SAND2014-19724C

Exceptional service in the national interest

Sandia National Laboratories



Using Vertex-Centric Programming Platforms to Implement SPARQL Queries on Large Graphs

Eric Goodman[1,2], Dirk Grunwald[2]
1 – Sandia National Laboratories, 2 – University of Colorado at Boulder

U.S. DEPARTMENT OF ENERGY

National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Main Idea

Sandia National Laboratories

Semantic Web

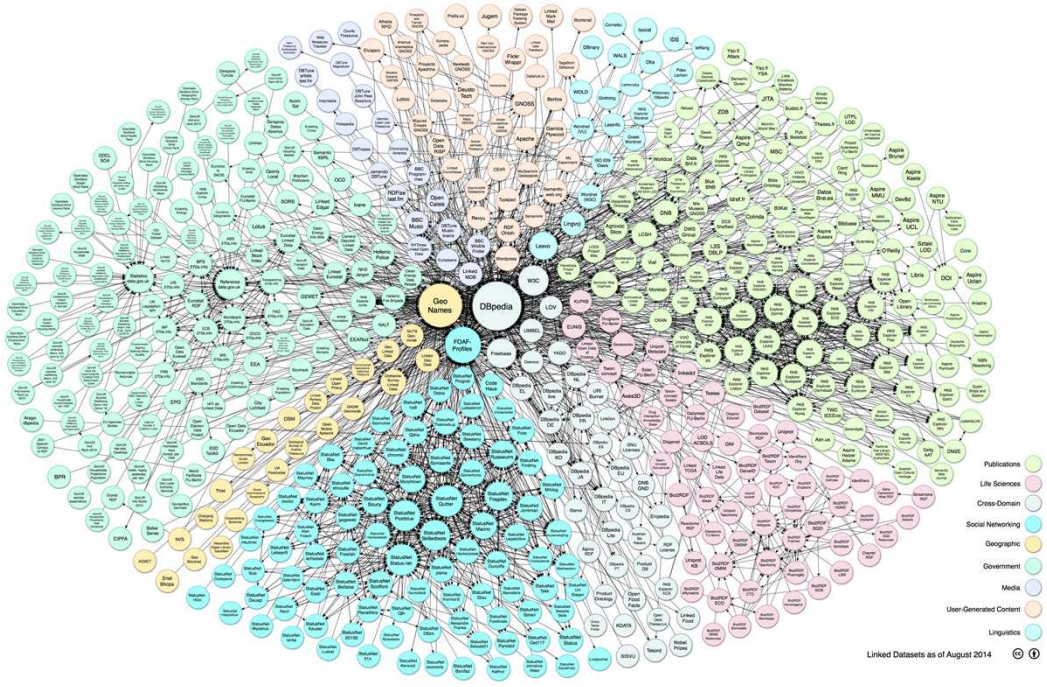

Linking Open Data cloud diagram 2014, by Max Schmachtenberg, Christian Bizer, Anja Jentzsch and Richard Cyganiak. http://lod-cloud.net/

Vertex-centric Programming

double **gather**(icontext_type& context, const vertex_type& vertex, edge_type& edge)

void **apply**(icontext_type& context, vertex_type& vertex, const gather_type& total)

edge_dir_type **scatter_edges**(icontext_type& context, const vertex_type& vertex)

What is the Semantic Web?

- Term coined by Berners-Lee, Hendler, and Lassila in 2001
- Data formatted in way such that computers can reason about data much the way we do.

