



# Granula: Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms

Wing Lung Ngai, Tim Hegeman, Stijn Heldens,  
and Alexandru Iosup

# Large-scale Graph Processing



Towards trillion vertices and edges



OpenG

PGX

GraphMat



Powergraph

Graph processing platforms

# Performance Evaluation

Performance evaluation is the process to

- quantify performance,
- explain differences,
- identify overheads
- recommend improvements.

## Performance Studies




User Developer

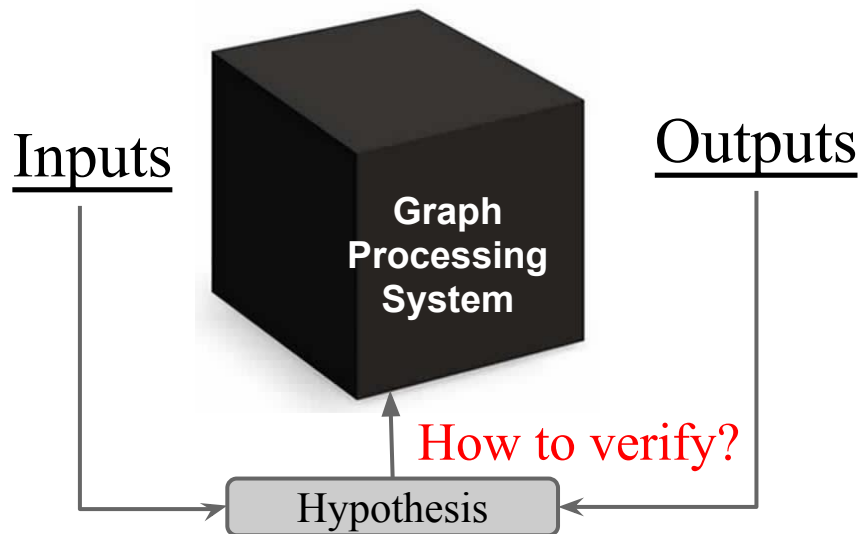
Researcher

# Identified issues

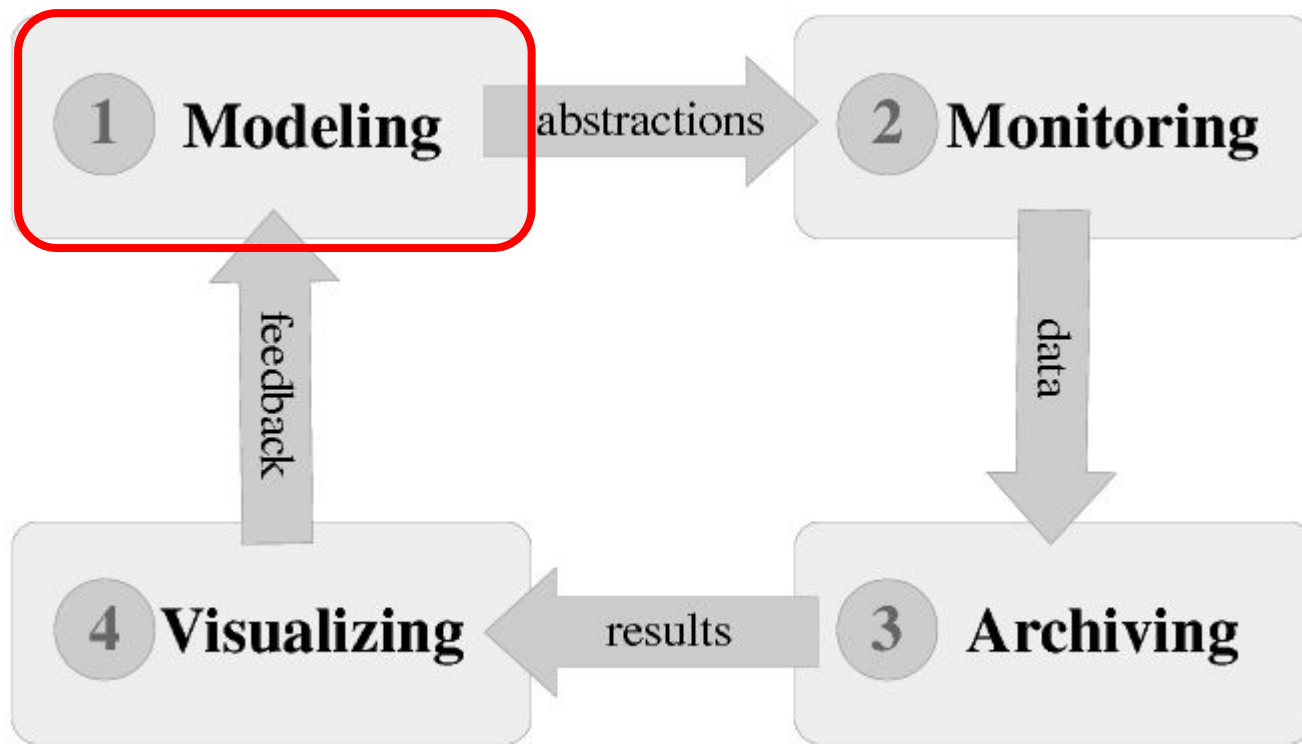
- coarse-grained vs fine-grained
- limited reusability of studies
- lack of an end-to-end process



