The Art of Measurement[†]

IN4390 Quantitative Evaluation of Embedded Systems Koen Langendoen





[†]Original slides by Mitra Nasri, now at TU/e

Course outline

• Lect 1: terminology, non-functional properties

• Lect 2+3: Design of Experiments

• Lect 4: The Art of Measurements





Agenda



• The art of measurement

- 1. Pillars of measurement-based evaluation
- 2. Goals and metrics (book chapter 3)
- 3. Workload selection/generation (book chapter 4)
- 4. Measurement: practical considerations
- 5. The art of data presentation (book chapter 10, 11)
- 6. Confidence level and confidence interval (book chapter 13)
- 7. Summarizing measured data (book chapter 12)

• Self study

- Common mistakes (book chapter 1, 2)
- Summarizing measured data (book chapter 12)
- The art of data presentation (book chapter 10, 11)



Source:

• Raj Jain, "The Art of Computer Systems Performance Analysis"





An overview of evaluation techniques

Analytical modeling

- Based on a rigorous mathematical models
- Lectures 7, 9-13 (MC + QT)

Simulation

- Simulate the system operation (usually only small parts thereof)
- Not covered in this course

Measurement

- Implement the system in full and measure its performance directly
- This lecture

- **Positive**: Provides the best insight into the effects of different parameters and their interaction
 - example: is it better to configure the system with one fast disk or with two slow disks?
- **Positive**: Can be done before the system is built and takes a short time
- Negative: these models are rarely accurate
 - Usually needs many simplifying assumptions
 - Depends on the quality and correctness of these assumptions

read more here: https://tu-dresden.de/zih/ressourcen/dateien/lehre/ws1112/lars/vorlesungen/lars_lecture_02_requirements-metrics-techniques.pdf?lang=en





An overview of evaluation techniques

Analytical modeling

- Based on a rigorous mathematical models
- Lectures 7, 9-13 (MC + QT)

Simulation



- Simulate the system operation (usually only small parts thereof)
- Not covered in this course

Measurement

- Implement the system in full and measure its performance directly
- This lecture

- **Positive**: full control of simulation model, parameters, level of detail (high flexibility).
- Positive: Can be done before the system is built
 - Simulation of a full system is infeasible
 - Simulation of the system parts does not take everything into account

• Negative:

• It may still include simplifying assumptions

read more here: <u>https://tu-dresden.de/zih/ressourcen/dateien/lehre/ws1112/lars/vorlesungen/lars_lecture_02_requirements-metrics-techniques.pdf?lang=en</u>





An overview of evaluation techniques

Analytical modeling

- Based on a rigorous mathematical models
- Lectures 7, 9-13 (MC + QT)

Simulation

- Simulate the system operation (usually only small parts thereof)
- Not covered in this course

Measurement

- Implement the system in full and measure its performance directly
- This lecture



Positive: The most convincing

- Negative: has high costs
 - Requires a full implementation of the system (buying hardware)

Numbers don't lie (but statistics do ⁽ⁱⁱⁱ⁾)

• What if you want to change the hardware in case the performance is bad?

read more here: <u>https://tu-dresden.de/zih/ressourcen/dateien/lehre/ws1112/lars/vorlesungen/lars_lecture_02_requirements-metrics-techniques.pdf?lang=en</u>





Pillars of Measurement-Based Evaluation







Pillars of measurement-based evaluation







Goal What do you want to evaluate?



Goal: no human will ever be harmed by my system

Question: What should you measure to ensure this property?

Assume that your object-detection algorithm is 100% accurate and never does any mistake!

Question: What type of non-functional property is this?





Metrics

Metric types



- a count of how many times an event occurs
- the duration of some time interval
- the size of some parameter

Metrics can have conditions

- Number of successfully processed images in one hour
- Total execution time of the image processing task for serving 100 images
- Time until the next failure







Metric and goal: example











System parameters (configurations)

 System parameters are configurations and parameters that impact a system's performance independent of the system's input workload

Example:

Metric is the round trip time of the object tracking

Question: what are the system parameters in our example?



System parameters for our example:

- Speed of: base station's CPU, robot's microcontroller, and network.
- Scheduling policy on both ends.
- Network protocol.
- Operating system overhead for interfacing with the network.
- Reliability of the network affecting the number of retransmissions required.
- •







Workload model

- Workload model defines the model of input workload to the system
- It is usually a representation of the actual usage of the system after deployment



Workload parameters for this example:

- Number of IoT devices (user)
- Frequency of sending requests to the base station
 - and distribution (periodic, Poisson, ...)
- Complexity of the navigation task at different locations
 - Example: around cross roads is harder than straight roads
 - Proximity (presence) of other devices when it reaches to a cross road
- Other load on the robot's microcontroller and base station CPUs
- Other load on the network









Factors

- Factors: Parameters that are varied during the evaluation
- Levels: Values of a factor

Example:

- Image size: {10KB, 100KB, 1MB}
- Frequency of calls to service:
 - Periodic: 1 every {1, 10, 100, 1000} milliseconds
 - Exponential distribution with rate: {0.1, 0.2, 0.3, 0.4, 0.5}

Assumption:

- Fixed: type of CPU and operating system.
- Measure under no other load on the hosts and the network.