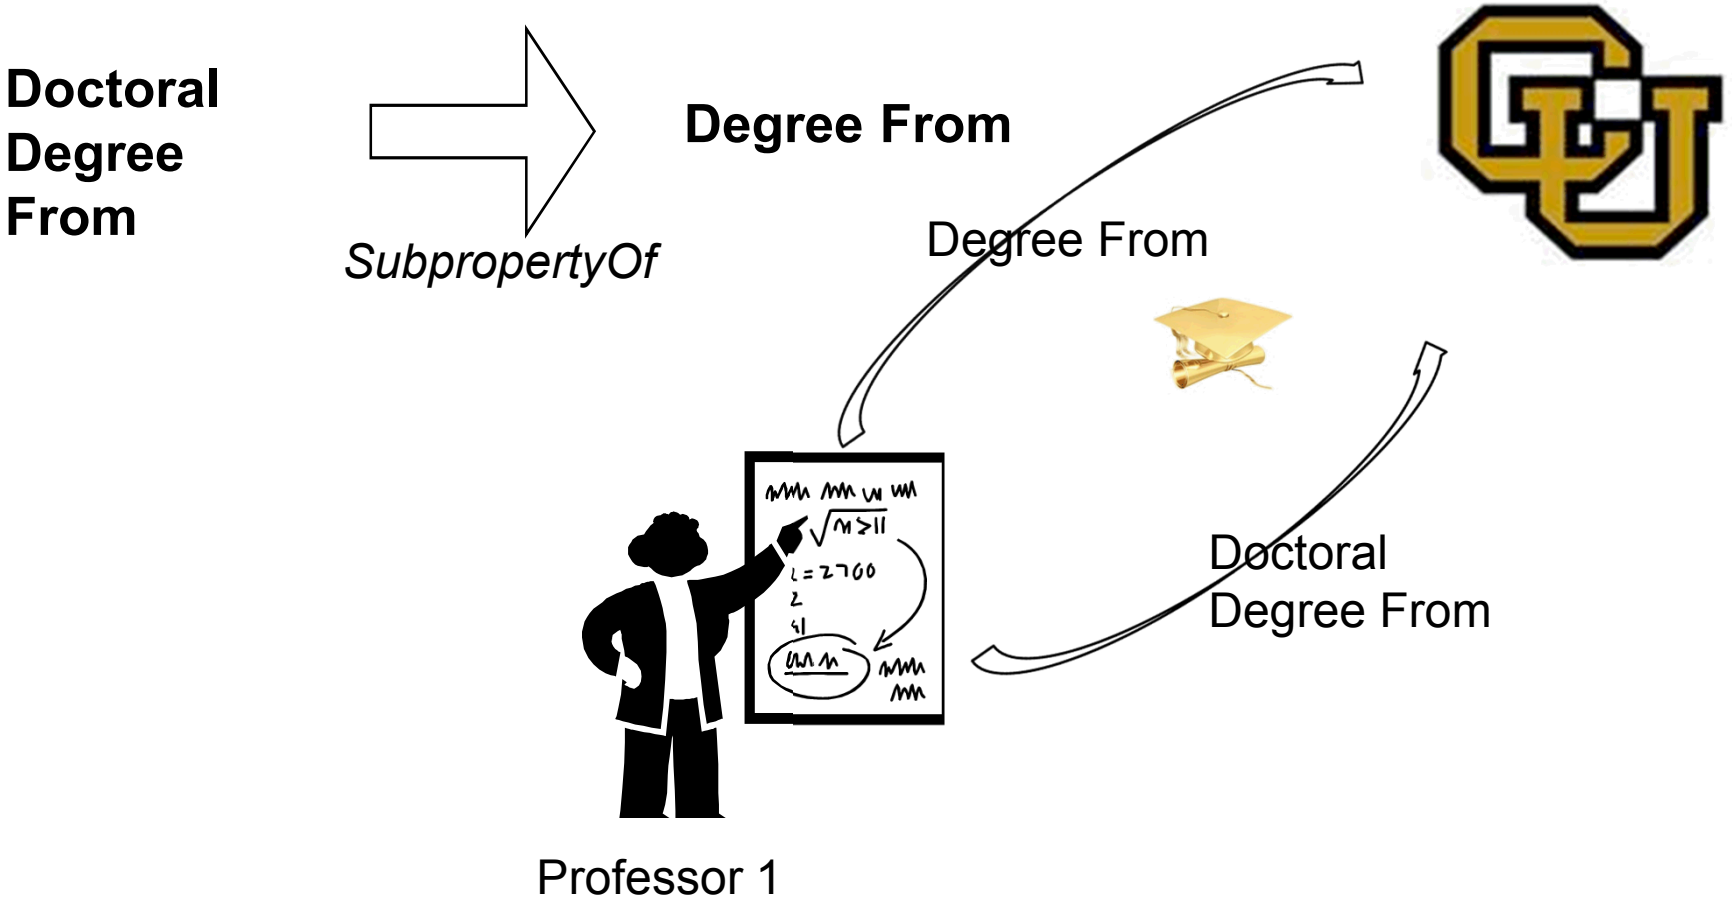

Semantic Web as a Graph

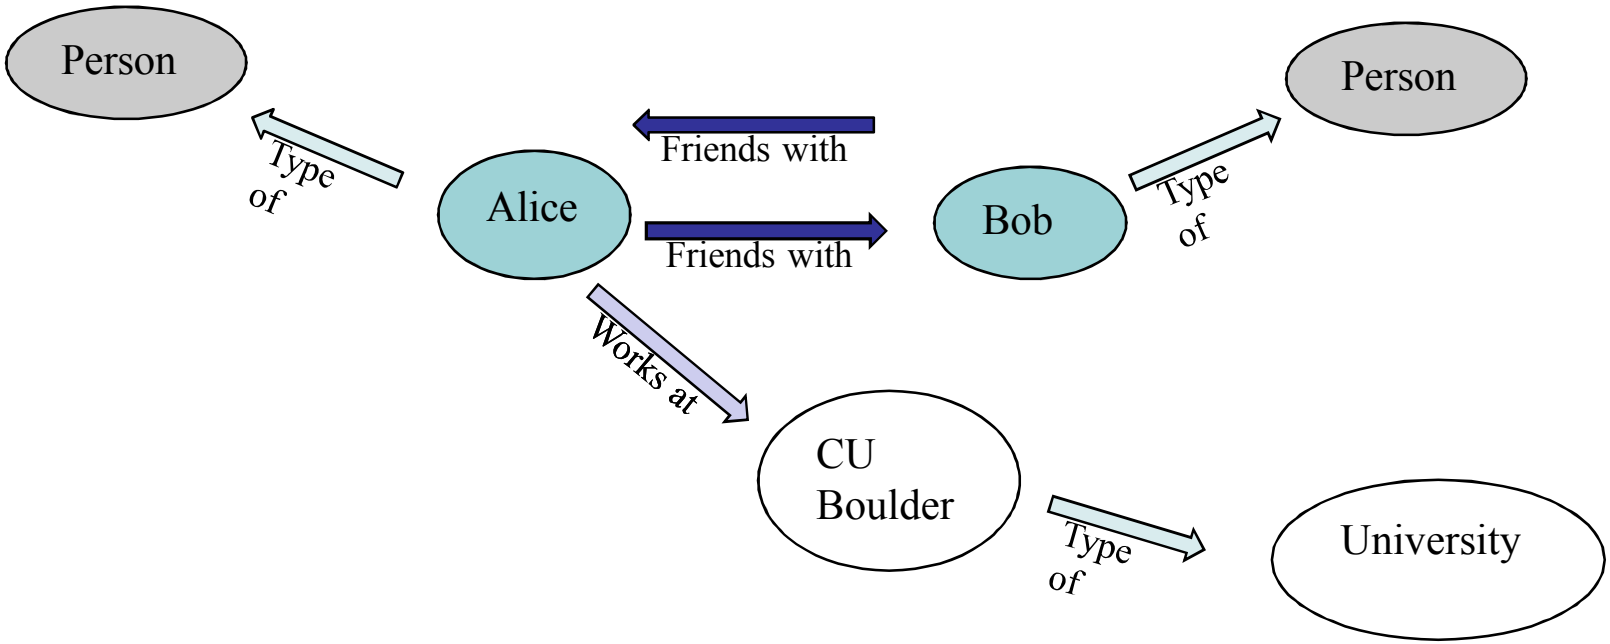

RDF

Sandia National Laboratories

- Resource Description Format
- Format: **Subject Predicate Object**

<http://www.Department9.University3272.edu/UndergraduateStudent354>

<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#emailAddress>

"UndergraduateStudent354@Department9.University3272.edu”

SPARQL

```
SELECT ?X1 ?X2 ?X3
WHERE {
?X1 foaf:knows ?X2 .
?X2 foaf:knows ?X3 .
?X3 foaf:knows ?X1 }
```

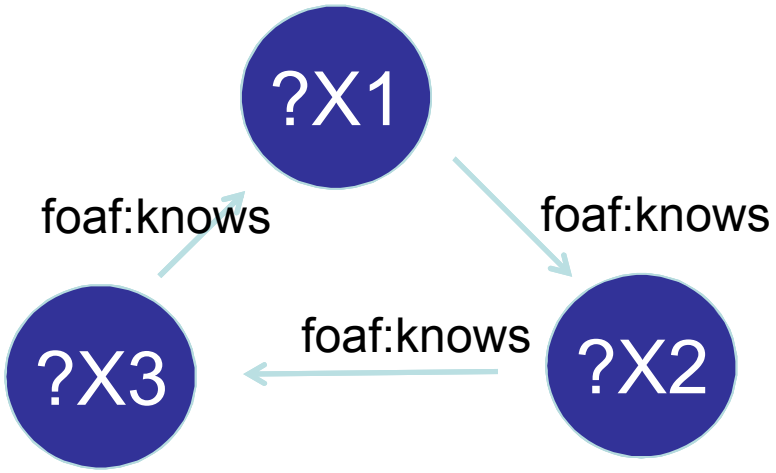

Vertex-Centric Computing

Sandia National Laboratories

- Pregel
- Giraph
- GraphX
- GraphLab
- GraphChi
- Graph Processing System (GPS)

Bulk Synchronous Parallel

Sandia National Laboratories

- Has three phases:

1. Concurrent computation: Each node performs computation independent of all other nodes. Memory access occurs only from the nodes' local memory.
2. Communication: The nodes communicate with each other.
3. Barrier Synchronization: All nodes are stalled until the communication phase is over.

Vertex-centric Computation

Sandia National Laboratories

- For example PageRank:

1. Each vertex takes its current pagerank score, divides by the number of outgoing edges, and then sends that value to all of its neighbors. This is the communication phase of BSP.
2. Progress halts until all vertices have received messages with updated page rank values from their neighbors. This is the synchronization phase of BSP.
3. Each vertex re-computes the pagerank score based upon the updated values from its neighbors. This is the concurrent computation phase of BSP.

General Approach

Sandia National Laboratories

- Take a walk along the subgraph of interest, accumulating results as we go.

