RESEARCH

Programming your network at run-time for big data applications

Guohui Wang, *IBM TJ Watson Research Center*

T. S. Eugene Ng, *Rice University*

Anees Shaikh, *IBM TJ Watson Research Center*

Application-aware networking – a long history

- *Protocol parameter tuning*
- *Overlay routing*
- *Per-application routing instances*

edge-based, limited network control

- *Reconfigurable topologies (e.g., hybrid optical)*

coarse-grained, limited app-level visibility

- *Embedding code in the network (active networking)*

ultimate control, risky, limited deployability

Revisiting application-aware networking

Software Defined Networking / OpenFlow

- *flexible network control using rules on packet headers*
- *rule installation in milliseconds*

Reconfigurable optical circuit switching

- *reconfigurable physical topology using OCS*
- *link reconfiguration in milliseconds*

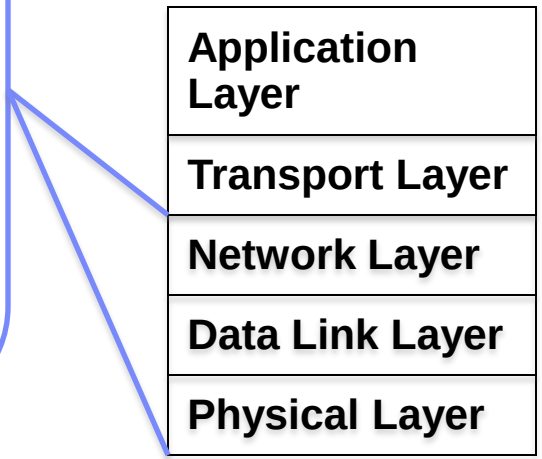

Flexible control over the network

How and which applications can benefit?

Characteristics of big data applications (e.g., MapReduce)

Big data applications

- *well-defined communication structure*
- *centralized control architecture*
- *high bandwidth requirements*

