Precise Low-Overhead Performance Measurements

Lode Vandevenne and Jan Wassenberg

Highwayhash Profiler

- https://github.com/google/highwayhash
- profiler.h
- tsc_timer.h

Google Research Europe

Rapidly accelerating use of Deep Learning at Google

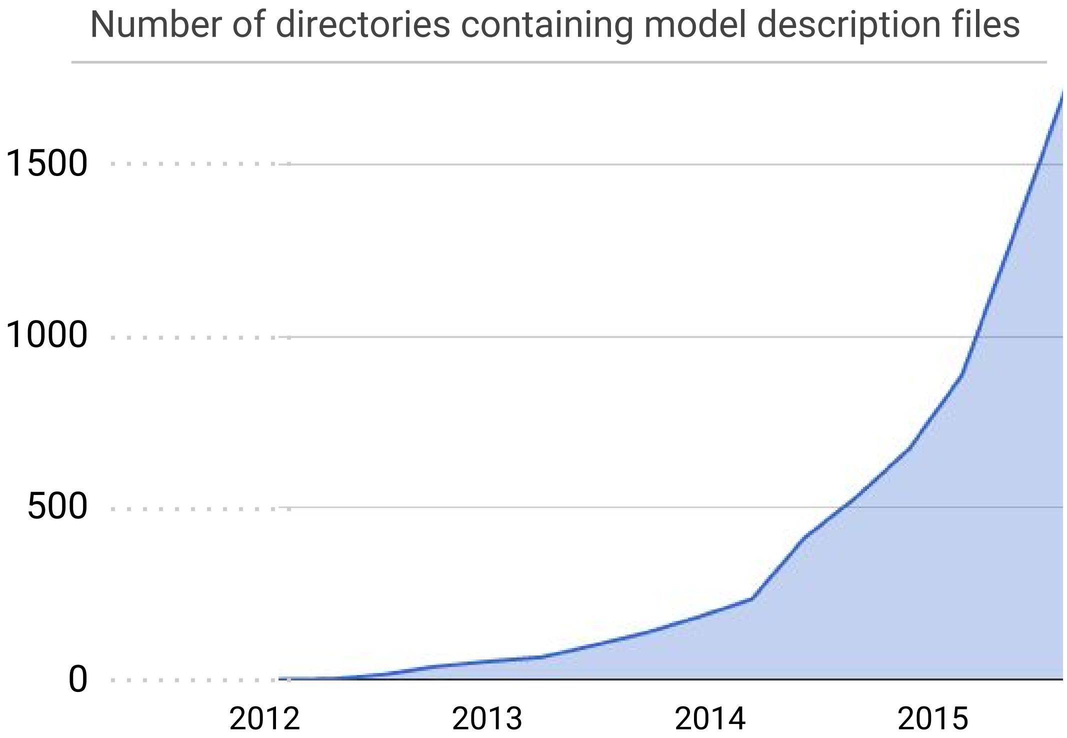

Used across products:

Android
Apps
Gmail
Maps
Photos
Speech
Search
Translation
YouTube

and many others ...

Highwayhash Profiler

- Software profiler for C++
- Flat or call-graph output (cumulative elapsed time per zone)
- Instrumentation based (not statistical like gprof or simulating like cachegrind)
- Manual: PROFILER_ZONE at each scope to measure, PROFILER_PRINT_RESULTS() at end of program

Why another?

- Realistic
- Fast
- High resolution
- Multi-core
- Measure arbitrary zones

Implementation

- Event-driven tracing
- Write combining
- Fences
- Time Stamp Counter
- Delta-coding string literals

Event-driven tracing

Write combining

Fences

Time Stamp Counter

Delta-coding string literals

Results

Event-driven tracing

- Zone entry and exit events
- Small fast event packets (64 bits)
- Call-graph computation deferred to the end
- Self-organizing list

Event-driven tracing

Capture events

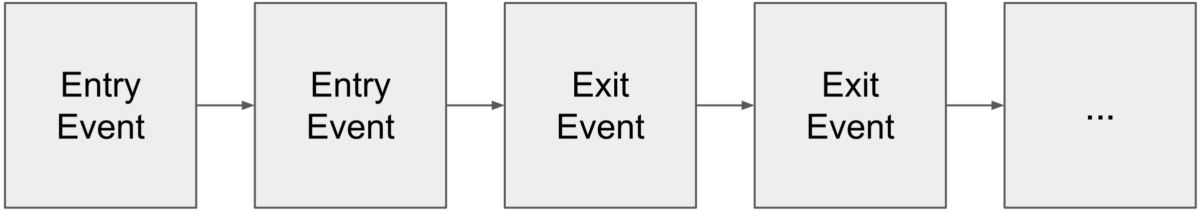

Entry Event

Entry Event

Exit Event

Exit Event

...

