Profiling with OProfile and Intel Core 2 performance counters

- How to use OProfile.
- How to profile with performance counters to see instruction parrallelism, cache misses, etc.
- Examples of how to interpret results.

Kristian Nielsen, independent software professional.
Free Software since 1990: Wine, MySQL, McStas, and others.

Why profile?

- CPU-bottlenecked appliations (only).
- "Premature optimization is the root of all evil."

Learn *where* and *what* to optimise

OProfile

- Very good profiler for Linux
- *Full system* profiler
- Unmodified binaries (debug symbols help)
- Both kernel and user-space
- *Performance counters* (cache misses, ...)
- *Statistical profiling*
- Limitation: native code only

```
# apt-get install oprofile
# opcontrol --vmlinux=/boot/vmlinux-dbg
# opcontrol --reset
# opcontrol --start
# opcontrol --stop
$ opreport
$ opreport -l <program>
$ opannotate --source <program>
$ opannotate --assembly <program>
```

Example: LAMP stack

```
CPU: Core 2, speed 2401 MHz (estimated)
Counted CPU_CLK_UNHALTED events (Clock cycles when not halted) with a unit mask of 0x00 (Unhalted
core cycles) count 100000
CPU_CLK_UNHALT...|
  samples|      %|
------------------
   544844 60.0413 libphp5.so
   177137 19.5203 vmlinux-debug-2.6.24-19-generic
    64145  7.0687 libc-2.7.so
    29056  3.2019 mysqld
    15854  1.7471 libmysqlclient.so.15.0.0
    11754  1.2953 apache2
        CPU_CLK_UNHALT...|
          samples|      %|
        ------------------
            11675 99.3279 apache2
               10  0.0851 [vdso] (tgid:25751 range:0x7fff69ffe000-0x7fff6a000000)
                9  0.0766 [vdso] (tgid:26185 range:0x7fff69ffe000-0x7fff6a000000)
                7  0.0596 [vdso] (tgid:26075 range:0x7fff69ffe000-0x7fff6a000000)
                6  0.0510 [vdso] (tgid:25753 range:0x7fff69ffe000-0x7fff6a000000)
                6  0.0510 [vdso] (tgid:25755 range:0x7fff69ffe000-0x7fff6a000000)
                6  0.0510 [vdso] (tgid:25837 range:0x7fff69ffe000-0x7fff6a000000)
                6  0.0510 [vdso] (tgid:26143 range:0x7fff69ffe000-0x7fff6a000000)
                5  0.0425 [vdso] (tgid:25752 range:0x7fff69ffe000-0x7fff6a000000)
                5  0.0425 [vdso] (tgid:26142 range:0x7fff69ffe000-0x7fff6a000000)
                5  0.0425 [vdso] (tgid:26183 range:0x7fff69ffe000-0x7fff6a000000)
                4  0.0340 [vdso] (tgid:25754 range:0x7fff69ffe000-0x7fff6a000000)
                4  0.0340 [vdso] (tgid:26078 range:0x7fff69ffe000-0x7fff6a000000)
                4  0.0340 [vdso] (tgid:26100 range:0x7fff69ffe000-0x7fff6a000000)
                2  0.0170 [vdso] (tgid:26184 range:0x7fff69ffe000-0x7fff6a000000)
    10292  1.1342 libapr-1.so.0.2.11
     9965  1.0981 libpthread-2.7.so
     6105  0.6728 mysql.so
```

