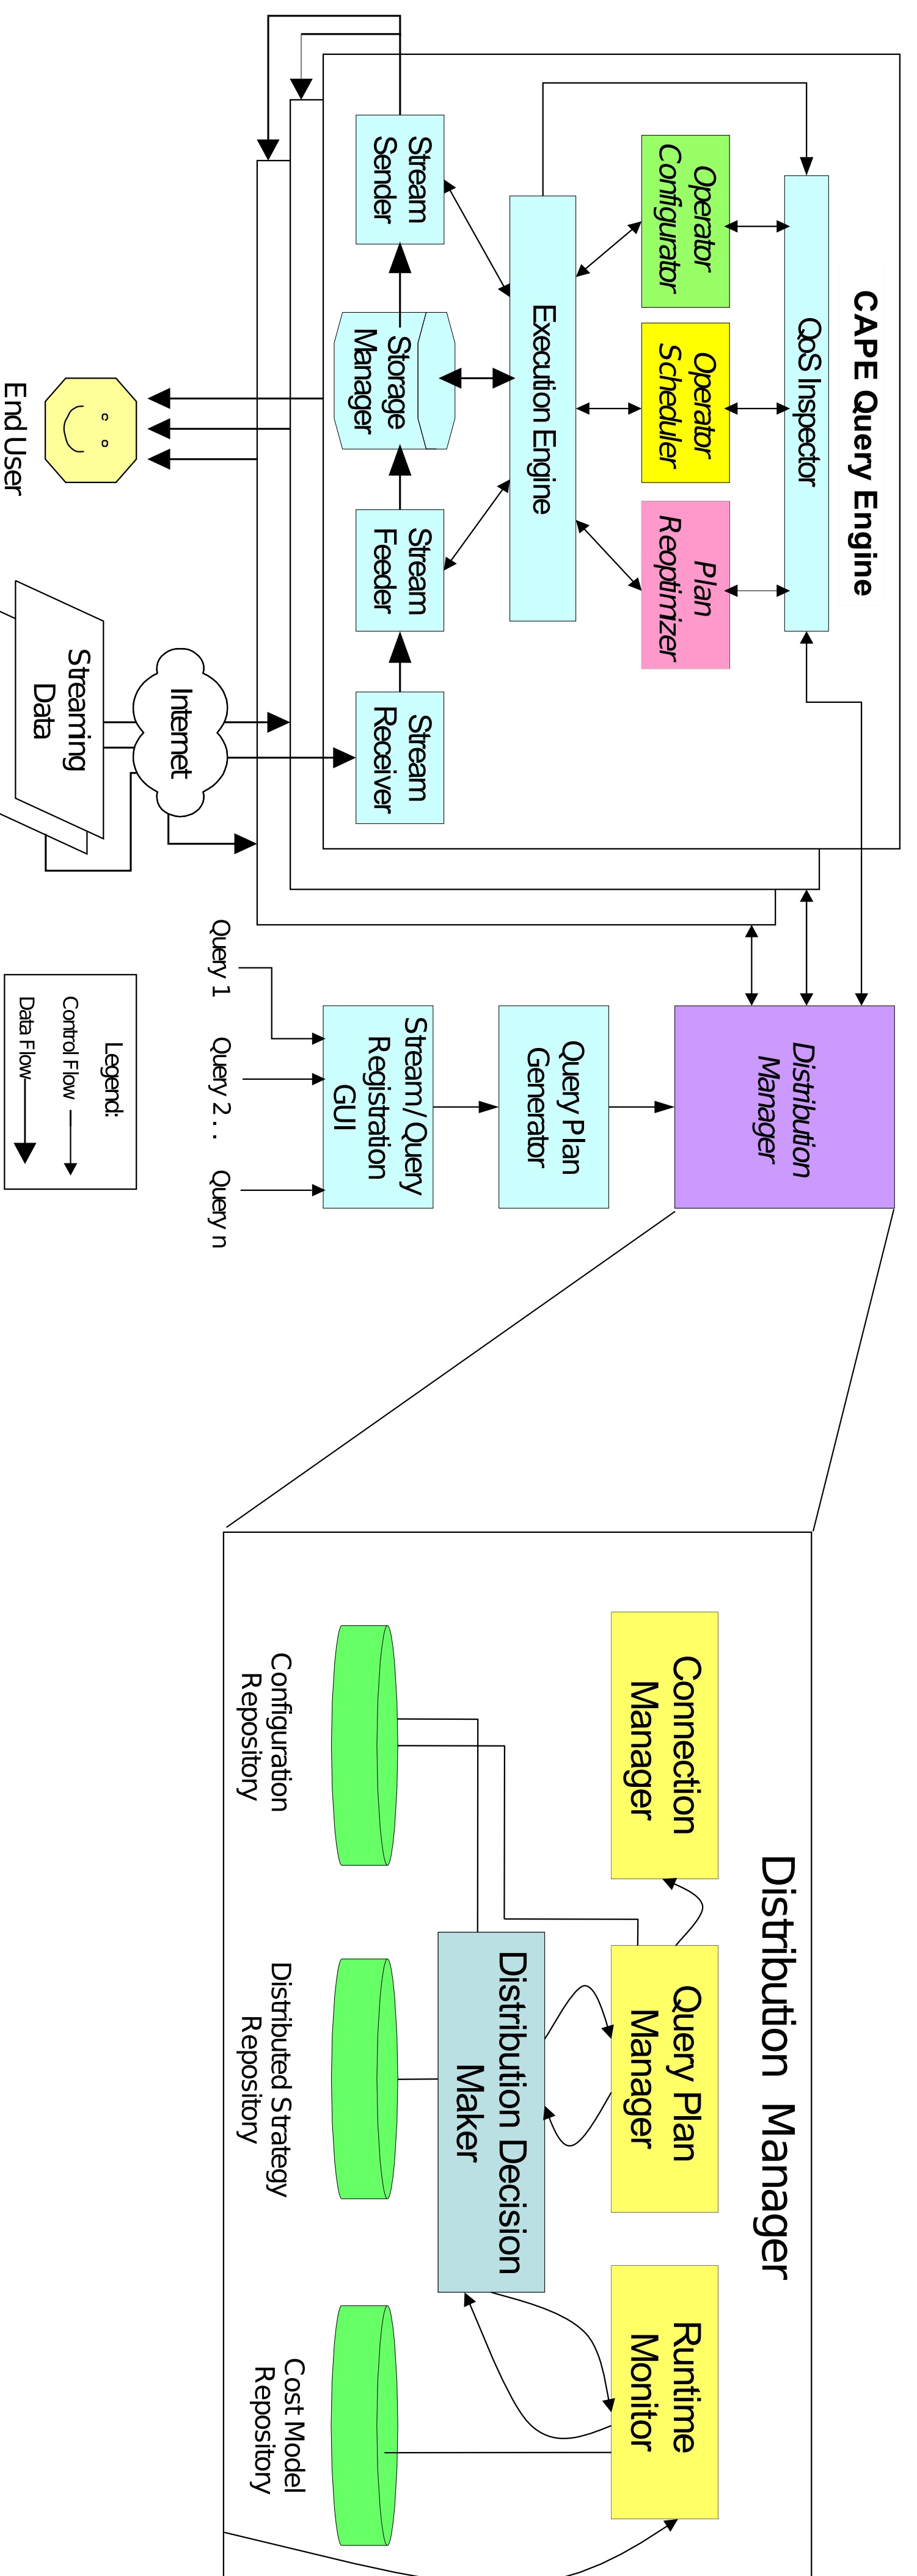


High-performance Data Stream Processing with Constraints and Parallel Network Execution in the Distributed CAPE Processing Engine

