


IRTG Lecture Week Bergen 2009

Programming of Multicore Systems

-

Programming Tools: Profiling



- **Often programs are too slow**
 - **Too much processing time taken up by parts of code**
- **These processing intensive parts have to be identified**
- **Profilers are programs to help in that task**

- **Profilers allow to**
 - **Determine call graph hierarchy of a program**
 - Which function was called by another function
 - By which other function was a function called
 - How often a function called another function
 - **See how much processing time was used for (almost) each source line**


- **What is profiling**
 - **Profiling, is the investigation of a program's behavior using information gathered as the program executes (i.e. a form of dynamic program analysis).**
 - (Paraphrased from Wikipedia)
- **What is a profiler**
 - **A profiler is a performance analysis tool that measures the behavior of a program as it executes, particularly the frequency and duration of function calls.**
 - (Paraphrased from Wikipedia)
 - **A profiler can be used to show you where your programs take up the most time.**
 - (Paraphrased from me)

Find Your Inefficient Programming For You

**But it can show
you likely locations
of bottlenecks**

- Profilers typically require same information about program as for a debugger
 - Where variables and functions are located in memory
 - Link between generated object code and original source code

```
movl    $0, %eax
subl    %eax, %esp
subl    $12, %esp
pushl   -4(%ebp)
call    answer
addl    $16, %esp
movl    $0, %eax
```



```
answer( question );
```

- Called Debugging Information
 - Provided by gcc/g++ with `-ggdb3` switch

- **Profiling tool covered here:**
 - valgrind
- **Powerful instrumentation framework**
- **Multiple tools available**
 - **Memory checker to detects memory leaks and incorrect accesses (e.g. array boundaries)**
 - **Cache profiler to determine cache hits/misses**
 - **Call graph instruction based profiler to determine execution times for code**
 - Covered here
 - **Race condition detector for multi-threaded programs**
 - ...

- **Version 3.3.0 installed in HLT cluster**
 - `/opt/HLT/tools/valgrind/valgrind-3.3.0`
 - **Script to set environment variables provided**
 - `. /opt/HLT/tools/valgrind/valgrind-3.3.0/bin/setenv.sh`
 - (Has to be sourced)
 - **Call graph profiler called as valgrind tool callgrind**
 - `valgrind --tool=callgrind`

Calling Valgrind



- `valgrind --time-stamp=yes --log-file=profiler-test-%p-`date +%Y%m%d.%H%M%S`-valgrind.log \`
`--tool=callgrind --callgrind-out-file=profiler-\`
`test-%p-`date +%Y%m%d.%H%M%S`-callgrind.out \`
`examined-program program-arguments`
 - `--time-stamp=yes`: **Include timestamps in log files**
 - `--log-file=profiler-test-%p-`date +%Y%m%d.%H%M%S`-valgrind.log`: **Name for log file**
 - `--tool=callgrind`: **Use the callgrind tool of valgrind**
 - `--callgrind-out-file=profiler-test-%p-`date +%Y%m%d.%H%M%S`-callgrind.out`: **Name for output measurement file**
 - `examined-program program-arguments`: **Name of executable to be run with arguments**
 - **Output file names include PID (`%p`) and current timestamp**

- **Valgrind internally emulates CPU**
 - (Can use arbitrary cache hierarchies for cache profiling if needed)
 - Can trap/profile essentially any event
 - Runs significantly slower than in non-profiling mode
 - Factor 50 observed
 - No special compilation information necessary

- **In text form:**

- `callgrind_annotate profiler-test-<PID>-<YYYYMMDD.HHMMSS>-callgrind.out`
 - Replace <PID> by PID, <YYYYMMDD.HHMMSS> by date in filename
 - Gives list of functions sorted by exclusive CPU time
 - Not including time of called functions
- `callgrind_annotate -inclusive=yes profiler-test-<PID>-<YYYYMMDD.HHMMSS>-callgrind.out`
 - Gives list of functions sorted by inclusive CPU time
 - Does include time of functions called by listed function

- **More intuitive using GUI**
 - `kcachegrind`
- **Provides exclusive & inclusive time**
- **Call graph hierarchy**
- **Direct source code view**
 - **(Relative) Time spent in each line of code**

Kcachegrind GUI



The screenshot displays the KCachegrind GUI interface. On the left, a table lists functions with columns for 'Incl.', 'Self', 'Called', 'Full', and 'Function'. A blue callout bubble labeled 'Inclusive Time' points to the 'Incl.' column, and another labeled 'Exclusive Time' points to the 'Self' column.

On the right, the 'Source' tab shows the C++ code for the function `Count16BitWords`. A blue callout bubble labeled 'Annotated Source' points to the source code.

Below the source code, the 'Call Graph' tab is active, showing a hierarchical tree of function calls. A blue callout bubble labeled 'Call Graph' points to this view. The graph shows `main` calling `Process`, which calls `Count16BitWords`. `Count16BitWords` then branches into several calls to `std::vector<StatData16Bit, std::allocator<StatData16Bit>>::push_back`, `std::vector<StatData16Bit, std::allocator<StatData16Bit>>::begin`, `std::vector<StatData16Bit, std::allocator<StatData16Bit>>::end`, and `std::vector<StatData16Bit, std::allocator<StatData16Bit>>::size`.

At the bottom, the status bar shows the total instruction fetch cost: 7 773 991 105.

- **Need to specify source directory/-ies**
- **Menu *Settings***
 - ***Configure KcacheGrind...***
 - *Dialog Tab Annotations*
 - *Button Add*
- **Function unfortunately unstable in version installed in cluster**

- <http://valgrind.org/>
- <http://valgrind.org/docs/manual/QuickStart.html>
- <http://valgrind.org/docs/manual/manual.html>