Freie Universität Berlin

Data Location Optimization for a Self-Organized Storage System

Hannes Mühleisen, Tilman Walther and Robert Tolksdorf

[A. Bockoven]

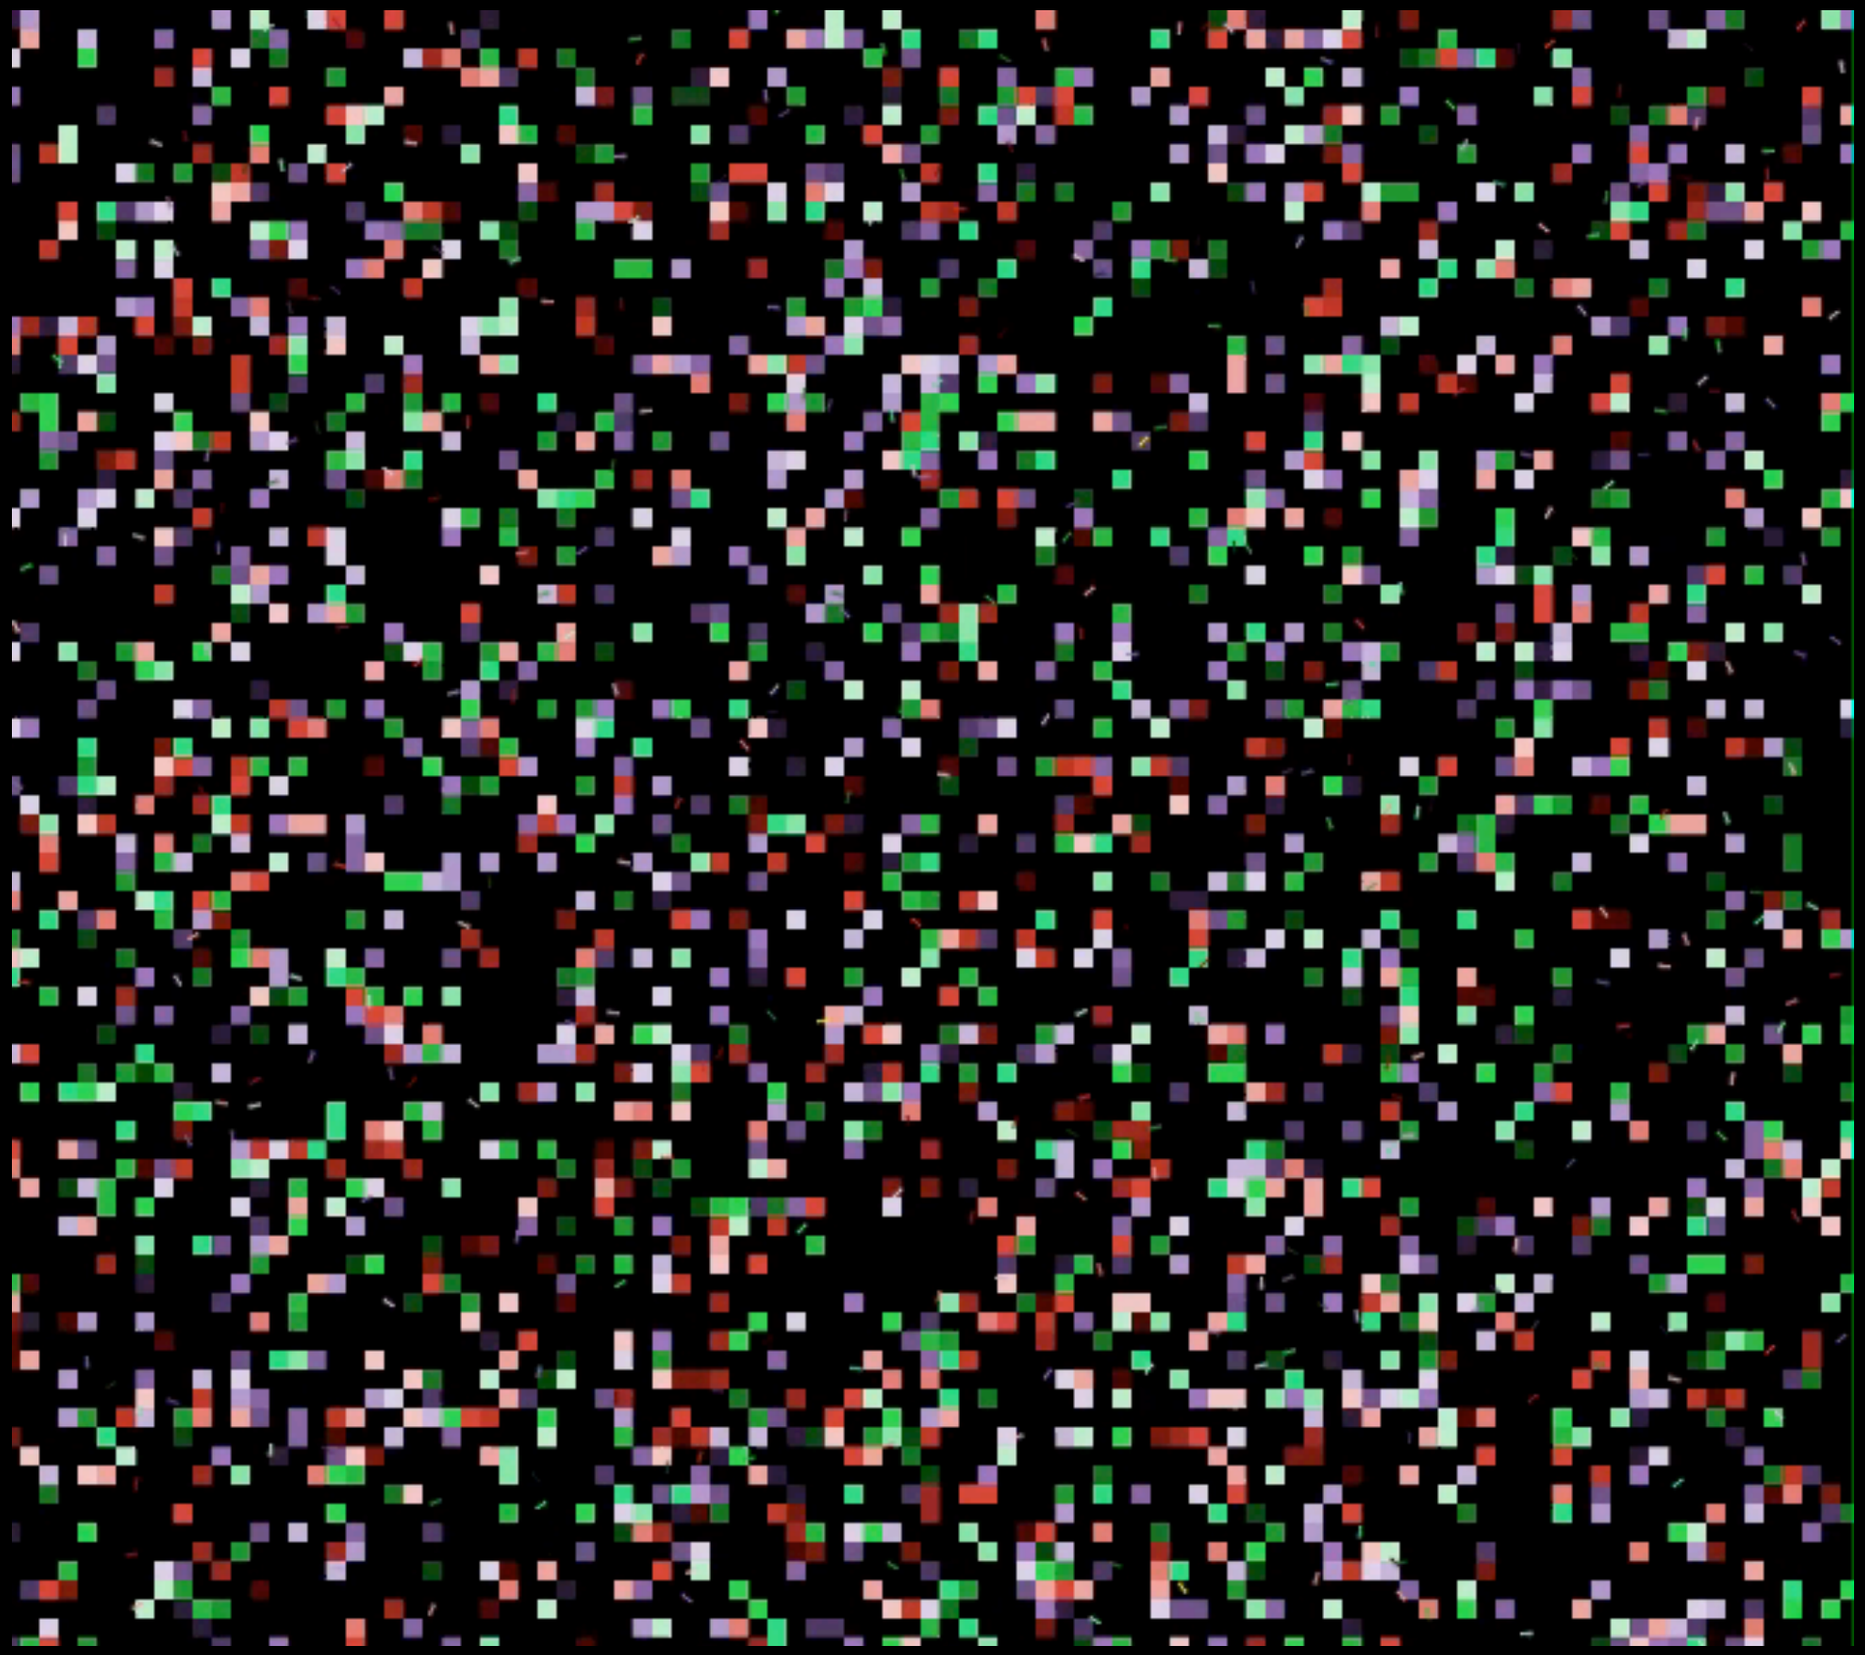

[Thomas Schmickl]

Brood Sorting - Algorithm

```
item = null;
while (true)
    if (item != null)
        if (similarity(item,nearbyItems()) > α)
            drop(item)
            item = null
    else
        item = min(similarity(nearbyItems()²))
        pickup(item)
    move()
```

