



Jon Weissman and Abhishek Chandra
(+ great UMn students)

Department of CS&E
University of Minnesota

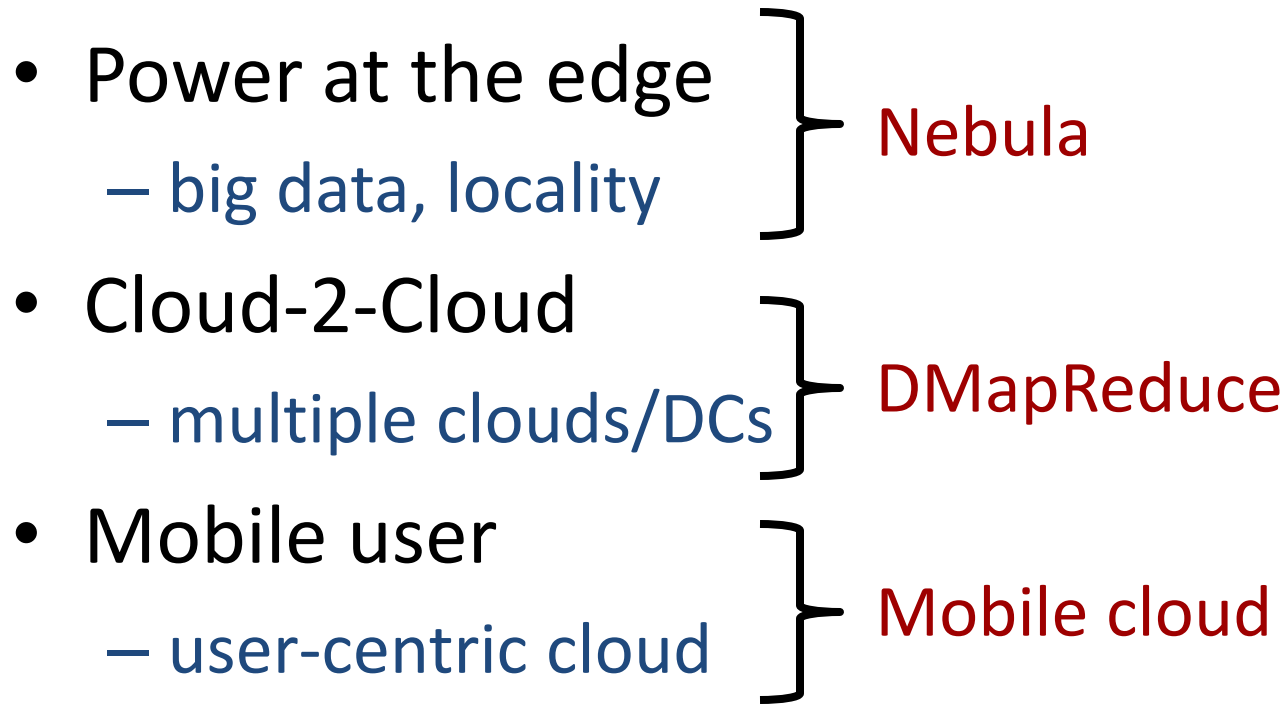
Key Trends

- Client technology
 - devices: smart phones, ipods, tablets, sensors
- Big data
 - located at the network edge
- Privacy/trust
 - local clouds
- Multiple DCs/clouds
 - global services

Minnesota Cloud Research

- Eye towards cloud evolution
- Projects
 - Nebula
 - DMapReduce
 - Mobile cloud
 - Proxy cloud
 - Active cloud storage
 - Virtualization

Project Clusters

- Power at the edge
 - big data, locality
 - Cloud-2-Cloud
 - multiple clouds/DCs
 - Mobile user
 - user-centric cloud
- Nebula
- DMapReduce
- Mobile cloud
- 

Big Data Trend


- Big data is distributed
 - earth science: weather data, seismic data
 - life science: GenBank, NCI BLAST, PubMed
 - health science: GoogleEarth + CDC pandemic data
 - web 2.0: user multimedia blogs
 - “everyone is a sensor”
- Cost in moving data to the cloud

Big Data Trend: Nebulas


- Process data “close by”
 - fully and/or on-route to the central cloud
 - cost: save time and money
 - privacy (think: patient records)
- Close by
 - network distance
 - trusted peers