Analyze: Stack,
Self-Organizing List

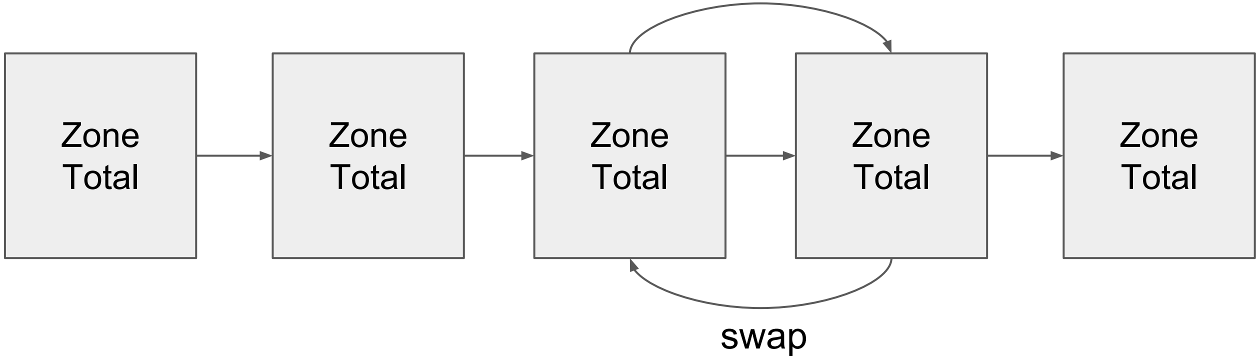

Zone Total

Zone Total

Zone Total

Zone Total

Zone Total

swap

Event-driven tracing

Write combining

Fences

Time Stamp Counter

Delta-coding string literals

Results

Write combining

- When writing events
- Avoid polluting the cache (only 64 bytes used)
- Write-Combine buffer (512-bit)
- Requires aligned pointers to cache line size

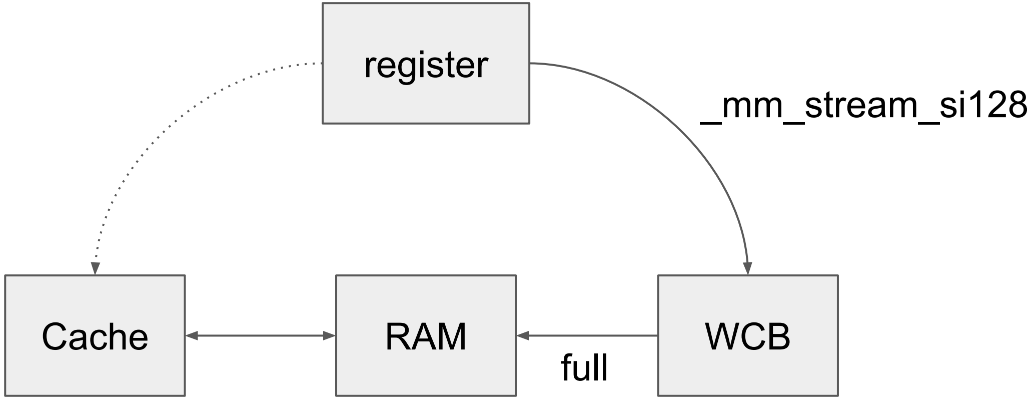

Event-driven tracing

Write combining

Time Stamp Counter

Fences

Delta-coding string literals

Results

Fences (Barriers)

