

EVALUATING THE USE OF PROXIES IN MG-RAST

Caitlin Ross^{*#}, Misbah Mubarak⁺, Wei Tang⁺, John Jenkins⁺,
Wolfgang Gerlach⁺, Phil Carns⁺, Christopher Carothers^{*}, Robert
Ross⁺

^{*}Rensselaer Polytechnic Institute, Troy, NY

⁺Argonne National Laboratory, Argonne, IL

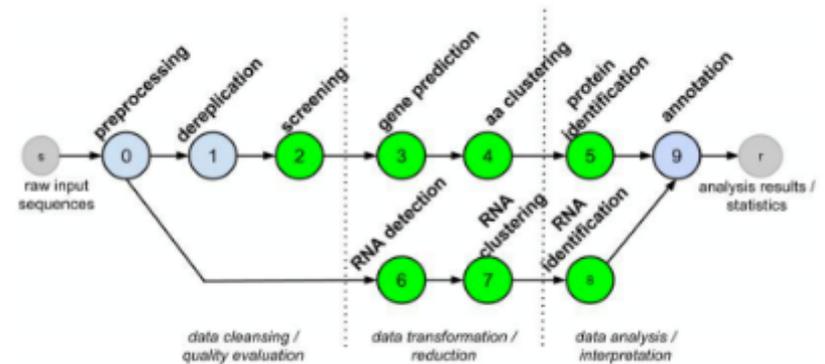
rossc3@rpi.edu

Summer of CODES

13 July 2015

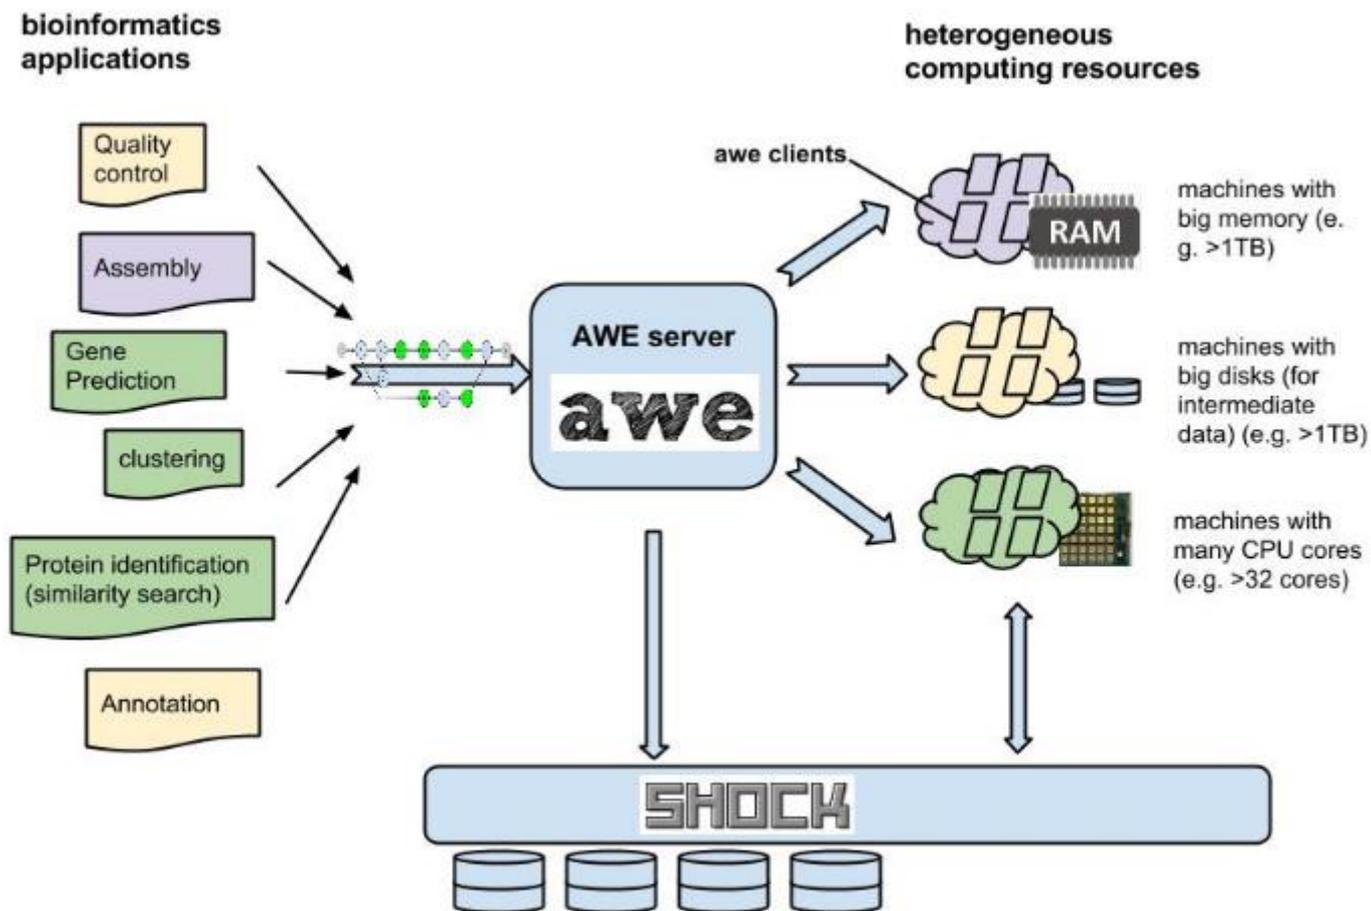
MG-RAST

- A metagenomic analysis service developed at ANL
- Takes in raw DNA sequence data for analysis
- Job: analyzes single metagenomic data set
- Comprised of:
 - AWE workflow management system
 - Shock data management system
- clients at one or more sites get assigned workunits to compute and receive their data from centralized Shock server