Example: Dispersed-Data-Intensive Services

- Data is geographically distributed
 - Costly, inefficient to move to central location




Example Instance: Blog Analysis




Nebula: A New Cloud Model


- Make the cloud more “distributed”
 - exploit the rich collection of edge computers
 - volunteers (P2P, @home), commercial (CDNs)
 - enormous computing potential, network diversity
 - lower latency: “on demand”, native client sandboxing



Example: Blog Analysis



Blog Results




Current Status

- Prototype running on PlanetLab
- Distributed Data-Store Service
- Network Dashboard Service
- Node software stack
 - native client

Nebula Going Forward

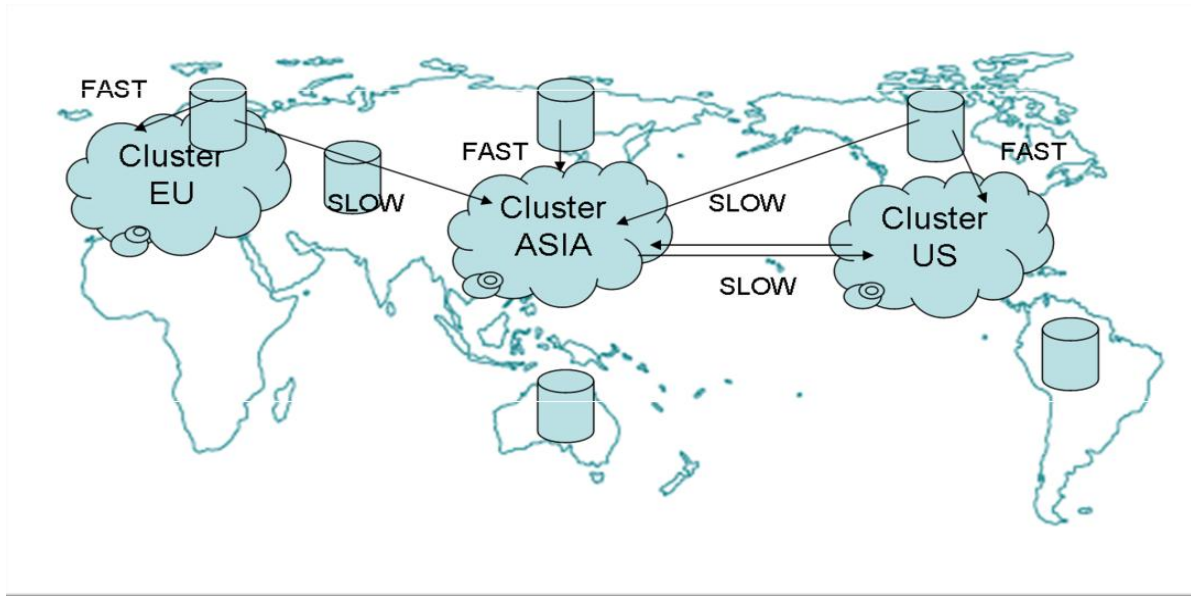
- Organize Nebulas
 - around trusted peers
 - social groups
 - communities of interest
 - local resources
- Nebula + commercial cloud
 - “use the edge opportunistically”

