# Smaller and Faster: Parallel Processing of Compressed Graphs with Ligra+

Julian Shun

Based on joint work with Guy Blelloch and Laxman Dhulipala

# Ligra Graph Processing Framework

EdgeMap

VertexMap

| | |
|---|---|
| Breadth-first search | Single-source shortest paths |
| Betweenness centrality | Eccentricity estimation |
| Connected components | (Personalized) PageRank |
| Triangle counting | Local graph clustering |
| K-core decomposition | Biconnected components |
| Maximal independent set | Collaborative filtering |
| Set cover | … |

*Simplicity, Performance, Scalability*

# Steps for Graph Traversal

- Operate on a subset of vertices
- Map computation over subset of edges in parallel
- Return new subset of vertices
- Map computation over subset of vertices in parallel

Graph

VertexSubset

EdgeMap

VertexMap

# Large Graphs

## Amazon EC2

| | vCPU | ECU | Memory (GiB) | Instance Storage (GB) | Linux/UNIX Usage |
|---|---|---|---|---|---|
| x1e.xlarge | 4 | 12 | 122 | 1 x 120 SSD | $0.834 per Hour |
| x1e.2xlarge | 8 | 23 | 244 | 1 x 240 SSD | $1.668 per Hour |
| x1e.4xlarge | 16 | 47 | 488 | 1 x 480 SSD | $3.336 per Hour |
| x1e.8xlarge | 32 | 91 | 976 | 1 x 960 | $6.672 per Hour |
| x1e.16xlarge | 64 | 179 | 1952 | 1 x 1920 SSD | $13.344 per Hour |
| x1e.32xlarge | 128 | 340 | 3904 | 2 x 1920 SSD | $26.688 per Hour |

