




Apache Hama v0.4

@ Korea Telecom



Edward J. Yoon, Jan 31, 2012
<edwardyoon@apache.org>

Index

- ▶ **What is Hama?**
 - ▶ Bulk Synchronous Parallel
 - ▶ Schematic diagram of a superstep
- ▶ **Hama Characteristics**
 - ▶ Internals
- ▶ **Why Hama and BSP?**
- ▶ **Hama In Korea Telecom**
- ▶ **Performance Evaluation**
 - ▶ SSSP
 - ▶ K-means clustering
 - ▶ PageRank



About Me

- ▶ Founder of Apache Hama.
- ▶ Employee for Korea Telecom.



What Is Hama?

- ▶ Apache Incubator Project.
- ▶ BSP (Bulk Synchronous Parallel) for massive scientific computations.
- ▶ Written In Java.
- ▶ Currently 3 releases, 3 main committers and 2 more active contributors.




Bulk Synchronous Parallel?

- ▶ Parallel programming model introduced by Valiant.
- ▶ Consist of a sequence of supersteps.
- ▶ Conceptually simple and intuitive from a programming standpoint.
- ▶ Used for a variety of applications e.g., scientific computing, genetic programming, ...



Schematic diagram of a superstep



Hama Characteristics

- ▶ Provides a Pure BSP.
 - ▶ Job submission and management interface.
 - ▶ Multiple tasks per node.
 - ▶ Input/Output Formatter.
 - ▶ Checkpoint recovery.
- ▶ Supports to run in the Clouds using Apache Whirr.
- ▶ Supports to run with Hadoop YARN.



Internals

- ▶ Hadoop RPC is used for BSP tasks to communicate each other.
- ▶ Collection and bundling of messages as a technique to reduce network overheads and contentions.
- ▶ Zookeeper is used for Barrier Synchronization.



Why Hama and BSP?

- ▶ Supports message passing paradigm style of application development.
- ▶ Provides a flexible, simple, and easy-to-use small APIs.
- ▶ Enables to perform better than MPI for communication-intensive applications.
- ▶ Guarantees impossibility of deadlocks or collisions in the communication mechanisms.



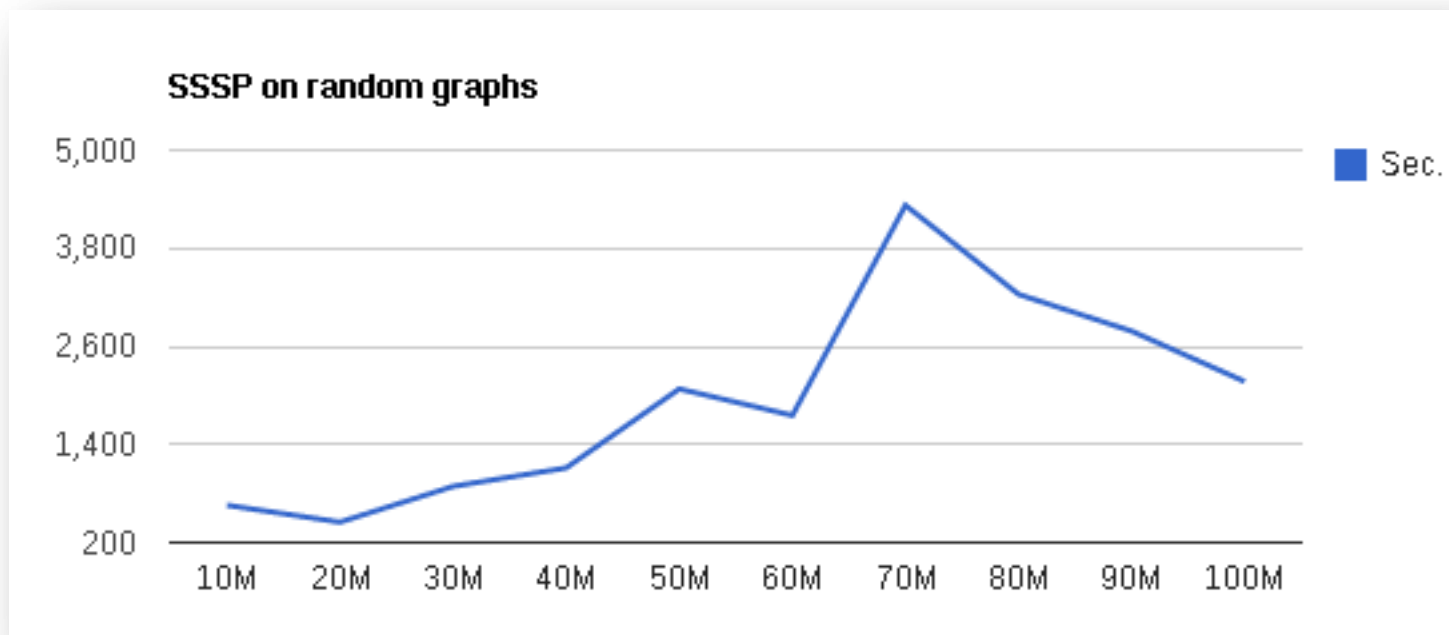
Hama In Korea Telecom

- ▶ **Structural Analysis of Network Traffic Flows.**
 - ▶ traffic mining, engineering, anomaly detection, traffic forecasting and capacity planning
 - ▶ Currently BSP jobs are experimentally running on 512 multi-cores machines.




Performance: SSSP algorithm

- ▶ A SSSP for a random graph (100 million vertices, 1 billion edges) is computed in **30 minutes** on 512 cores machines.



Performance: PageRank

- ▶ A PageRank for 30 million random web pages (10 anchors per page) is computed in 20 minutes on 256 cores Hama cluster!



What's Next?

- ▶ Fault Tolerant System.
- ▶ Message Compression for High Performance.
- ▶ Add some frameworks on top of Hama.
- ▶ Add other RPC libs.
- ▶ Elegant Web UI.