- Service: image rotation
- Input: image and angle
- Output: rotated image







Measuring Time: Practical Considerations



Image source: https://phys.org/news/2017-03-scientists-animals-year.html





16 | 67

Measuring time

Don't measure by hand!

Measuring by hand: "looking" at a clock, launching the program, "looking" at the clock again when the program terminates

Approaches

- 1. Measuring a process's execution time
 - UNIX's "time" command

2. Measuring time inside a program

• Programming language / OS support

3. Measuring time from outside of the system

- @inside: when the event (of interest) happens, trigger an I/O pin.
- @outside: monitor the I/O pin for changes to record the event time



UNIX time command



- 2.11 seconds of user time
- 1.94 seconds of system time
- 1 minutes and 8.72 seconds of wall-clock time
- 2% of CPU was used
- 0+0k memory used (text + data)
- 0 input, 0 output (file system I/O)
- 1155 minor pagefaults
- O swaps

See more here: https://en.wikipedia.org/wiki/Time_(Unix)





- time spent executing user code
- System Time
 - time spent executing kernel code (inside OS)
- Wall-Clock Time
 - time from start to end



UNIX time command

```
surf:~$ /usr/bin/X11/time ls -la -R ~/ > /dev/null
2.11user 1.94system 1:8.72elapsed 2%CPU
(0avgtext+0avgdata 0maxresident)k
0inputs+0outputs (0major+1155minor)pagefaults 0swaps
```

- User Time
 - time spent executing user code
- System Time
 - time spent executing kernel code (inside OS)
- Wall-Clock Time
 - time from start to end

Question: Why wall-clock ≥ User + System?

- 1. because the process can be preempted by other processes that run concurrently on the system
- 2. because the process can be blocked waiting for I/O



System time

Question: When do you get a large "system time"?

Larger system time may be a sign of higher calls to system APIs such as fopen(), fread(), fwrite(), etc.

To pinpoint which system calls are more often used, you need to track them via separate tracing tools such as **ftrace**



KernelShark

- A visualization tool to track "ftraces"
- Docs: <u>https://kernelshark.org/Documentation.html</u>
- **Git:** <u>https://git.kernel.org/pub/scm/utils/trace-cmd/trace-cmd.git/</u>

