Scalability!
But at what COST?

Frank McSherry Michael Isard Derek G. Murray

Scalable Systems

Graph processing

This is an intervention

20xPR	cores	twitter_rv	uk_2007_05
Spark	128	857s	1759s
Giraph	128	596s	1235s
GraphLab	128	249s	833s
GraphX	128	419s	462s

from Gonzalez *et al*., OSDI 2014

```
fn pagerank<G: Graph>(graph: &G, nodes: usize, alpha: f32)
{
    let mut src = vec![0f32; nodes];
    let mut dst = vec![0f32; nodes];
    let mut deg = vec![0f32; nodes];

    graph.map_edges(|x, _| { deg[x] += 1f32 });

    for _iteration in (0 .. 20) {
        println!("Iteration: {}", _iteration);
        for node in (0 .. nodes) {
            src[node] = alpha * dst[node] / deg[node];
            dst[node] = 1f32 - alpha;
        }

        graph.map_edges(|x, y| { dst[y] += src[x]; });
    }
}
```

```rust
fn pagerank<G: Graph>(graph: &G, nodes: usize, alpha: f32)
{
    let mut src = vec![0f32; nodes];
    let mut dst = vec![0f32; nodes];
    let mut deg = vec![0f32; nodes];

    graph.map_edges(|x, _| { deg[x] += 1f32 });

    for _iteration in (0 .. 20) {
        println!("Iteration: {}", _iteration);
        for node in (0 .. nodes) {
            src[node] = alpha * dst[node] / deg[node];
            dst[node] = 1f32 - alpha;
        }

        graph.map_edges(|x, y| { dst[y] += src[x]; });
    }
}
```

```
fn pagerank<G: Graph>(graph: &G, nodes: usize, alpha: f32)
{
    let mut src = vec![0f32; nodes];
    let mut dst = vec![0f32; nodes];
    let mut deg = vec![0f32; nodes];

    graph.map_edges(|x, _| { deg[x] += 1f32 });

    for _iteration in (0 .. 20) {
        println!("Iteration: {}", _iteration);
        for node in (0 .. nodes) {
            src[node] = alpha * dst[node] / deg[node];
            dst[node] = 1f32 - alpha;
        }

        graph.map_edges(|x, y| { dst[y] += src[x]; });
    }
}
```

```rust
fn pagerank<G: Graph>(graph: &G, nodes: usize, alpha: f32)
{
    let mut src = vec![0f32; nodes];
    let mut dst = vec![0f32; nodes];
    let mut deg = vec![0f32; nodes];

    graph.map_edges(|x, _| { deg[x] += 1f32 });

    for _iteration in (0 .. 20) {
        println!("Iteration: {}", _iteration);
        for node in (0 .. nodes) {
            src[node] = alpha * dst[node] / deg[node];
            dst[node] = 1f32 - alpha;
        }

        graph.map_edges(|x, y| { dst[y] += src[x]; });
    }
}
```

20xPR	cores	twitter_rv	uk_2007_05
Spark	128	857s	1759s
Giraph	128	596s	1235s
GraphLab	128	249s	833s
GraphX	128	419s	462s
Laptop	1		

20xPR	cores	twitter_rv	uk_2007_05
Spark	128	857s	1759s
Giraph	128	596s	1235s
GraphLab	128	249s	833s
GraphX	128	419s	462s
Laptop	1	300s	651s

20xPR	cores	twitter_rv	uk_2007_05
Spark	128	857s	1759s
Giraph	128	596s	1235s
GraphLab	128	249s	833s
GraphX	128	419s	462s
Laptop	1	~~300s~~ 110s	~~651s~~ 256s

Connectivity	cores	twitter_rv	uk_2007_05
Spark	128	1784s	8000s+
Giraph	128	200s	8000s+
GraphLab	128	242s	714s
GraphX	128	251s	800s

from Gonzalez *et al*., OSDI 2014

Connectivity	cores	twitter_rv	uk_2007_05
Spark	128	1784s	8000s+
Giraph	128	200s	8000s+
GraphLab	128	242s	714s
GraphX	128	251s	800s
Laptop	1		

Connectivity	cores	twitter_rv	uk_2007_05
Spark	128	1784s	8000s+
Giraph	128	200s	8000s+
GraphLab	128	242s	714s
GraphX	128	251s	800s
Laptop	1	153s	417s

Connectivity	cores	twitter_rv	uk_2007_05
Spark	128	1784s	8000s+
Giraph	128	200s	8000s+
GraphLab	128	242s	714s
GraphX	128	251s	800s
Laptop	1	~~153s~~ 15s	~~417s~~ 30s

How did this happen?