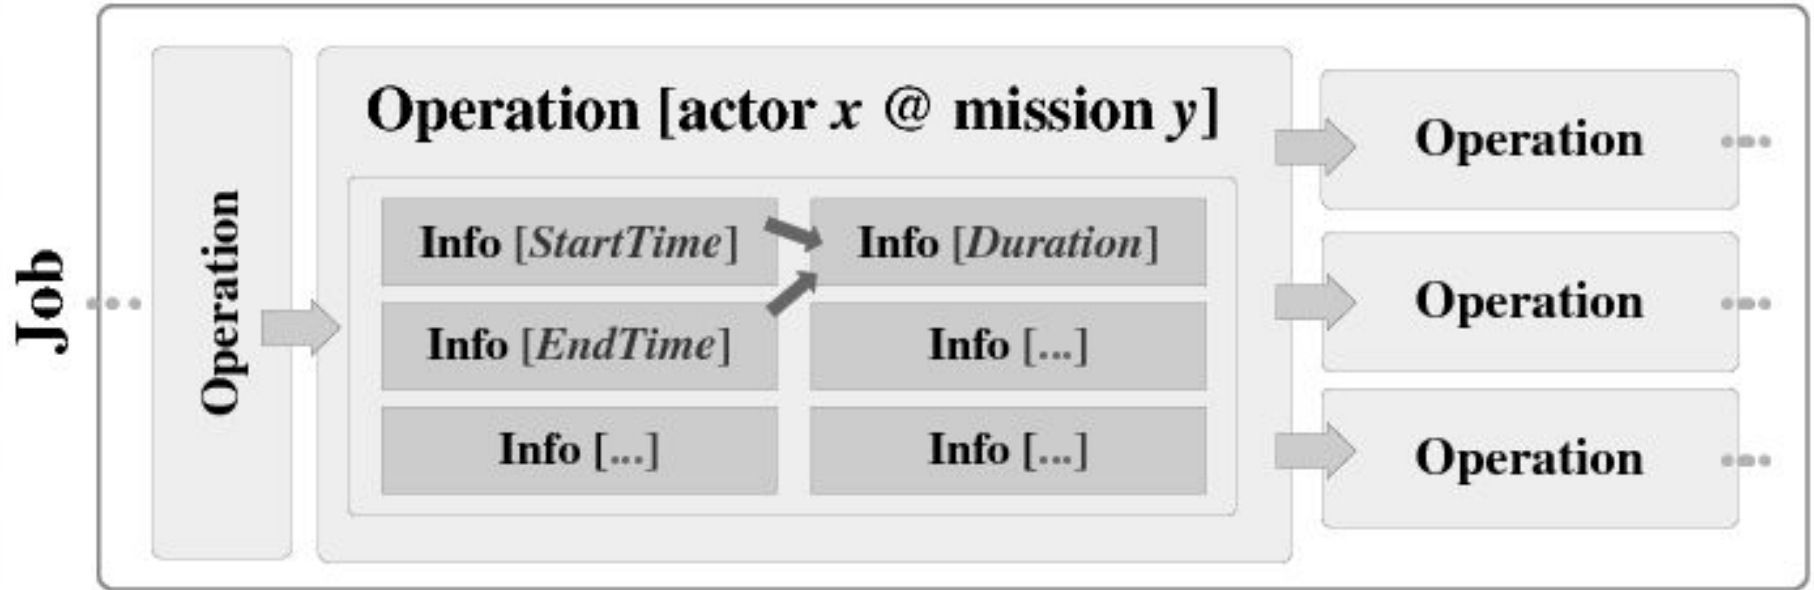
Inefficiency of evaluation



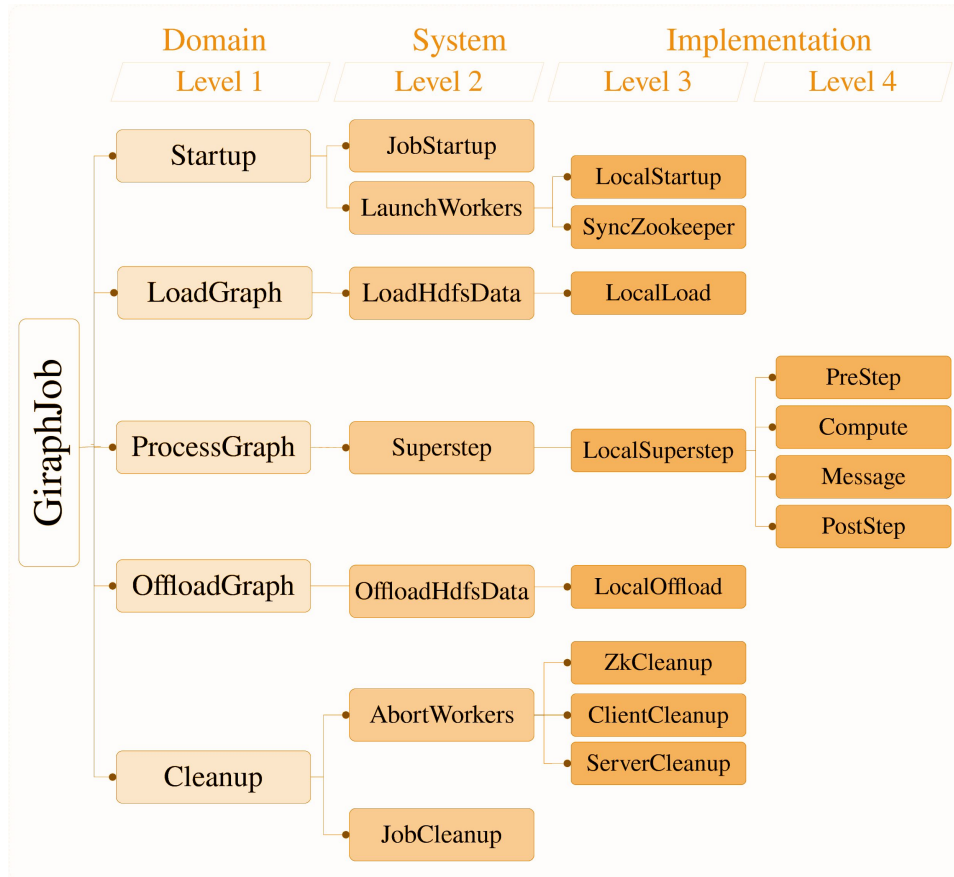