Distributed CAPE System Architecture

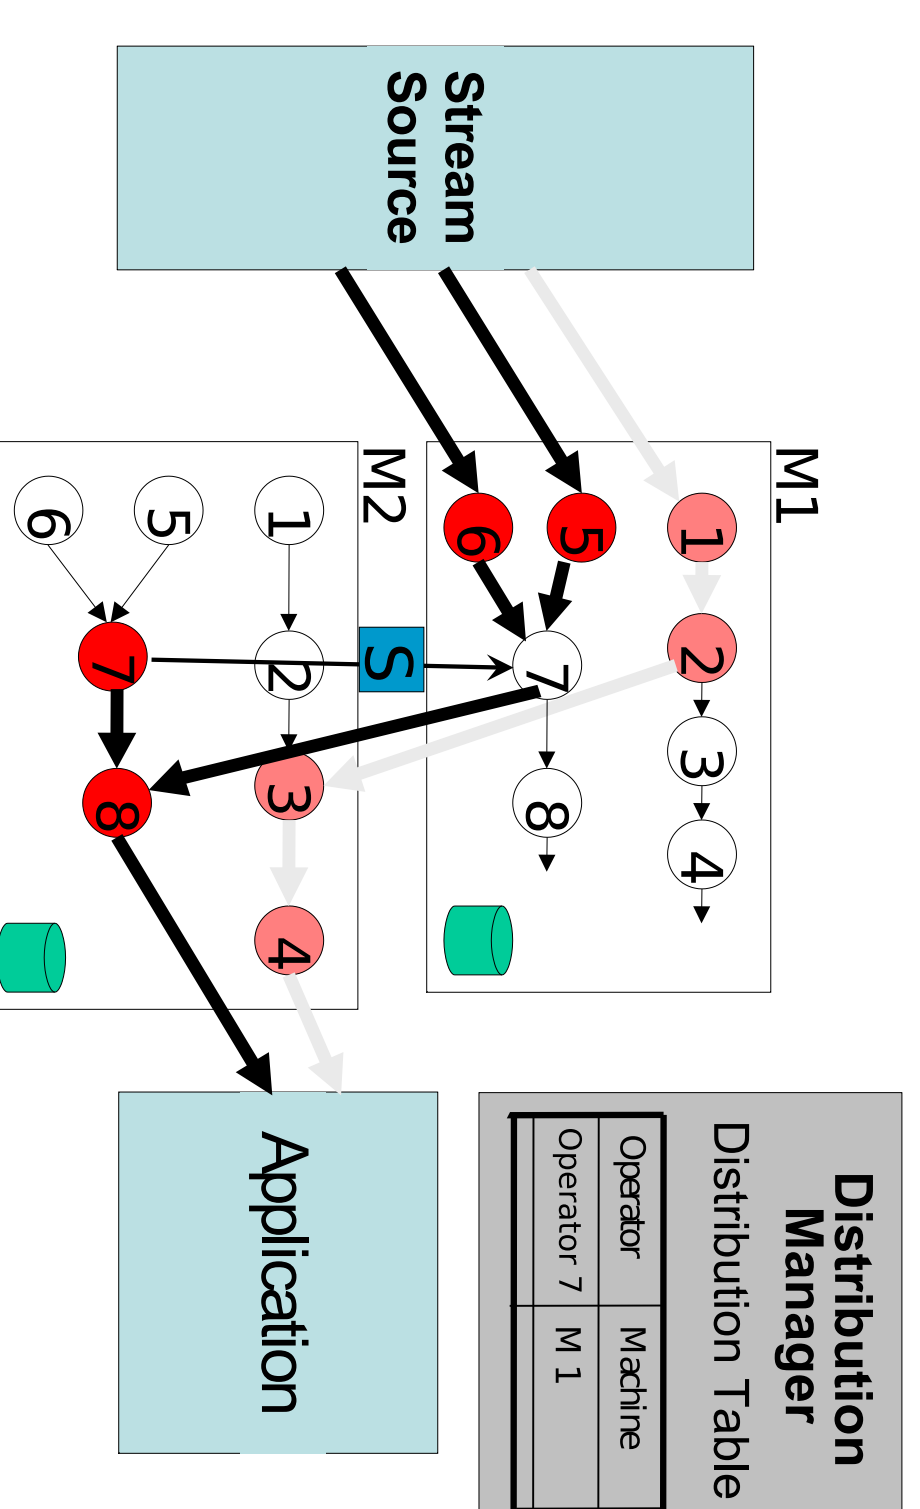



Initial Distribution Process

Step 1

Distribution Manager	
Operator	Machine
Operator 1	M 1
Operator 2	M 1
Operator 3	M 2
Operator 4	M 2
Operator 5	M 1
Operator 6	M 1
Operator 7	M 2
Operator 8	M 2