- Most can fit on commodity shared memory machine

Example
Dell PowerEdge R930:
Up to 96 cores and 6 TB of RAM

# What if you don't have or can't afford that much memory?

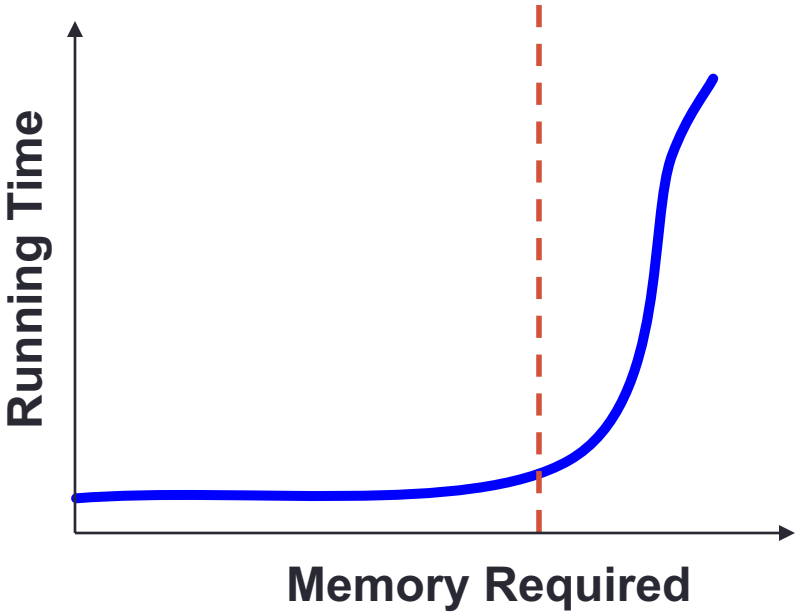


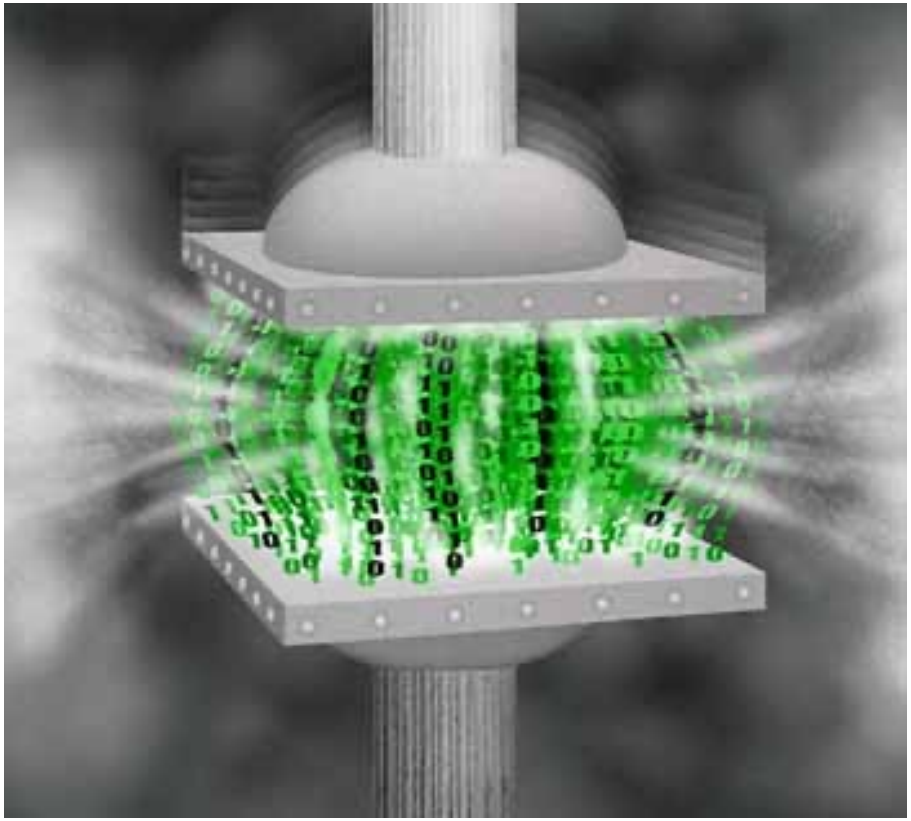

*Graph Compression*

# Ligra+: Adding Graph Compression to Ligra

# Ligra+: Adding Graph Compression to Ligra

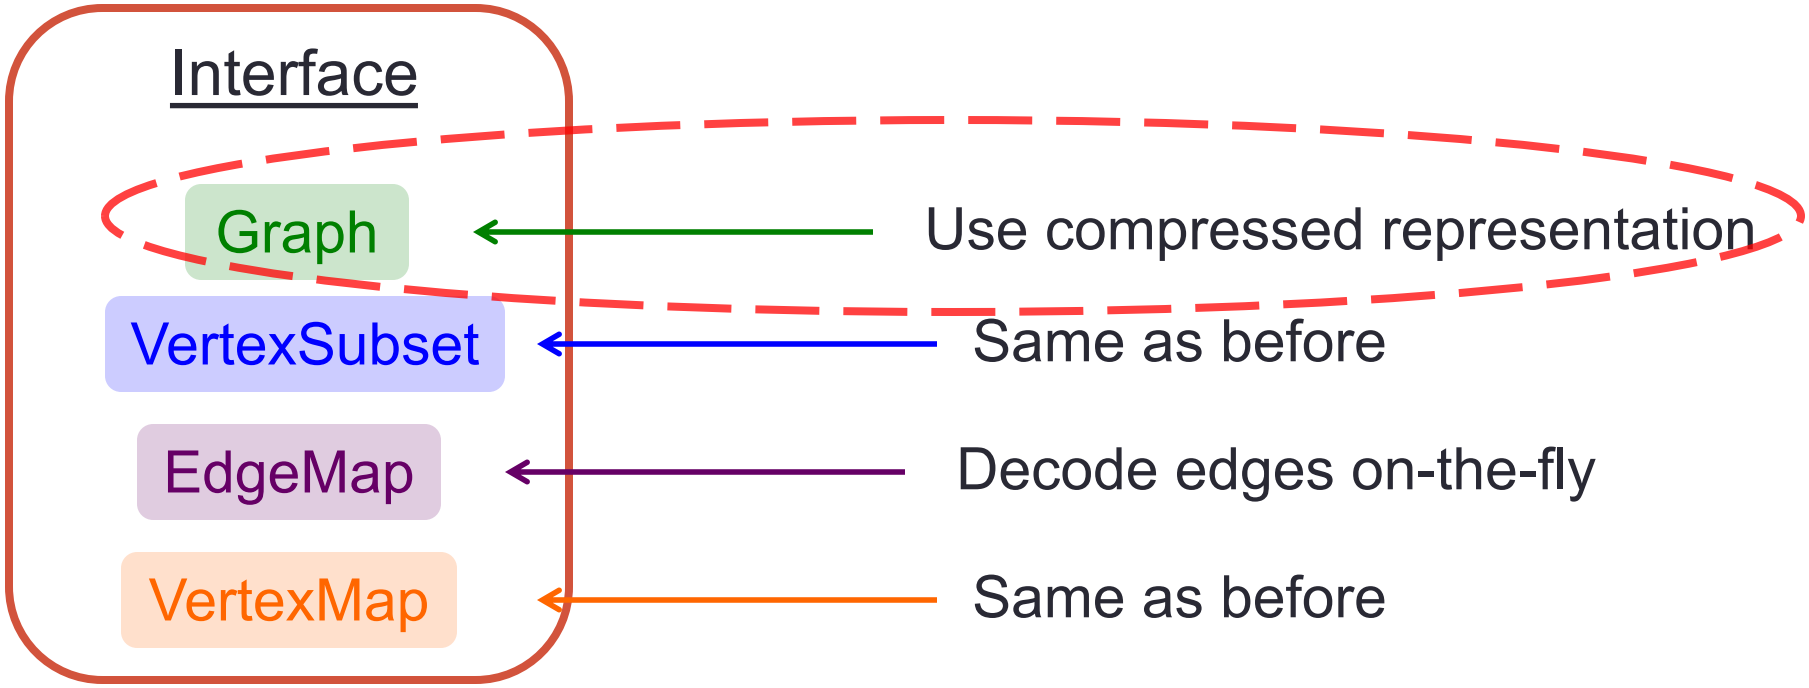


- Same interface as Ligra
- All changes hidden from the user!