Example: LAMP stack

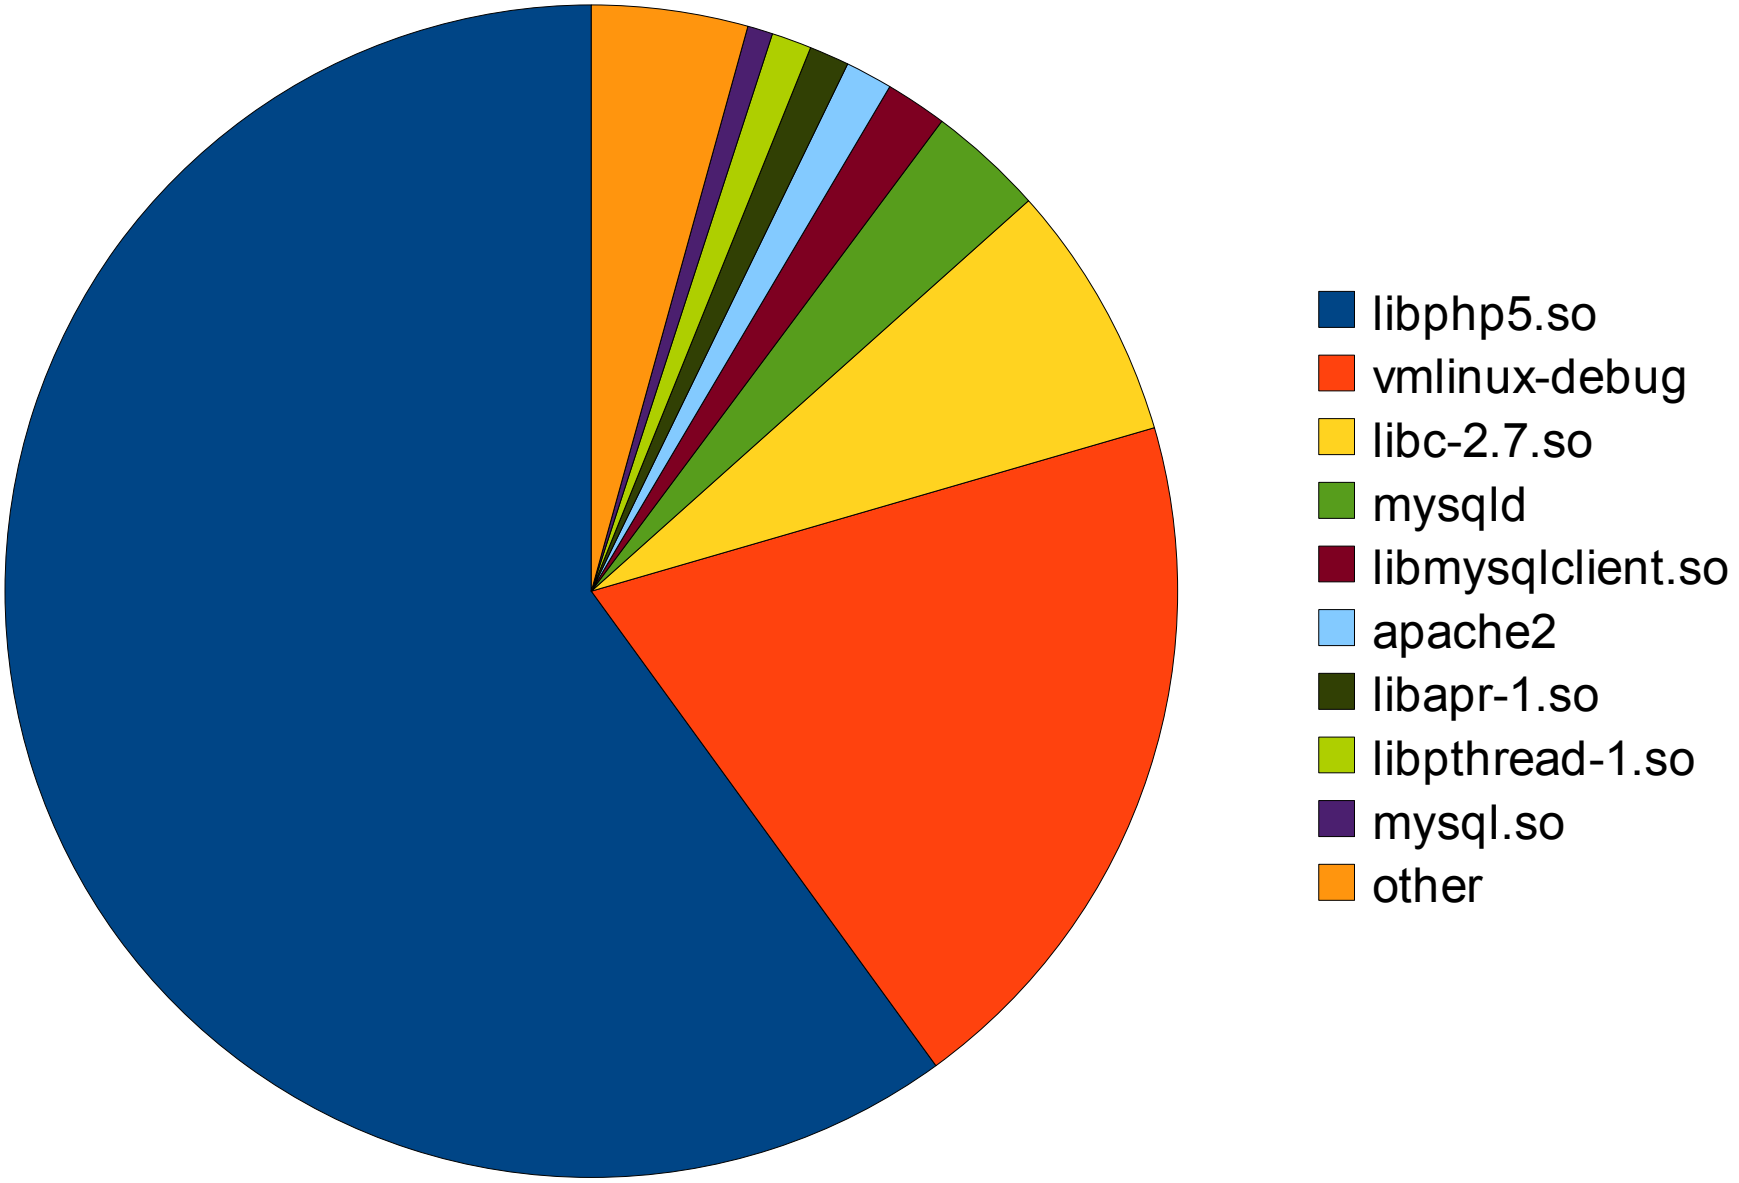

opcontrol --separate=kernel --start

opreport