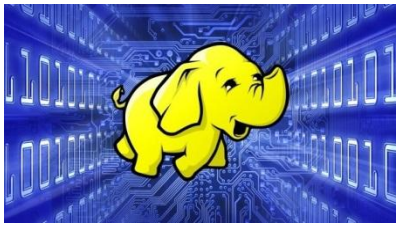

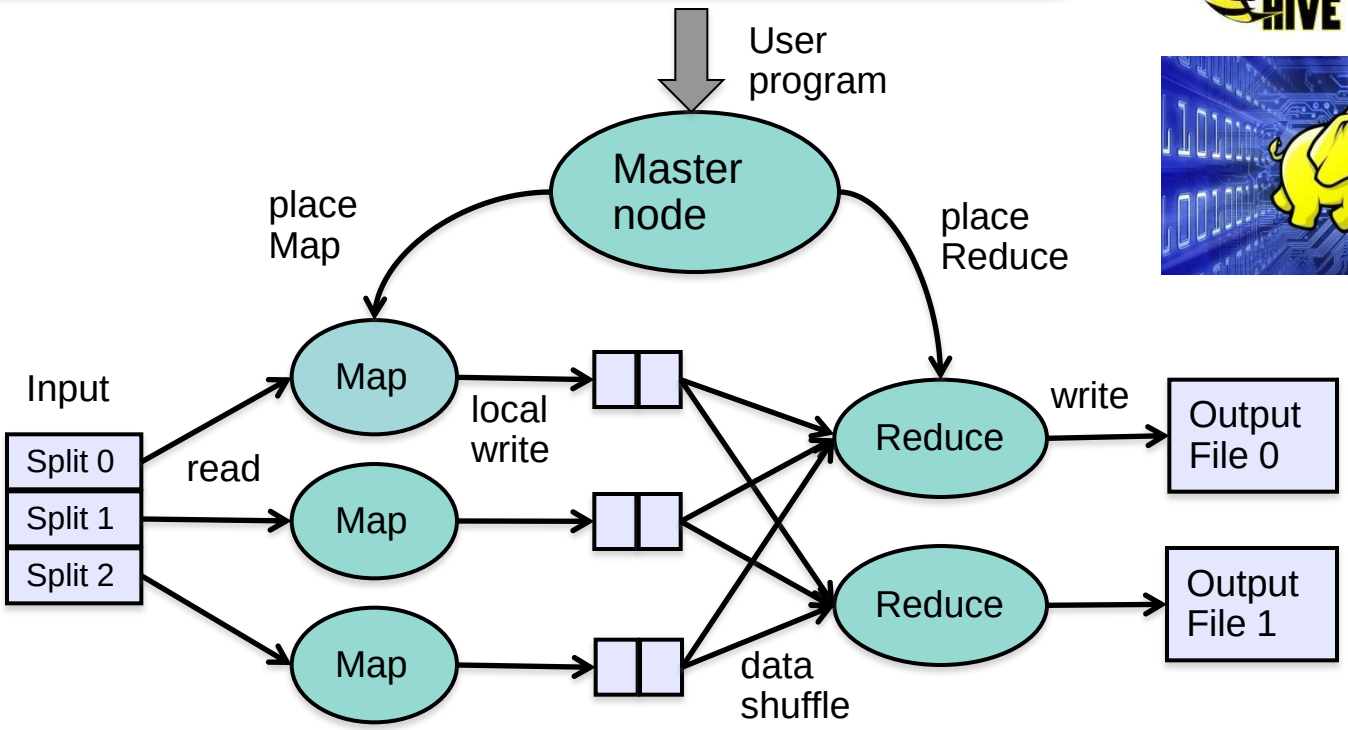

New opportunity to exploit SDN and reconfigurable topologies for better application performance

Cross-layer control architecture

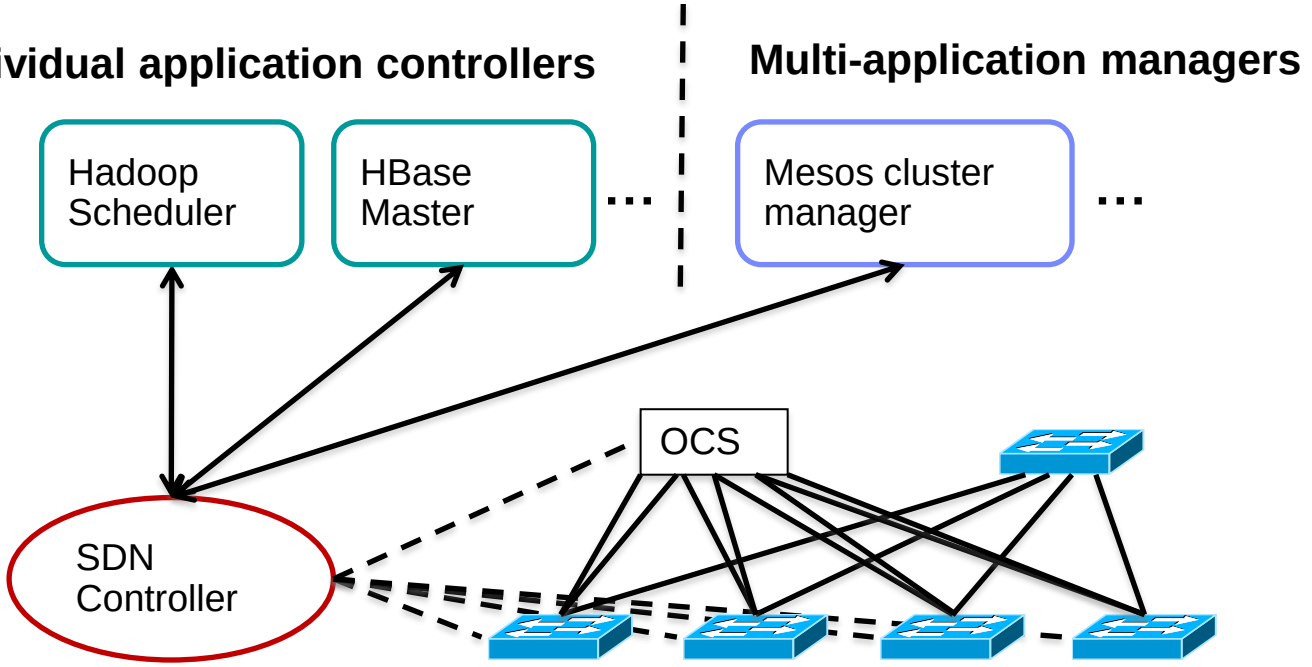

Software Defined Networking for big data applications

- *cross-layer network control* plane that combines application controller and SDN controller
- run-time topology construction and routing configuration based on application specific traffic structure

Workload optimized networking

- *tight application and network integration*
- *jointly optimize network utilization and application performance*

Application controller operation (Hadoop example)

- ***Job placement:***
  - *rack-based bin packing placement of map/reduce tasks*
  - *batch processing (correlated tasks in one batch)*
- ***Traffic demand estimation:***
  - *job tracker traffic volume analysis*