mistaking scalability for performance

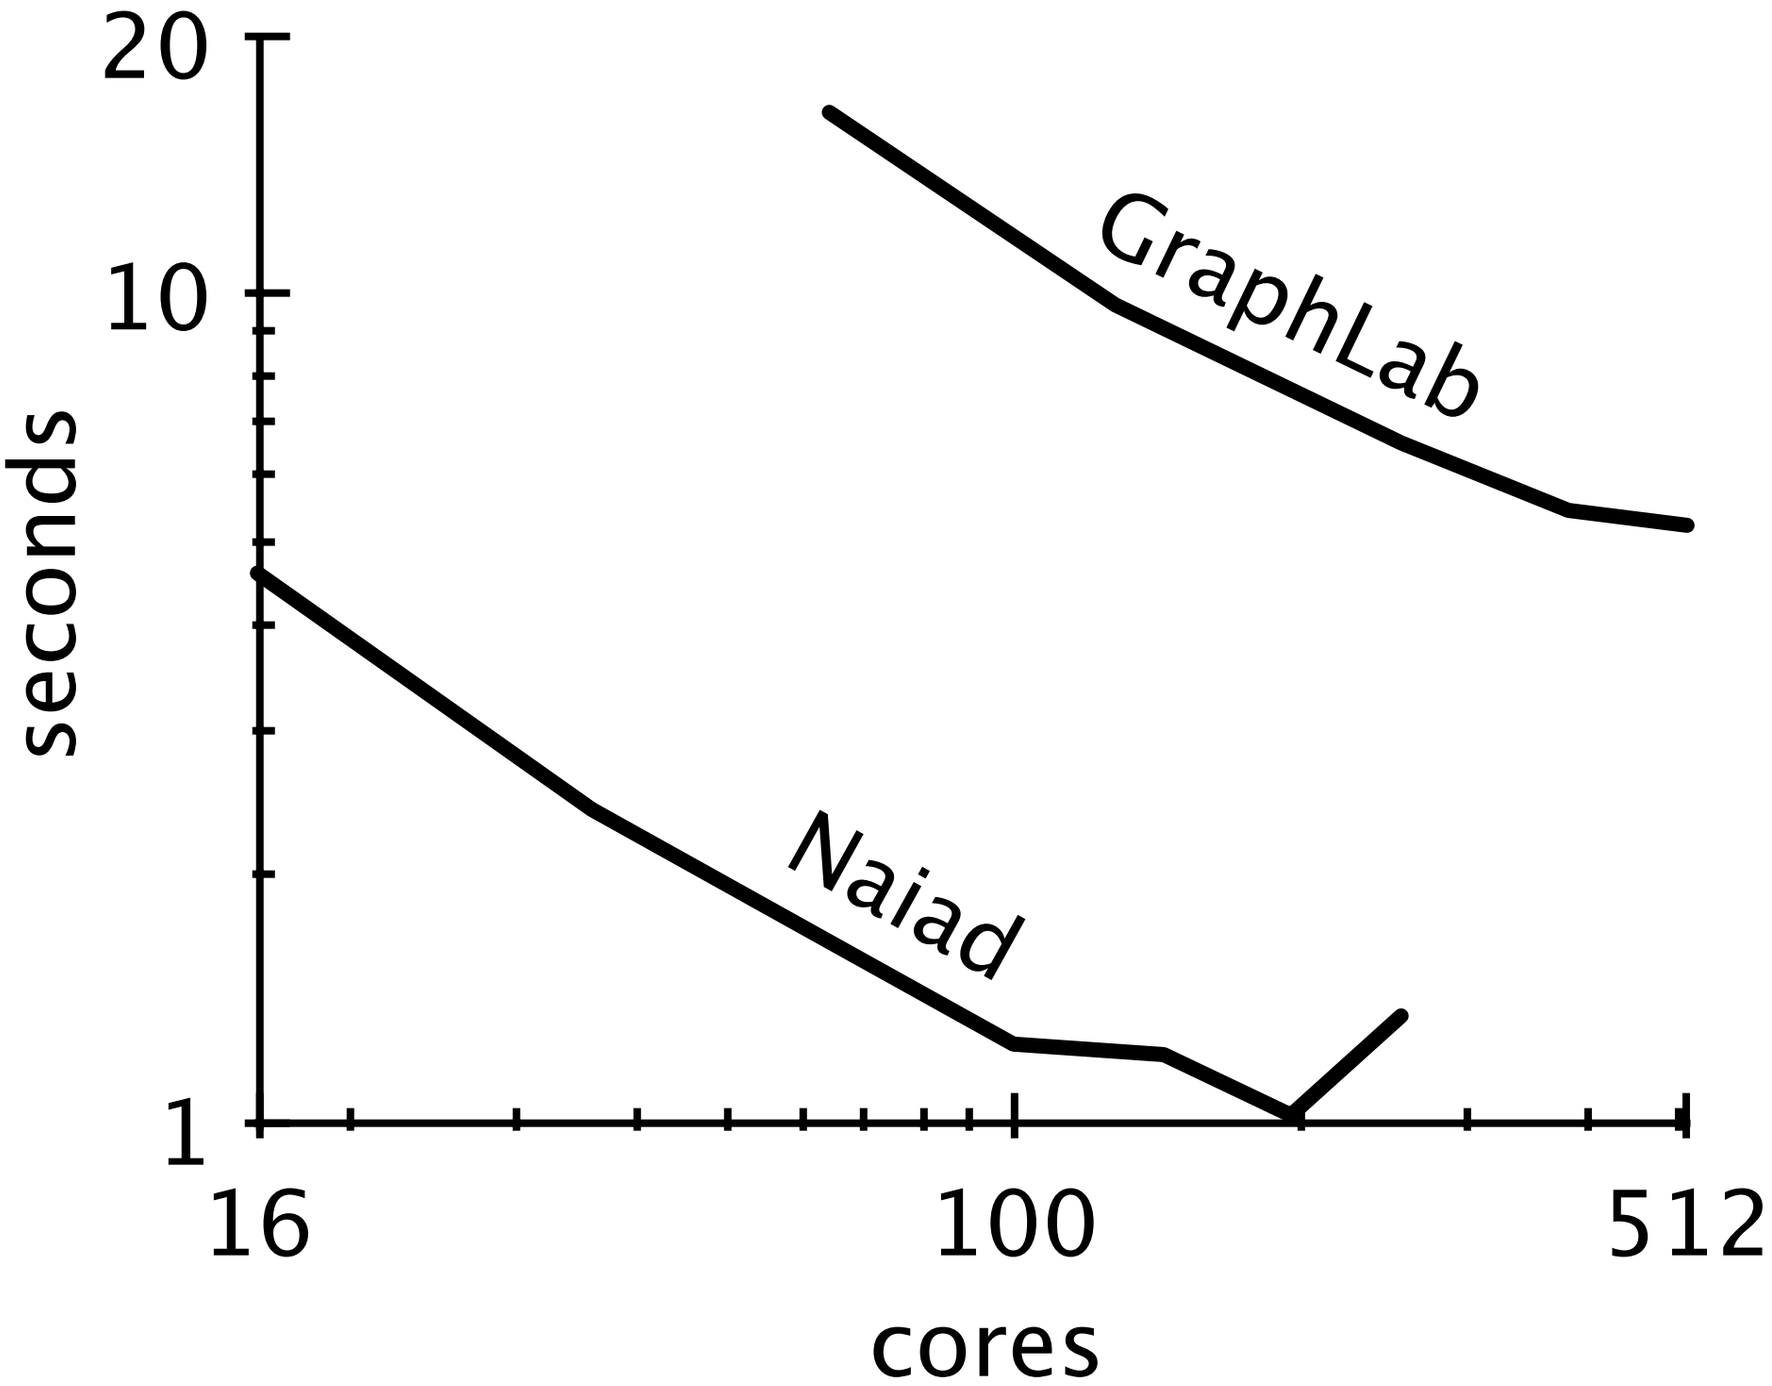
20
10
1
seconds
16
100
512
cores
GraphLab
Naiad