# Graph representation

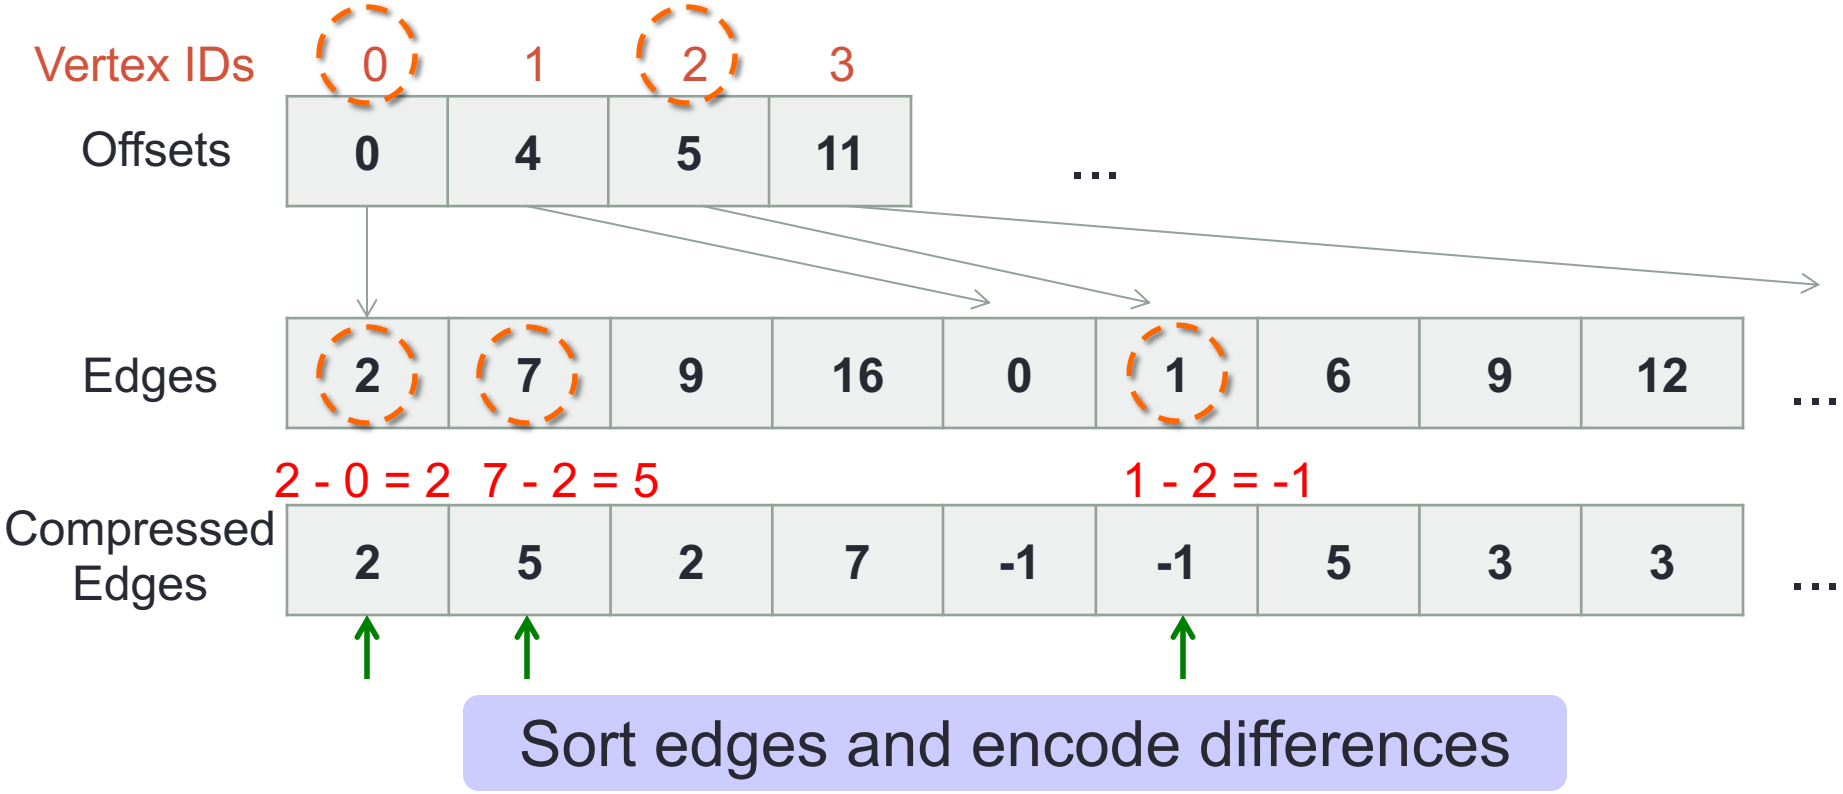

# Variable-length codes

- k-bit codes
  - Encode value in chunks of k bits
  - Use k-1 bits for data, and 1 bit as the “continue” bit
- Example: encode “401” using 8-bit (byte) code
- In binary:

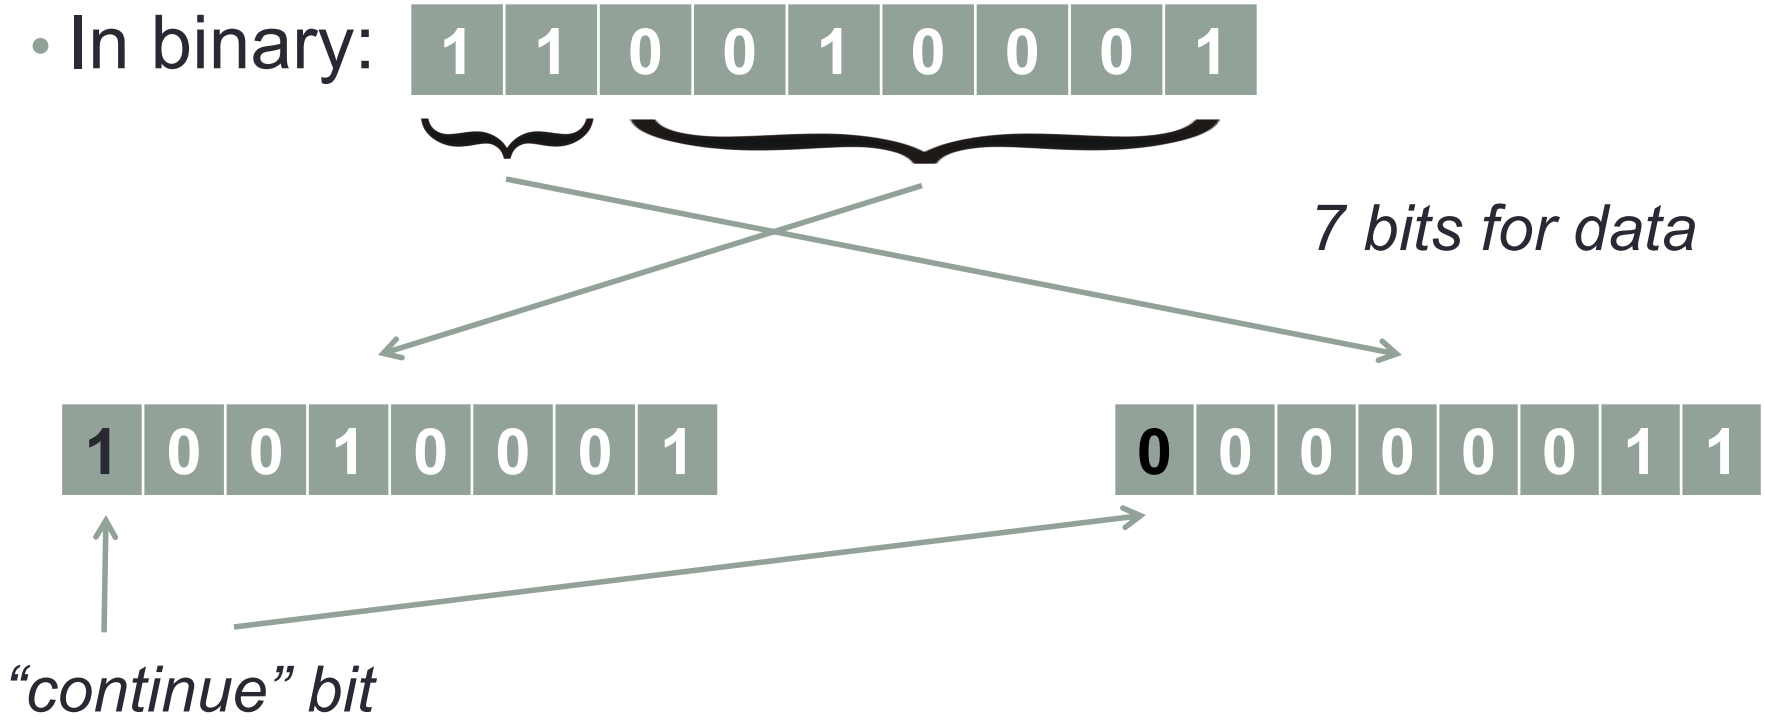

# Encoding optimization

- Another idea: get rid of “continue” bits

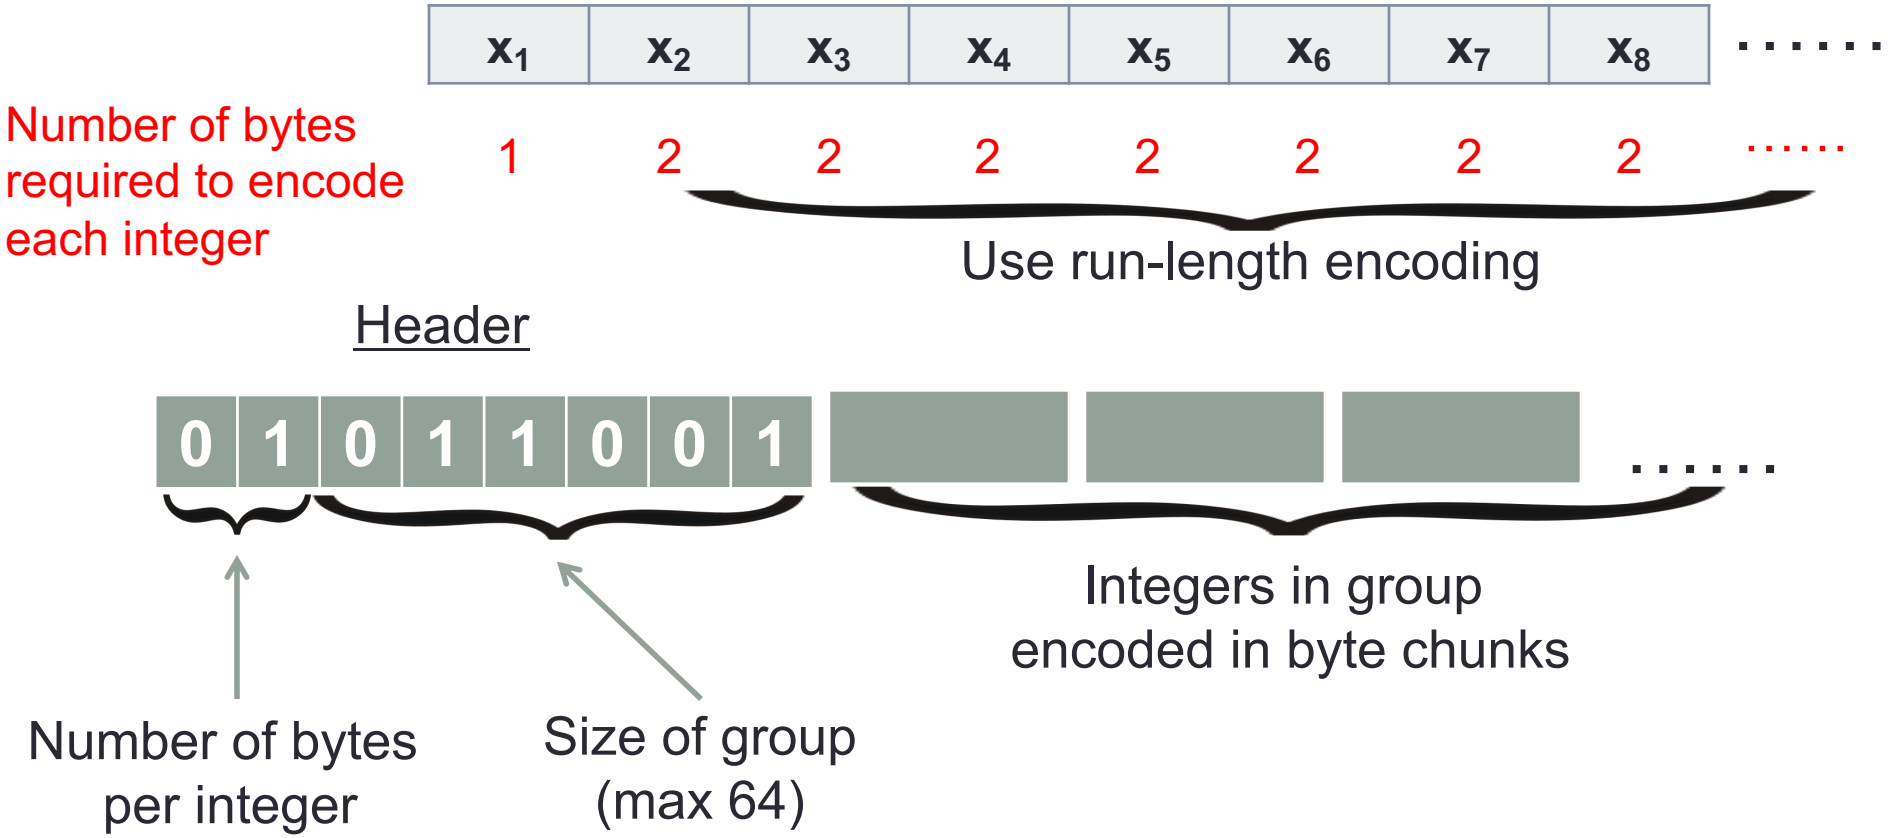


- Increases space, but makes decoding cheaper (no