The Pattern

- Gather
- Apply
- Scatter

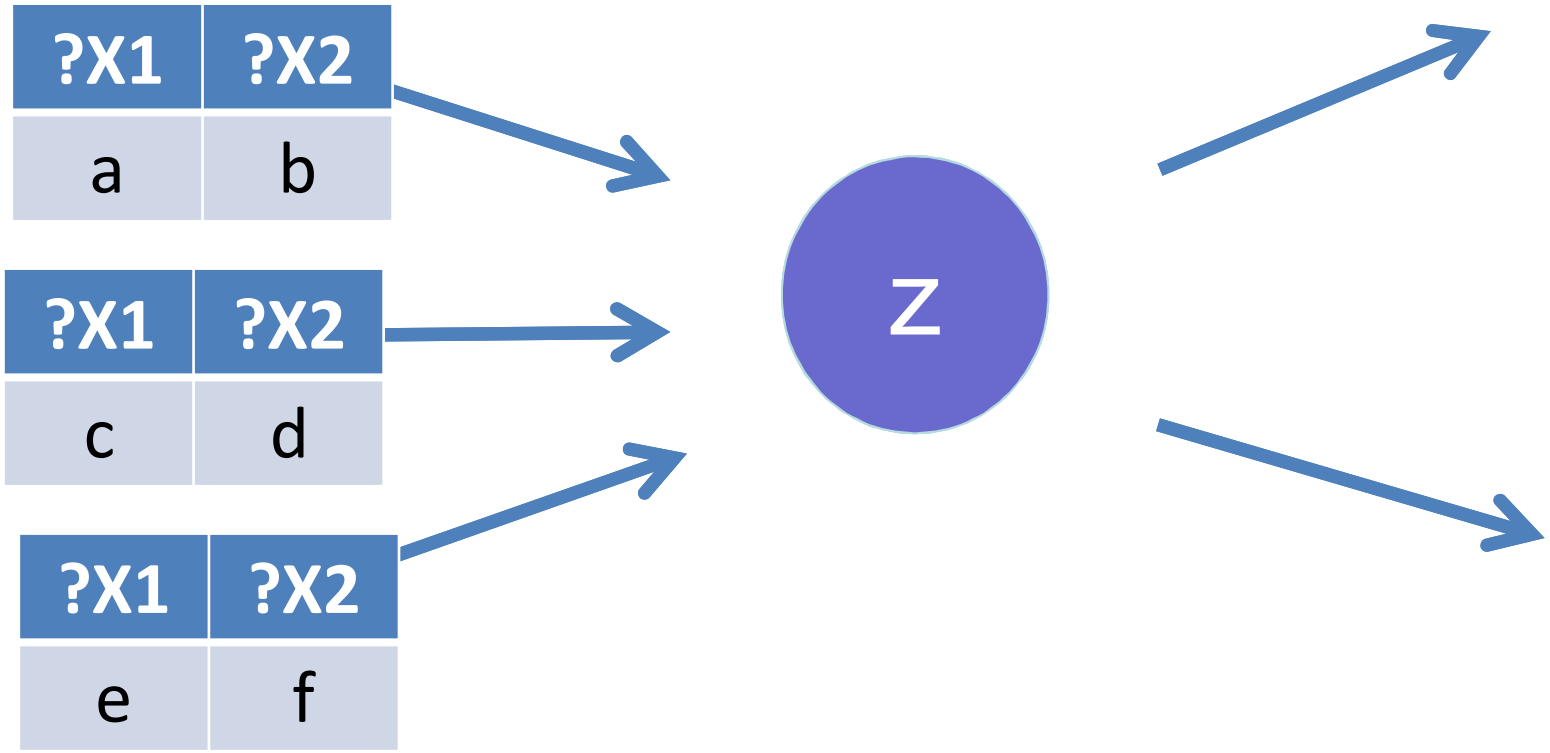

The Pattern

- **Gather**
- Apply
- Scatter

?X1	?X2
a	b
c	d
e	f

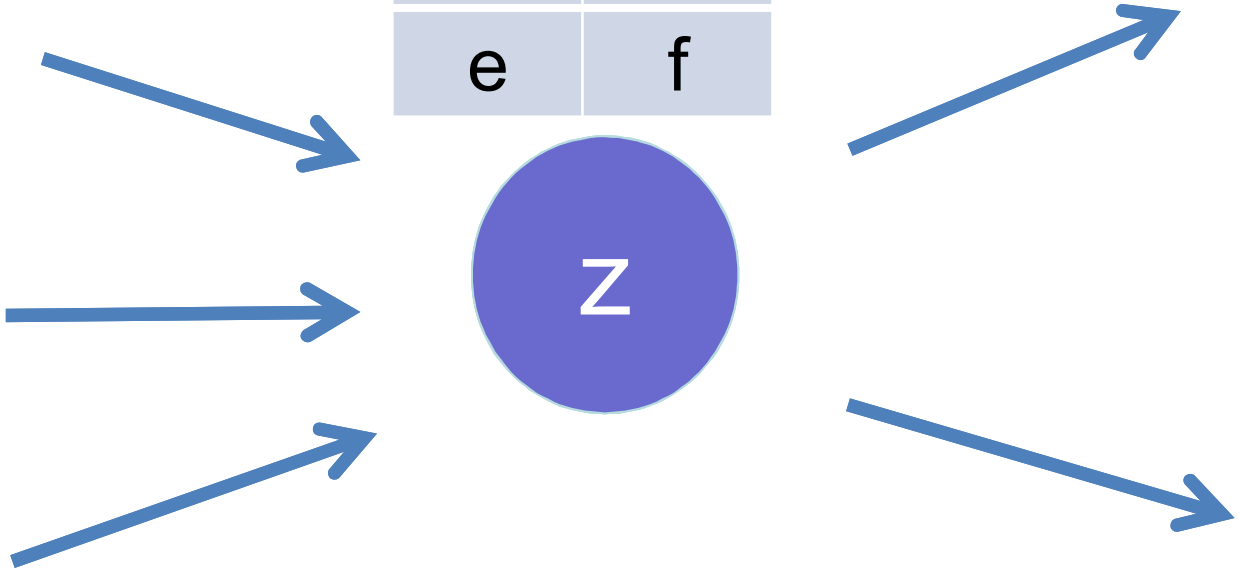

The Pattern

- Gather
- **Apply**
- Scatter

?X1	?X2
a	b
c	d
e	f

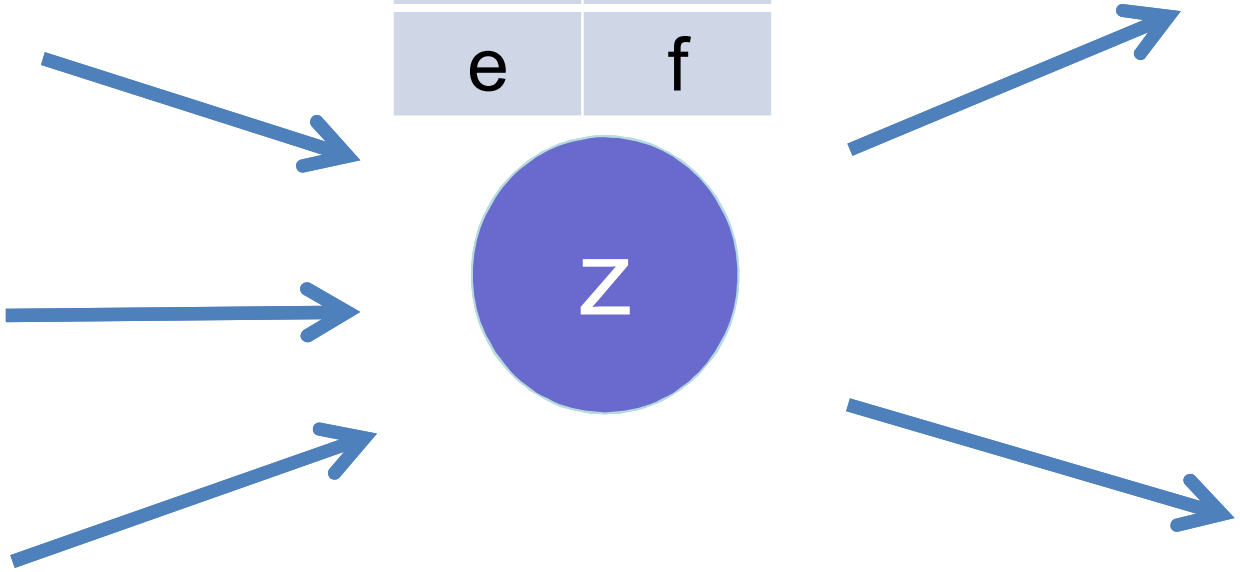

The Pattern

- Gather
- Apply
- **Scatter**

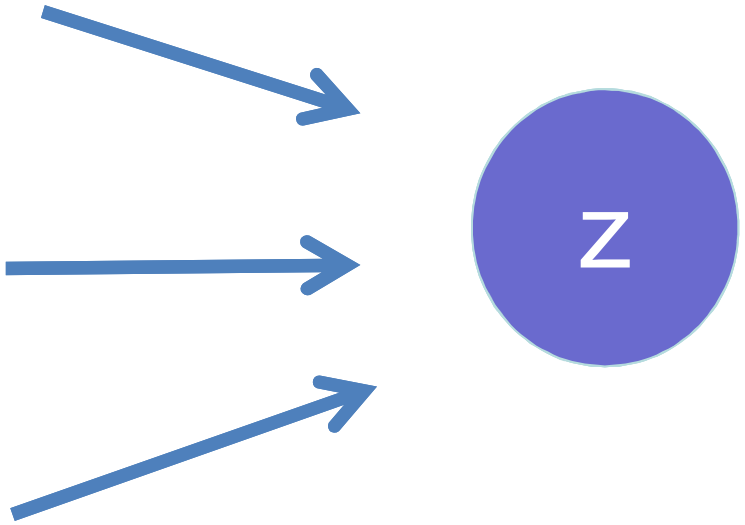

?X1	?X2	?X3
a	b	z
c	d	z
e	f	z

?X1	?X2	?X3
a	b	z
c	d	z
e	f	z

Example Query

Sandia National Laboratories

```
SELECT ?X1 ?X2 ?X3
WHERE {
?X1 P1 ?X2 .
?X2 P2 ?X3 .
?X3 P3 ?X1 }
```