Probabilistic Request Routing

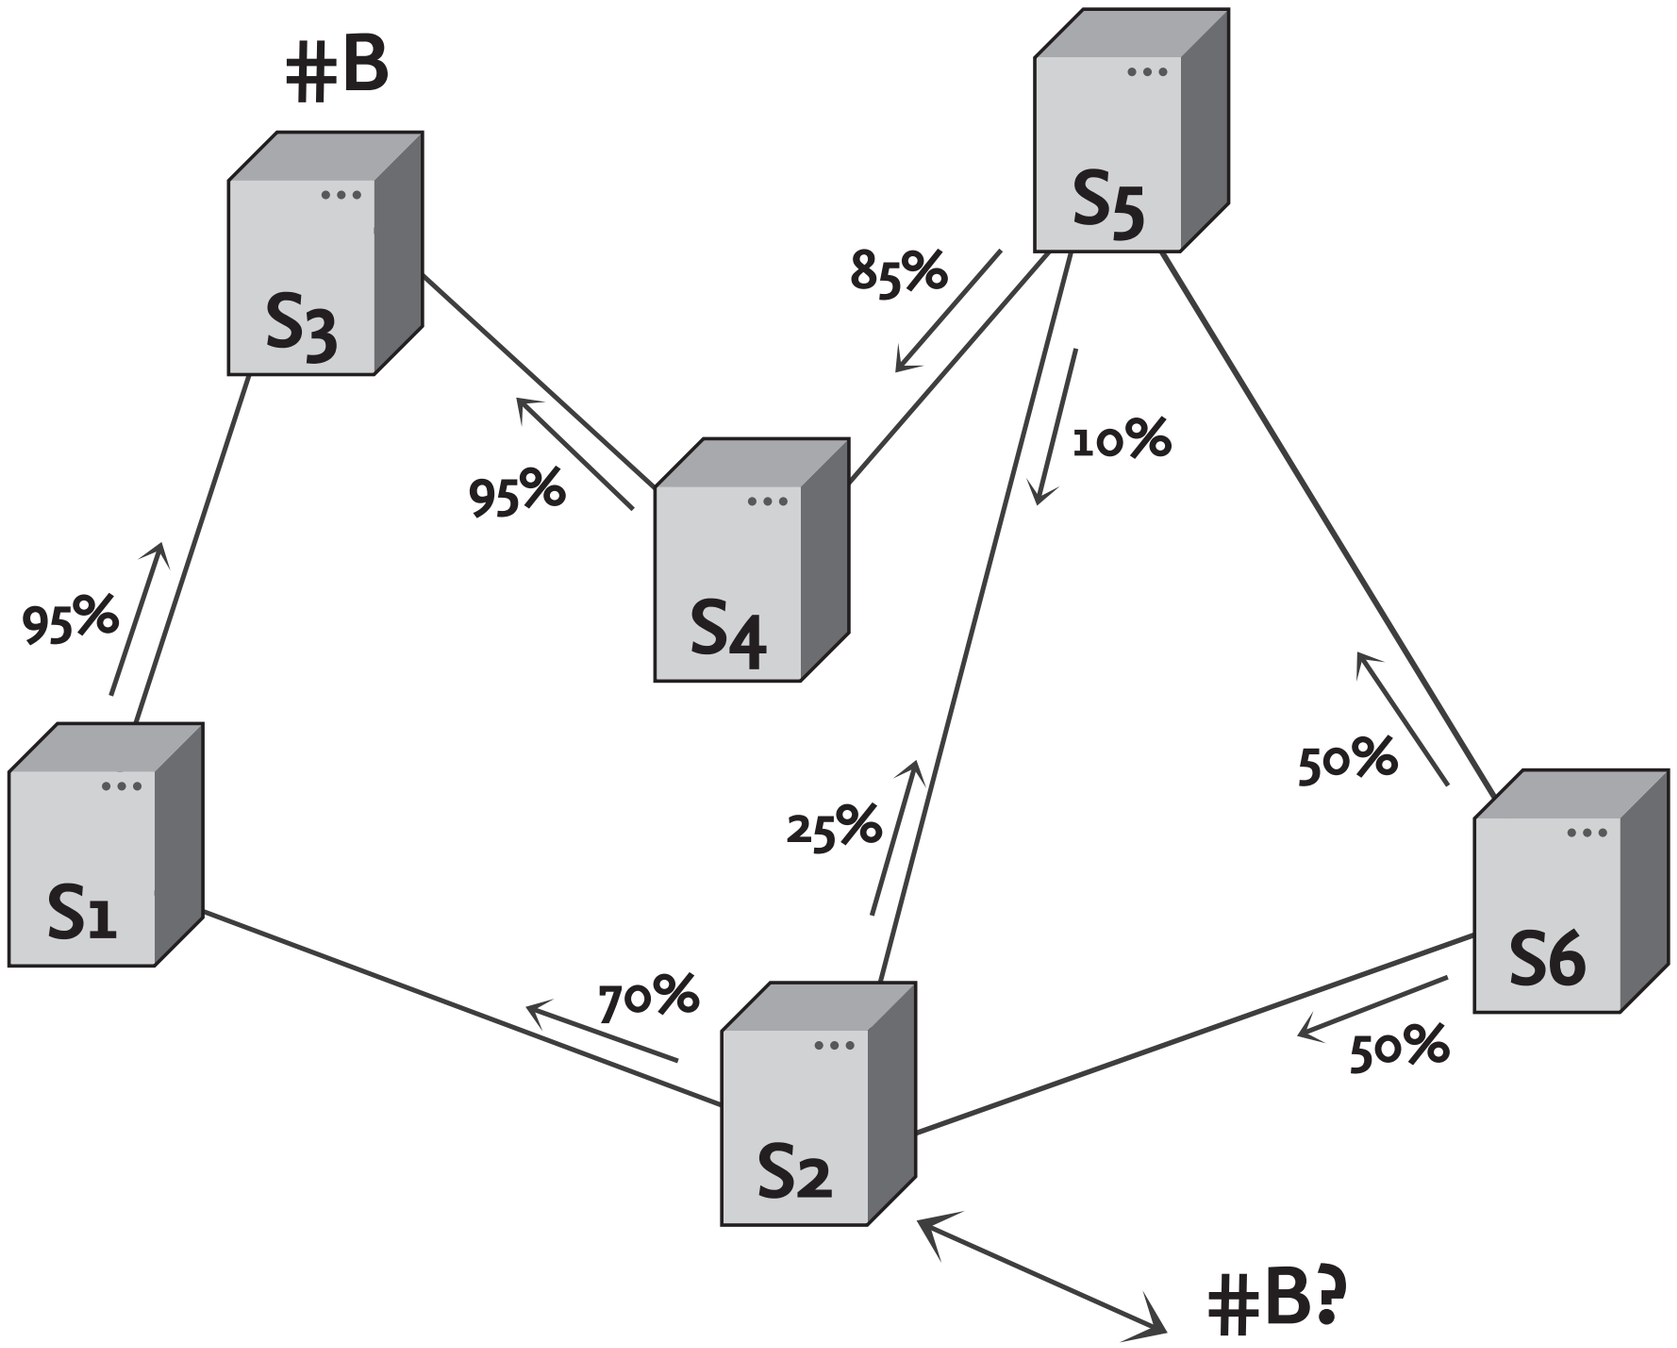

[Lindgreno3]

Research Question

Can brood sorting improve data placement in a large-scale distributed storage system based on probabilistic routing?

Some Adaptions

- Data is clustered into a limited amount of “buckets”
- Movement split up into two phases:
  - Search phase: Every node periodically generates “profile” of locally stored data and sends it on its way
  - Response phase: Nodes compare incoming profiles to local stored data, generating movement responses

(1)

