ESTD 1845

QUEEN'S
UNIVERSITY
BELFAST

QUEEN'S
UNIVERSITY
BELFAST

Graph Analytics

HANS VANDIERENDONCK

Graph Analytics and Cyber Security

- Insight from graph-structured data
- Representing interactions between entities as graphs
- Centrality metrics
- Anomaly detection
  - OddBall [PAKDD’10], Recursive structural features [KDD’11]
  - Outlier detection [KDD’10], PICS [SDM’12]
  - YAGADA [CIKM’11]
- Subgraph matching and enumeration [VLDB’12]
- Extrapolating labeled information to other vertices [*Pimplikar* et al, CIKM’14]

QUEEN’S UNIVERSITY BELFAST

High-Performance Graph Analytics

- Timeliness: high-performance analytics
- Open source tools facilitate adoption:
  - Giraphe, GraphLab/PowerGraph, Spark+GraphX
- Tools focus on scalability, not performance
  - Scale-out; small clusters; data centers
  - McSherry: Scalability! But at what COST?
    - Compare against laptop!

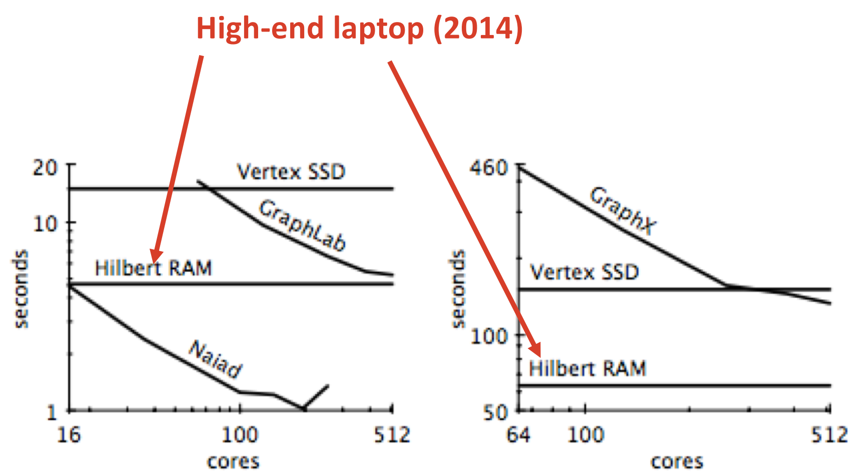

Performance of 20 iterations of PageRank twitter_rv graph – full 2010 snapshot of Twitter: 41.7M vertices, 1.47G edges
Source: McSherry USENIX HotOS '15

QUEEN'S UNIVERSITY BELFAST

The GraphGrind Framework: Fast Graph Analytics on Shared-Memory Systems

- The issue with scale-out systems
  - Workloads characterised by frequent communication and synchronisation
  - Can fall-back to processing from disk; likely competitive to distributed memory processing
- The alternative: a scaled-up server
  - Plenty of working memory possible; up to 16TB in shared memory configuration
  - Future: die stacking (e.g., hybrid memory cube), non-volatile memory (e.g., PCM, ReRAM)
  - But: must be a multi-socket (NUMA, non-uniform memory architecture) setup

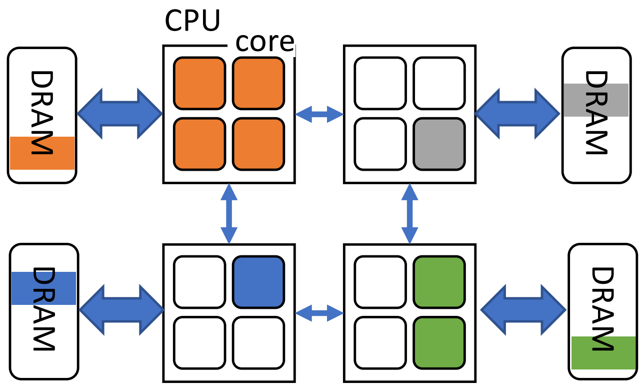

QUEEN'S UNIVERSITY BELFAST

GraphGrind

- Challenges:
  - Balance threads and data over NUMA nodes
  - Load balancing depends on graph algorithm, varies throughout computation
    - Tension between balancing #vertices vs #edges
  - Difficulty: memory locality and memory access latency
- Solution:
  - Partition graphs finely: improves memory locality, unit of work for scheduling, load balancing
  - Co-locate partitions and the main code operating on them on the same NUMA node
  - Vertex and Edge Balanced Reordering (VEBO):
    - Relabel vertices in order to have #vertices/partition AND same #edges/partition

