Design and Engineering of Computer Systems
Professor Mythili Vutukuru
Computer Science and Engineering
Indian Institute of Technology, Bombay
Lecture 50 (Week 7, Tutorial 1)
Basics of Perf

(Refer Slide Time: 0:16)

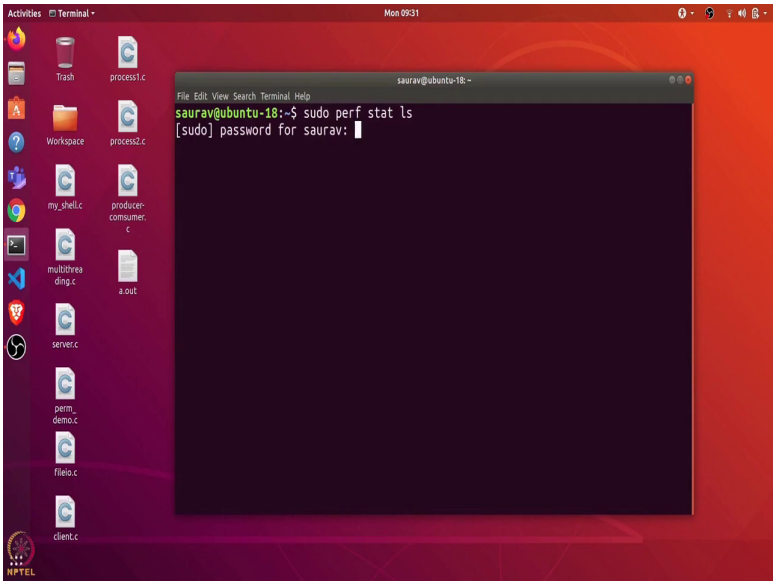

Hi everyone. In this video we will learn about Perf. Perf is a tool which offers a rich set of commands to collect and analyze performance data. If we want to know which parts of the programs are taking a lot of CPU time, then we can use perf. And let us open a terminal. And let us go through some of the very basic commands of perf. So, I will start with the stat command. Perf stat is used to collect statistics about any program. So, we can give it any command, let us say ls, and it will show us the statistics when ls is run. So, we will use sudo.

(Refer Slide Time: 0:55)

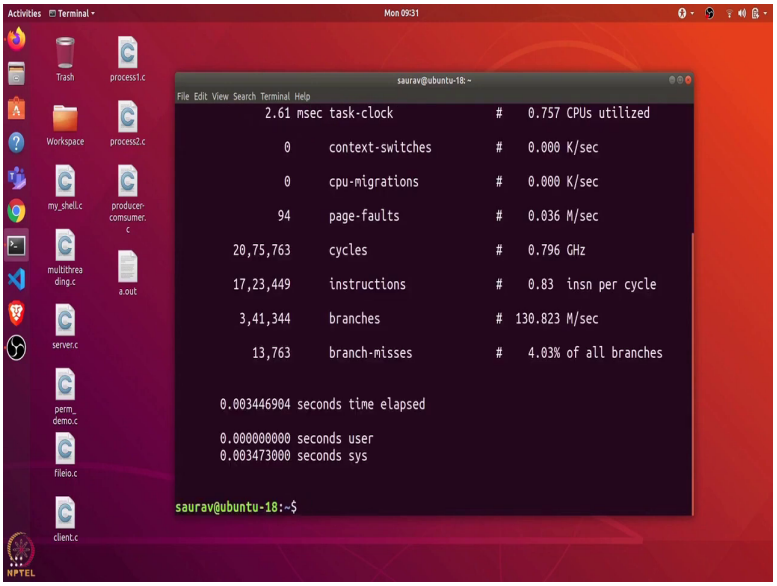

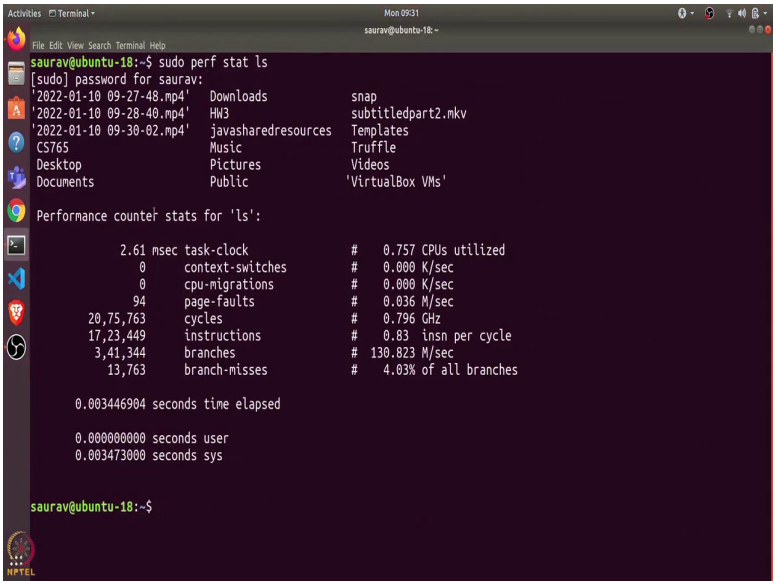

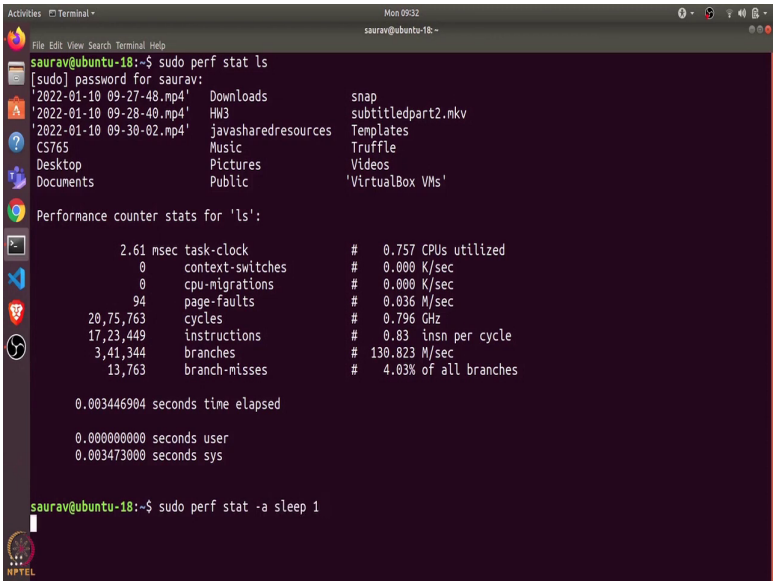

First, it executed ls and then it shows the counter statistics for ls, gives us how many context switches are there, how many CPU migrations page fault cycles, instructions et cetera. This is very basic use of perf to collect statistics about any program. Let us say we want to collect stats for the complete system. So, we want to collect statistics about every process that is running on any processor, we can use the -a flag to collect the statistics. And let us say we want to collect the statistics for one second, perf stat -a and then we can give it command sleep 1.

(Refer Slide Time: 1:36)

