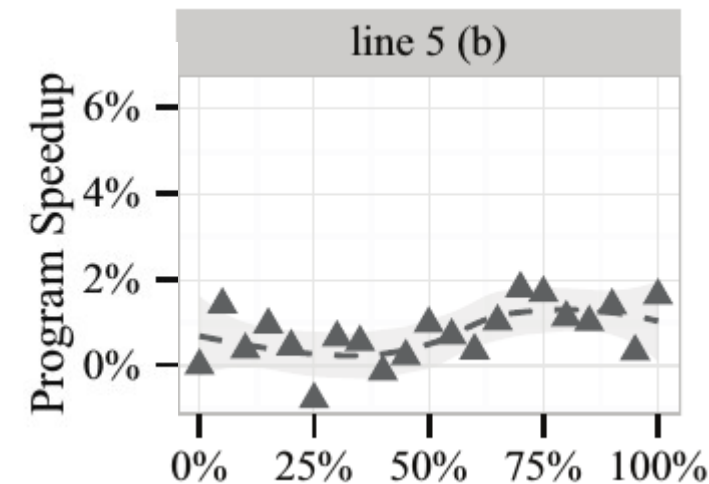

```
1 void a() { // ~6.7 seconds
2   for(volatile size_t x=0; x<2000000000; x++) {}
3 }
4 void b() { // ~6.4 seconds
5   for(volatile size_t y=0; y<1900000000; y++) {}
6 }
7 int main() {
8   // Spawn both threads and wait for them.
9   thread a_thread(a), b_thread(b);
10  a_thread.join(); b_thread.join();
11 }
```

Conventional Profile for `example.cpp`

% time	cumulative seconds	self seconds	calls	self Ts/call	total Ts/call	name
55.20	7.20	7.20	1			a()
45.19	13.09	5.89	1			b()

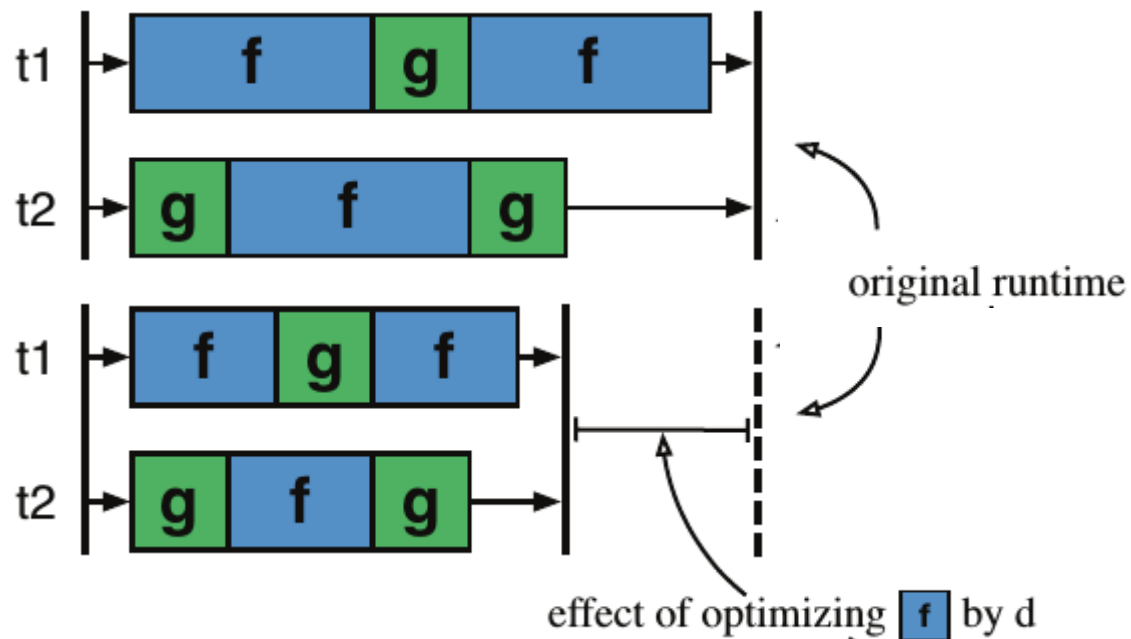
% time	self	children	called	name
				<spontaneous>
55.0	7.20	0.00		a()

				<spontaneous>
45.0	5.89	0.00		b()