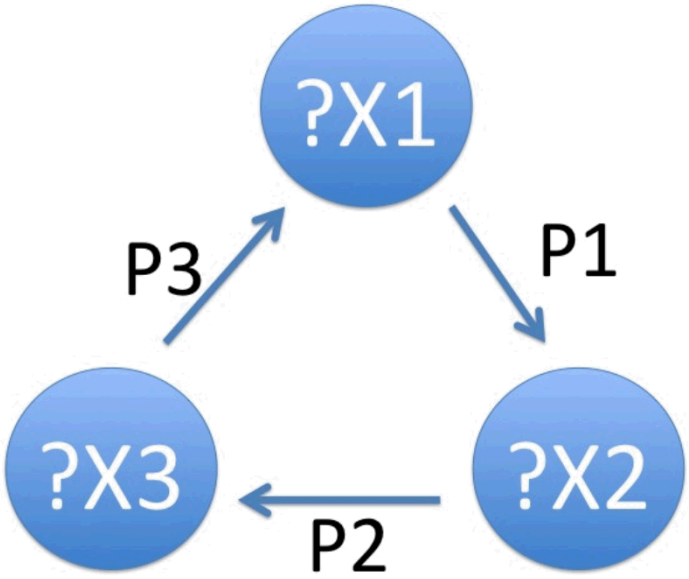

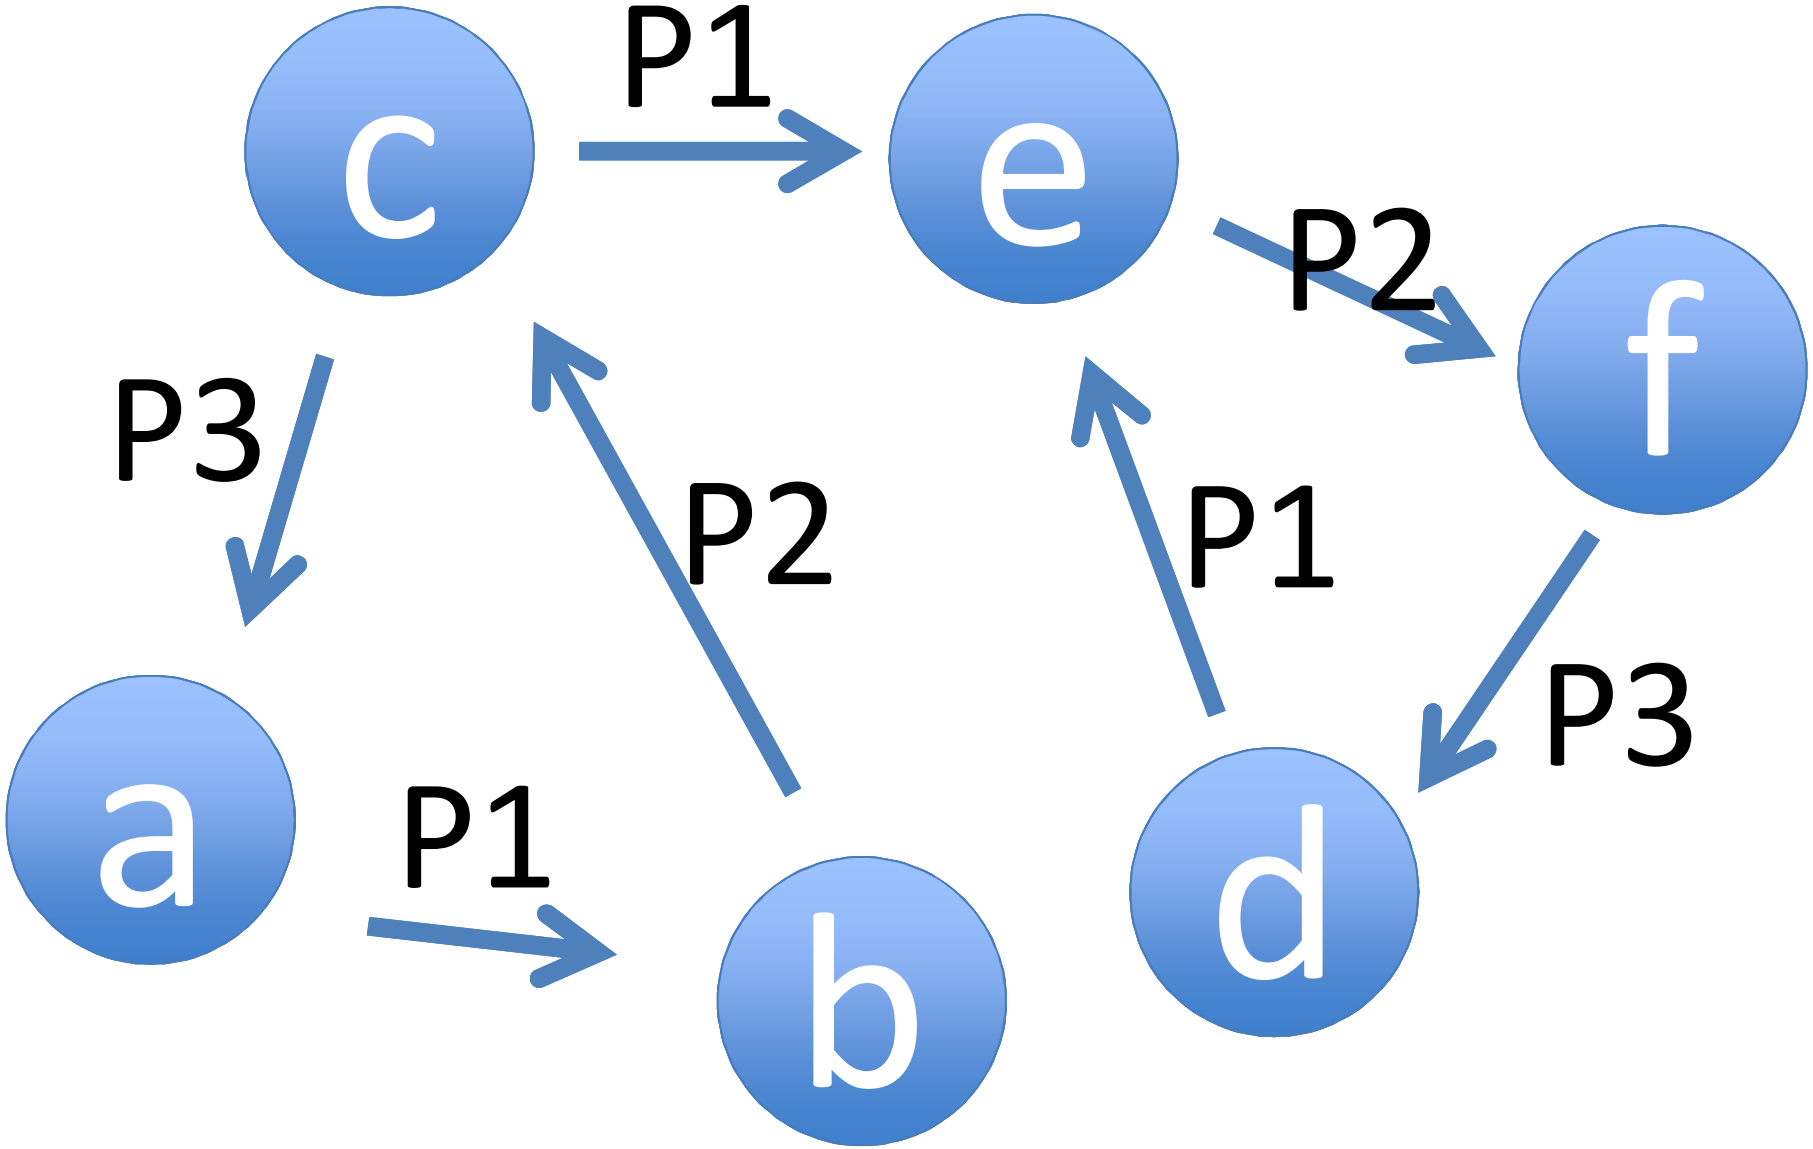
c
P1
e
P2
f
P3
P2
P1
P3
a
P1
b
d

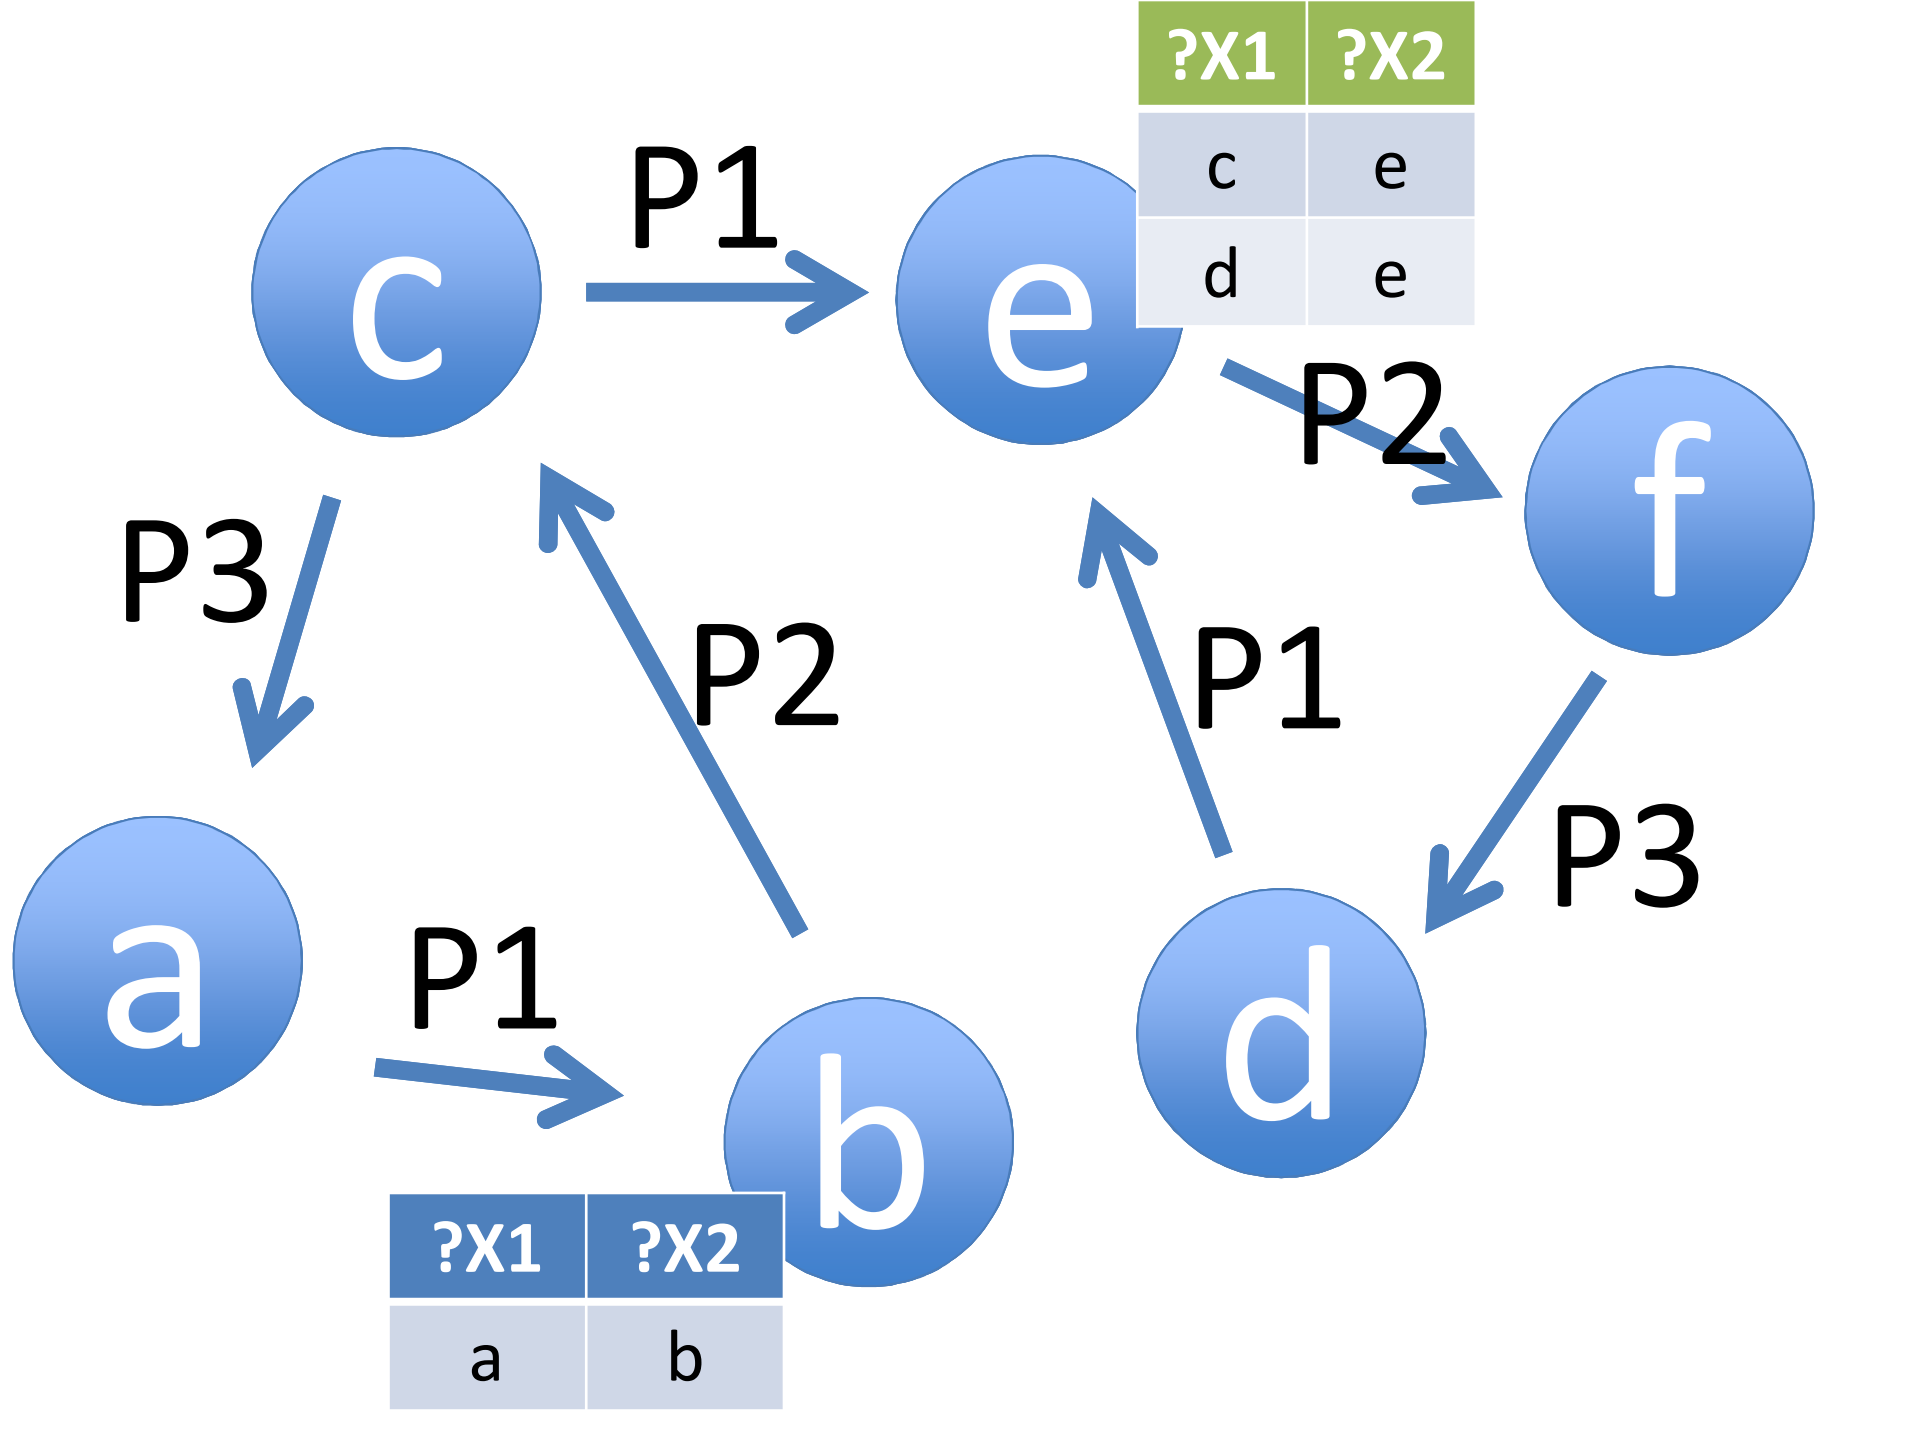
?X1
?X2
c
e
d
e
P1
P2
P3
P2
P1
P3
P1
?X1
?X2
a
b
a
b
c
d
e
f