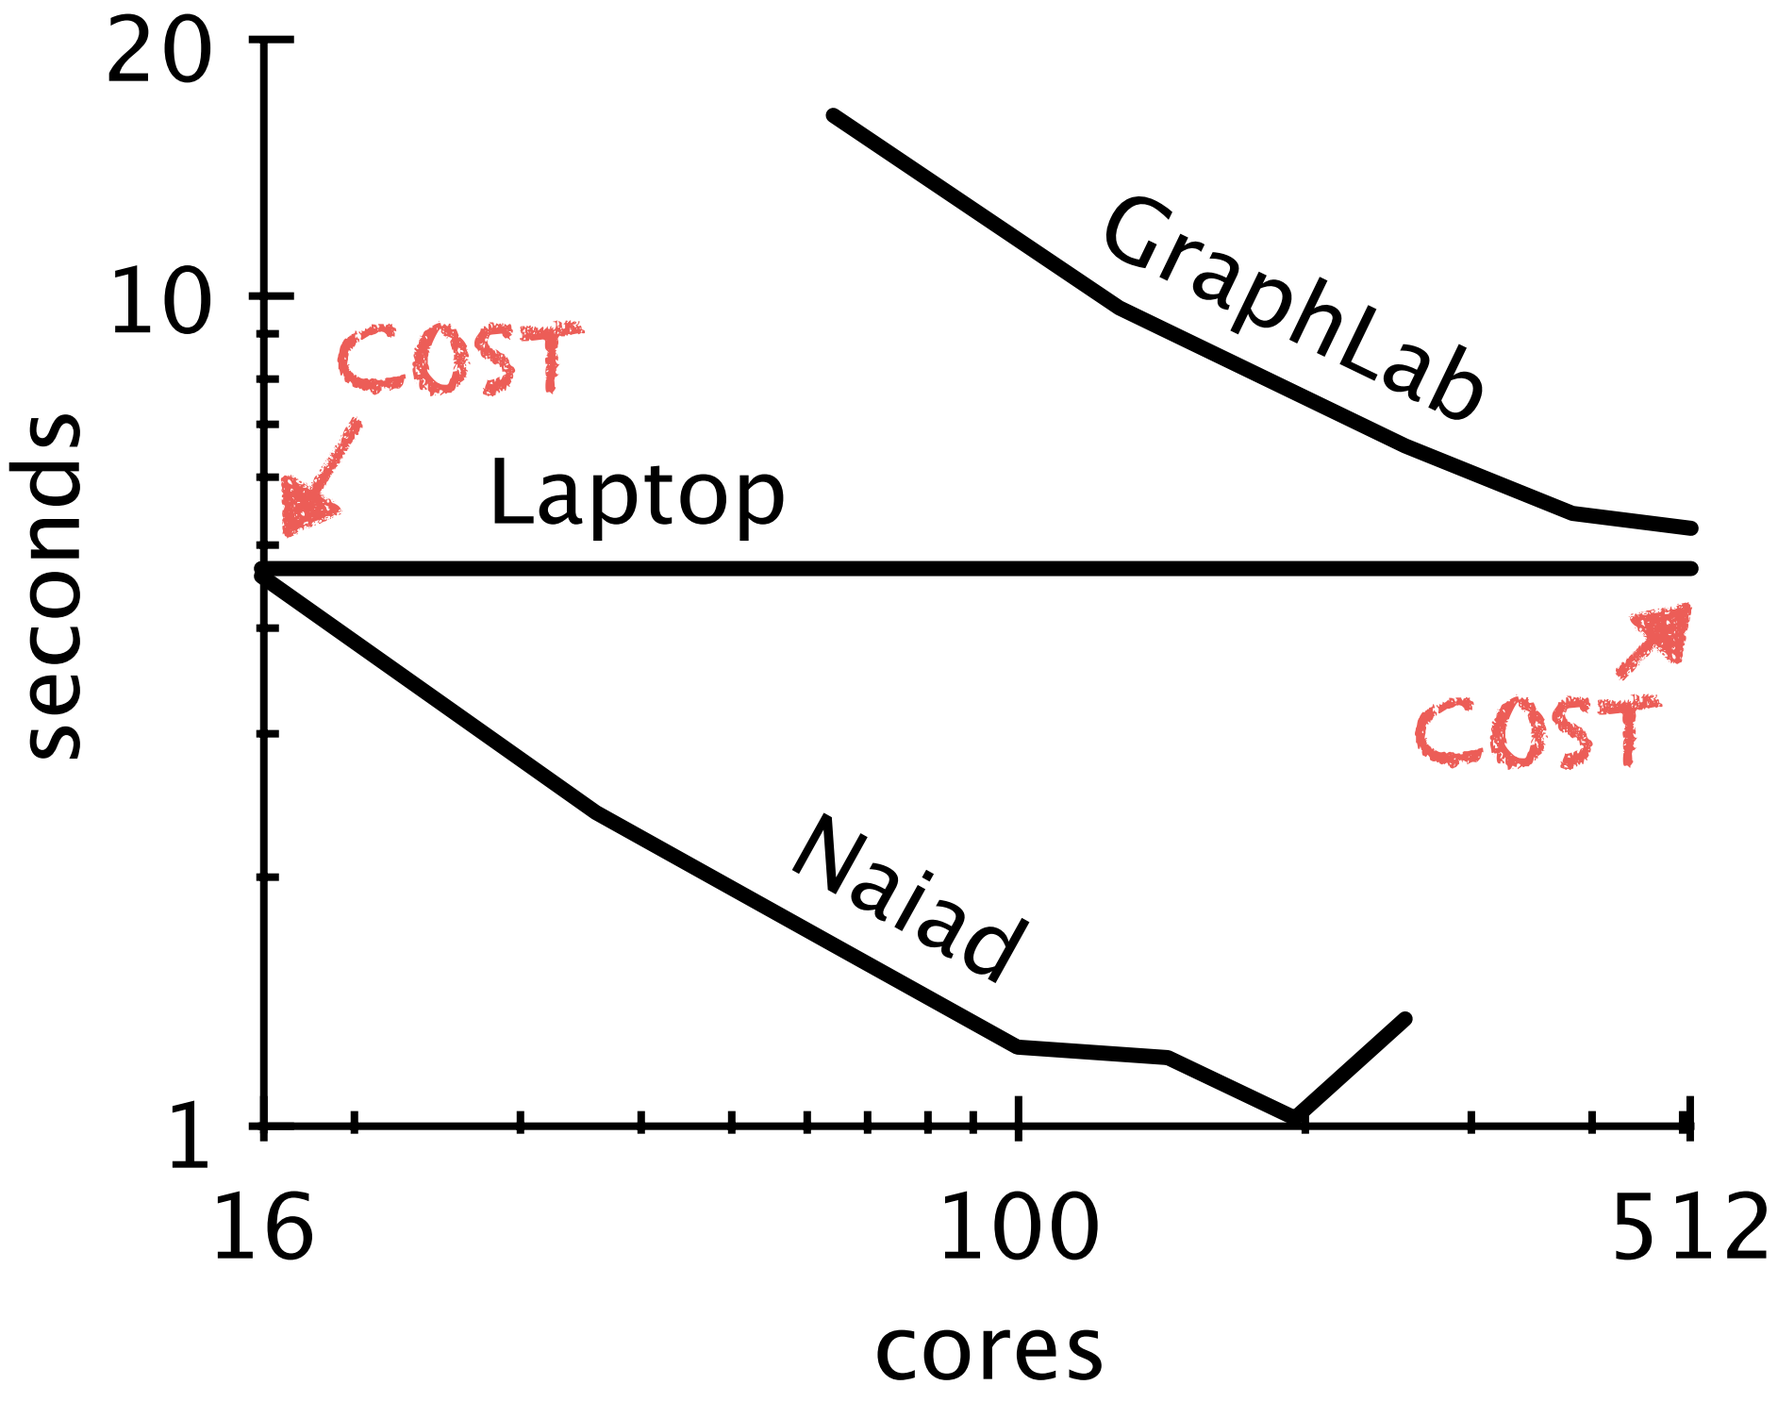
20
10
1
seconds
COST
Laptop
GraphLab
Naiad
COST
16
100
512
cores