More information: https://www.kernel.org/doc/Documentation/trace/ftrace.txt







Visualizing ftrace outputs using KernelShark

·	sho	ws the executi multiple CPU	ion on Is Kernel Shark (/home/rostedt/trace.o	dat)	• . f
le Filter Plots Tools Pointer: 21 CPUs Tasks < + - > 2121452.512234	Help Help Help	121452.5122693 Marker 8 2121452	.5122731 A.B Delta: 0.0000039 2121452.512279	provides very nice ZOOM option	2121452.5123
сри о			արդի ուստա		
CPU 1				1 4 11	
CPU 2					
CPU 3					
			I I		
search: Column #	* contains *	Next Prev	✓ Graph follows	Tells you when the system calls t	hat
iearch: Column # CPU Time Stam	 contains Task PID 	Latency Event	✓ Graph follows	Tells you when the system calls t	hat
Search: Column #	 contains Task PID 2269 migrate 288 	Latency Event 3 d sched_wakeup	Graph follows Info migrate:28802 [91] success=1 CPU:000	Tells you when the system calls the	hat
Search: Column # CPU Time Stam; 197442 1 2121452.51: 197443 1 2121452.51:	 contains r p Task PID 2269 migrate 288 2269 migrate 288 	Latency Event 13 d sched_wakeup 3 sys_exit	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0	Tells you when the system calls the you have tracked are called	hat
Search: Column # CPU Time Stam 97442 1 2121452.51 97444 0 2121452.51	v contains v p Task PID 2269 migrate 288/ 2269 migrate 288/ 2269 migrate 287	Next Prev Latency Event 3 d systemit 3 systemit 6 dNh. reschedule_entry	✓ Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253	Tells you when the system calls the system calls the system called by th	hat
CPU Time Stam 9 CPU Time Stam 197442 1 197443 1 2121452.51 97444 1 2121452.51 97445 1 2121452.51	v contains v p Task PID 2269 migrate 288	Next Prev Latency Event 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit_futex	✓Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0	Tells you when the system calls the system calls the system called by th	hat
Search: CPU Time Stam 9 CPU 2121452.51 197442 1 2121452.51 197443 1 2121452.51 197445 1 2121452.51 197445 1 2121452.51 197445 1 2121452.51 197446 0 2121452.51 197445 1 2121452.51	v contains v p Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 287 2269 migrate 287 2270 migrate 287	Next Prev 3 sched_wakeup 3 sys_exit 6 dth. reschedul_entry 3 sys_exit 6 dth. reschedul_entry 6 dth. reschedul_exit		Tells you when the system calls the you have tracked are called	hat
CPU Time Stam 9 CPU Time Stam 197442 1 2121452.51 97443 1 2121452.51 97444 0 2121452.51 197445 1 2121452.51 197446 0 2121452.51 197440 0 2121452.51 197440 0 2121452.51	contains PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2269 migrate 287 2270 migrate 287 2270 migrate 287	Next Prev Latency Event 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit_futex 6 dNh. reschedule_exit 6 dNh. reschedule_exit	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 0x0 vector=253 Start context switch	Tells you when the system calls the you have tracked are called	hat
Column # CPU Time Stam, 10742 1 2121452.51 2121452.51 97443 1 2121452.51 97444 0 2121452.51 97445 1 2121452.51 97446 0 2121452.51 97447 0 2121452.51 97448 0 2121452.51 97448 0 2121452.51	Contains PID Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2269 migrate 288 2270 migrate 287 2270 migrate 287 2270 migrate 287	Next Prev Latency Event 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit_futex 6 dNh. reschedule_exit 6 dNh. rcu_utilization 6 dN. rcu_utilization	✓ Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 0x0 vector=253 Start context switch End context switch	Tells you when the system calls the you have tracked are called gives you details	hat
CPU Imme F P CPU Time Stam 197442 1 2121452.51 Stat21452.51 197443 1 2121452.51 Stat21452.51 197444 0 2121452.51 Stat21452.51 197445 1 2121452.51 Stat21452.51 197447 0 2121452.51 Stat21452.51 197448 0 2121452.51 Stat21452.51 197449 1 2121452.51 Stat21452.51	v contains v p Task PID 2269 migrate 288 2270 migrate 287 2270 migrate 287 2270 migrate 287 2270 migrate 288	Next Prev 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit_futex 6 dNh. reschedule_exit 6 dNh. rcu_utilization 6 dN. rcu_utilization 3 sys_enter	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 vector=253 Start context switch End context switch End context switch NR 1 (3, 7f7accba06ff, 1a, 0, 0, 7f7accb9fb26)	Tells you when the system calls the you have tracked are called gives you details	hat
Search: CPU Time Stam 0 CPU 2121452.51 107442 1 2121452.51 107443 1 2121452.51 107444 1 2121452.51 107445 1 2121452.51 107446 0 2121452.51 107447 0 2121452.51 107448 0 2121452.51 107449 0 2121452.51 107449 1 2121452.51 107449 1 2121452.51 107449 1 2121452.51 107449 1 2121452.51	contains PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2269 migrate 288 2269 migrate 287 2270 migrate 287 2270 migrate 287 2271 migrate 288 2271 migrate 288	Next Prev 13 d sched_wakeup 13 sys_exit 14 dth.n reschedule_entry 15 dth.n reschedule_exit 15 dth.n reschedule_exit 16 dth.n reschedule_exit 16 dth.n res_utilization 16 dth.n res_utilization 16 sys_enter 13 sys_enter_write		Tells you when the system calls the you have tracked are called gives you details	hat