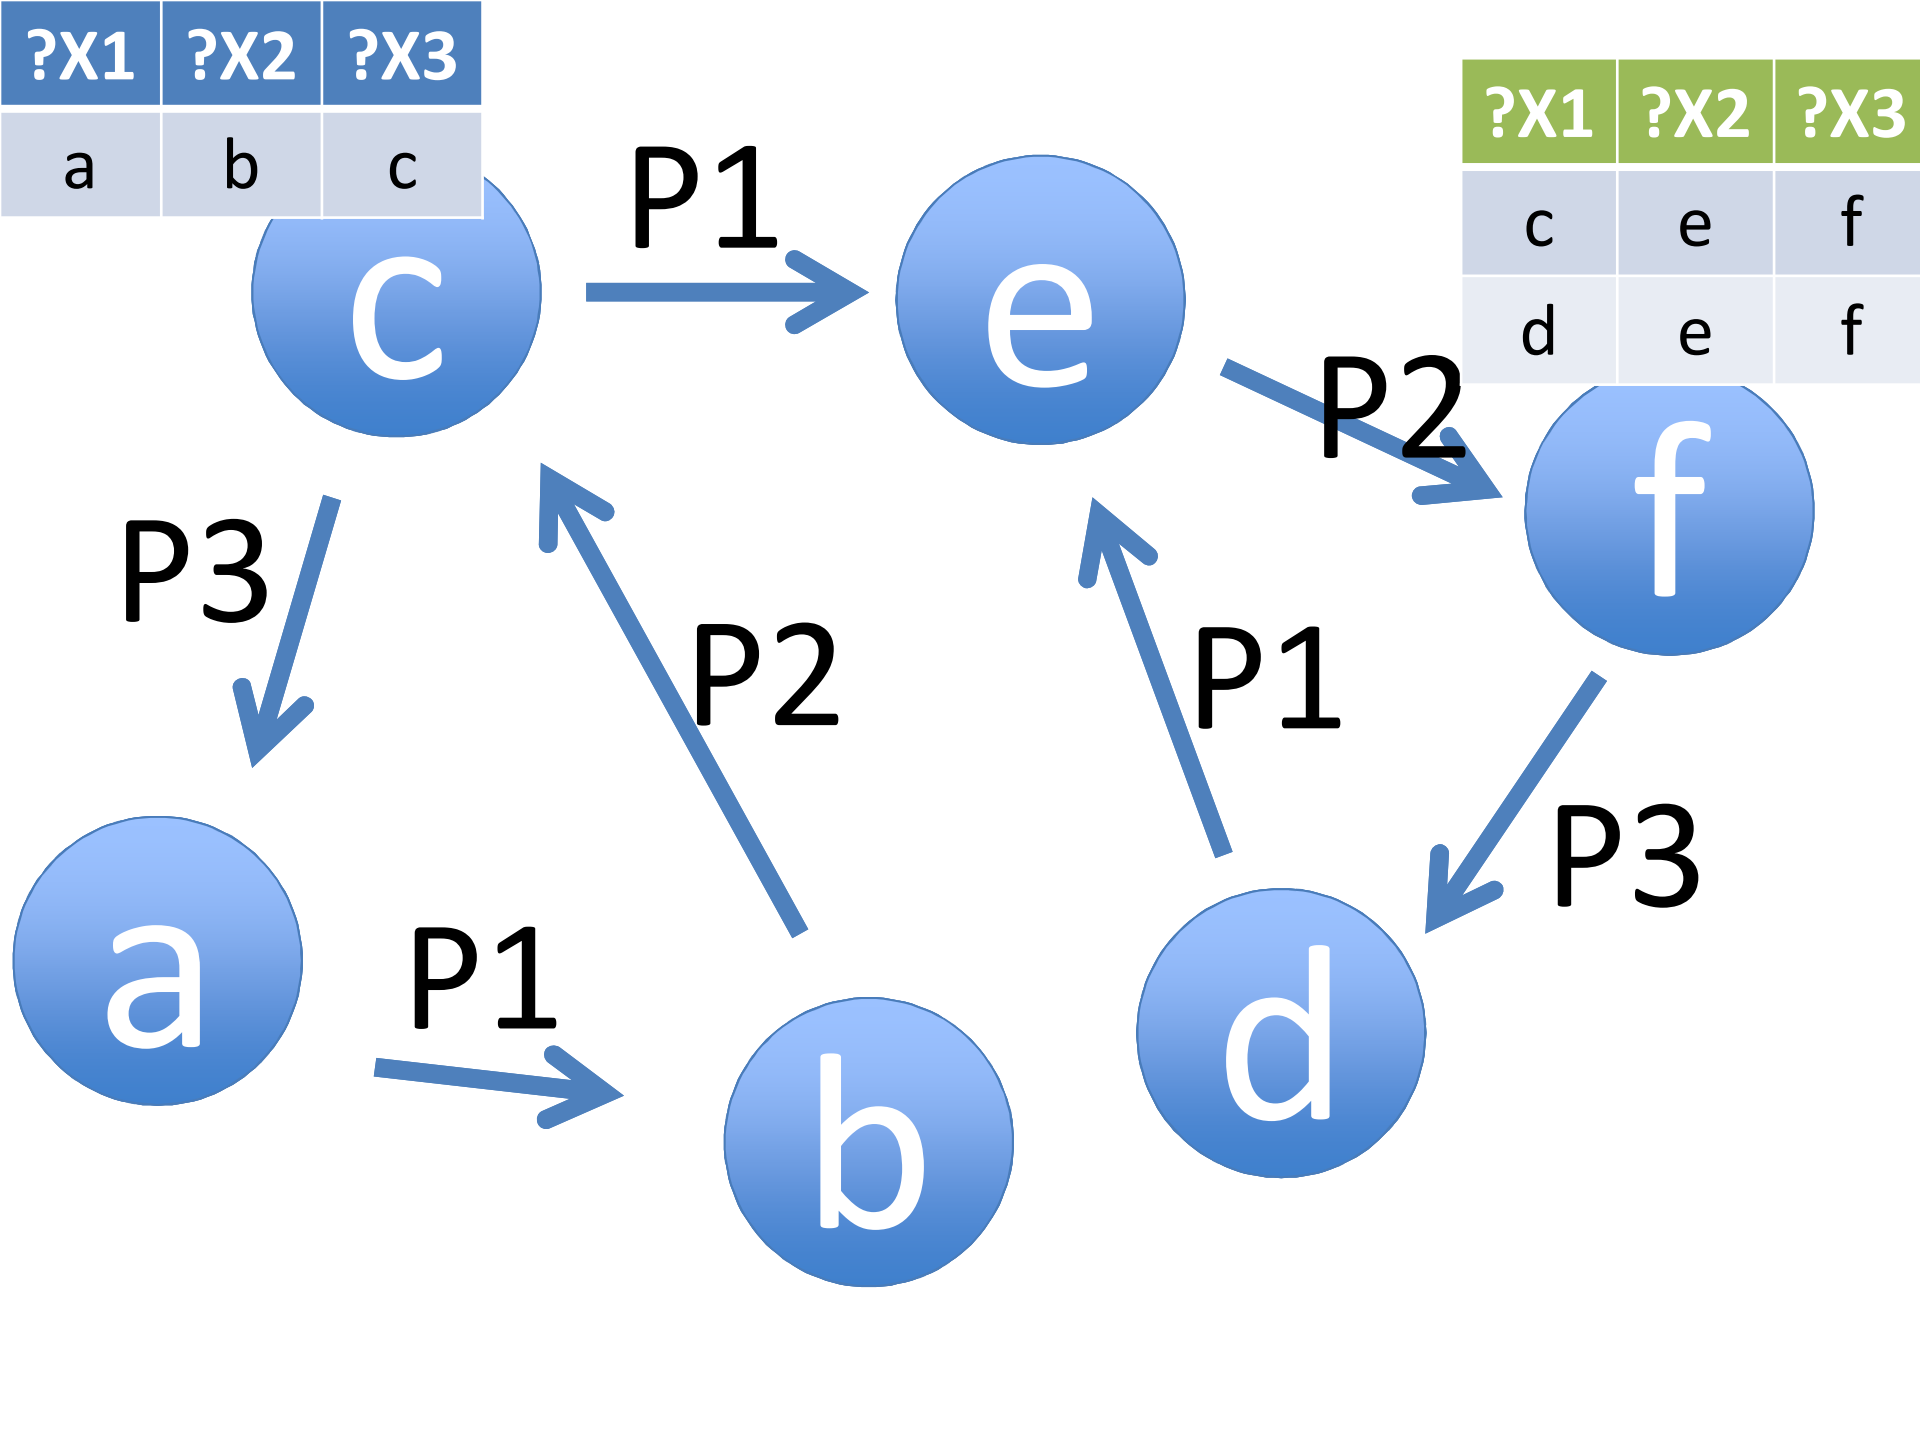
?X1 ?X2 ?X3
a b c
?X1 ?X2 ?X3
c e f
d e f
c
e
f
a
b
d
P1
P2
P3
P2
P1
P3
P1
P3