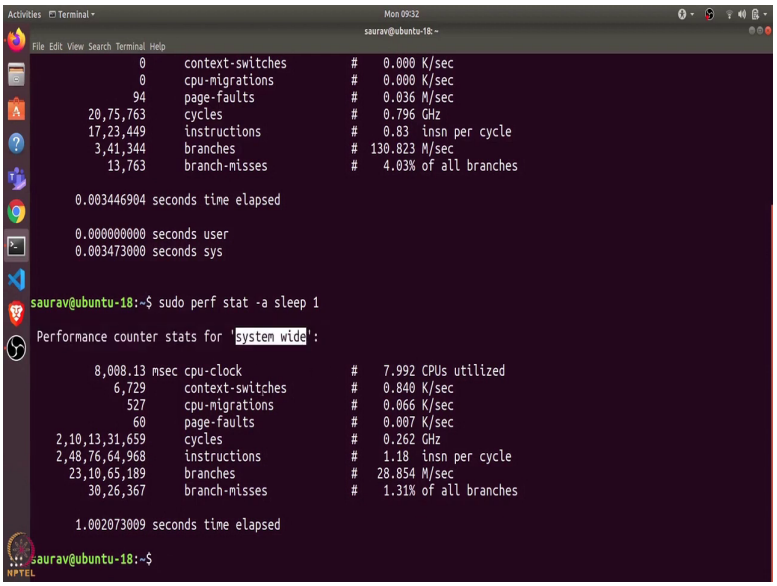

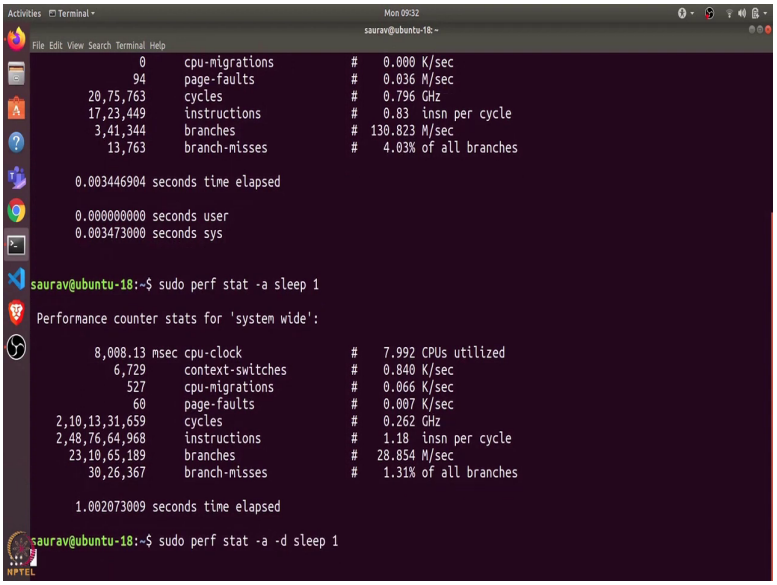

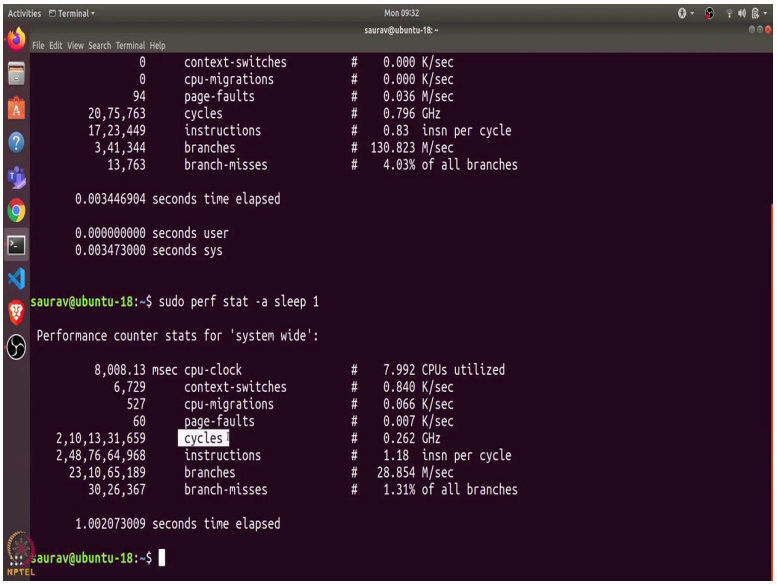

So, this will count the statistics for 1 second. So, it says statistics for system wide. So, we have these many contracts, which is in one second, these many page faults, etc. And we can use the -d flag, get some detailed statistics.

(Refer Slide Time: 1:50)

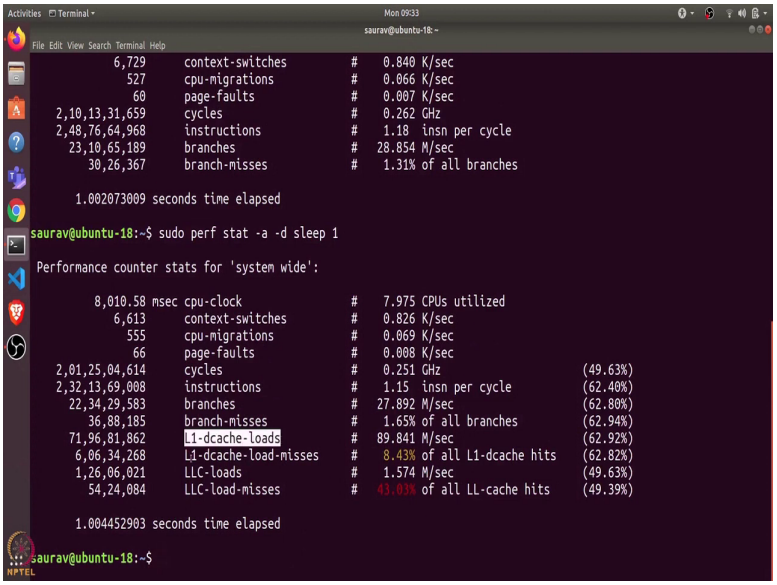
Activities
Terminal
Mon 09:33
saurav@ubuntu-18: ~
File Edit View Search Terminal Help
6,729 context-switches # 0.840 K/sec
527 cpu-migrations # 0.066 K/sec
60 page-faults # 0.007 K/sec
2,10,13,31,659 cycles # 0.262 GHz
2,48,76,64,968 instructions # 1.18 insn per cycle
23,10,65,189 branches # 28.854 M/sec
30,26,367 branch-misses # 1.31% of all branches
1.002073009 seconds time elapsed
saurav@ubuntu-18:~$ sudo perf stat -a -d sleep 1
Performance counter stats for 'system wide':
8,010.58 msec cpu-clock # 7.975 CPUs utilized
6,613 context-switches # 0.826 K/sec
555 cpu-migrations # 0.069 K/sec
66 page-faults # 0.008 K/sec
2,01,25,04,614 cycles # 0.251 GHz (49.63%)
2,32,13,69,008 instructions # 1.15 insn per cycle (62.40%)
22,34,29,583 branches # 27.892 M/sec (62.80%)
36,88,185 branch-misses # 1.65% of all branches (62.94%)
71,96,81,862 L1-dcache-loads # 89.841 M/sec (62.92%)
6,06,34,268 L1-dcache-load-misses # 8.43% of all L1-dcache hits (62.82%)
1,26,06,021 LLC-loads # 1.574 M/sec (49.63%)
54,24,084 LLC-load-misses # 43.03% of all LL-cache hits (49.39%)
1.004452903 seconds time elapsed
saurav@ubuntu-18:~$
NPTEL

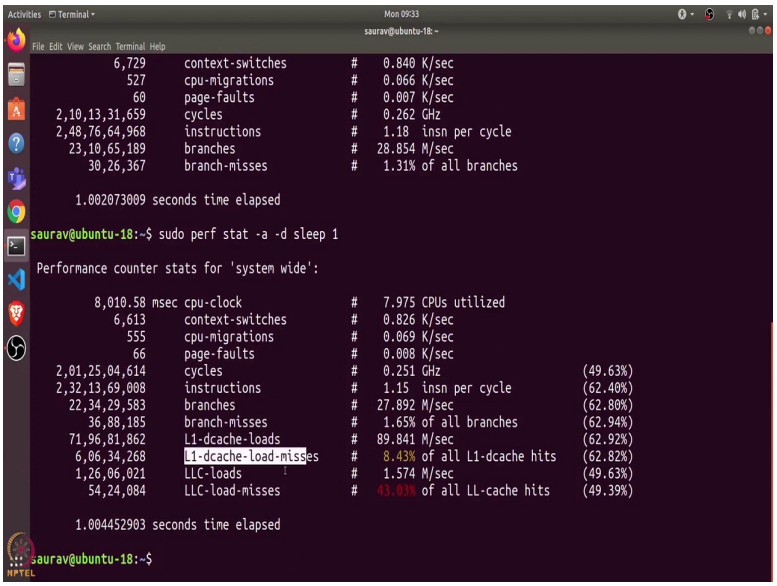
Activities
Terminal
Mon 09:33
saurav@ubuntu-18: ~
File Edit View Search Terminal Help
6,729 context-switches # 0.840 K/sec
527 cpu-migrations # 0.066 K/sec
60 page-faults # 0.007 K/sec
2,10,13,31,659 cycles # 0.262 GHz
2,48,76,64,968 instructions # 1.18 insn per cycle
23,10,65,189 branches # 28.854 M/sec
30,26,367 branch-misses # 1.31% of all branches
1.002073009 seconds time elapsed
saurav@ubuntu-18:~$ sudo perf stat -a -d sleep 1
Performance counter stats for 'system wide':
8,010.58 msec cpu-clock # 7.975 CPUs utilized
6,613 context-switches # 0.826 K/sec
555 cpu-migrations # 0.069 K/sec
66 page-faults # 0.008 K/sec
2,01,25,04,614 cycles # 0.251 GHz (49.63%)
2,32,13,69,008 instructions # 1.15 insn per cycle (62.40%)
22,34,29,583 branches # 27.892 M/sec (62.80%)
36,88,185 branch-misses # 1.65% of all branches (62.94%)
71,96,81,862 L1-dcache-loads # 89.841 M/sec (62.92%)
6,06,34,268 L1-dcache-load-misses # 8.43% of all L1-dcache hits (62.82%)
1,26,06,021 LLC-loads # 1.574 M/sec (49.63%)
54,24,084 LLC-load-misses # 43.03% of all LL-cache hits (49.39%)
1.004452903 seconds time elapsed
saurav@ubuntu-18:~$
NPTEL