QUEEN'S UNIVERSITY BELFAST

Memory accesses per 1000 instructions (MPKI) twitter graph

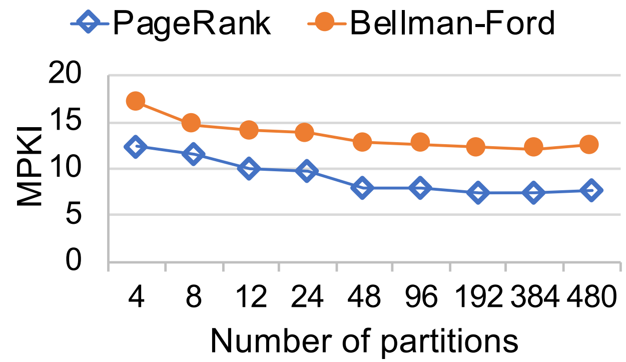

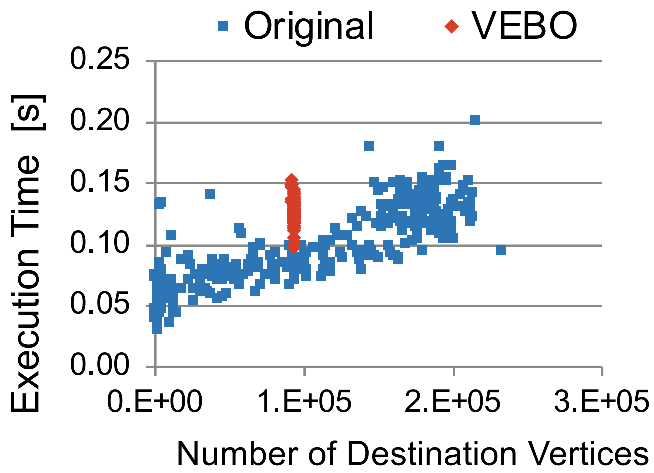

GraphGrind: Performance Comparison

Runtime on twitter graph for
BP: Belief Propagation
PR: 10 iterations of PageRank
PRDelta: approximating PageRank

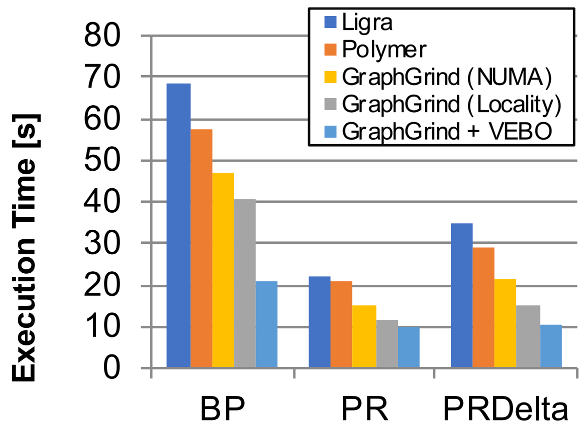

Time per iteration for PageRank on UK-Union 2007 snapshot of internet pages in the .uk domain

Framework	#cores	Time/PageRank iteration
Giraphe (Apache)	128	62s
GraphLab (Apache)	128	42s
GraphX (Apache)	128	23s
Ligra (CMU/MIT)	48*	3.8s
GraphGrind	48*	1.2s

*4-socket Xeon E7-4860 v2

QUEEN'S UNIVERSITY BELFAST

Questions

- The right hardware for the job. What is it?
  - Storage class memory
  - Scale-out or scale-up?
- What are the key challenges in processing industrial graph datasets?
- What workloads are speed-sensitive, which aren’t?

QUEEN'S UNIVERSITY BELFAST