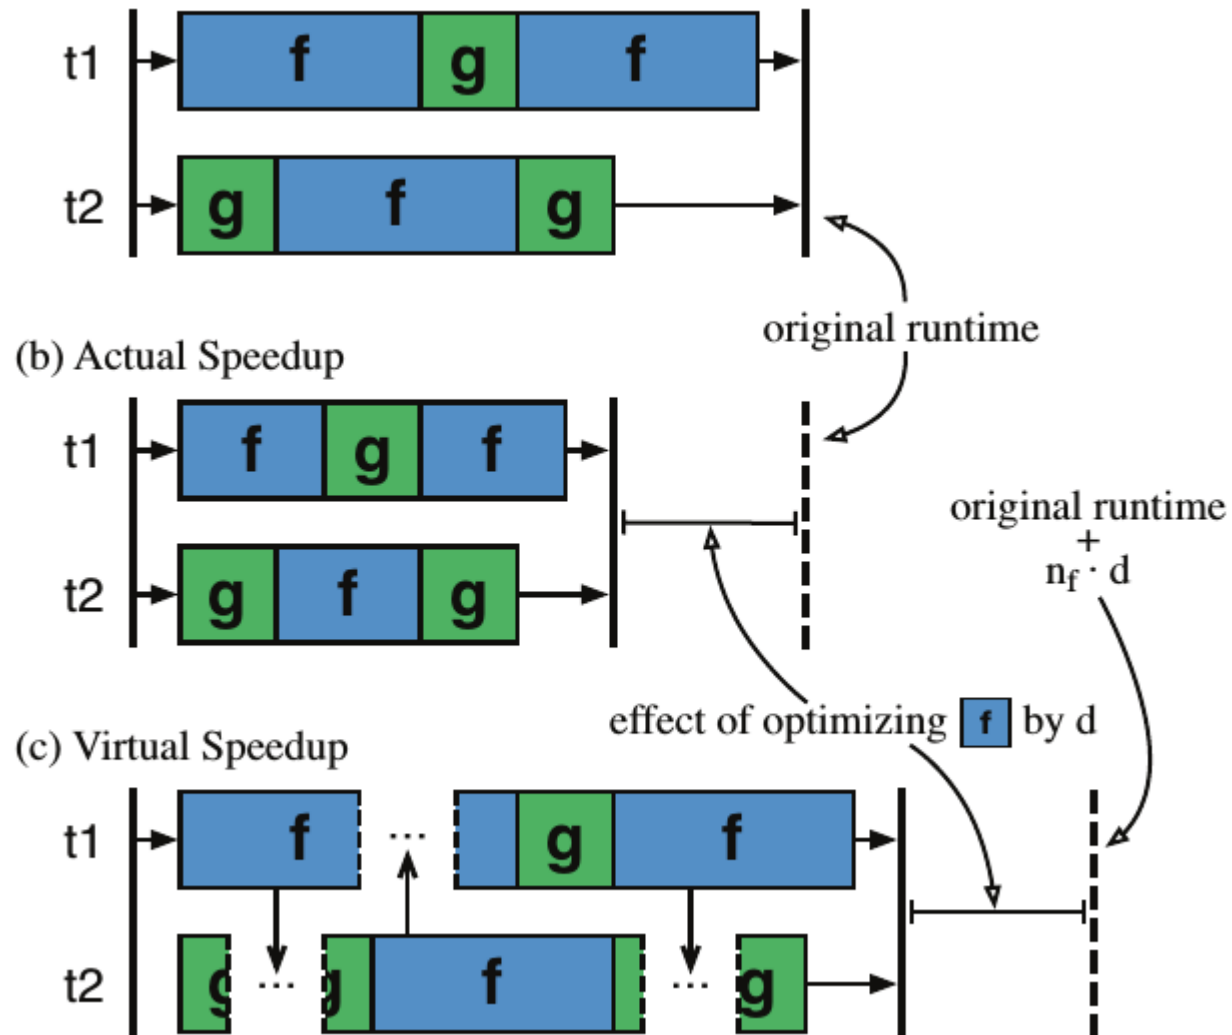


Idea

Perform *performance experiments* during program execution that measure how fast a program gets if you slow down a specific line of code



Idea



Implementation

1) Startup

coz run --- <program> <args>

2) Experiment Initialization

randomly choose source line and speedup (5% steps)

3) Apply virtual speedup

sample each thread every ms, apply virtual speedup when selected source line runs