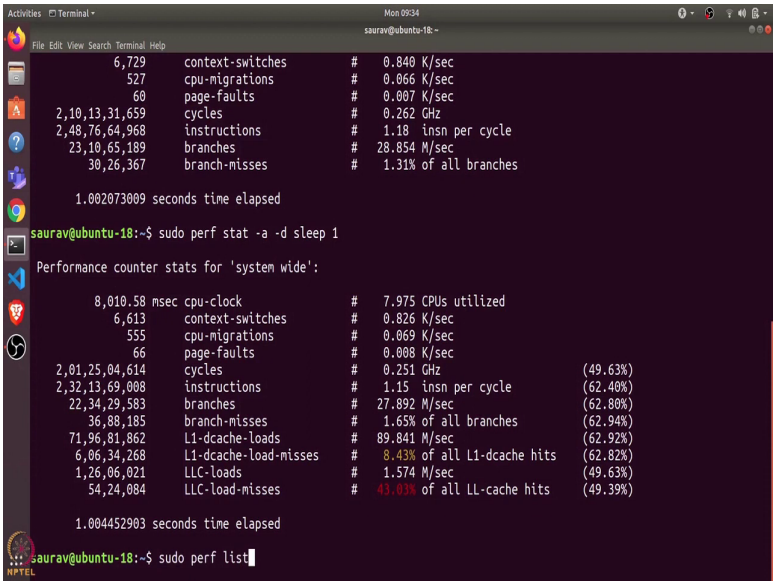
Activities
Terminal
Mon 09:34
saurav@ubuntu-18: ~
File Edit View Search Terminal Help
6,729 context-switches # 0.840 K/sec
527 cpu-migrations # 0.066 K/sec
60 page-faults # 0.007 K/sec
2,10,13,31,659 cycles # 0.262 GHz
2,48,76,64,968 instructions # 1.18 insn per cycle
23,10,65,189 branches # 28.854 M/sec
30,26,367 branch-misses # 1.31% of all branches
1.002073009 seconds time elapsed
saurav@ubuntu-18:~$ sudo perf stat -a -d sleep 1
Performance counter stats for 'system wide':
8,010.58 msec cpu-clock # 7.975 CPUs utilized
6,613 context-switches # 0.826 K/sec
555 cpu-migrations # 0.069 K/sec
66 page-faults # 0.008 K/sec
2,01,25,04,614 cycles # 0.251 GHz (49.63%)
2,32,13,69,008 instructions # 1.15 insn per cycle (62.40%)
22,34,29,583 branches # 27.892 M/sec (62.80%)
36,88,185 branch-misses # 1.65% of all branches (62.94%)
71,96,81,862 L1-dcache-loads # 89.841 M/sec (62.92%)
6,06,34,268 L1-dcache-load-misses # 8.43% of all L1-dcache hits (62.82%)
1,26,06,021 LLC-loads # 1.574 M/sec (49.63%)
54,24,084 LLC-load-misses # 43.03% of all LL-cache hits (49.39%)
1.004452903 seconds time elapsed
saurav@ubuntu-18:~$ sudo perf list
NPTEL

If I run it again, there, it shows us some other strategies such as L1 cache loads, L1 cache, load misses et cetera. Another important command that you can use with perf is perf list. So, if I run sudo perf list.

(Refer Slide Time: 2:03)

