


# SCALABLE, SECURE ANALYSIS OF SOCIAL SCIENCES DATA ON AZURE

Yogesh Simmhan

Electrical Engineering Department  
University of Southern California




Estimated U.S. Energy Use in 2010: ~98.0 Quads



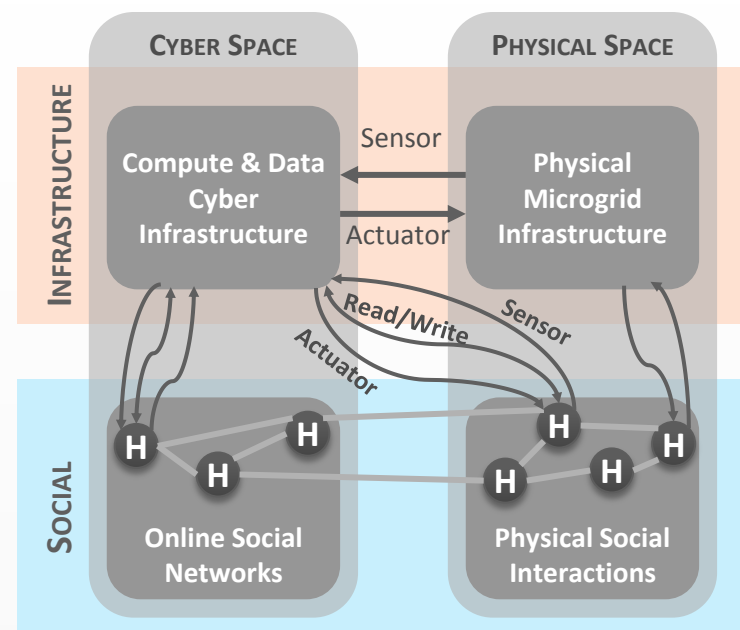
Source: LLNL 2011. Data is based on DOE/EIA-0384(2010), October 2011. If this information or a reproduction of it is used, credit must be given to the Lawrence Livermore National Laboratory and the Department of Energy, under whose auspices the work was performed. Distributed electricity represents only retail electricity sales and does not include self-generation. EIA

# Energy Informatics

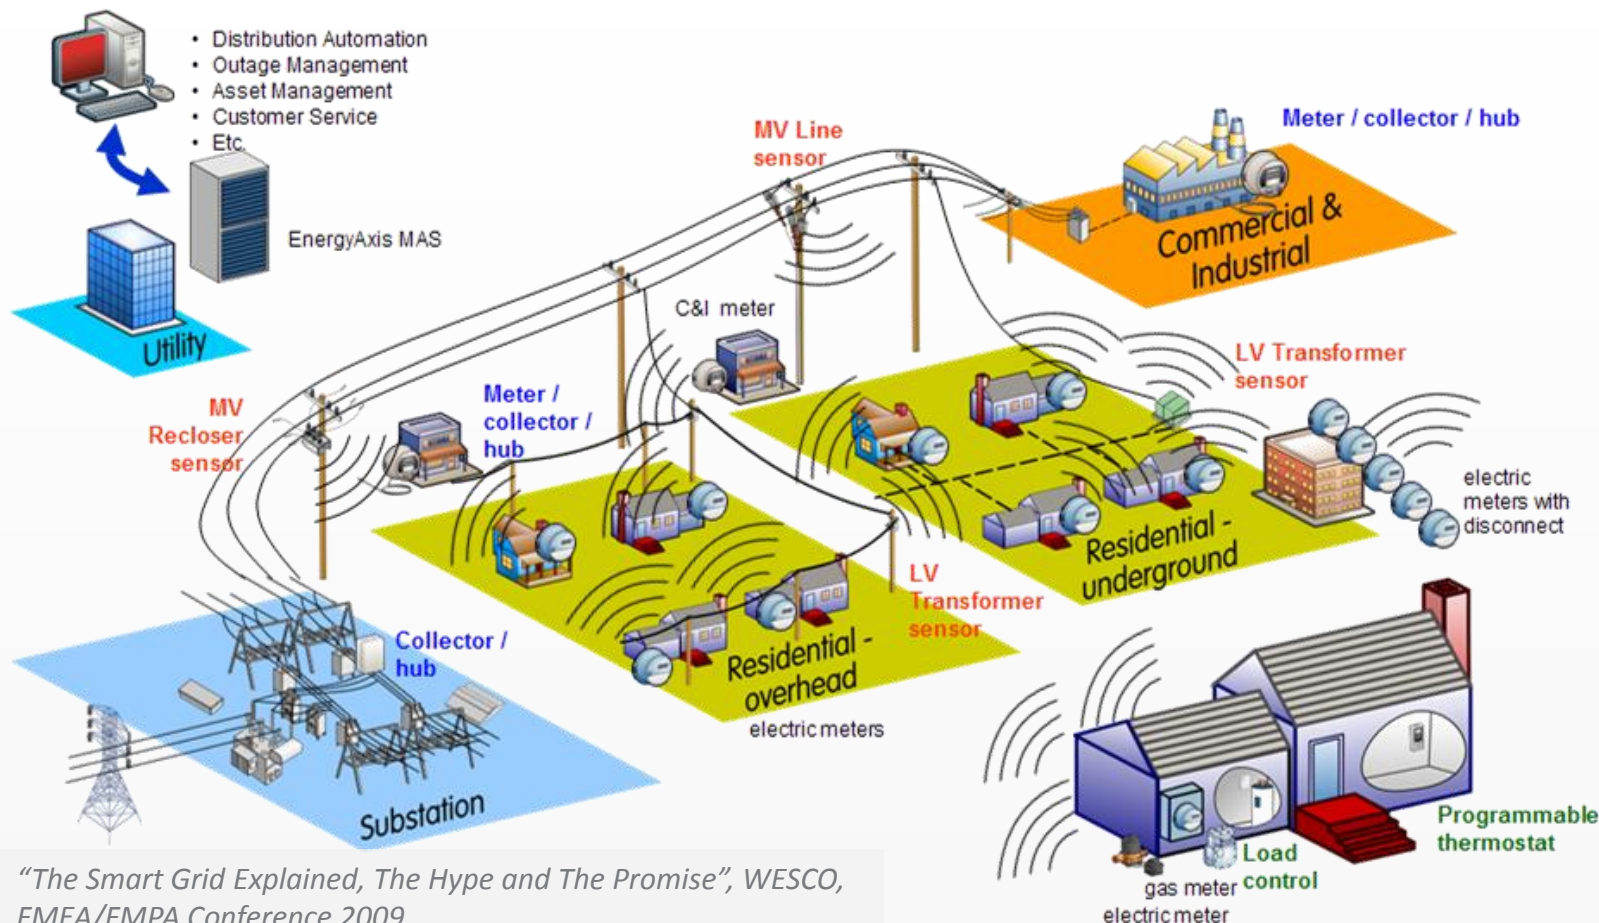


# Cyber-Physical-Social Systems

- Physical & Cyber Infrastructure
- Social interactions
- Transportation, Energy, Health, ...
- Impact on sustainable living
- *Internet of Things!*



# Smart Power Grids



*"The Smart Grid Explained, The Hype and The Promise", WESCO, FMEA/FMPA Conference 2009*

# Los Angeles Smart Grid Demo Project

- ▶ LA is the largest public utility in US
  - ▶ 1% of US power consumed by 1.4M customers
- ▶ Collaborative effort
  - ▶ USC engaged with DR, CB & EV



USC





UCLA

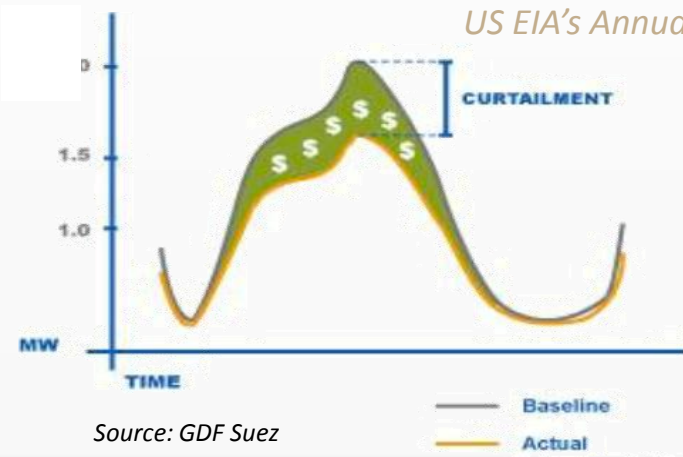
JPL

# Demand Response Optimization

- More renewables
  - ▶ *Intermittent sources*
- More hybrid & electric vehicles
  - ▶ *Increase dynamic load on power grid*
- Efficient management of demand & supply
  - ▶ *Shift, Shape, Shave power demand*



