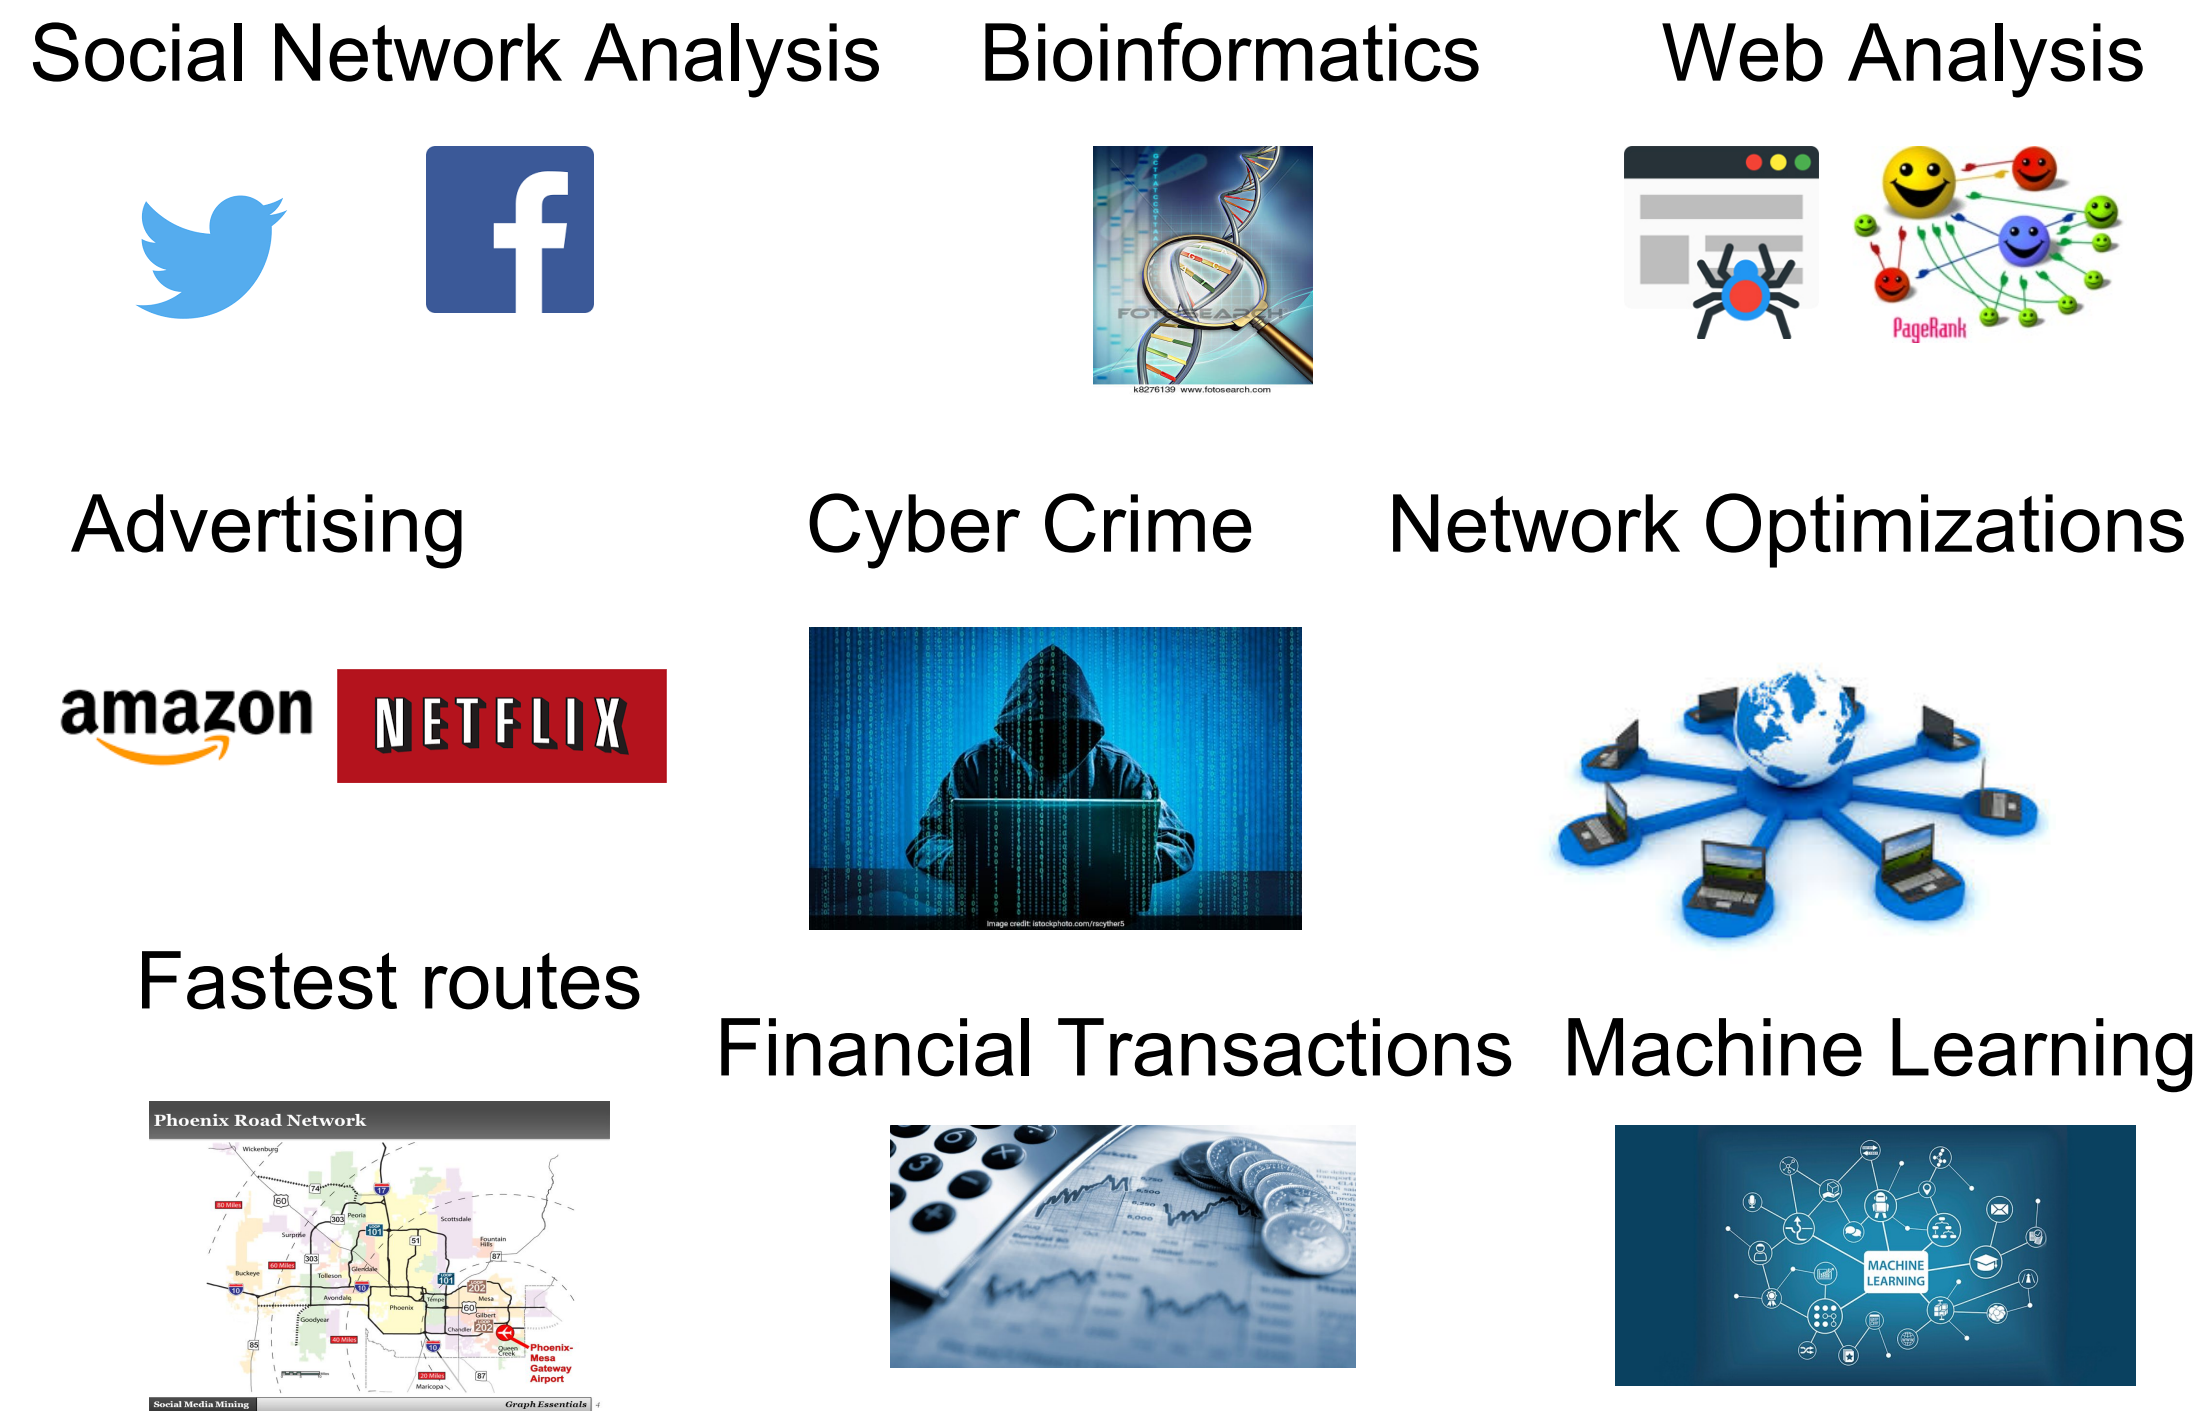


# D-SAGA: A Domain Specific Architecture