Rethinking Adaptability in Wide-Area Stream Processing Systems

Albert Jonathan, Abhishek Chandra, Jon Weissman

University of Minnesota - Twin Cities

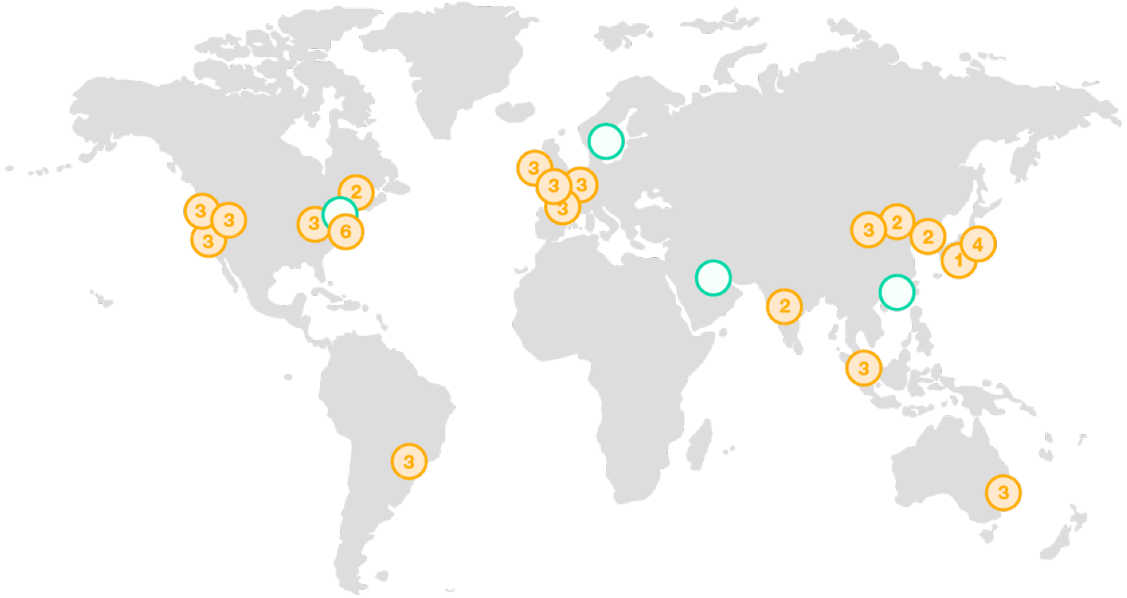


UNIVERSITY OF MINNESOTA

Wide-Area Streaming Analytics

- Lots of data are continuously "born" geo-distributed
  - CDN log updates
  - Microblogs from social network apps
  - Sensor readings from IoT

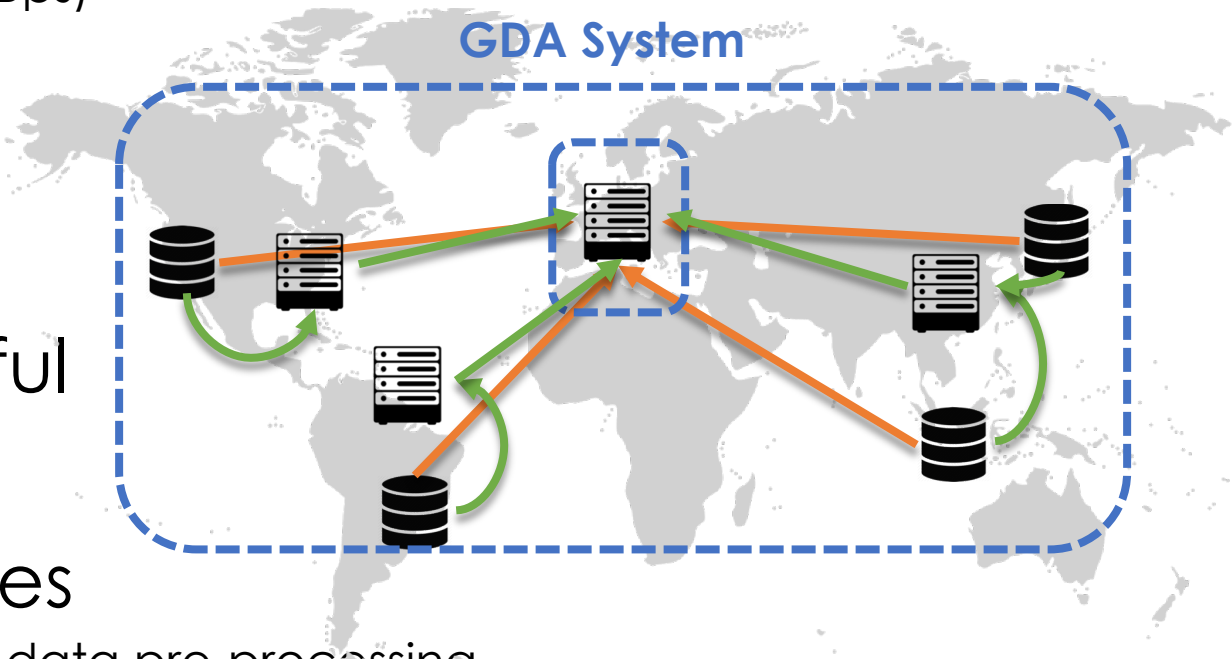
- Applications:
  - Monitoring CDN logs for Internet services
  - Global trend detection
  - Monitoring video surveillance



