Parallel and Distributed Algorithms and Programs
TD n°7 - Scheduling (3)

Hadrien Croubois
hadrien.croubois@ens-lyon.fr

Aurélien Cavelan
aurelien.cavelan@ens-lyon.fr

8/1/2016

All documents are available on my website: `http://hadriencroubois.com/#Teaching`

The DAG

In today's practice session, we will consider the following DAG of tasks that we want to schedule on 3 identical processors.

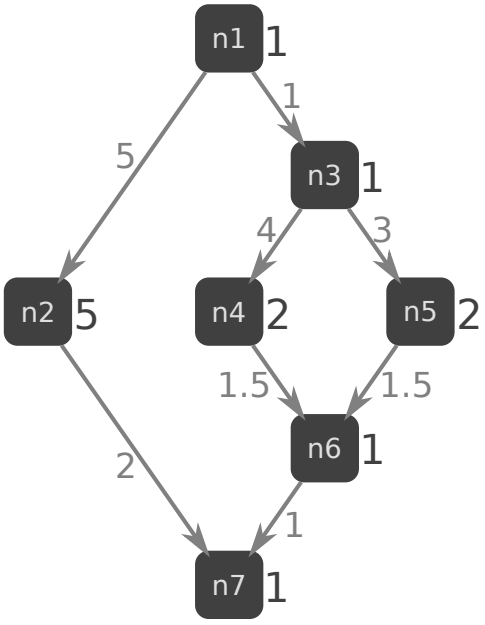

Figure 1: The Dag

Part 1

Scheduling without communications

In this section, we do not consider the communications costs.

Question 1

a) Compute the bottom level for each node.

b) Schedule the DAG on our 3 processors using a list heuristic. What is the make-span of our schedule ? Is it optimal ?

Part 2

Scheduling with communications

From now on we will consider the communication costs which have to be accounted for when two adjacent tasks are scheduled on different processors.

2.1

Critical path scheduling

By modifying the computation of the bottom level to consider communications, we can apply our heuristic to consider communications.

Question 2

a) What communications should be taken into consideration when computing the bottom level ?

b) Compute the bottom level for each node.

c) Schedule the DAG on our 3 processors using a list heuristic. What is the make-span of our schedule ?

2.2

Modified critical path scheduling

Sometimes, it is worth waiting to schedule a task on a busy processor rather then using the first processor available. This is what the modified critical path proposes.

Question 3

a) Which wrong decision, made in the previous section, would be avoided by using this new heuristic ?

b) Using this approach, propose a new schedule for our DAG. What is it's make-span ?

2.3

Clustering : Sarkar's algorithm

List heuristics are simple but not very efficient. In order to increase the performance of distributed systems, many clustering have been developed. By analysing the graph, Sarkar's algorithm proposes to remove the communication cost of some edges by forcing the neighbouring tasks to be executed on the same processor. Clustering algorithms may vary in how they decide which edge to remove (or which cluster to merge).

A very simple clustering algorithm is Sarkar's algorithm. Considering all edges in decreasing order of cost, we consider removing them if it doesn't increase the execution time.

Question 4

a) Apply Sarkar's algorithm to our DAG of tasks. which cluster do you get ?

b) Schedule the execution of those clusters on our 3 processors. What is the make-span of this schedule ?

c) Is this schedule optimal ?