Step 2



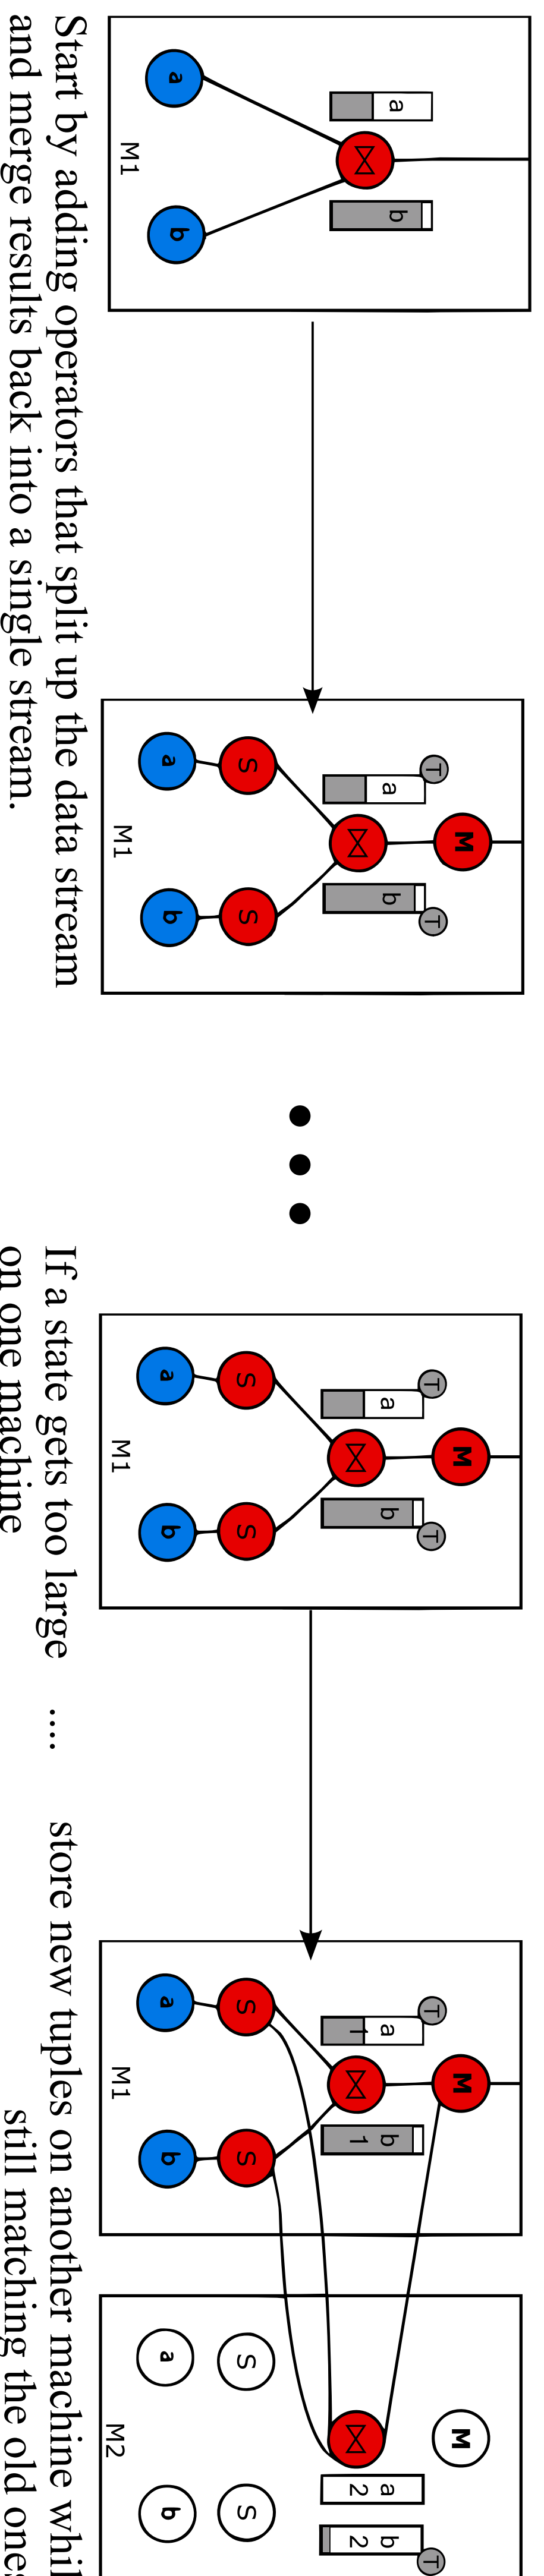
Operator Reallocation

Create distribution table using initial distribution algorithm.

Send distribution information to processing machines (nodes).

Move operator and any associated state wholesale from one machine to another.

SPREADFIRE: Unhashed State Partitioning



If a state gets too large on one machine


... store new tuples on another machine while still matching the old ones.

Bradley Moberger
Yali Zhu
Bin Liu
Mariana Jbantova
Elke Rundensteiner

WPI Database Systems Research Laboratory

Partitioned Parallel Processing

Partition State-Intensive Operators into Hashed "Buckets" on Multiple Machines




Initially processing all data on one machine

After split, process only partial data in each machine


Dealing with Overload

State Relocation



Relocated state partitions are active in another machine

State Spill



Push state partitions temporarily to disk