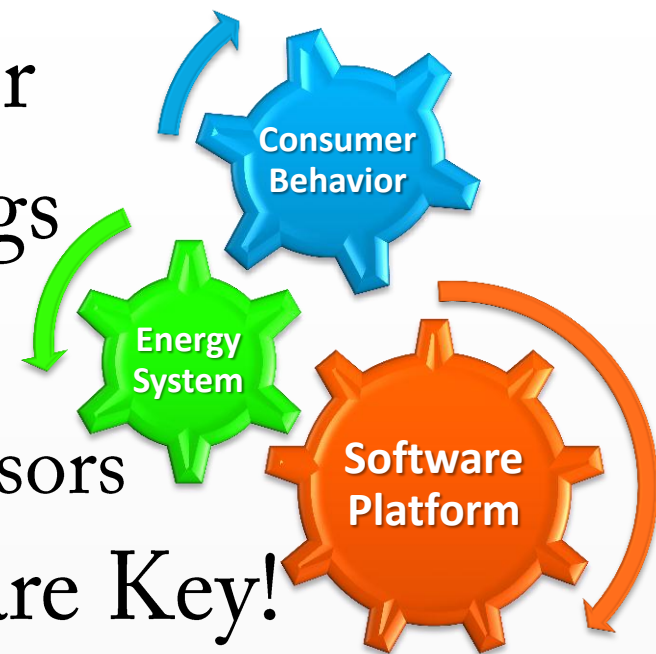
US EIA's Annual Energy Outlook 2011



Source: GDF Suez

# USC Campus MicroGrid

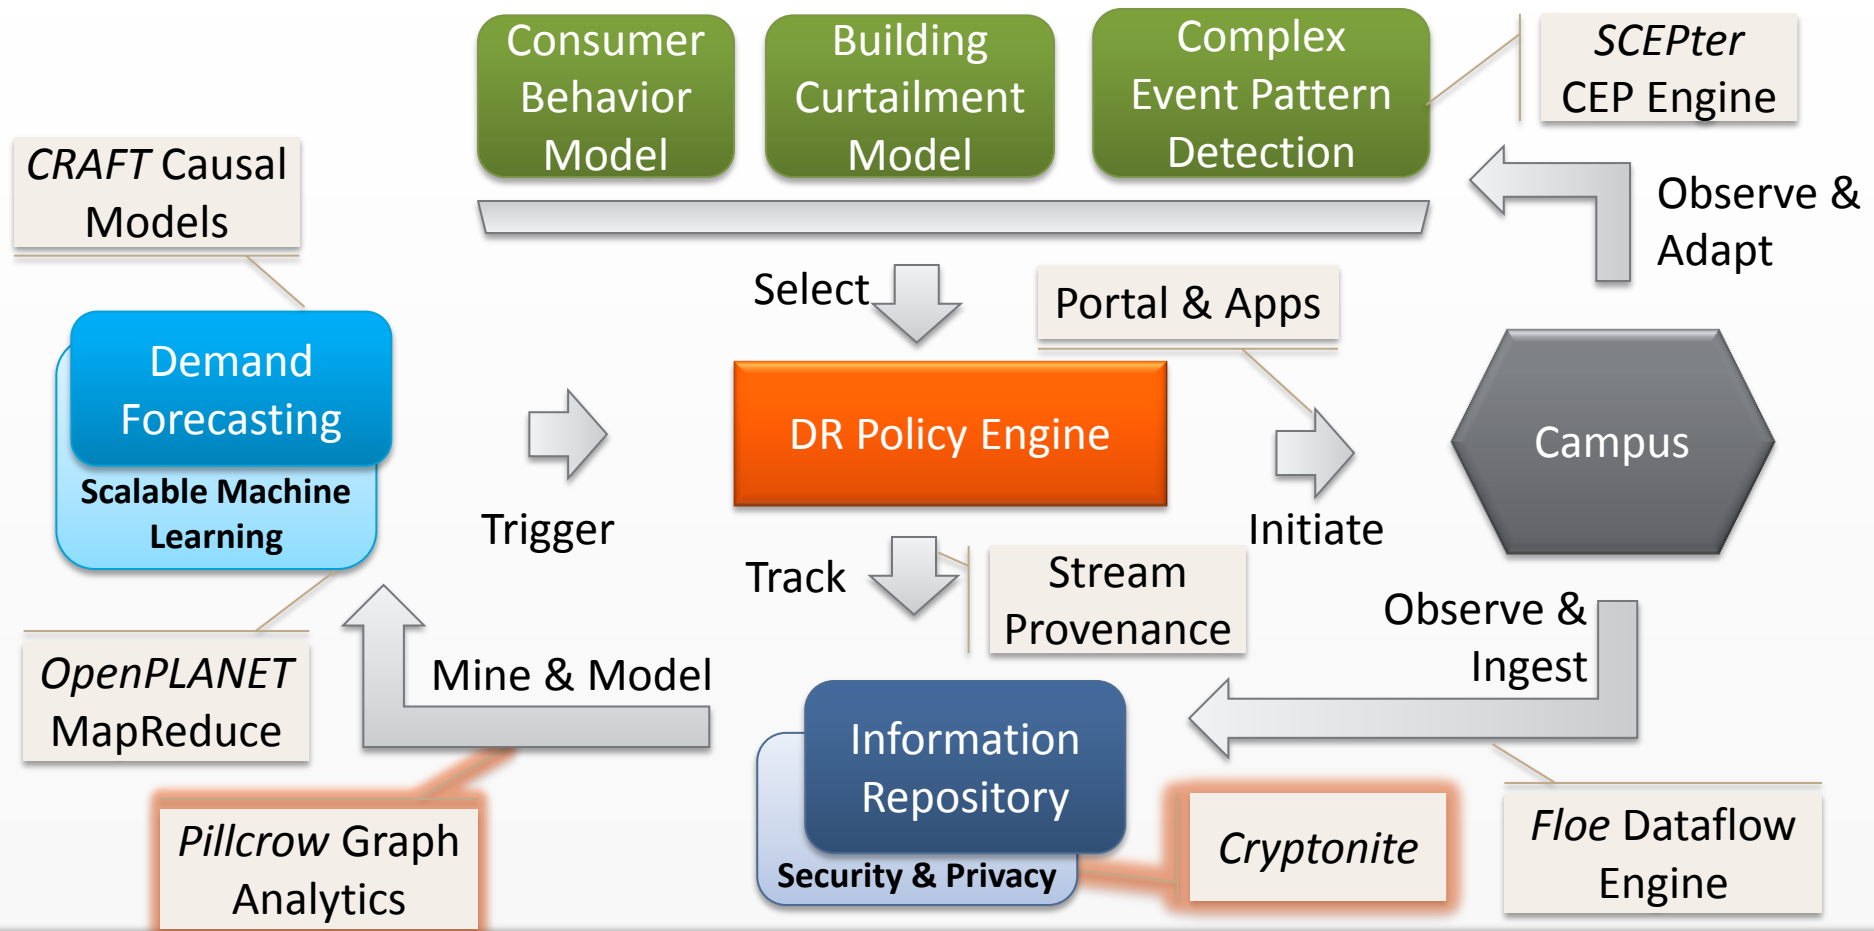
- ▶ LA's largest private consumer
- ▶ 60,000 people, 100+ buildings
- ▶ Well instrumented
  - ▶ 50K power, HVAC, temp sensors
- ▶ **Consumers & Information** are Key!
  - ▶ CS, Facilities, Sociologists, Psychologists, ...