Big Data Trend: DMapReduce

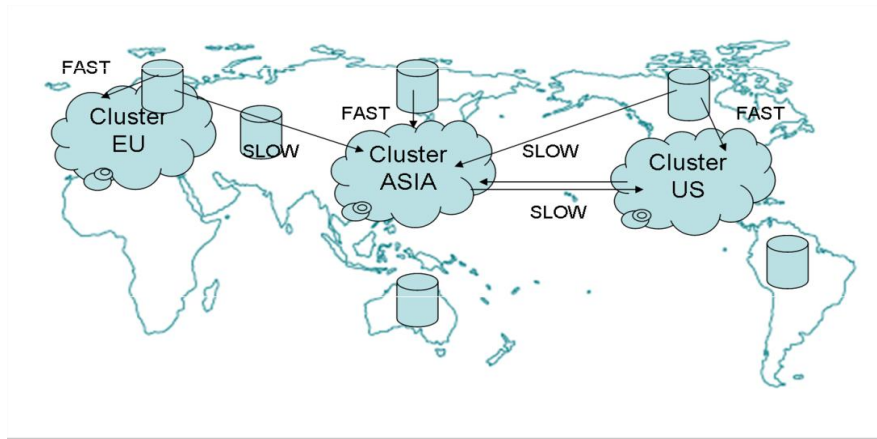


Big Data Trend: Distribution

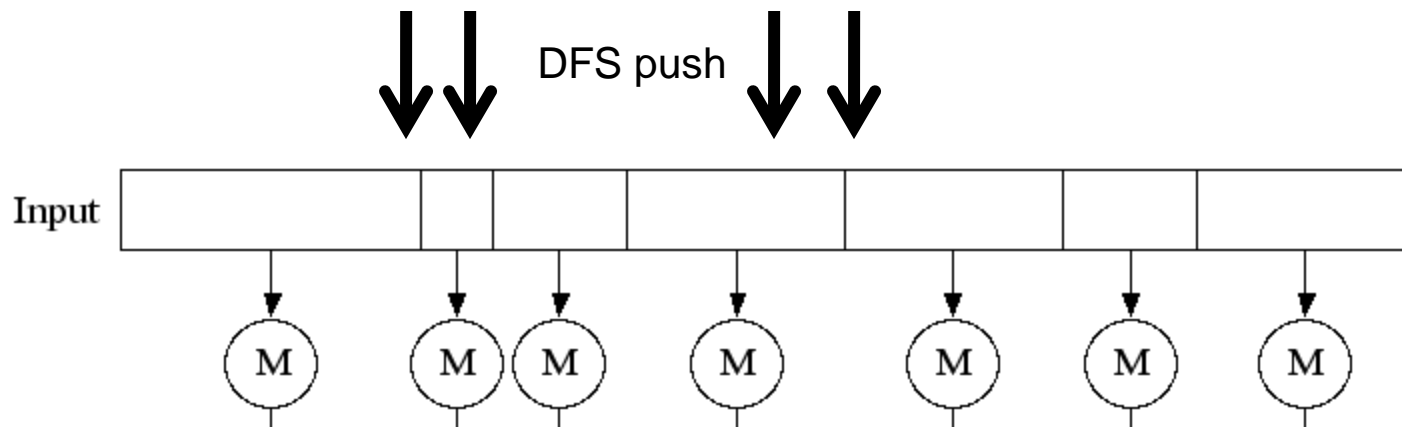
- Big data is distributed
 - earth science: weather data, seismic data
 - life science: GenBank, NCI BLAST, PubMed
 - health science: GoogleEarth + CDC pandemic data
 - web 2.0: user multimedia blogs