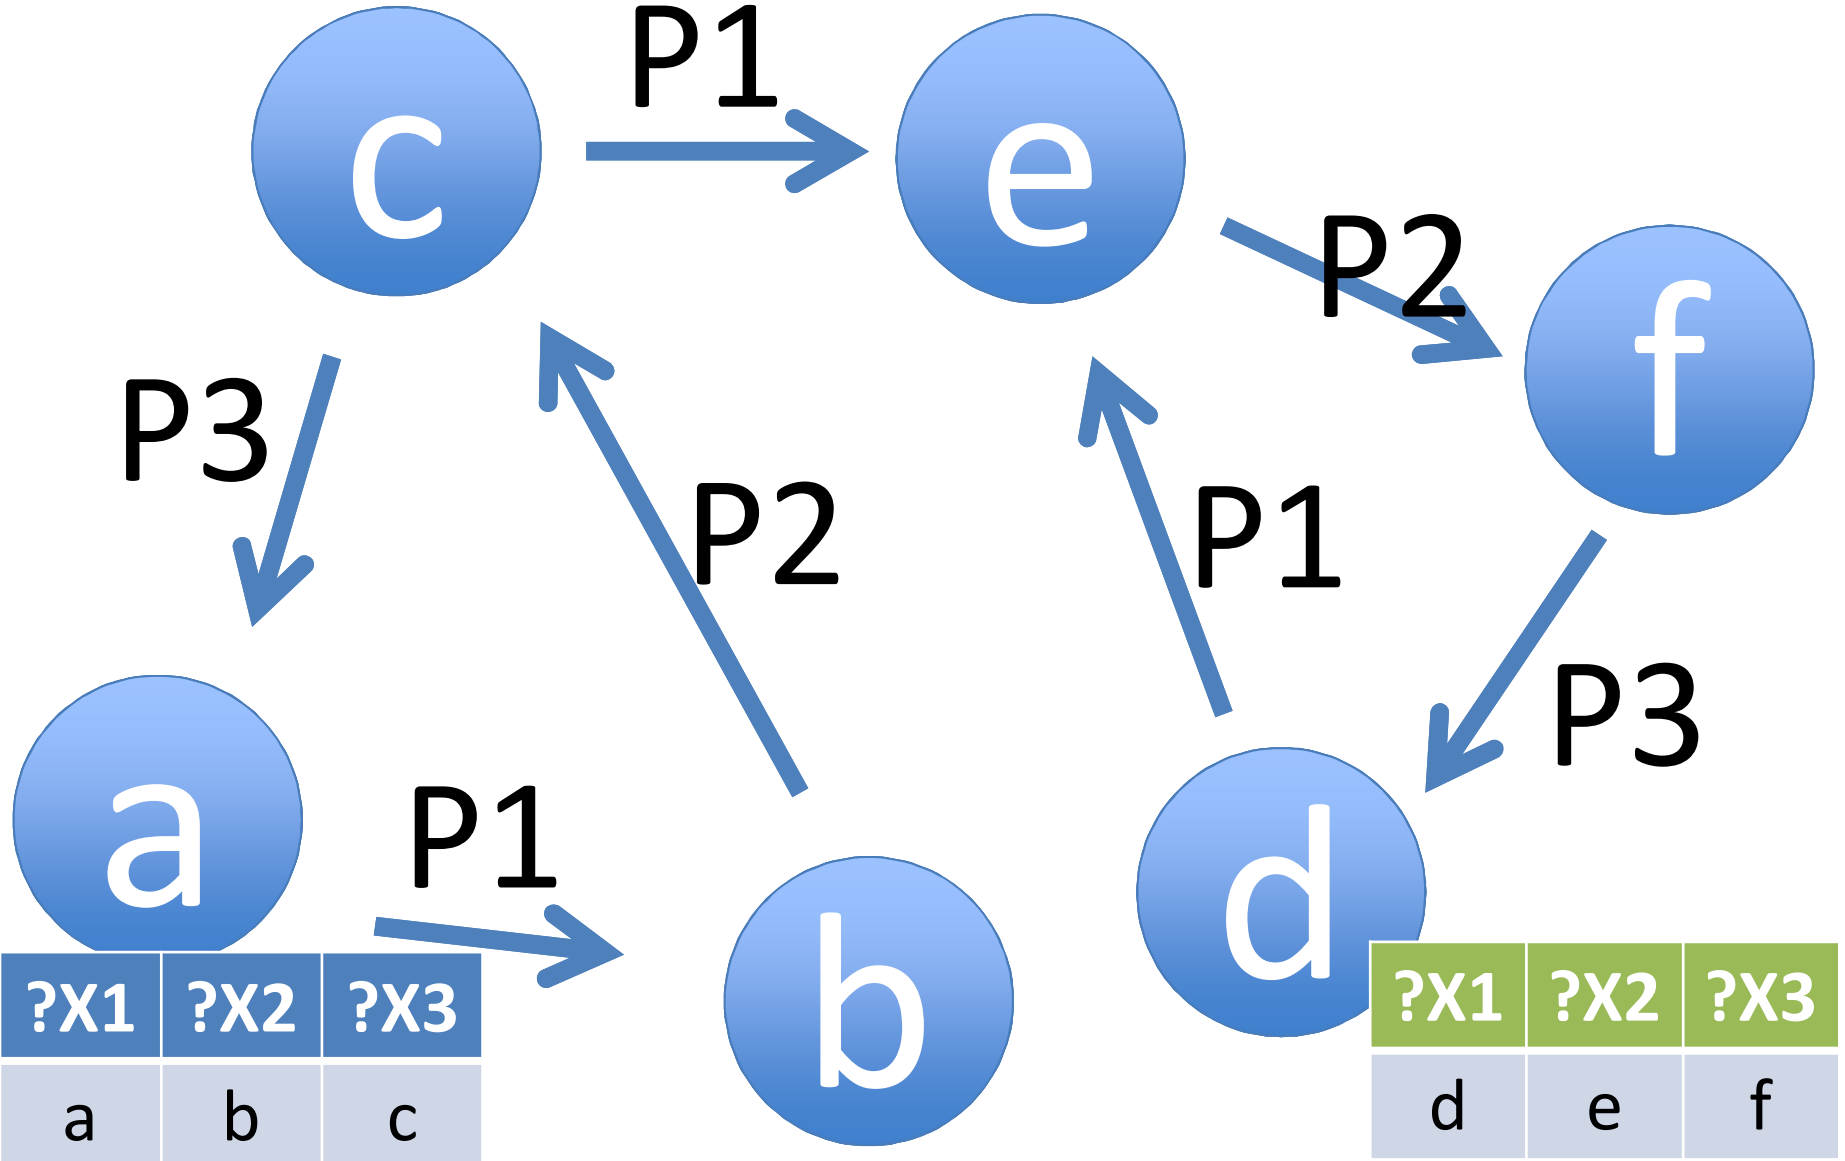
c
P1
e
P2
f
P3
P2
P1
P3
a
P1
b
d
?X1
?X2
?X3
a
b
c
?X1
?X2
?X3
d
e
f

Test Dataset

Sandia National Laboratories

- Lehigh University Benchmark
- Artificial data set
- Creates
  - Universities
  - Students
  - Teachers
  - Departments
  - Classes
  - Etc
- LUBM(8000) – 8000 universities and ~1.35 billion edges

Platform: Plato

Sandia National Laboratories

- 51 nodes with
  - 2x Intel E5-2470 2.30GHz 8-core
  - 96 GB memory
  - 15x 2TB 6G SAS HDD
  - 10Gb Ethernet

