


Brendan Gregg
Intel Fellow, Office of the CTO
Nov 2024

Brendan Gregg
Intel Fellow, Office of the CTO
Nov 2024



Globally Critical Issues in 2024



AI Datacenter Cost

AI Flame Graphs

New advanced analysis tool
available as a preview in the
Intel® Tiber® AI Cloud



CPU Flame Graphs


Prior CPU Profilers

```
# perf record -F 99 -p 13204 -g -- sleep 30
# perf report -n --stdio
[...]
```


#	Overhead	Samples	Command	Shared Object	Symbol
#	20.42%	605	bash	[kernel.kallsyms]	[k] xen_hypcall_xen_version
					--- xen_hypcall_xen_version
					check_events
					--44.13%-- syscall_trace_enter
					tracesys
					--35.58%-- __GI___libc_fcntl
					--65.26%-- do_redirection_internal
					do_redirections
					execute_builtin_or_function
					execute_simple_command
					execute_command_internal
					execute_command
					execute_while_or_until
					execute_while_command
					execute_command_internal
					execute_command
					reader_loop
					main
					__libc_start_main

```
[...]
```


Prior CPU Profilers...full output



...same data, as a Flame Graph



E.g.: \$16M/year compute



CPU Flame Graphs

Shows code
proportional to cost

Used by everyone

(developers, ops, sales support,
end-users)

Interactive

CPU Flame Graph Adoption

Implementations: **>100**


Related open source projects: **>500**

Commercial product adoptions: **>30**

New startups: **6** (so far)

Startups exits: **2** (so far)

Included in Intel® VTune™ Profiler and is the basis for Intel gProfiler.




