

Big Data Management and Analytics Assignment 5

(a) Given a Hadoop setup with several DataNodes.

One Data Node d is broken. How is this node failure being detected by Hadoop?

- DataNodes send periodically Heartbeat messages to the NameNode.
- Absence of Heartbeats \rightarrow DataNode is no longer reachable \rightarrow Marked as dead

How does the NameNode react to d with regards to I/O requests?

- The NameNode will no longer send any new I/O-requests to the ,dead'-labeled DataNode d

(a) Given a Hadoop setup with several DataNodes.

What can be stated regarding the availability of the data which was stored on the DataNode d?

- The data which was registered to the dead DataNode is no longer available to the HDFS

What can be the effects regarding the replication factor and what action is taken by the NameNode?

- A DataNode failure may lead to a drop of the replication factor of some blocks below their specified value

Assignment 5-1

(b) The free space on one DataNode d falls below a certain threshold. What could be a potential scheme in context of rebalancing available space among DataNodes?

- One potential scheme could be to move data from one particular DataNode to another

Assignment 5-1

(c) Given the case of a highly frequent access on a specific file. What could be a scheme in order to rebalance the frequent file access?

- The dynamic creation of additional replicas could counter the highly frequent file access

(d) From a DataNode a data block is fetched. How is checked if the arrived data block is corrupted?

- An implemented checksum checking on the contents of the HDFS files is used for detecting failures in data block integrity

(e) By some failure the NameNode is no longer available. What are the implications regarding the functionality of the Hadoop setup?

- Given the case that the NameNode fails, the entire HDFS cluster is no longer functional.

Is any manual intervention necessary? If yes, explain why.

- Due to the fact that no automatic restart and failover is available for the NameNode, a manual intervention is required.

(f) What are FsImage and EditLog?

- FsImage is a file which stores the complete file system namespace which encompasses the mapping of blocks to files and the file system properties.
- The NameNode stores persistently every modification to the filesystem metadata in the EditLog

What would be the implications if any of these files are corrupted?


- The entire HDFS instance becomes non-functional

What is done in order to be resilient against the case of corrupted FsImage and EditLog files?

- Several copies of both files are maintained

Assignment 5-1

(g) In a data center we have two racks r1 and r2. Each of the racks contains several DataNodes. The client is located on a DataNode d1 on r1. The DataNode with the client fails. Which of the replicas according to Hadoop's replica placement strategy are available?



- (a) Given $k \in \mathbb{N}$ number of tasks which require at least $b \in \mathbb{N}$ GB of RAM, and $n \in \mathbb{N}$ nodes which provide $a_1, a_2 \dots a_n$ GB of memory. How can the memory of n nodes be distributed among k tasks such that no task has less than ε_t GB of RAM? If a task does not get the required amount of memory, the task waits until sufficient memory is available. What is the runtime complexity?

BINPACKING problem:

Given a set of bins $S_1, S_2 \dots$ where all bins have the same volume V and a list of n items with sizes $a_1, \dots a_n$ to pack. Find an integer number of bins B such that the sum of items per bin does not exceed V .

→ Combinatorial problem is NP-complete!

BINPACKING problem:

Given a set of bins $S_1, S_2 \dots$ where all bins have the same volume V and a list of n items with sizes a_1, \dots, a_n to pack. Find an integer number of bins B such that the sum of items per bin does not exceed V .

Reformulate problem:

Given a set of bins $S_1, S_2 \dots$ where all bins have the same **minimum volume** V_{min} and a list of n items with sizes a_1, \dots, a_n to pack. Find an integer number of bins B such that the sum of items per bin **does exceed or at least equals** V_{min} .

- Is known as the ‚dual version of bin packing‘ or ‚bin covering‘
- Also NP-complete!


Assignment 5-2

(b) In which field(s) of computer science does the resource scheduling problem occur?