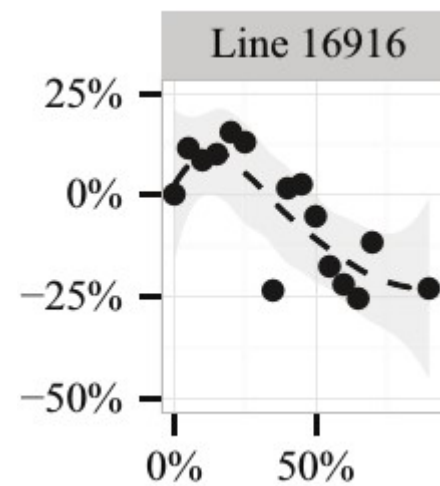
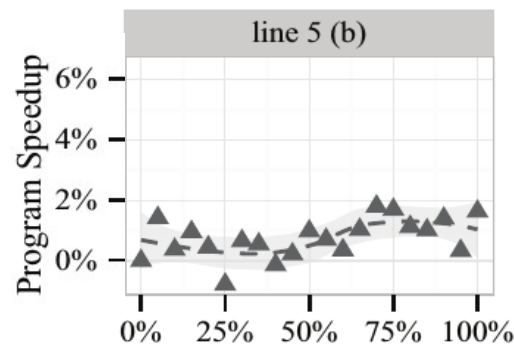

4) Ending virtual experiment

after predetermined time, at least five visits to selected line, else double time

5) Produce profile

analyze all logged results, calculate speedups

6) Interpret casual profiles



Challenges

Blocking: Who has to wait when?

- I/O simple (delay after io finishes)
- Synchronization primitives: hard, don't punish twice