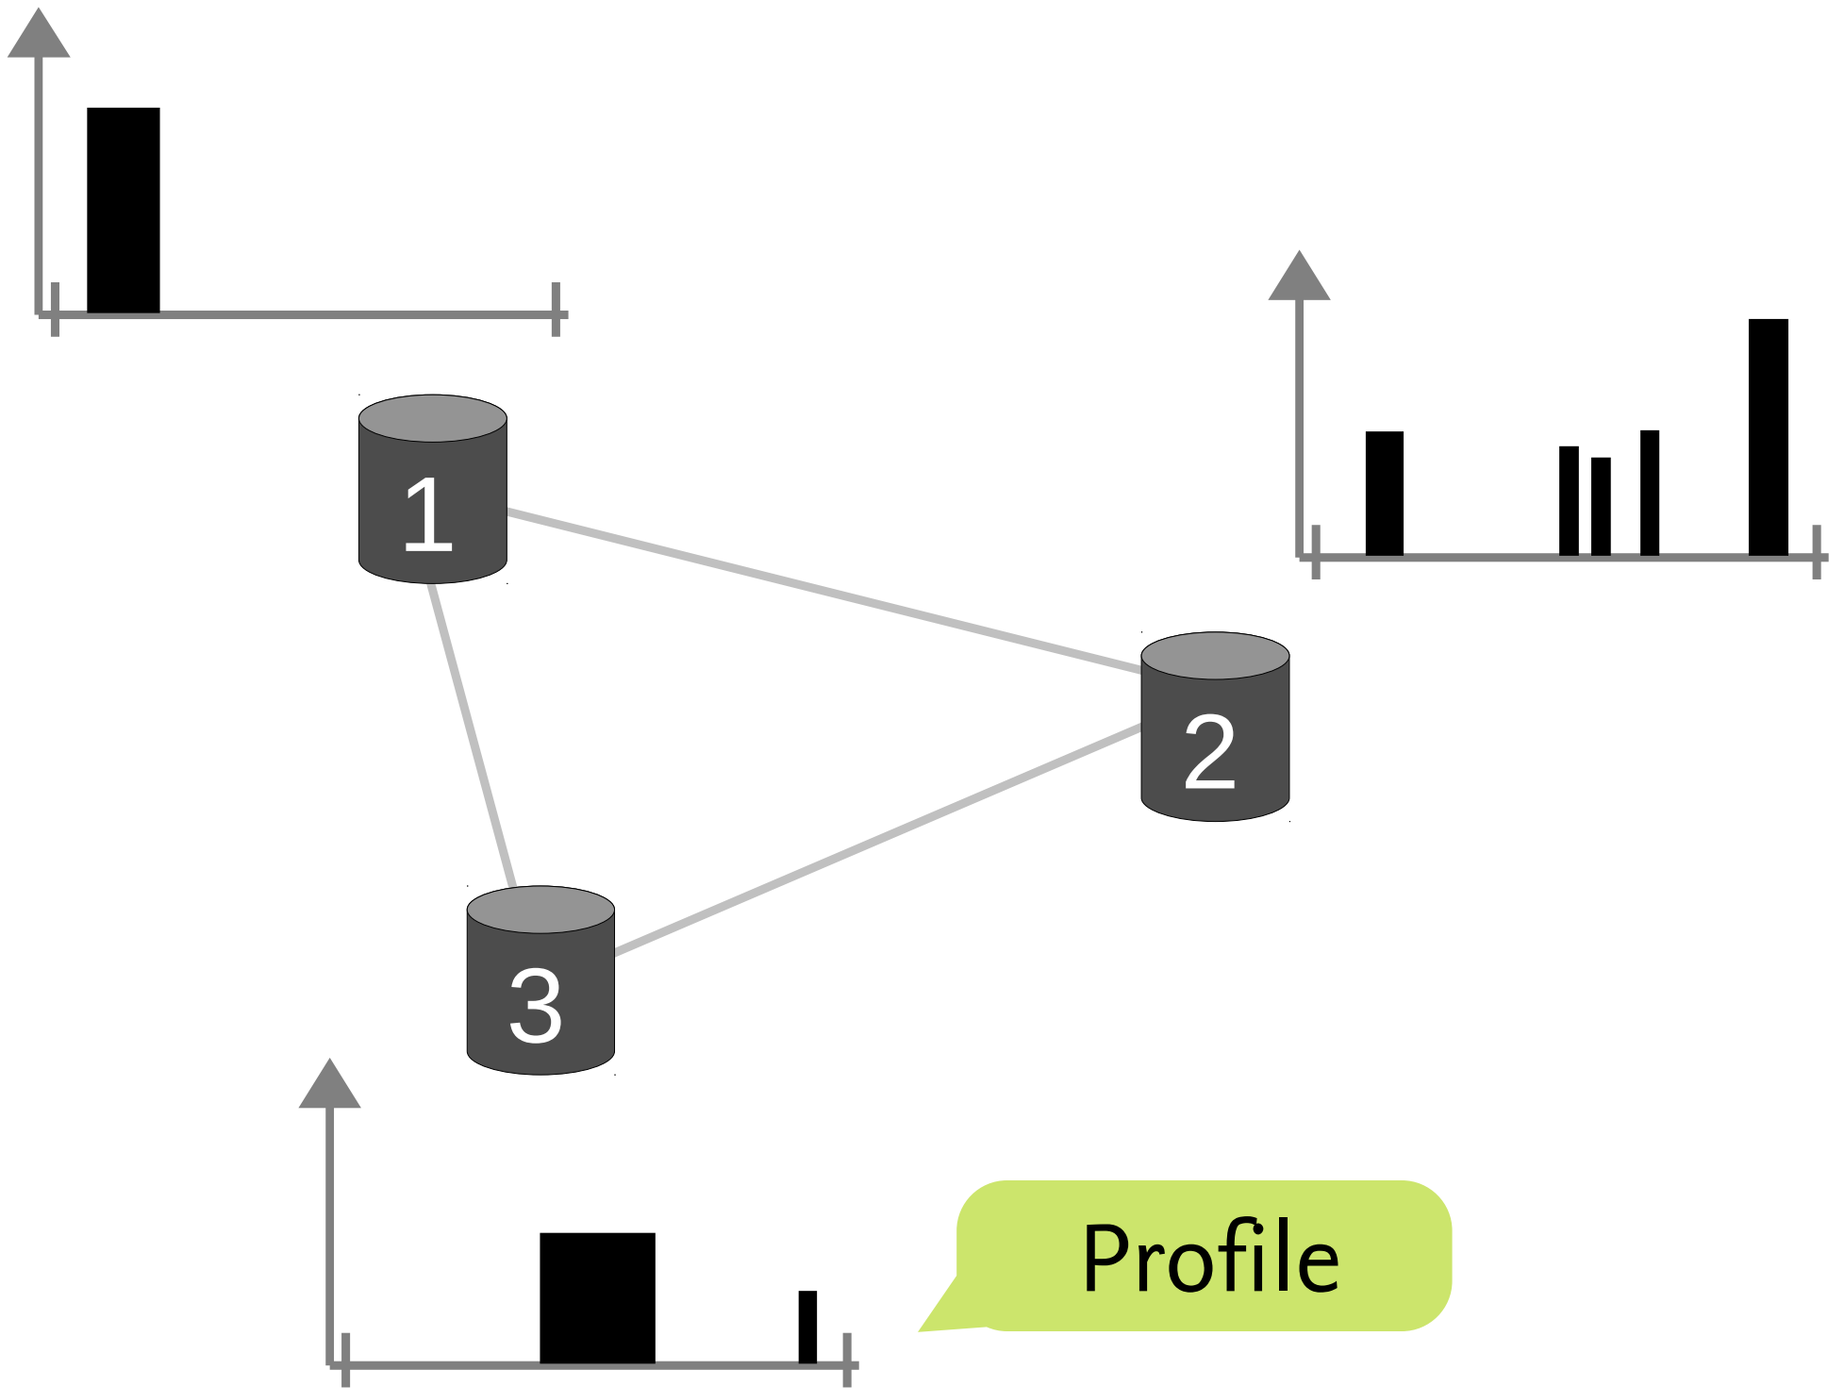

(2)

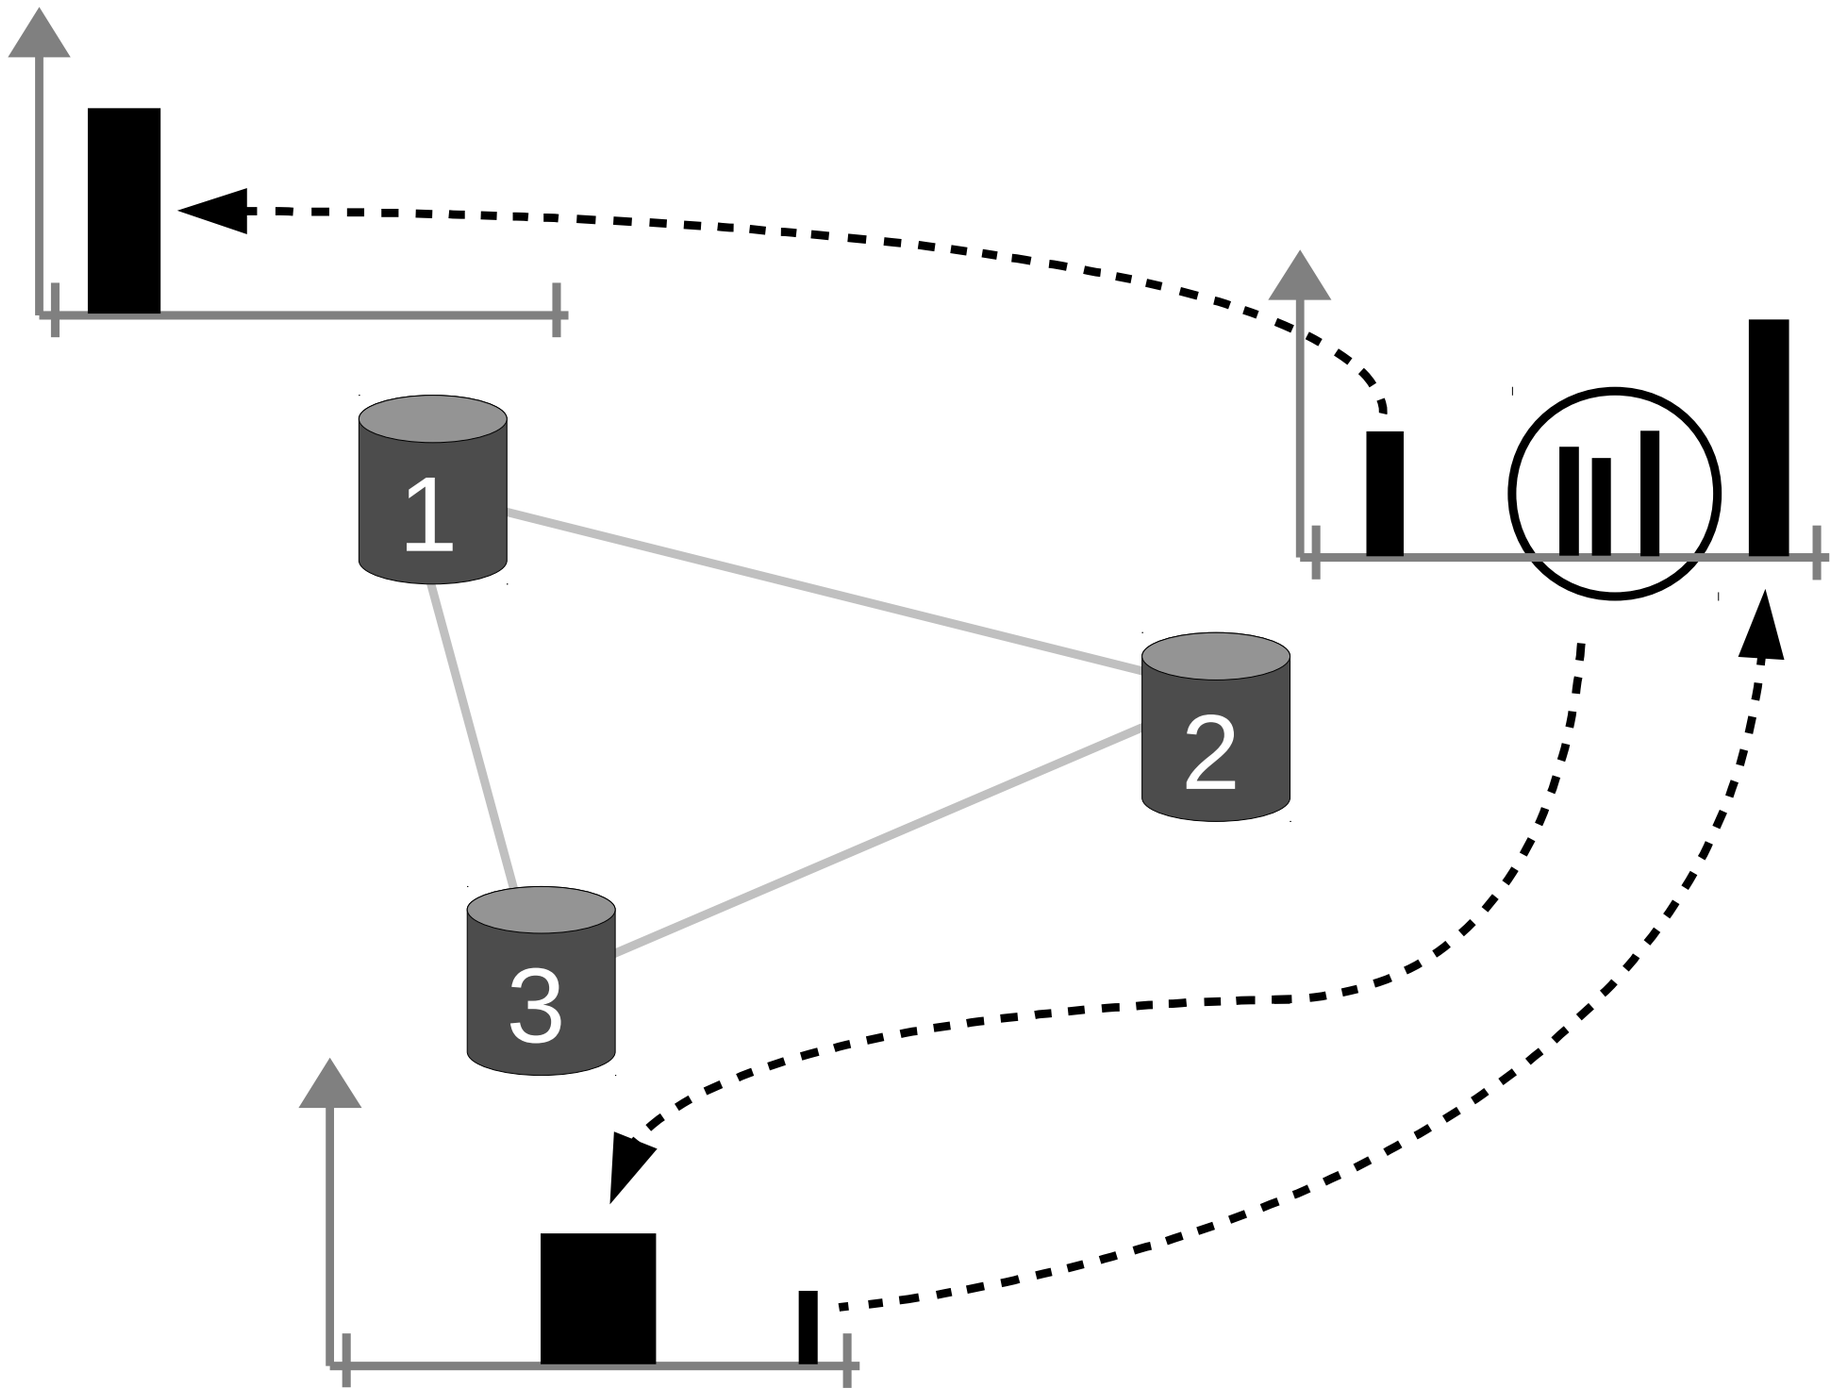

(3)