# End-to-end evaluation process



# Performance Modeling Language

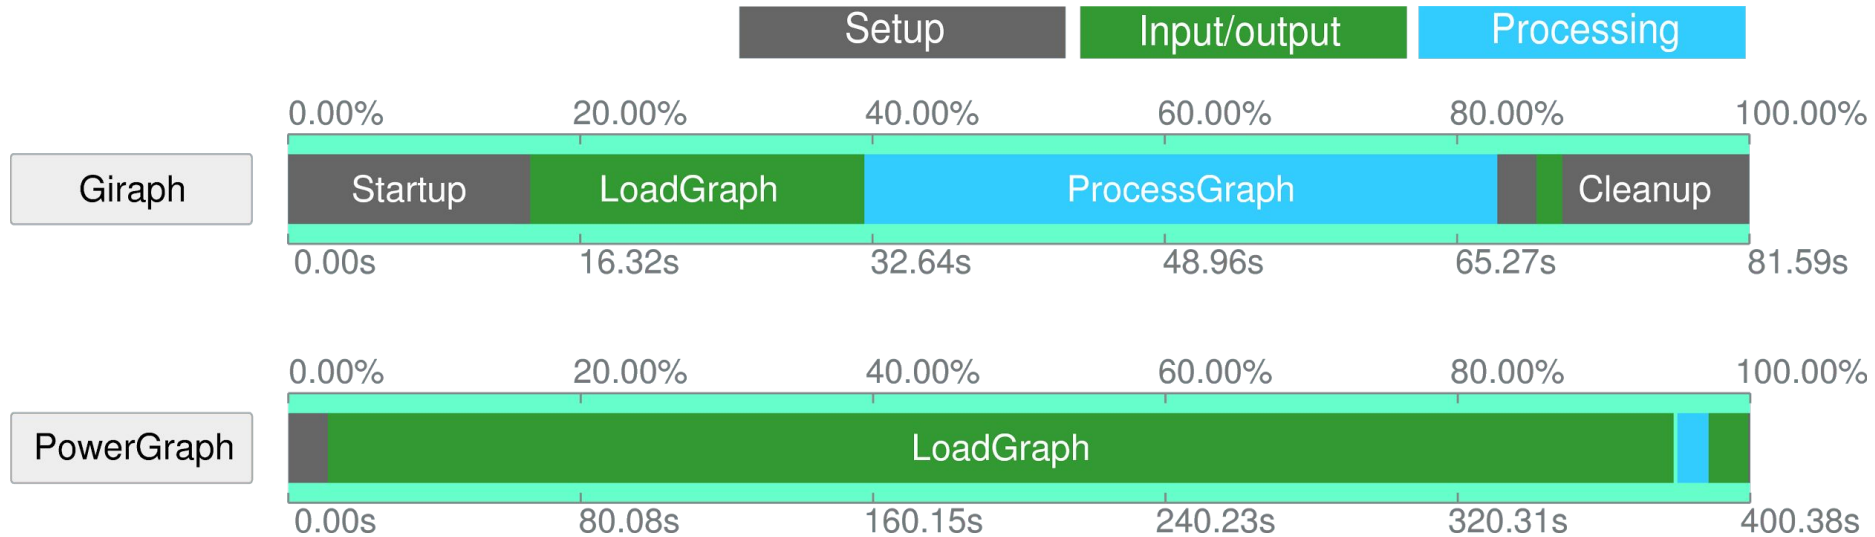


# Building Performance Model



# Quantifying System Performance


Running BFS on dg1000 (giraph and powergraph)