for Fast Processing of Large Graph Applications


Presented by: Rashmi Agrawal  
Advisor: Prof. Michel Kinsy

## Graph Applications



## Properties of Real-time Graphs

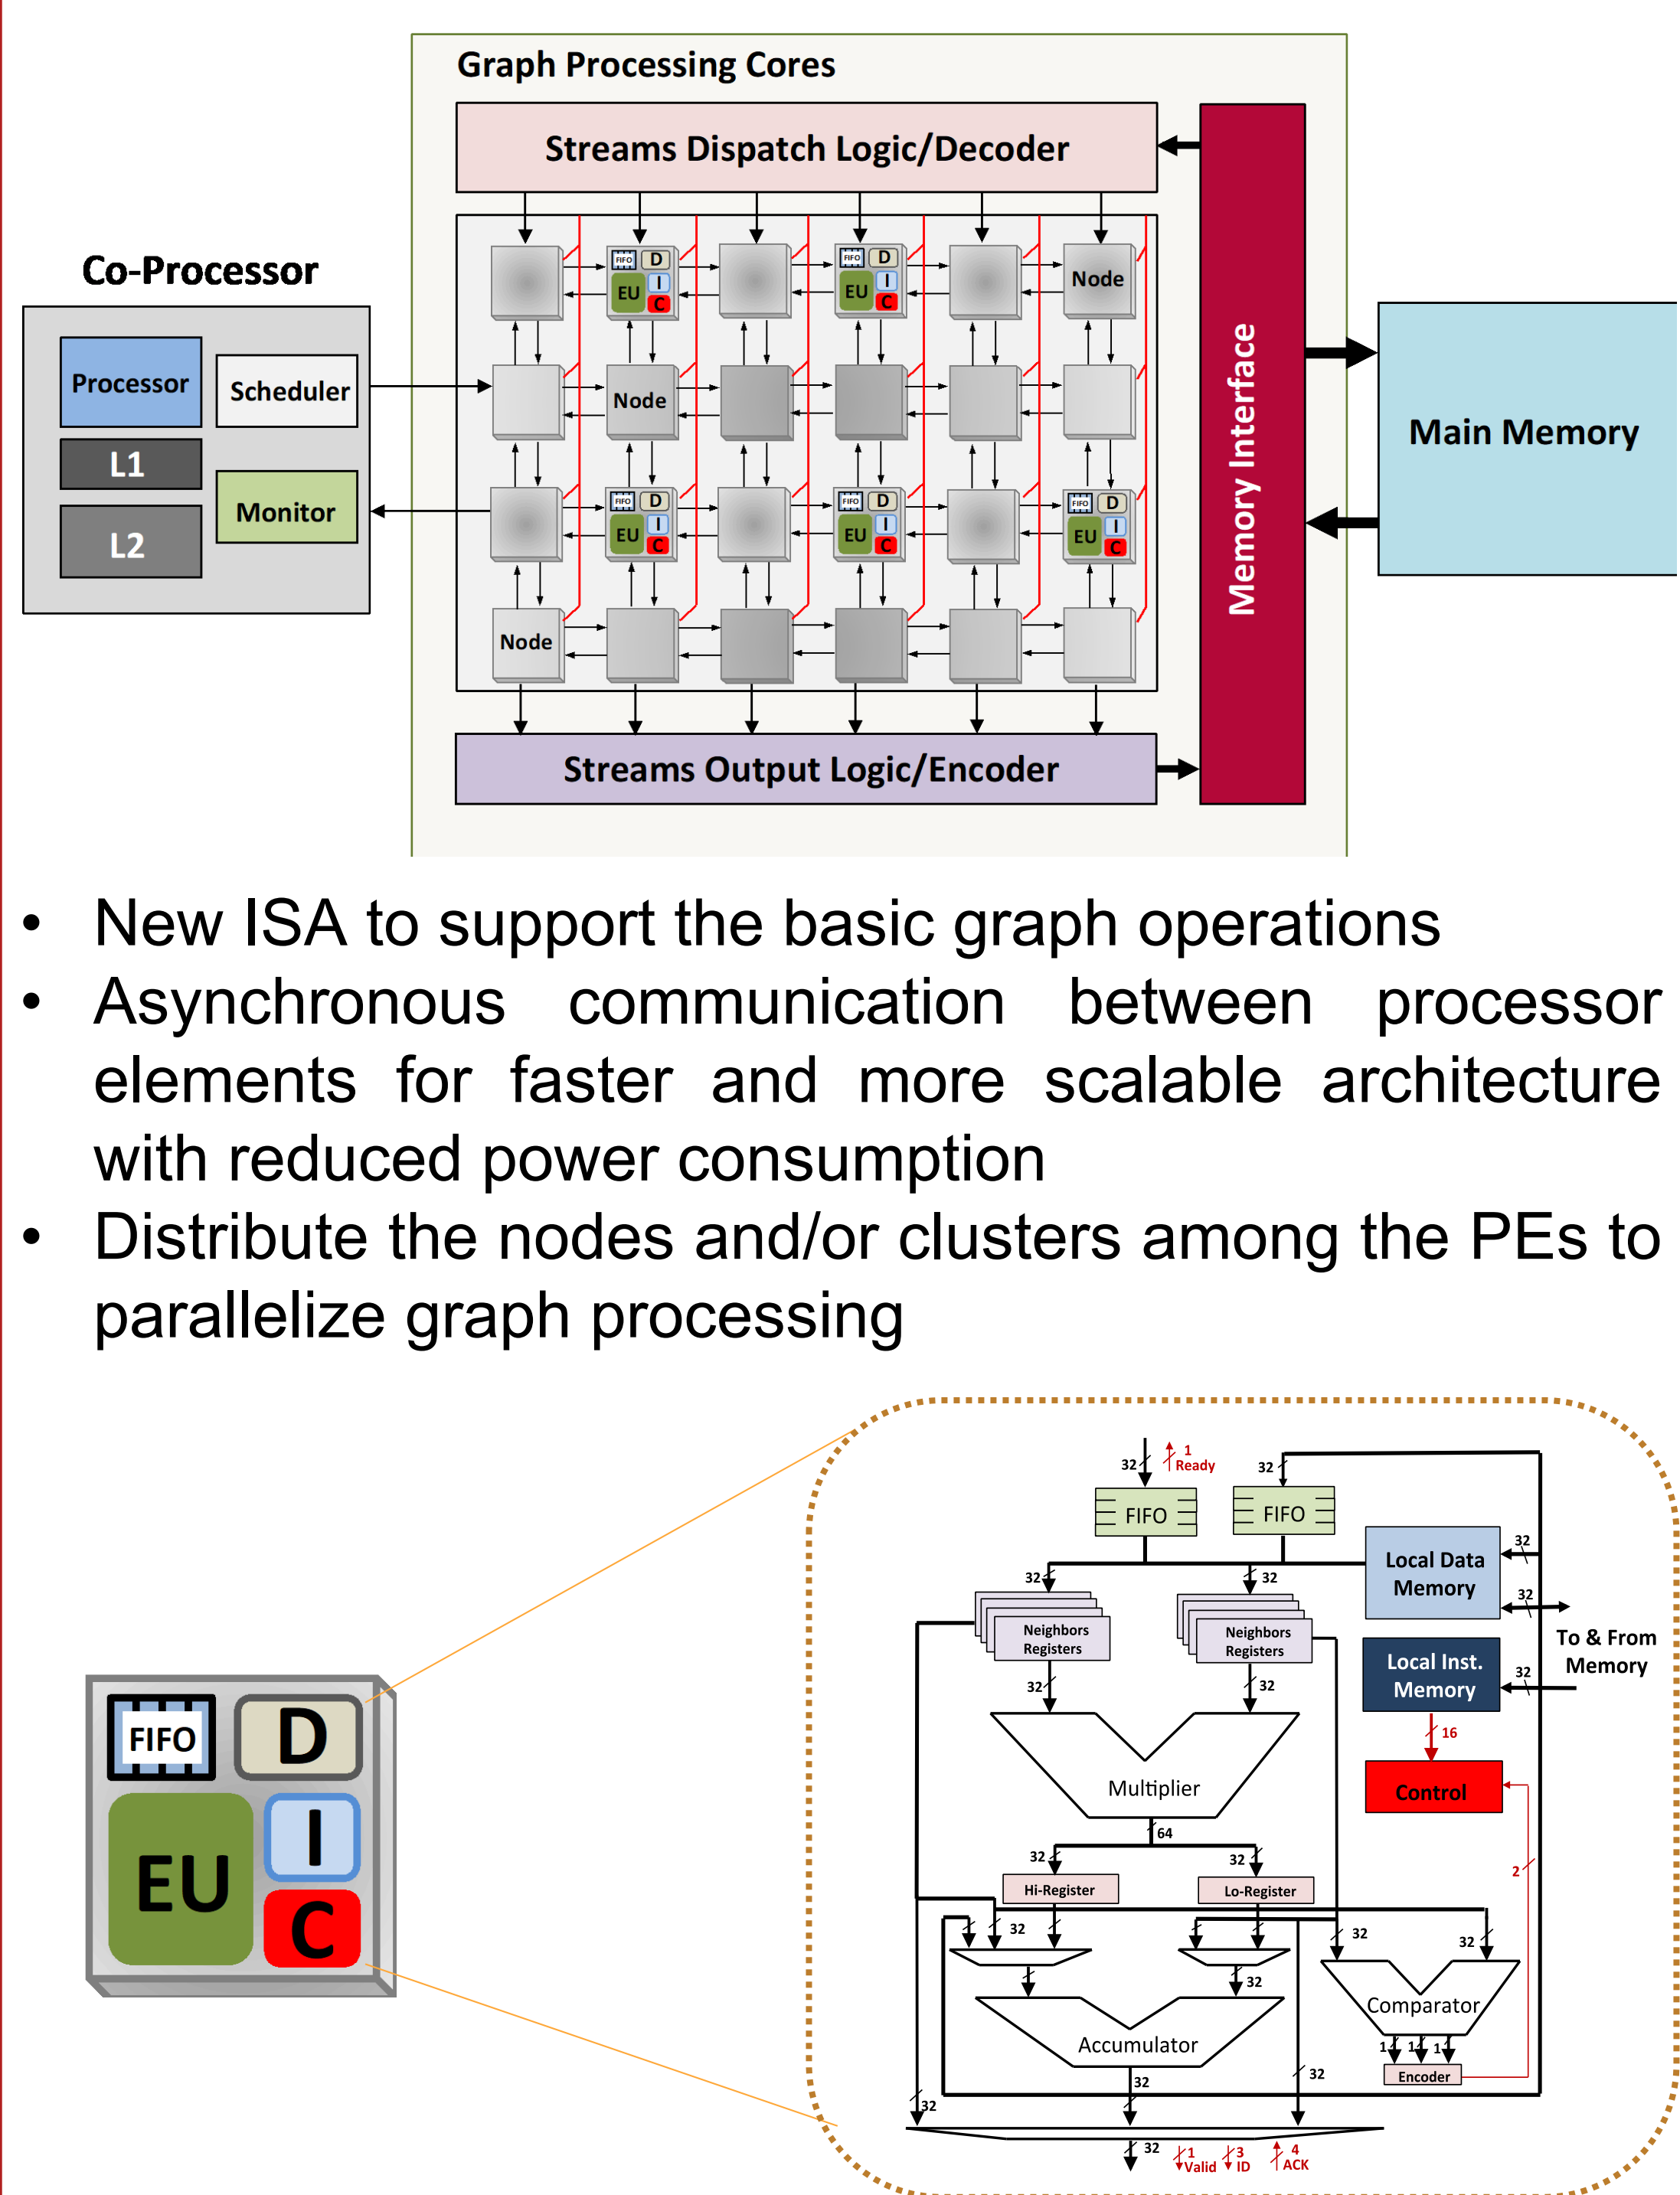
- Large graphs with billions of vertices and edges
- Have rich metadata
- Power law distribution
  - Sparse
  - Highly irregular