*A Living, Learning Laboratory for DR experiments*



# Platform for Smart Grid Analytics

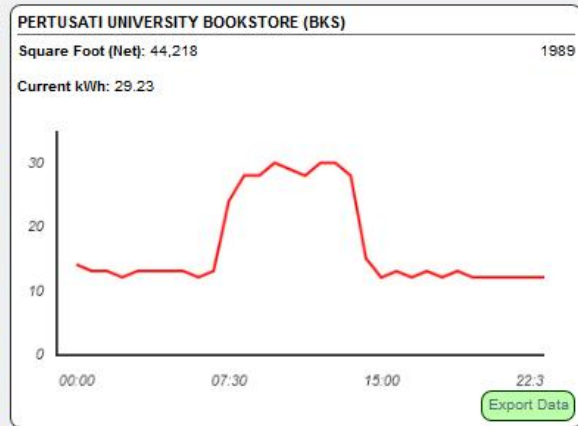
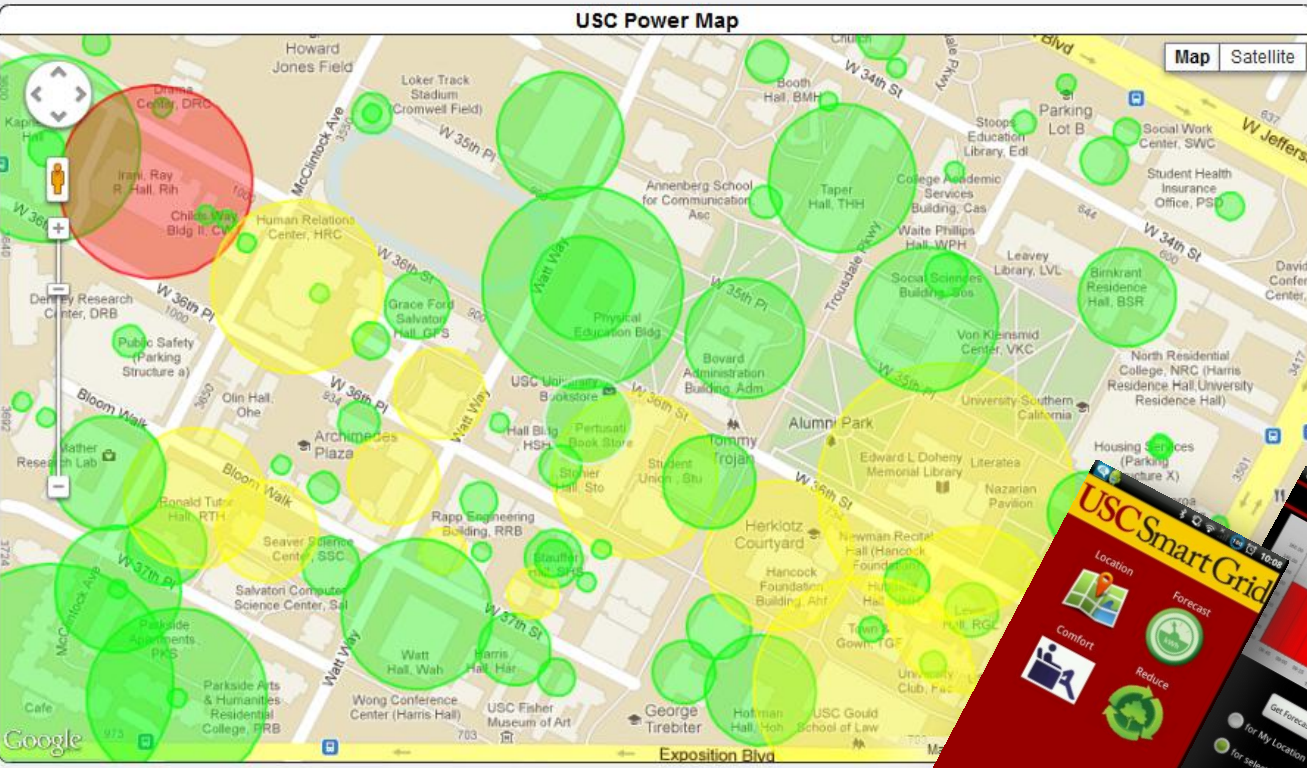


*Learn from the Microgrid. Scale to the Mega-city grid.*



Web Portal Admin Panel

USC Smart Grid Web Portal



**USC Smart Grid**

Energy Forecast

Comfort

Reduce Your Energy Usage

Sustainability Newsfeed

# Social Sciences Challenges on Clouds

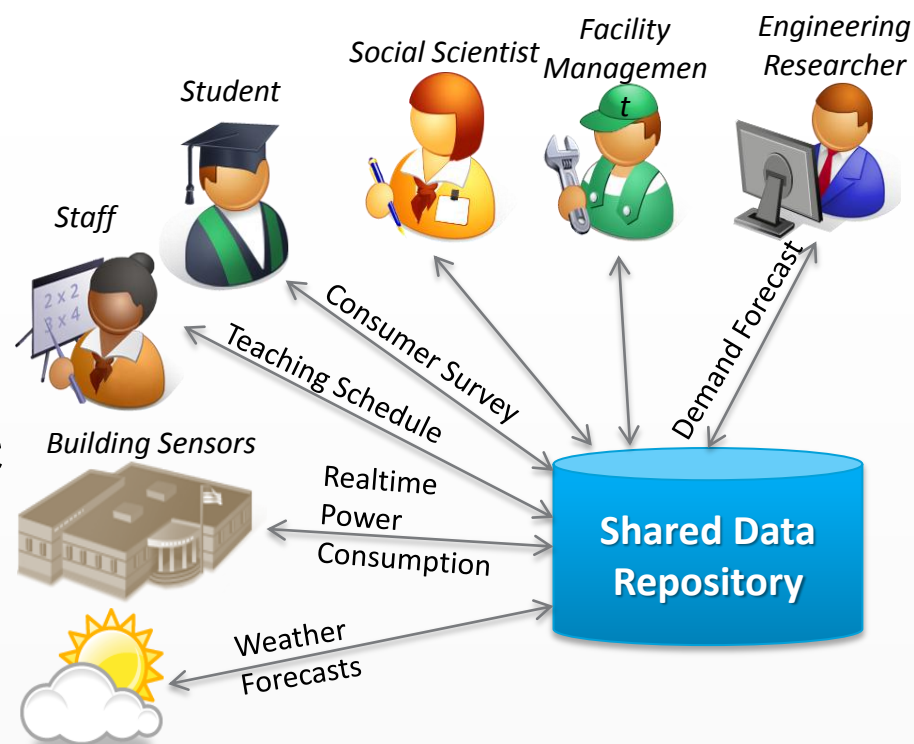
1. *How does a multi-disciplinary team securely share data and collaborate?*
2. *How do we analyze social interaction graphs at large scales?*

*Cryptonite*

# SECURE DATA SHARING ON PUBLIC CLOUDS

# Data Sharing on Clouds

- Ease of management
- Reliability, persistence
- Scalable sharing
- Co-located with Cloud Apps
- *But...*



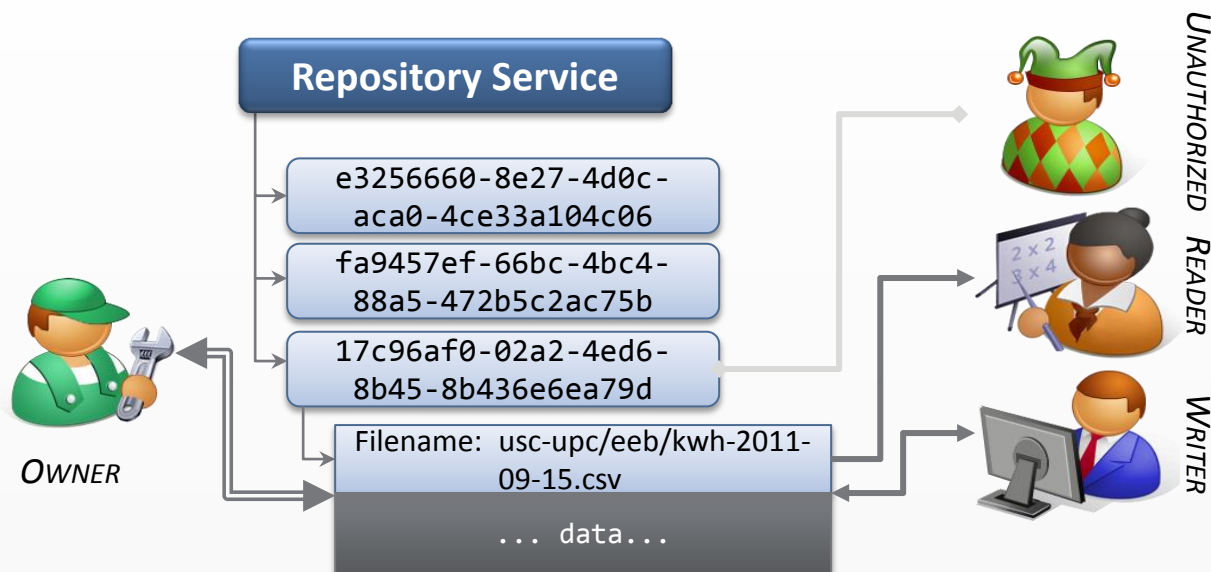
<i>Data</i>	Social Sci.	Power Engr.	Facility Mgmt.	Univ. Staff	Extern
Cons. Survey	O	R*	-	-	-
Staff Schedule	R*	R*	-	O	-
Consumption	R	R	O	-	-
Weather	R	R	O	R	W
Campus Power Forecast	R	W <sup>13</sup>	O	R	-

# Storage Concerns on Public Clouds

- Wider attack surface
- Loss of physical control
- Not a major concern for most applications, ...
- Use Cloud provider Access Control
  - ▶ *Challenge: What if you don't trust Cloud provider?*
- Plain text data may potentially leak
  - ▶ Encrypt data before uploading
  - ▶ *Challenge: Key management & sharing*

- **Data storage service** which
  - retains the **persistence & availability** offered by public Cloud storage,
  - but with **security & encryption controlled by clients**,
  - a **low key management overhead**,
  - **compatibility** with existing storage service interfaces,
  - while **not sacrificing performance & scalability**.

