Opportunities and Challenges in Scaling Up Graph Analytics

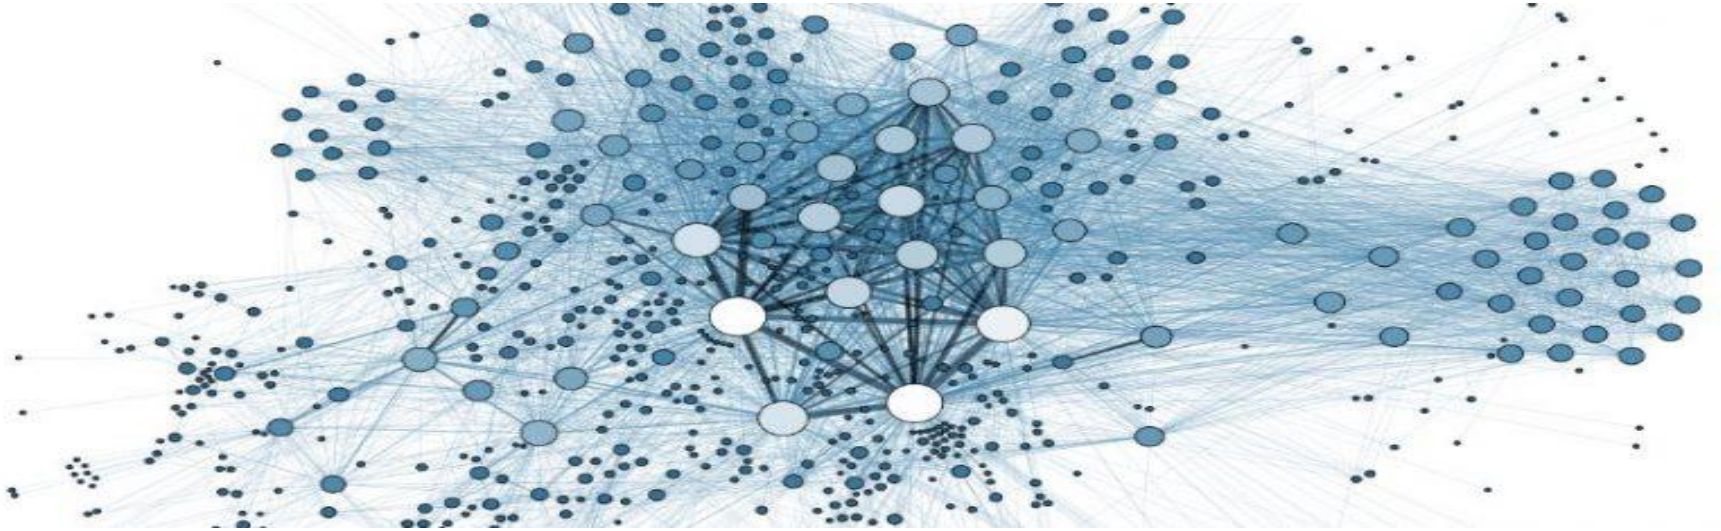

Priyank Faldu and **Boris Grot**

The University of Edinburgh

* This work is partially supported by **Oracle Labs**

Why Graphs?

Graphs are a natural way to represent pair-wise relationships among objects in the real world

- Each object is a vertex
- Relationship among a pair of objects represented with an edge

Why Graphs?

Graphs are a natural way to represent pair-wise relationships among objects in the real world

- Each object is a vertex
- Relationship among a pair of objects represented with an edge

Example:

Facebook Social Network

Why Graphs?

Graphs are a natural way to represent pair-wise relationships among objects in the real world

- Each object is a vertex
- Relationship among a pair of objects represented with an edge

Example:

Facebook Social Network

UoE

Boris Grot

Priyank Faldu

Why Graphs?

Graphs are a natural way to represent pair-wise relationships among objects in the real world

- Each object is a vertex
- Relationship among a pair of objects represented with an edge

Example:

Facebook Social Network

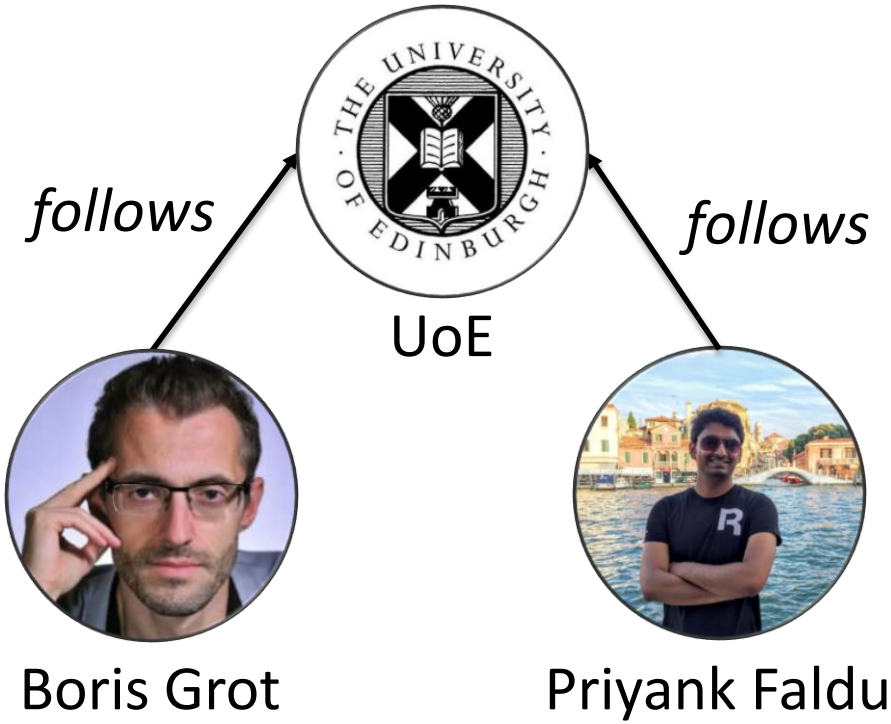

Why Graphs?

Graphs are a natural way to represent pair-wise relationships among objects in the real world

- Each object is a vertex
- Relationship among a pair of objects represented with an edge

Example:

Facebook Social Network

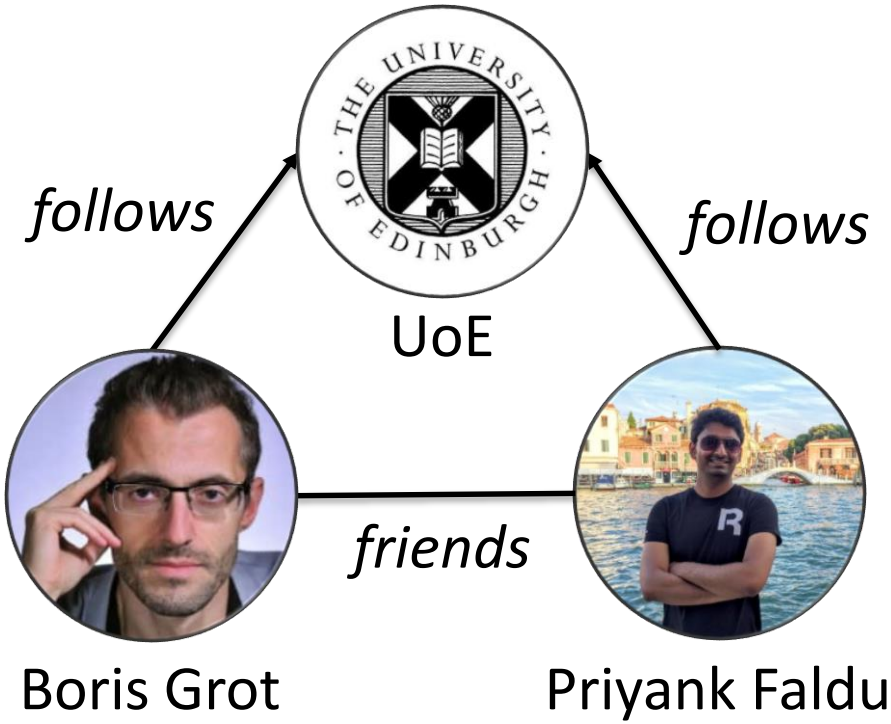

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

- Applications

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

- Applications

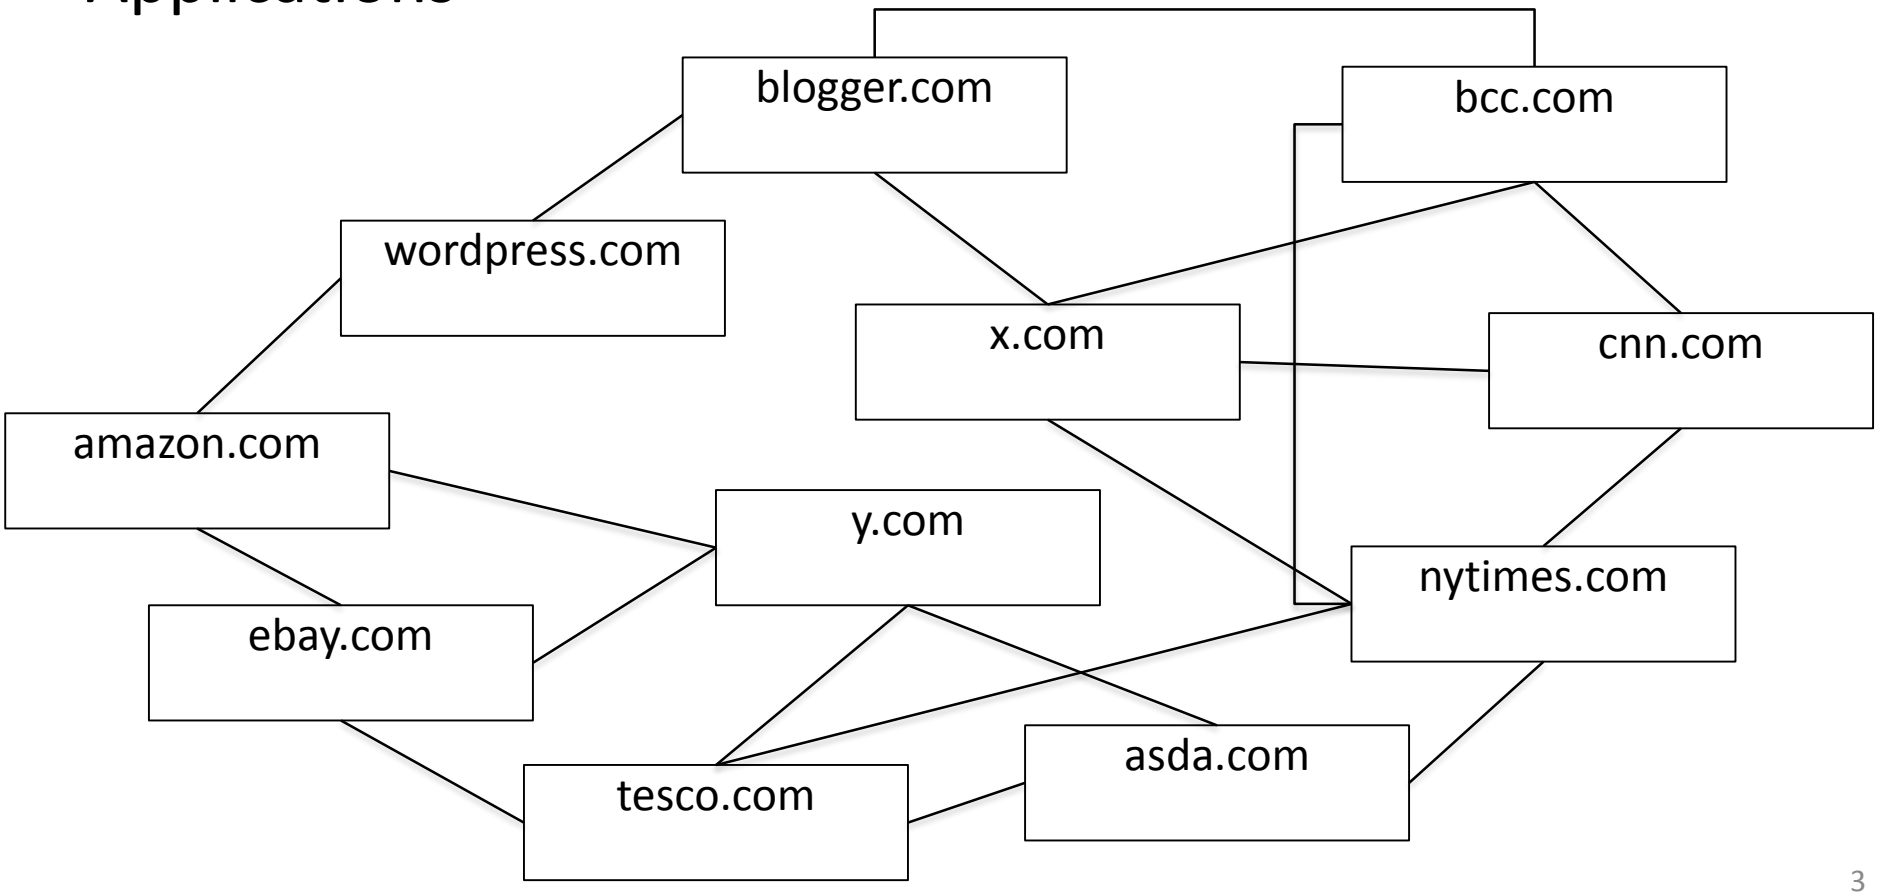

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

- Applications

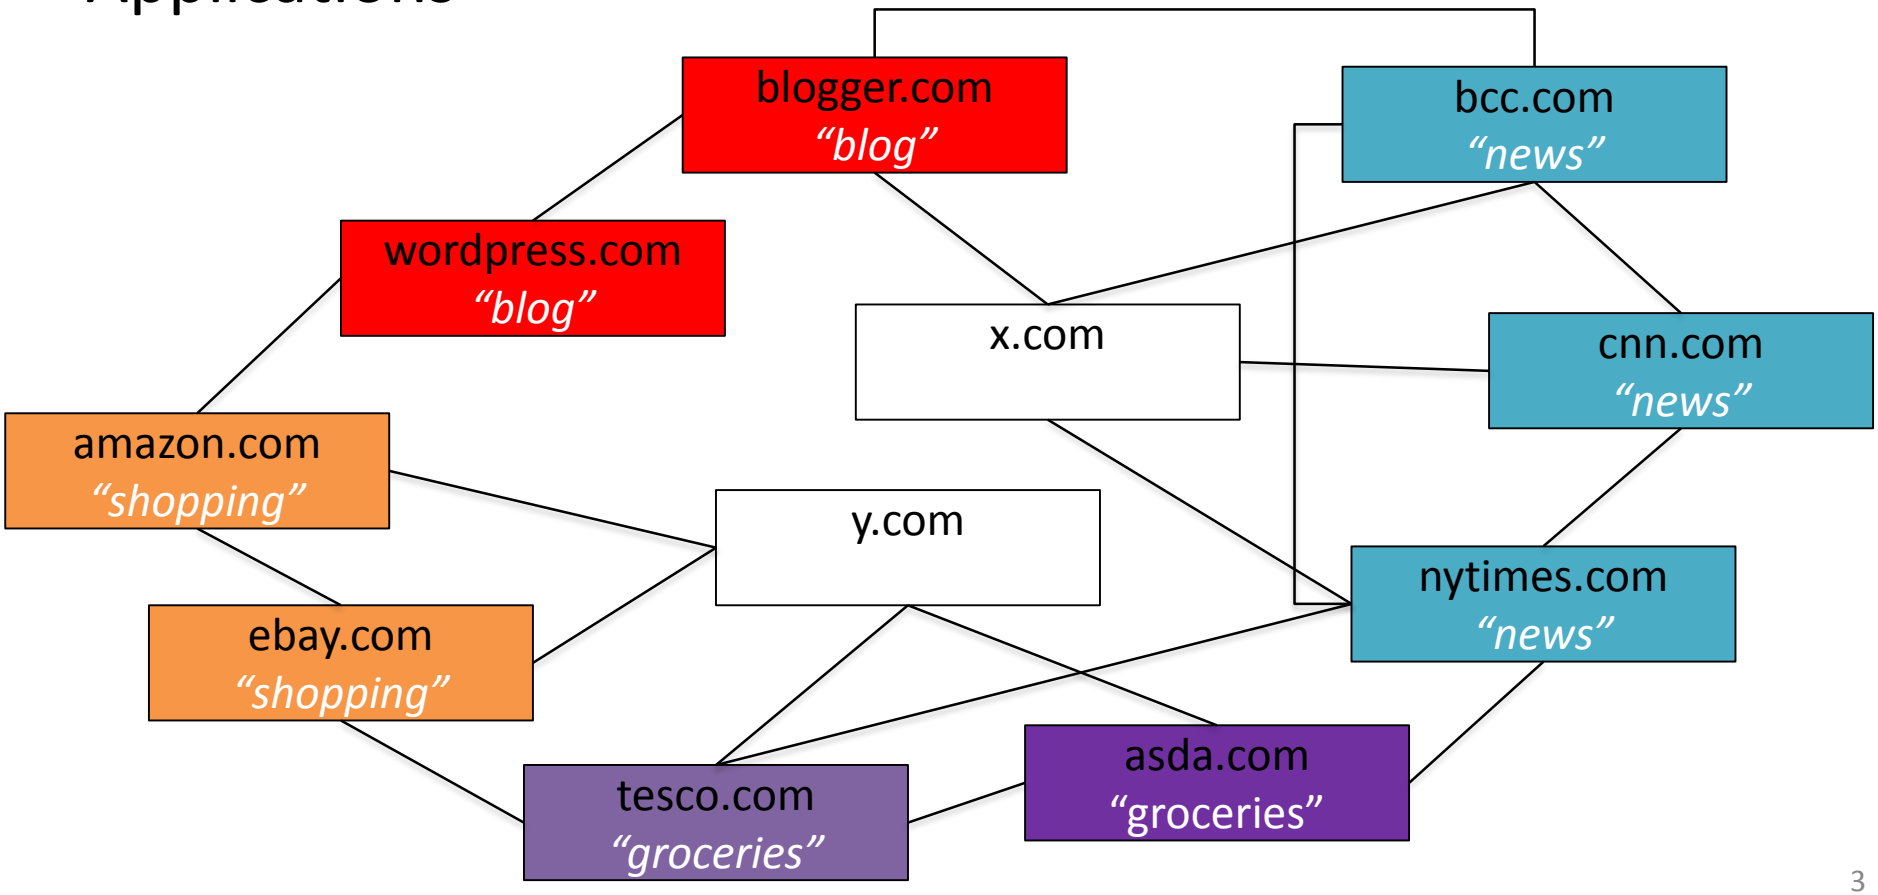

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

- Applications

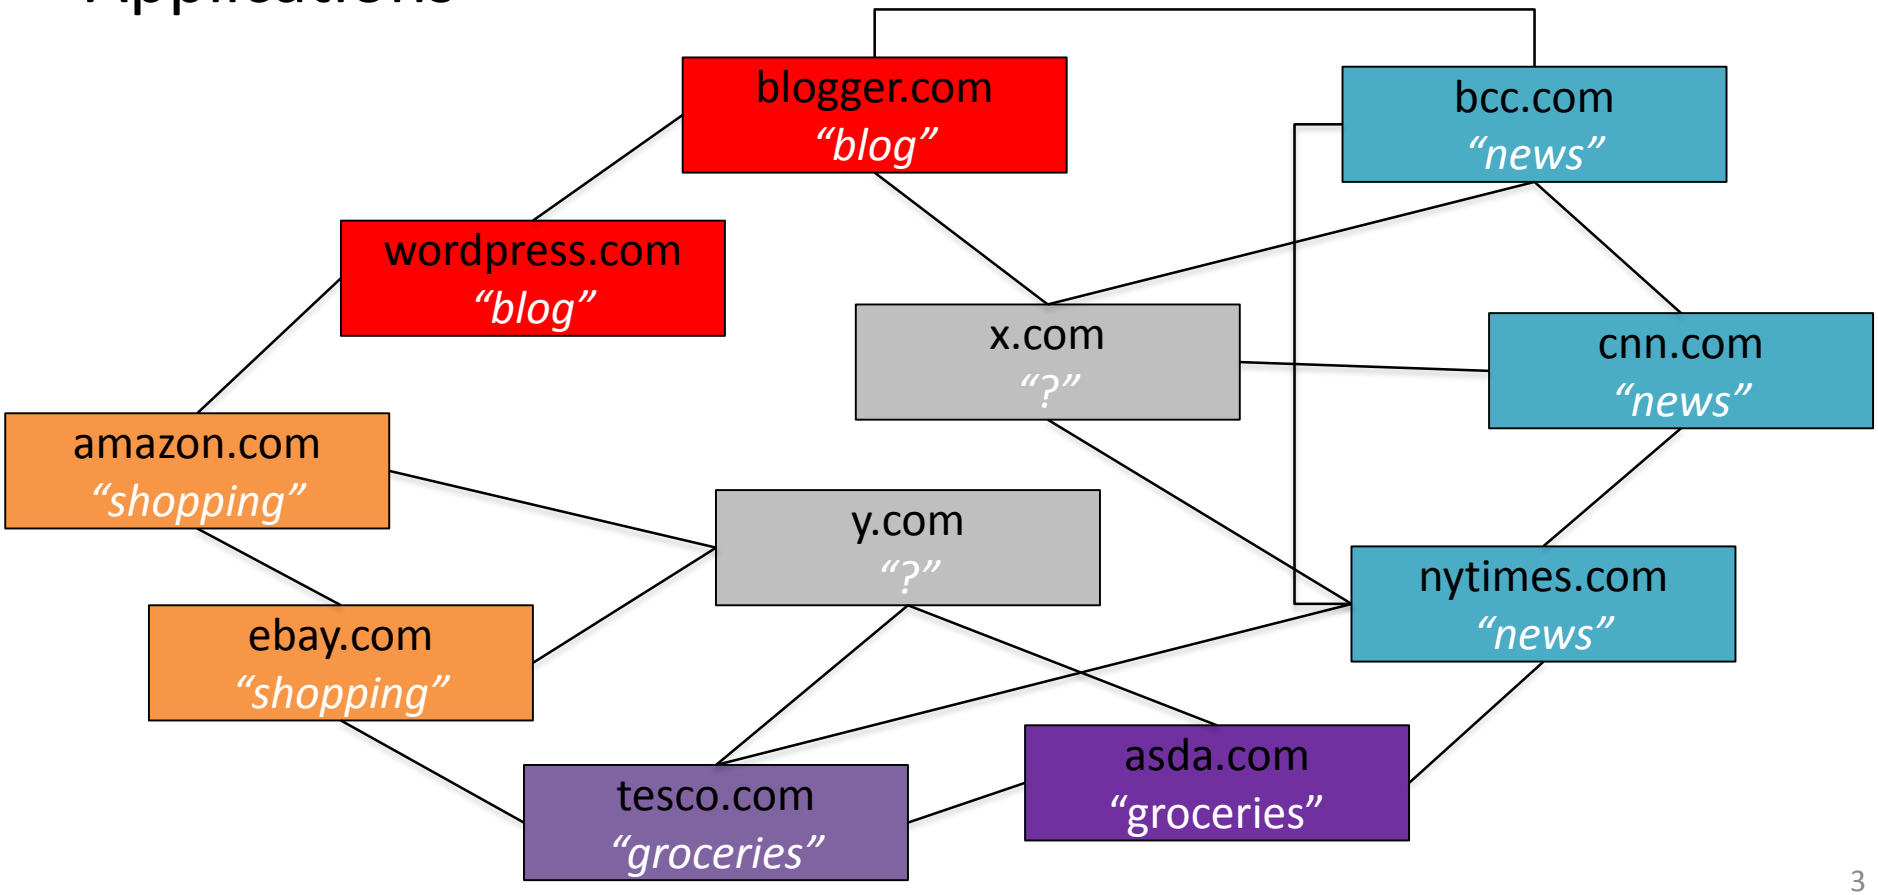

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