- Avoid instruction reordering

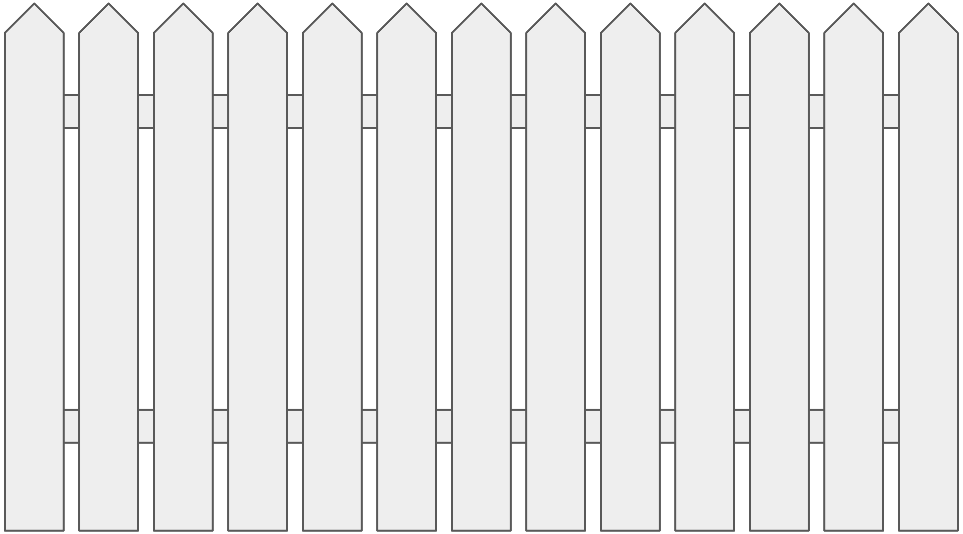

Compiler Fence

- Compiler can reorder code
- Don't let measured code escape!
- Compiler specific keywords, e.g. gcc:
  asm volatile("" : : : "memory")

LFENCE

- CPU can reorder instructions: out of order execution
- Don't let measured code escape!
- Avoid with serializing instruction
- Load Fence: Serializes load operations
- Faster and less variability than CPUID
- Platform specific

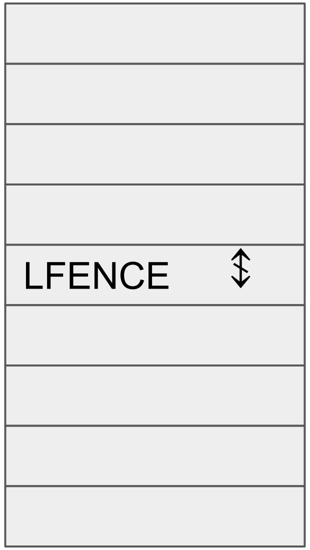

SFENCE

- Store Fence: Serializes store operations
- To flush write-combine buffers: read correctly when printing results

Event-driven tracing

Write combining

Fences

Time Stamp Counter

Delta-coding string literals

Results

Time Stamp Counter

- Precise: fixed at nominal CPU frequency
- RDTSC, RDTSCP
- tsc_timer.h
- Per socket, not per core