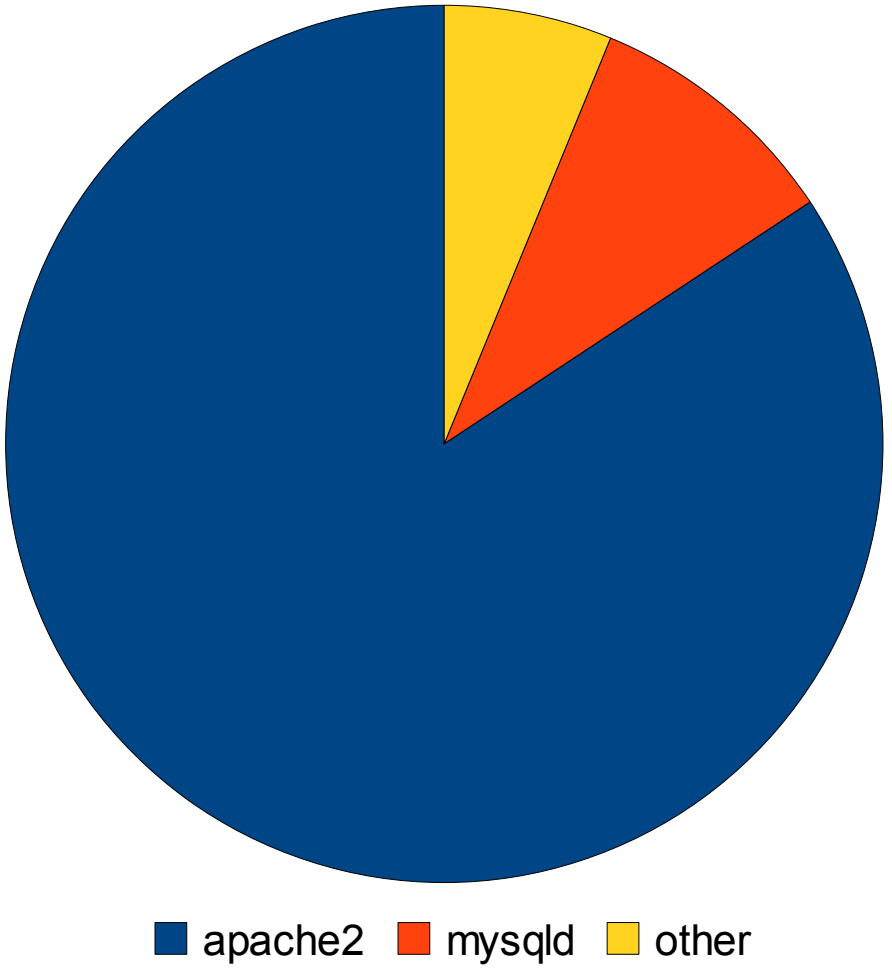

opreport -l apache2

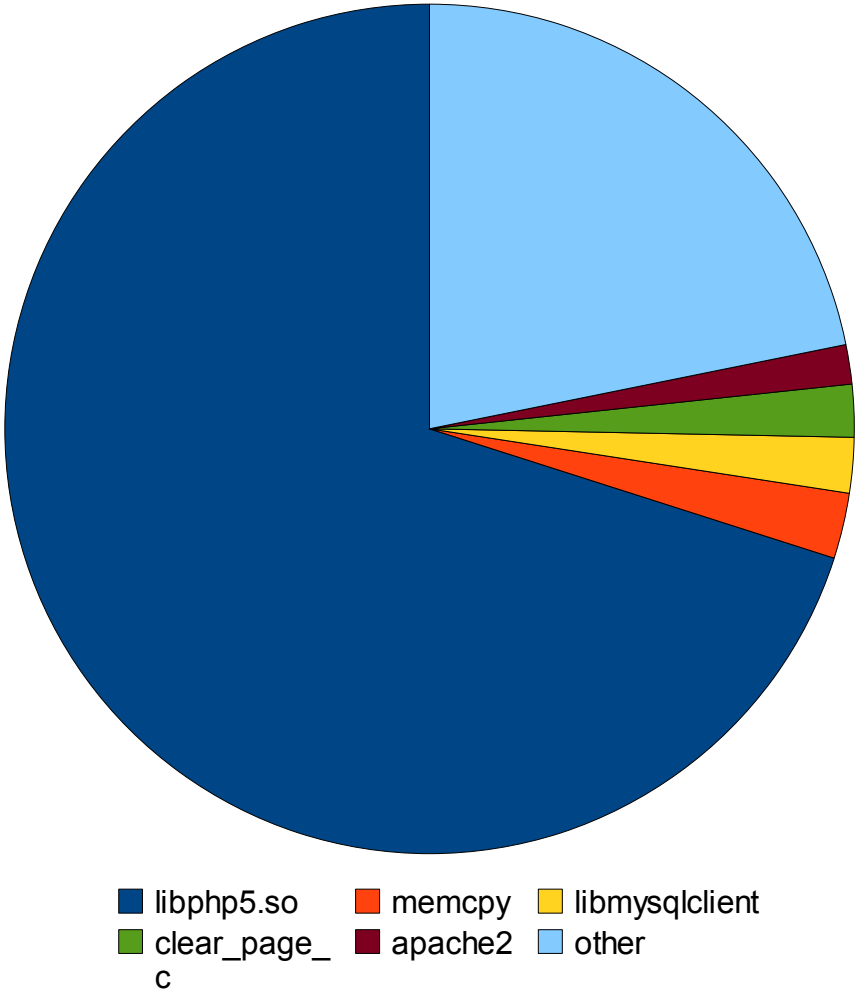

```
opcontrol
    --event=CPU_CLK_UNHALTED:1000000:0:1:1
    --event=UOPS_RETIRED:1000000:0x0f:1:1
```