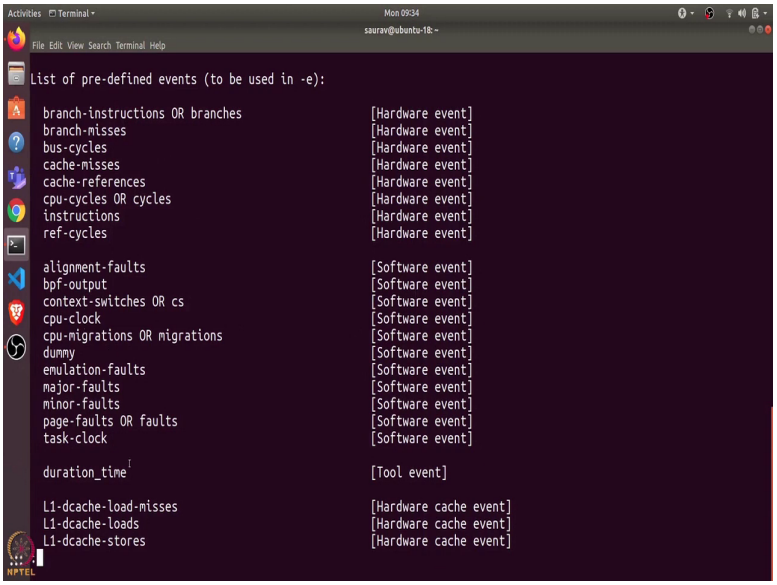

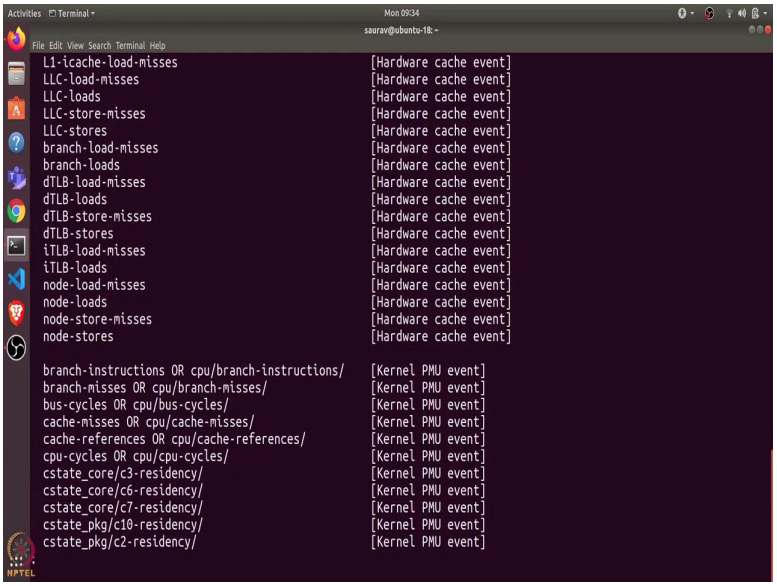

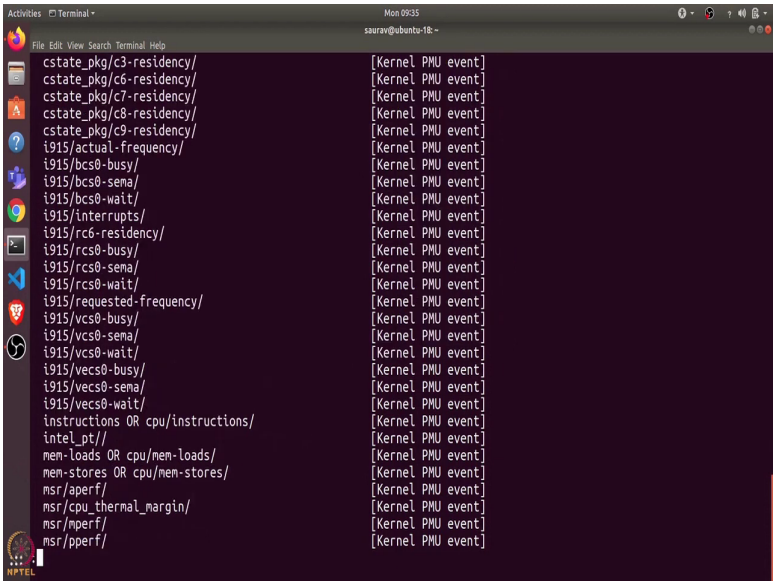

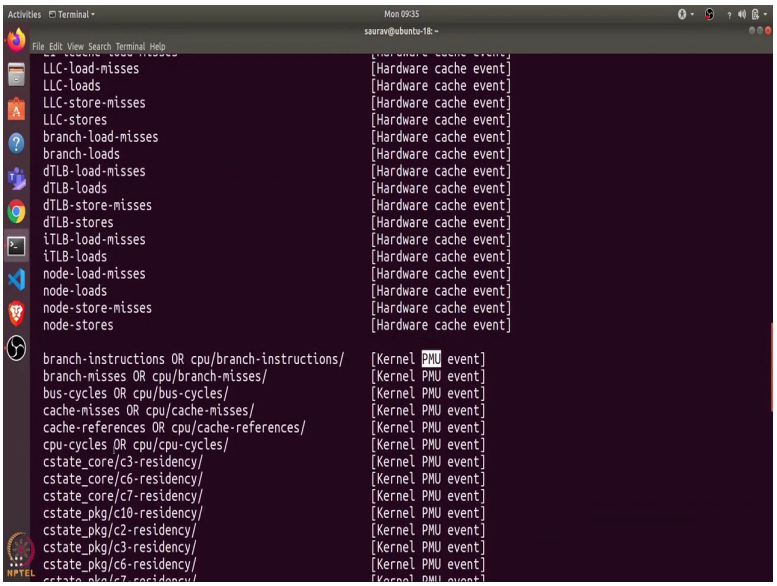

It will list all the events that we can count using perf. So, it shows us some hardware events, such as branch instruction, branch misses, then there are some software events for which kernel maintains a counter, such as CPU migration, context switches. And then we have some cache events. This PMU refers to performance monitoring unit. So, there is a separate unit with the processor, which counts various hardware events, which include these branch instructions, or cache misses, et cetera.

So, there are a lot of events, let us say want to measure certain event. For instance, let us try to measure these branch instructions.

(Refer Slide Time: 2:42)

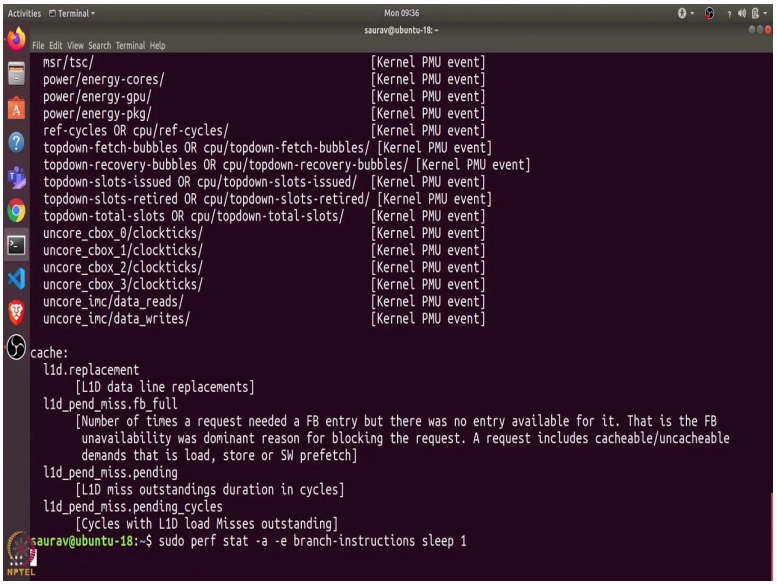

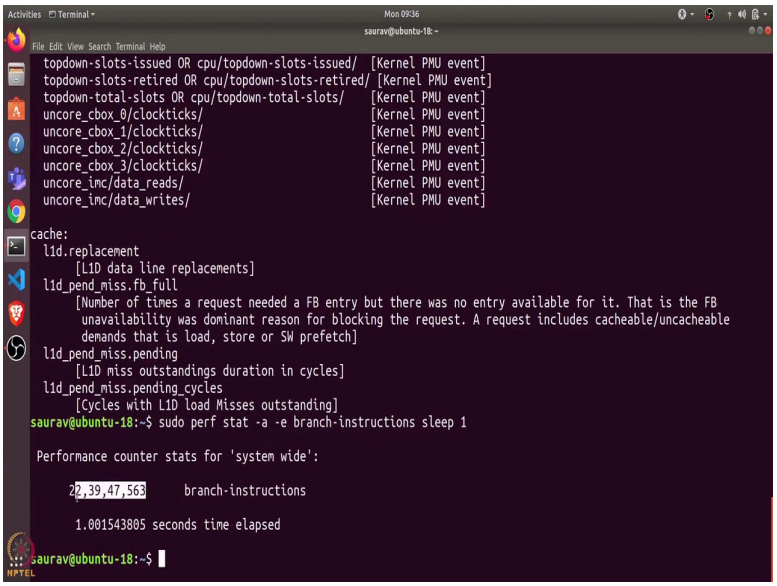

So, what we can do is we can use the -e flag. Let us say you want to measure branch instructions for 1 second, across all the processes, so I can use -e branch instructions. And then it will give me the count for the branch instructions for 1 second.

(Refer Slide Time: 3:04)

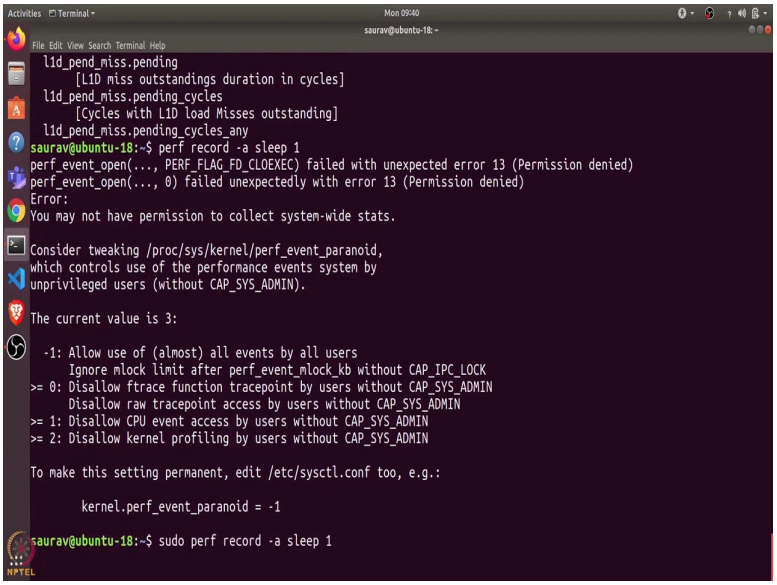

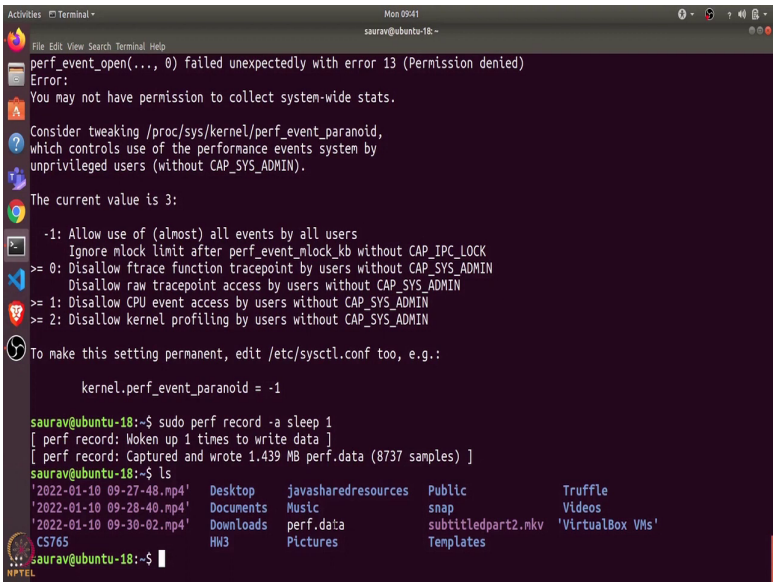

The next command that we will see is perf report. Perf report is used to collect the profile data about a certain program in a separate file. So, let us say I want to record the data for complete system for 1 second, I need to use sudo. If we do ls, then we can see there is this perf.data file. This is a binary file, which stores the profiling data for our complete system.