Potentially *unblocking* calls

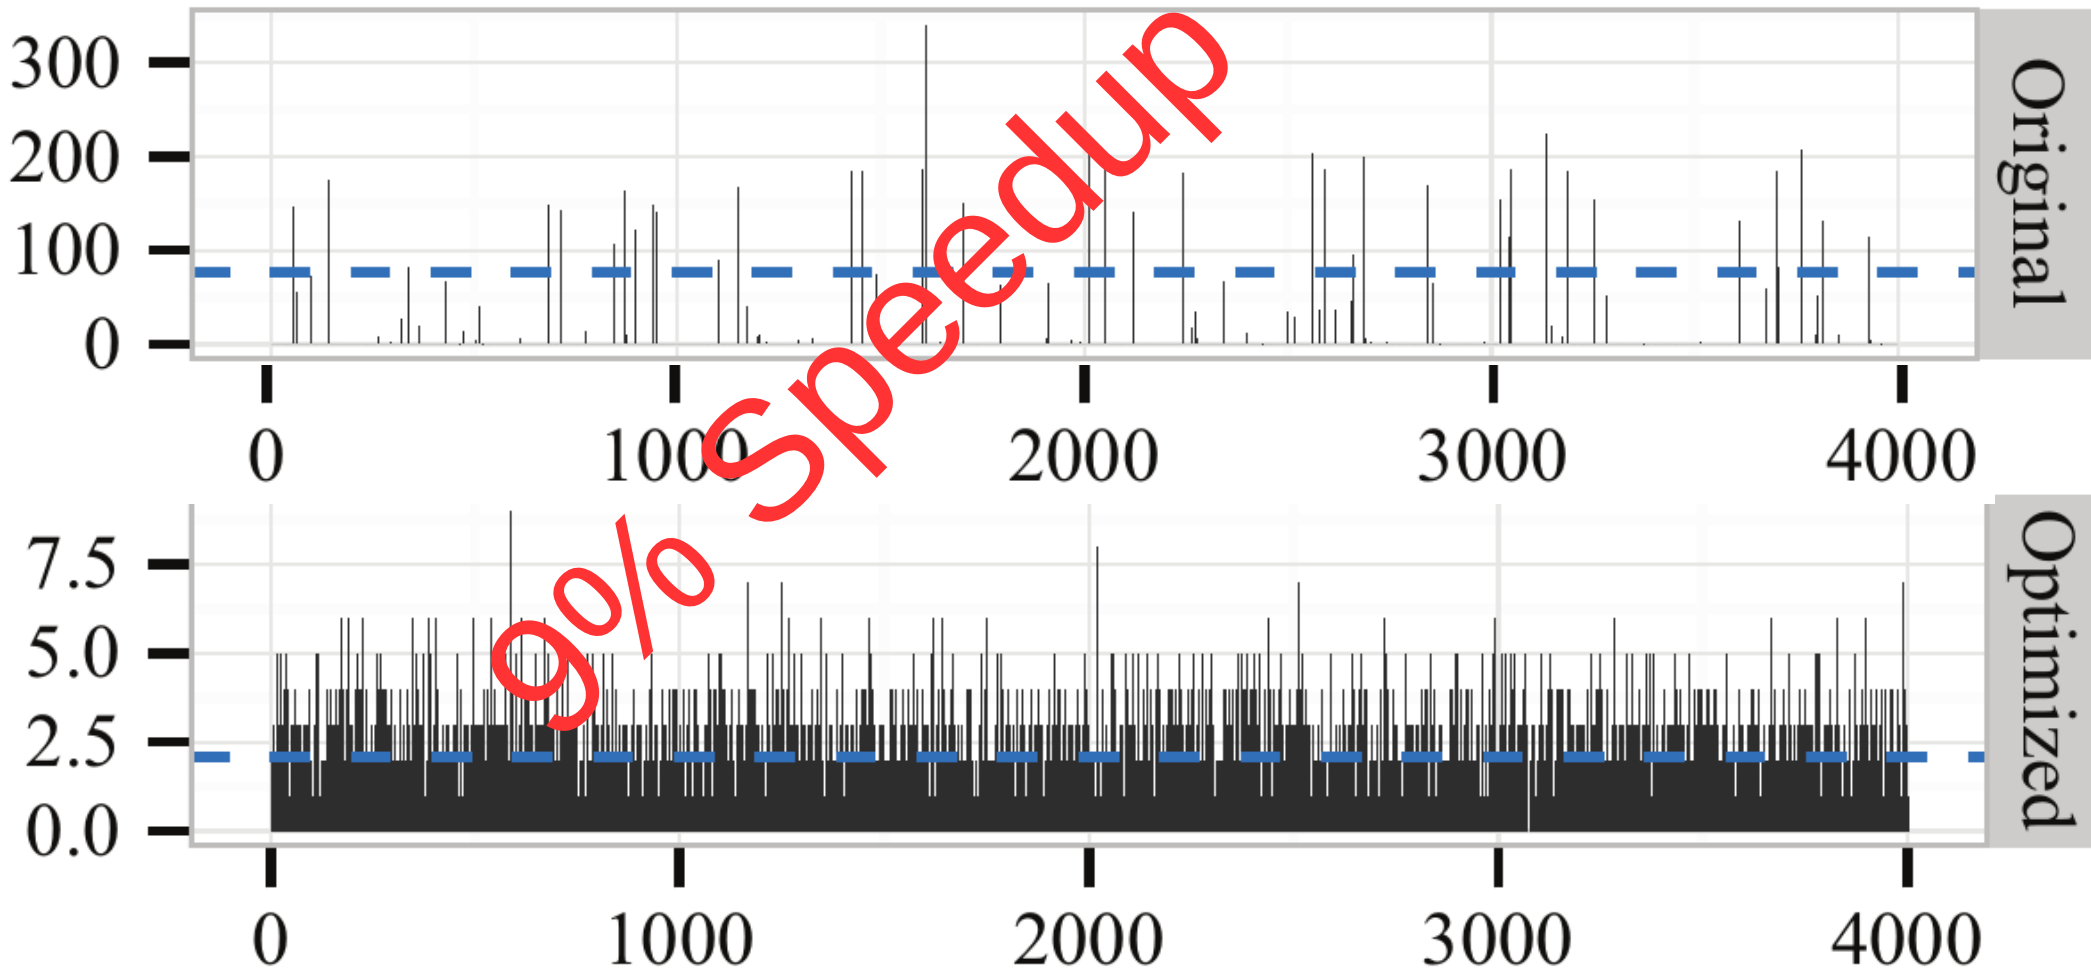
<code>pthread_mutex_unlock</code>	unlock a mutex
<code>pthread_cond_signal</code>	wake one waiter on a c.v.
<code>pthread_cond_broadcast</code>	wake all waiters on c.v.
<code>pthread_barrier_wait</code>	wait at a barrier
<code>pthread_kill</code>	send signal to a thread
<code>pthread_exit</code>	terminate this thread