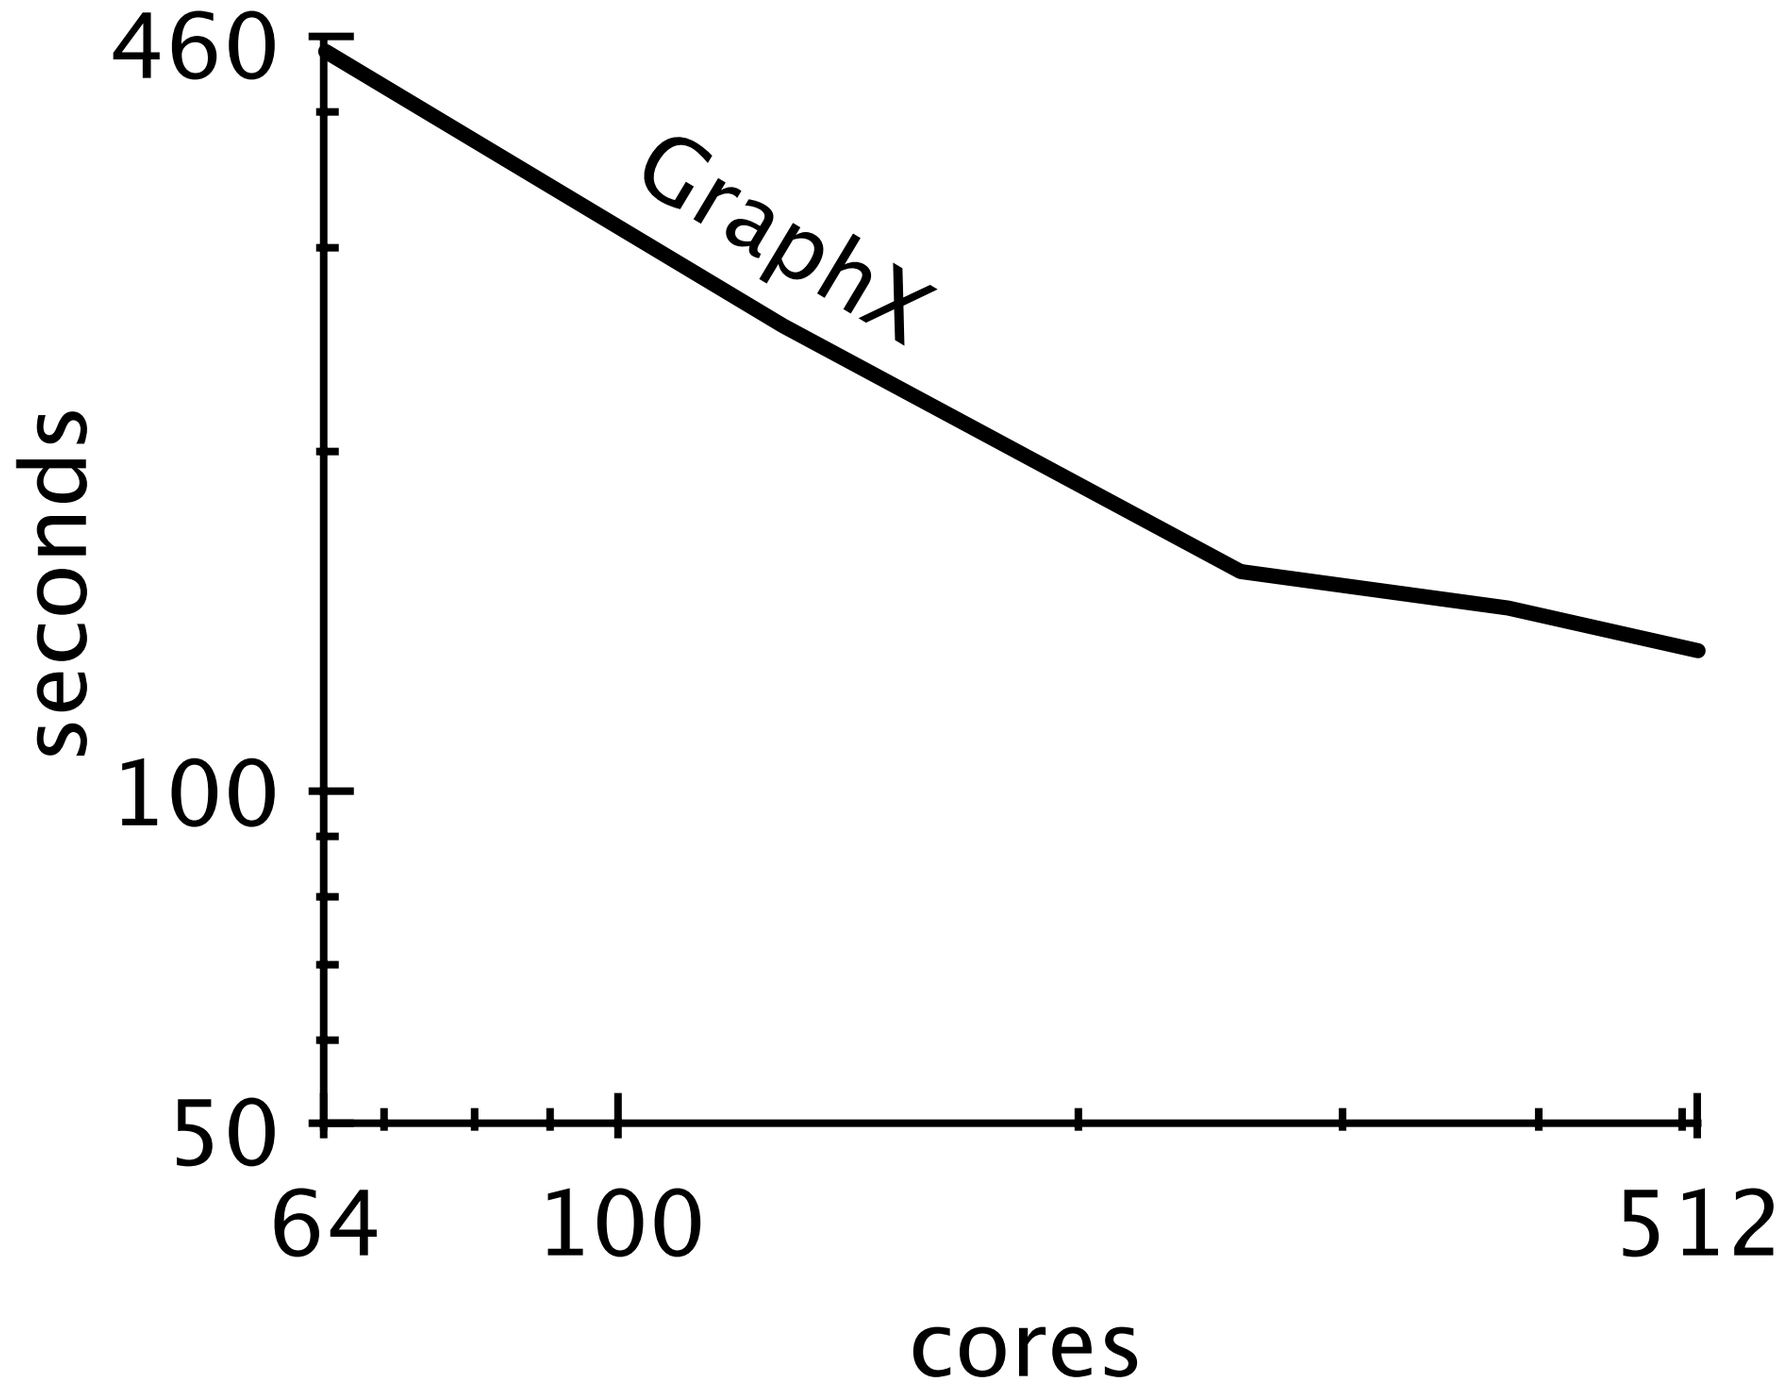
460
100
50
seconds
GraphX
64
100
512
cores