Sandia National Laboratories

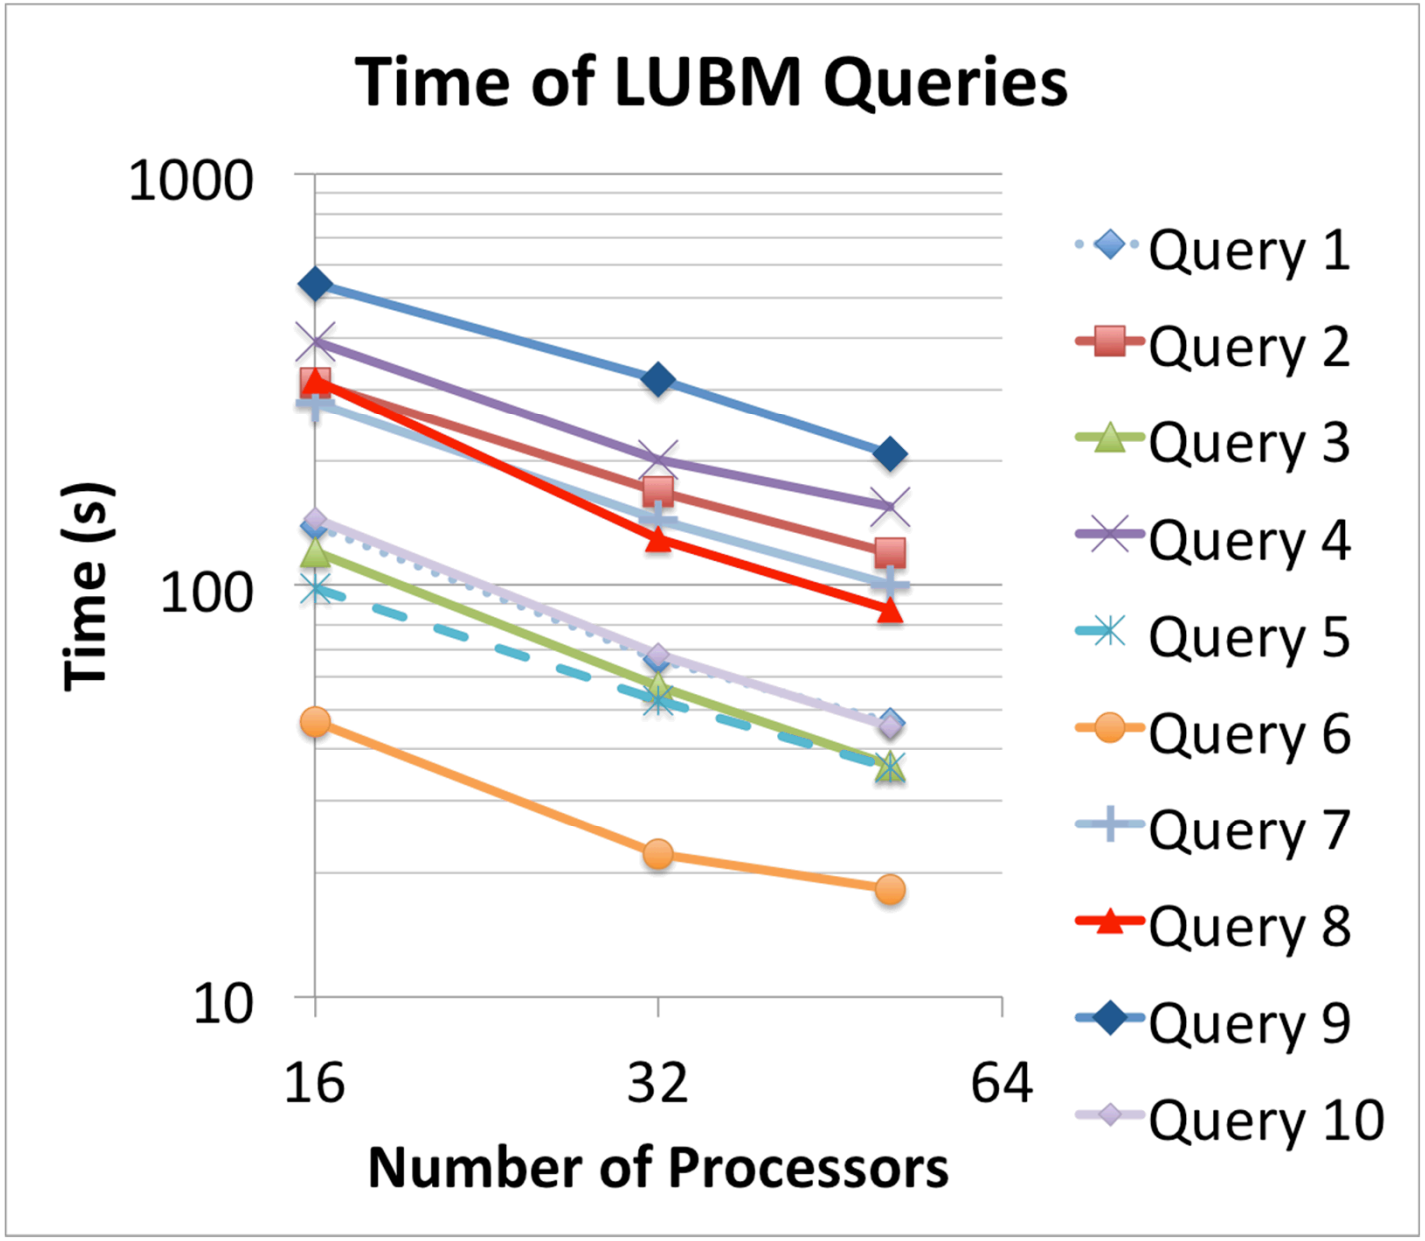

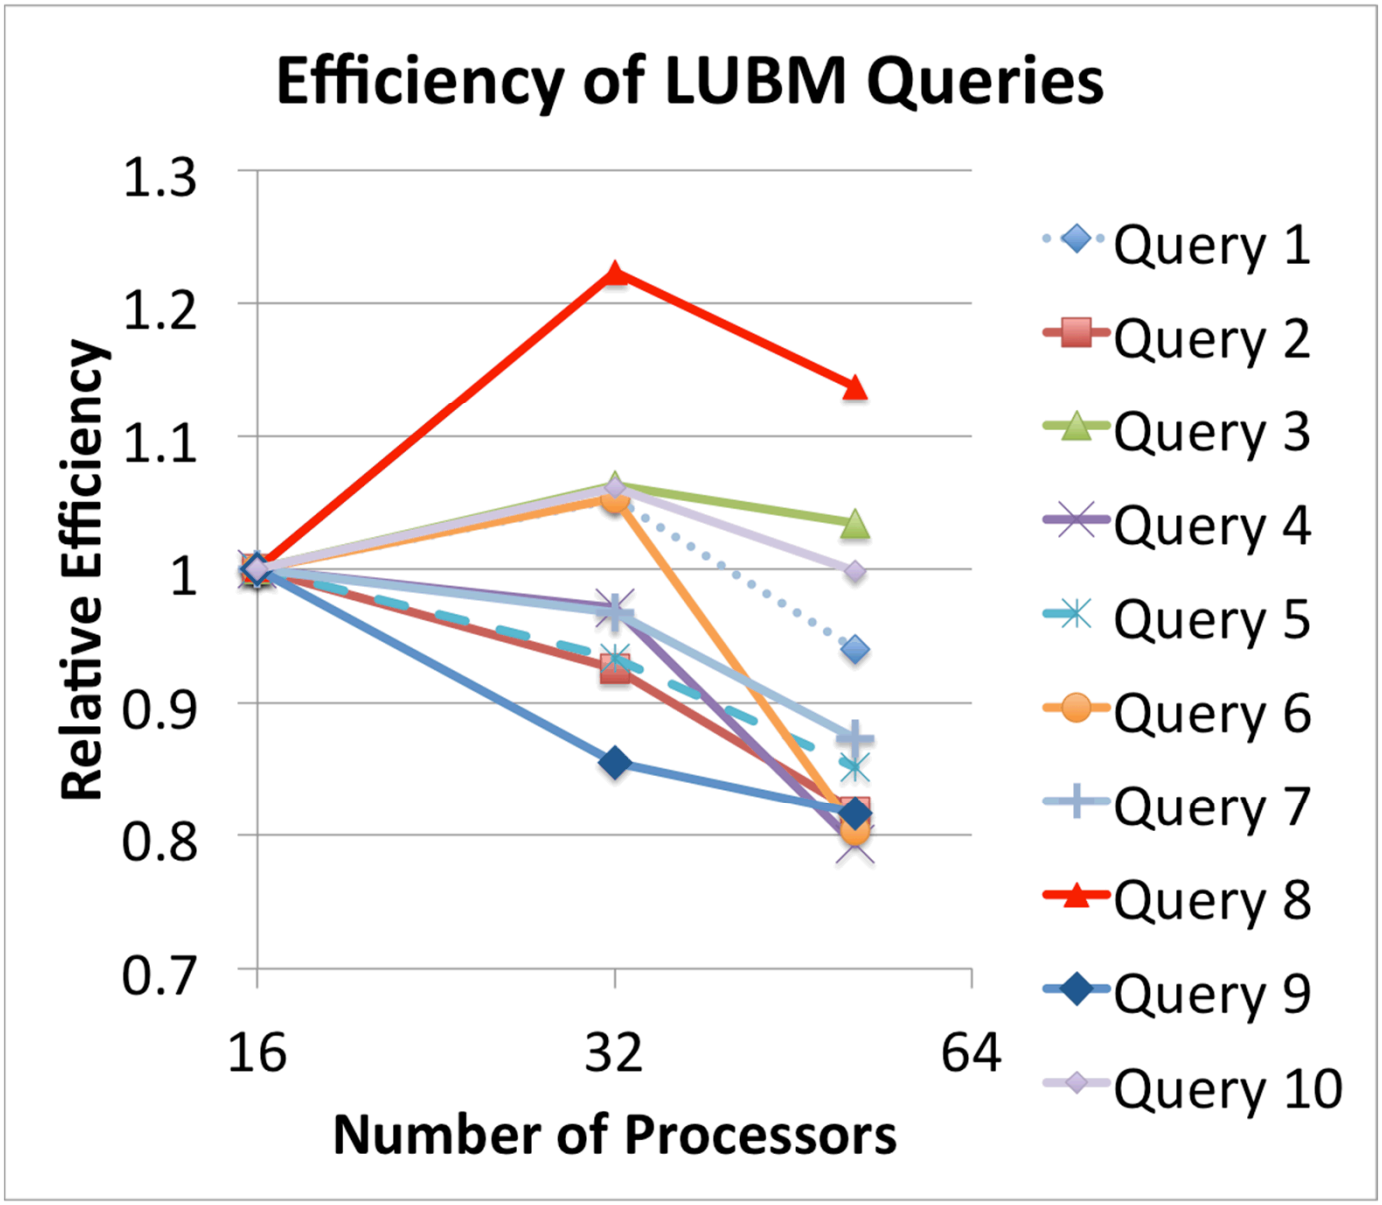
Efficiency of LUBM Queries
Relative Efficiency
1.3
1.2
1.1
1
0.9
0.8
0.7
16
32
64
Number of Processors
Query 1
Query 2
Query 3
Query 4
Query 5
Query 6
Query 7
Query 8
Query 9
Query 10