(Refer Slide Time: 3:44)

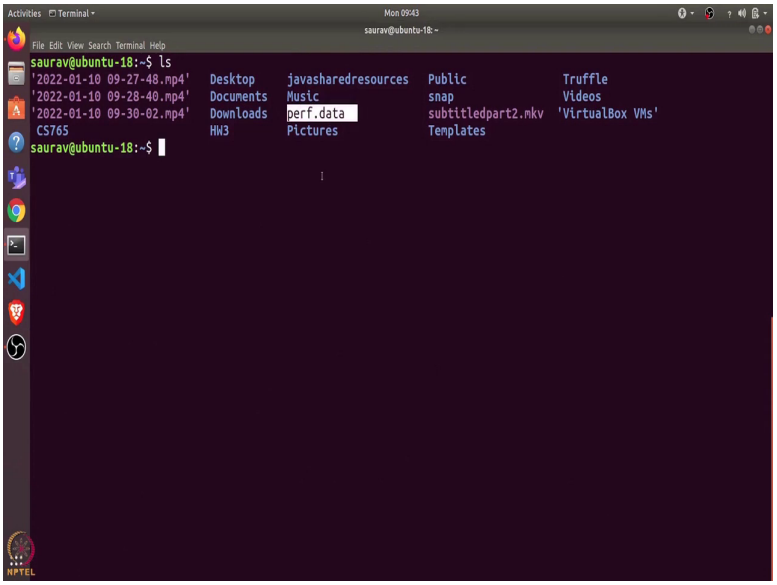
Activities
Terminal
Mon 09:43
saurav@ubuntu-18: ~
File Edit View Search Terminal Help
saurav@ubuntu-18:~$ ls
'2022-01-10 09-27-48.mp4' Desktop javasharedresources Public Truffle
'2022-01-10 09-28-40.mp4' Documents Music snap Videos
'2022-01-10 09-30-02.mp4' Downloads perf.data subtitledpart2.mkv 'VirtualBox VMs'
CS765 HW3 Pictures Templates
saurav@ubuntu-18:~$

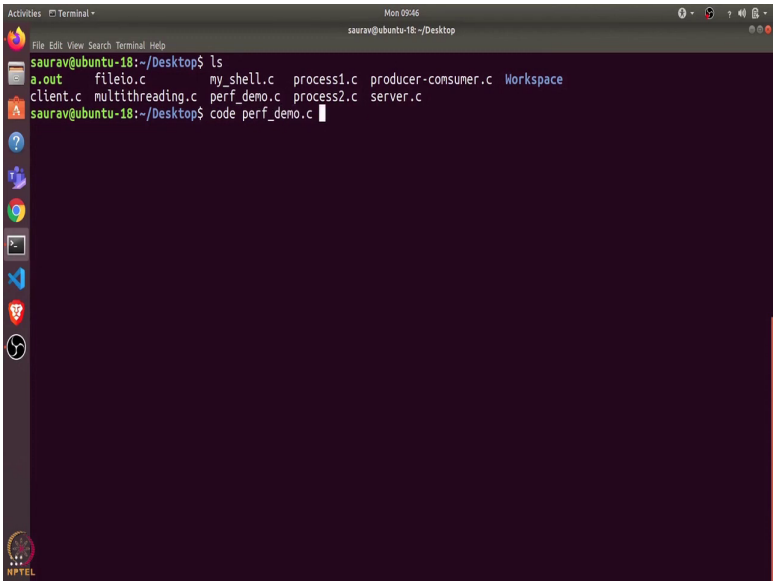
Activities
Terminal
Mon 09:46
saurav@ubuntu-18: ~/Desktop
File Edit View Search Terminal Help
saurav@ubuntu-18:~/Desktop$ ls
a.out fileio.c my_shell.c process1.c producer-comsumer.c Workspace
client.c multithreading.c perf_demo.c process2.c server.c
saurav@ubuntu-18:~/Desktop$ code perf_demo.c

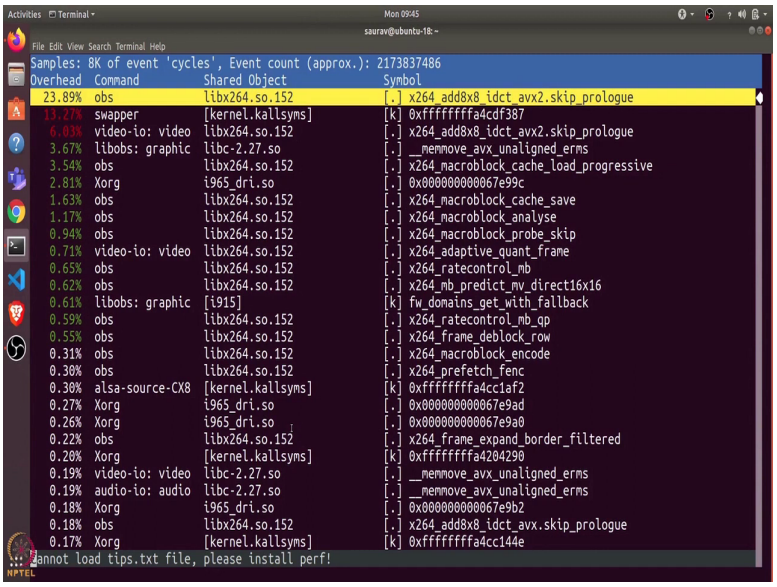
Activities
Terminal
Mon 09:45
saurav@ubuntu-18: ~
File Edit View Search Terminal Help
Samples: 8K of event 'cycles', Event count (approx.): 2173837486
Overhead Command Shared Object Symbol
23.89% obs libx264.so.152 [.] x264_add8x8_idct_avx2.skip_prologue
13.27% swapper [kernel.kallsyms] [k] 0xffffffffa4cdf387
6.03% video-io: video libx264.so.152 [.] x264_add8x8_idct_avx2.skip_prologue
3.67% libobs: graphic libc-2.27.so [.] __memmove_avx_unaligned_erms
3.54% obs libx264.so.152 [.] x264_macroblock_cache_load_progressive
2.81% Xorg i965_dri.so [.] 0x000000000067e99c
1.63% obs libx264.so.152 [.] x264_macroblock_cache_save
1.17% obs libx264.so.152 [.] x264_macroblock_analyse
0.94% obs libx264.so.152 [.] x264_macroblock_probe_skip
0.71% video-io: video libx264.so.152 [.] x264_adaptive_quant_frame
0.65% obs libx264.so.152 [.] x264_ratecontrol_mb
0.62% obs libx264.so.152 [.] x264_mb_predict_mv_direct16x16
0.61% libobs: graphic [i915] [k] fw_domains_get_with_fallback
0.59% obs libx264.so.152 [.] x264_ratecontrol_mb_qp
0.55% obs libx264.so.152 [.] x264_frame_deblock_row
0.31% obs libx264.so.152 [.] x264_macroblock_encode
0.30% obs libx264.so.152 [.] x264_prefetch_fenc
0.30% alsa-source-CX8 [kernel.kallsyms] [k] 0xffffffffa4cc1af2
0.27% Xorg i965_dri.so [.] 0x000000000067e9ad
0.26% Xorg i965_dri.so [.] 0x000000000067e9a0
0.22% obs libx264.so.152 [.] x264_frame_expand_border_filtered
0.20% Xorg [kernel.kallsyms] [k] 0xffffffffa4204290
0.19% video-io: video libc-2.27.so [.] __memmove_avx_unaligned_erms
0.19% audio-io: audio libc-2.27.so [.] __memmove_avx_unaligned_erms
0.18% Xorg i965_dri.so [.] 0x000000000067e9b2
0.18% obs libx264.so.152 [.] x264_add8x8_idct_avx.skip_prologue
0.17% Xorg [kernel.kallsyms] [k] 0xffffffffa4cc144e
Cannot load tips.txt file, please install perf!

So, how can we do this perf dot data file? We need to use perf report to read this binary file. Let us use perf report where input file perf.data. So, this shows us what is the percentage of CPU cycles that were used for this particular command and the shared object and so on. So, let us quit this using Q and let us try to profile some C program. So, I have written this perf demo dot c file. Let us first have a look at its contents.