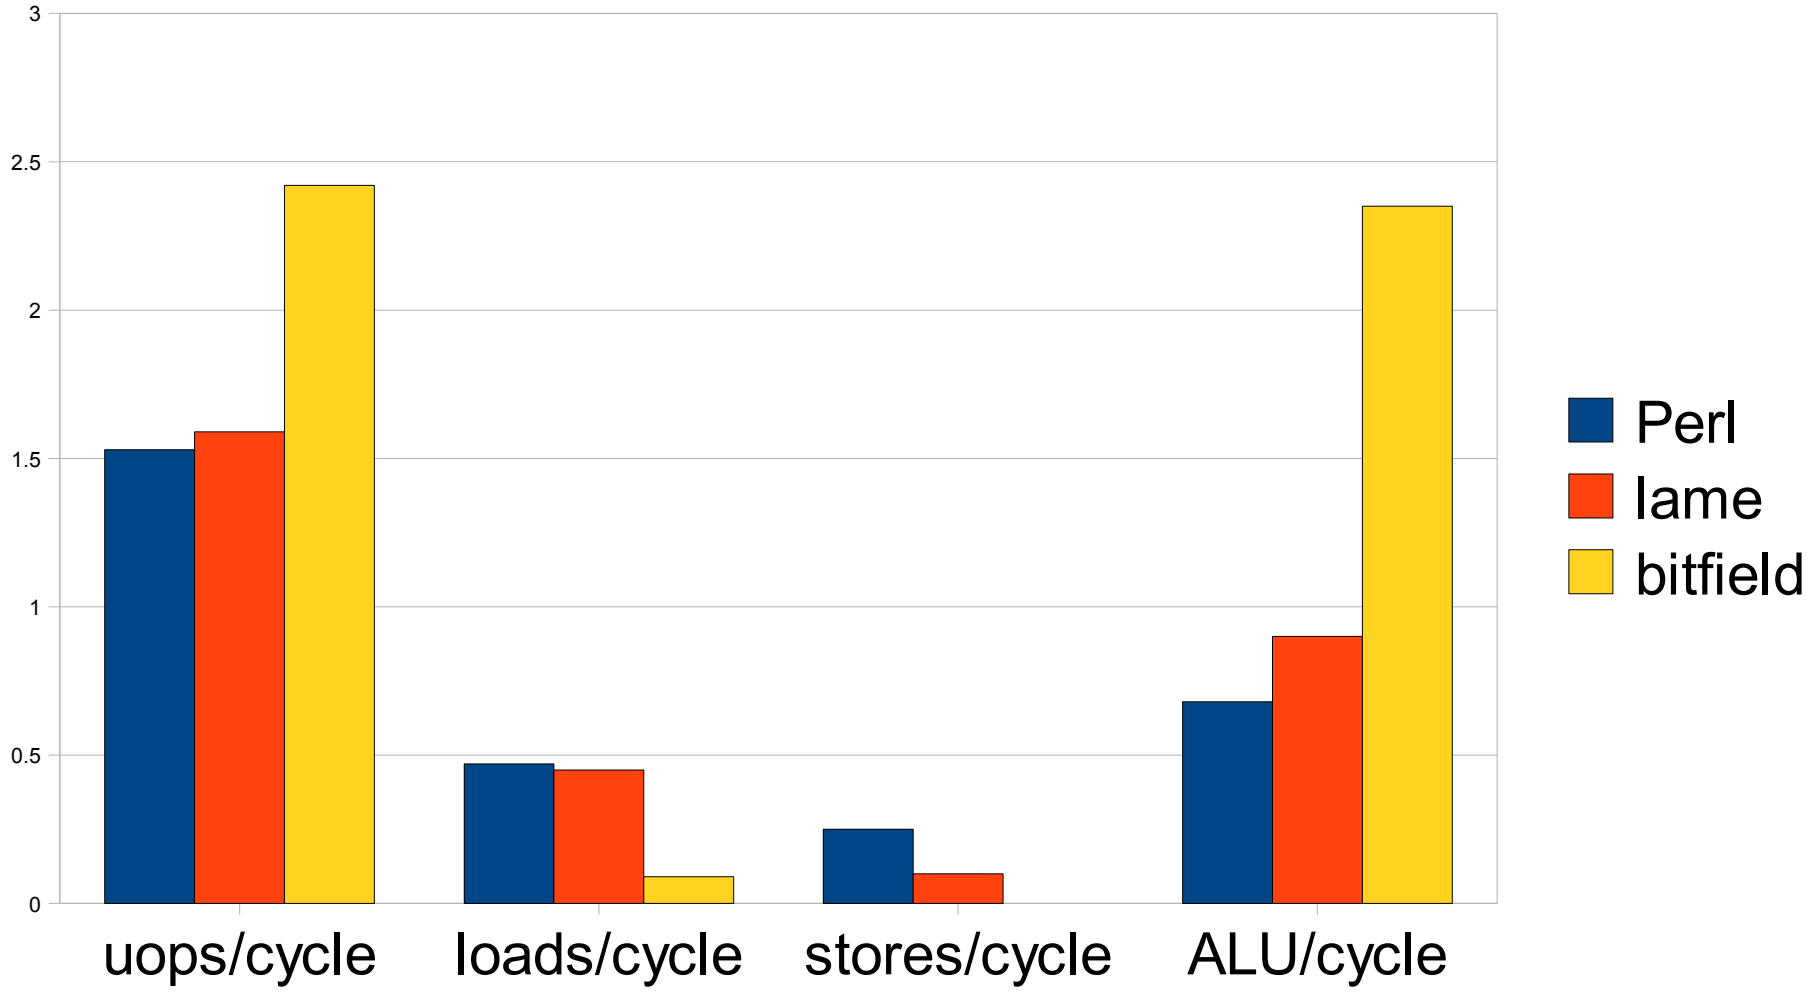

```
opcontrol
  --event=MEM_LOAD_RETIRED:100000:0x01:1:1
```

```
opreport -l
```