AI Flame Graphs

Intel Innovation

World's first easy AI profiler

- A. Code visualization
- B. Low overhead
- C. No setup required

...same as CPU profiling



AI/
GPU

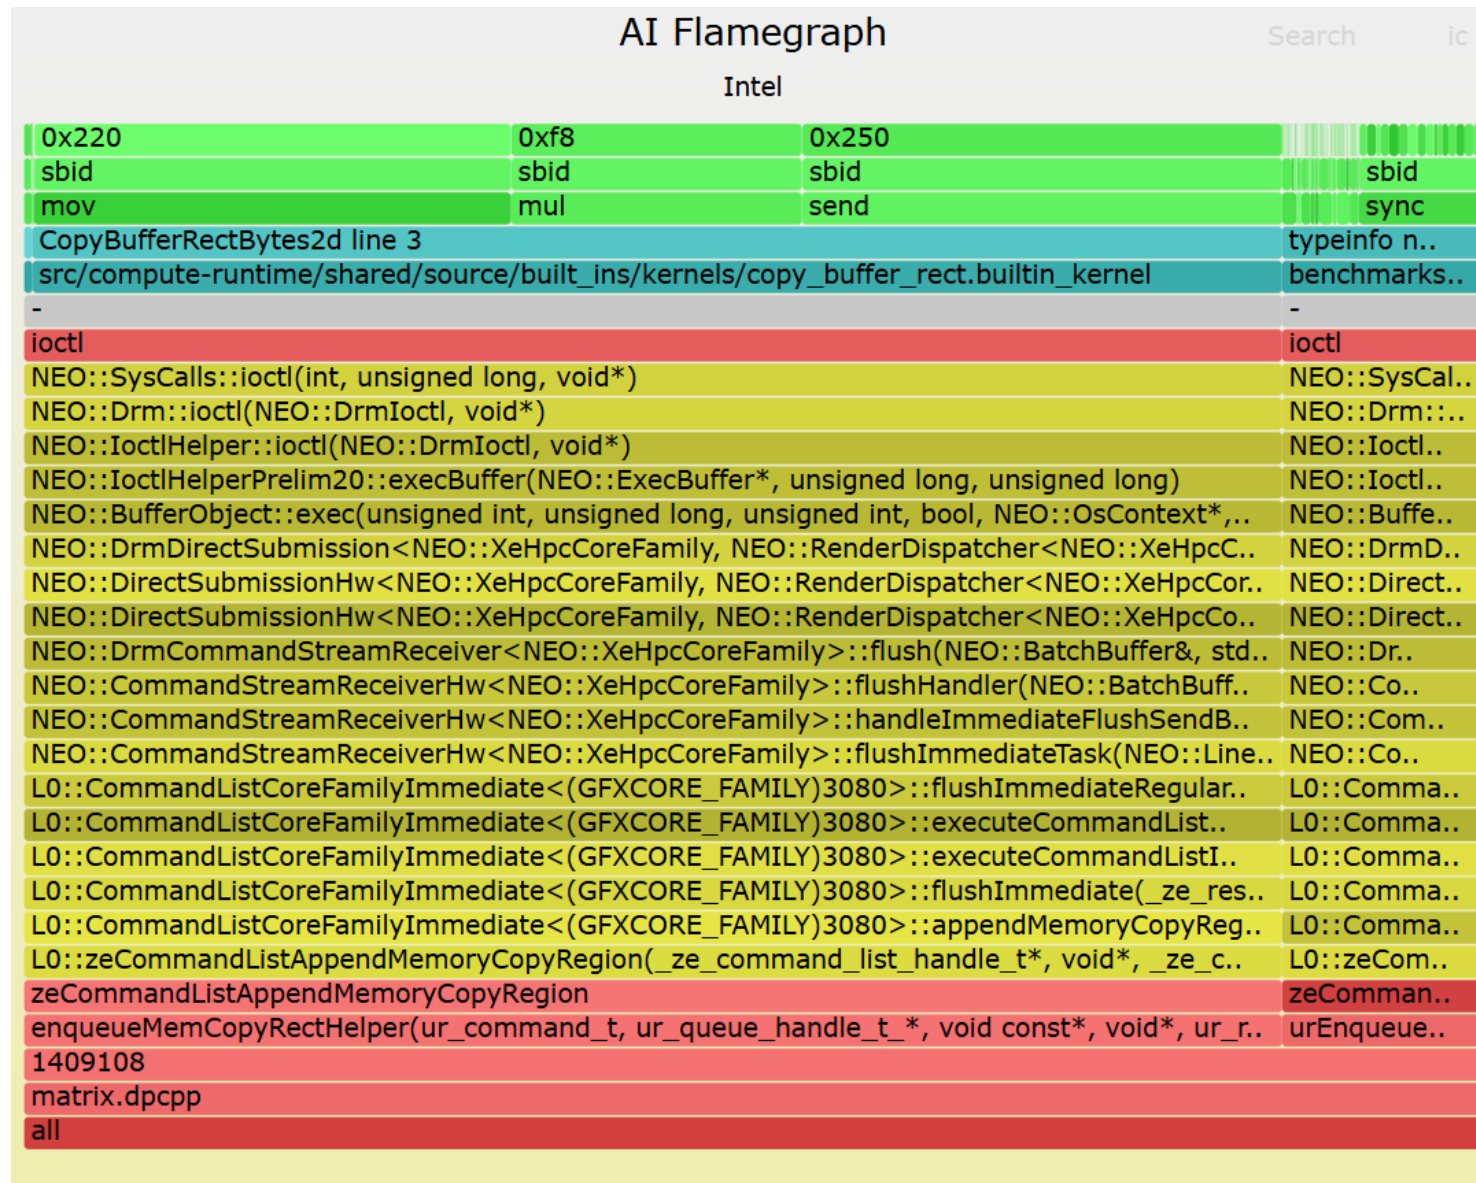
Instruction offset
Stall reason
Instruction mnemonic
Stack/source functions
Source file