(Refer Slide Time: 4:18)

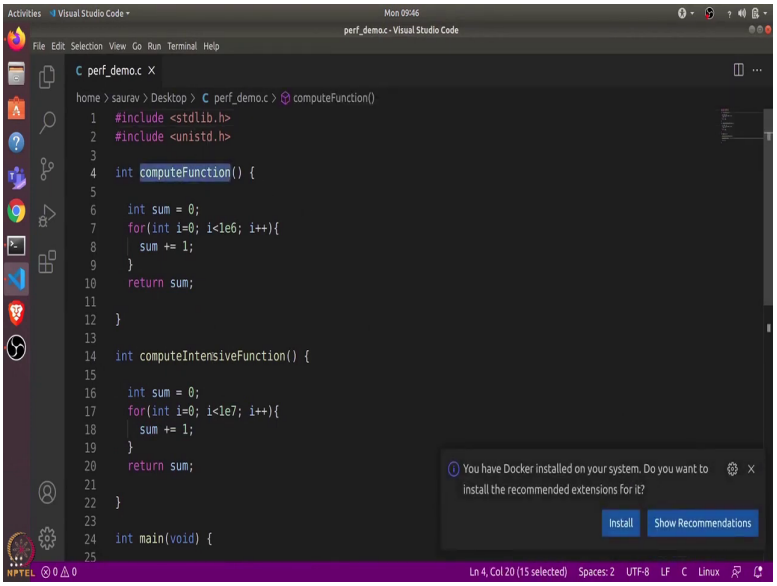

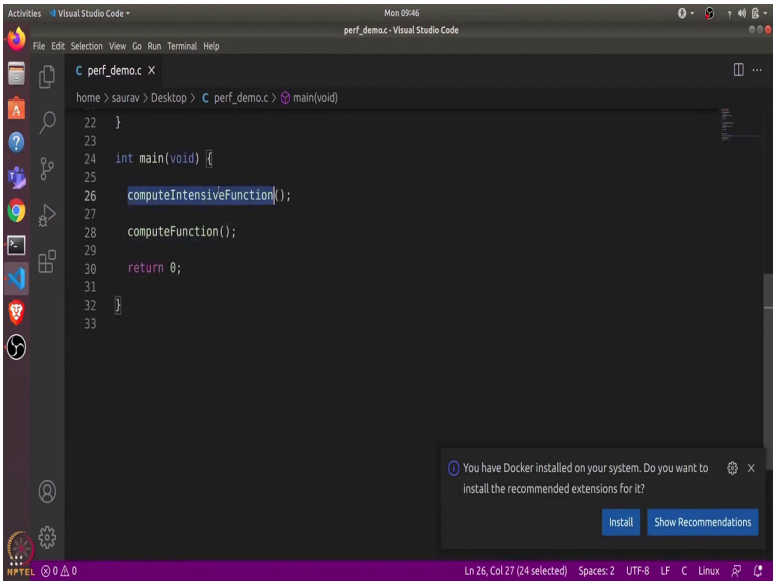

So, this is a very simple program. It has two functions. One is computeFunction, another is computeIntensiveFunction. And in the main function, we just call these two functions.

(Refer Slide Time: 4:31)

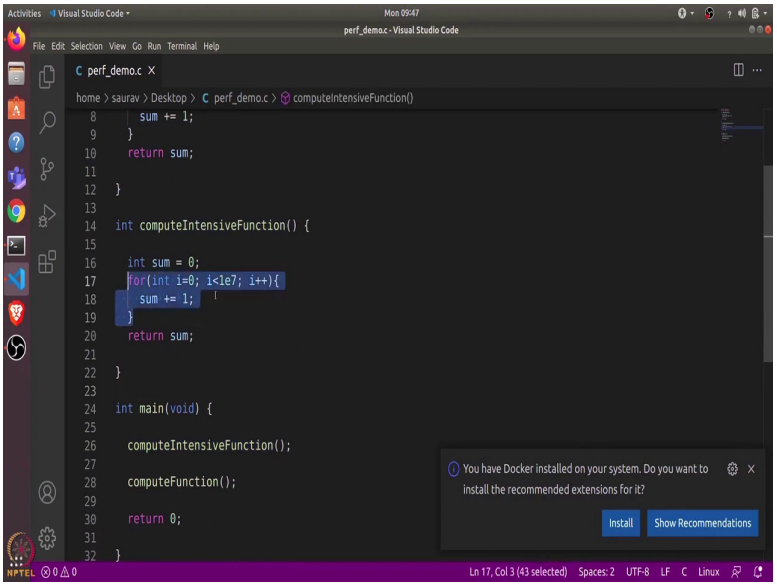

In the computeIntensiveFunction, we iterate 10 raise to 7 times, and we just add one to the sum a dummy variable and return that sum.

(Refer Slide Time: 4:38)

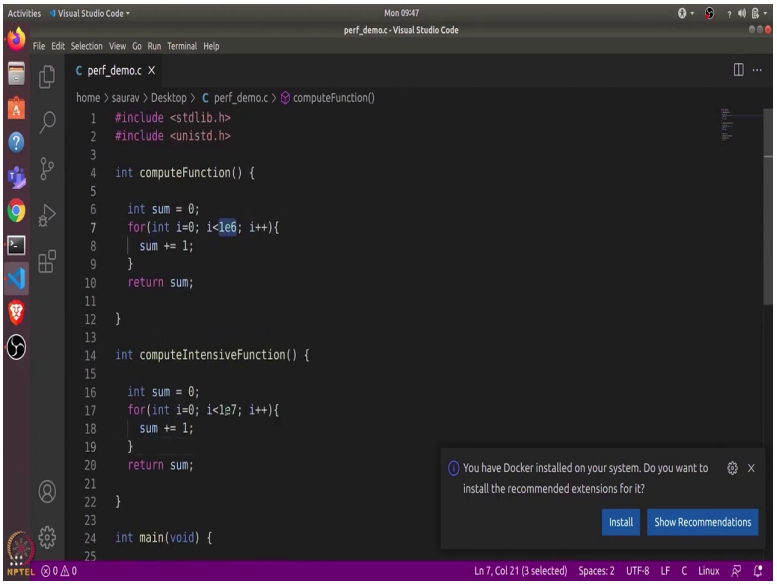

In computeFunction, we iterate 10 raise to 6 times. So, computeIntensiveFunction uses 10 times more number of iterations then computeFunction.

(Refer Slide Time: 4:46)

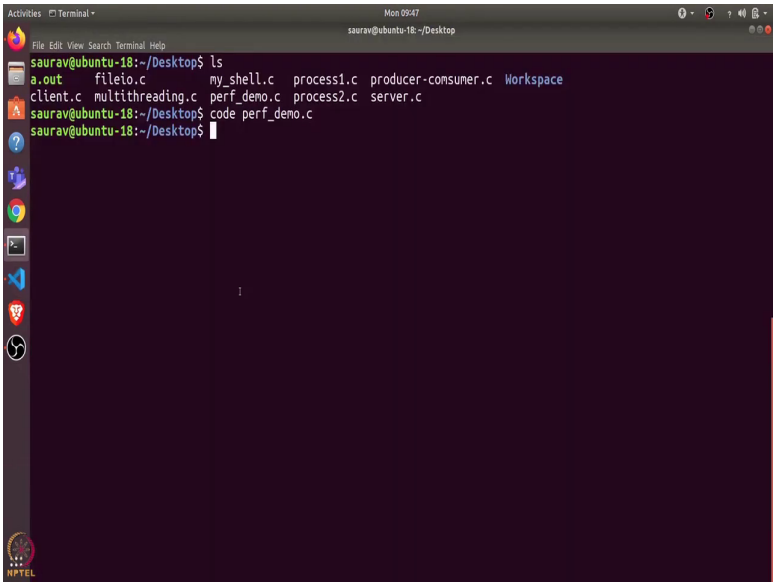

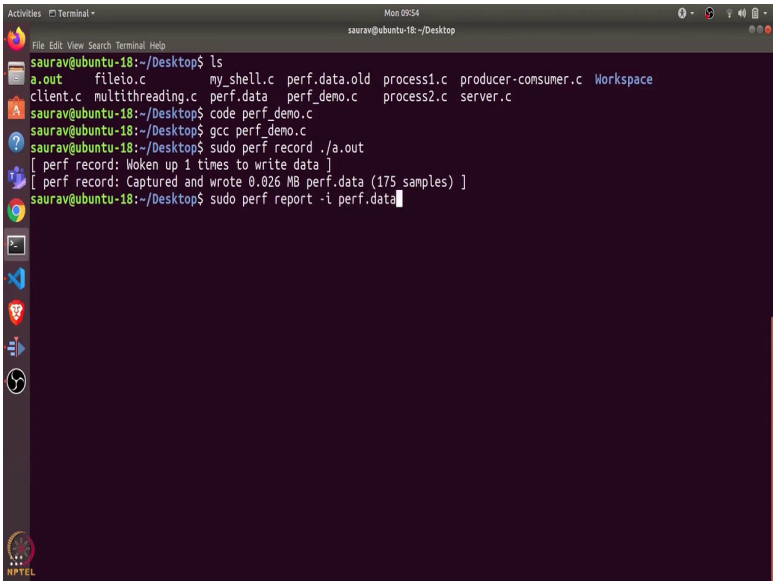

Let us try to compile this gcc perf_demo.c. And let us run the executable and try to profile it using perf. So, I will use sudo perf record ./a.out and this will generate the perf.data file. Now let us try to see the profile using sudo perf report -i perf.data.