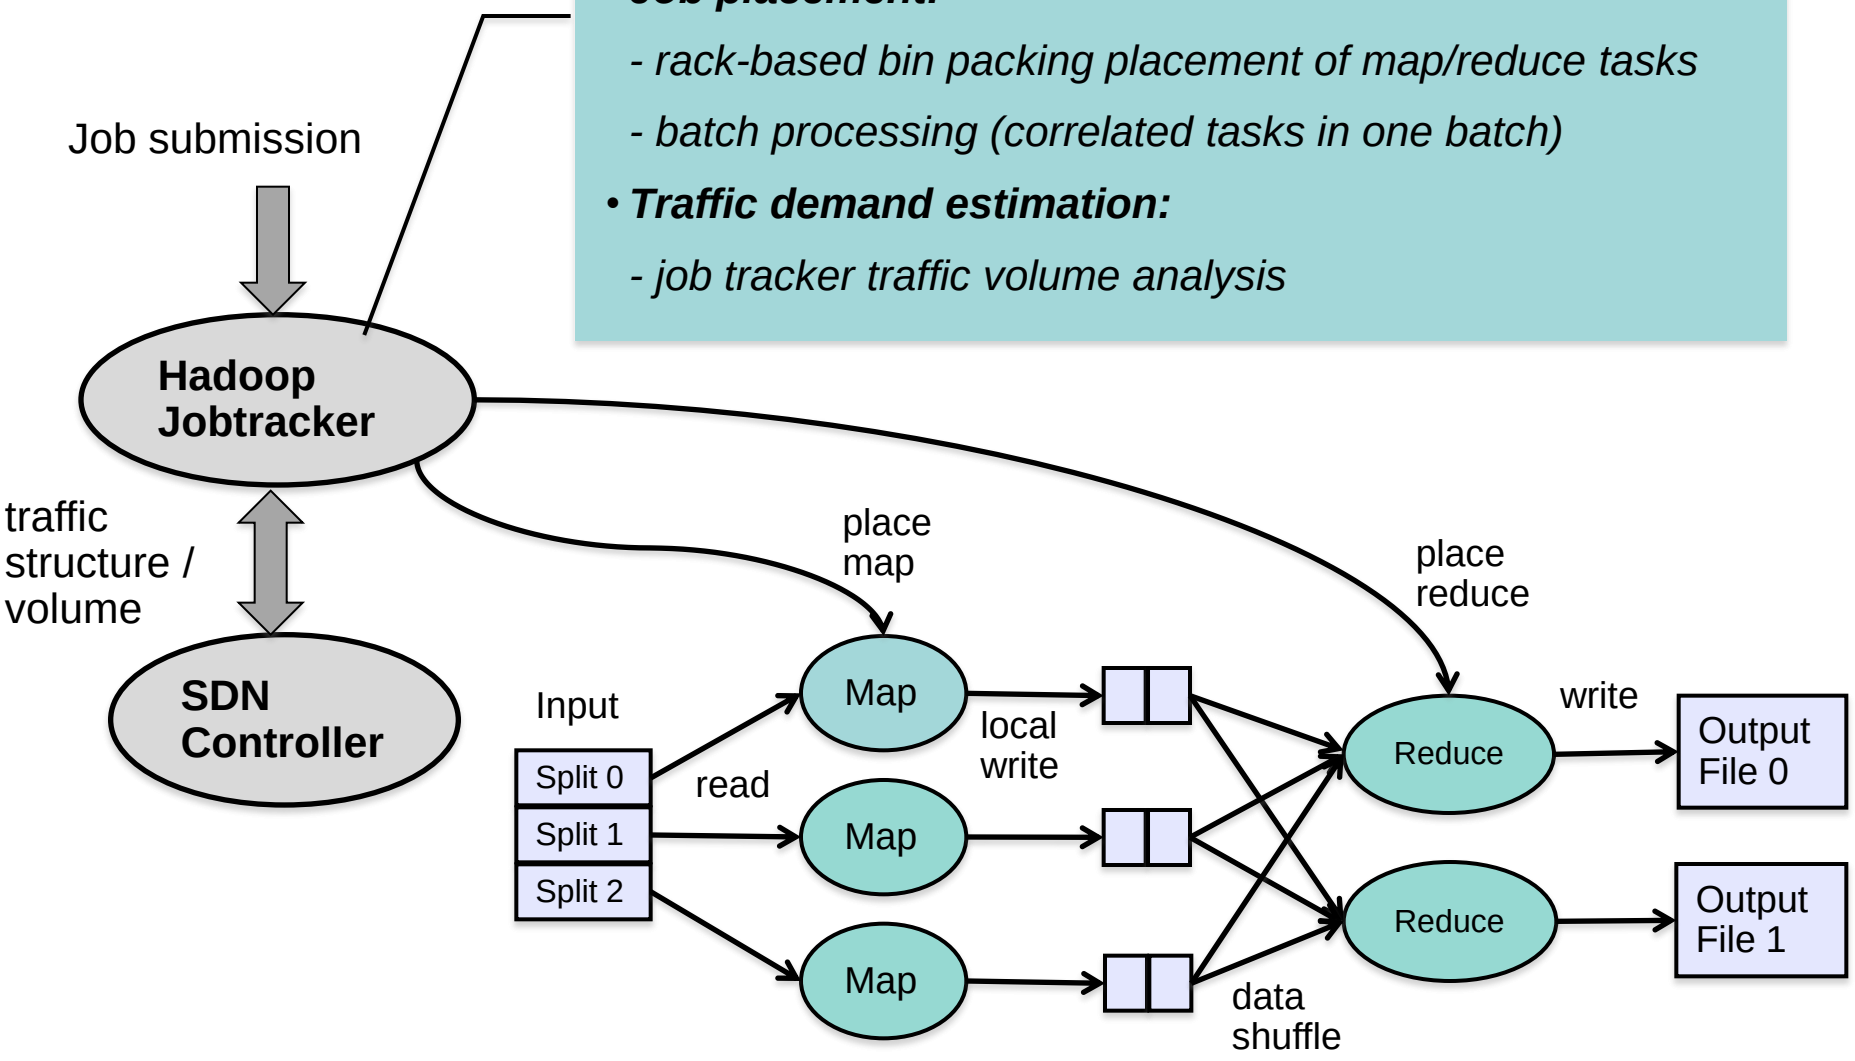

Network controller operation (Hadoop example)