# Entities and Roles



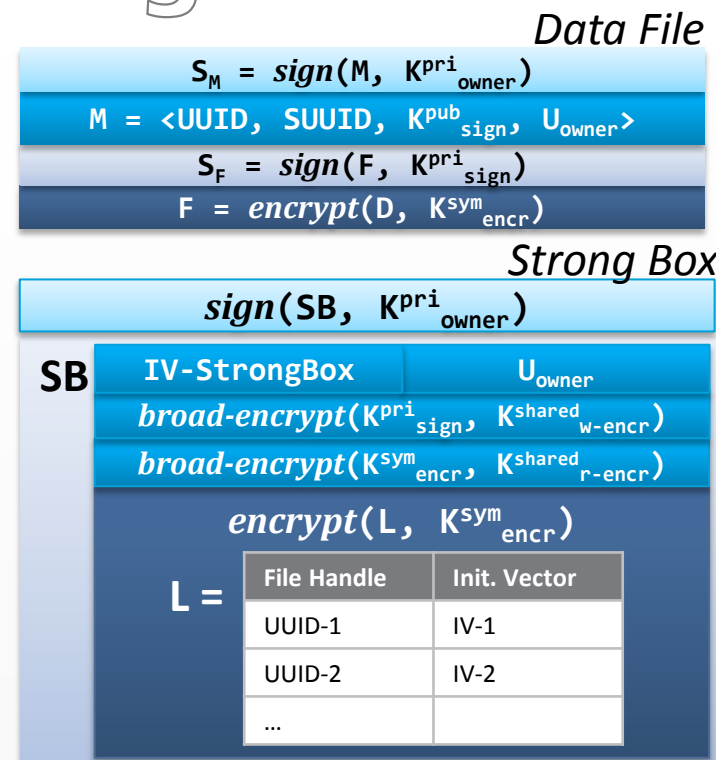


# Entities and Roles

- **Users (*USC Community*)**
  - ▶ Community of users who require a secure, shared storage
  - ▶ Owners, Readers, Writers
  - ▶ **Can include 'Malicious' attackers**
- **Cloud Storage Provider (*Azure*)**
  - ▶ Provides persistent scalable storage space
  - ▶ **Trusted with 'availability' (SLA)**
  - ▶ **Not trusted with data security**
- **Secure, Shared Data Repository Service (*Cryptonite*)**
  - ▶ Run and Managed by user community on the Cloud
  - ▶ **Not trusted with plain text data**
  - ▶ **Partial trusted to perform requested operations**
  - ▶ **All operations should be verifiable**

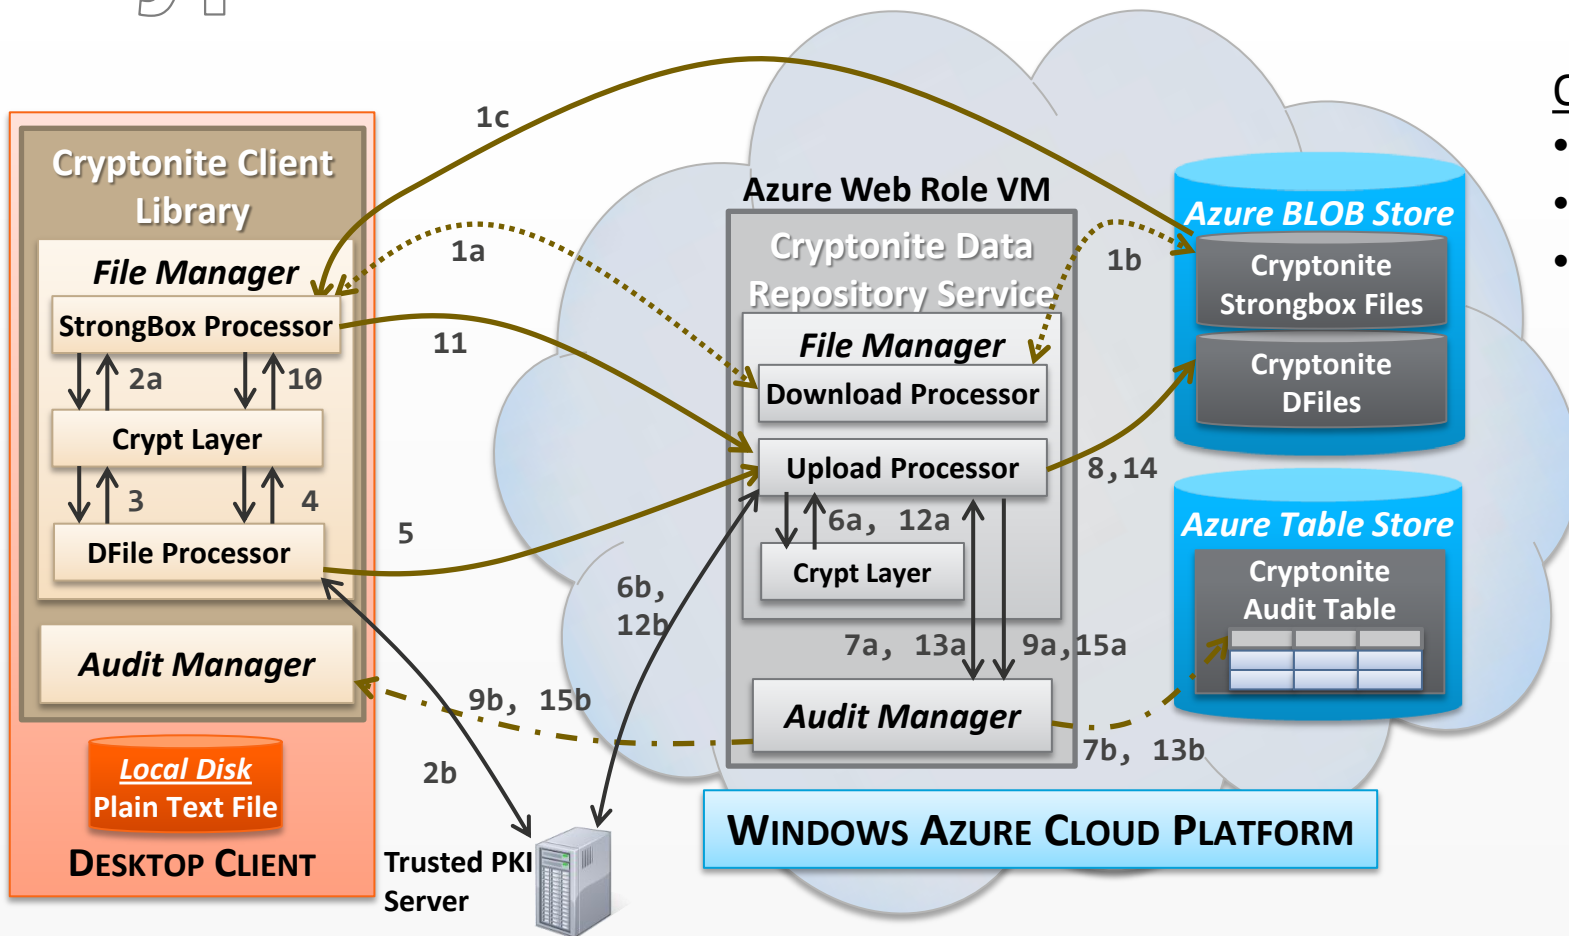
# StrongBox for Key Management

- Every user has one global private:public key pair
- File is **encrypted** using a **symmetric key** unique to ACL
- File is **signed** using a **key pair** unique to ACL
- Symmetric key** is **broadcast encr.** using public keys of **readers**
- Private key** of signature is **BE** using public keys of **writers**
- StrongBox** is signed by **owner pvt key**



Designing a Secure Storage Repository for Sharing Scientific Datasets using Public Clouds, A. Kumbhare, Y. Simmhan & V. Prasanna, *DataCloud*, 2011