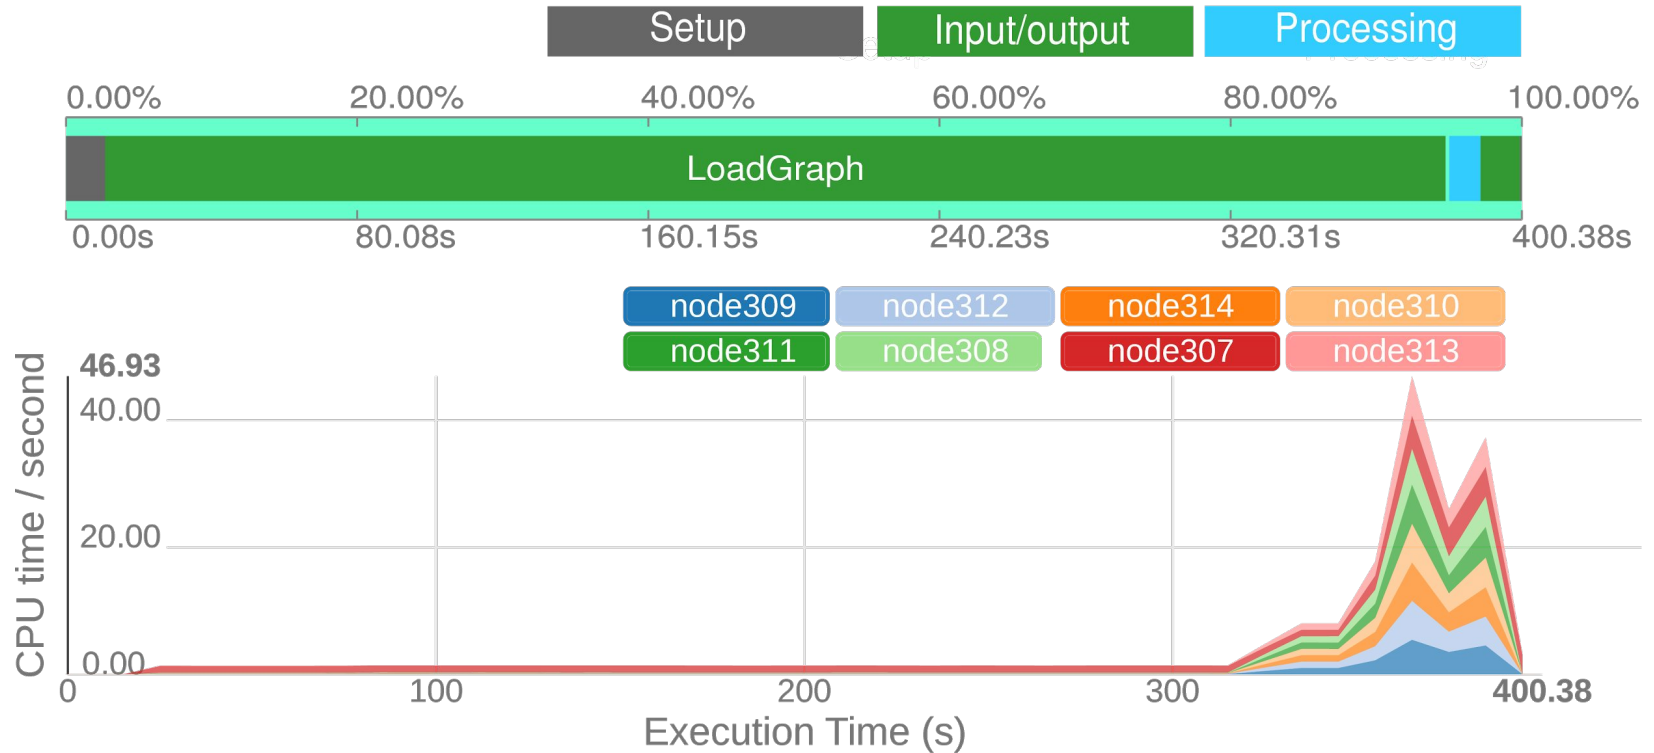
# Monitoring Resource Usage

Running BFS on dg1000 (giraph)



