

Excercise Sheet 6

(To be processed until 28.11.)

Lecture Distributed Systems Winter Term 2024/25

Exercise 1: Load balancing strategies

- How do static, dynamic and preemptive dynamic load balancing differ from each other? Which advantages and disadvantages do the methods have and which prerequisites must be given in each case in order to be able to apply one of the methods?
- All load distribution strategies base their decisions on some *load metric*. The simplest, often considered load metric for theoretical analysis of algorithms is the number of (ready) processes per node. Why is this metric sometimes not sufficient in practice? What would be a better metric and what new problems arise then?

Exercise 2: Graph partitioning and list scheduling


In the lecture (chapter 5.1.1) two types of distributed scheduling algorithms based on graphs were presented:

- load distribution through graph partitioning and
- list scheduling.

What similarities and differences can you find between these two classes of algorithms? Under what conditions are they applicable?

Exercise 3: Loda Distribution via Graph Partitioning (**Mandatory exercise for 6 CP, submit via moodle!**)


Remember the load distribution by graph partitioning (see lecture chapter 5.1.1). Assume that in the figure below the process H changes from node 2 to node 3. How large is the resulting (total) network traffic between the nodes?



Exercise 4: List Scheduling (**Mandatory exercise for 6 CP, submit via moodle!**)

The lecture presented different approaches to distributed scheduling (see section 5.1.1). Two frequently used list scheduling algorithms are: *High Level First with Estimated Time* (HLFET) and *Earliest Task First* (ETF). The first one was already shown in the lecture (see animation in section 5.1.1: List Scheduling with HLFET). For the second algorithm you can find a detailed description in the article by Hagraš and Janacek (accessible via the lecture page or via the URL <https://ojs.cvut.cz/ojs/index.php/ap/article/download/490/322>) on page 19.

Given is a system with 3 processors and the DAG shown below, which models a program consisting of several dependent tasks. The nodes contain tasks with execution times and the edges show the necessary communication with transmission time.



How are the tasks scheduled on the three processors (with HLFET / with ETF)? Draw one Gantt chart for each. Suppose a local communication doesn't cost any time.

Exercise 5: Process Migration

Some multicomputers allow the migration of running processes from one node to another. Is it sufficient to stop the process, freeze the memory image, and transfer the whole thing to another node? Specify two non-trivial problems that must be solved for the model to work.

Exercise 6: Policies of Dynamic Load Balancing Systems

- a) Describe the four policies (strategies) that are usually components of a dynamic load balancing system. Specify for each strategy:
 - an intuitive description of the purpose of this strategy,
 - an example of a specific implementation of this strategy.
- b) Why are receiver-initiated procedures for dynamic load balancing sometimes undesirable?