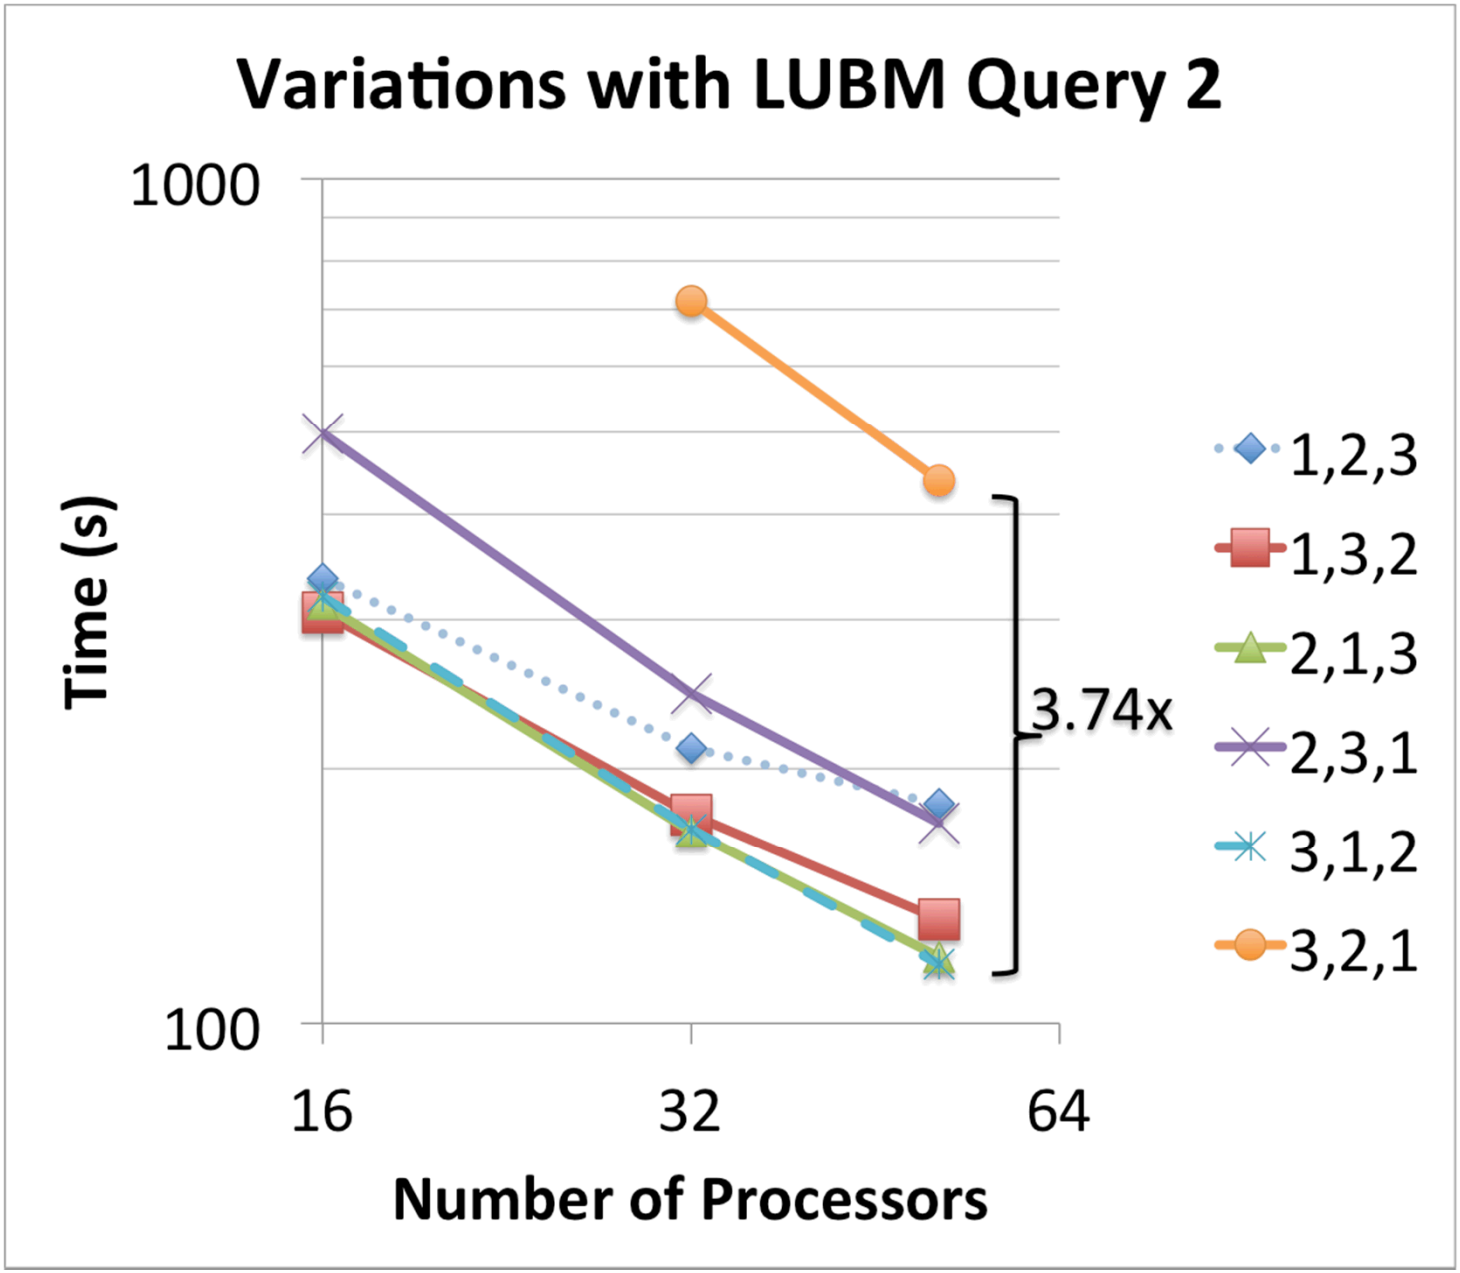

Sandia National Laboratories

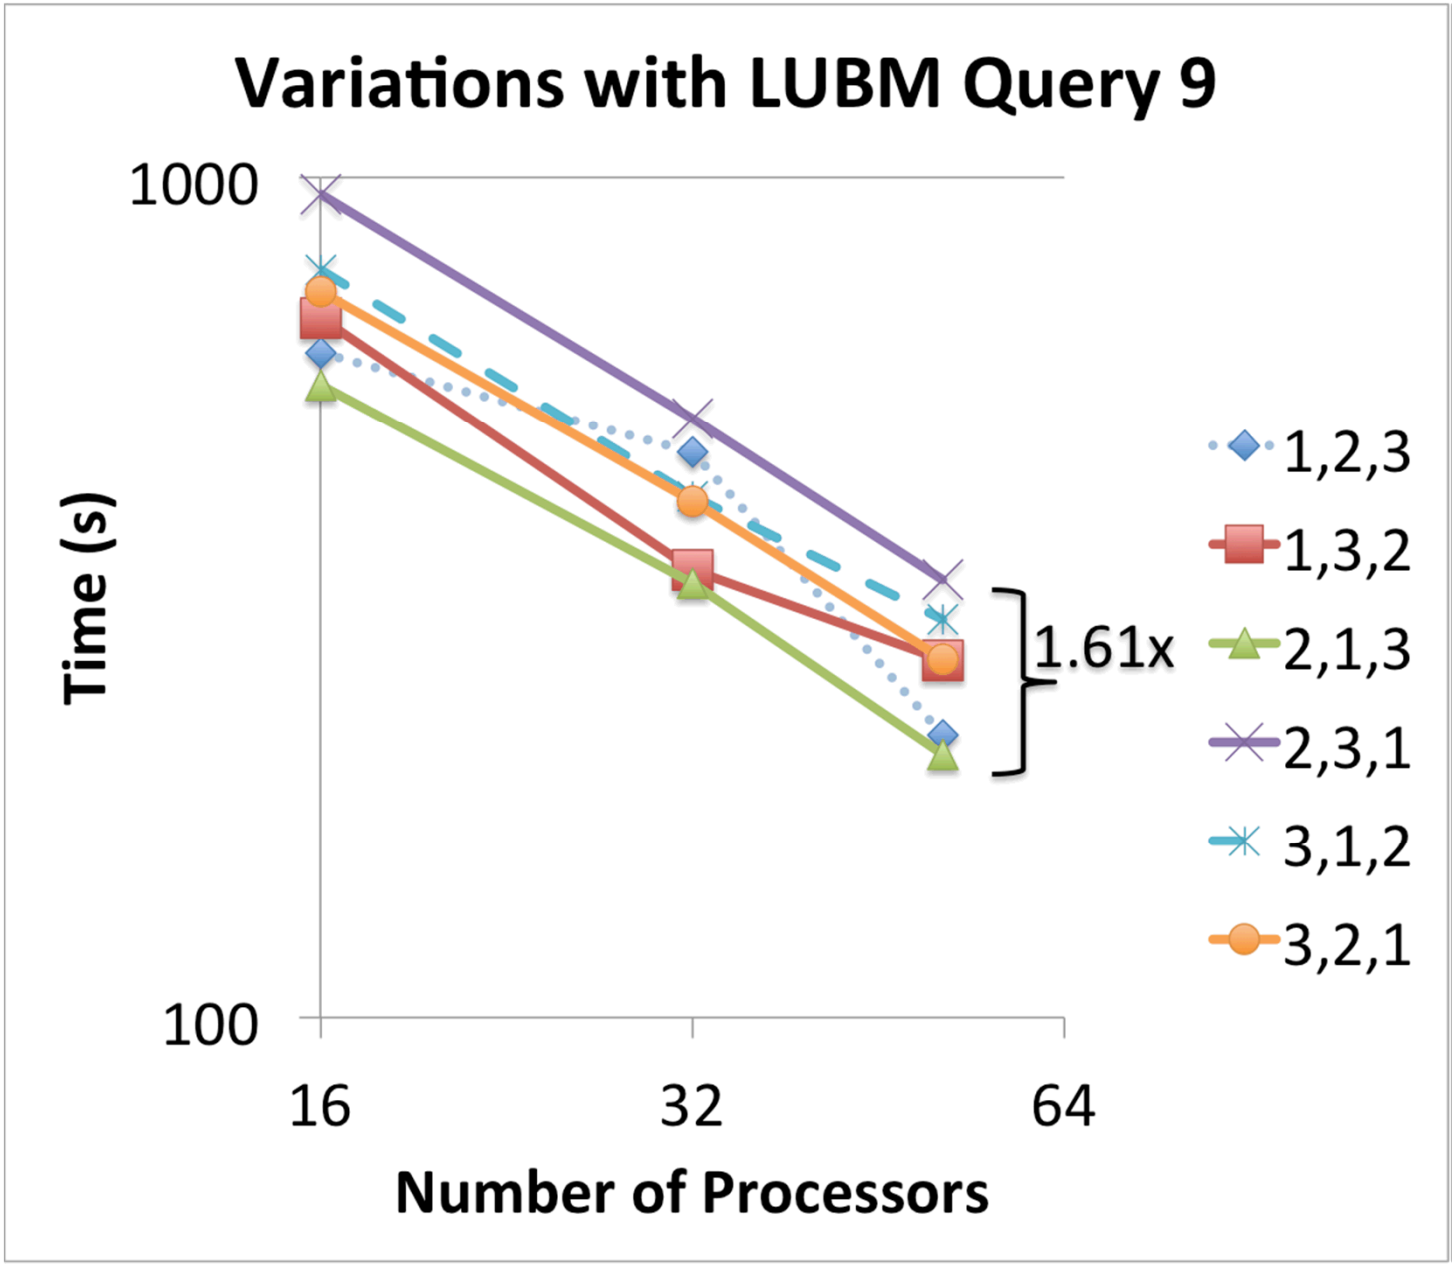
Variations with LUBM Query 9
1000
100
Time (s)
16
32
64
Number of Processors
1.61x
1,2,3
1,3,2
2,1,3
2,3,1
3,1,2
3,2,1

Limitations

Sandia National Laboratories

- No execution plan
  - Need to add indices/statistics
- Looks at every edge the first iteration
- No dictionary encoding
- Limited set of SPARQL

Conclusions

Sandia National Laboratories

- Vertex-centric computing is a viable option for parallelizing SPARQL queries
  - Straightforward implementation
  - Good scalability
- Semantic Web can benefit from additional functionality of vertex-centric paradigm