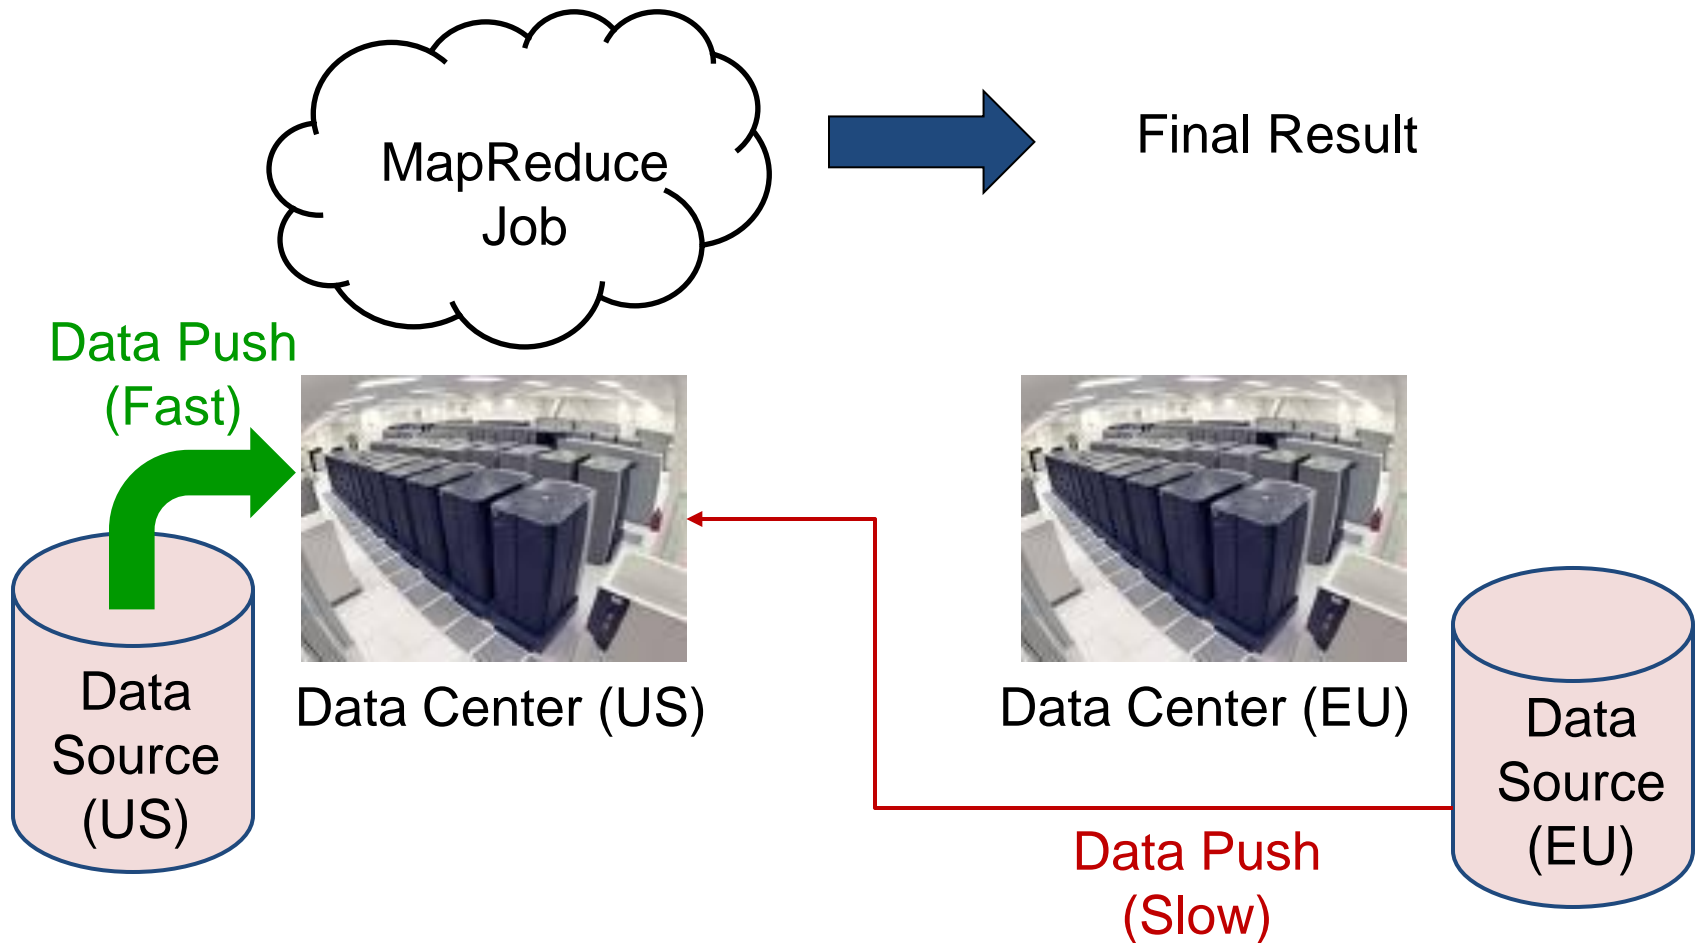
Wide-Area MapReduce