Column # CPU Time Stam, 107442 1 2121452.51 107443 1 2121452.51 107444 0 2121452.51 107445 1 2121452.51 107446 0 2121452.51 107447 0 2121452.51 107448 0 2121452.51 107449 0 2121452.51 107449 1 2121452.51 107450 1 2121452.51 107451 0 2121452.51	contains PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2269 migrate 287 2270 migrate 287 2270 migrate 287 2270 migrate 287 2270 migrate 287 2271 migrate 288 2271 migrate 287	Next Prev Latency Event 3 sched_wakeup 3 4 dNb. 7 reschedule_entry 3 6 dNn. 7 reschedule_exit 6 dNn. 7 reschedule_exit 6 dNn. 7 reschedule_exit 6 dNn. 7 sys_enter 3 8 sys_enter 3 9 senter_write 6 d	Graph follows Infe migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 0x0 vector=253 Start context switch End context switch End context switch NR 1 (3, 7/7accba06ff, 1a, 0, 0, 7/7accb9fb26) fd: 0x0000003, buf: 0x7/7accba96ff, count: 0x000000 migrate:28796 [97] 5 ==> migrate:28802 [91]	Tells you when the system calls the you have tracked are called gives you details	hat
CPU Time Stam, 9742 1 2121452.51 19742 1 2121452.51 19744 0 2121452.51 19744 1 2121452.51 19744 1 2121452.51 19744 0 2121452.51 19744 0 2121452.51 19744 0 2121452.51 19744 1 2121452.51 19745 1 2121452.51 19745 1 2121452.51 19745 1 2121452.51 19745 1 2121452.51 19745 1 2121452.51 19745 1 2121452.51 19745 1 2121452.51	contains v p Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 287	Next Prev 3 sched_wakeup 3 sys_exit 6 dxlh. reschedule_entry 13 sys_exit_fletex 6 dxlh. reschedule_exit 6 dNn. reschedule_exit 6 dN. rcu_utilization 13 sys_enter 13 sys_enter 13 sys_enter 14 spis_enter 15 spis_enter 16 rewithe	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 ox0 vector=253 Start context switch End context switch End context switch NR 1 (3, 7/7accba06ff, 1a, 0, 0, 7/7accb9fb26) fd: 0x00000003, buf: 0x777accba06ff, count: 0x000000 migrate:28796 [97] S ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping	Tells you when the system calls the you have tracked are called gives you details	hat
Search: CPU Time Stam 0 CPU 2121452.51 107442 1 2121452.51 107443 1 2121452.51 107444 1 2121452.51 107445 1 2121452.51 107446 0 2121452.51 107447 0 2121452.51 107448 0 2121452.51 107449 1 2121452.51 107451 0 2121452.51 107452 1 2121452.51 107451 0 2121452.51 107451 1 2121452.51 107451 1 2121452.51	contains PID 2269 migrate 288 2269 migrate 288 2269 migrate 289 2269 migrate 289 2269 migrate 287 2270 migrate 289 2271 migrate 287 2271 migrate 288 2271 migrate 289	Next Prev 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 13 sys_exit futex 6 dNh. reschedule_exit 6 dNh. reschedule_exit 6 dNh. res_utilization 6 dN rcs_utilization 13 sys_enter 13 sys_exit	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0 vector=253 Start context switch End context switch NR 1 (3, 77a.ccba06ff, 1a, 0, 0, 777accb9fb26) fd: 0x00000003, buf: 0x717accba06ff, count: 0x000000 migrate:28796 (97) 5 ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping NR 1 = 26	Tells you when the system calls the you have tracked are called gives you details	hat
Search: Column # P CPU Time Stam, 197442 1 2121452.51 197443 1 2121452.51 197444 0 2121452.51 197445 1 2121452.51 197446 0 2121452.51 197447 0 2121452.51 197448 1 2121452.51 197449 1 2121452.51 197450 1 2121452.51 197451 0 2121452.51 197452 1 2121452.51 197451 0 2121452.51 197452 1 2121452.51 197453 1 2121452.51 197454 0 2121452.51	contains PID 2269 migrate 288 2269 migrate 289 2269 migrate 289 2269 migrate 289 2269 migrate 289 2270 migrate 287 2270 migrate 287 2270 migrate 287 2271 migrate 288	Next Prev Latency Event 3 sched_wakeup 3 sys_exit 6 dN. reschedule_entry 7 sys_exit 6 dN. reschedule_exit 6 dN. recutilization 6 dN. rcu_utilization 3 sys_enter_write 6 d sched_switch 3 sys_enter_write	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 Ox0 vector=253 Start context switch End context switch End context switch End context switch If: 0x0000003, buf: 0x717accba06ff, count: 0x000000 migrate:28796 [97] 5 ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping NR 1 = 26 x86ifpu: 0xfff90e60ac7f600 initialized: 1 xfeatures: 2 x	Tells you when the system calls the you have tracked are called	hat
Column # CPU Time Stam, 1 2121452.51 197443 1 197444 2121452.51 197445 1 197446 2121452.51 197445 1 197446 2121452.51 197447 0 197448 0 197449 1 197449 1 197449 1 197450 1 197451 1 197452 1 197453 1 197454 1 197455 1 197455 1 197454 1 197455 1 197455 1 197455 1 197455 1 197455 1 197455 1 197455 1 197455 1 197455 1 197455 1 19	contains v p Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 287 2271 migrate 288 2271 migrate 288 2271 migrate 288	Next Prev 13 d sched_wakeup 13 sys_exit 16 dNh. reschedule_entry 13 sys_exit 16 dNh. reschedule_entry 13 sys_exit_futex 16 dNn. reschedule_exit 16 dNn. reschedule_exit 16 dNn. reschedule_exit 13 sys_enter 13 sys_enter 13 sys_enter 13 sys_entith 14 sys_entith 15 sys_entith 16 d x86_fpu_regs_deactivated 15 sys_exit, write	Graph follows Infe migrate:28802 [91] success=1 CPU:000 NR 202 = 0 vector=253 0x0 vector=253 Start context switch End context switch End context switch NR 1 (3, 77accba06ff, 1a, 0, 0, 77accb9fb26) fd: 0x00000000, buf: 0x7f7accba06ff, count: 0x0000001 migrate:28796 [97] 5 ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping NR 1 = 26 x86/fpu: 0xfff90660ac7f600 initialized: 1 xfeatures: 2 xm 0x1a	Tells you when the system calls the you have tracked are called	hat