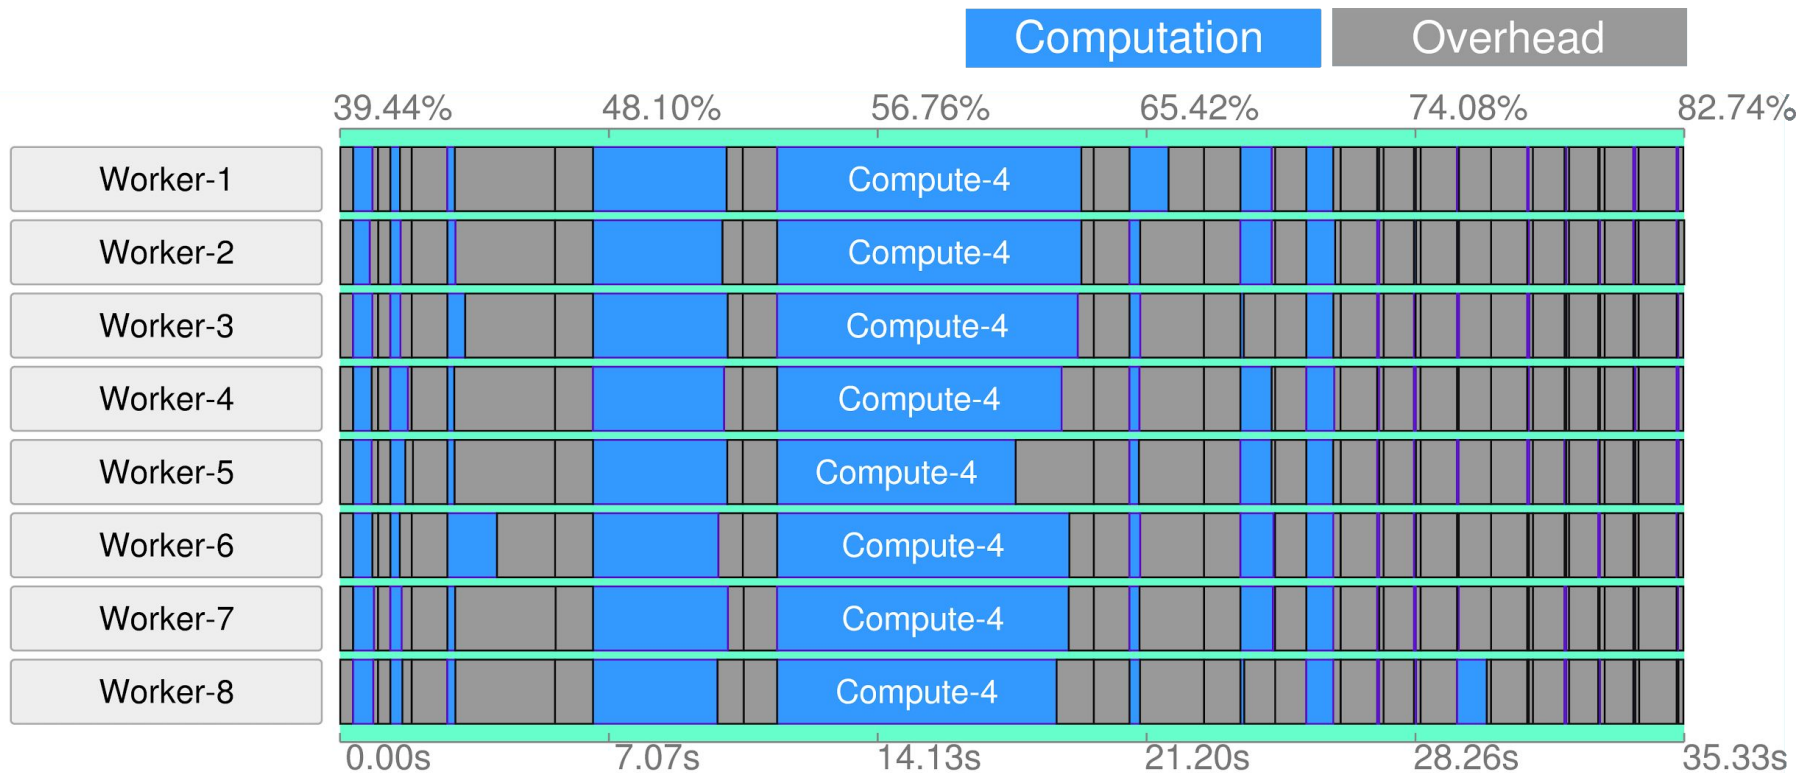
# Monitoring Resource Usage

Running BFS on dg1000 (powergraph)



# Visualizing System Behavior

Running BFS on dg1000 (giraph)



# Visualizing System Behavior

⇒ SyncOverhead@[BspMaster @ Bsplteration] ⇒ SyncOverhead@[BspWorker-M @ MasterTask]

## Target Information

[SyncOverhead] is aggregated from [SyncOverhead]s of all children operation with mission type GlobalSuperstep.

### Information [SyncOverhead]

Name	Type	Value	Information Owner
SyncOverhead	BasicInfo	1758	Operation [BspWorker-M @ MasterTask]