branch misprediction from checking “continue” bit)

# Ligra+: Adding Graph Compression to Ligra

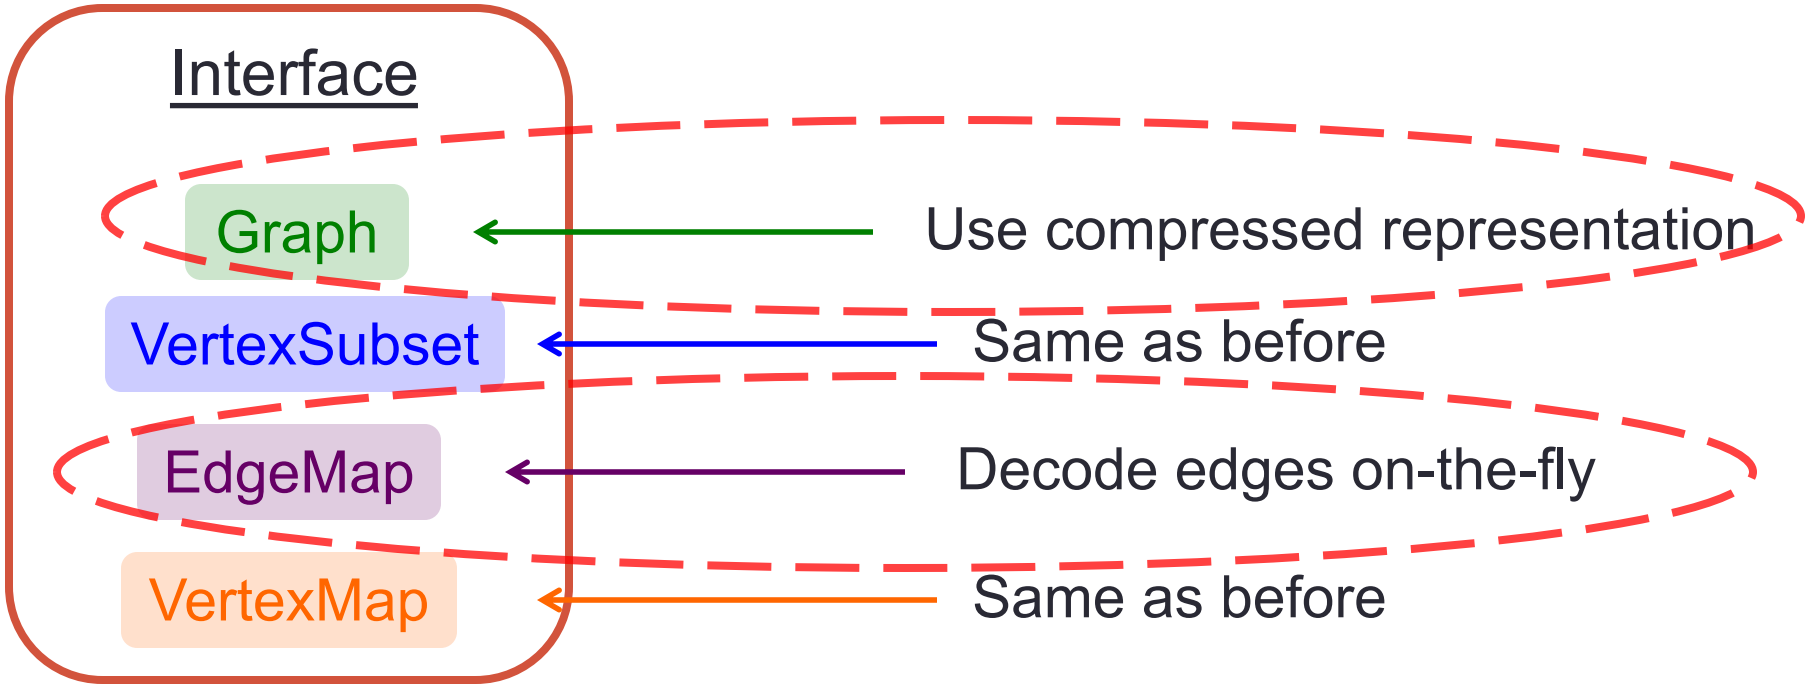


- Same interface as Ligra
- All changes hidden from the user!