- Applications
  - Label Propagation

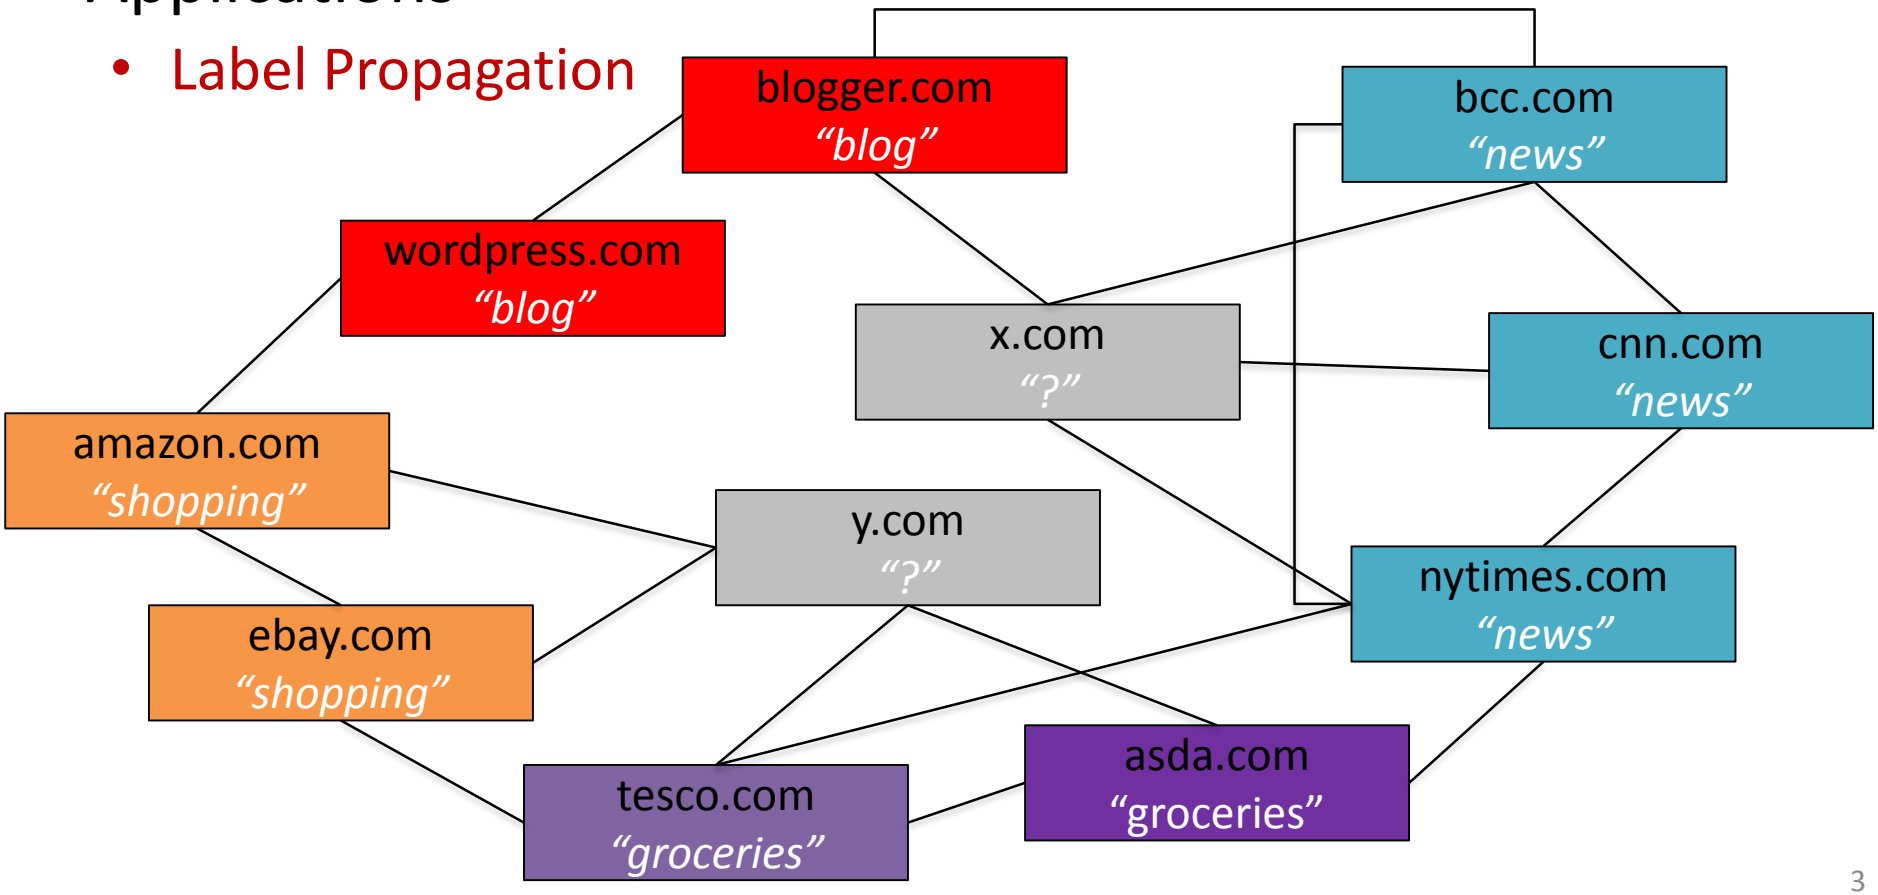

Graph Analytics

Extracting meaningful information out of complex many-to-many relationships among entities in a graph

- Applications
  - Label Propagation
  - Centrality Analysis
    - Most influential people and information in social media
  - Community Analysis
    - Identify customers with similar interests
  - Connectivity Analysis
    - Find weakness in a network
  - Path Analysis
    - Route optimization for distribution and supply chain