Search: CPU Time Stam 0 CPU 2121452.51 107442 1 2121452.51 107443 1 2121452.51 107444 1 2121452.51 107445 1 2121452.51 107446 0 2121452.51 107447 0 2121452.51 107448 0 2121452.51 107449 1 2121452.51 107451 0 2121452.51 107452 1 2121452.51 107453 1 2121452.51 107454 0 2121452.51 107455 1 2121452.51 107454 0 2121452.51 107455 1 2121452.51 107455 1 2121452.51 107456 0 2121452.51 107456 1 2121452.51	contains PID 2269 migrate 286 2269 migrate 286 2269 migrate 286 2269 migrate 286 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 288 2271 migrate 288 2271 migrate 288 2271 migrate 287 2272 migrate 287 2272 migrate 287	Next Prev 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit 6 dNh. reschedule_entry 6 dNh. reschedule_exit 6 dNh. reschedule_exit 6 dN. rcu_utilization 6 dN. rcu_utilization 6 dN. rcu_utilization 6 dN. sys_enter 3 sys_enter 3 sys_enter 3 sys_enter 3 sys_exit 6 sched_switch 3 sys_exit 6 sd6_fbu_regs_deactivated 2 sys_exit_write 6 d write_mar		Tells you when the system calls the you have tracked are called	hat
Column P CPU Time Stam, 107421 1 2121452.51 107443 1 2121452.51 107444 1 2121452.51 107445 1 2121452.51 107446 0 2121452.51 107447 0 2121452.51 107448 1 2121452.51 107449 1 2121452.51 107445 1 2121452.51 107445 1 2121452.51 107451 0 2121452.51 107452 1 2121452.51 107453 1 2121452.51 107454 1 2121452.51 107455 1 2121452.51 107454 0 2121452.51 107455 1 2121452.51 107456 0 2121452.51 107456 0 2121452.51 107457 0 2121452.51	contains PID 2269 migrate 288 2269 migrate 288 2269 migrate 288 2269 migrate 287 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 288 2272 migrate 288 2271 migrate 288 2272 migrate 288 2272 migrate 288 2272 migrate 288	Next Prev 13 d sched_wakeup 13 sys_exit 14 sched_wakeup status 15 sys_exit 16 dth. reschedule_exit 16 dth. reschedule_exit 16 dth. rcu_utilization 13 sys_enter_write 16 dth. sched_switch 13 sys_exit_write 16 dth x86_fpu_regs_deactivated 13 sys_exit_write 16 dth write_msr	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0 Vector=253 Start context switch End context switch End context switch Id x00000003, buf: 0x7f7accba06ff, count: 0x0000000 migrate:28796 [97] 5 ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping NR 1 = 26 x86/fpu: 0xfff90e60ac7f600 initialized: 1 xfeatures: 2 x 0x1a c0000100, value 7f7acd3a3700 x80(fpu: 0xfff90e8c4e35780 initialized: 1 xfeatures: 2 x	Tells you when the system calls the you have tracked are called	hat
Column # CPU Time Stam, 1 2121452.51 197442 1 2121452.51 197443 1 197444 1 197445 1 197445 1 197445 1 197446 1 197447 0 1212452.51 197448 0 1212452.51 197449 1 1212452.51 197450 1 1212452.51 197451 1 1212452.51 197451 1 1212452.51 197451 1 1212452.51 197451 1 1212452.51 197455 1 1212452.51 197455 1 1212452.51 197455 1 197456 1 197457 1 197457 1 197456 </td <td>contains PID P Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 289 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 288 2272 migrate 288 2271 migrate 288 2272 migrate 287</td> <td>Next Prev 13 d sched_wakeup 13 sys_exit 16 dNh. reschedule_entry 13 sys_exit_fitex 16 dNh. reschedule_exit 16 dNn. reschedule_exit 16 dNn. reschedule_exit 16 dN. reu_utilization 13 sys_enter 13 sys_enter 14 sched_switch 15 sys_enter 16 d x86_fpu_regs_deactivated 13 sys_exit, write 16 d x86_fpu_regs_activated 13 sys_enter</td> <td>Coraph follows Infe migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0 Vector=253 Start context switch End context switch End</td> <td>Tells you when the system calls the you have tracked are called</td> <td>hat</td>	contains PID P Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 289 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 288 2272 migrate 288 2271 migrate 288 2272 migrate 287	Next Prev 13 d sched_wakeup 13 sys_exit 16 dNh. reschedule_entry 13 sys_exit_fitex 16 dNh. reschedule_exit 16 dNn. reschedule_exit 16 dNn. reschedule_exit 16 dN. reu_utilization 13 sys_enter 13 sys_enter 14 sched_switch 15 sys_enter 16 d x86_fpu_regs_deactivated 13 sys_exit, write 16 d x86_fpu_regs_activated 13 sys_enter	Coraph follows Infe migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0 Vector=253 Start context switch End	Tells you when the system calls the you have tracked are called	hat