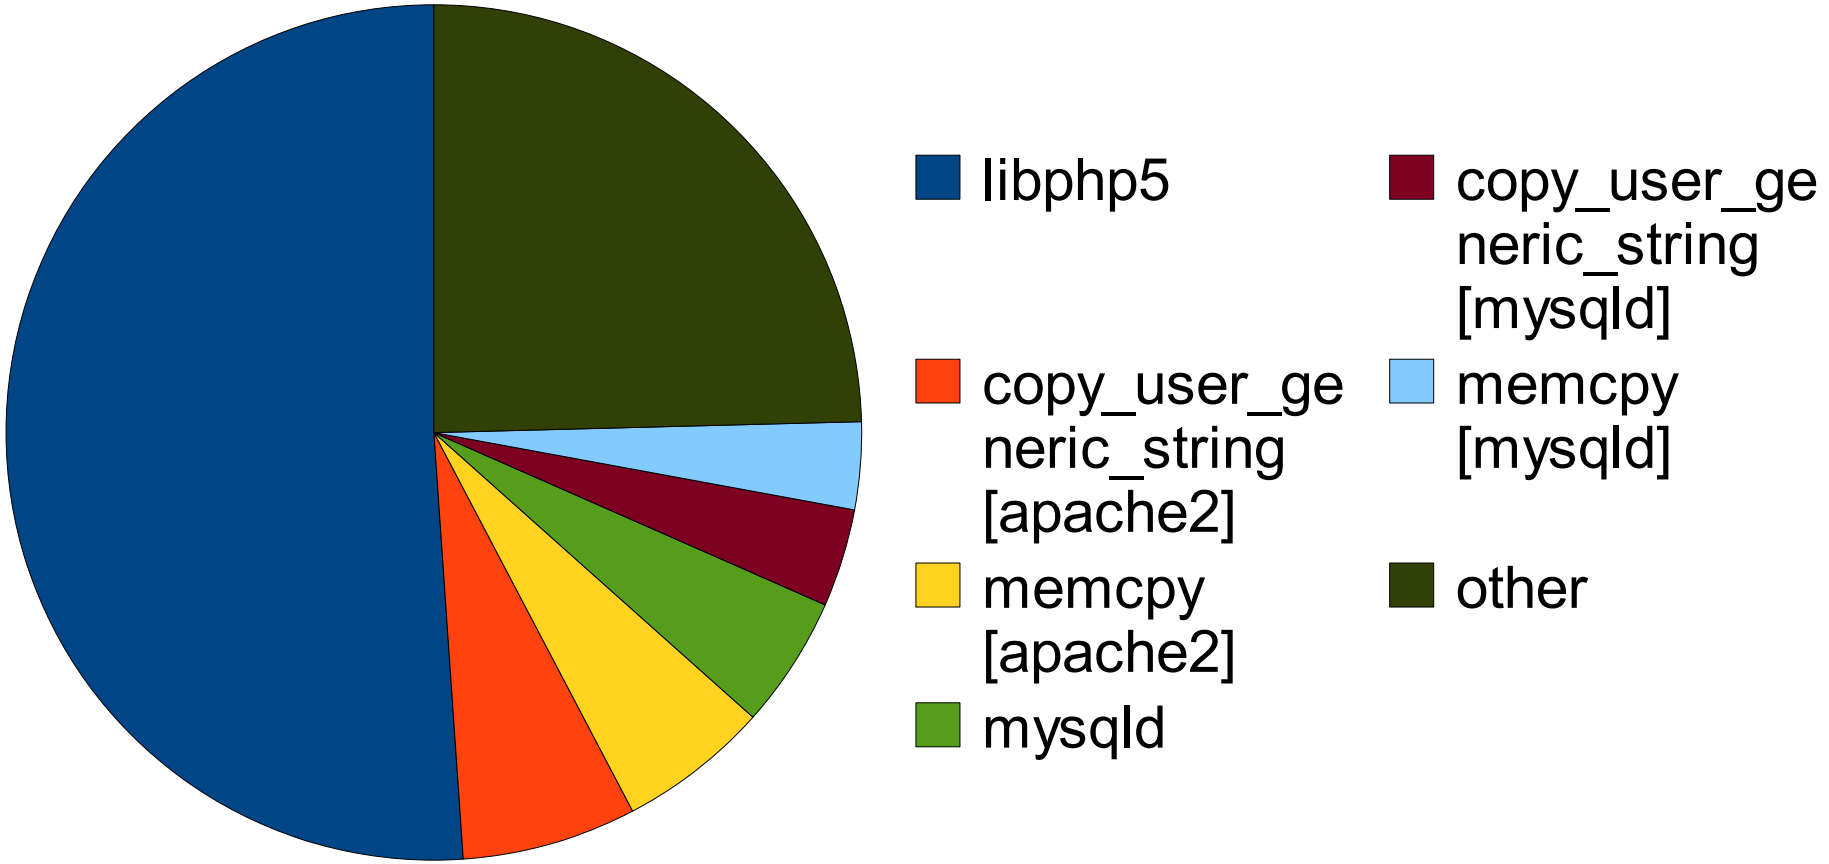

Understanding the results

- Very detailed information, but complicated to learn and use.
- `opcontrol --list-events`
- Intel® 64 and IA-32 Architectures Optimization Reference Manual
- Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3 A/B: System Programming Guide

Conclusion

- Oprofile is easy to install and run.
- Oprofile can give a full profile (kernel + all processes) of an unmodified system.
- Very detailed CPU info available with advanced usage.

http://oprofile.sourceforge.net/

http://kristiannielsen.livejournal.com/