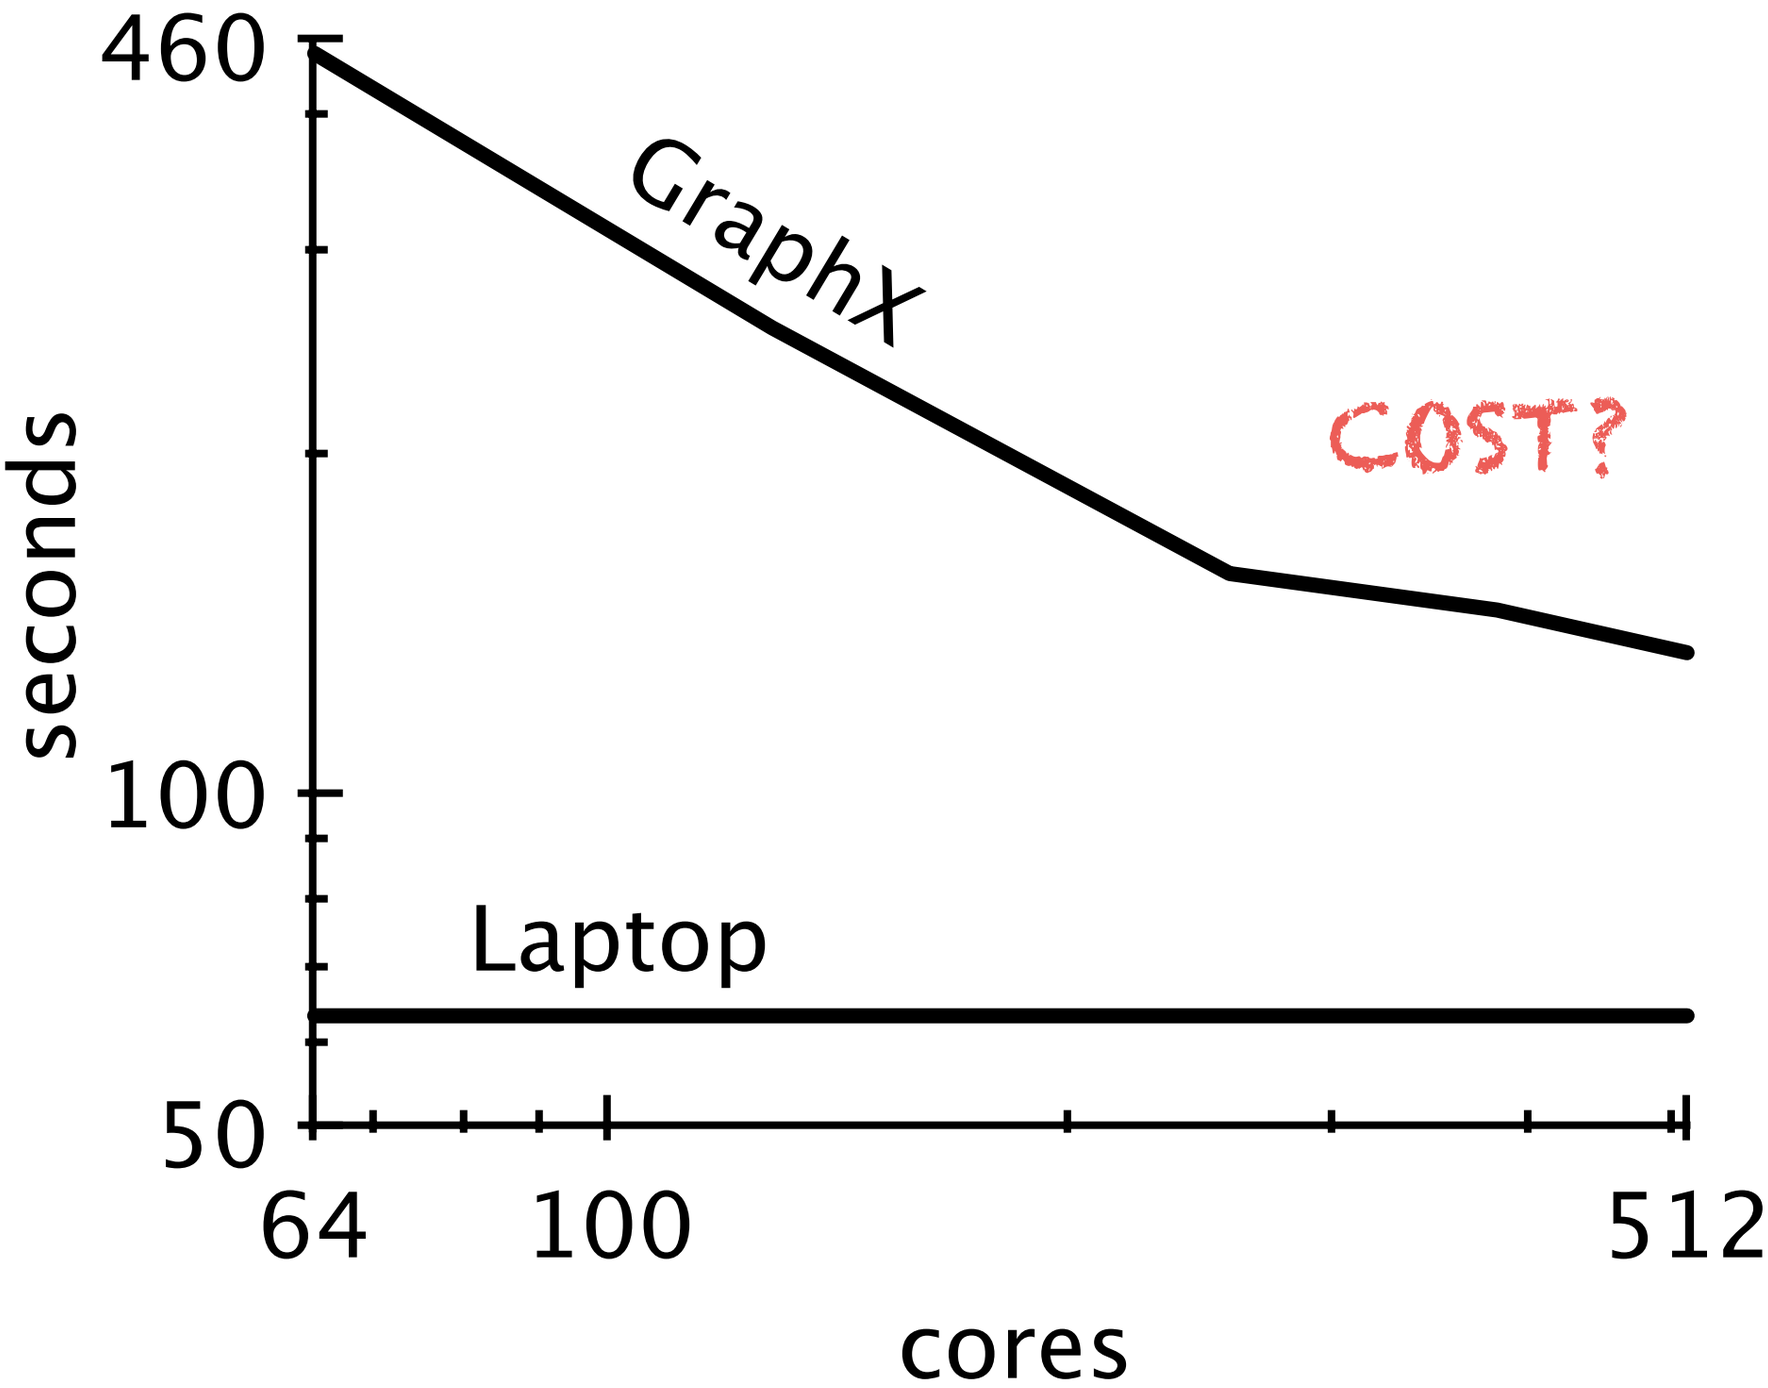
460
100
50
seconds
GraphX
COST?
Laptop
64
100
512
cores