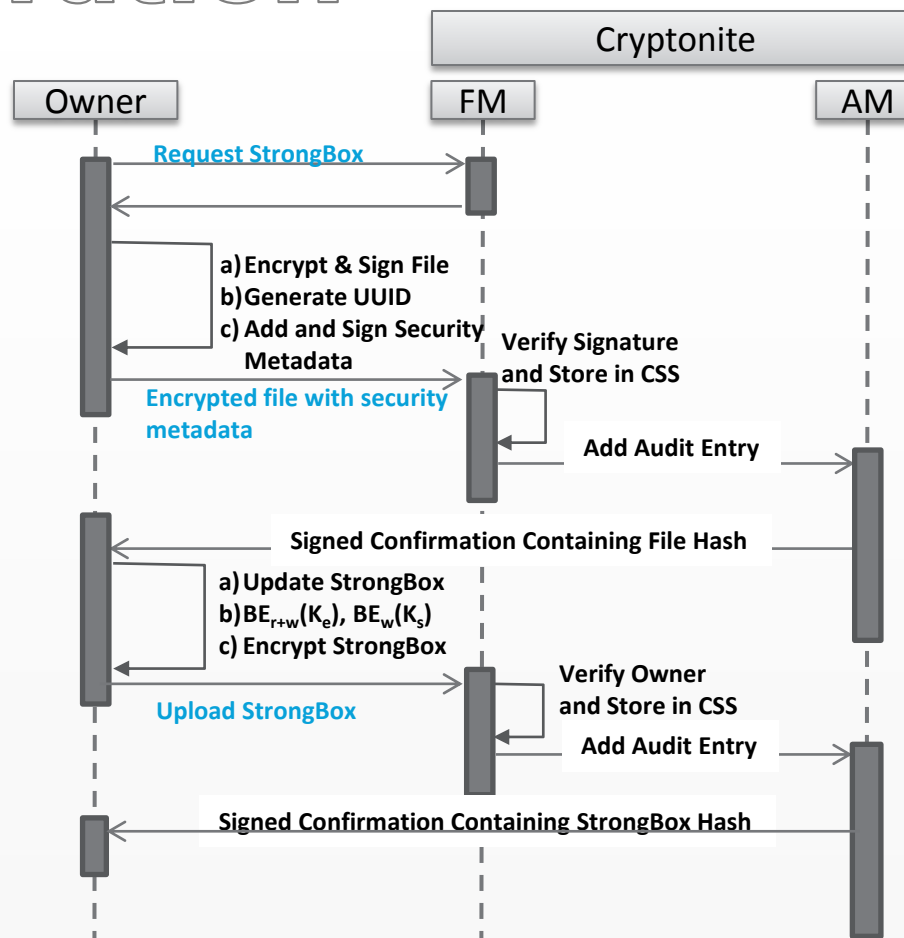
# Cryptonite Architecture



## Operations

- GET
- PUT
- POST

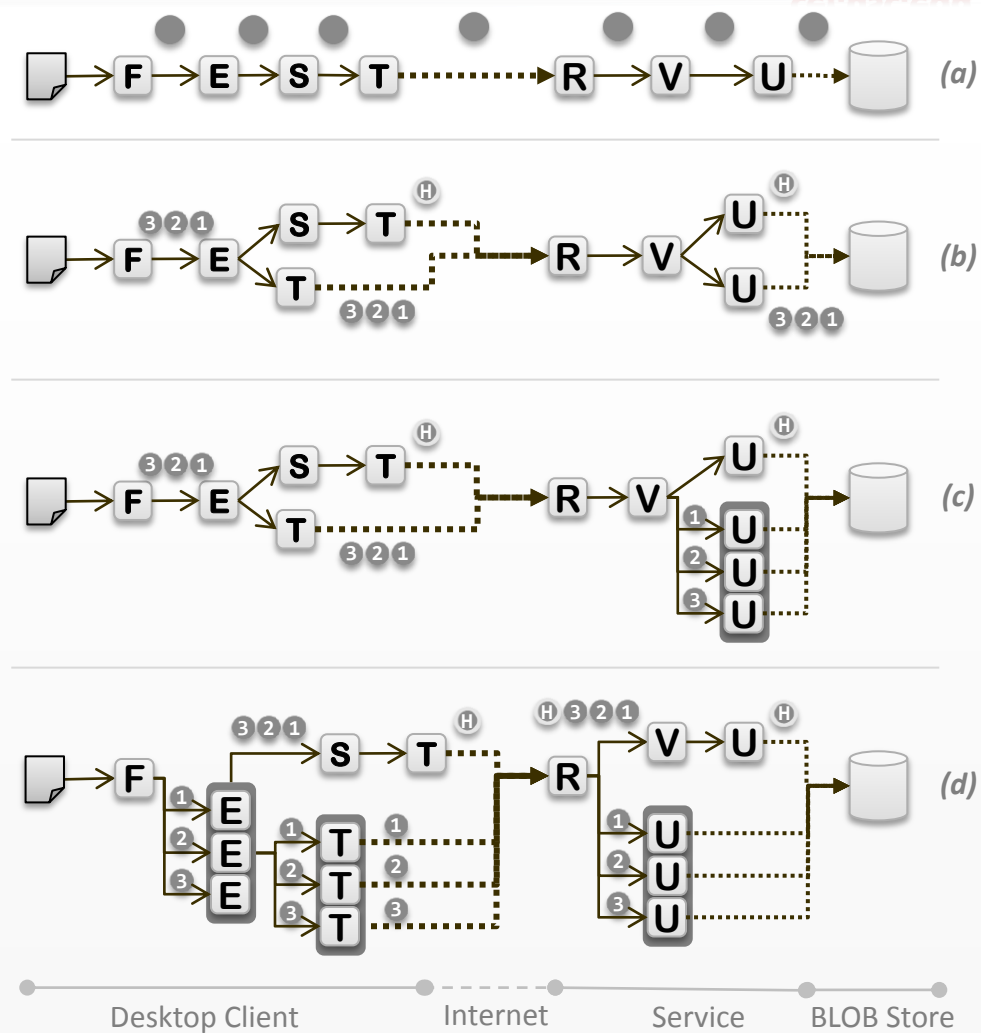
# PUT Operation




# Optimizations

- File read, Encrypt, Strongbox, Transfer
- NW Read, Verify, Blob Upload
- Pipelining
- Data Parallel


Cryptonite: A Secure and Performant Data Repository on Public Clouds, A. Kumbhare, Y. Simmhan & V. Prasanna, *IEEE CLOUD*, 2012




# Baseline Performance



(a) POST



(b) PUT




(c) GET


- PUT/POST comparable for medium sized files.  
Poor for large/small files.
  - ▶ Azure Client API does parallel upload >32M
- GET comparable for medium-large files

# Optimized Perf.

- ▶ PUT outperforms
  - ▶ Cryptonite service buffers, stores async
- ▶ GET outperforms
  - ▶ Azure API does not do parallel download
- ▶ Scalability: 400Mb/s B/W Saturation in VM



(a) POST



(b) GET

# Future

- ▶ Scale out of VMs
- ▶ Auditing & verification support
- ▶ Leveraging native storage fabric B/W?
- ▶ Table oriented data sharing

[ganges.usc.edu/wiki/Cryptonite](http://ganges.usc.edu/wiki/Cryptonite)



*Pillcrow: Daytona MR, Pregel.NET, ...*

# GRAPH ANALYTICS ON AZURE

# Graph Data Analysis

- ▶ Understand interconnections between energy consumers, physical space
- ▶ Mine for patterns of interesting energy (mis)use activity
- ▶ Analyze key influencers of sustainability in peer group