# Modifying EdgeMap

- Processes outgoing edges of a subset of vertices

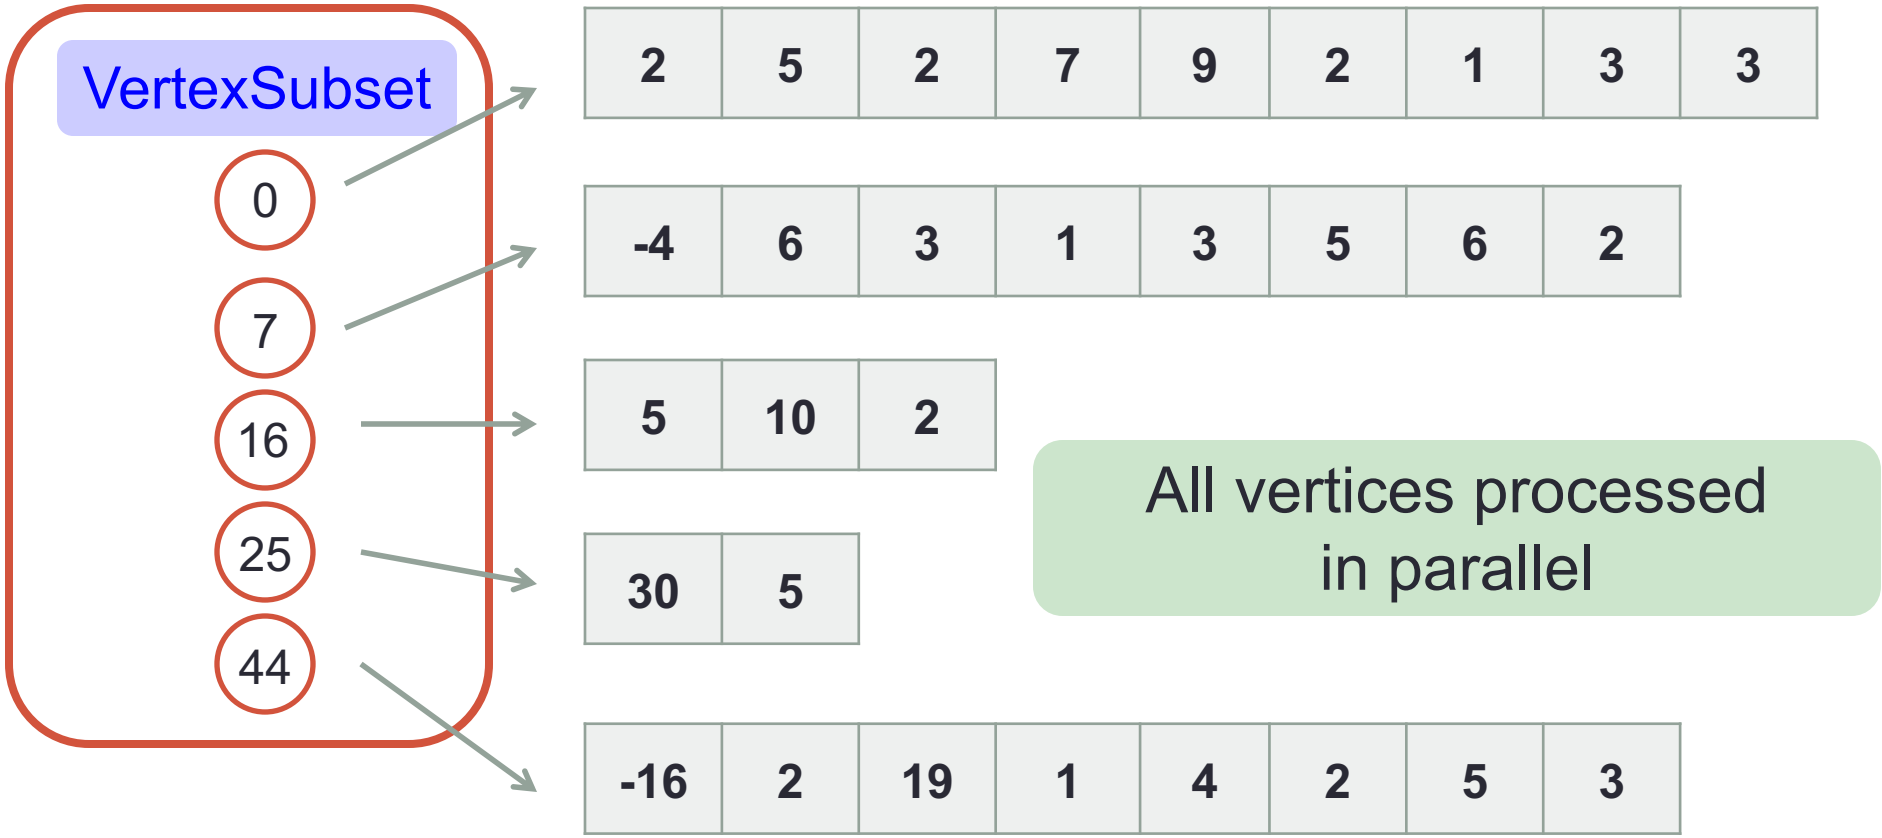


*What about high-degree vertices?*

# Handling high-degree vertices

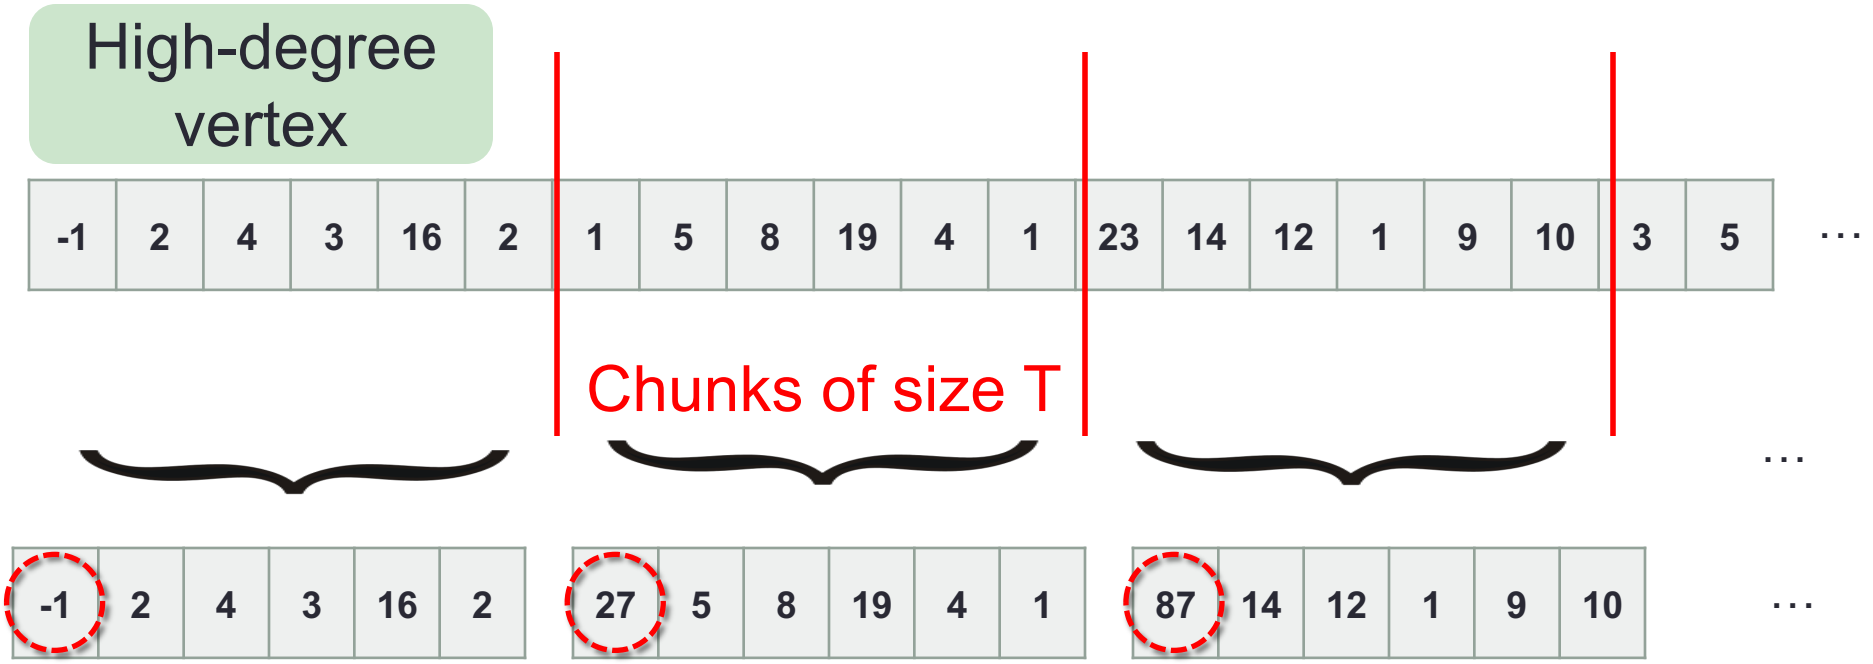


All chunks can be decoded in parallel!

- We chose T=1000