How can we help?

make demands

Demand Baselines

Ligra	Shared-memory
FlashGraph	Solid-state drives
Naiad	Distributed systems

Ask to see the cost

iPhone 6 Plus

Available in silver, gold, and space gray, iPhone 6 Plus features an A8 chip, Touch ID, faster LTE wireless, a new 8MP iSight camera with Focus Pixels, and iOS 8.

Buy now ›

iPhone 6

Available in silver, gold, and space gray, iPhone 6 features an A8 chip, Touch ID, faster LTE wireless, a new 8MP iSight camera with Focus Pixels, and iOS 8.

Buy now ›

iPhone 5s

Available in silver, gold, and space gray, iPhone 5s features an A7 chip, Touch ID, LTE wireless, an 8MP iSight camera, and iOS 8.

Buy now ›

iPhone 5c

Available in green, blue, yellow, pink, and white, iPhone 5c features an A6 chip, LTE wireless, an 8MP iSight camera, and iOS 8.

Buy now ›

http://www.commoncrawl.org

128B edges, 3.6B nodes

http://www.commoncrawl.org

128B edges, 3.6B nodes

20xPR	cores	twitter_rv	uk_2007_05	common crawl
Laptop	1	110s	256s	46,600s
Connectivity	**cores**	**twitter_rv**	**uk_2007_05**	**common crawl**
Laptop	1	15s	30s	1,700s

Demand Research

https://github.com/frankmcsherry/COST
/timely-dataflow
/differential-dataflow
/dataflow-join

“You can have a second computer once you’ve shown you know how to use the first one.”

–Paul Barham