# Comparing Prog. Models

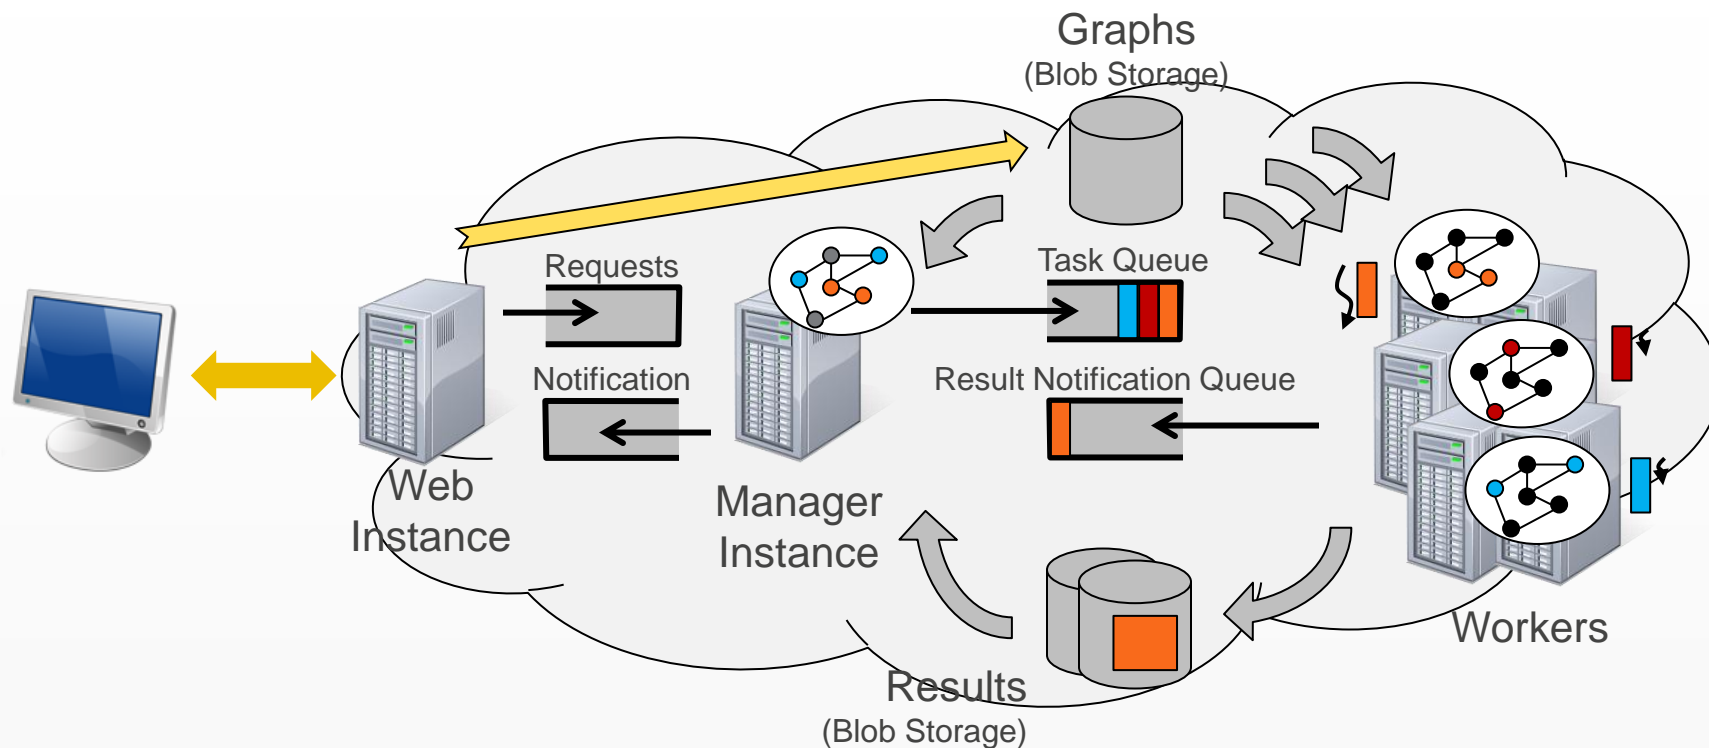
- ▶ Loosely-coupled
  - ▶ Azure Task queues
- ▶ MapReduce
  - ▶ MSR Project Daytona
- ▶ Pregel
  - ▶ Pregel.NET
- ▶ Use Betweenness Centrality for evaluation

# Betweenness Centrality (BC)

- Determines “centrality” as # of shortest paths from all-pairs on which a vertex lies
- Perform  $|V|$  Breadth-first traversals
  - ▶ Gives shortest path from root to other vertices
- Reverse traversal from leaf to root, summing % of shortest-paths each vertex belongs to
- Sum vertex’s scores across traversals gives BC

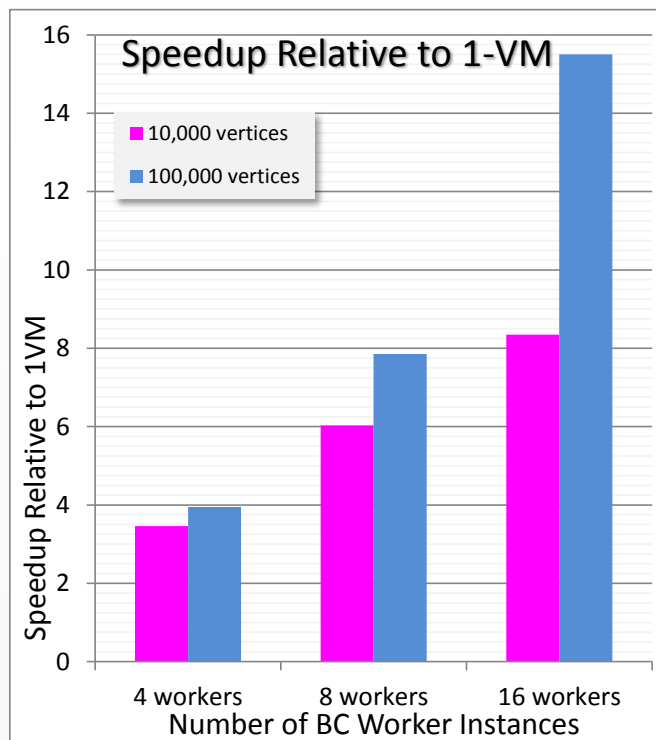
Brandes, U.: A faster algorithm for betweenness centrality.  
*J. Mathematical Sociology*, 25(2), 163–177 (2001)

# BC using Loosely Coupled Tasks




Performance Analysis of Vertex-centric Graph Algorithms on the Azure Cloud Platform, M. Redekopp, Y. Simmhan & V. Prasanna, *ParGraph*, 2011

# Perf. Of Loosely Coupled



For large graphs performance is scalable over number of workers





For large graphs and 16 workers we get about 8x speedup compared to a single local core

# Task Sizes and Resilience

Set timeout to average task time+12%

Analyze performance penalties for different MTBF

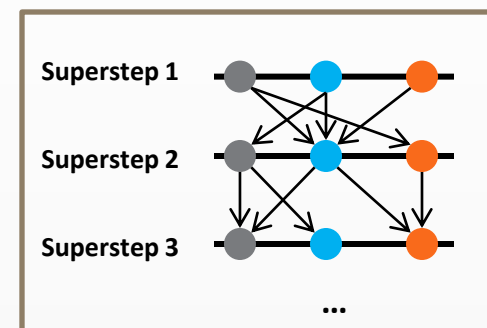


For  $10^5$  vertex graph, all tasks finish with in 12% variation from the average

More tasks per worker (smaller task size) helps to mitigate performance impact of failures