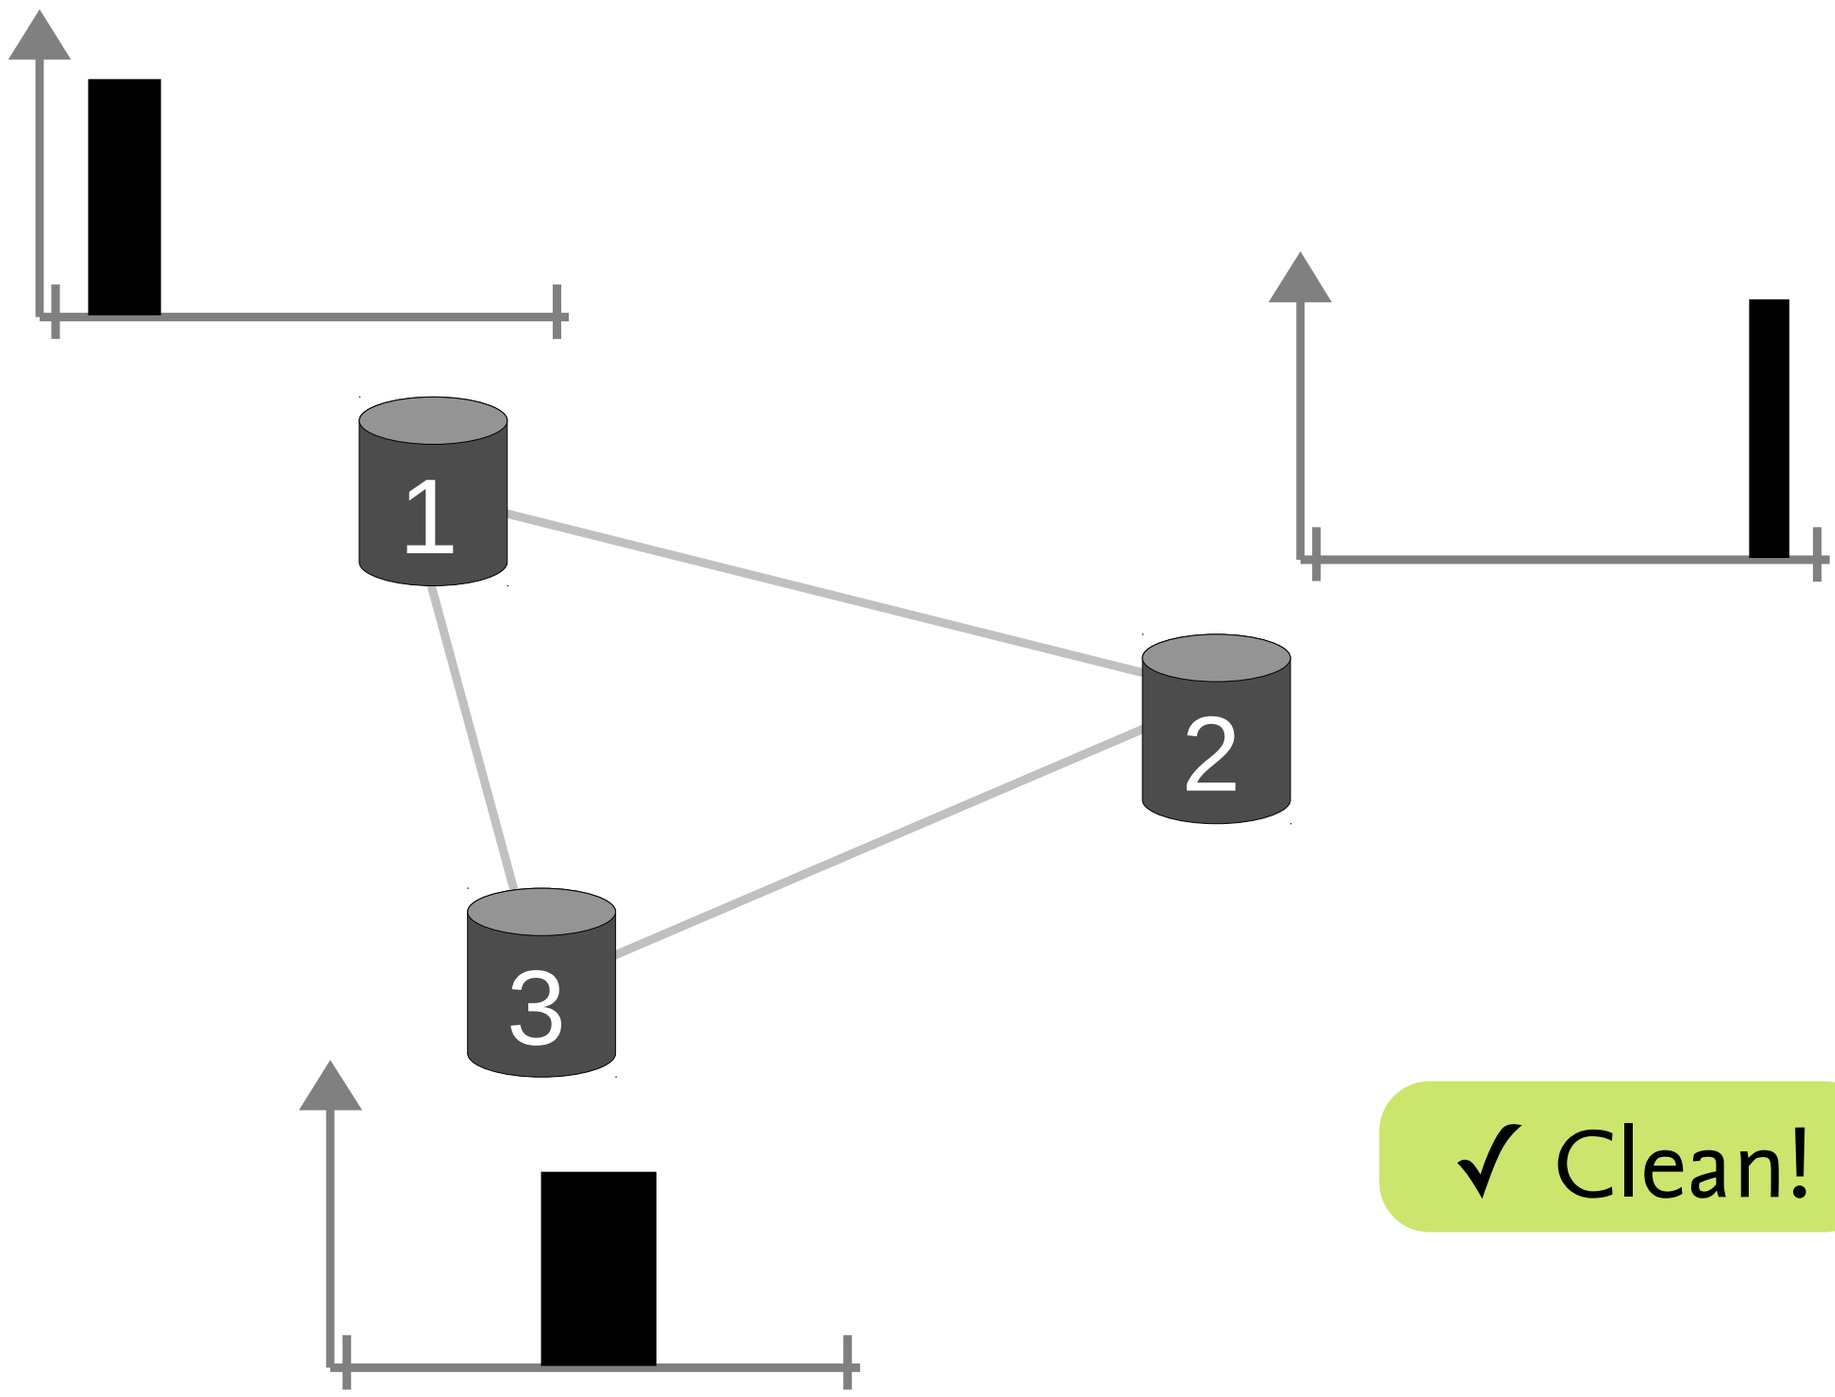

✓ Clean!

