

Continuous MapReduce on Heterogeneous Parallel Hardware for Stream Processing

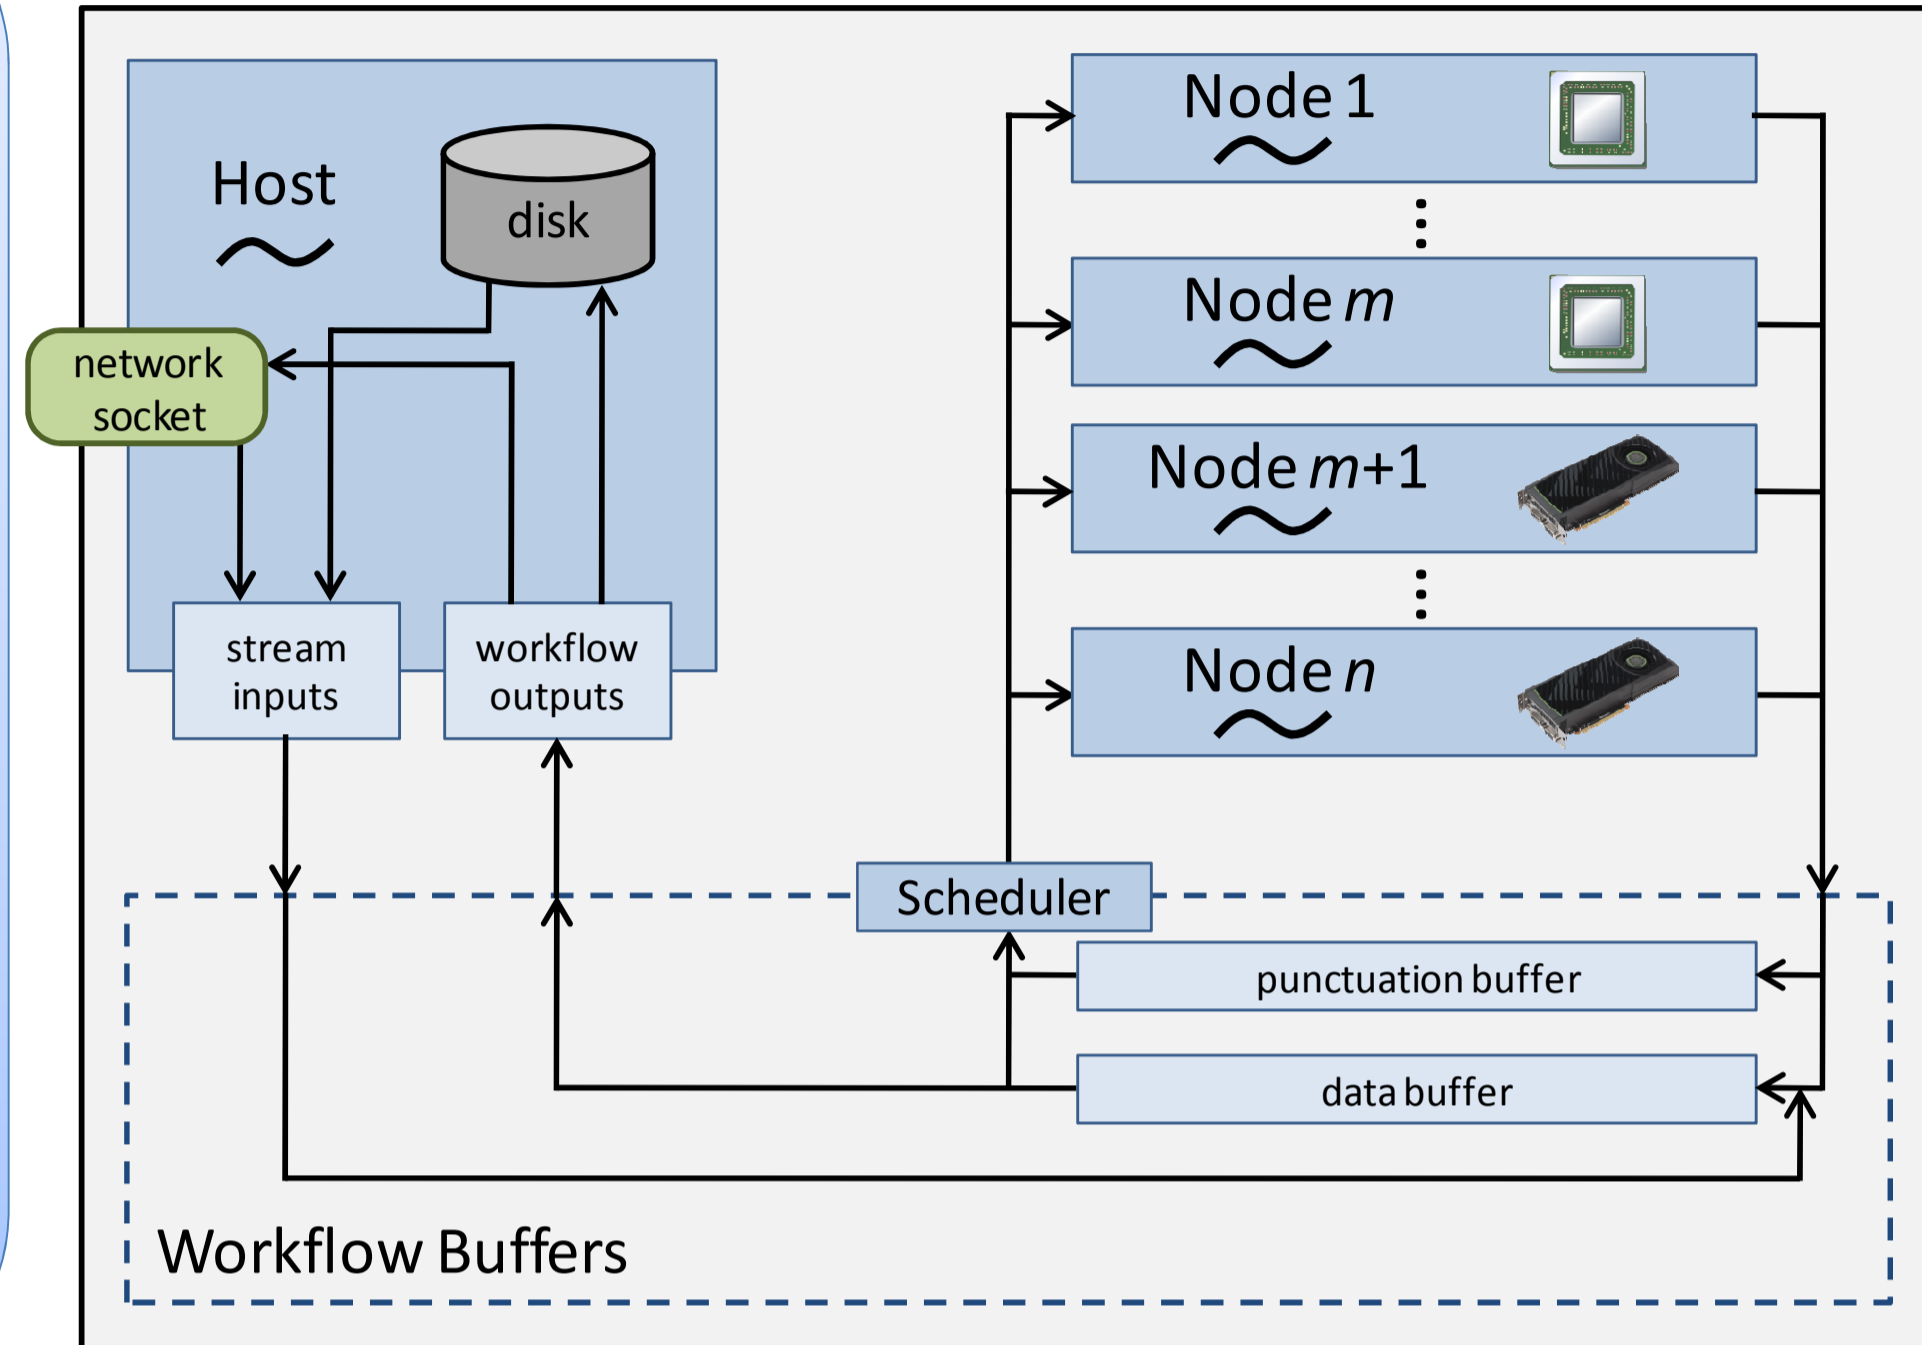
Nathan Backman, Karthik Pattabiraman, Uğur Çetintemel

- CPUs and GPUs are abstracted away into Nodes. They implement the Node's processData function which calls the appropriate code.


