

# Bulk-Synchronous Parallel (BSP) Computing



• A Bridging Model between abstract PRAM model and real-world parallel computers, to support algorithm development







# BSP Remarks Local variables of BSP processes persist to next superstep In superstep *s*, contents of messages sent (and received) in step *s*-1 are accessible Two-sided and/or one-sided message passing possible BSP implementations in the 1990s include BSP implementations in the 1990s include BSP implementations (put, get) PUB – library for C atop MPI [Bonorden et al.'03] One-sided communication NestStep [K. 2000, 2004]

 Partitioned global address space (PGAS) language extension of C

















### Pregel

### [Malevicz et al. 2010]

II,Ų IN

- A framework to process/query large distributed graphs
- Proprietary (by Google)
- Pregel is the "MapReduce" for graphs
  - Iterative computations
    - Sequence of BSP supersteps
  - Each BSP superstep is basically a composition of the MapReduce phases (Map, Combine, Sort, Reduce)
- Attempts to utilize all servers available by partitioning and distributing the graph
  - · Good for computations touching all vertices / edges
  - Bad for computations touching only few vertices / edges

lı.v Pregel **Programming Model** States of a vertex Graph Vertices Inactive • Each with unique identifier (String) • Each with a user-defined value • Each with a state in { Active, Inactive } • Initially (before superstep 1), every vertex is active Graph Edges Each edge can have a user-defined value (e.g., weight) One (conceptual) BSP process assigned to each vertex • Not to the edges, by the way... • Iteratively executing supersteps while active • Two-sided communication with send() and receive() calls Graph can be dynamic Vertices and/or edges can be added or removed in each superstep Algorithm termination When all vertices are simultaneously inactive and there are no messages in transit

- Otherwise, go for another superstep











# Pregel C++ API (1)

Vertex<VertexValue,EdgeValue,MessageValue> class

- User overrides virtual Compute() method for superstep behavior
- vertex\_id() returns vertex identifier (string)
- VertexValue: user-specified datatype for value
- Get value() reads, Mutable value() sets the value
- GetOutEdgeIterator() get an OutEdgeIterator
- int64 superstep() queries the superstep number.
- Vertex and edge values are the only values that persist to the next superstep
- All messages sent to vertex in previous superstep are available
  - Each message contains a value of type MessageValue
  - void SendMessageTo( dest\_vertex\_id, msg\_value )
  - void VoteToHalt()



# Pregel C++ API (3)

LU UNCP

.

### Graph topology modifications

- Superstep execution can add or remove vertices or edges
- Could lead to conflicts across parallel BSP processes, e.g.
  - Multiple processes try to create a vertex with same name in same superstep
  - ... Or: with same name but different initial values
  - · One wants to add an edge from/to a vertex that another wants to remove
- ...
- Conflict resolution policy:
  - Edge removals always done first, then vertex removals
  - · Then do vertex additions, then edge additions
  - Then do Compute() for the vertex
- Purely vertex-local modifications (e.g. adding/removing own outgoing edges) are conflict-free



- (Combine and) Aggregate messages to vertices on each other node Tell the master how many local vertices remain active for next • superstep
- Fault tolerance through checkpointing, master reassigns lost partitions





## References

# I,U UNIVERS

### BSP

- L. Valiant: A bridging model for parallel computation. *Comm. ACM* 33(8), 103-111, 1990
   L. Valiant: A bridging model for multi-core computing. *J. Comp. and Syst. Sciences* 77(1), 154-166, 2011.
- D.B. Skillicorn, Jonathan M.D. Hill, and W.F. McColl: Questions and Answers about BSP. Scientific Programming, vol. 6, no. 3, pp. 249-274, 1997. doi:10.1155/1997/532130

### Pregela

 Grzegorz Malewicz, Matthew H. Austern, Aart J. C. Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski: Pregel: A System for Large-Scale Graph Processing. Proc. SIGMOD'10, June 6–11, 2010, pp. 135-145, ACM.

### BSP on Hadoop / Cloud:

- Apache HAMA: https://hama.apache.org/
   K Siddigue et al.: Apache Hama: An Emerging Bulk Synchronou
- K. Siddique et al.: Apache Hama: An Emerging Bulk Synchronous Parallel Computing Framework for Big Data Applications. IEEE Access 4:8879 - 8887, Nov. 2016. http://ieeexplore.ieee.org/document/7752866/
- Redekopp, M., Simmhan, Y., Prasanna, V.K.: Optimizations and analysis of BSP graph processing models on public clouds. IPDPS 2013



- How could Aggregators be implemented in a Pregel cluster?
- How could removals of edges speed up the vertex value maximization algorithm for large distributed graphs?
- Find further graph algorithms that are similar in structure to distributed connected components or PageRank.