Potentially *blocking* calls

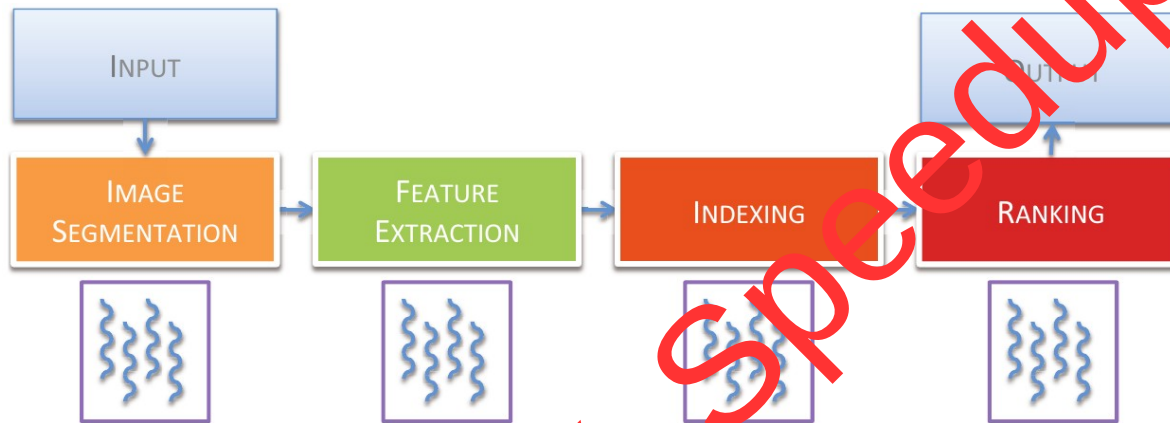
<code>pthread_mutex_lock</code>	lock a mutex
<code>pthread_cond_wait</code>	wait on a condition variable
<code>pthread_barrier_wait</code>	wait at a barrier
<code>pthread_join</code>	wait for a thread to complete
<code>sigwait</code>	wait for a signal
<code>sigwaitinfo</code>	wait for a signal
<code>sigtimedwait</code>	wait for a signal (with timeout)
<code>sigsuspend</code>	wait for a signal

Results (dedup)


While toop in hashtable_search function



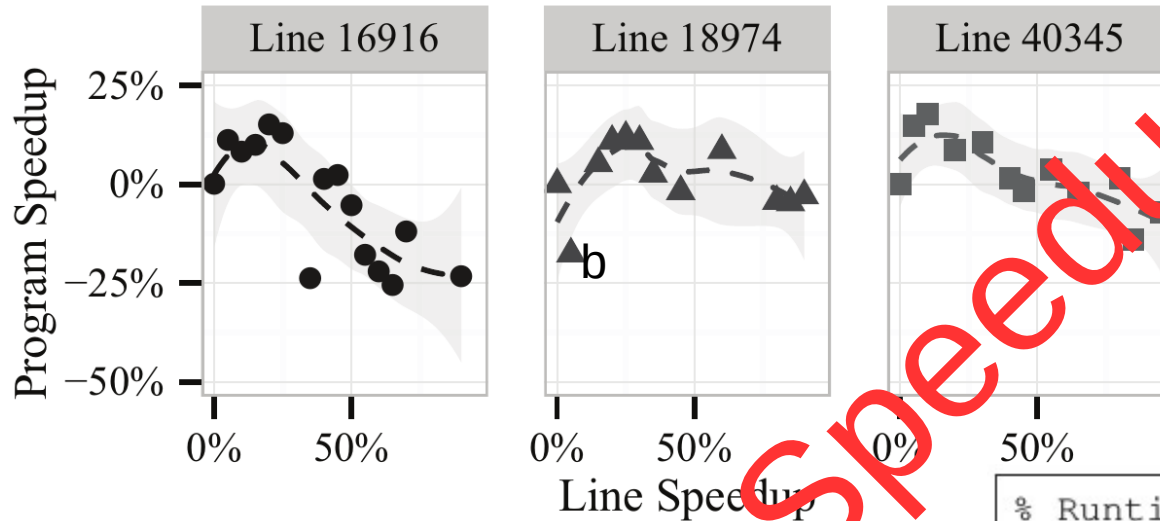
Results (ferret)



Improvement opportunities in three of the four threads



Results (SQLite)




Remove unnecessary
'vtable' calls


% Runtime	Symbol
85.55%	<code>_raw_spin_lock</code>
1.76%	<code>x86_pmu_enable_all</code>
...	30 lines ...
0.10%	<code>rcu_irq_enter</code>
0.09%	<code>sqlite3MemSize</code>
0.09%	<code>source_load</code>
...	26 lines ...
0.03%	<code>__queue_work</code>
0.03%	<code>pcache1Fetch</code>
0.03%	<code>kmem_cache_free</code>
0.03%	<code>update_cfs_rq_blocked_load</code>
0.03%	<code>pthreadMutexLeave</code>
0.03%	<code>sqlite3MemMalloc</code>

Results (more)

Benchmark	Cause	Improvement
fluidanimate	Custom barrier	37.5%
streamcluster	Custom barrier	68%
memcached	'false' Lock sharing	9%
blackscholes	Common subexpressions	2.6%
swaptions	Inefficient array handling	15%



Efficiency



Discussion

- Cool Idea
- How long do you have to sample for reasonable coverage of 148k lines of sqlite? (1ms, each 5 times, so 500s? How is the coverage of this?)