- Nodes asynchronously consume data from the Scheduler as they are able for *any* workflow operation.

- Fast Nodes consume more than slow or externally burdened Nodes.

C-MR Architecture



Expanded Workflow w/ Combine

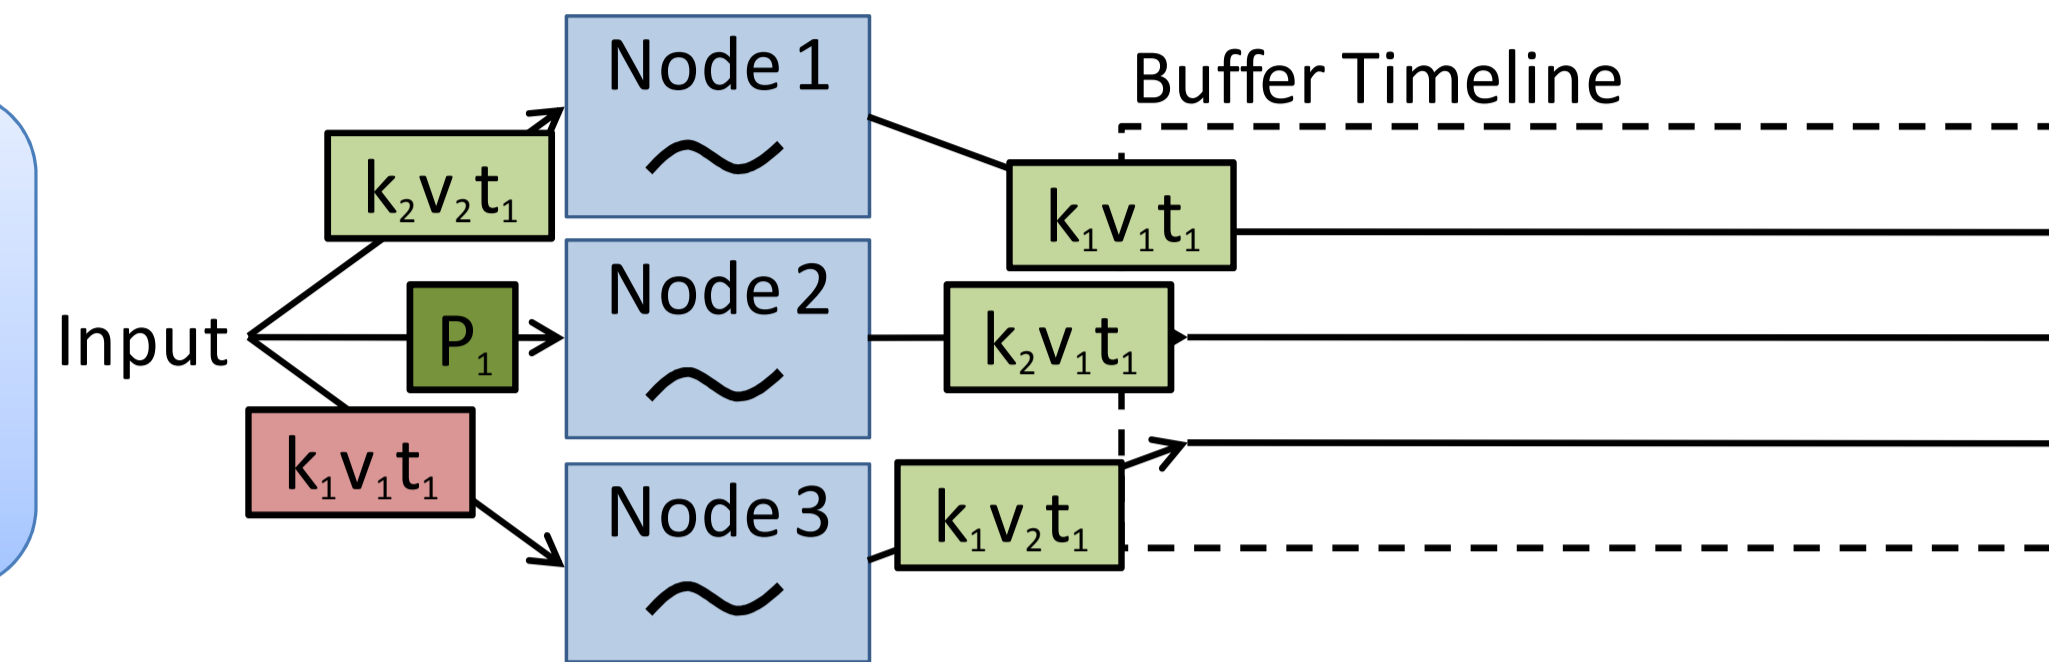


- Combine enables incremental processing while reducing redundant computations.

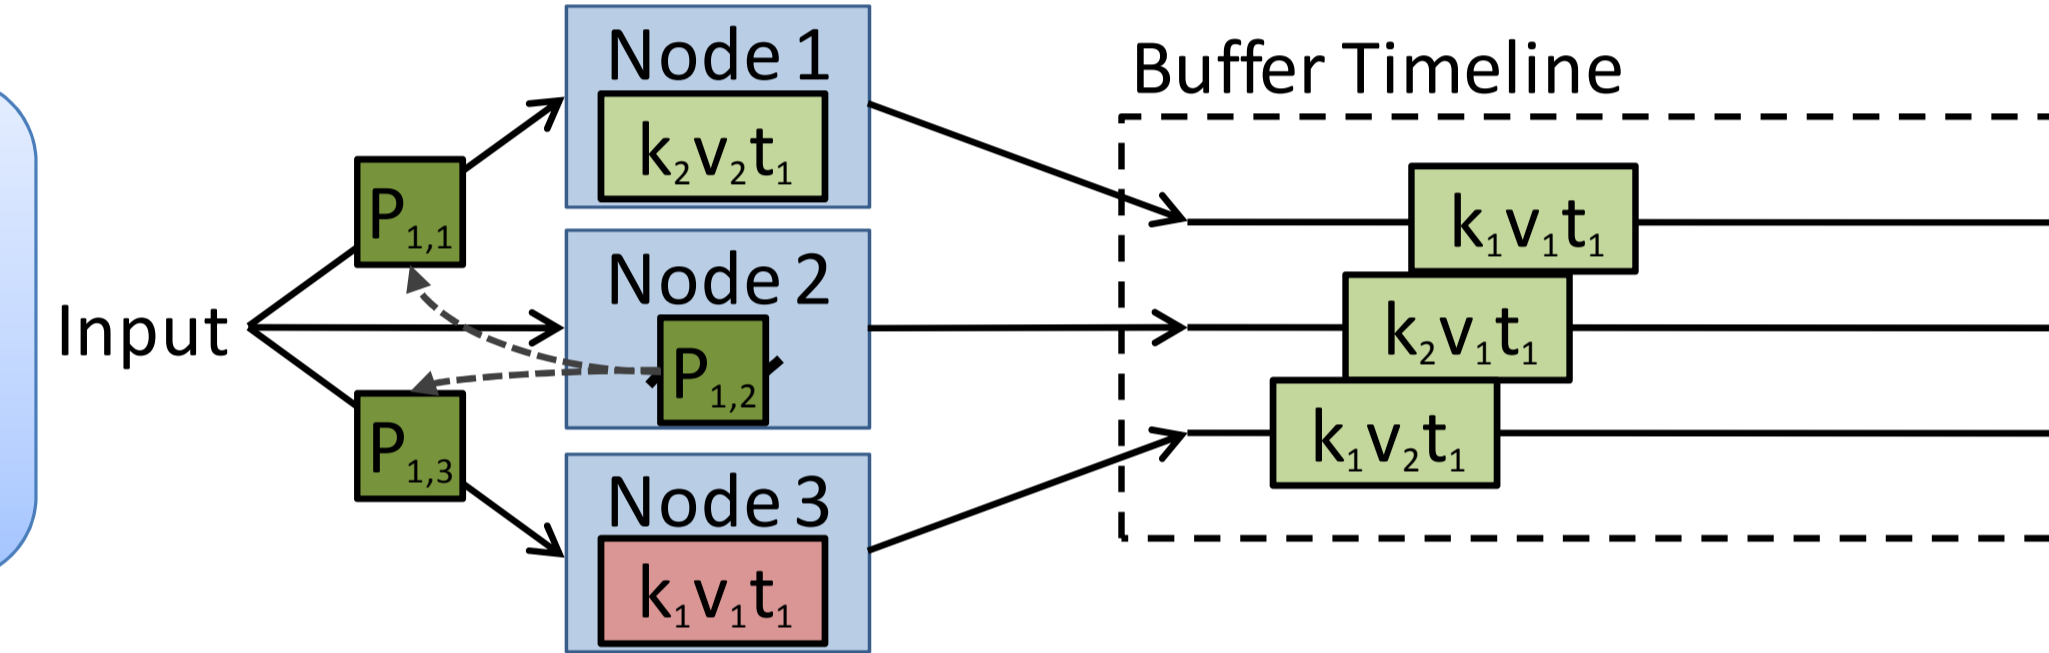
- Combine sub-windows can be shared across overlapping Reduce windows and operators.