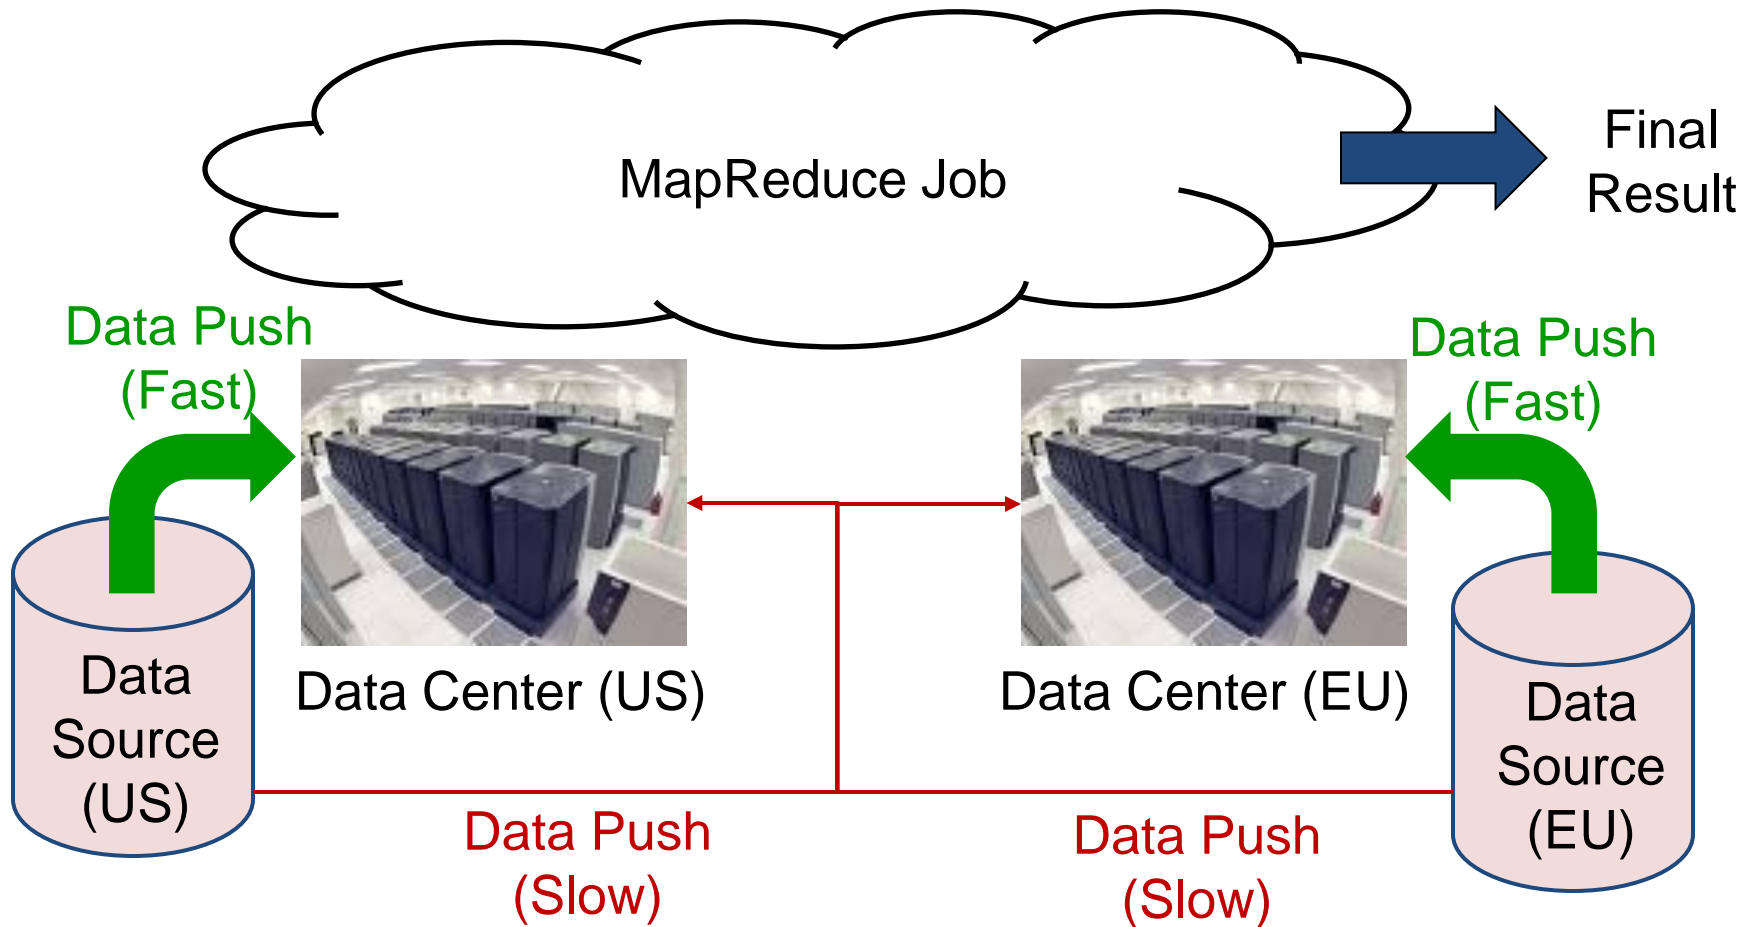
- Data in different data-centers
- Run MapReduce across them
- Data-flow spanning wide-area networks