- Similar performance and space usage for a wide range of T

# Ligra+ Space Savings

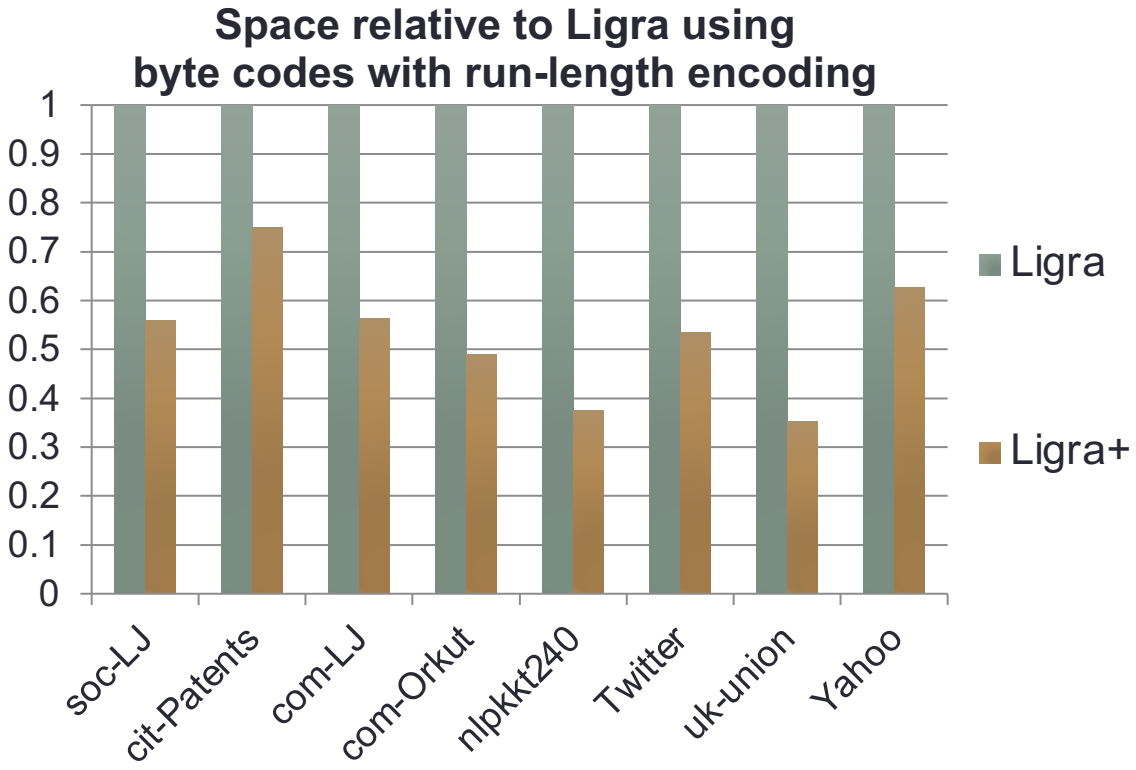


- Space savings of about 1.3—3x
- Could use more sophisticated schemes to further reduce space, but more expensive to decode
- Cost of decoding on-the-fly?