Generative CPU Time Stam 9 CPU 2121452.51 107443 1 2121452.51 107444 1 2121452.51 107445 1 2121452.51 107444 0 2121452.51 107445 0 2121452.51 107447 0 2121452.51 107448 0 2121452.51 107449 1 2121452.51 107451 0 2121452.51 107452 1 2121452.51 107453 1 2121452.51 107454 0 2121452.51 107455 1 2121452.51 107454 0 2121452.51 107455 1 2121452.51 107456 0 2121452.51 107457 0 2121452.51 107458 1 2121452.51 107458 1 2121452.51 107458 1 2121452.51 107458 1 <t< td=""><td>contains PID 2269 migrate 284 2269 migrate 284 2269 migrate 284 2269 migrate 284 2270 migrate 287 2270 migrate 287 2271 migrate 284 2272 migrate 287 2272 migrate</td><td>Next Prev 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit 6 dNh. reschedule_exit 6 dNh. reschedule_exit 6 dNh. reschedule_exit 6 dN. rcu_utilization 6 dN. rcu_utilization 6 dN. rcu_utilization 6 d sys_enter 3 sys_enter 3 sys_enter 3 sys_exit 6 d sd6.fpu_regs_deactivated 3 sys_exit 6 d x86.fpu_regs_activated 3 sys_enter 4 d x95_enter 3 sys_enter</td><td></td><td>Tells you when the system calls the you have tracked are called gives you details</td><td>hat</td></t<>	contains PID 2269 migrate 284 2269 migrate 284 2269 migrate 284 2269 migrate 284 2270 migrate 287 2270 migrate 287 2271 migrate 284 2272 migrate 287 2272 migrate	Next Prev 3 sched_wakeup 3 sys_exit 6 dNh. reschedule_entry 3 sys_exit 6 dNh. reschedule_exit 6 dNh. reschedule_exit 6 dNh. reschedule_exit 6 dN. rcu_utilization 6 dN. rcu_utilization 6 dN. rcu_utilization 6 d sys_enter 3 sys_enter 3 sys_enter 3 sys_exit 6 d sd6.fpu_regs_deactivated 3 sys_exit 6 d x86.fpu_regs_activated 3 sys_enter 4 d x95_enter 3 sys_enter		Tells you when the system calls the you have tracked are called gives you details	hat
CPU Time Stam, 0 CPU Time Stam, 97442 1 2121452.51 97443 1 2121452.51 97444 0 2121452.51 97445 1 2121452.51 97446 0 2121452.51 97447 0 2121452.51 97448 0 2121452.51 97449 1 2121452.51 97449 1 2121452.51 97445 1 2121452.51 97450 1 2121452.51 97451 1 2121452.51 97452 1 2121452.51 97453 1 2121452.51 97454 0 2121452.51 97455 1 2121452.51 97456 0 2121452.51 97457 0 2121452.51 97458 1 2121452.51 97459 1 2121452.51 97459 1 2121452.51 97450 1	contains PID 2269 migrate 288 2269 migrate 289 2269 migrate 289 2269 migrate 289 2270 migrate 287 2271 migrate 287 2271 migrate 289 2272 migrate	Next Prev 13 sched_wakeup 13 sys_exit 14 sys_exit sys_exit 15 dth. reschedule_entry 13 sys_exit_futex 16 dth. reschedule_exit 16 dth. rcu_utilization 13 sys_enter 13 sys_enter 13 sys_enter 14 sys_enter 15 sys_enter 16 dt. sched_switch 17 sys_enter 18 sys_enter 19 sys_enter 10 sys_enter 13 sys_enter 14 sys_enter 16 16 17 sys_enter 18 <td>Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0 vector=253 Start context switch End context switch End context switch End context switch fd: 0x00000003, buf: 0x7f7accb0fb26) fd: 0x00000003, buf: 0x7f7accb0ff, count: 0x0000001 migrate:28796 [97] 5 ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping NR 1 = 26 x86/fpu: 0xfff90e60ac7f600 initialized: 1 xfeatures: 2 x8 Ox1a co000100, value 7f7acd3a3700 x86(fpu: 0xfff90e8c4e35780 initialized: 1 xfeatures: 2 x8 NR 35 (7f7accb9fda0, 0, 0, 77accb9fdd0, 0, 7f7accb9fb rqtp: 0x7f7accb9fda10.20000000 httime=0xfff85d501612090 clockid=CLOCK MONOTOR</td> <td>Tells you when the system calls the you have tracked are called gives you details</td> <td>hat</td>	Graph follows Info migrate:28802 [91] success=1 CPU:000 NR 202 = 0 Vector=253 Ox0 vector=253 Start context switch End context switch End context switch End context switch fd: 0x00000003, buf: 0x7f7accb0fb26) fd: 0x00000003, buf: 0x7f7accb0ff, count: 0x0000001 migrate:28796 [97] 5 ==> migrate:28802 [91] tracing_mark_write: thread 7 iter 42 sleeping NR 1 = 26 x86/fpu: 0xfff90e60ac7f600 initialized: 1 xfeatures: 2 x8 Ox1a co000100, value 7f7acd3a3700 x86(fpu: 0xfff90e8c4e35780 initialized: 1 xfeatures: 2 x8 NR 35 (7f7accb9fda0, 0, 0, 77accb9fdd0, 0, 7f7accb9fb rqtp: 0x7f7accb9fda10.20000000 httime=0xfff85d501612090 clockid=CLOCK MONOTOR	Tells you when the system calls the you have tracked are called gives you details	hat
Column Image: Column (Column) OPU Time Stam 109744 1 2121452.51 107444 1 2121452.51 107444 1 2121452.51 107444 1 2121452.51 107447 0 2121452.51 107448 0 2121452.51 107449 1 2121452.51 107449 0 2121452.51 107450 1 2121452.51 107451 1 2121452.51 107452 1 2121452.51 107453 1 2121452.51 107454 1 2121452.51 107455 1 2121452.51 107456 1 2121452.51 107456 1 2121452.51 107456 1 2121452.51 107456 1 2121452.51 107456 1 2121452.51 107456 1 2121452.51 107456 1 2121452.51	▼ contains ▼ p Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 289 2270 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 288 2272 migrate 288 2271 migrate 288 2272 migrate </td <td>Next Prev 13 </td> <td>Coraph follows</td> <td>Tells you when the system calls the you have tracked are called gives you details</td> <td>hat</td>	Next Prev 13	Coraph follows	Tells you when the system calls the you have tracked are called gives you details	hat
CPU Time Stam P742 1 212145251 97443 1 212145251 97444 1 212145251 97445 1 212145251 97446 0 212145251 97447 0 212145251 97448 0 212145251 97449 1 212145251 97449 1 212145251 97450 1 212145251 97451 0 212145251 97452 1 212145251 97453 1 212145251 97454 0 212145251 97455 1 212145251 97456 0 212145251 97457 0 212145251 97458 1 212145251 97459 1 212145251 97459 1 212145251 97450 1 212145251 97450 1 212145251 97450 1	contains PI P Task PID 2269 migrate 288 2269 migrate 288 2269 migrate 289 2269 migrate 287 2270 migrate 287 2271 migrate 287 2271 migrate 288 2271 migrate 289 2271 migrate 289 2271 migrate 289 2272 migrate 289 2272 migrate 289 2272 migrate 289 2272 migrate 289 2273 migrate 289 2273 migrate 289	Next Prev 3		Tells you when the system calls the you have tracked are called gives you details	hat