AWS global infrastructure map

Geo-Distributed System Model

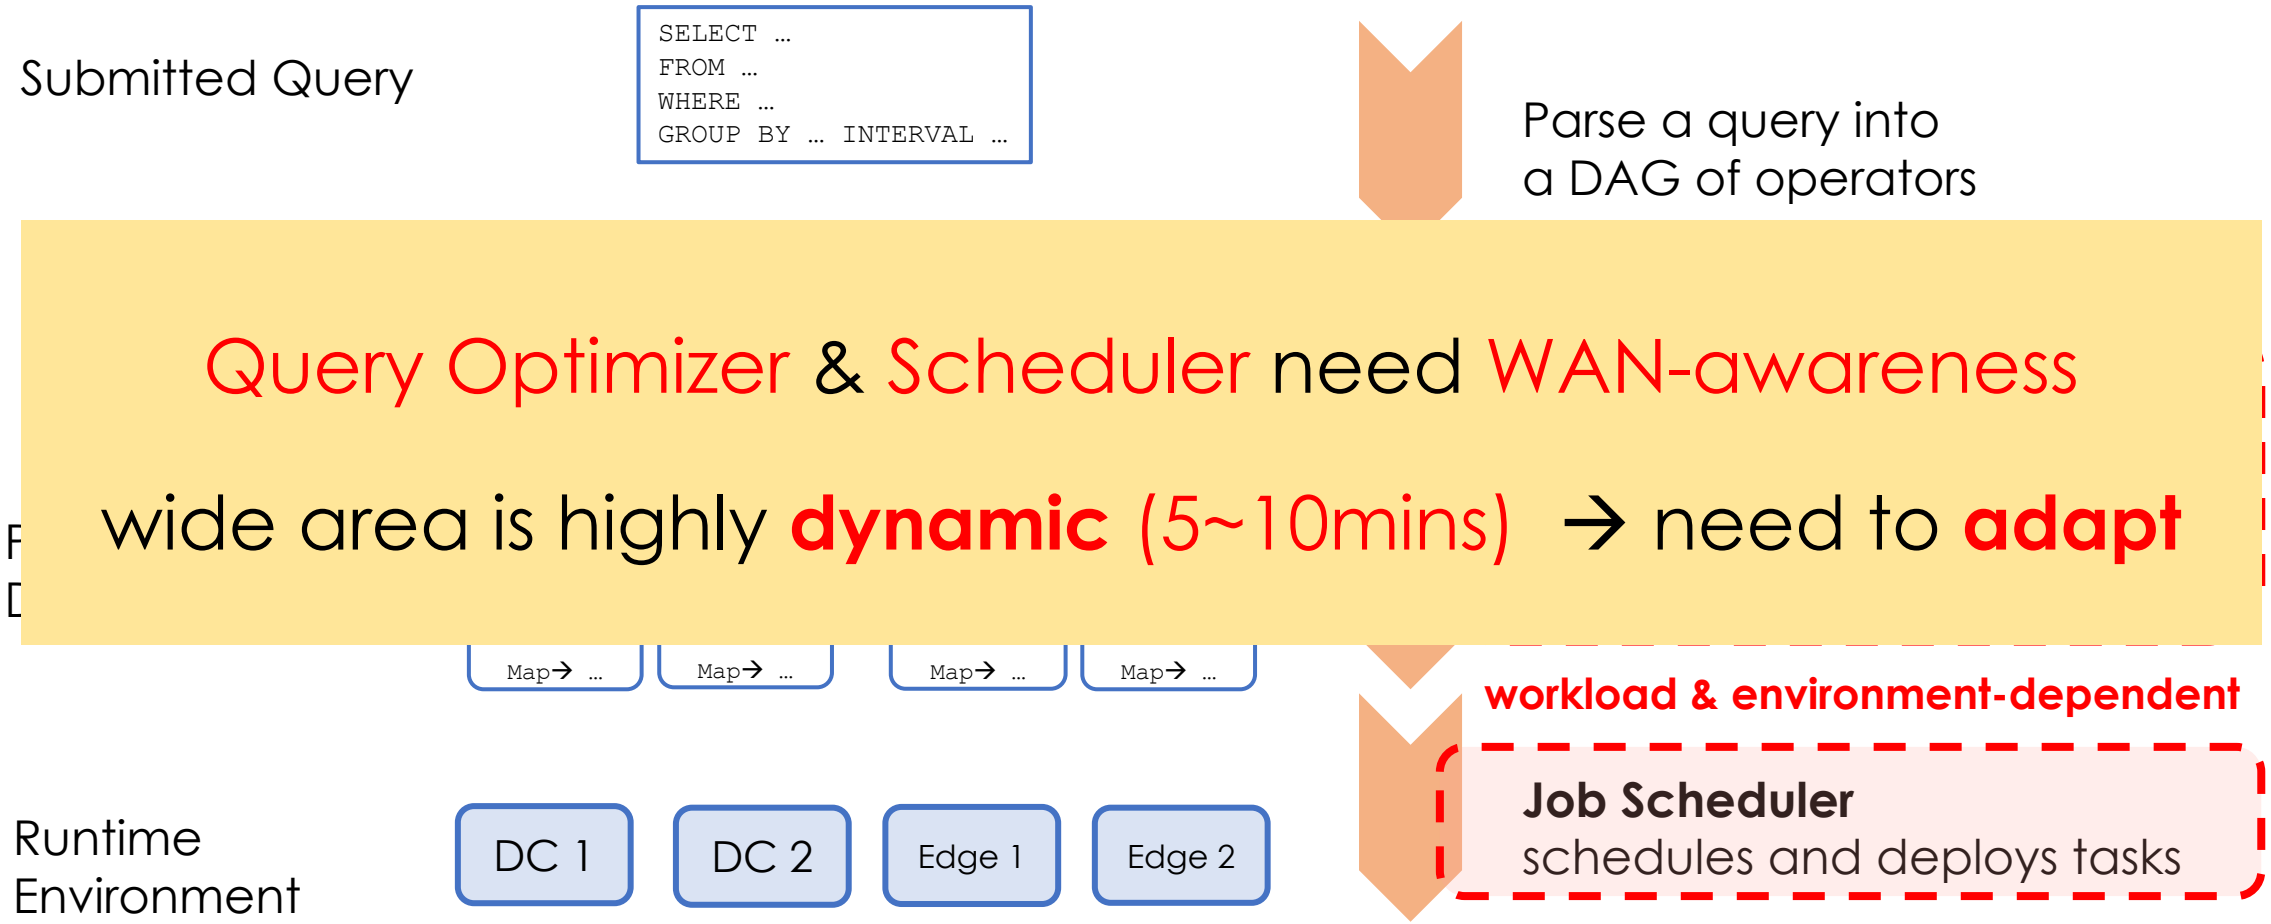
- Main challenges in wide-area data analytics
  - Limited WAN bandwidth (10s MBps)
  - Highly heterogeneous links (>20X)
  - Expensive WAN bandwidth ($$$)