Evaluation

- Cluster of 100 Linux nodes
- Two datasets, random & synthetic
- 1000 write operations, four phases
- Recorded data:
  - # Data items in network
  - # Successful movement operations
  - Bucket amount & size

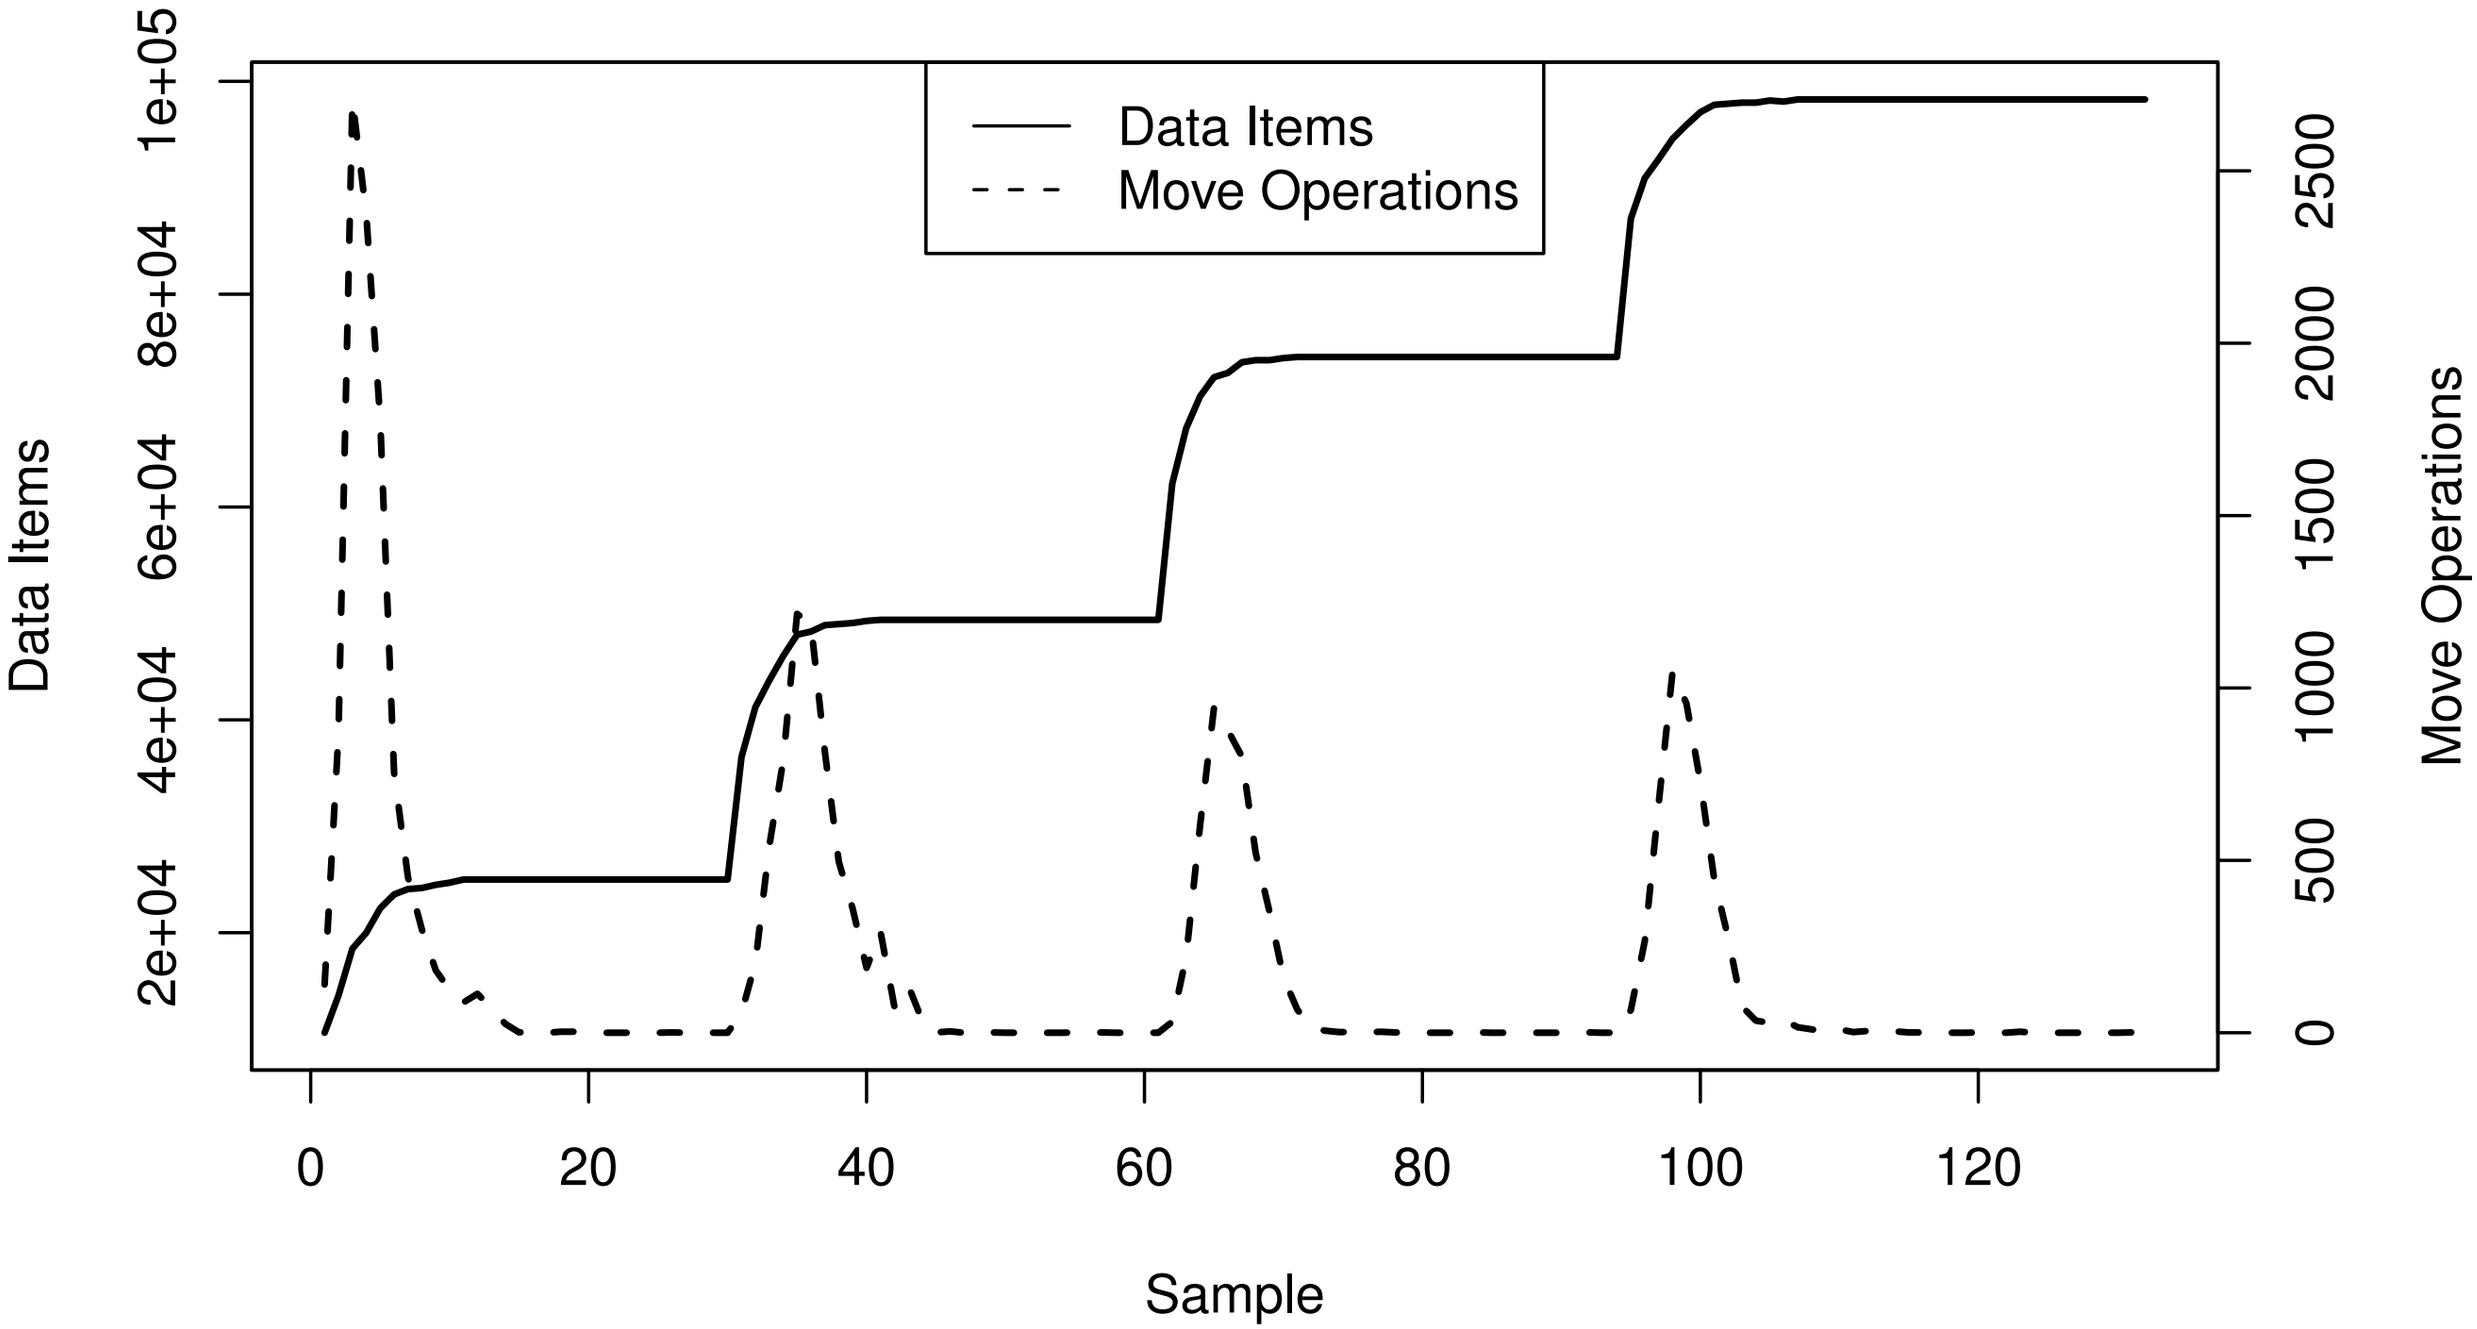
Data Items vs. Move Operations synthetic/100nodes
Data Items
Move Operations
Data Items
Move Operations
Sample
2e+04
4e+04
6e+04
8e+04
1e+05
0
500
1000
1500
2000
2500
0
20
40
60
80
100
120