(Refer Slide Time: 5:22)

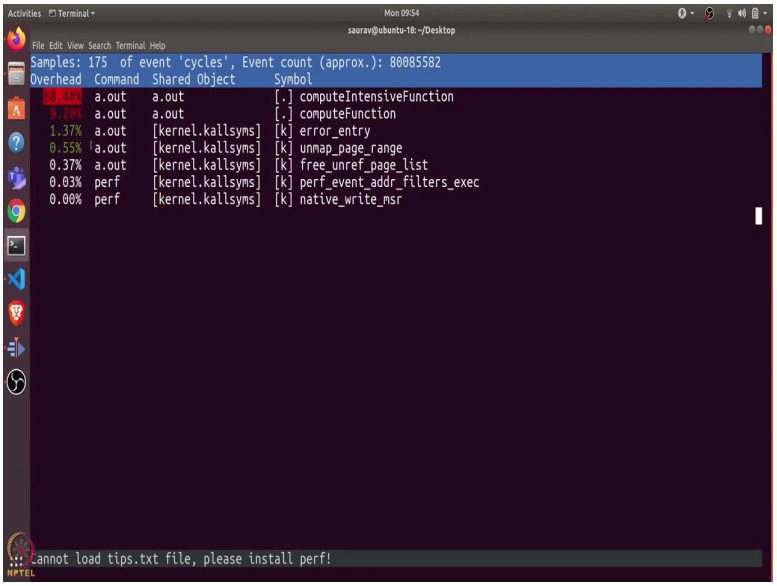

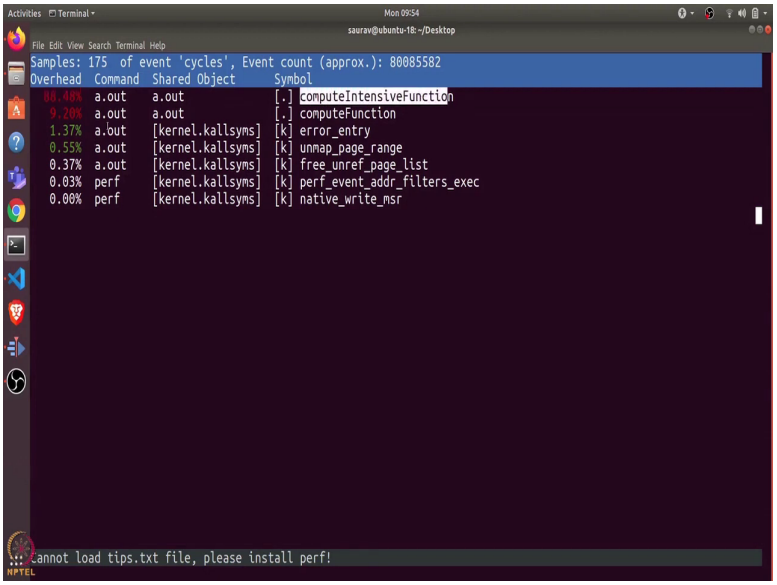

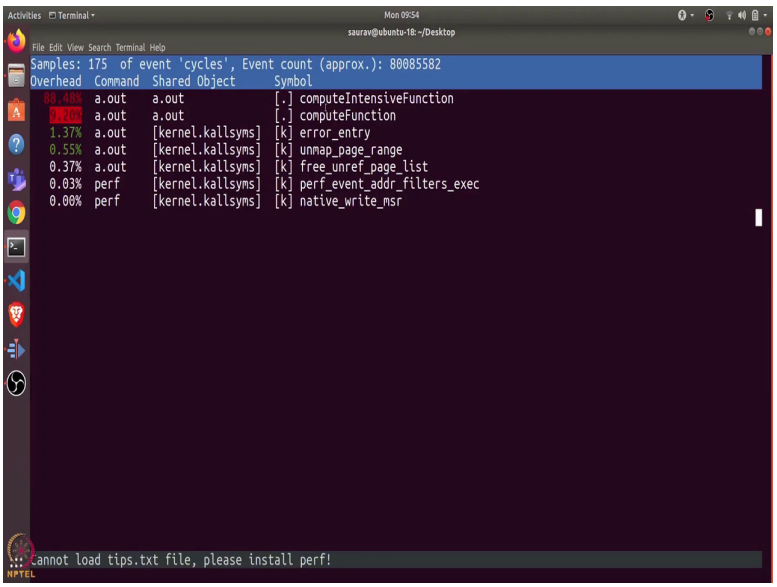

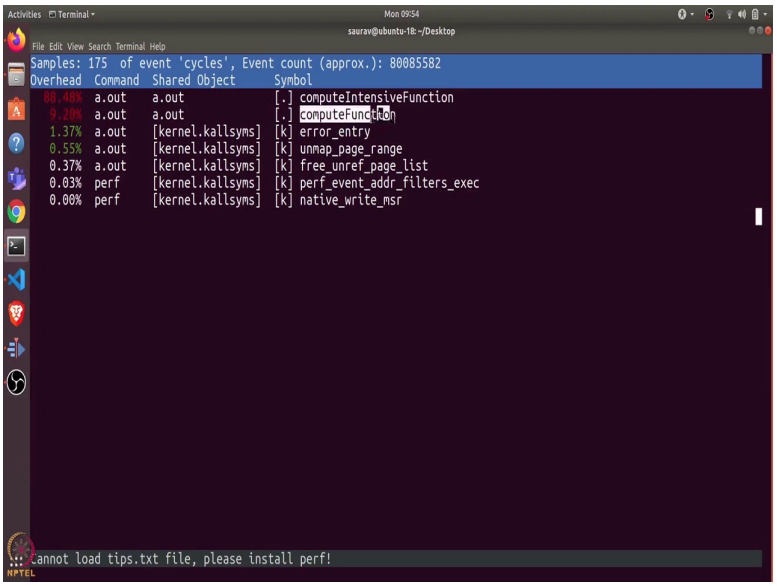

So, here it shows us the percentage of CPU that is used for various functions and the command that is a.out, and then the function name. Here we can see that almost 88 percent of CPU is used only in computeIntensiveFunction, and almost 9 percent is used in the computeFunction.

(Refer Slide Time: 5:46)

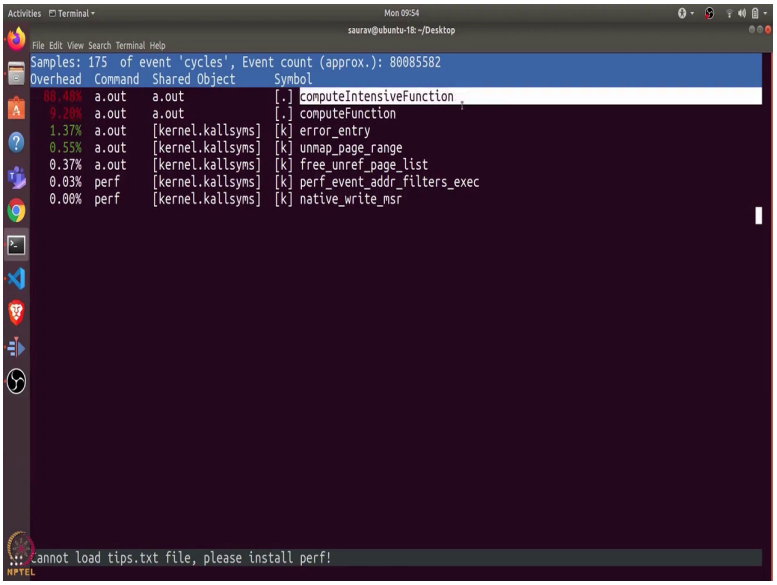

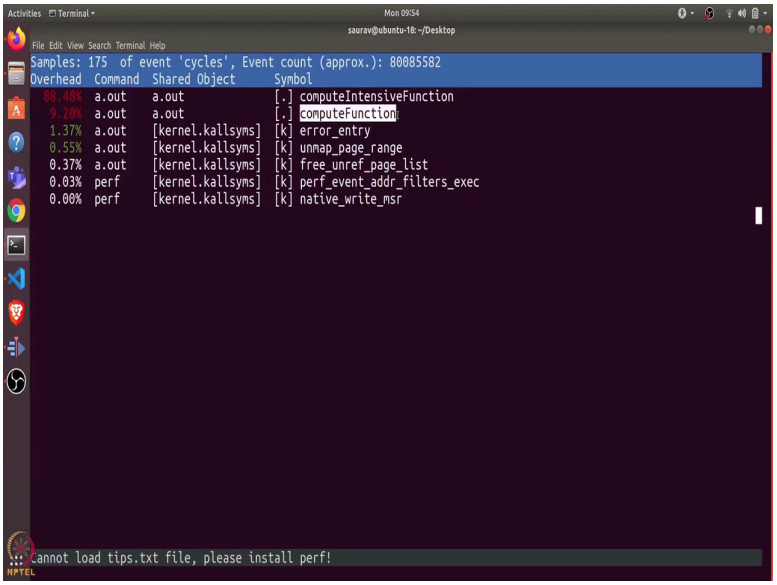

So, computeIntensiFunction is almost 10 times more expensive than compute function. And this also shows us what is the bottleneck of our program that is the computeIntensiveFunction.