- Central aggregation is wasteful

- GDA systems span multiple sites
  - High performance → locality execution & data pre-processing

Deploying Stream Queries in Wide Area

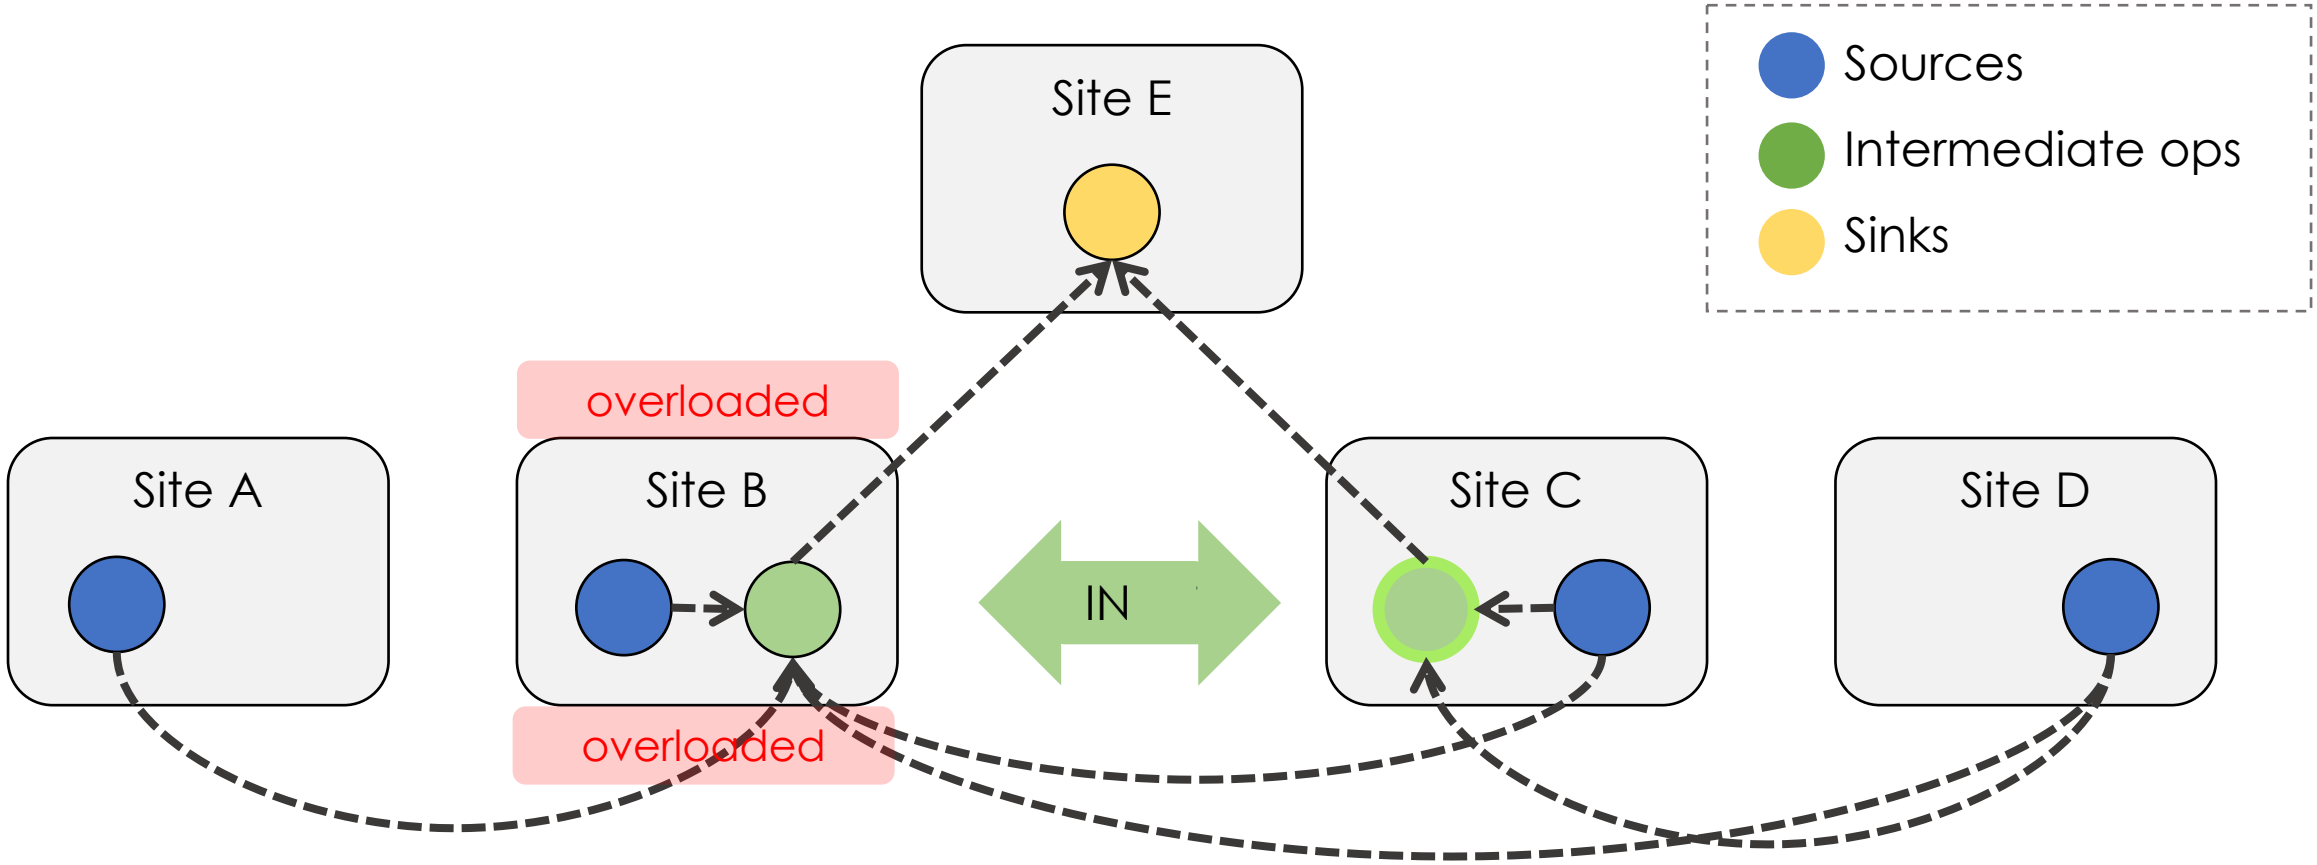


Adaptability in Wide-Area Stream Processing Systems

- Scale Up/Down
  - Increase/decrease task parallelism within a site
  - Provision more nodes to handle computing bottleneck
- Scale Out/In
  - Increase/decrease task parallelism across sites
  - Distribute workload across sites to handle WAN bottleneck
- Query Re-planning
  - Change logical query execution plan
  - Avoid WAN bandwidth contention