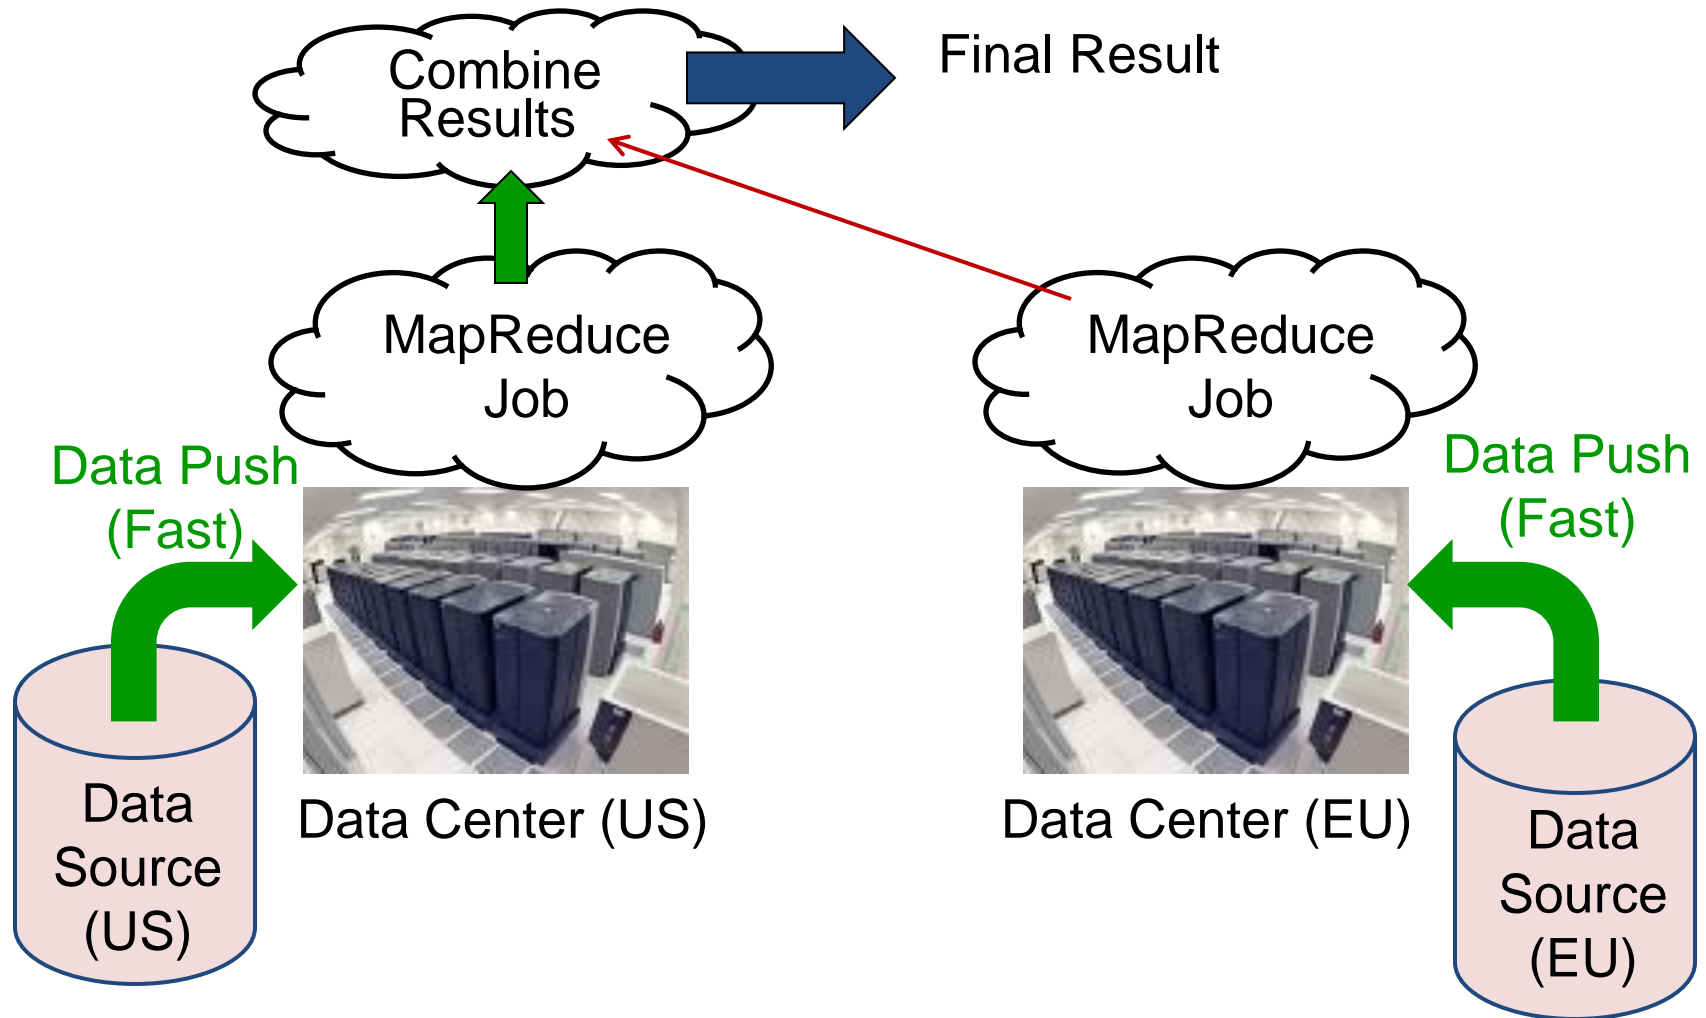
Option 1: Local MapReduce




Option 2: Global MapReduce



Option 3: Distributed MapReduce




PlanetLab: WordCount (text data)



- Distributed MR works best in presence of data aggregation


PlanetLab: WordCount (random data)




- Local MR works best in presence of data ballooning

EC-2 Results


WordCount (Text)



WordCount (Random)