(Refer Slide Time: 5:55)

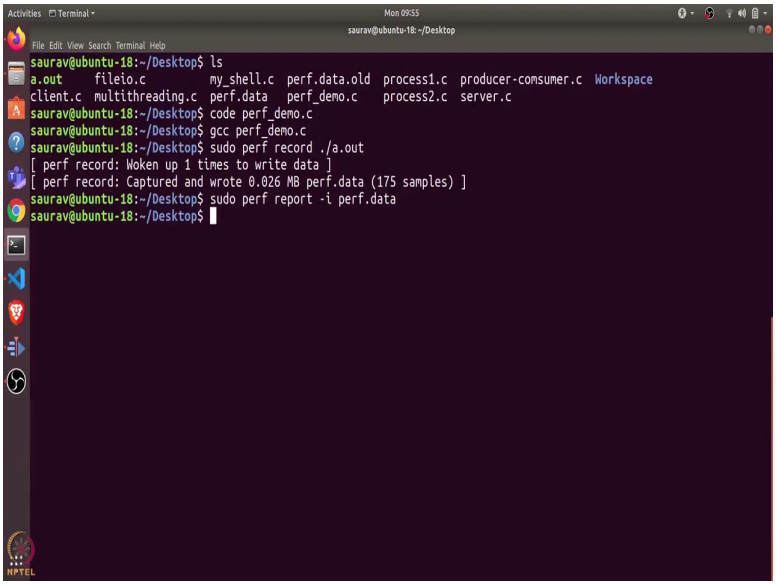
Activities
Terminal
Mon 09:55
saurav@ubuntu-18: ~/Desktop
File Edit View Search Terminal Help
saurav@ubuntu-18:~/Desktop$ ls
a.out fileio.c my_shell.c perf.data.old process1.c producer-comsumer.c Workspace
client.c multithreading.c perf.data perf_demo.c process2.c server.c
saurav@ubuntu-18:~/Desktop$ code perf_demo.c
saurav@ubuntu-18:~/Desktop$ gcc perf_demo.c
saurav@ubuntu-18:~/Desktop$ sudo perf record ./a.out
[perf record: Woken up 1 times to write data]
[perf record: Captured and wrote 0.026 MB perf.data (175 samples)]
saurav@ubuntu-18:~/Desktop$ sudo perf report -i perf.data
saurav@ubuntu-18:~/Desktop$

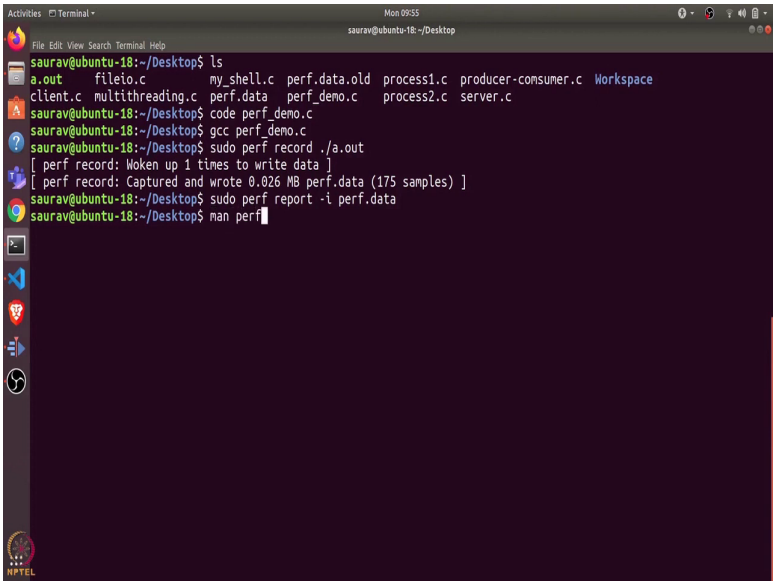
Activities
Terminal
Mon 09:55
saurav@ubuntu-18: ~/Desktop
File Edit View Search Terminal Help
saurav@ubuntu-18:~/Desktop$ ls
a.out fileio.c my_shell.c perf.data.old process1.c producer-comsumer.c Workspace
client.c multithreading.c perf.data perf_demo.c process2.c server.c
saurav@ubuntu-18:~/Desktop$ code perf_demo.c
saurav@ubuntu-18:~/Desktop$ gcc perf_demo.c
saurav@ubuntu-18:~/Desktop$ sudo perf record ./a.out
[perf record: Woken up 1 times to write data]
[perf record: Captured and wrote 0.026 MB perf.data (175 samples)]
saurav@ubuntu-18:~/Desktop$ sudo perf report -i perf.data
saurav@ubuntu-18:~/Desktop$ man perf

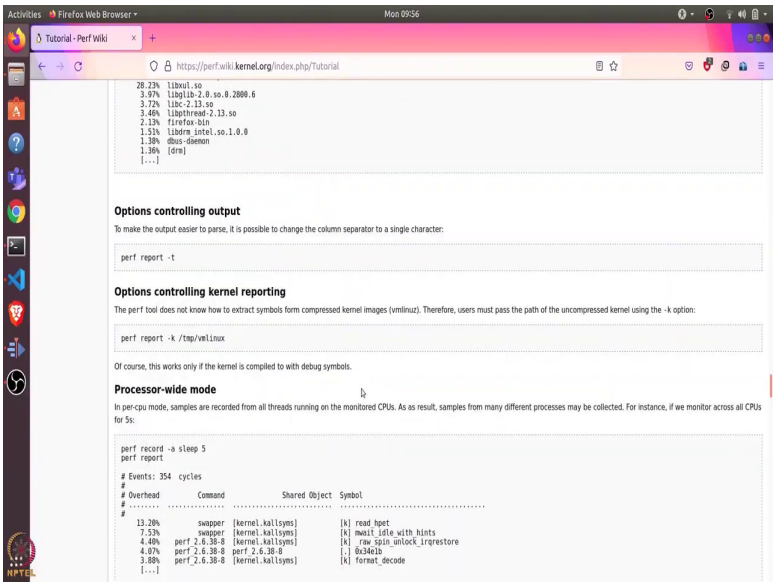
Activities
Firefox Web Browser
Mon 09:56
Tutorial - Perf Wiki
https://perf.wiki.kernel.org/index.php/Tutorial
28.23% libxul.so
3.97% libglib-2.0.so.0.2800.6
3.72% libc-2.13.so
3.46% libpthread-2.13.so
2.13% firefox-bin
1.51% libdrm_intel.so.1.0.0
1.38% dbus-daemon
1.36% [drm]
[...]
Options controlling output
To make the output easier to parse, it is possible to change the column separator to a single character:
perf report -t
Options controlling kernel reporting
The perf tool does not know how to extract symbols form compressed kernel images (vmlinuz). Therefore, users must pass the path of the uncompressed kernel using the -k option:
perf report -k /tmp/vmlinux
Of course, this works only if the kernel is compiled to with debug symbols.
Processor-wide mode
In per-cpu mode, samples are recorded from all threads running on the monitored CPUs. As as result, samples from many different processes may be collected. For instance, if we monitor across all CPUs for 5s:
perf record -a sleep 5
perf report
Events: 354 cycles
Overhead Command Shared Object Symbol
13.20% swapper [kernel.kallsyms] [k] read_hpet
7.53% swapper [kernel.kallsyms] [k] mwait_idle_with_hints
4.40% perf_2.6.38-8 [kernel.kallsyms] [k] _raw_spin_unlock_irqrestore
4.07% perf_2.6.38-8 perf_2.6.38-8 [.] 0x34e1b
3.80% perf_2.6.38-8 [kernel.kallsyms] [k] format_decode
[...]

So, that is how we can profile various C programs and try to optimize their bottlenecks. If you want to see more details about perf then you can have a look at the man page of perf. And also there is a very nice tutorial on perf which you can find online. So, this explains various commands of perf, and what else you can do using perf. So, that is it for this video. Thanks and have a nice day.