

Profiling memory and CPU usage

with Google PerfTools and Valgrind

Brett Viren

May 30, 2025

This is meant to help the DNNROI inference optimization work.

Overview

It is very difficult to understand memory and CPU usage by observing the the job as a **monolith**.

- Eg, by watching the program “`top`”.
- One can **profile** the job to better understand which parts deserve “blame”.

How do you “profile” a job?

- Run the job with some **profiling tool**.
- The tool generates some output.
- Run a provided analysis tool to make a visualization.
- Examine the visualization to gain understanding.

Google’s PerfTools and Valgrind

These are both simple, commonly available and can be very useful.

Run Google PerfTools

For shorthand:

```
$ export PROG=/path/to/your/program
```

CPU

```
$ export CPUPROFILE=cpu.prof  
$ LD_PRELOAD=/usr/lib/x86_64-linux-gnu/libprofiler.so.0 $PROG [args]  
$ ls -l cpu.prof*
```

Memory

```
$ export HEAPPROFILE=mem.prof # required  
$ LD_PRELOAD=/usr/lib/x86_64-linux-gnu/libtcmalloc.so.4 $PROG  
$ ls -l mem.prof*
```

Visualize PerfTools results

Here we use the **cpu** file. The **mem** file works identically but is named like `mem.prof.heap.XXXX`.

Basic graph generation

```
$ google-pprof --pdf $PROGRAM $CPUPROFILE > cpu.pdf
```

Focus on a particular function:

```
$ google-pprof --focus 'as_pctree' --pdf $PROGRAM $CPUPROFILE > cpu.pdf  
$ google-pprof --help
```



More ways to process see:

```
$ google-pprof --help
```

Example graphs - CPU and Memory usage

- Nodes sized by how much resource (cpu/mem) they use.
 - ▶ Numbers give used resource directly by the function or by functions the function calls.
- Edges show how much resource went for each (sub) function call.

Tips: Use a PDF viewer that allows easy zoom/pan and play with filtering and other google-pprof options.



How Valgrind works

Valgrind works similarly to Google Perftools. It has a number of “tools”

- `callgrind` makes snapshots of the call stack to make and analyze a call graph.

- `massif` checks heap memory by overriding `malloc` etc and taking snapshots. Generates **profile file** for later visualization.

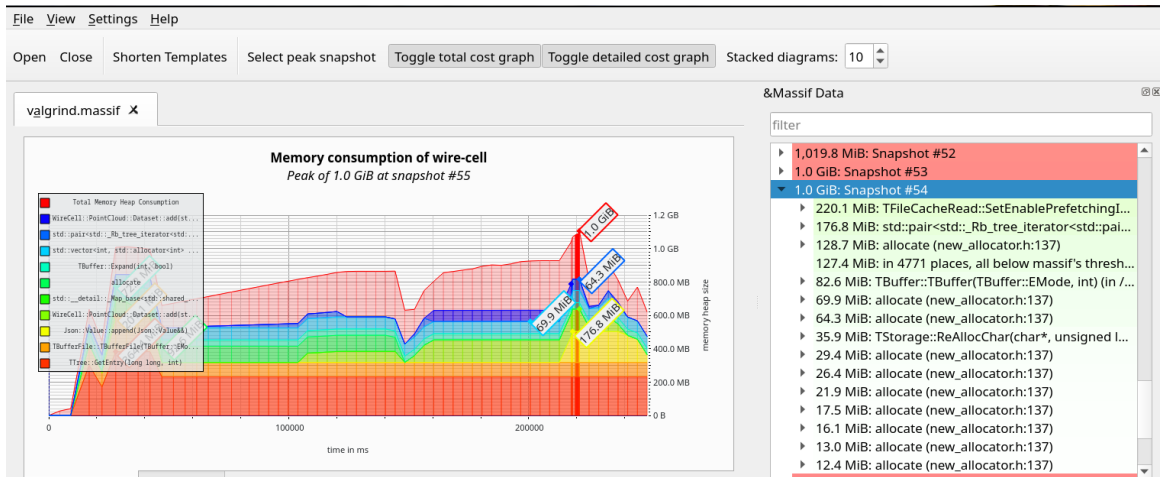
- `memcheck` looks for memory leaks and other errors. Generates **summary report** and error messages.

General usage:

```
valgrind --tool=TOOLNAME [tool options] your_program [program args]
```

Visualize massif results

```
$ valgrind --tool=massif --time-unit=ms --detailed-freq=1 --massif-out-file=v.  
myprogram [my program args]  
$ massif-visualizer valgrind.massif
```



Going further

GPU profiling has its own tools, usually specific to the hardware.

- NVIDIA: Nsight, nvprof, nvpp (nvidia-smi gives a “top” for the GPU)

Various profiling tools work at Linux kernel level.

- perf, eBPF, ...

FIN

What is a call graph

Assume a program like:

```
void c() { c1_work(); c2_work() }  
void b() { c(); b_work() }  
void a() { b(); a_work() }  
void main() { a(); }
```

A snapshot looks at CPU/Memory usage in any instance and records the **call stack**.

Some example stack snapshots:

- main() -> a() -> b() -> c() -> c1_work()
- main() -> a() -> b() -> b_work()
- main() -> a() -> a_work()

Google Perftools simply book keeps how much CPU and Memory is being used on each snapshot.