# Ligra+ Performance

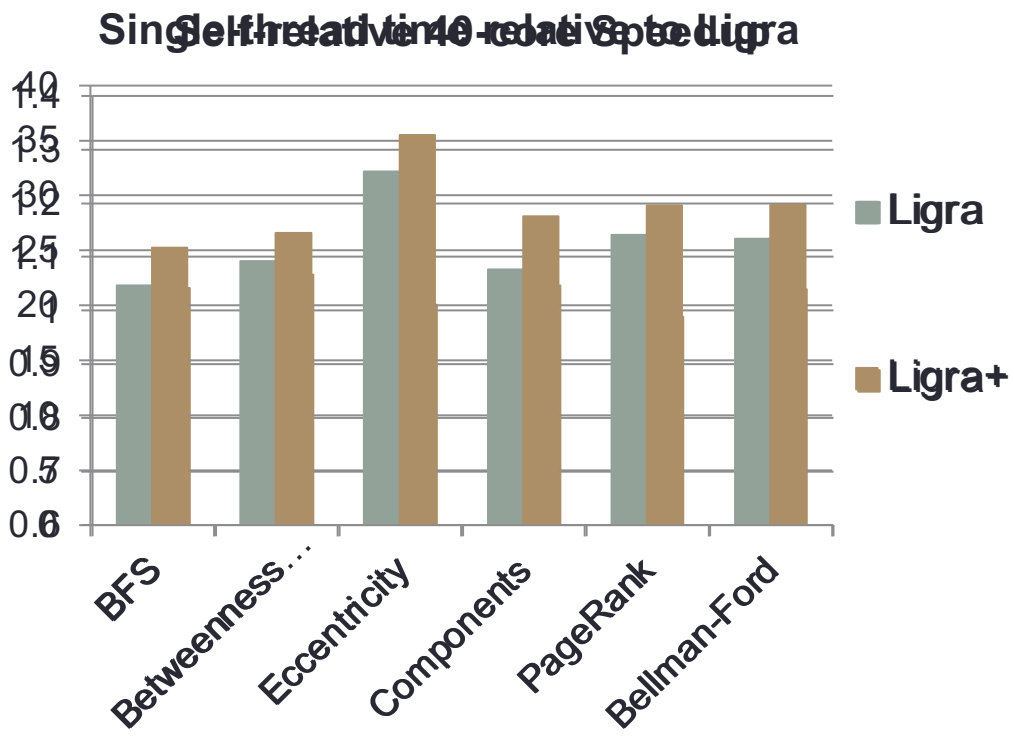


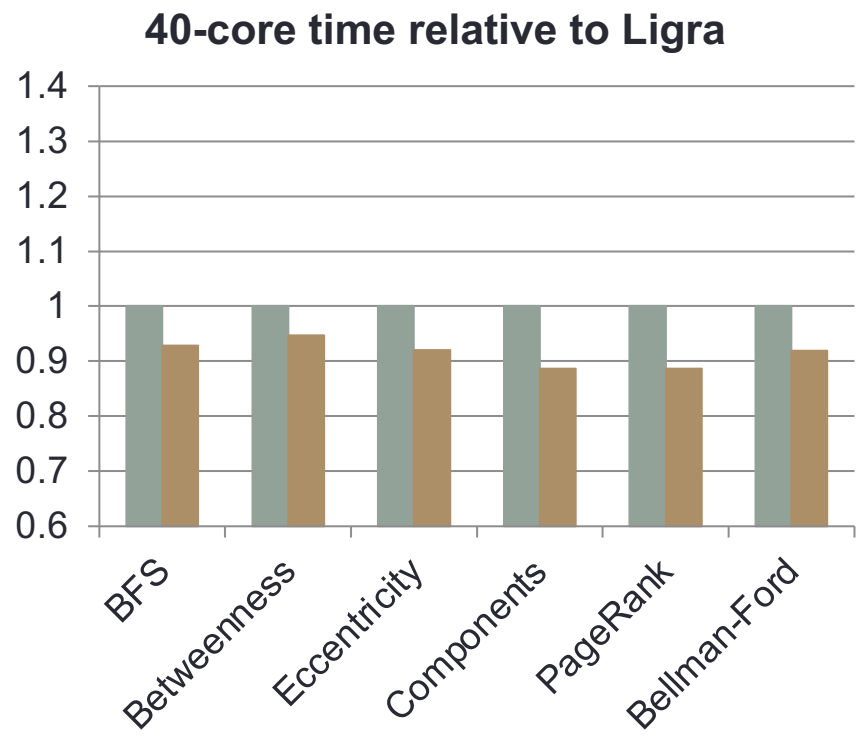


- Cost of decoding on-the-fly?
- Memory subsystem is a scalability bottleneck in parallel as these graph algorithms are memory-bound
- Ligra+ decoding gets better parallel speed up

# Ligra Summary

VertexSubset | VertexMap | EdgeMap

*Optimizations: Hybrid traversal and graph compression*

| | |
|---|---|
| Breadth-first search | Single-source shortest paths |
| Betweenness centrality | Eccentricity estimation |
| Connected components | (Personalized) PageRank |
| Triangle counting | Local graph clustering |
| K-core decomposition | Biconnected components |
| Maximal independent set | Collaborative filtering |
| … | … |

*Simplicity, Performance, Scalability*

# Thank you!

J. Shun and G. E. Blelloch. Ligra: *A Lightweight Graph Processing Framework for Shared Memory*, Principles and Practice of Parallel Programming, 2013.

J. Shun, L. Dhulipala and G. E. Blelloch. Smaller and Faster: *Parallel Processing of Compressed Graphs with Ligra+*, Data Compression Conference, 2015.

Code: https://github.com/jshun/ligra/