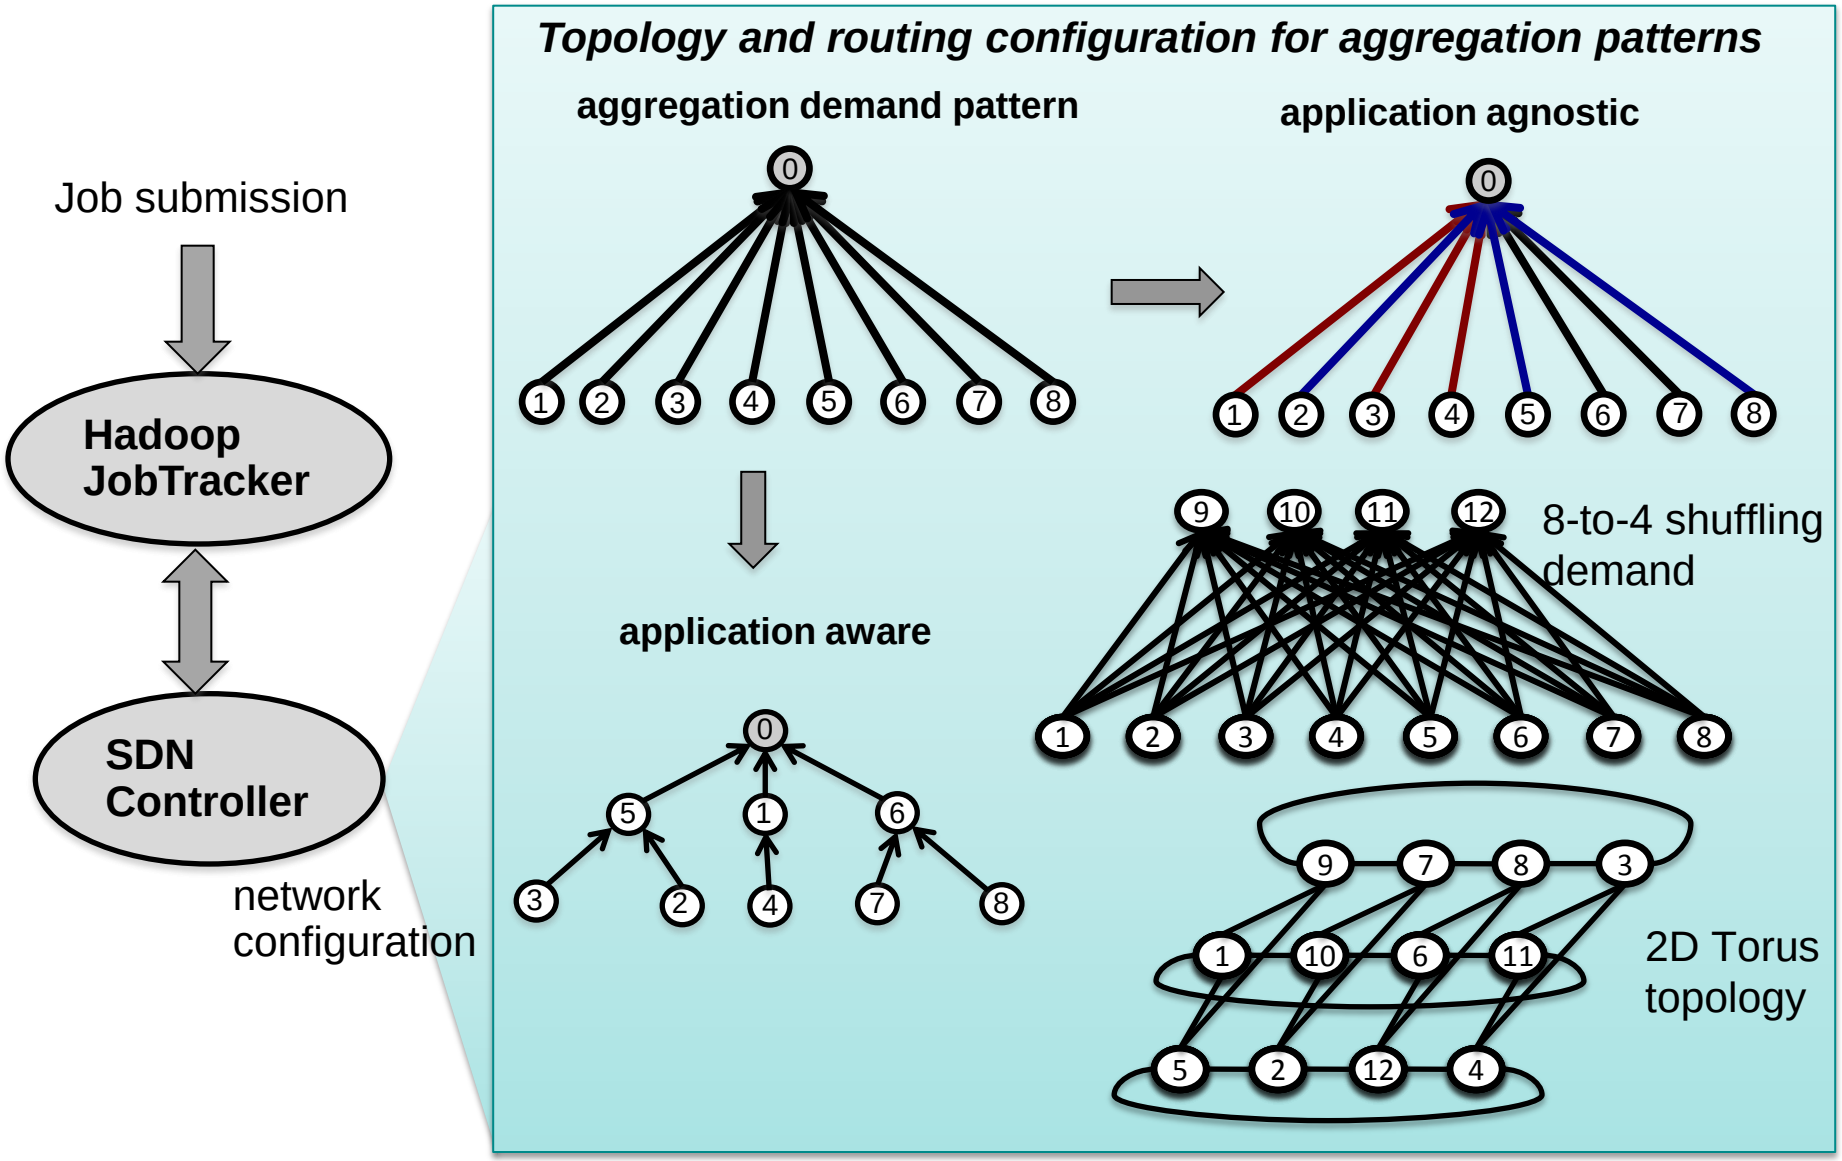

Design implications and future work

Implications on the SDN controller

- how to update network-wide state with low delay requirements?
  - update network topology and routing in tens of ms
  - tag packets and use proactive rule installation to reduce the number of rule updates
- how to maintain state consistency for fast updates?
- how to coordinate configuration requests among different kinds of applications?

Ongoing / future work

- implementation / eval with MEMS-based OCS and OF switches
  - traffic volume estimation, placement strategies, flow rule installation
- flow-level traffic engineering for big data apps (e.g., alternative to OCS)
- handling fairness and prioritization with multiple application controllers