Time Stamp Counter

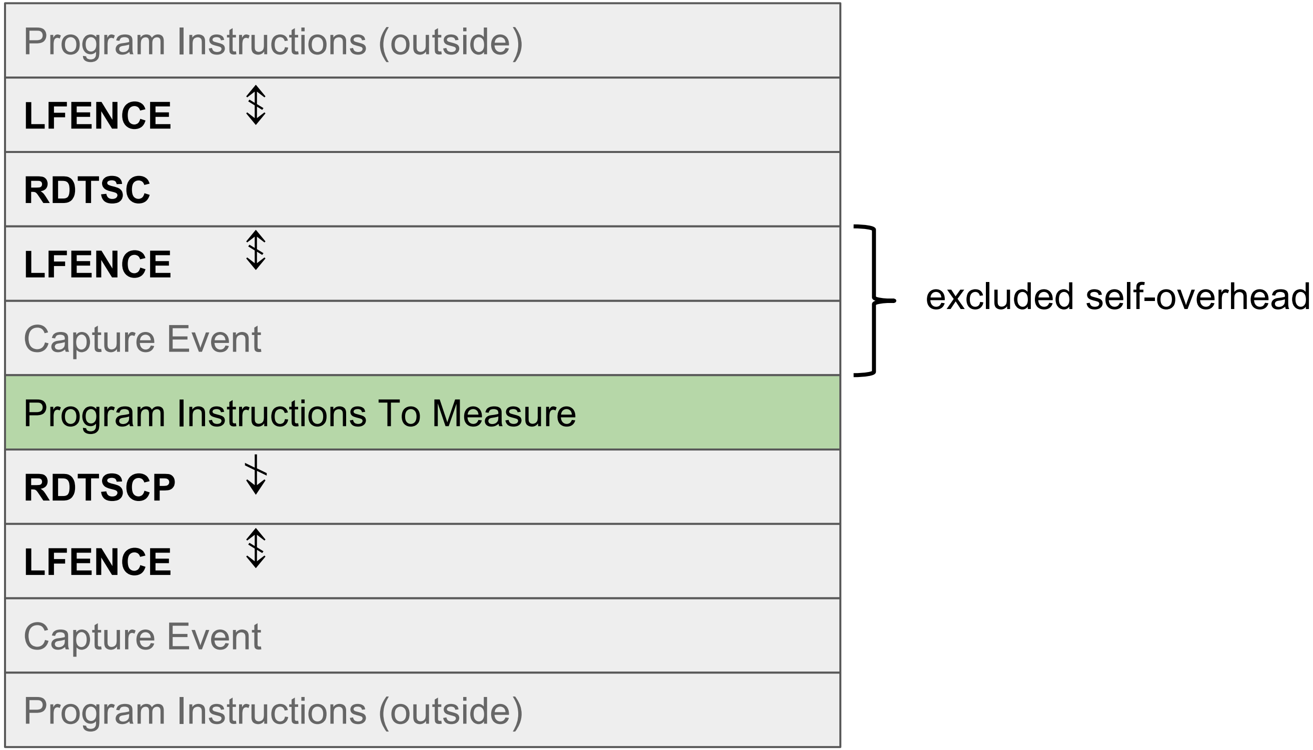

Program Instructions (outside)

LFENCE

RDTSC

LFENCE

Capture Event

Program Instructions To Measure

RDTSCP

LFENCE

Capture Event

Program Instructions (outside)

excluded self-overhead

Event-driven tracing

Write combining

Time Stamp Counter

Fences

Delta-coding string literals

Results

Delta-coding string literals

- Macros using __func__ in zone names
- Need small fast key for zone
- Offset of address of string to a known string literal
- Assume all others are close to it in the binary
- Event: 25 bits offset, 39 bits timestamp

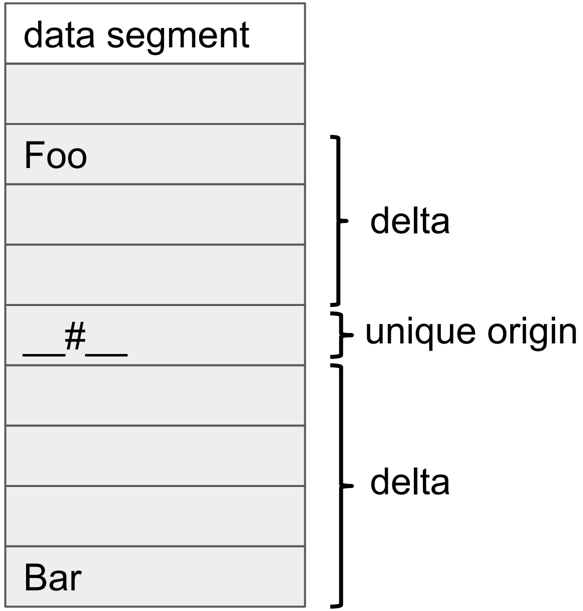

data segment

Foo

__#__

Bar

delta

unique origin

delta

Event-driven tracing

Write combining

Time Stamp Counter

Fences

Delta-coding string literals

Results

Results

```
Mixed:         10000 x     996177 = 3831.453102 ms
RunNoInline:       1 x     136544 = 0.052517 ms
RunInline:         1 x     123894 = 0.047652 ms
```

```
 %   cumulative  self          self   total
time  seconds   seconds calls s/call s/call  name
55.06   2.10     2.10    5001   0.00   0.00  RunInline()
45.10   3.83     1.72    5000   0.00   0.00  Mixed()
 0.00   3.83     0.00       1   0.00   1.72  RunNoInline()
```

Results - Variability

```
ZopfliFindLongestMatch:    2673533 x             166 = 170.947735 ms
ZopfliFindLongestMatch:    2673533 x             215 = 221.913315 ms
ZopfliFindLongestMatch:    2673533 x             171 = 176.444056 ms
ZopfliFindLongestMatch:    2673533 x             174 = 179.001032 ms
ZopfliFindLongestMatch:    2673533 x             161 = 165.676642 ms
```

```
  %   cumulative   self              self     total
 time   seconds   seconds    calls   s/call   s/call  name
 44.45      0.20      0.20  6874417     0.00     0.00  ZopfliFindLongestMatch
 25.00      0.26      0.13  6874417     0.00     0.00  ZopfliFindLongestMatch
 31.38      0.16      0.16  6874417     0.00     0.00  ZopfliFindLongestMatch
 26.54      0.29      0.13  6874417     0.00     0.00  ZopfliFindLongestMatch
 31.92      0.15      0.15  6874417     0.00     0.00  ZopfliFindLongestMatch
```