- Operating systems (process scheduler)
- Parallel computing

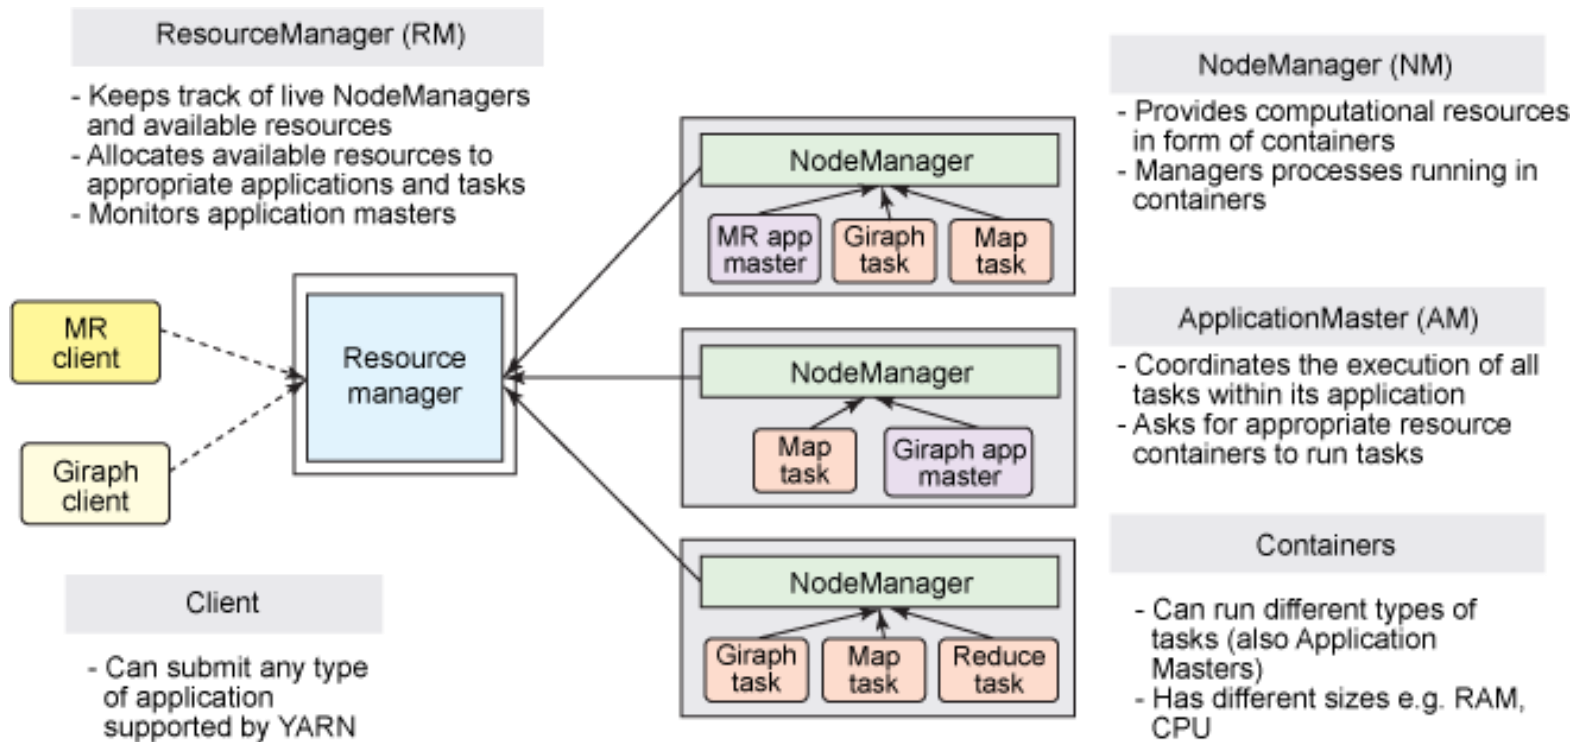
- (c) What is the purpose of YARN and in how far does it relate to the mentioned resource distribution problem?
- Is the ,new' resource manager for Hadoop
 - Decouples core functionalities of the JobTracker into two deamons:
 - ResourceManager
 - Application Master

(c) What is the purpose of YARN and in how far does it relate to the mentioned resource distribution problem?