## Challenges in Graph Processing

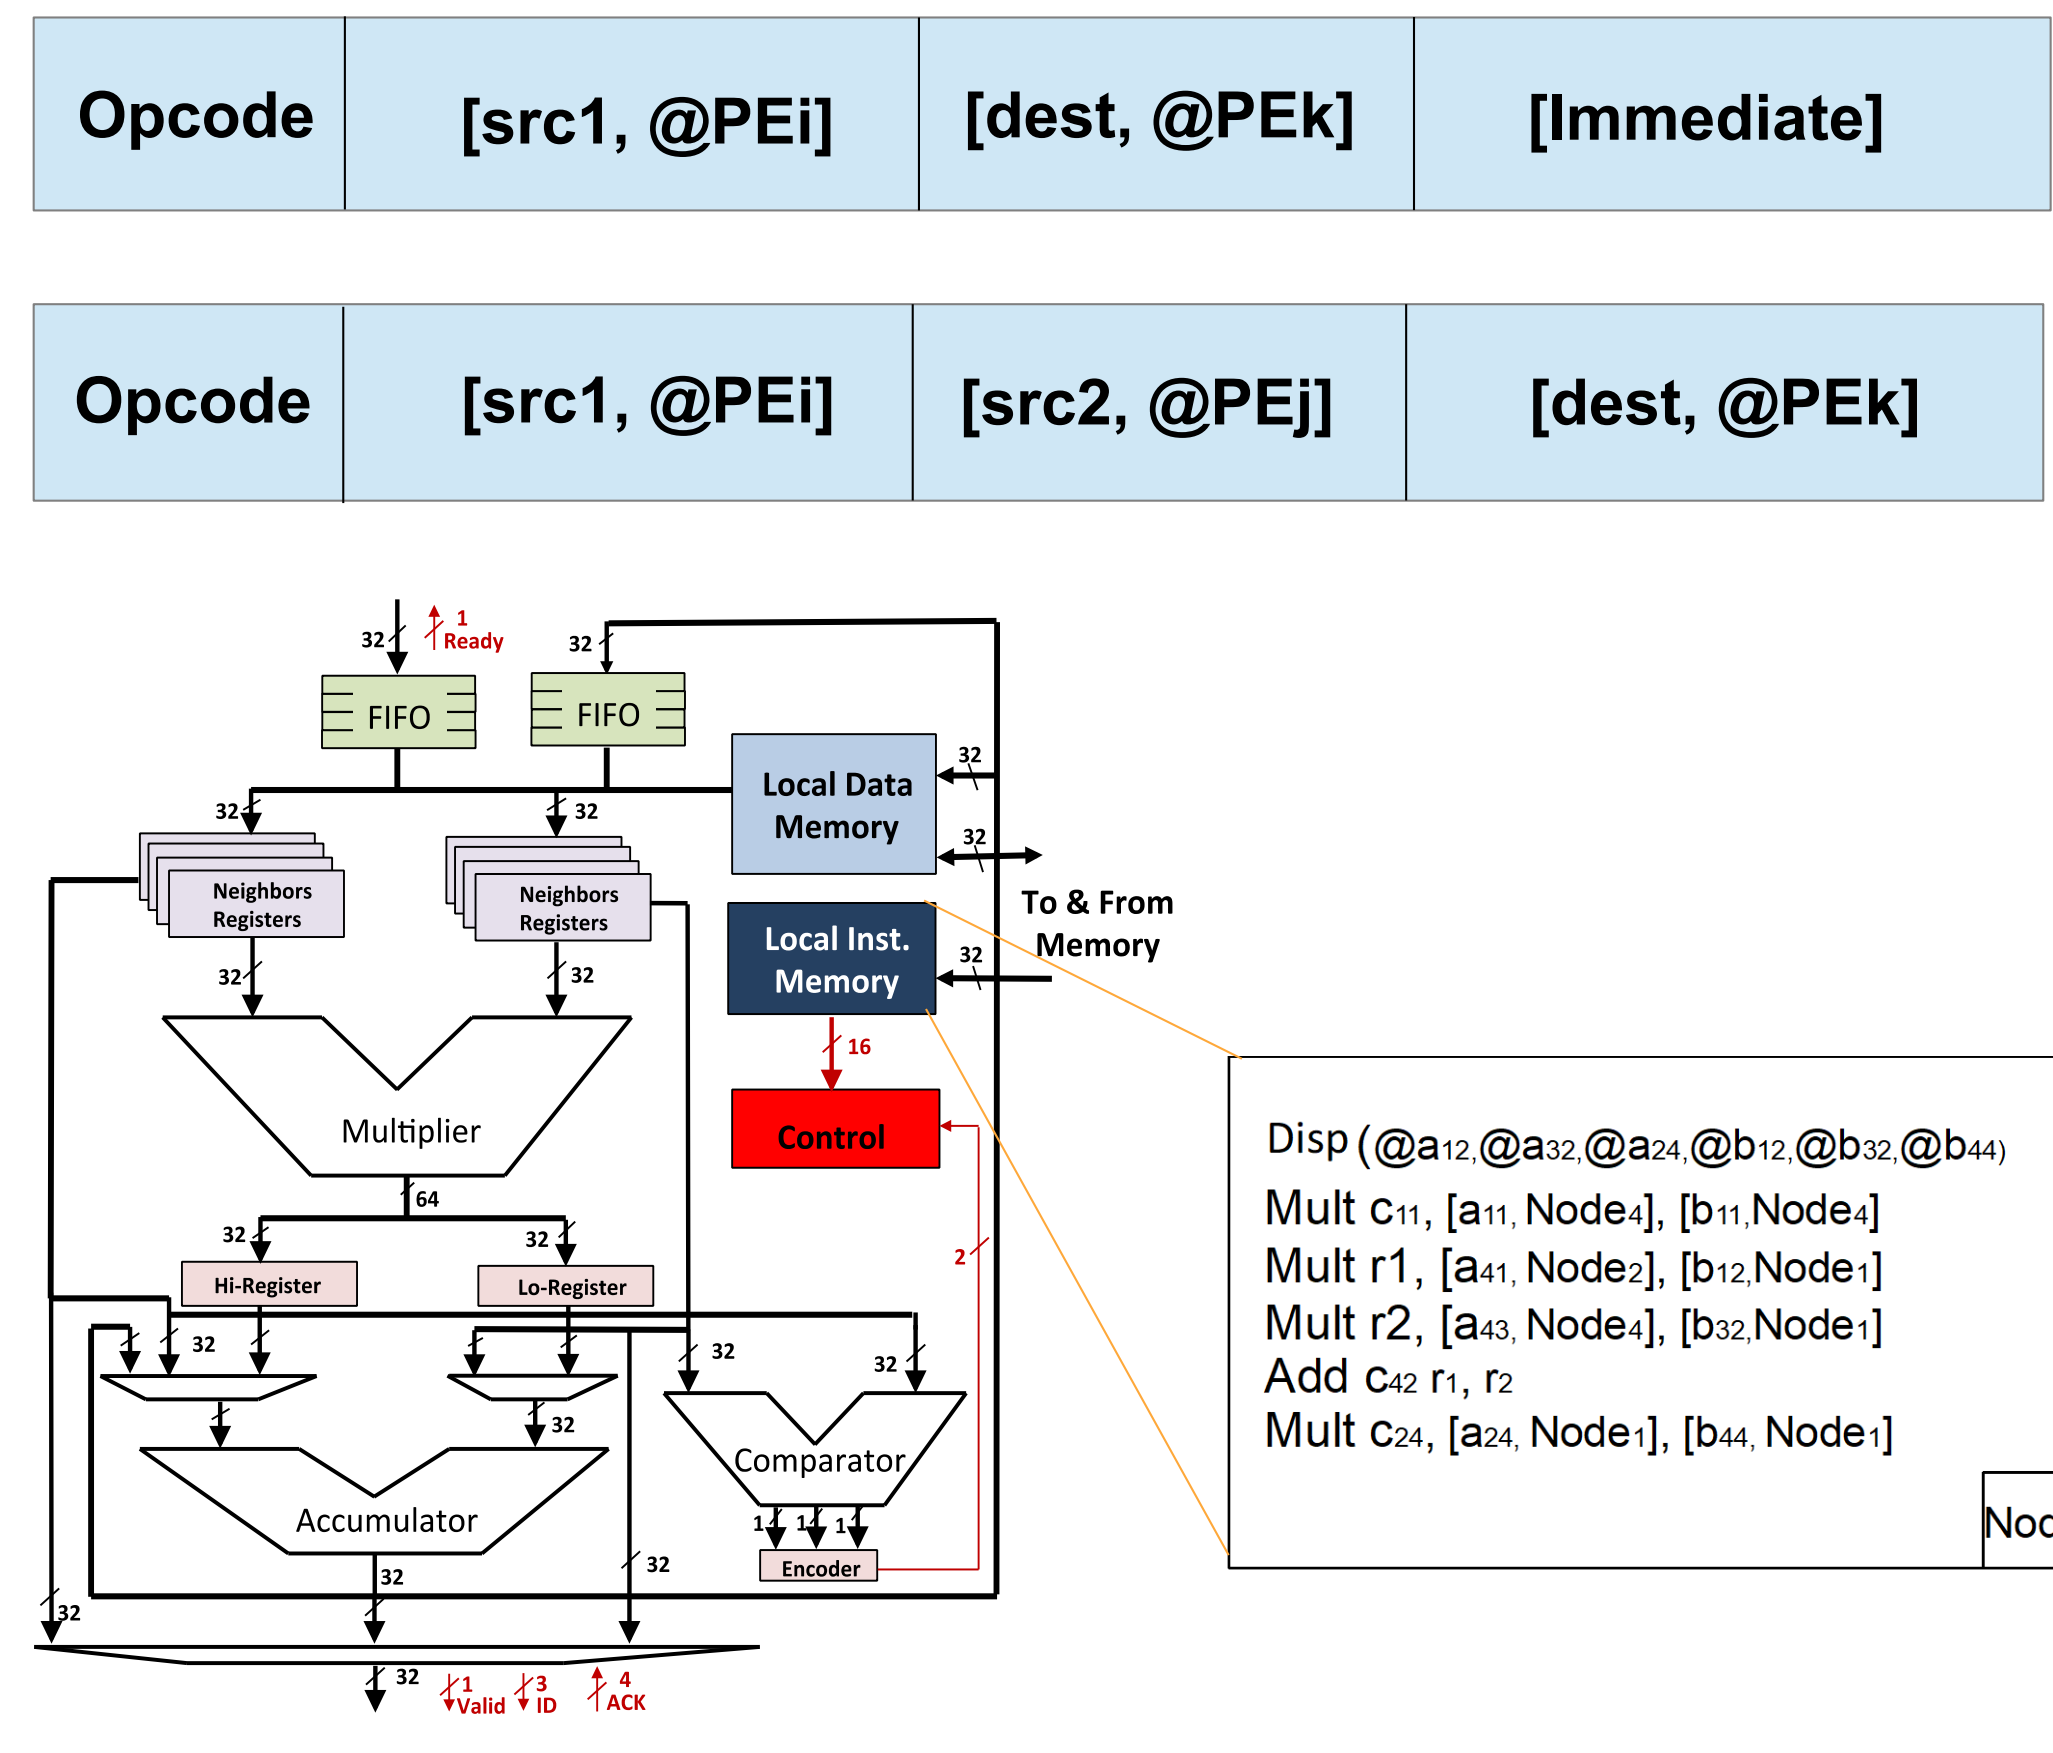
- Scattered memory accesses lead to:
  - Poor cache performance
  - Poor prefetching
  - Poor branch prediction
- Memory intensive operations lead to:
  - Low compute to memory ratio
  - High memory access latency
  - Low memory bandwidth utilization
- Use of Hyperthreading leads to:
  - Inefficient load balancing
  - Synchronization issues
  - Communication latency
- Inefficient graph partitioning is an issue
- Performance tuning required per input/algorithm basis

## Proposed Architecture



- New ISA to support the basic graph operations
- Asynchronous communication between processor elements for faster and more scalable architecture with reduced power consumption
- Distribute the nodes and/or clusters among the PEs to parallelize graph processing

## Graph Instructions



## Experimental Setup


Application	Number of Vertices	Number Of Edges	Average Degree
Road Networks California (CA)	1,965,206	2,766,607	1.41
Social Networks Facebook (FB)	2,937,612	41,919,708	14.3
Social Networks Livejournal (LJ)	4,847,571	85,702,475	17.6

### Algorithms:

- Single Source Shortest Path (SSSP)
- Breadth First Search (BFS)
- Depth First Search (DFS)
- PageRank (PR)
- Triangle Counting (TC)
- Connected Components (CC)



## Graph Storage

- We exploit spatial locality when storing graphs
- Edges/Nodes within the same cluster are stored on the same device
- Clauset-Newman-Moore cluster detection algorithm for very large graphs
- Growing neural gas network algorithm for midsize graphs





## Results

### Graph Benchmarks



### PARSEC and SPLASH-2 Benchmarks



## Illustrative Example: Mapping