Wall-clock time

Question: What else does impact the wallclock time other than system and user times?

Wall clock – system – user =

I/O + preemption + suspension + other interferences

- I/O: time waiting for performing an I/O interaction (including reading from or writing to files, ports, etc.)
- **Preemption**: time waiting for other tasks to execute
- Suspension: time being blocked (e.g., when a task wants to access a data that is protected by a lock) or being suspended due to an I/O operation
- **Other interferences:** time spent on accessing memory bus of a memory bank, etc.

Question: How can you control the impact of these delays on your measurements?





Wall-clock time

Wall clock – system – user =

I/O + preemption + suspension + other interferences

One approach is to try to run your process in isolation (when no other process is running)

Question: Would this be enough?

hey, I have designed a stabilizer that processes 1000 samples per millisecond on an Arduino Mega

Oh, even when your drone is flying?

No way! It happens only in the lab when I turn off every other process in the system!







m

Wall-clock time

surf:~\$ /usr/bin/X11/time ./parallelQuicksort2

9.76user 10.51system 0:06.11elapsed 331%CPU

(Oavgtext+Oavgdata 158268maxresident)k

Oinputs+Ooutputs (Omajor+7599minor)pagefaults Oswaps

Question: Why wall clock time < system + user time?

This system could have multiple cores and execute the program in parallel on multiple cores. In that case, the wall-clock time becomes smaller than user or system times.







Drawbacks of UNIX time command

• The time command has poor resolution

• "Only" milliseconds

"time" times the whole code

- Often ES code is not terminating (i.e. runs forever)
- Sometimes we're only interested in timing some part of the code, for instance the one that we are trying to optimize



Timing with gettimeofday

- gettimeofday from the standard C library
- Measures the number of microseconds since midnight, Jan 1st 1970, expressed in seconds and microseconds

Can be used to time sections of code

- Call gettimeofday at beginning of section
- Call gettimeofday at end of section
- Compute the time elapsed in microseconds: (end.tv sec*1000000.0 + end.tv usec start.tv sec*1000000.0 - start.tv usec) / 1000000.0)

Have a look at clock_gettime()

Measures the number of nanoseconds since midnight, Jan 1st 1970, expressed in seconds and nanoseconds. Obviously not precise at the nanosecond level but much better than gettimeofday

https://linux.die.net/man/2/gettimeofday

https://linux.die.net/man/3/clock gettime







Using timers

Assumption: There is no other process or an operating system in this example. The goal is to measure the runtime of the middle instruction that runs on an Arduino Mega.



- Timers are not infinitely accurate
 - All clocks have a granularity
 - The error in a time measurement, even if everything is perfect, may be the size of this granularity (sometimes called a clock tick)
- Always know your clock granularity
- Ensure that your measurement is for a long enough duration (say 100x the "tick")

source: http://wgropp.cs.illinois.edu/courses/cs598-s16/lectures/lecture06.pdf







Timers and rollover



Counter value is read from a memory location.



Timer Rollover

• Occurs when an n-bit counter undergoes a transition from its maximum value $2^n - 1$ to zero.







Timers and rollover



There is a trade-off between rollover time and granularity of the timer







The Art of Data Gathering



Image source: https://streetfightmag.com/2019/03/25/the-fundamental-ethical-stakes-of-data-collection-and-ad-targeting/#.Xc68AvlKhaQ





Why should you repeat your measurements?

Errors

Error is the difference between the **measured value** and the 'true value' of the thing being measured

Uncertainties

Uncertainty is a <u>quantification</u> of the *doubt about the measurement result*



- Whenever possible, we try to correct for any known errors
- But any error whose value we do not know is a source of uncertainty



Measurement process

Assumption: there is no other process in the system. There is no operating system. The measurements are the runtime of the only task in the system.





Cold start

- Code may still be on disk, and not even loaded into memory.
- Data may be in slow memory rather than fast, e.g., cache, (note: this may be wrong or right for what you are measuring)

Question: How to solve this issue if we want to measure the long-run average?

Let the program run for a while and then start measuring

Note: usually you need to put effort to ensure that the data in the intended part of the memory hierarchy.





The effect of hidden or ignored parameters

TUDelft



Embedded and

Networked Systems

35 67

Question: What is it?

Measurement-induced perturbations

- The system resources consumed by the measurement tool itself as it collects data may strongly affect the system's performance.
- Tracing produces the highest level of perturbation (both CPU and disk are used) on
 - time measurements,
 - spatial and temporal memory access (cache flush, different paging, . . .), or
 - system response time (and thus on workload characterization).
- The larger the overhead, the more likely the system behavior will be modified.

On embedded systems, writing outputs on the terminal will heavily impact the system's behavior

In Arduino Mega, storing time in an array takes about 2~3 microseconds, while writing it using printf() is about 70~80 microseconds







Question: What is it?