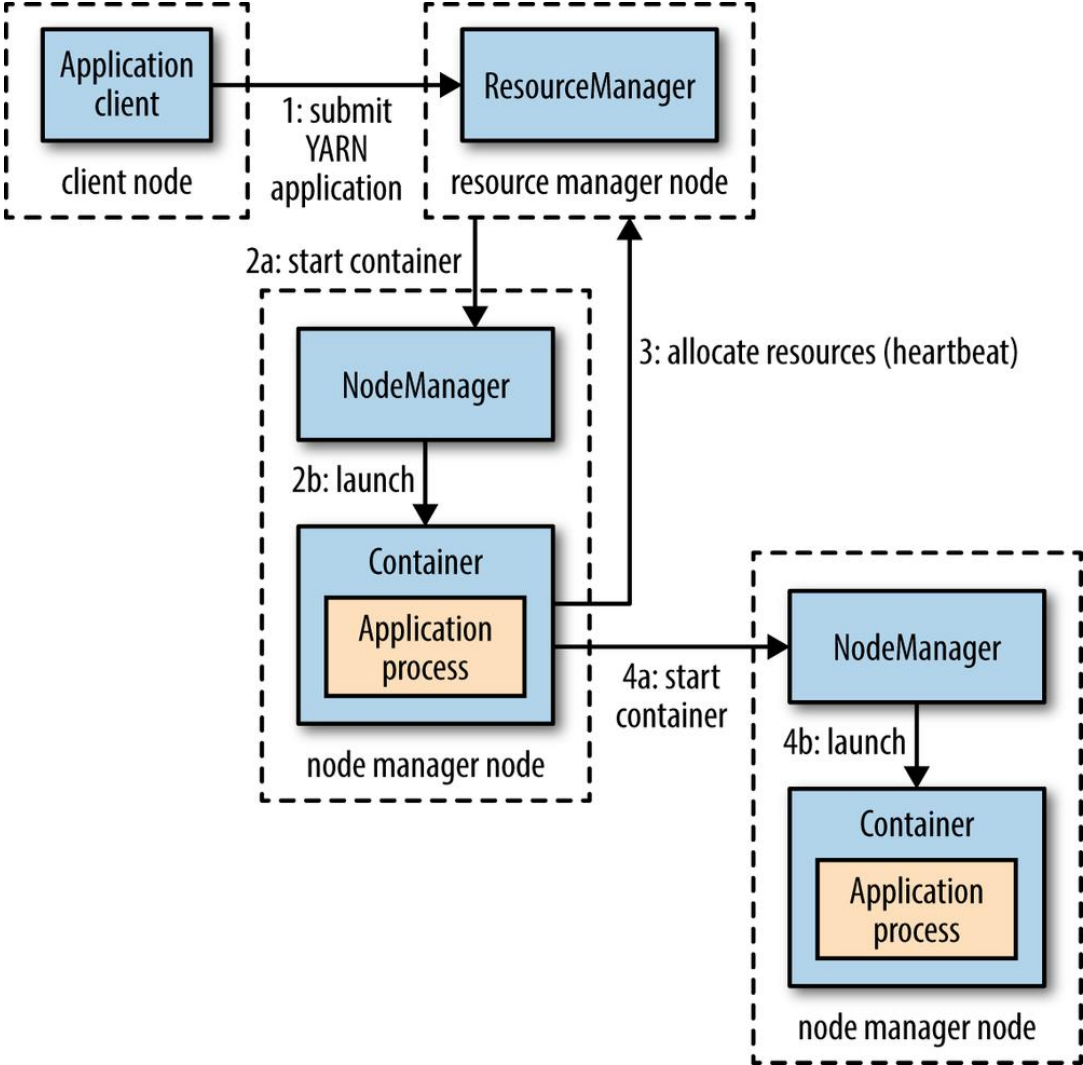


<https://www.safaribooksonline.com/library/view/hadoop-the-definitive/9781491901687/ch04.html>

(c) What is the purpose of YARN and in how far does it relate to the mentioned resource distribution problem?



Assignment 5-2




Assignment 5-2

MapReduce1	YARN
Jobtracker	ResourceManager ApplicationMaster TimelineServer
Tasktracker	NodeManager
Slot	Container
Can run up to 4000 nodes and 40000 tasks	Scales up to 10000 nodes and 100000 tasks

YARN – Scheduler options


i. FIFO Scheduler



- Requests for the first application in the queue are allocated first → if requests are finished the next application is served
- pro: doesn't require any configuration
- con: not suitable for shared clusters

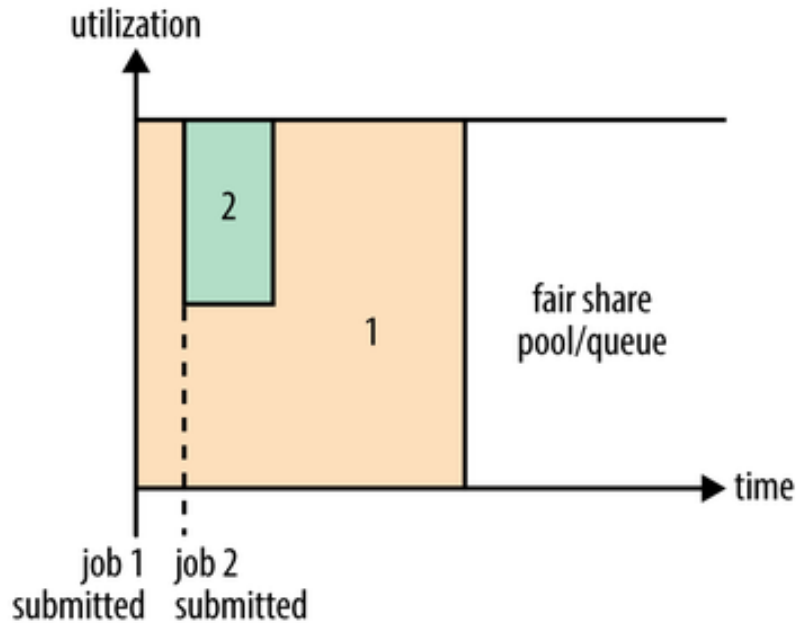
Within each queue applications are scheduled according to FIFO

ii. Capacity Scheduler



- Separate dedicated queue for small jobs
- pro: allows small jobs to start as soon as it is submitted
- con: large job finishes later than when using the FIFO Scheduler


iii. Fair Scheduler



- Dynamically balances resources between all running jobs
- pro: no capacity reservation required
- con: lag occurs → phase of waiting for resources to be freed up

Assignment 5-2

- Fair scheduling *between queues*



<https://www.safaribooksonline.com/library/view/hadoop-the-definitive/9781491901687/ch04.html>