CPU call stack

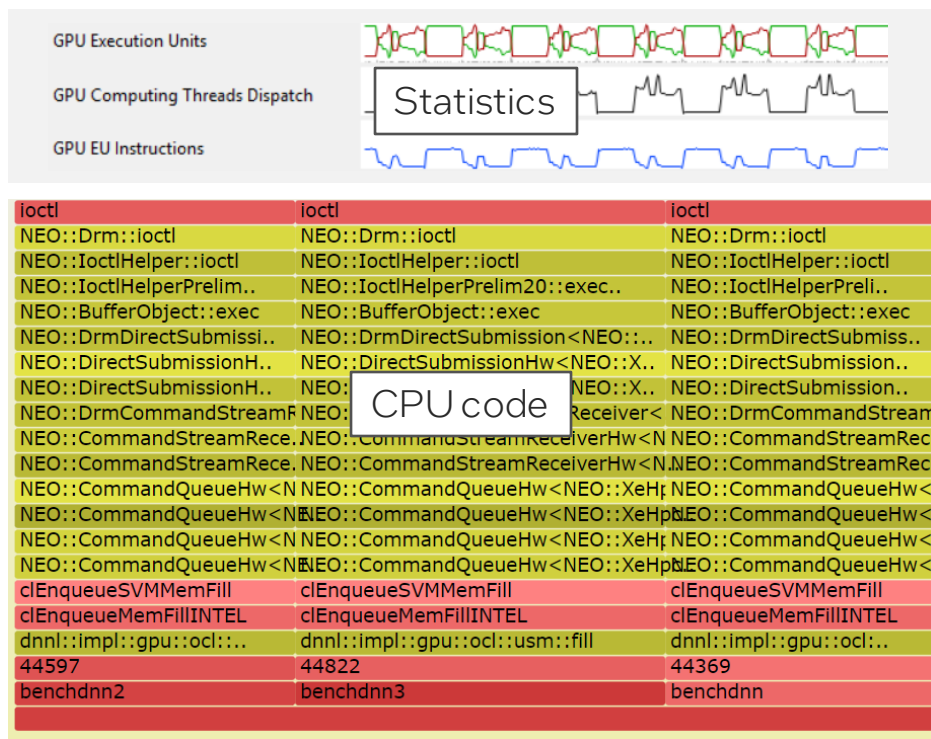
Legend:

GPU HW
GPU source
OS Kernel
CPU C++
CPU C
Python

PID
Program



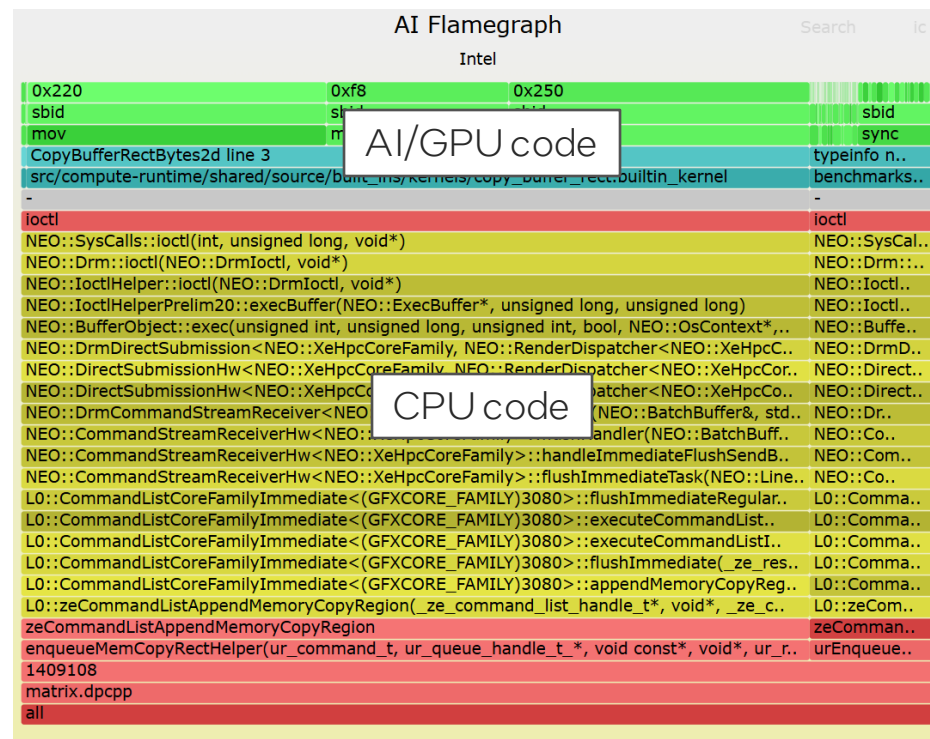
Prior AI Profilers



No code profiling on the AI accelerator or GPU

(Or partial code visibility via static instrumentation, or full accelerator code profiling but via binary instrumentation and prohibitively high overhead)

Intel AI Profiler



All your code in one visual, proportional to cost

(Statistics still available; E.g., in VTune)

The image displays a PyTorch Flame Graph, a visualization tool used to analyze the performance of PyTorch programs. The graph shows the execution flow of a program, with nodes representing different operations and edges representing the flow of execution. The nodes are color-coded and labeled with their respective operations. The graph is organized into a hierarchical structure, with the main entry point at the top and various sub-operations branching out. The nodes are labeled with their respective operations, such as 'PyTorch::Tensor::add', 'PyTorch::Tensor::mul', 'PyTorch::Tensor::div', etc. The edges represent the flow of execution, showing how the program progresses from one operation to the next. The graph is a detailed representation of the program's execution, showing the various operations and the flow of data between them.

- Linux (OS Kernel)

Intel NEO (Compute Runtime)

- Intel Level Zero API (User-mode driver)

- Intel SYCL (Runtime library)

Intel Extension for PyTorch (Framework)

Profiler is based on Intel EU Stalls, flamegraph.pl, and eBPF

Visualising an Open AI Ecosystem



Intel i915 (GPU kernel driver)

<https://github.com/torvalds/linux/tree/master/drivers/gpu/drm/i915>

Linux (OS Kernel)

<https://github.com/torvalds/linux>

Intel NEO (Compute Runtime)

<https://github.com/intel/compute-runtime>

Intel Level Zero API (User-mode driver)

<https://github.com/oneapi-src/level-zero>

Intel SYCL (Runtime library)

<https://github.com/intel/llvm/tree/sycl/sycl>

Intel Extension for PyTorch (Framework)

<https://github.com/intel/intel-extension-for-pytorch>

Profiler is based on Intel EU Stalls, flamegraph.pl, and eBPF

<https://github.com/brendangregg/FlameGraph>

<https://github.com/torvalds/linux/tree/master/kernel/bpf>




Open source technology

Safe and efficient kernel execution
environment

I pioneered its use for performance
observability and security monitoring




eBPF Observability




All observability tools are open source

eBPF Security Monitoring



My slide from a BSidesSF 2017 talk

Globally Critical Issues in 2024



AI Datacenter Cost

Your device ran into a problem and needs to restart. We're just collecting some error info, and then we'll restart for you.

0% complete



For more information about this issue and possible fixes, visit <https://www.windows.com/stopcode>

If you call a support person, give them this info:

Stop code: PAGE_FAULT_IN_NONPAGED_AREA

What failed: csagent.sys

Security & Reliability



Business outcomes

Bringing AI Everywhere


Enabling AI across every enterprise with world-changing technology that improves the life of every person on the planet.



Business outcome

Accelerate innovation

Put AI to work right now, even on sensitive and regulated data, with enterprise-ready solutions and open, optimised software tools preferred and securely trusted by AI developers.



The Intel logo is centered on a solid blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®).

intel®