6th Seminar – Performance Modelling, Visual Data Analysis

1 Server's performance

We measured the following performance metrics on a server:

Time of measure [ms]	500	600	700	800	900
Requests processed in the last 100 ms [request]	11	12	21	18	20
Average serving time in the last 100 ms [ms]	15	20	21	25	27
CPU utilization of the last 100 ms [%]	12	13	16	17	19
HDD I/O utilization of the last 100 ms [%]	55	63	87	61	73

a. Based on the available data, which server resource seems to be the bottleneck and why?
b. What is the server's throughput at the time of the first measure? What is the average and the median of the throughput based on these 5 measurements?
c. What estimate can we provide for the average number of requests being served at the same time based on these 5 measurements?

2 Social website

We operate a social web company. Due to its recent rising popularity, response times have increased greatly. The business goal is to have 1500 simultaneous user requests served with less than 4 seconds of response time in average.

a. What minimal throughput should the service infrastructure be designed for, if delays outside our infrastructure (network traffic latency, HTML rendering on the client side) can be estimated as 1 second?
b. According to measurements, an average user request in the redesigned web site takes 20 ms CPU time on the web server, and occupies the database server for 12,5 ms. Currently we have 15 web servers to handle the requests, while the database is replicated to 5 machines. Assuming linear scalability, how much additional units of each kind of server should we buy to meet the above goal?
c. (∗) Calculate the utilization of each kind of server in the extended system. If the goal is to push the average utilization of the servers below 50% even during peak hours, do we need to scale out further?
d. Let's consider only 2 webservers and 3 database servers. Create state-based models about the resources in the infrastructure that model the availability of the resources (available or in use). What design decisions do we face? What are the pros and cons of the choices?

3 Sensor network (previous exam exercise) – data analysis

We have an agriculture sensor network that helps us to track the states of our open-field, glasshouse and foil tent areas based on some measured values (temperature, humidity, luminous intensity, wind speed, detected pests, etc.).

Date	Temp. [°C]	Hum. [%]	Pests [piece]
2015. 05. 04. 08:00	18	66,00	3
2015. 05. 04. 09:00	20	65,75	6
2015. 05. 04. 10:00	20	65,75	8
2015. 05. 04. 11:00	20	65,50	9
2015. 05. 04. 12:00	20	65,50	5
2015. 05. 04. 13:00	21	65,00	12
2015. 05. 04. 14:00	21	64,70	5
2015. 05. 04. 15:00	21	64,70	6
2015. 05. 04. 16:00	21	64,60	7
2015. 05. 04. 17:00	22	64,00	2

a. Unfortunately the middle values (median) of Monday, May 4th are missing from the figures. Draw them based on the data in the table!
b. Interpret the diagrams: which variable's/variables' first quartile is strictly monotonic in time?
c. (Extra task.) We would like to compare the temperature values and pest numbers of Monday in a parallel coordinates diagram.

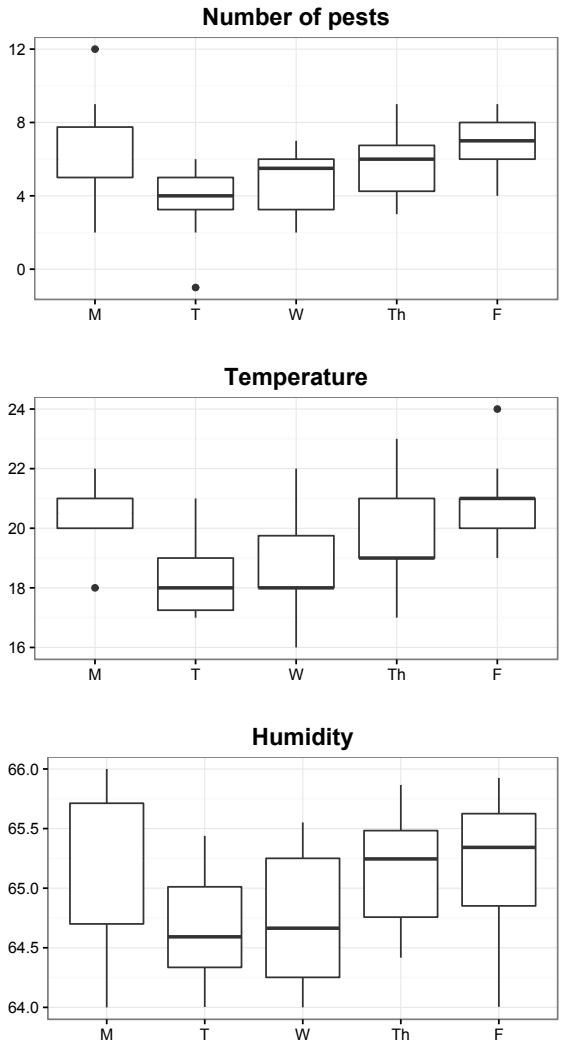

4 Sensor network (previous exam exercise) – perf. analysis (∗)

(Performance analysis exercises related to Exercise 3.) The different types of sensors provide data from a 100 meters radius around their location. The sensors forward their timestamped data to the central server through a radio communication-based network. The central server processes the requests then archives them to a storage unit. Our organization installed 4500 sensors and each one sends one measurement data in every minute. The system can successfully handle this load. The radio communication network can forward 100 measurement data every second. The central server's CPU is idle (not doing anything) in 75% of the time. Writing a measurement data to the storage unit takes 8 ms.

a. How many measurement data in a second is the current throughput of the system?
b. What is the throughput, maximum throughput and utilization of the radio network, CPU and storage?
c. How many more sensors can we install (to improve the measurement accuracy) without upgrading our infrastructure? Assume linear scaling!
d. The radio network uses smart encoding, so more than one sensor can forward data at the same time. How many sensors are forwarding data at the same time (overlapping) over the network currently and during maximum load, if a forwarding takes 40 ms?