Graphs are huge and growing …

Graphs are huge and growing …

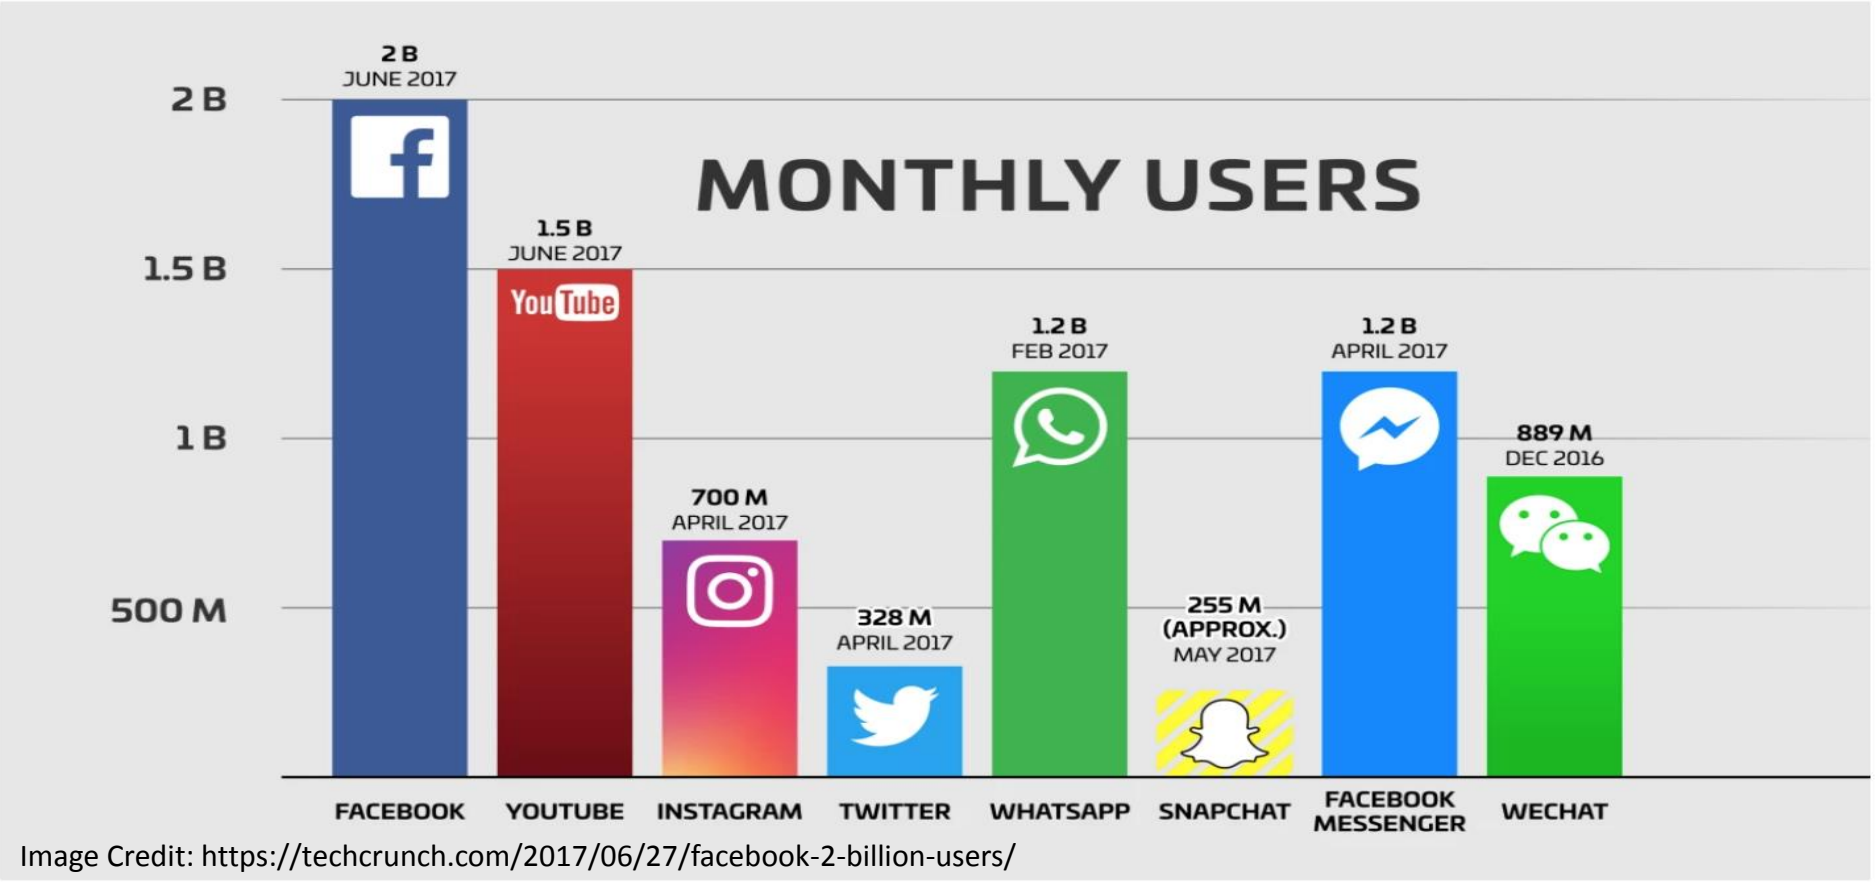

Image Credit: https://techcrunch.com/2017/06/27/facebook-2-billion-users/

Graph … owing …

Billions of Vertices
Trillions of Edges

Memory Footprint in Terabytes

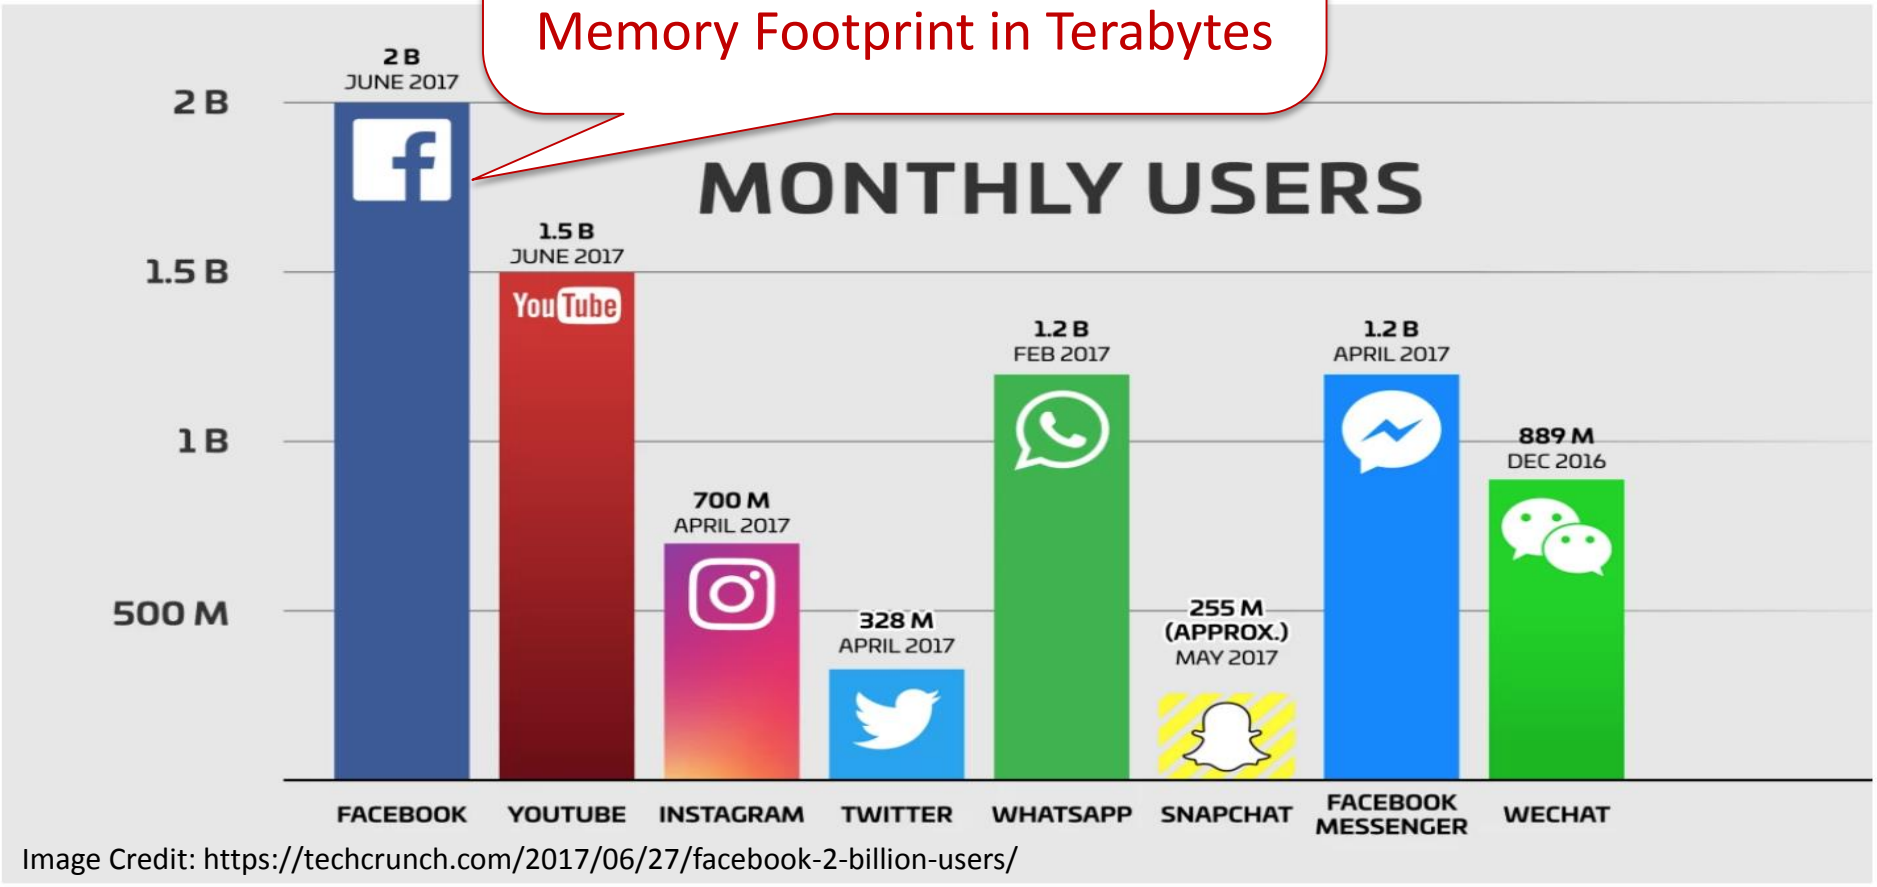

Image Credit: https://techcrunch.com/2017/06/27/facebook-2-billion-users/

Graph … owing …

Billions of Vertices
Trillions of Edges

Memory Footprint in Terabytes

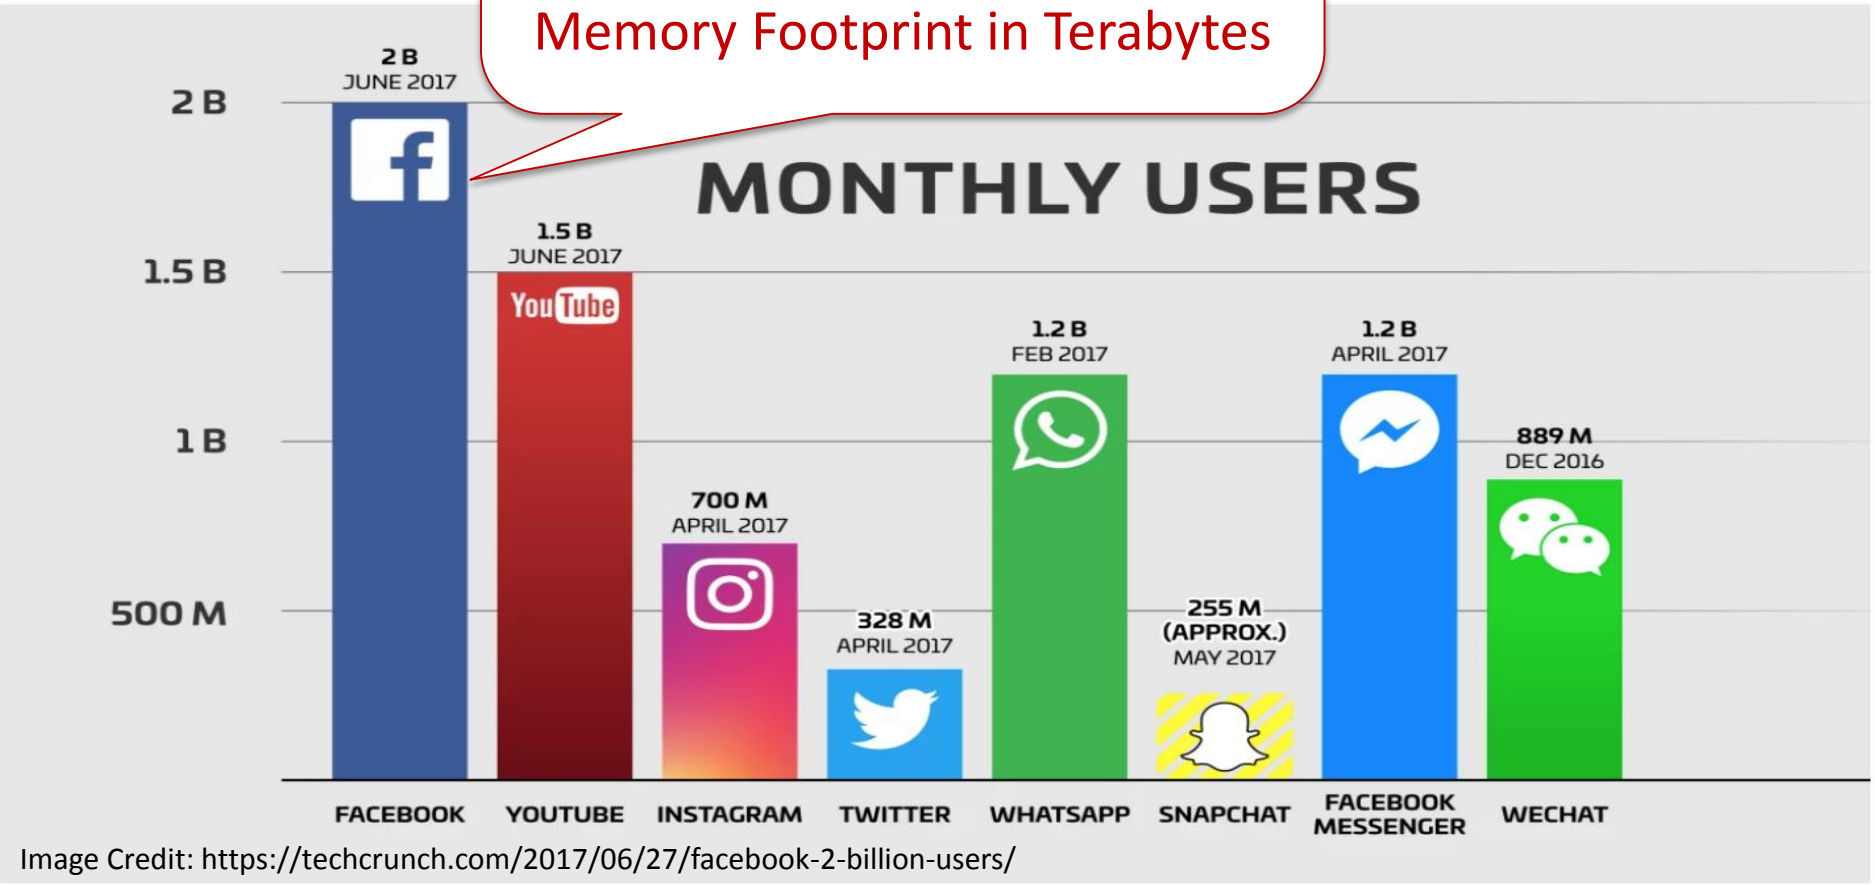

Image Credit: https://techcrunch.com/2017/06/27/facebook-2-billion-users/

Graphs don't fit in main memory of a single server

To Scale-Out
Or
To Scale-Up

Scale-Out Graph Analytics

- Graph is partitioned across a number of nodes
  - Graph is stored in combined memory of all nodes → Enables in-memory processing

Scale-Out Graph Analytics

- Graph is partitioned across a number of nodes
  - Graph is stored in combined memory of all nodes → Enables in-memory processing

PEGASUS [ICDM'09]

Pregel [SIGMOD'10]

PowerGraph [OSDI'12]

GraphLab [VLDI'12]

GraphX [OSDI'14]

GPS [SSDBM'13]

PGX.D [SC'15]

GRAM [SoCC'15]

Gemini [OSDI'16]

Scale-Out Graph Analytics

- Graph is partitioned across a number of nodes
  - Graph is stored in combined memory of all nodes → Enables in-memory processing

PEGASUS [ICDM'09]

Pregel [SIGMOD'10]

PowerGraph [OSDI'12]

GraphLab [VLDI'12]

GraphX [OSDI'14]

GPS [SSDBM'13]

PGX.D [SC'15]

GRAM [SoCC'15]

Gemini [OSDI'16]

Is scale-out graph analytics a good idea?

Scale-Out Graph Analytics Not Desirable

