


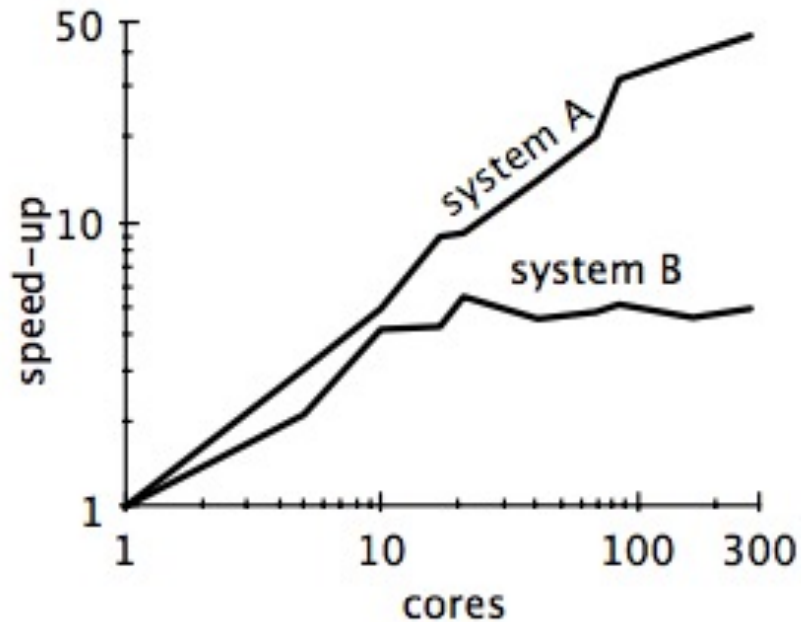
Scability! But at what COST?

Frank McSherry, Michael Isard, and Derek G. Murray

HotOS 2015

Presented by Julian Shun

Scalability



- System A scales much better despite being slower
- Many published big data systems resemble System A

COST

- COST: Configuration that Outperforms a Single Thread
- A system has unbounded COST if no configuration outperforms a single thread

PageRank (20 iterations)

scalable system	cores	twitter	uk-2007-05
GraphChi [12]	2	3160s	6972s
Stratosphere [8]	16	2250s	-
X-Stream [21]	16	1488s	-
Spark [10]	128	857s	1759s
Giraph [10]	128	596s	1235s
GraphLab [10]	128	249s	833s
GraphX [10]	128	419s	462s

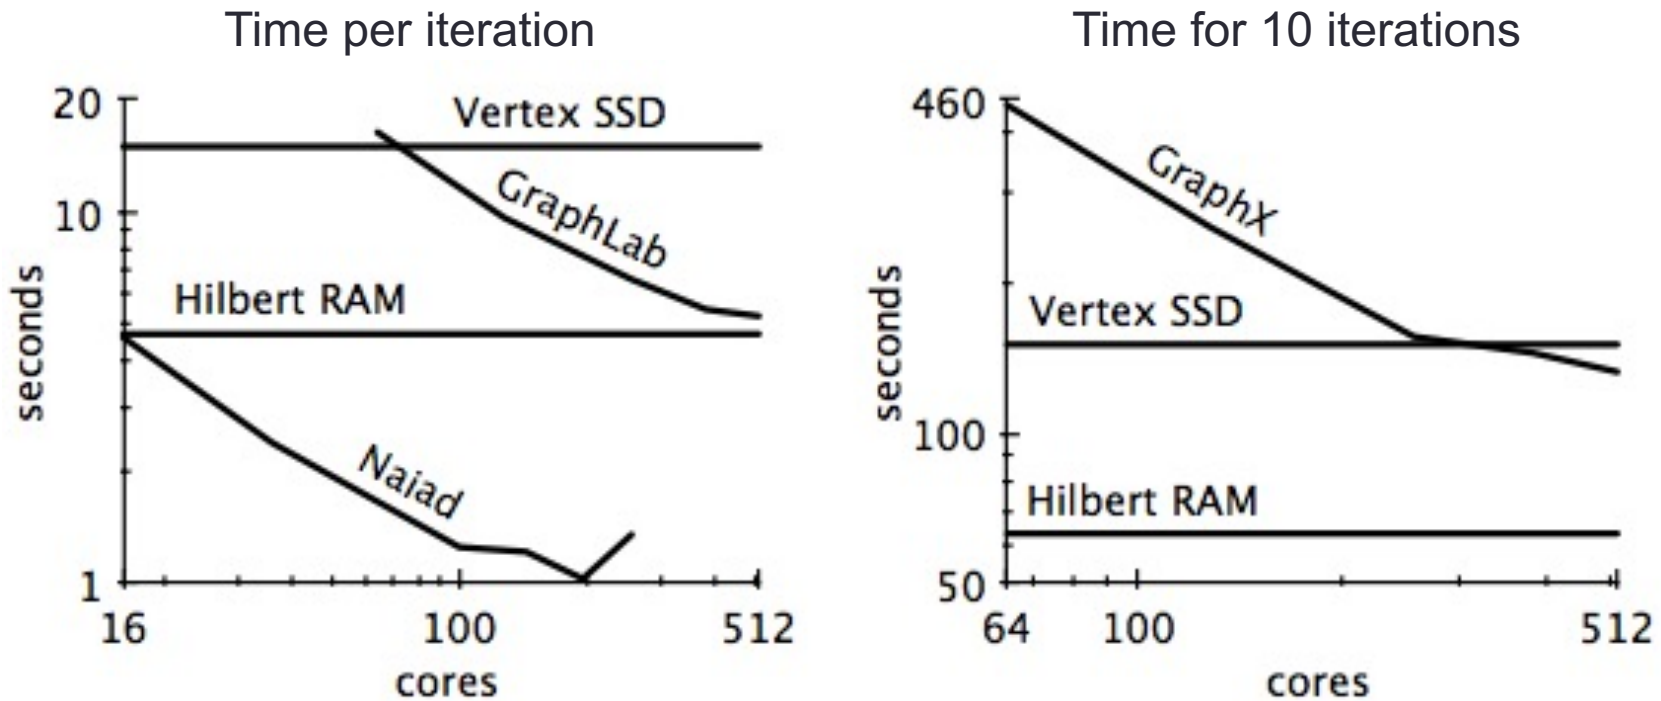
Hilbert order (SSD)	1	242s	256s
Hilbert order (RAM)	1	110s	-

Connected Components

scalable system	cores	twitter	uk-2007-05
Stratosphere [8]	16	950s	-
X-Stream [21]	16	1159s	-
Spark [10]	128	1784s	$\geq 8000s$
Giraph [10]	128	200s	$\geq 8000s$
GraphLab [10]	128	242s	714s
GraphX [10]	128	251s	800s

Union-Find (SSD)	1	15s	30s
------------------	---	-----	-----

PageRank Scaling



- Naiad has a COST of 16 cores
- GraphLab has a COST of 512 cores
- GraphX has unbounded COST

Conclusions

- Always compare to a good single-threaded baseline
- Scalability is not the only important performance factor
- *“You can have a second computer once you’ve shown you know how to use the first one.”*
- Paul Barham