Measurement-induced perturbations



When measuring time:

- Do not write the measured data on the output while gathering samples
 - **Solution:** store the data in an array and then output the array after gathering samples
- Do not use dynamic memory allocation for storing samples (e.g., malloc())







Question: how to solve it?

Problem: smart compiler

If the result of the computation is not used, the compiler may eliminate the code.









The Art of Data Summarization









Data summarization

- 1. How to report the performance as a **single number**?
 - Is specifying the mean the correct way?
 - How to report the variability of measured quantities?





Data summarization

- 1. How to report the performance as a **single number**?
 - Is specifying the mean the correct way?
 - How to report the variability of measured quantities?





Interpreting results









Data summarization









Confidence interval



Mean = 4.27

measured value







Histogram





Box plots





Further data representations



source: https://serialmentor.com/dataviz/boxplots-violins.html







Common Mistakes









Common mistakes

Biased goals

- Example:
 - To show that OUR system is better than THEIRS
- **Mistake:** Finding the metrics and workload such that OUR system turns out better rather than finding the right metrics and workloads.

Unsystematic Approach

• Selecting system parameters , factors, metrics, and workload arbitrarily

Incorrect Performance Metrics

- Example: Comparing Two CPUs based on the throughput (MIPS)
 - CISC
 - RISC







Common mistakes

Unrepresentative workload

• Example: if packet in the network are generally mixture of long and short, workload should consist of short and long packet sizes, not just long ones.

Not performing a sensitivity analysis on the factors

- Sensitivity analysis answers the question, "if the factors deviate from the expectations, what will the effect be on the system being analyzed".
- Putting too much emphasis on the results of the analysis, **presenting it as fact rather than** evidence.
- Ignoring variability of results when reporting the outcome
 - Only reporting mean value

Omitting Assumptions and limitations

- Assumption and limitations of the analysis are often omitted in the final report.
- This may lead the user to apply the analysis to another context where assumption will not be valid anymore.



Common mistakes in performance evaluation







A systematic approach to performance evaluation

- 1. State goals and define the system
- 2. List services/<u>functionalities</u> that you want to evaluate
- 3. Select <u>metrics</u>
- 4. List workload parameters
- 5. Select <u>factors</u> and their values
- 6. Select/build the workload
- 7. Design the experiment
- 8. Gather measurement data
- 9. Analyze the data
- 10. Present the results

effect model regression model









The Emb. Syst. perspective

- ES is about making the right trade off
 - choose the best parameter setting(s), system config, ...
- Performance vs. cost





Pareto Optimality

Definition 1 (Dominance relation) Let $f, g \in \mathbb{R}^m$. Then f is said to dominate g, denoted as $f \succ g$, iff



Definition 2 (Pareto set)

Let $F \subseteq \mathbb{R}^m$ be a set of vectors. Then the Pareto set $F^* \subseteq F$ is defined as follows: F^* contains all vectors $g \in F$ which are not dominated by any vector $f \in F$, i.e.

$$F^* := \{ g \in F \mid \not \exists f \in F : f \succ g \}$$

$$\tag{1}$$





Step 1: open www.kahoot.it in your browser (phone or laptop)
Step 2: enter the pin code and then a nickname

Assignment 1









Assignment 1

- Is focused on evaluating ROS 2.0's communication overheads
- ROS 2.0 is the latest generation of ROS (Robotics operating system)
- The assignment will be released Monday
- Lab starts from tomorrow $\textcircled{\odot}$







ROS 2.0 – example





ROS 2.0 – under the hood





Case study for Assignment 1



• Factors

- ROS configurations
 - 3 DDS choices: FastRTPS, OpenSplice, Connext

• OS configurations

- process priority
- Smart CPU frequency scaling: active, inactive





• Size of message

• 4KB, 16KB, ..., 4MB (in total 8 values)

68 67

Questions?









Sources used to make the slides

- Jean-Marc Vincent and Arnaud Legrand
 - Performance measurements of computer systems: tools and analysis, M2R PDES, 2015
 - <u>http://polaris.imag.fr/arnaud.legrand/teaching/2015/M2R_EP_measurements.pdf</u>
- Holger Brunst and Matthias S. Mueller
 - Performance Analysis of Computer Systems
 - Center for Information Services and High Performance Computing (ZIH), TU Dresden
 - <u>https://tu-dresden.de/zih/ressourcen/dateien/lehre/ws1112/lars/vorlesungen/lars_lecture_02_requirements-metrics-techniques.pdf?lang=en</u>
- William Gropp
 - Course: Designing and Building Applications for Extreme Scale Systems, Illinois, 2016
 - <u>http://wgropp.cs.illinois.edu/courses/cs598-s16/index.htm</u>
- Raj Jain
 - The Art of Computer Systems Performance Analysis
 - <u>https://www.cse.wustl.edu/~jain/books/perf_sli.htm</u>
- Other slides and sources:
 - <u>https://slideplayer.com/slide/7709208/</u>
 - <u>https://www.slideshare.net/ShubraBansal/uncertainity-in-measurement</u>
 - <u>https://serialmentor.com/dataviz/boxplots-violins.html</u>

Thank you