Maintaining Stream Order

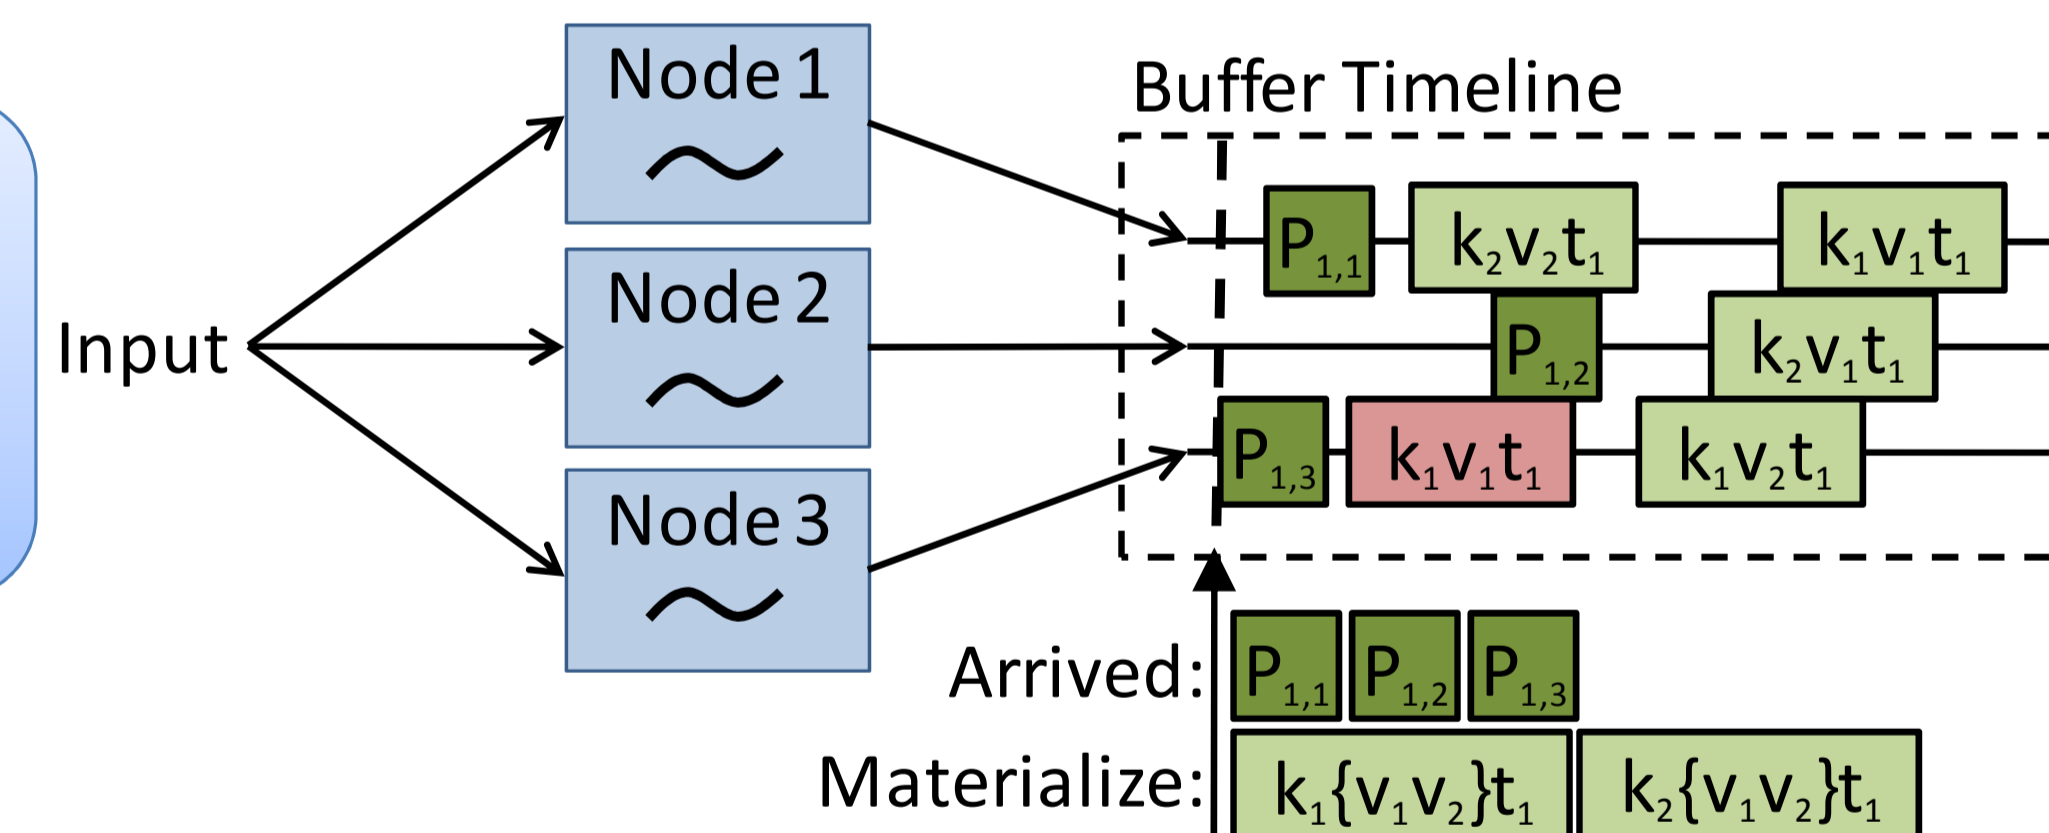
1 Nodes process data as they are received and place the results into the workflow buffer.