In-memory scale-out processing with 10s-100s nodes can be outperformed by disk-based scale-up graph processing

[GraphChi OSDI'12]

Scale-Out Graph Analytics Not Desirable

In-memory scale-out processing with 10s-100s nodes can be outperformed by disk-based scale-up graph processing

[GraphChi OSDI'12]

Reason for inefficiency of scale-out graph analytics:

- High inter-node communication & synchronization overhead
  - "*Small World*" property of graphs → Difficult to partition

Scale-Out Graph Analytics Not Desirable

In-memory scale-out processing with 10s-100s nodes can be outperformed by disk-based scale-up graph processing

[GraphChi OSDI'12]

Reason for inefficiency of scale-out graph analytics:

- High inter-node communication & synchronization overhead
  - "*Small World*" property of graphs → Difficult to partition

So, why not scale-up graph analytics?

Need large per-node memory capacity

- DRAM: a limiting factor in scaling up
  - Poor technology scaling ☹
  - High cost for large capacity ☹

Need large per-node memory capacity

- DRAM: a limiting factor in scaling up
  - Poor technology scaling ☹
  - High cost for large capacity ☹
- Alternative technology for main memory?

Need large per-node memory capacity

- DRAM: a limiting factor in scaling up
  - Poor technology scaling ☹
  - High cost for large capacity ☹
- Alternative technology for main memory?
  - Emerging solution: Storage Class Memory (SCM)
  - **Terabytes of capacity at affordable price** ☺

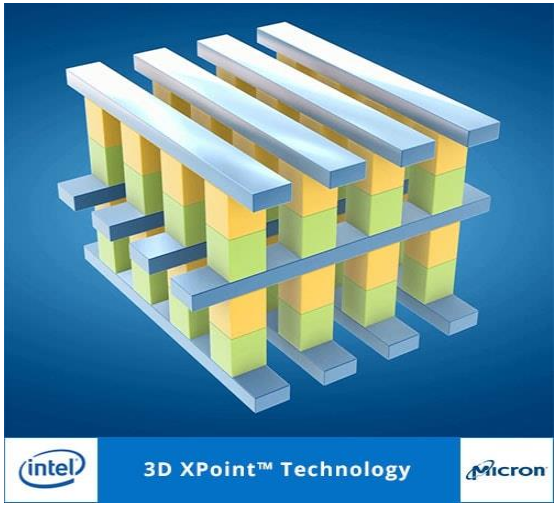

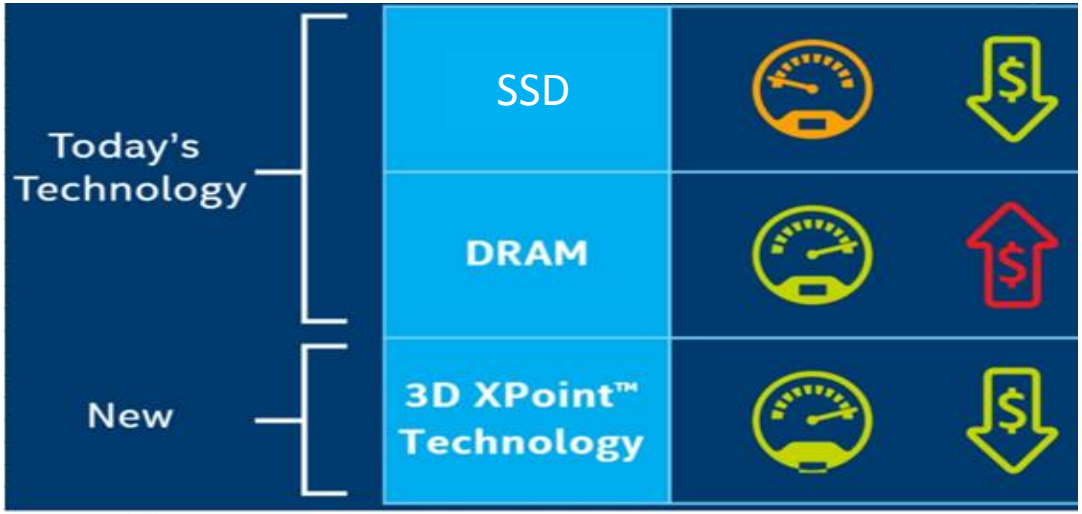

Image Credit: https://marketrealist.com/2016/03/microns-3d-xpoint-launch-stands-now