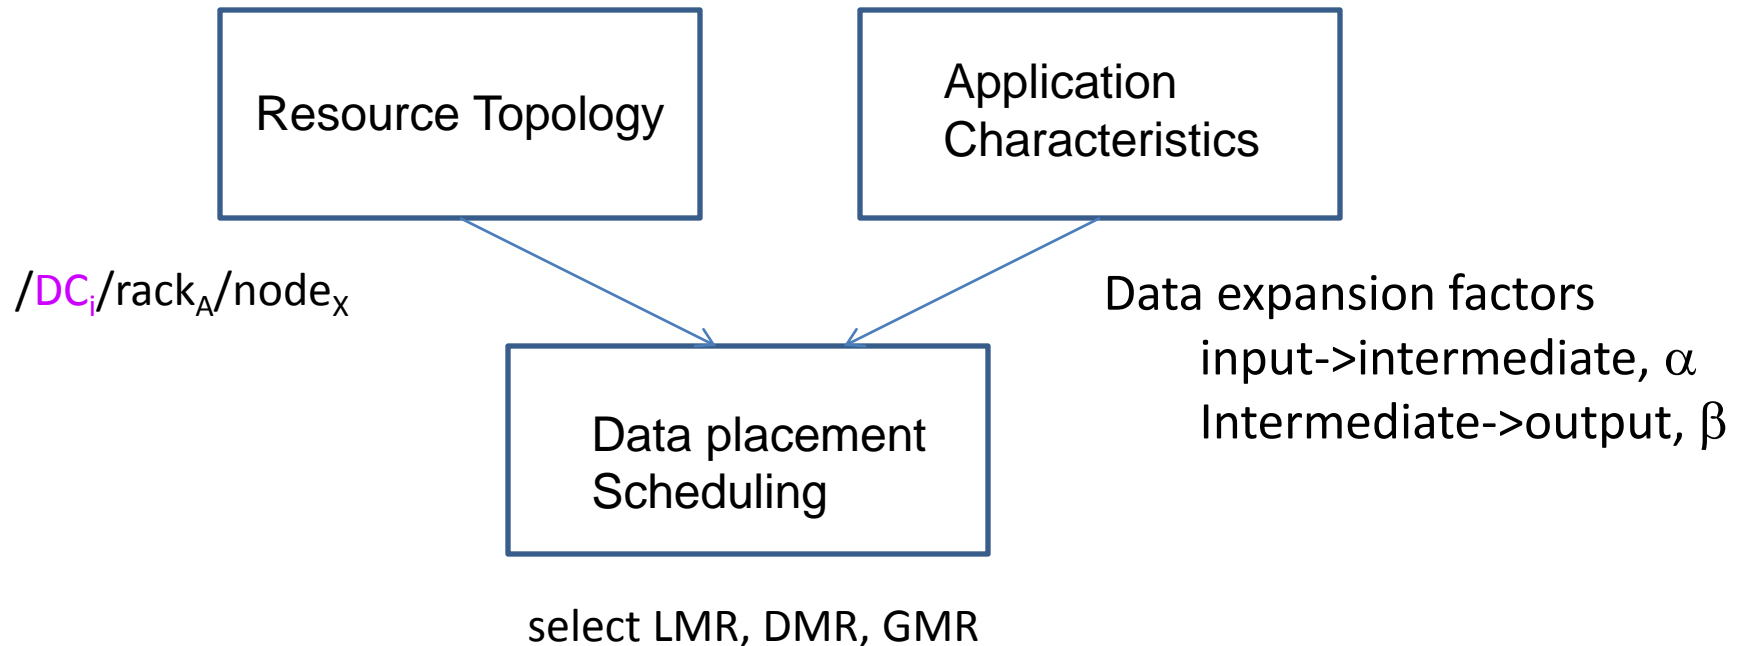


Sort



Intelligent Data Placement

- HDFS push
 - local node, same rack, random rack



Generalize beyond HDFS/MapReduce!

Mobility Trend: Mobile Cloud


- Mobile users/applications: phones, tablets
 - resource limited: power, CPU, memory
 - applications are becoming↑sophisticated
- Improve mobile user experience
 - performance, reliability, fidelity
 - tap into the cloud **dynamically** based on current resource state, preferences, interests, privacy

Cloud Mobile Opportunity

- Dynamic outsourcing
 - move computation, data to the cloud dynamically
- User context
 - exploit user behavior to pre-fetch, pre-compute, cache
- Multi-user sharing
 - discover implicit cloud sharing based on interests, social ties



Outsourcing

- Partitioning across (Mobile \leftrightarrow Server \leftrightarrow Cloud)
 - local data capture + cloud processing
 - images/video, speech, digital design, aug. reality

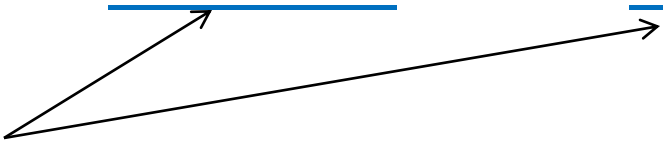


Experimental Results -Image Sharpening


- Response time
 - both WIFI & 3G
 - up to 27× speedup
 - 219K, WIFI
- Power consumption
 - save up to 9× times
 - 219K, WIFI



Cloud Speculation

- Dynamic user *profile*
 - contains *activities* in time and space
 - “read nytimes.com bet. 9-9:30am on the train; likes technology articles”
 - associate with *context*
- Patterns are relationships between activities
 - repetitive, sequential, concurrent, time-bounded
 - “user always does X and then does Y”
- Exploiting patterns: pre-fetching, pre-computing, caching in the cloud

Smart cloud



Back-end speculation and optimization

Summary

- Cloud Evolution
 - mobile users, big data, privacy/trust, global
- Our vision of the Cloud
 - locality of users, data, other clouds/data centers, user-centric behavior