Scale Out/In in Wide-Area Settings

- The "optimal" physical execution depends on WAN condition

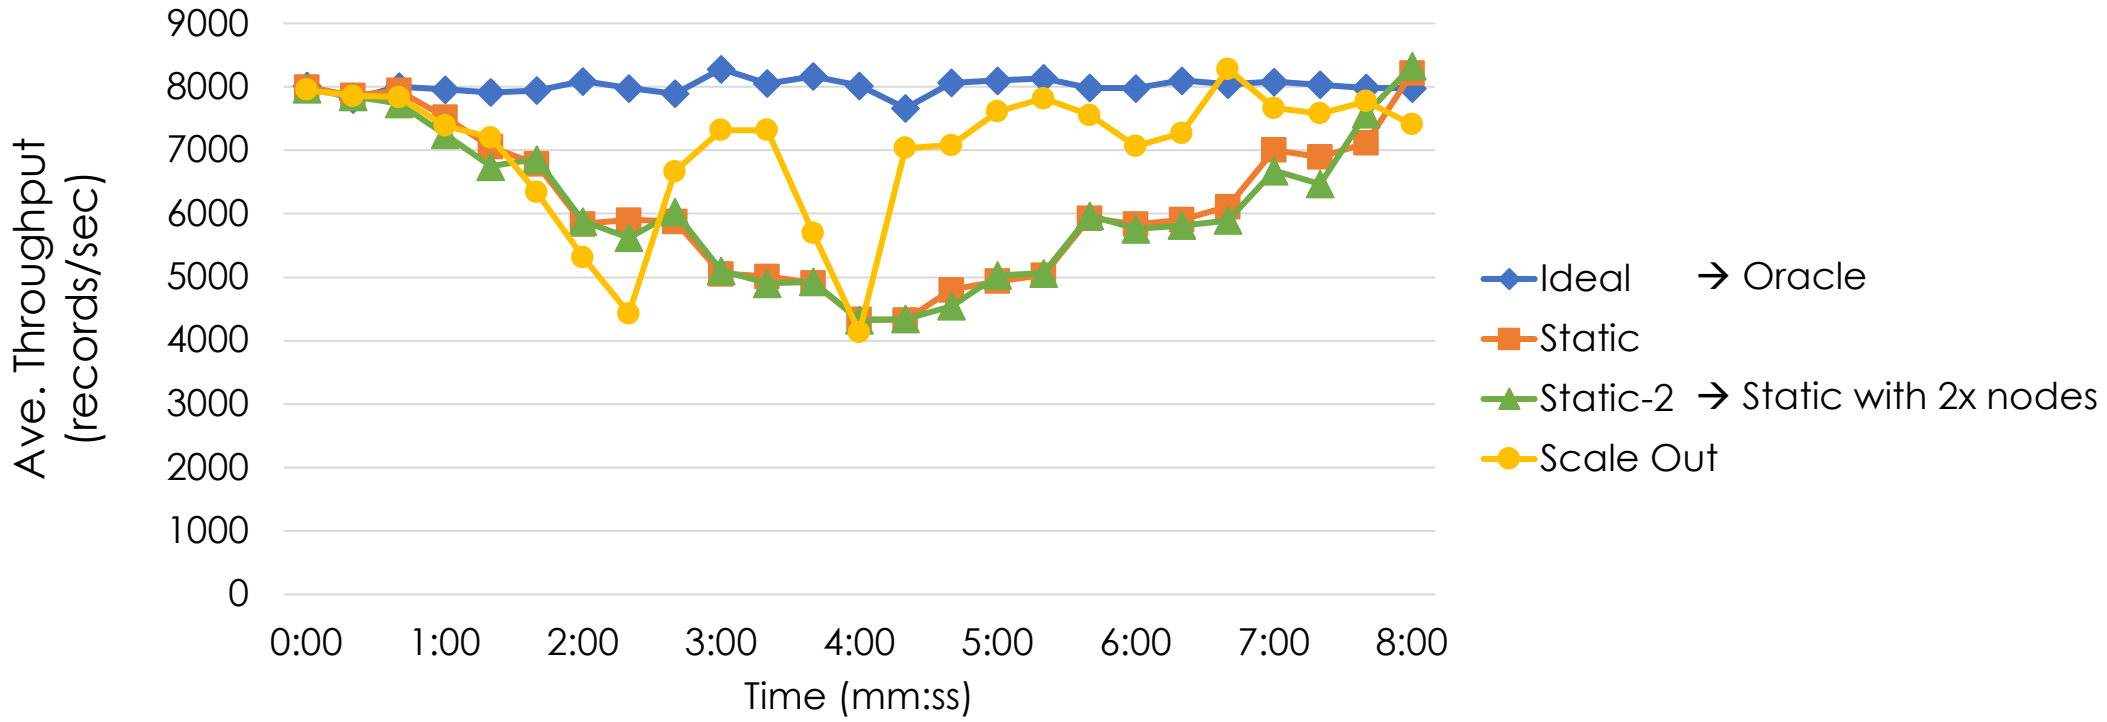


Scale Out/In in Wide-Area Settings

- Where to scale?
  - EC2: inter-site bandwidth within-region = ~5X compared to inter-region
- How to partition the streams across WAN?
  - Need to consider WAN bandwidth heterogeneity and availability
- When to scale?
  - Application specific (e.g., min. tolerable delay, throughput threshold)