Need large per-node memory capacity

SCM: Enabler for scale-up graph analytics

- Alternative technology for main memory?
  - Emerging solution: Storage Class Memory (SCM)
  - **Terabytes of capacity at affordable price** ☺

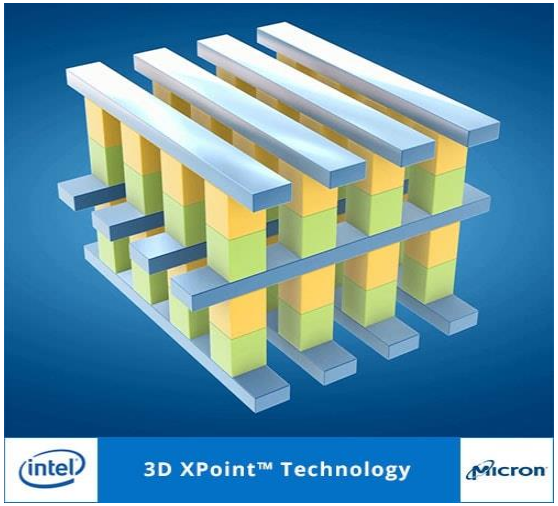

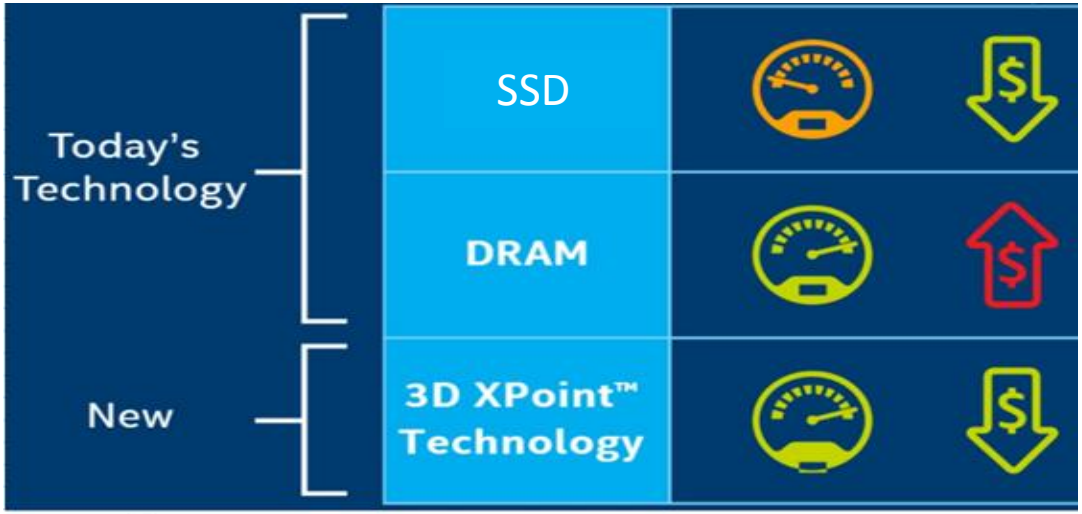

Image Credit: https://marketrealist.com/2016/03/microns-3d-xpoint-launch-stands-now

SCM: No Free Lunch

SCM: No Free Lunch

- 2x-4x slower access latency than DRAM ☹
- 4x-8x lower bandwidth than DRAM ☹
- Multiple orders of magnitude lower write endurance than DRAM ☹

SCM: No Free Lunch

- 2x-4x slower access latency than DRAM ☹
- 4x-8x lower bandwidth than DRAM ☹
- Multiple orders of magnitude lower write endurance than DRAM ☹