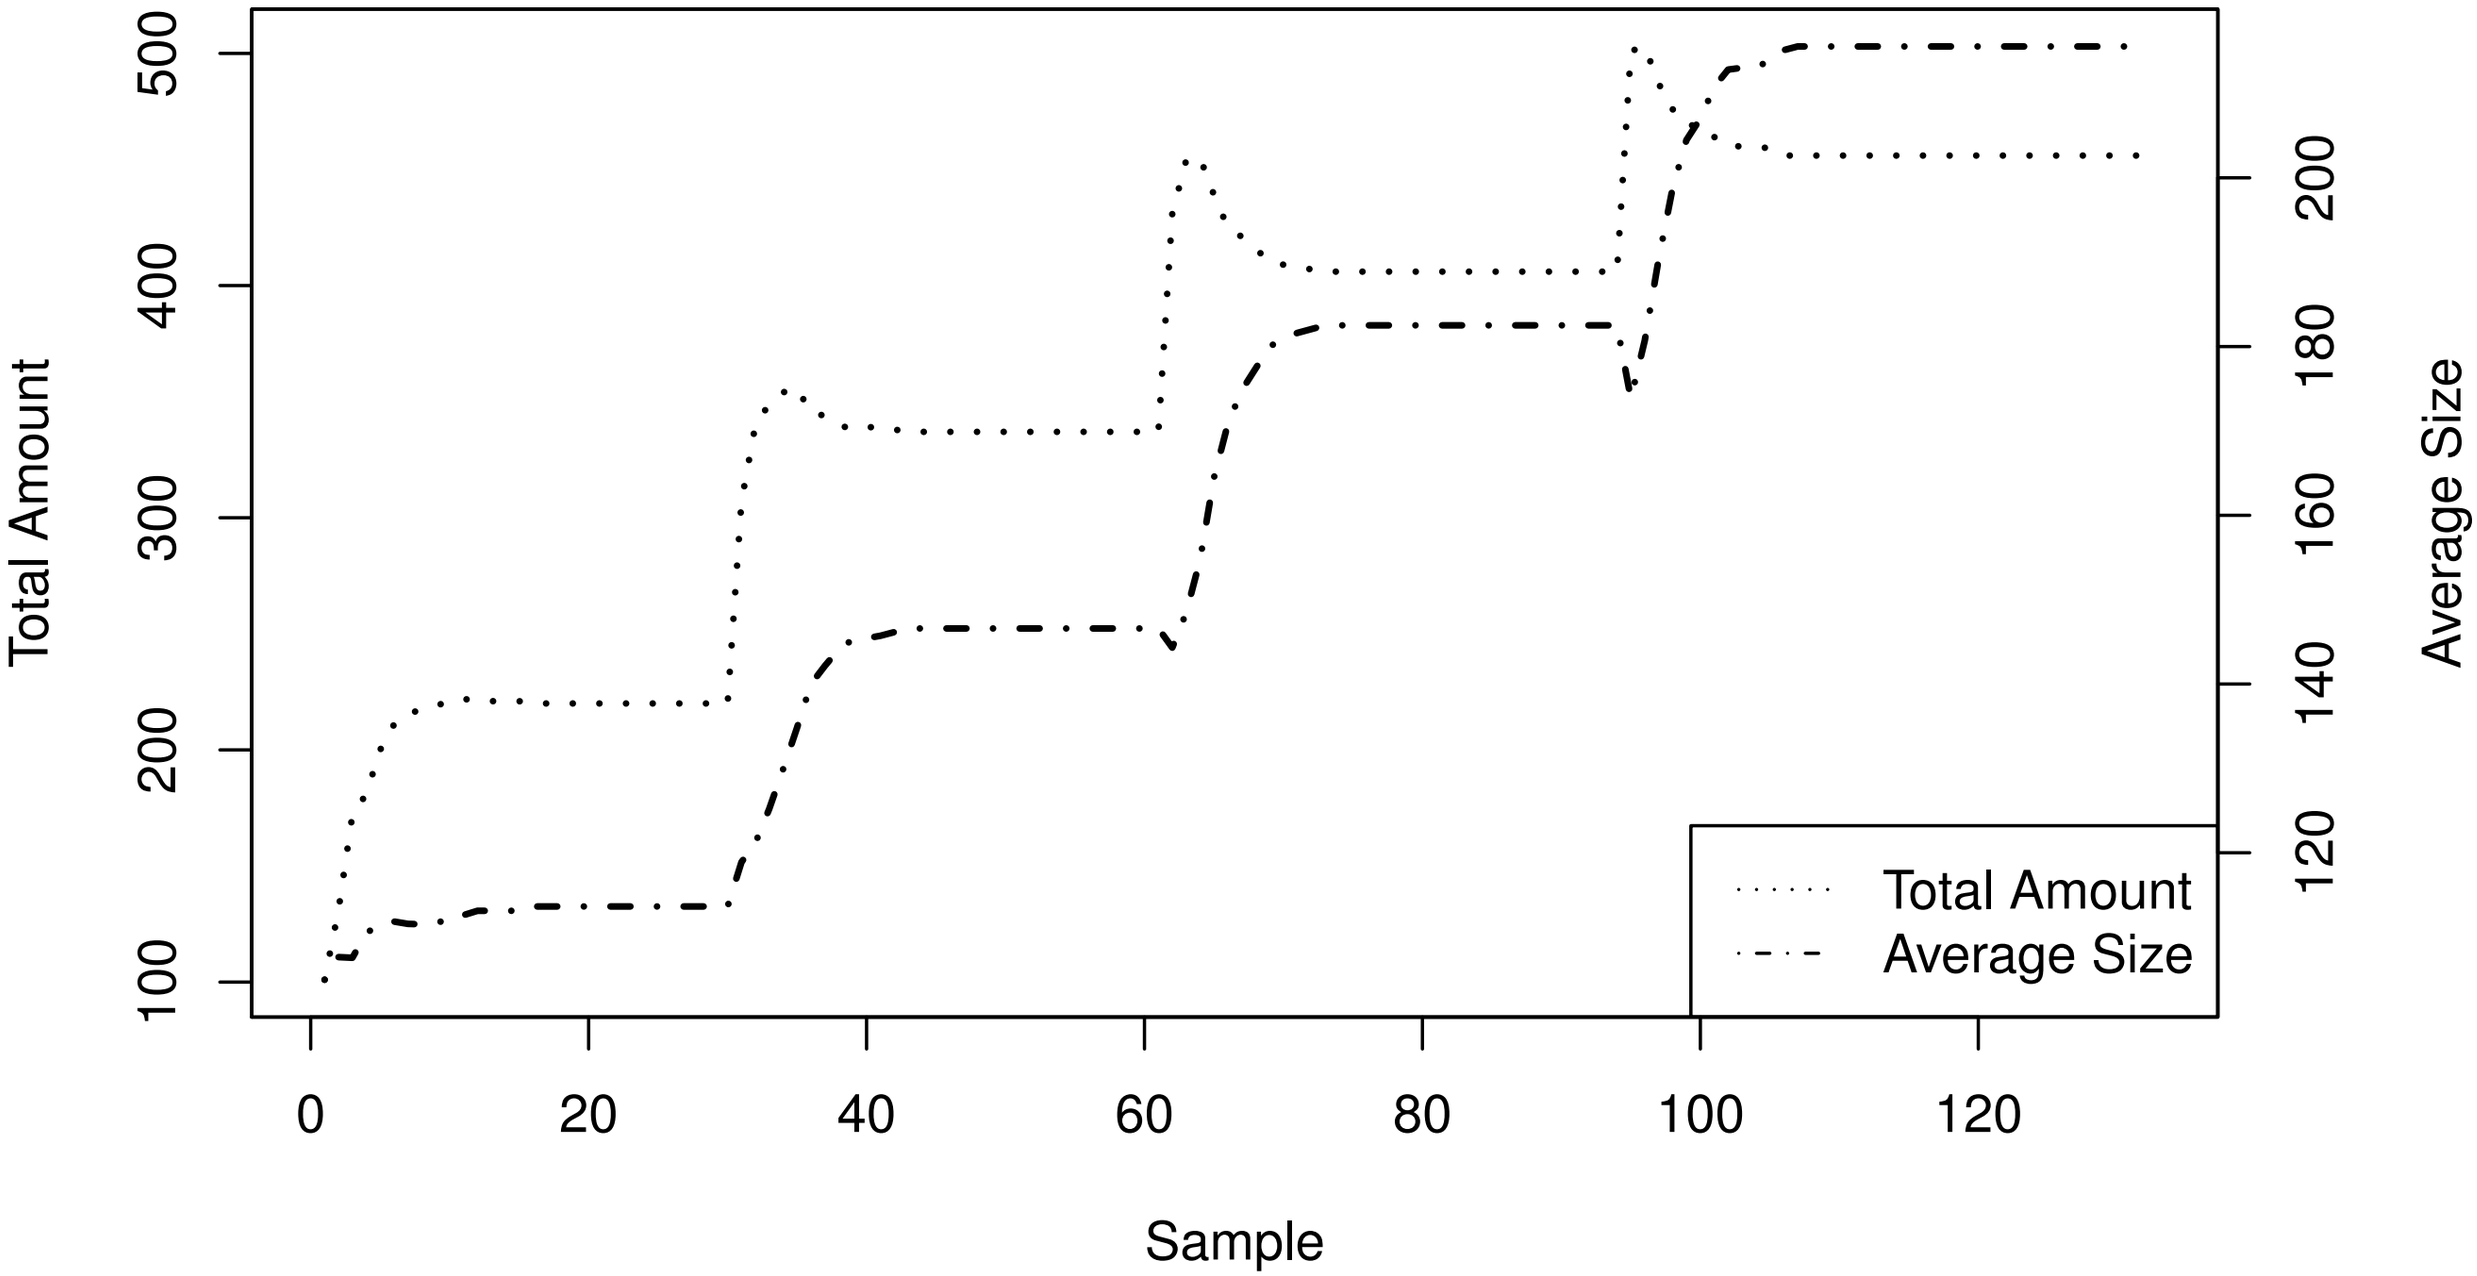
Bucket Amount vs. Average Size synthetic/100nodes
Total Amount
Average Size
100
200
300
400
500
120
140
160
180
200
0
20
40
60
80
100
120
Sample
Total Amount
Average Size