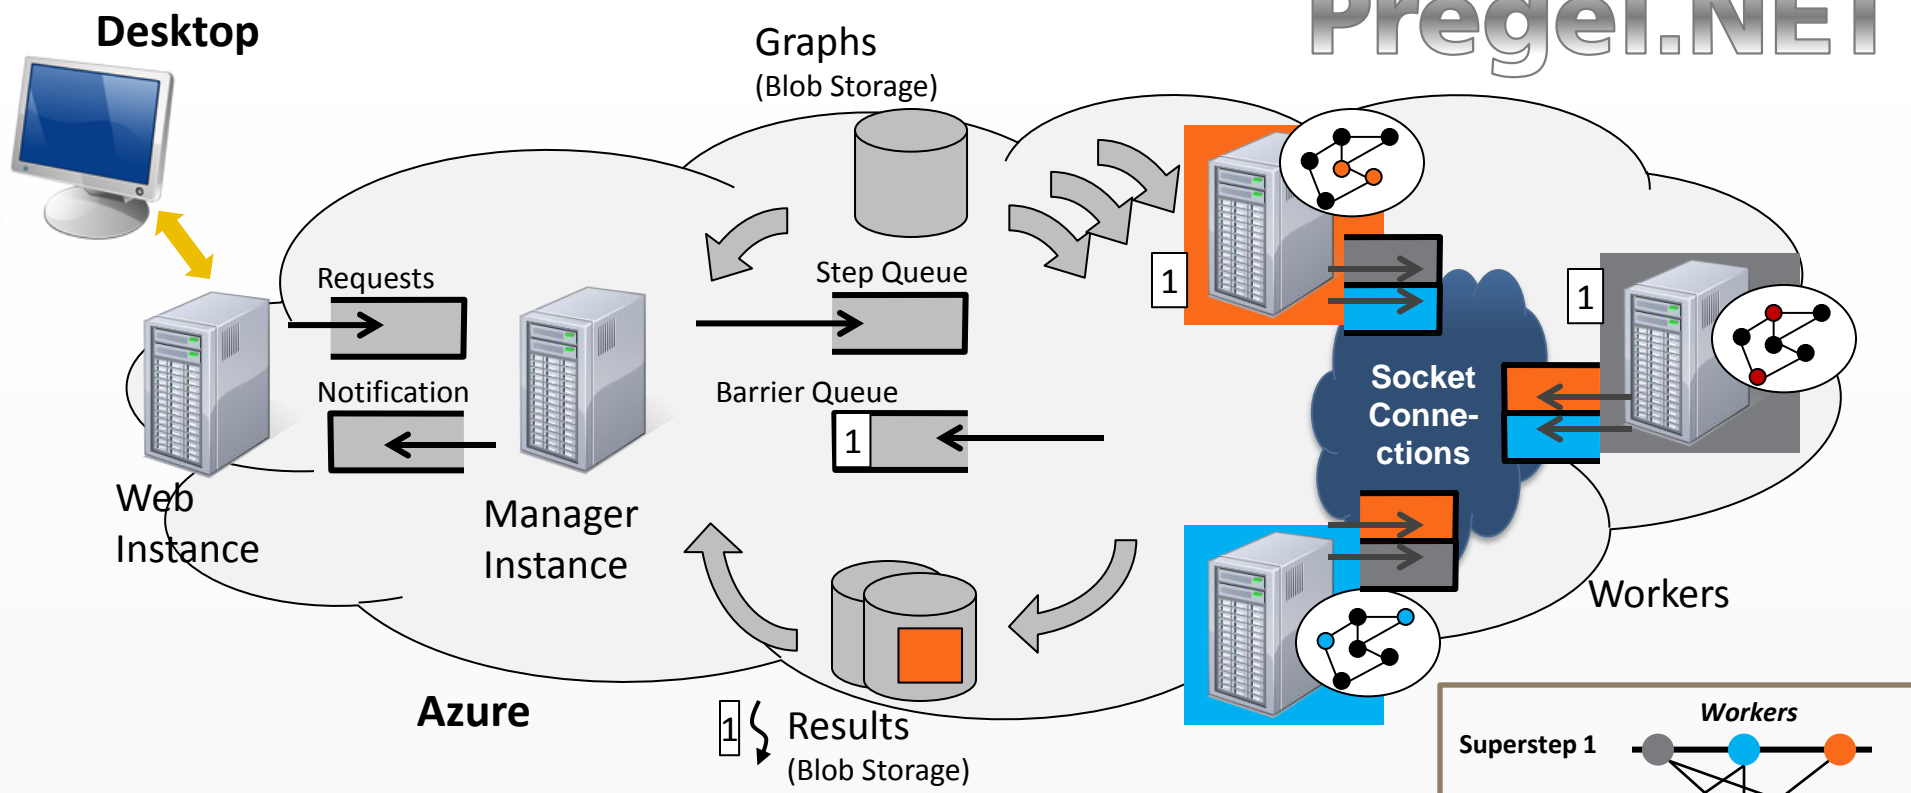
# Pregel

- ▶ Graph vertices distributed among workers
- ▶ Computation in synchronized supersteps
  - ▶ Bulk Synchronous Parallel model
- ▶ In each superstep
  - ▶ Vertex accepts messages in edges,
  - ▶ Processes messages, and
  - ▶ Emits messages to neighbors
- ▶ Stops when all “vote” to halt...M-R-R-R...