Implications for scale-up graph analytics

CPI Stack of DRAM based Scale-Up Graph Analytics

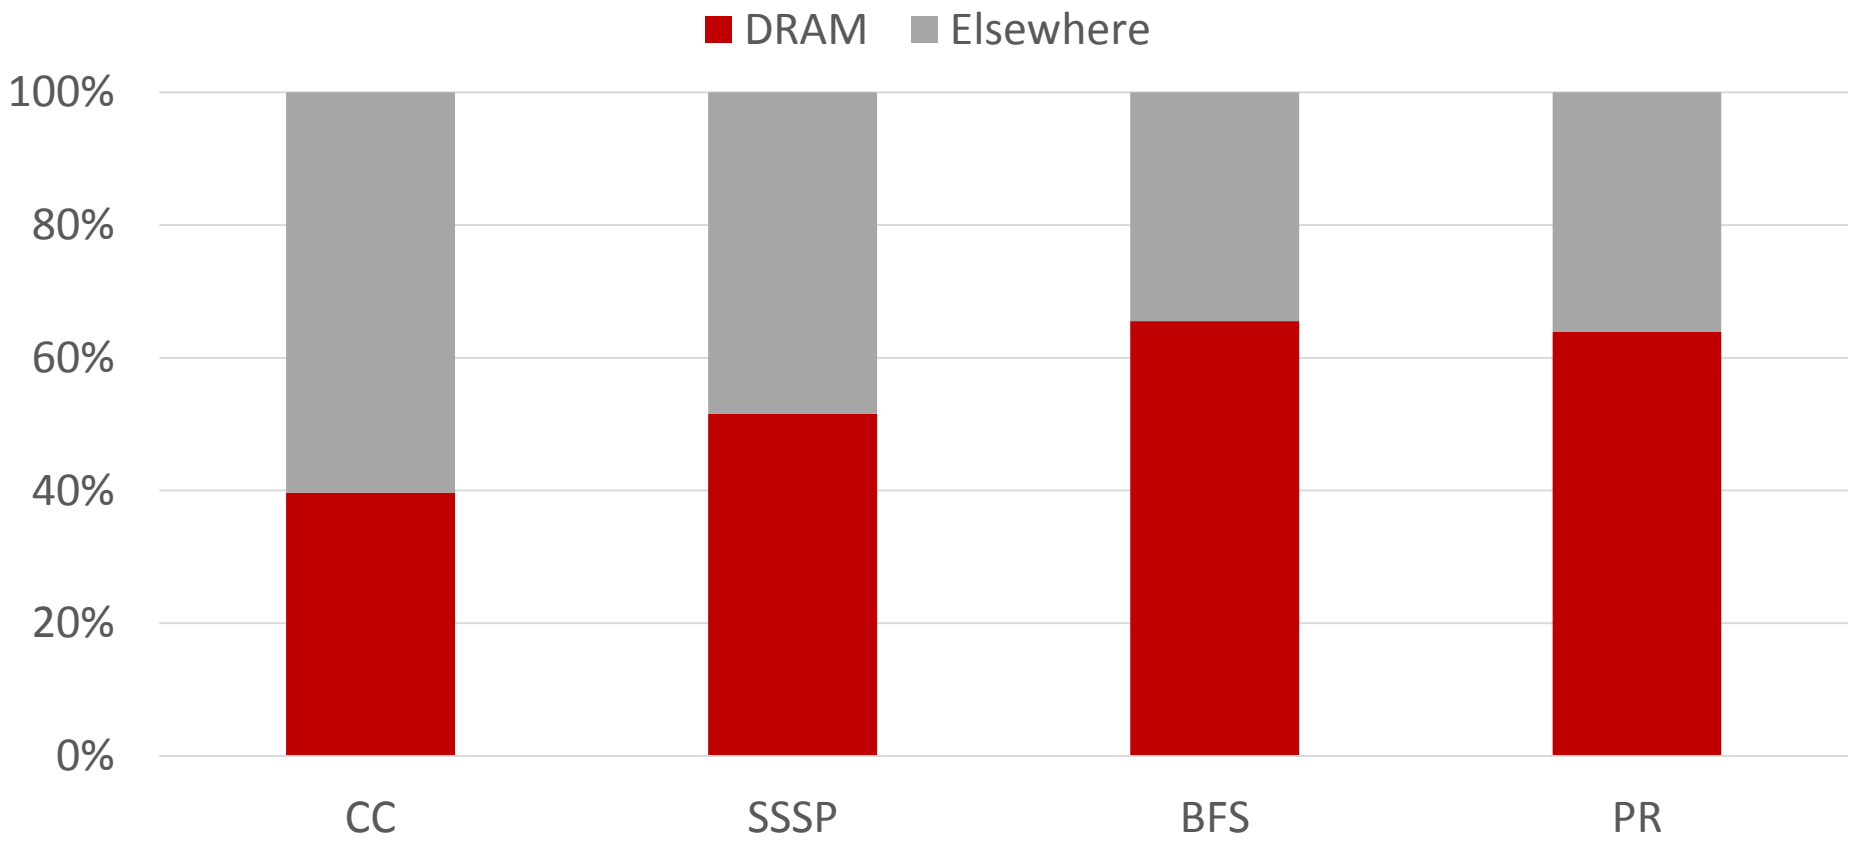

CPI Stack of DRAM based Scale-Up Graph Analytics

Cycles spent in DRAM vs Elsewhere

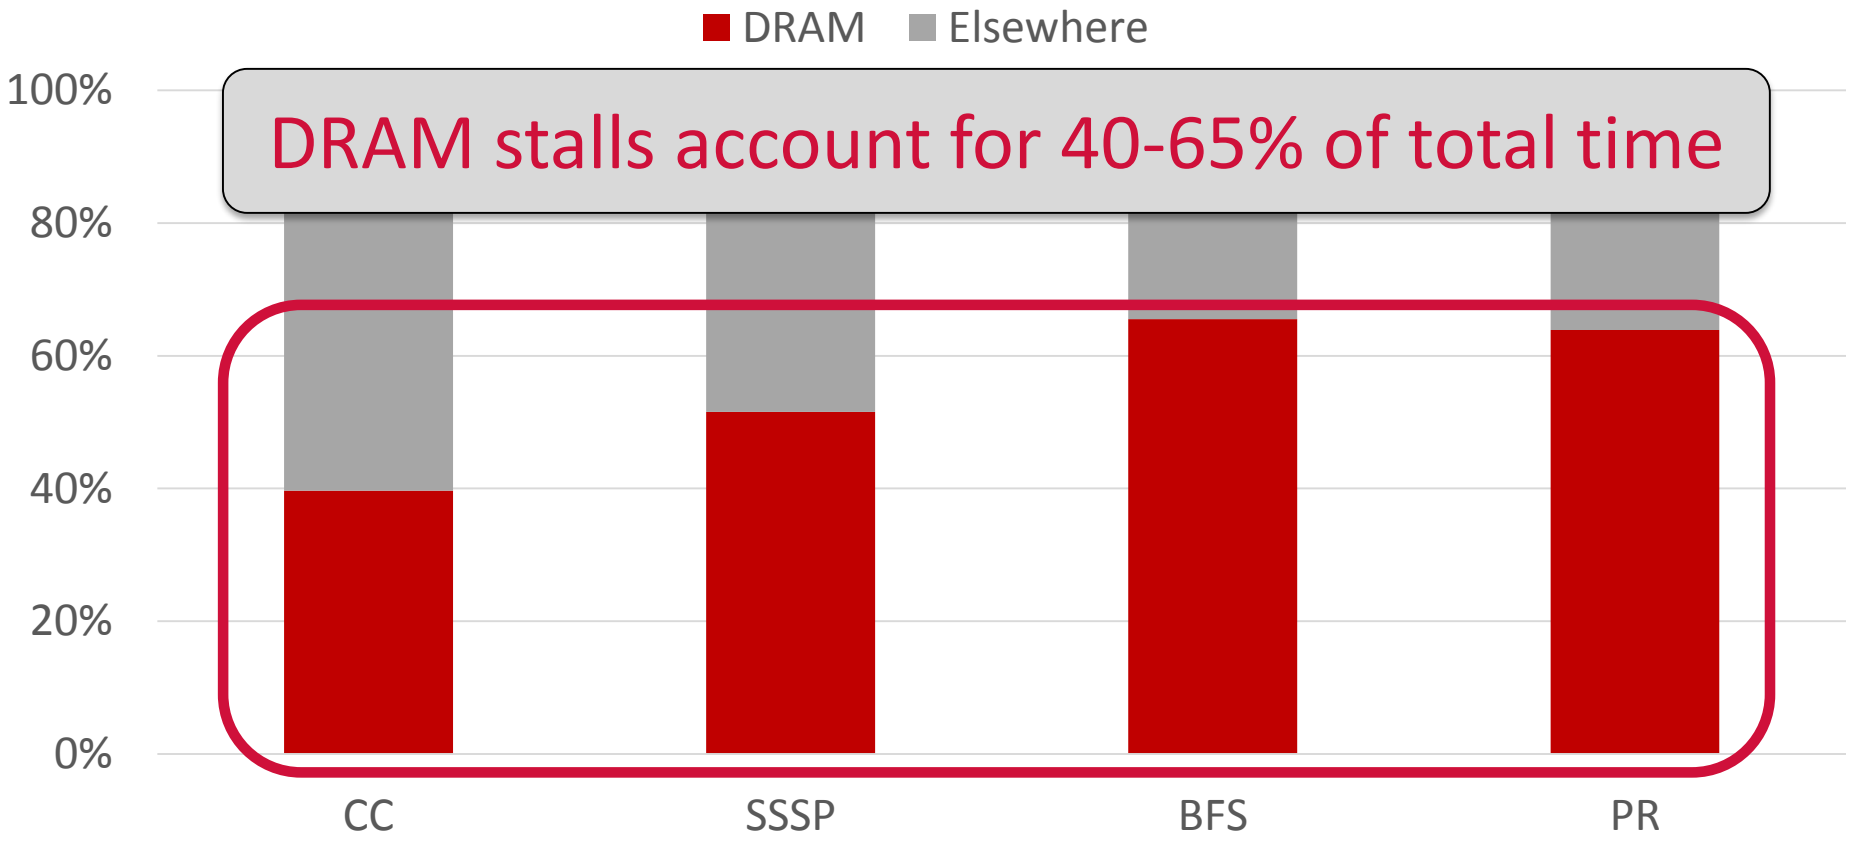

CPI Stack of DRAM based Scale-Up Graph Analytics

Cycles spent in DRAM vs Elsewhere

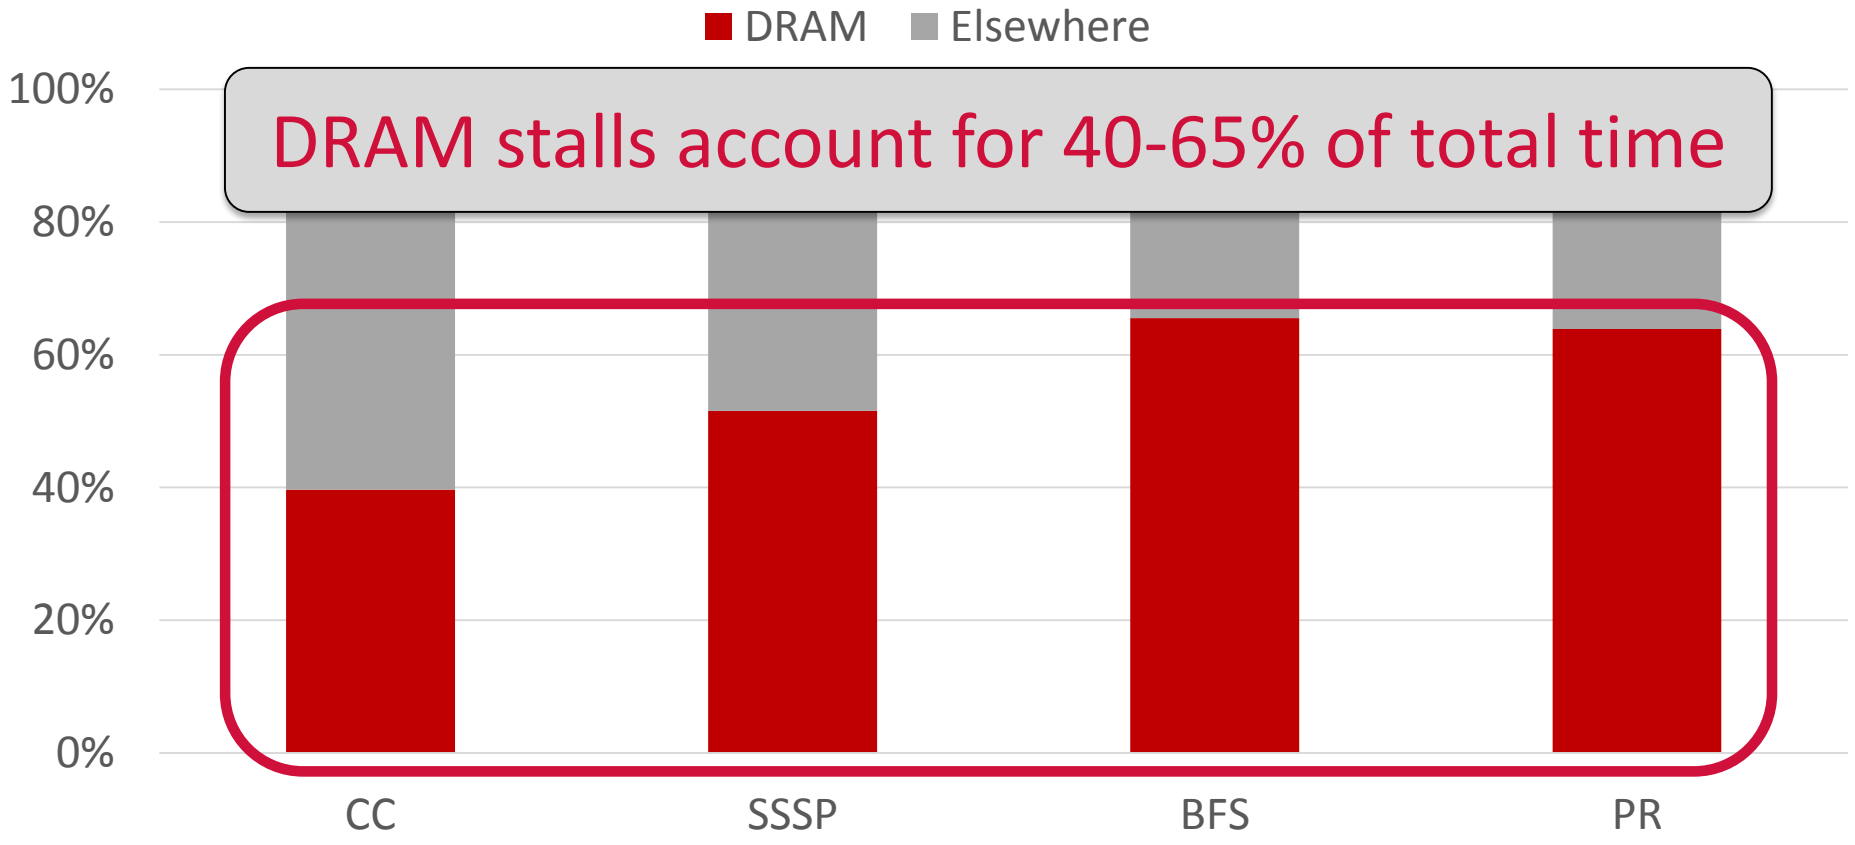

SCM will exacerbate the problem

Mitigating Slower Latency of SCM

Data Prefetching for latency hiding

Challenges:

- Graphs exhibit random access patterns
- Simple hardware prefetchers are inadequate for graphs

Mitigating Slower Latency of SCM

Data Prefetching for latency hiding

Challenges:

- Graphs exhibit random access patterns
- Simple hardware prefetchers are inadequate for graphs

Potential Solutions:

Hardware

- Graph specific hardware prefetcher [Ainsworth et al. ICS'16]

Software

- Employ software prefetcher in the framework
  - CPU often idle waiting on memory → Plenty of idle cycles to burn on extra instructions