Summary

- Low-overhead, fast, precise, flexible profiler
- Achieved with several low-level techniques

Questions? Comments?

https://github.com/google/highwayhash

lode@google.com
janwas@google.com

Child Overhead

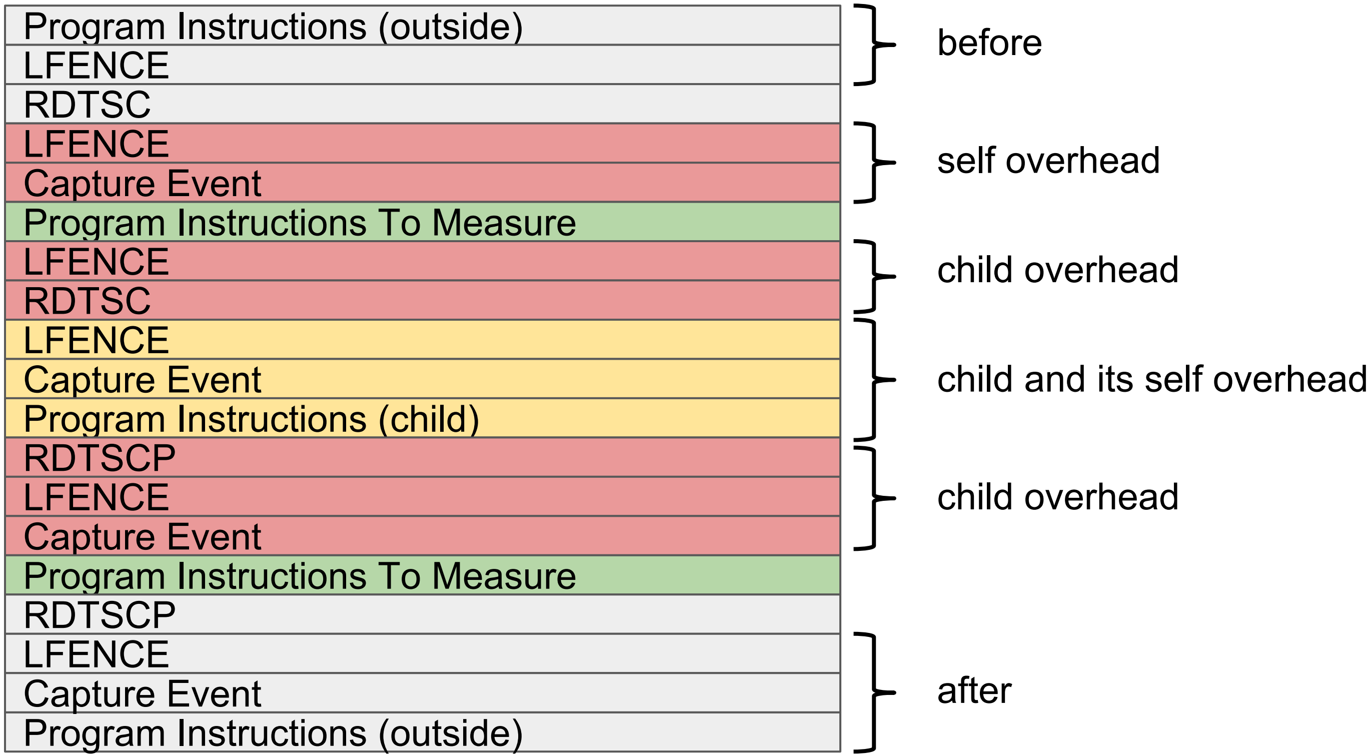

Program Instructions (outside)	before
LFENCE	
RDTSC	
LFENCE	self overhead
Capture Event	
Program Instructions To Measure	
LFENCE	child overhead
RDTSC	
LFENCE	child and its self overhead
Capture Event	
Program Instructions (child)	
RDTSCP	child overhead
LFENCE	
Capture Event	
Program Instructions To Measure	
RDTSCP	
LFENCE	after
Capture Event	
Program Instructions (outside)	