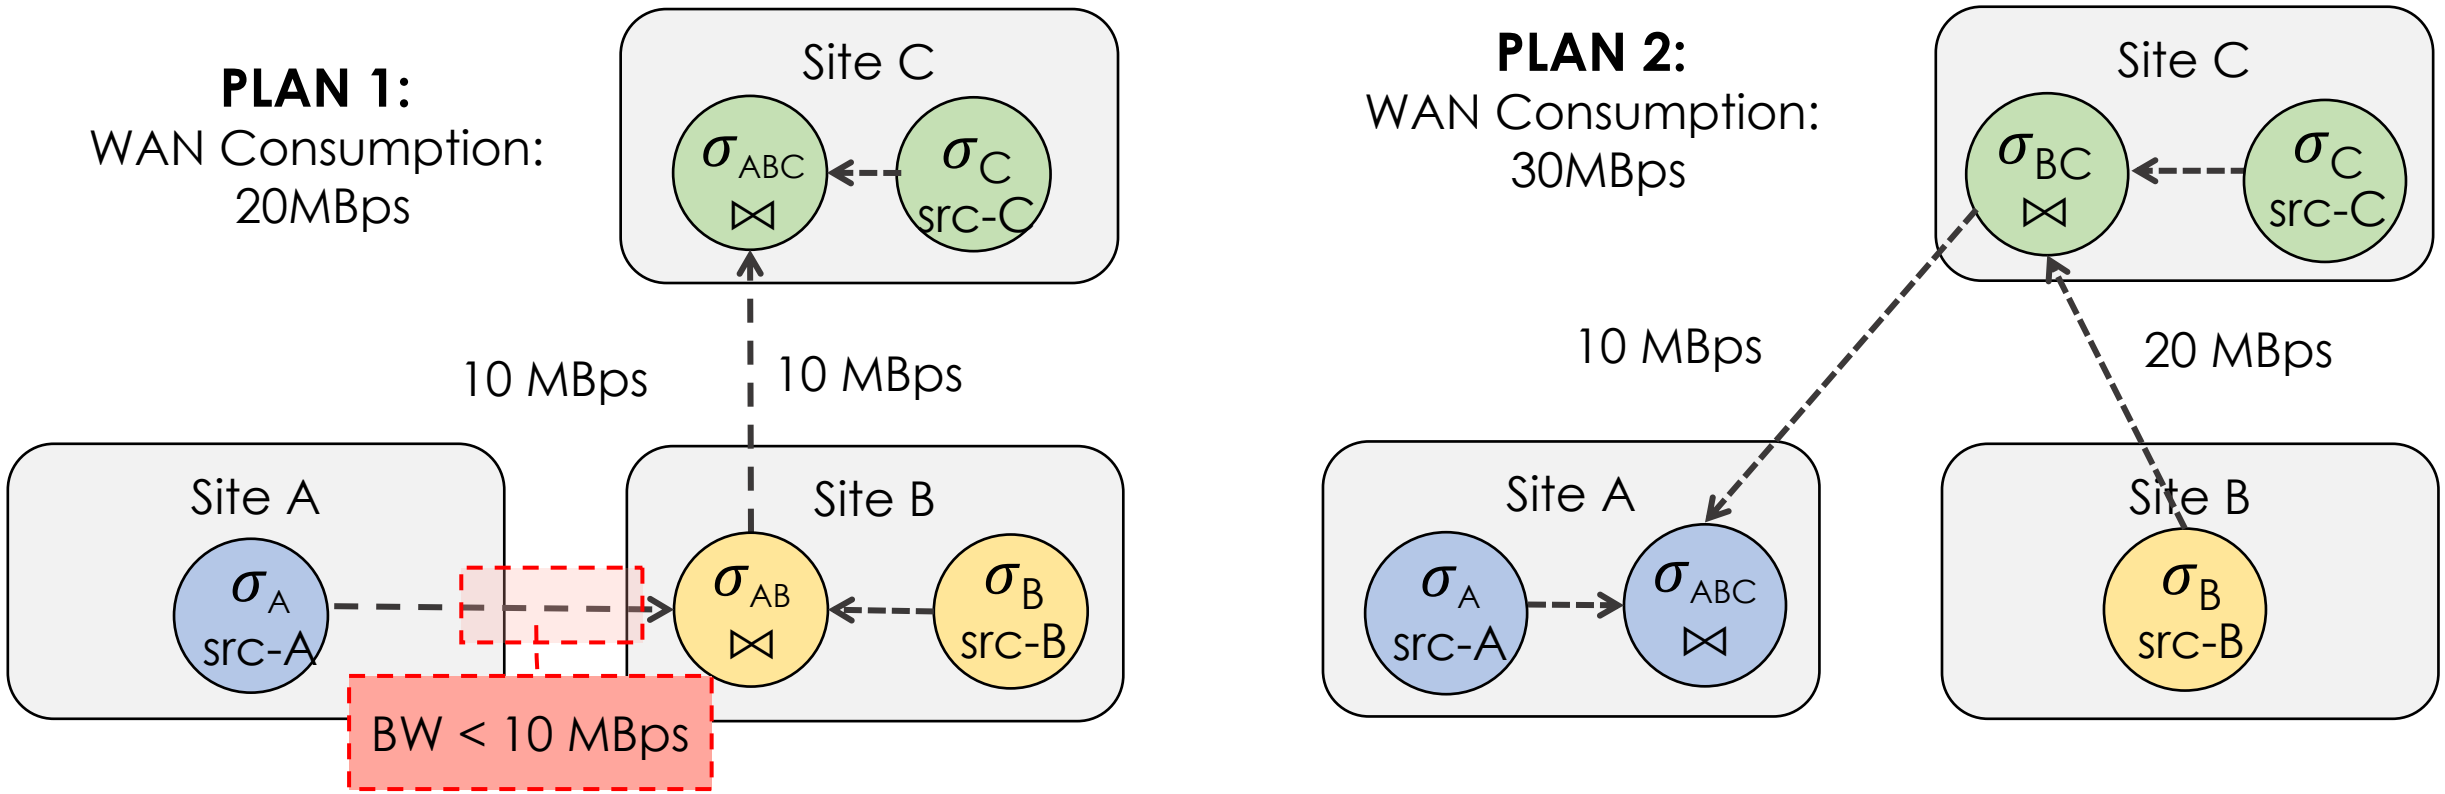
Scale Out: Preliminary Experiment

- Top-K tending topic on Twitter trace (~8000 tweets/sec) distributed based on locations
- Varying network load by 10%-30% → scale when throughput < 5000 tweets/sec