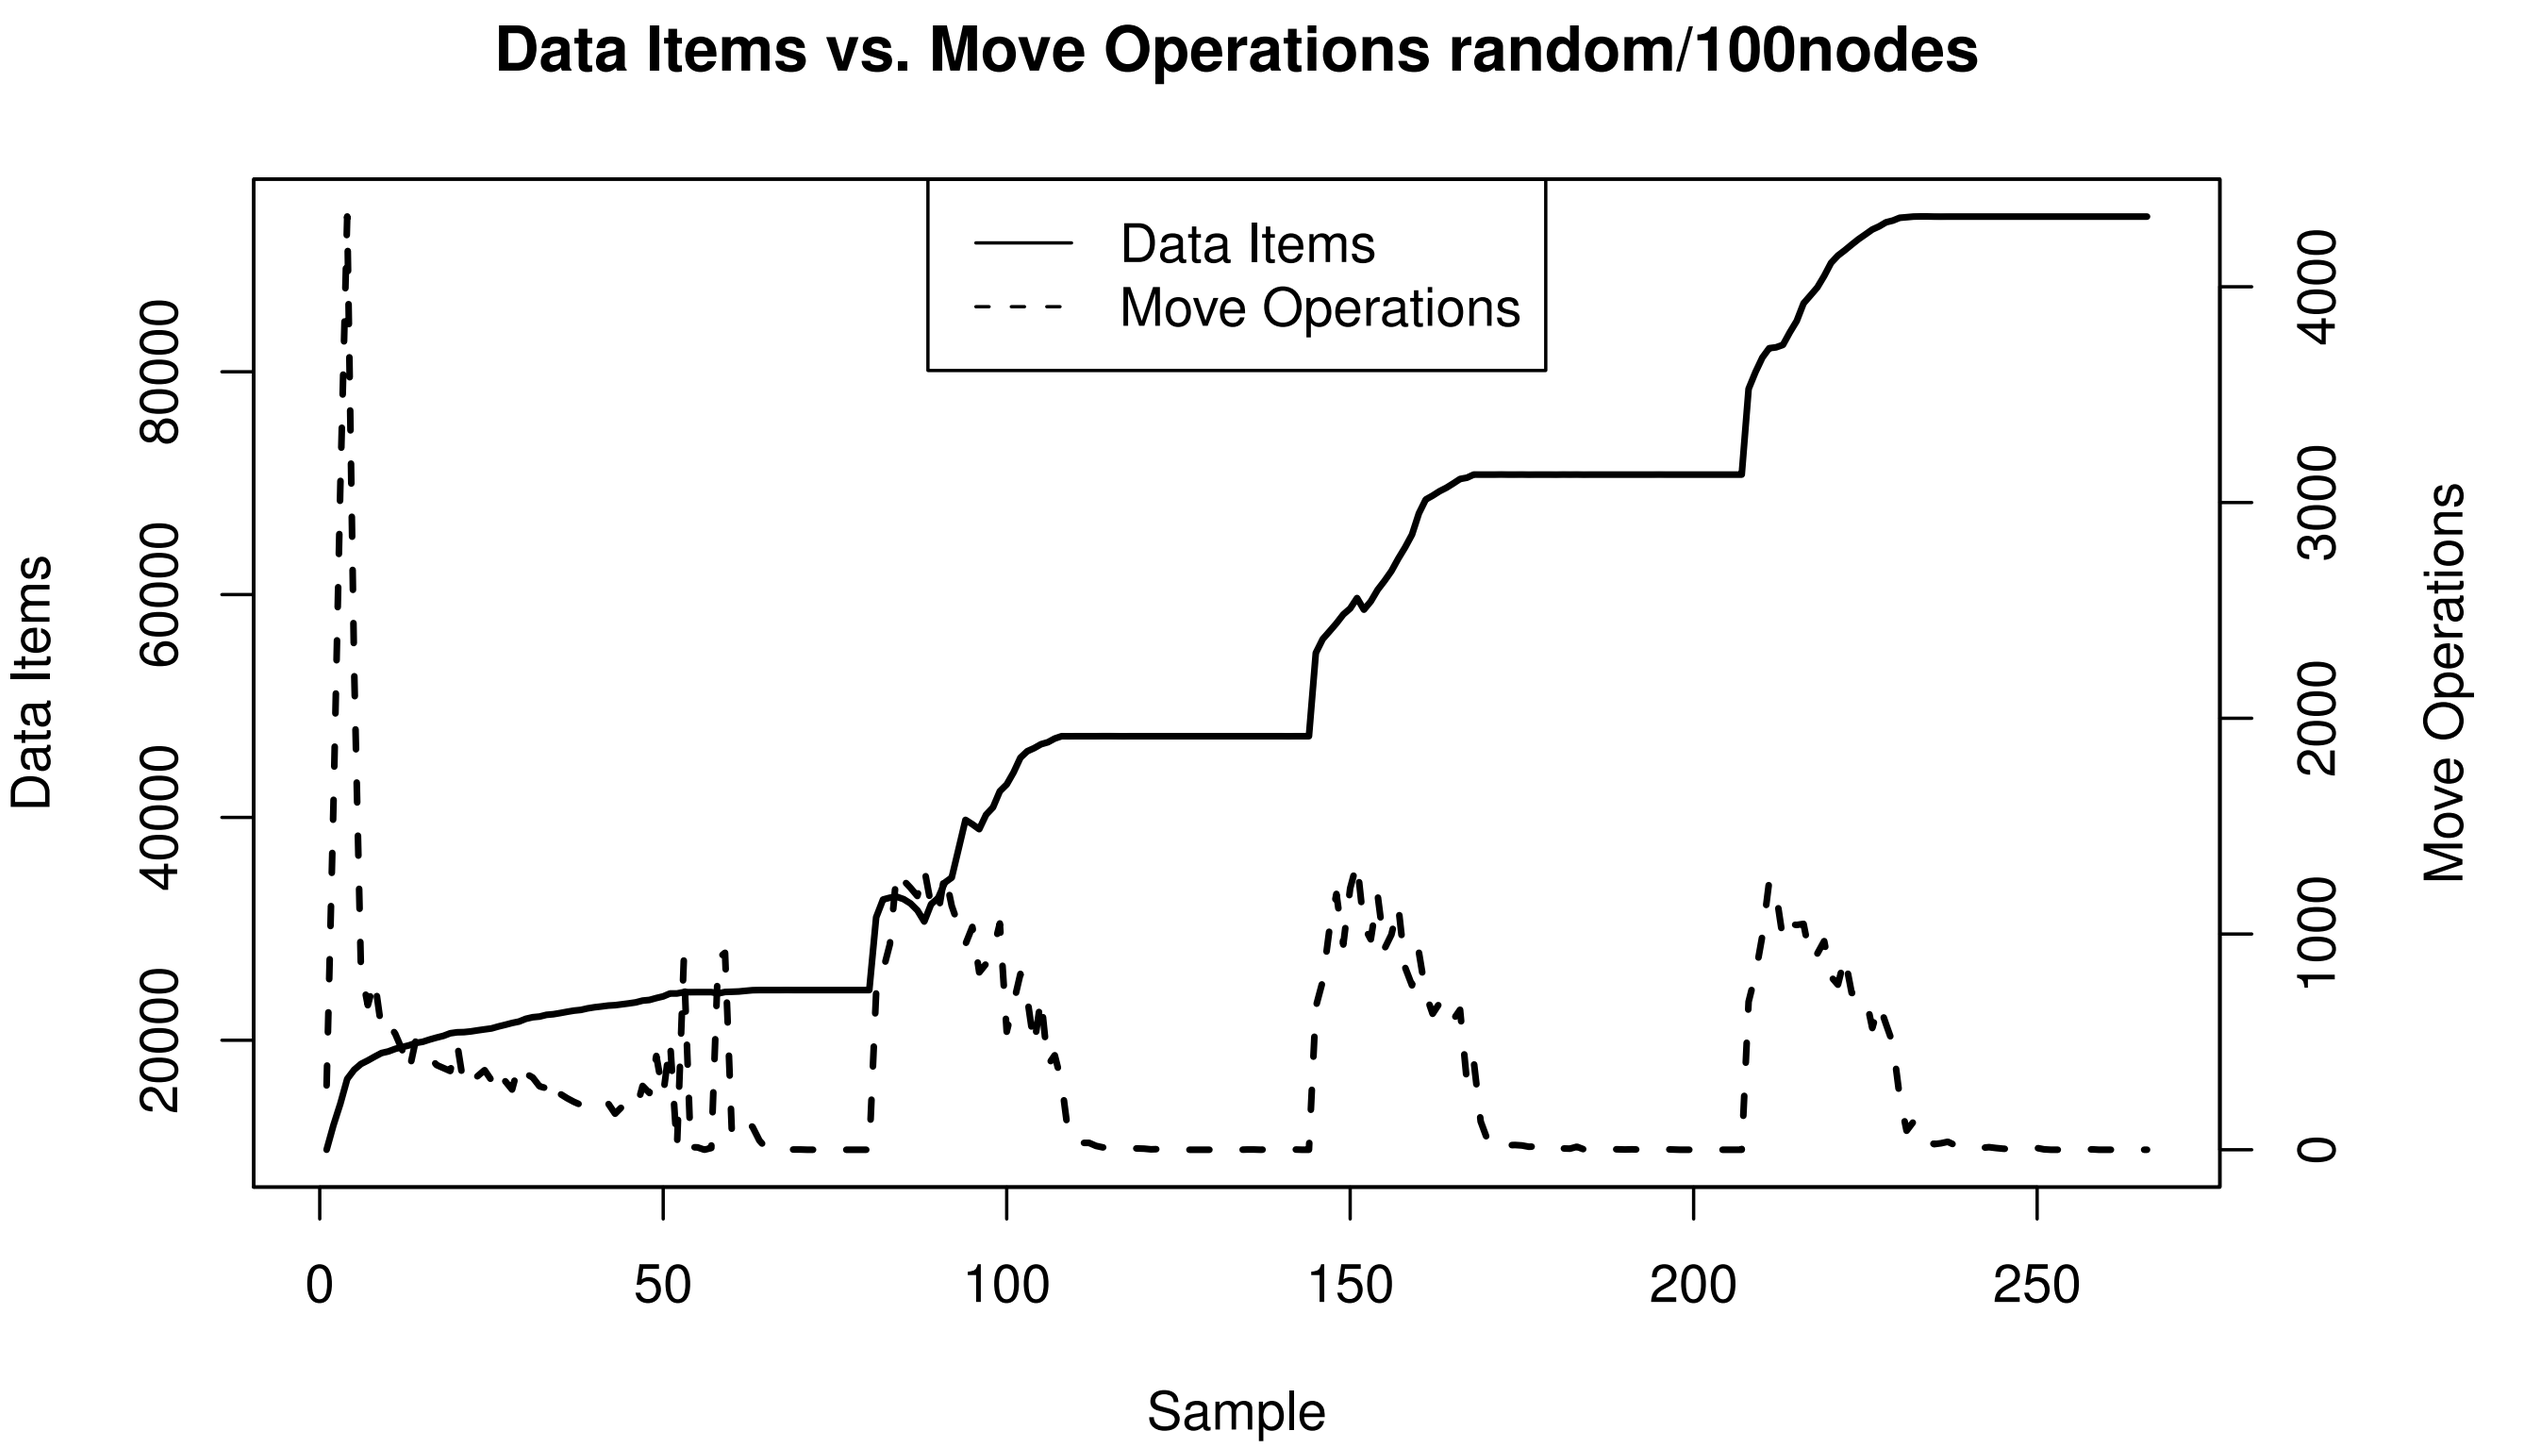
Data Items vs. Move Operations random/100nodes
Data Items
Move Operations
Data Items
Move Operations
20000
40000
60000
80000
0
1000
2000
3000
4000
0
50
100
150
200
250
Sample