2 As a node receives a punctuation on the input, it is replicated to the other nodes.

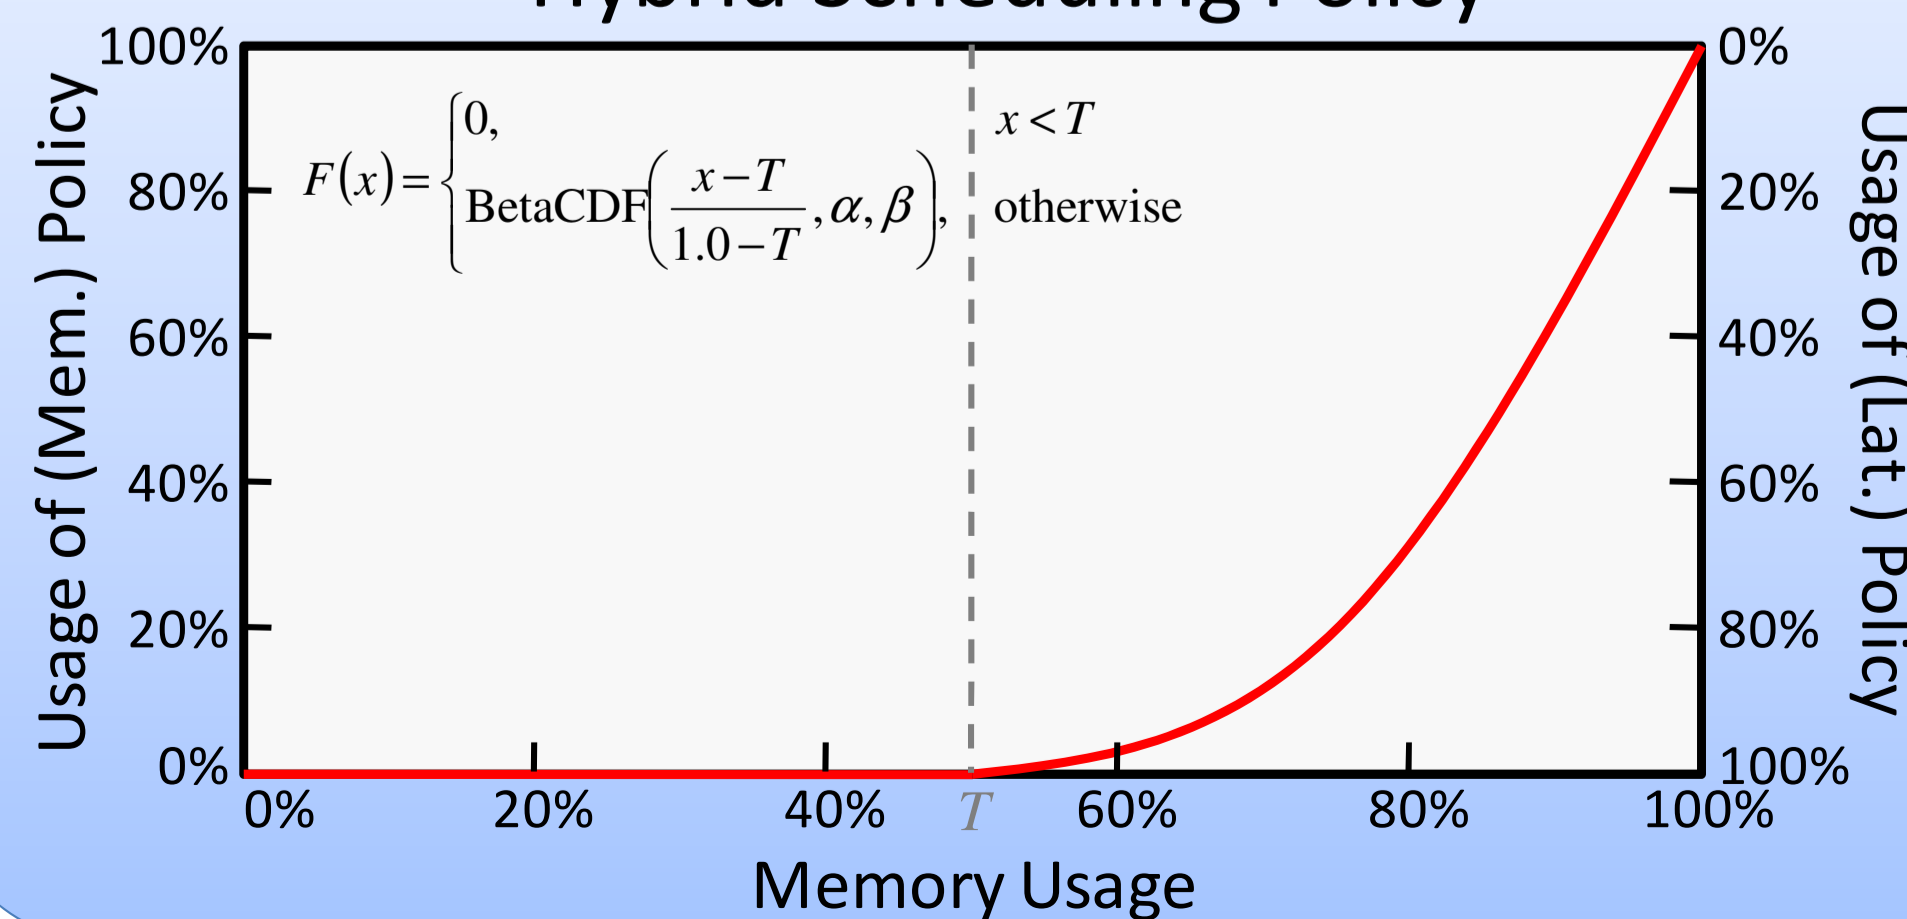


3 Once all replicated punctuations are received, then the corresponding window is formed.



Operator Scheduling

Hybrid Scheduling Policy

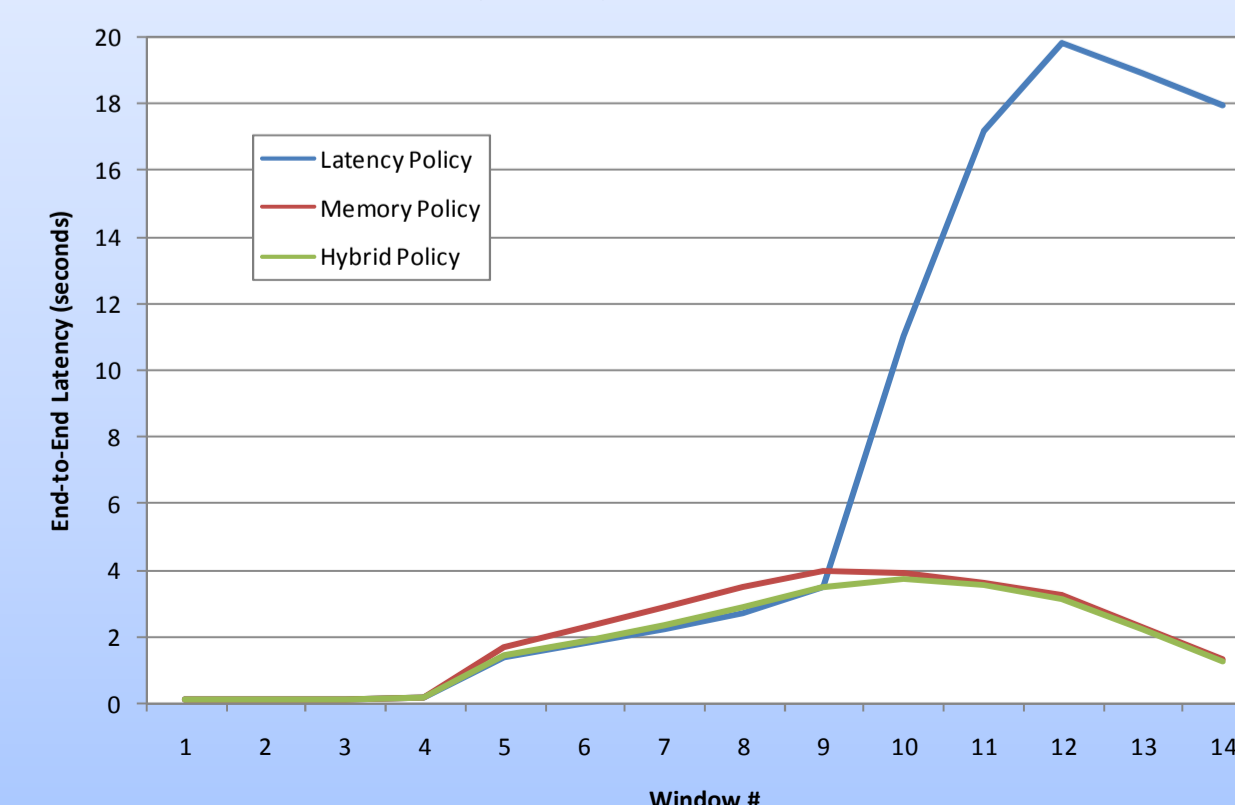


- Nodes request data through the scheduler which can assign data from any operator in the workflow.

- The scheduler can enact different scheduling policies for each request.

- Our hybrid policy transitions between latency and memory oriented policies as we approach running into swap space.

Latency Comparison of Policies



Memory Comparison of Policies