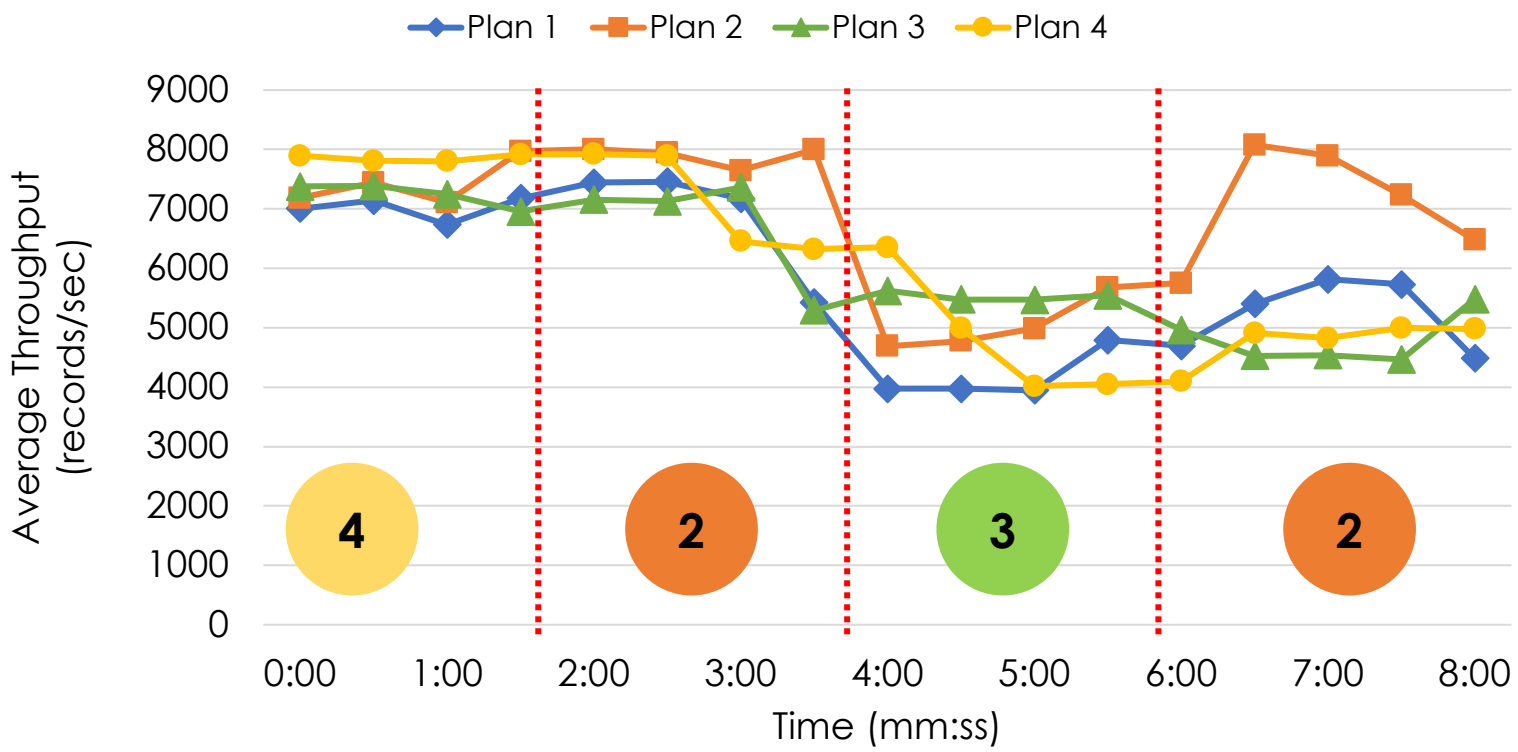
Query Re-Planning

- The "optimal" logical plan depends on WAN condition



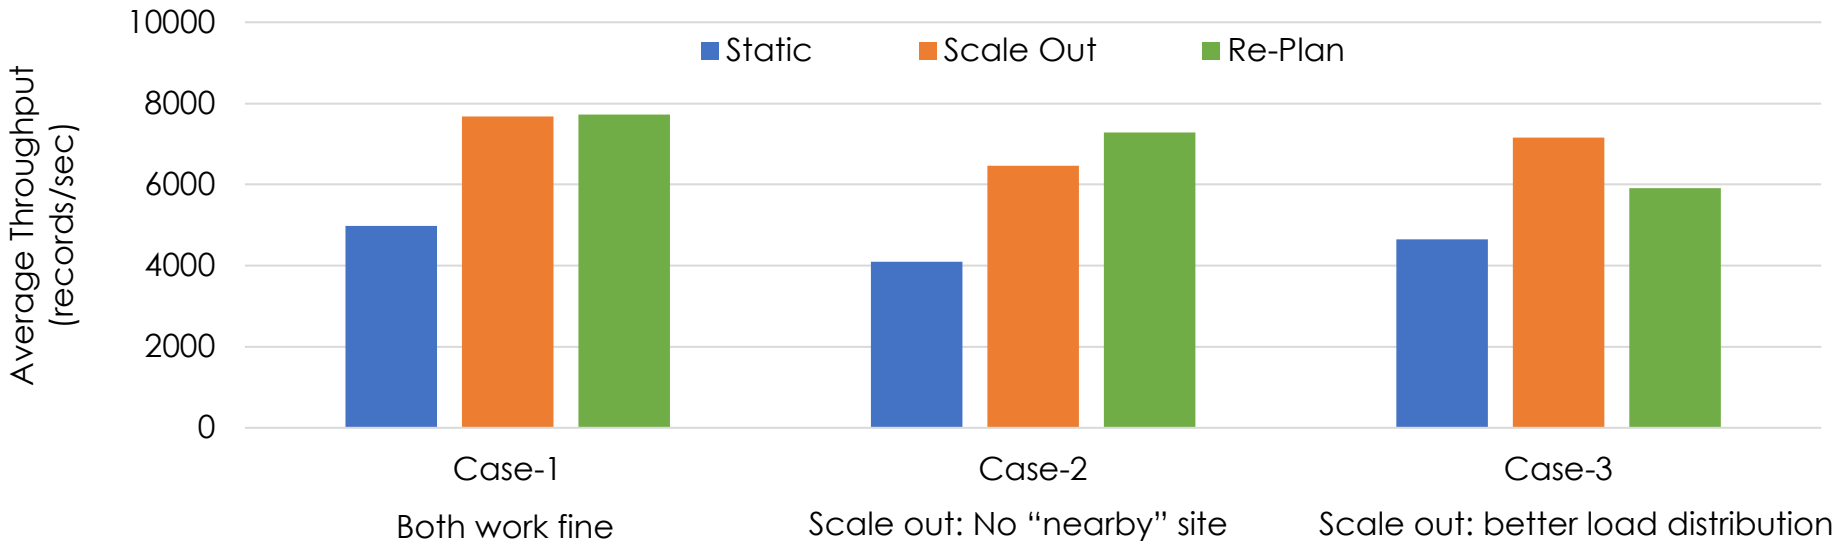
Re-Planning: Preliminary Experiment

- Same setup as the previous result
- 1 query, 4 plans: different aggregation orders



Scale Out vs. Query Re-Planning

- When to scale out/in and when to re-plan?
  - Type of operators (stateful, commutative)
  - Type of dynamics (short/long term)
  - Environment (node, bandwidth availability)



Conclusion & Next Steps

- Scale out and re-planning can be adapted to handle WAN dynamics
  - Scale out/in changes physical executions
  - Query re-planning changes logical executions
- Scale out vs. re-planning vs. approximation
- Adaptability to optimize multiple concurrent streaming analytics queries in wide-area

Thank You!