Mitigating Lower Bandwidth of SCM

- Data transfer between DRAM & CPU happens at cache-line granularity
- Observation: Graphs exhibit poor spatial locality

Mitigating Lower Bandwidth of SCM

- Data transfer between DRAM & CPU happens at cache-line granularity
- Observation: Graphs exhibit poor spatial locality

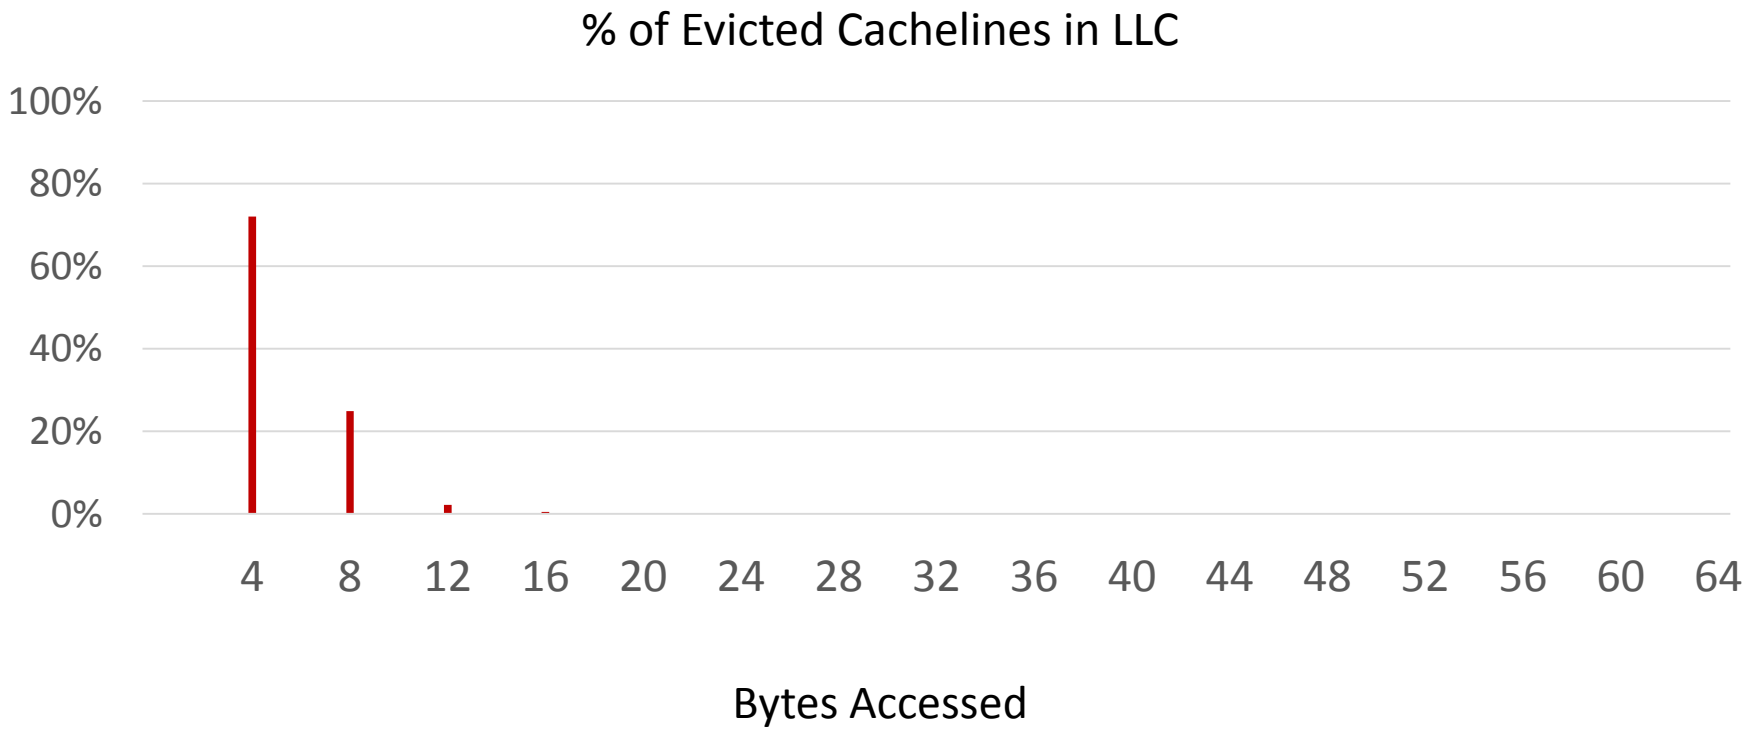

Mitigating Lower Bandwidth of SCM

- Data transfer between DRAM & CPU happens at cache-line granularity
- Observation: Graphs exhibit poor spatial locality

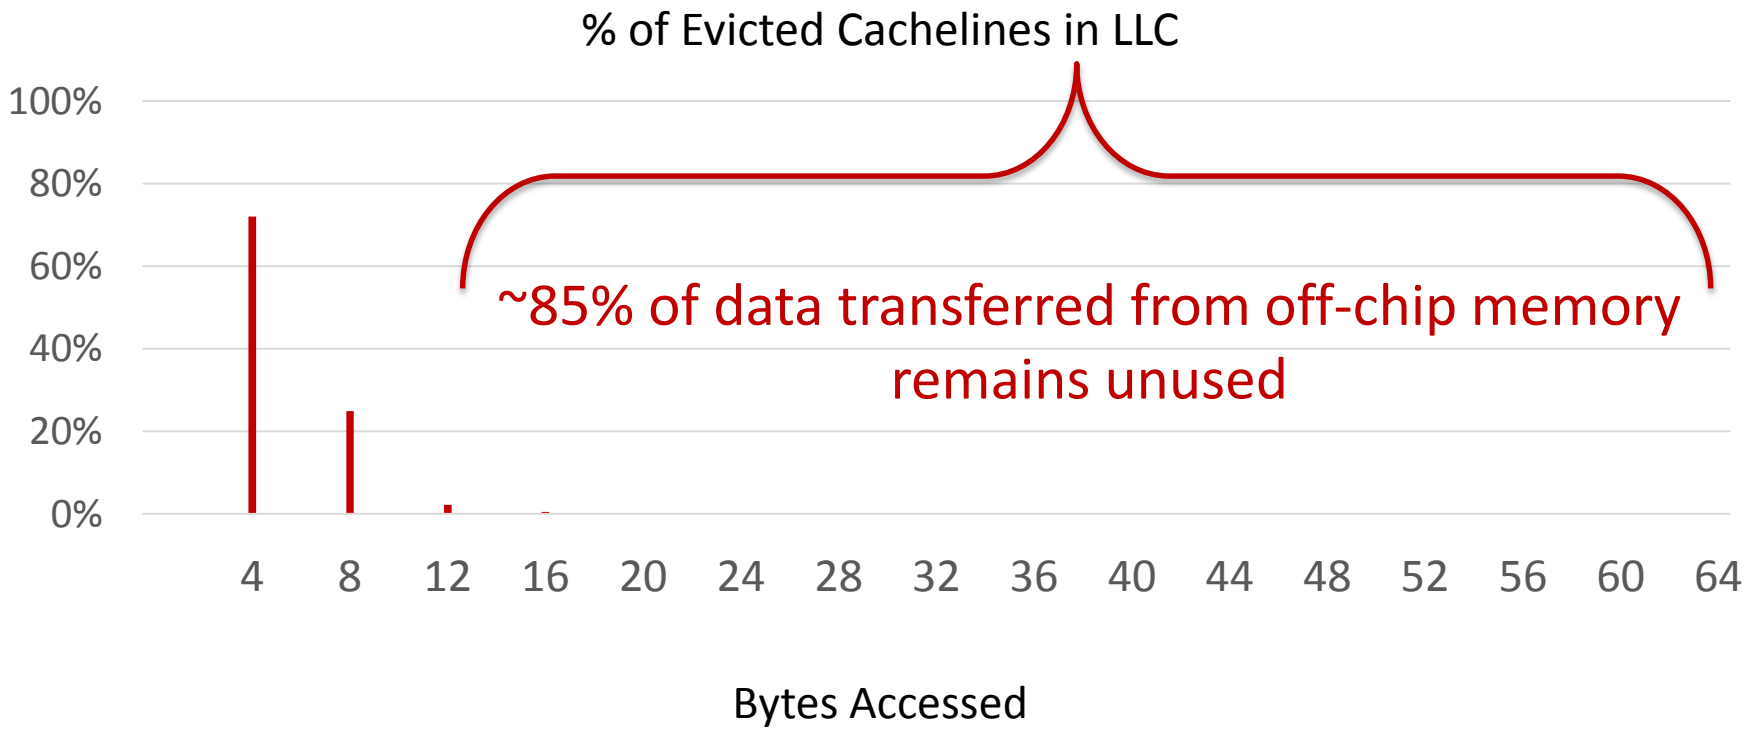

Mitigating Lower Bandwidth of SCM

Potential Solutions:

Software

- Exploit graph topology to reorder vertices → Improved spatial locality

Hardware

- Revisit *sectored caches* → Only fetch the required words into on-chip caches

Mitigating Lower Write Endurance of SCM

New design goal: Reduce off-chip write traffic

Potential Solutions:

Software:

- Approaches to improve write locality in on-chip caches
  - E.g., pull-based vs push-based approach

Hardware:

- Aggressively retain cachelines that accumulate writes
  - Even at the expense of short temporal reuse of some cachelines

Conclusion

- Introduction of SCM will provide large capacity main memory at affordable price in commodity server

- SCM will enable in-memory scale-up graph analytics even for extremely large graphs
  - Open research questions on how to address weaknesses of SCM to achieve high performance for scale-up graph analytics

Scale-Out

Conclusion

- Introduction of SCM will provide large capacity main memory at affordable price in commodity server

- SCM will enable in-memory scale-up graph analytics even for extremely large graphs
  - Open research questions on how to address weaknesses of SCM to achieve high performance for scale-up graph analytics

Scale-Up ~~Out~~

Thank You

Priyank Faldu

PhD Student

The University of Edinburgh

www.faldupriyank.com

Boris Grot

Lecturer

The University of Edinburgh

http://homepages.inf.ed.ac.uk/bgrot