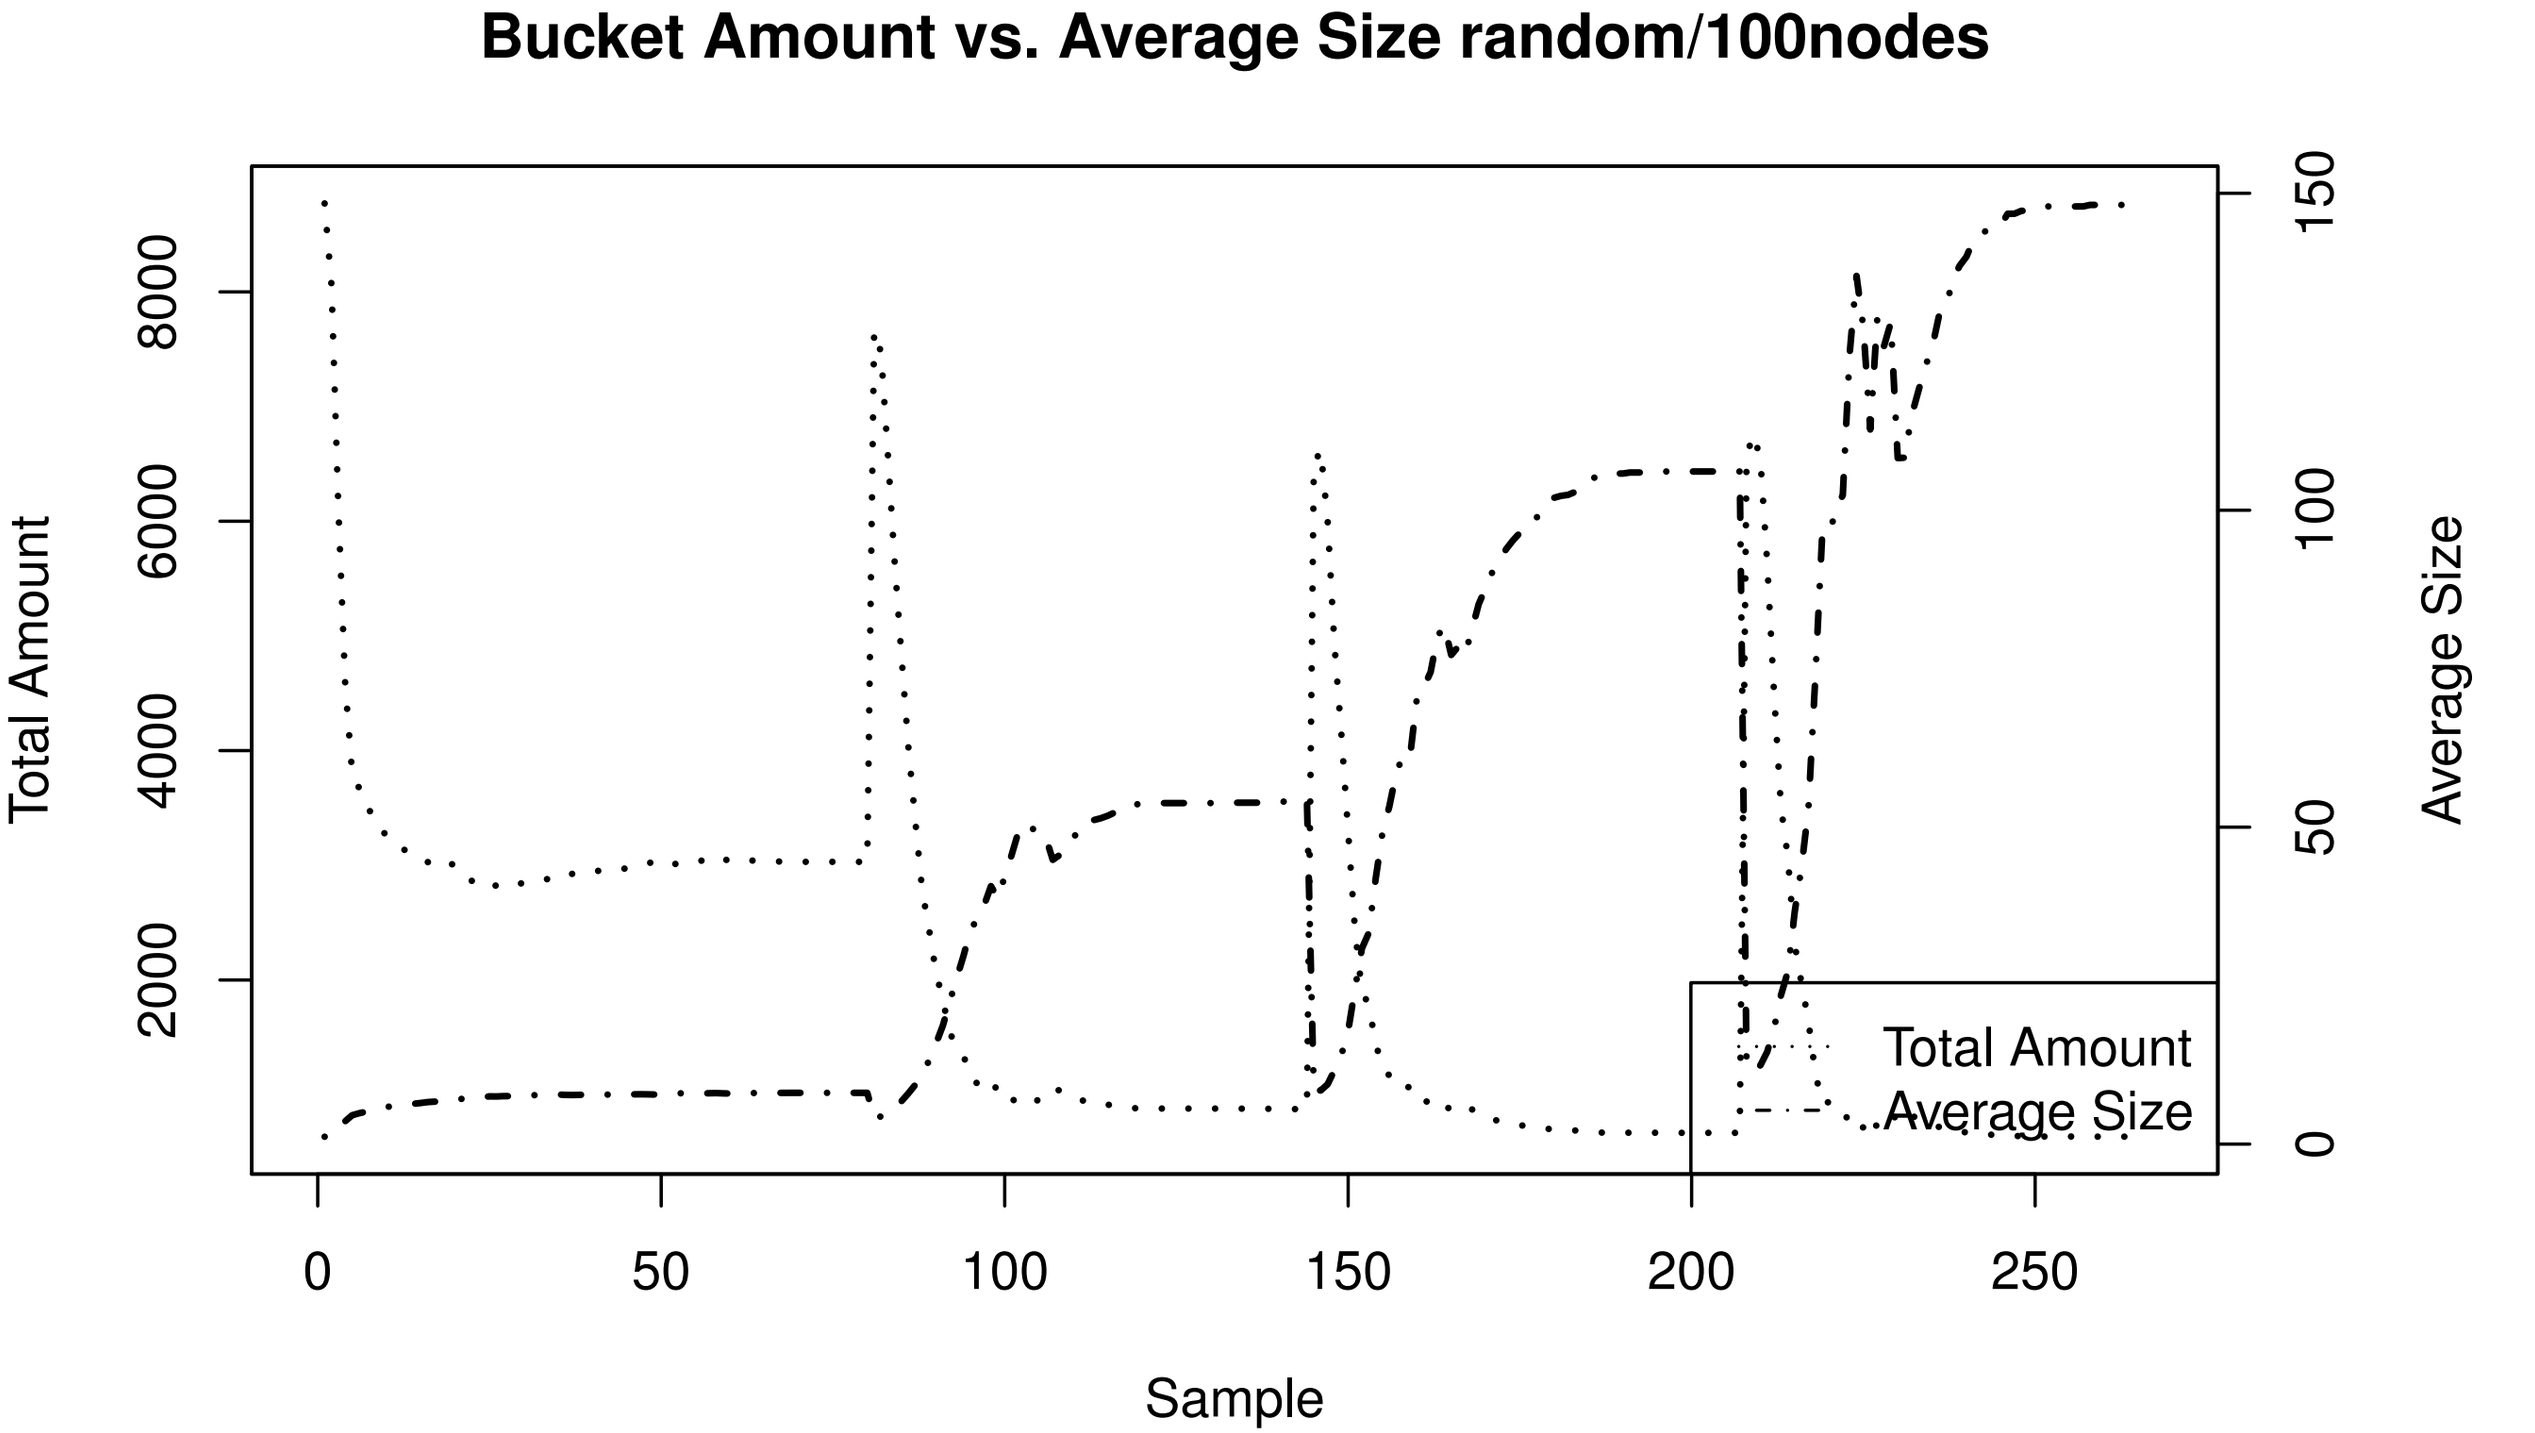
Bucket Amount vs. Average Size random/100nodes
Total Amount
Average Size
2000
4000
6000
8000
0
50
100
150
0
50
100
150
200
250
Sample
Total Amount
Average Size

Conclusion

- Brood Sorting works! *

* YMMV

Thank You!

Questions?

Web Page: http://hannes.muehleisen.org