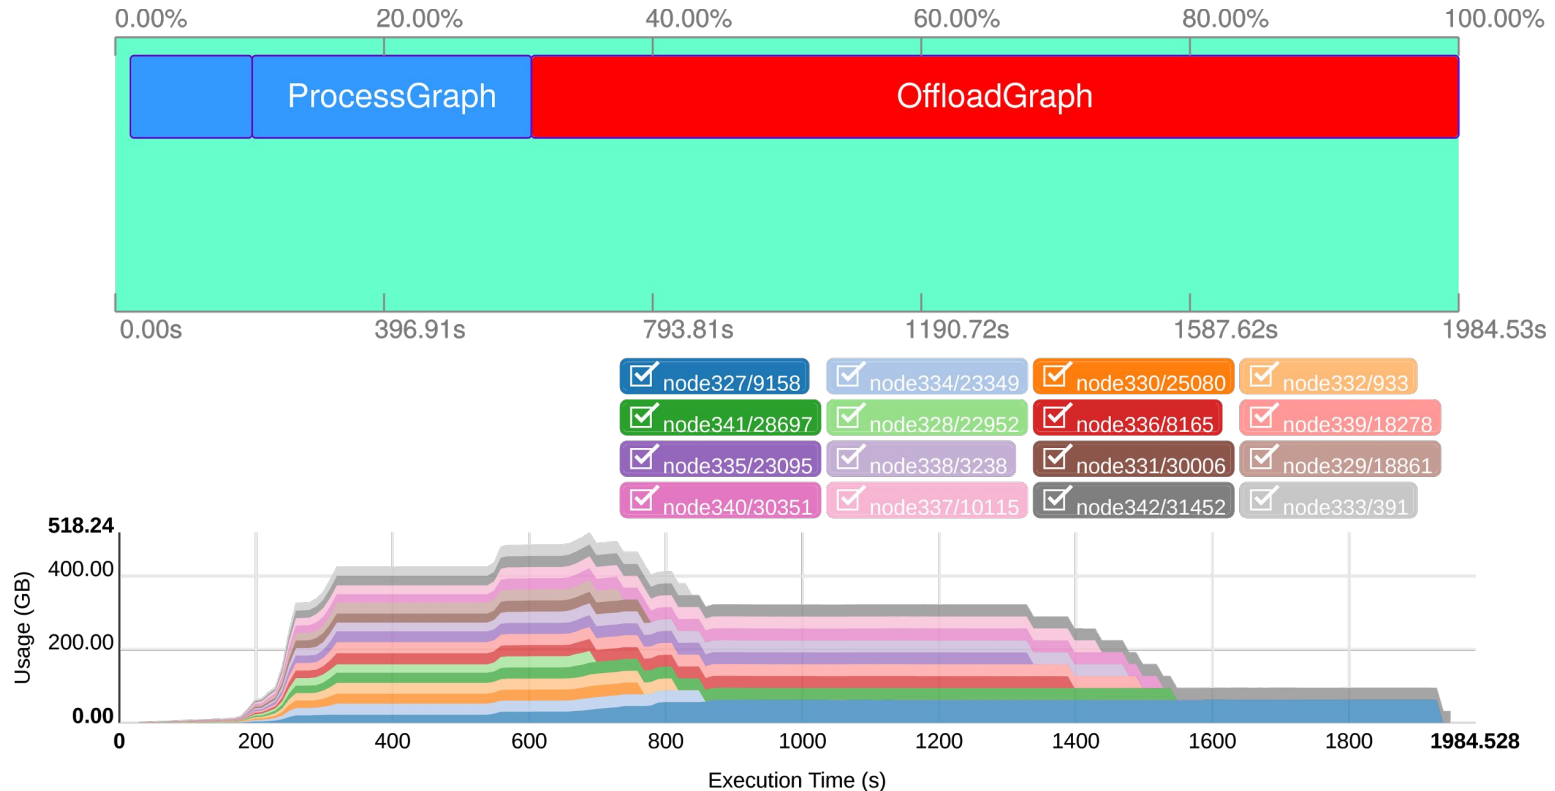
## Source List

### Source [SyncOverhead]

Name	Value	Information Owner
SyncOverhead @	132	Operation [GlobalCoordinator @ GlobalSuperstep-0]
SyncOverhead @	45	Operation [GlobalCoordinator @ GlobalSuperstep-1]
SyncOverhead @	25	Operation [GlobalCoordinator @ GlobalSuperstep-2]

# Diagnosing Job Failure

Running LCC on dg1000 (powergraph)



# Fine-grained Performance Analysis

- Comprehensive: end-to-end evaluation process.
- Standardized: reusable across platforms.
- Incremental: iterative performance modeling.
- Automatable: automated repetitive procedures.

Granula: a fine-grained performance analysis system for Big Data platforms

# Conclusion

We propose **Granula**:

- a fine-grained performance analysis system for Big Data platforms
- that facilitates modeling, monitoring, archiving, and visualization

Future Work:

- continue the development of our research prototype,
- apply our techniques on other types of Big Data platforms,
- better support for analysts on, e.g, failure diagnosis, regression tests,
- integrate performance analysis into standard software engineering practices