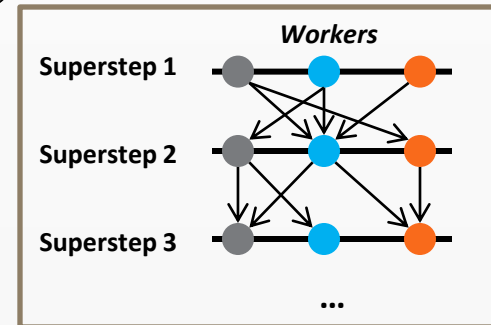




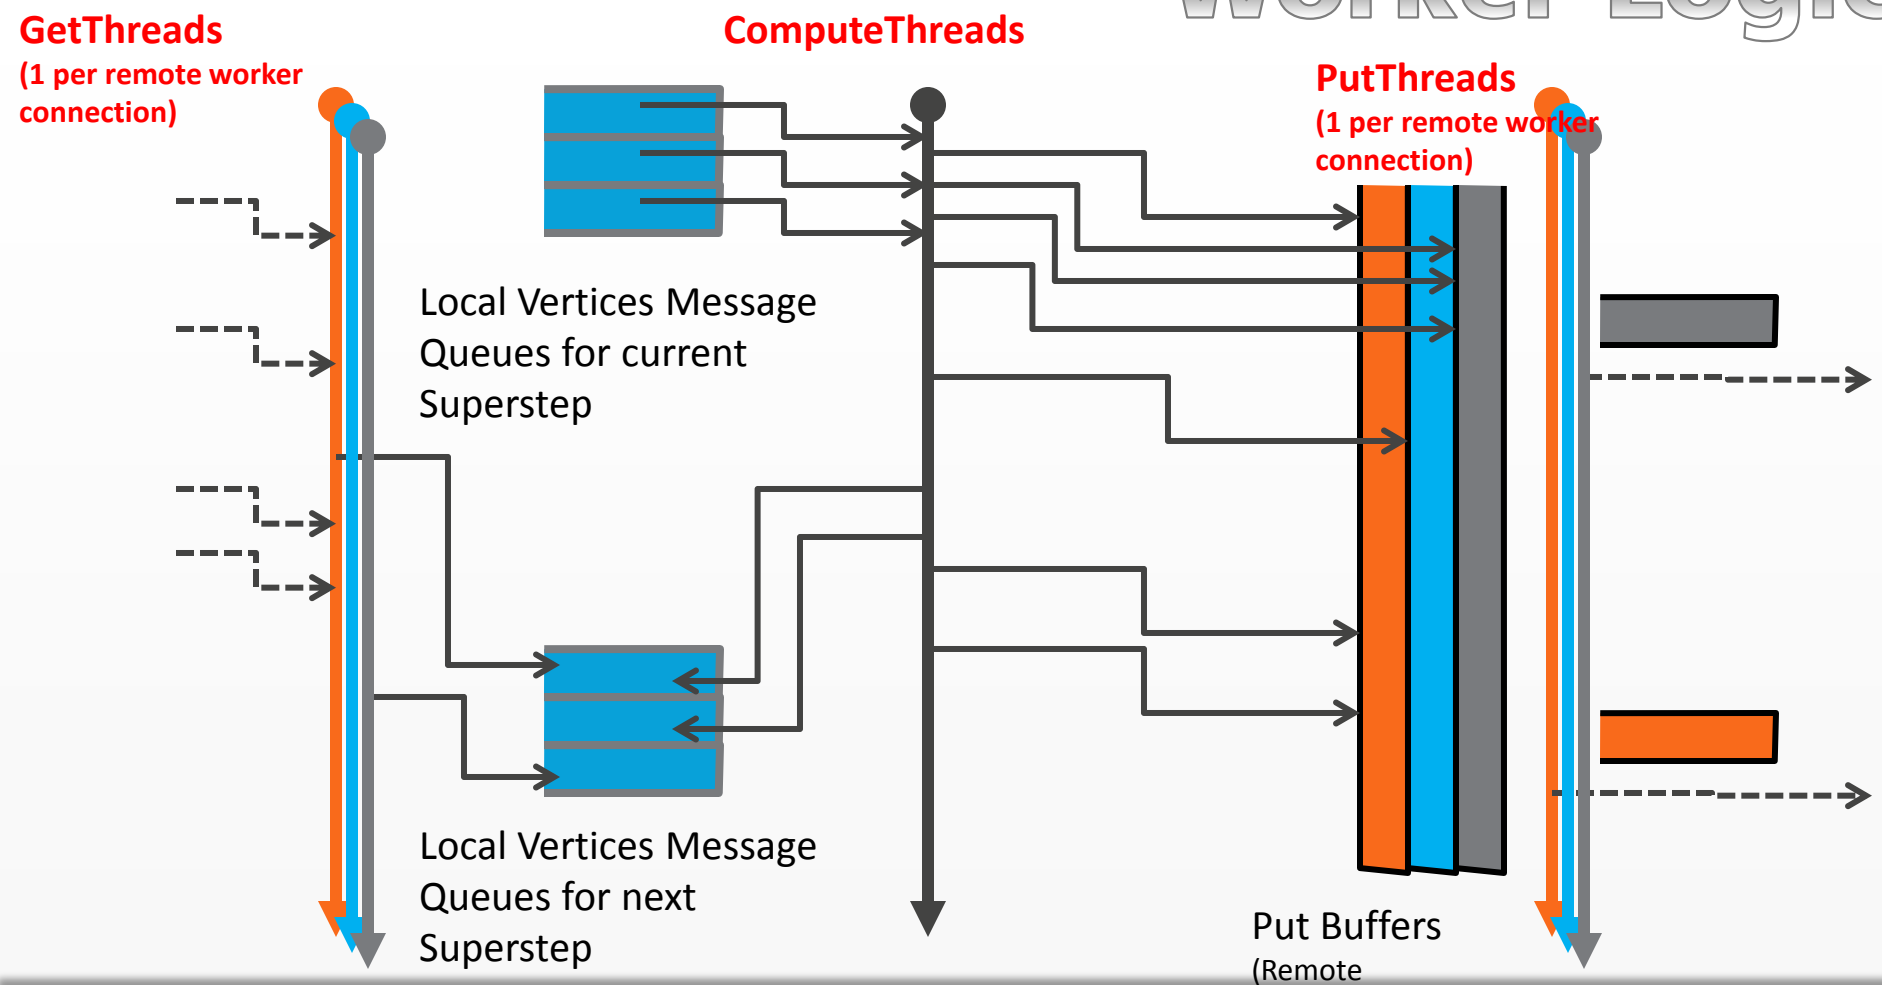
# Pregel.NET



## Pregel.NET Worker Role Implementation on Azure

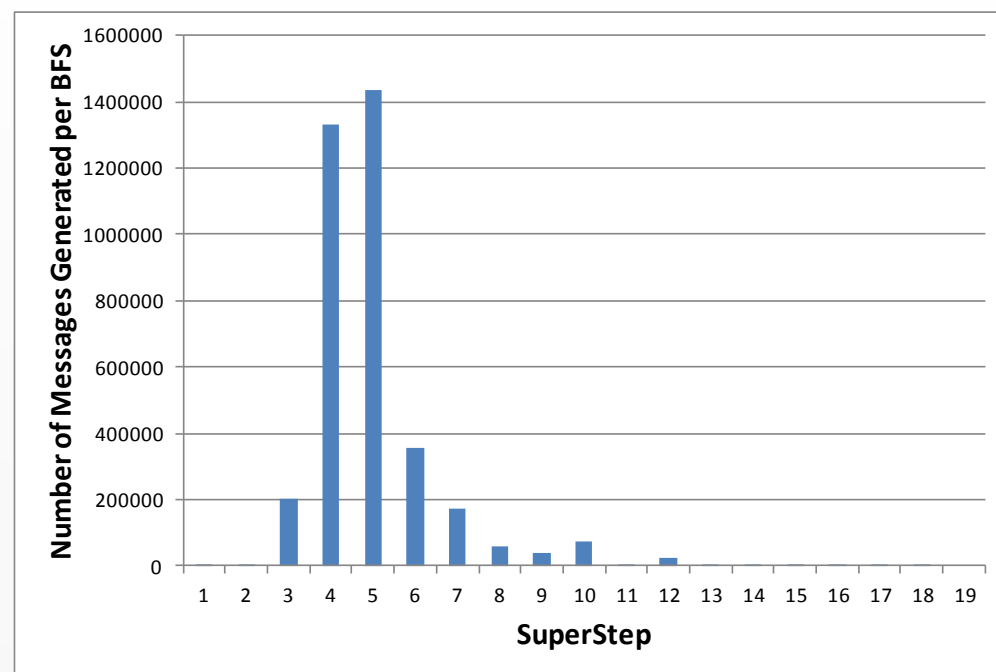


# Worker Logic




# Pregel Characteristics

- Communication Latency
  - ▶ Messages flow between workers on internal endpoints (sockets)
- Memory Pressure
  - ▶ Typical Pregel model does  $|V|$  BFS traversals in parallel
  - ▶ Large memory buffering for input messages
- Variability of message traffic & memory across supersteps ... long tail




# Pregel Optimizations

- Launch a swath of  $k$  BFSs at a time ( $k < |V|$ )
  - $k$  based on predicted peak memory requirements (How?)
- Launch swaths at intervals for constant message traffic, memory
  - Heuristics for selecting initiation interval

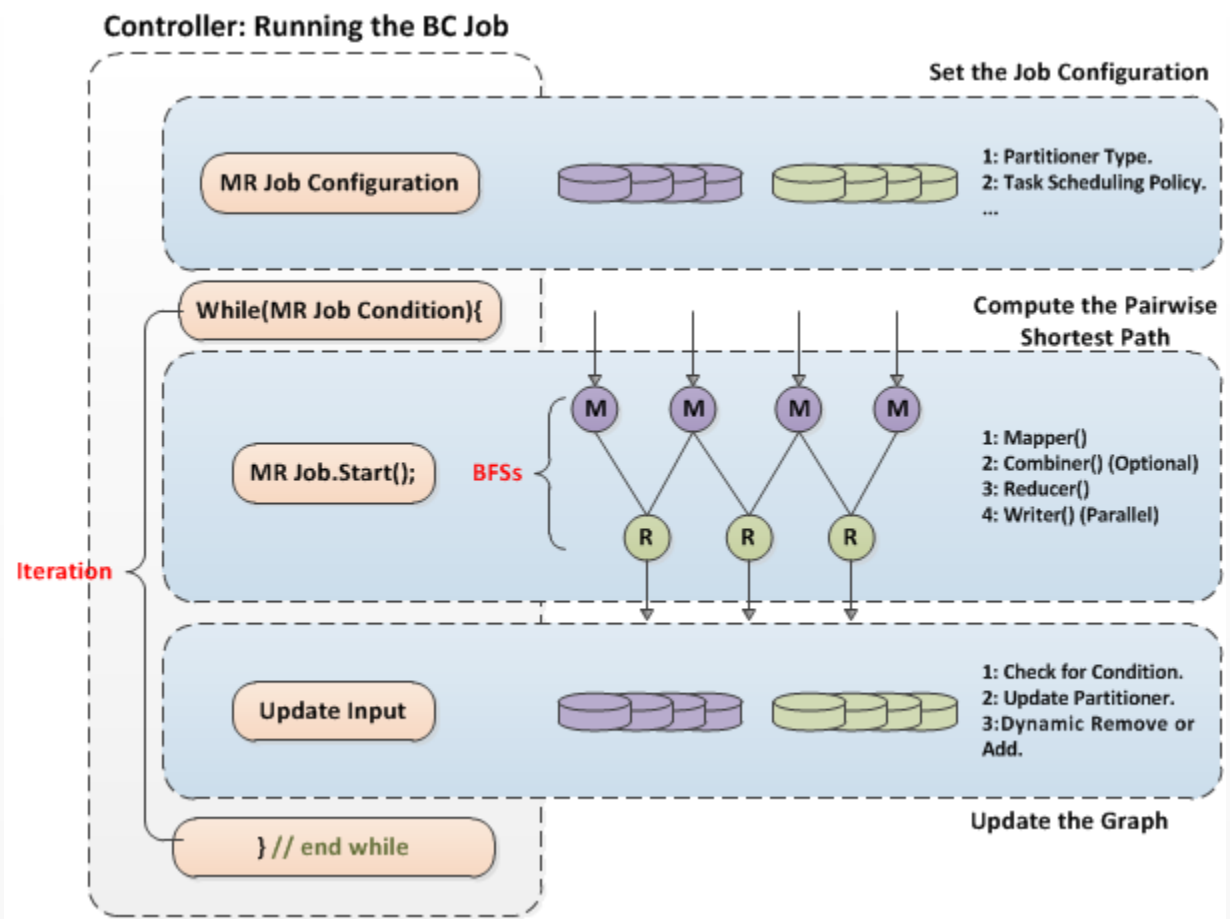


Message traffic from 4 swaths of BFS launched at an interval of 2 supersteps



Cumulative message traffic for the 4 swaths

# BC on Daytona's iMapReduce



# Future work

- Comprehensive evaluation
- Initiate parallel swaths in “waves”
  - ▶ Analytical model
- Dynamic graphs, incremental updates
- Other graph applications

# Plugs...

- ▶ **Workshop on Massive Data Analytics on Scalable Systems (DataMASS)**
  - ▶ Co-located with the 19<sup>th</sup> IEEE International Conference on High Performance Computing (HiPC), *December 18 - 21, 2012, Pune, India (TBA)*
- ▶ **Call for Articles: Cloud Computing for Data-driven Science and Engineering**
  - ▶ Concurrency and Computation: Practice and Experience Journal, *Deadline: August 15, 2012*

*This material is based upon work supported by the **National Science Foundation** and **Microsoft** under Award Number CCF-1048311 of the Computing in the Cloud initiative, and by the **Department of Energy** under Award Number DEOE0000192.*

*The views and opinions expressed herein do not necessarily state or reflect those of the US Government or any agency thereof.*



Los Angeles  
Department of  
Water & Power

**Microsoft®**

**THANK YOU!**

[ceng.usc.edu/~simmhan](http://ceng.usc.edu/~simmhan)  
[smartgrid.usc.edu](http://smartgrid.usc.edu)