W. Tang, J. Bischof, N. Desai, et. al, "Workload Characterization for MG-RAST Metagenomic Data Analytics Service in the Cloud," *in Proc. of IEEE International Conference on Big Data, 2014*

Current MG-RAST infrastructure



W. Tang. "Building an Integrated Genomic Big Data Platform and Addressing Cloud Resource Management Challenges," 2014

Motivation for using Shock proxies

- Centralized Shock server that receives all data download and upload requests
- Adding in proxies to provide load balancing of data requests
- Proxies store the large amount of intermediate data generated by clients
- Use LRU to evict data and send to Shock server
- Goals of this work:
 - Explore configurations to improve MG-RAST performance
 - Distribute data movement requests
 - Improved job response time

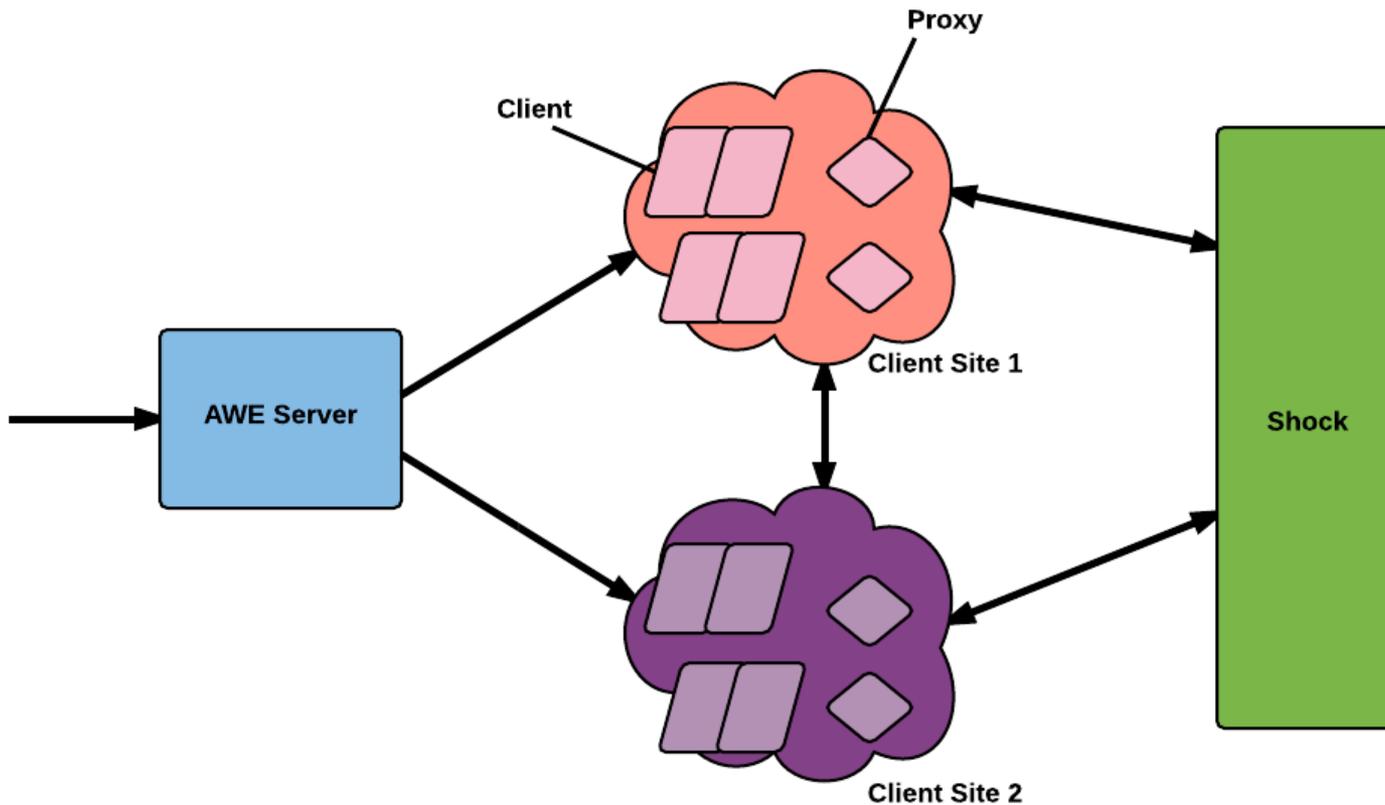
AweSim

- Coarse grained simulation of MG-RAST workflow
- Built using ROSS and CODES
- Developed to evaluate data-aware scheduling in MG-RAST for multisite setups
 - Best fit: schedule only most computationally expensive task in a job at the remote site
 - Greedy: schedule most computationally expensive task type at remote site. If none queued, look for the next most computationally expensive task type
- LPs: AWE Server, AWE clients, Shock server, router
- Peer to peer network configuration
- Uses MG-RAST production traces

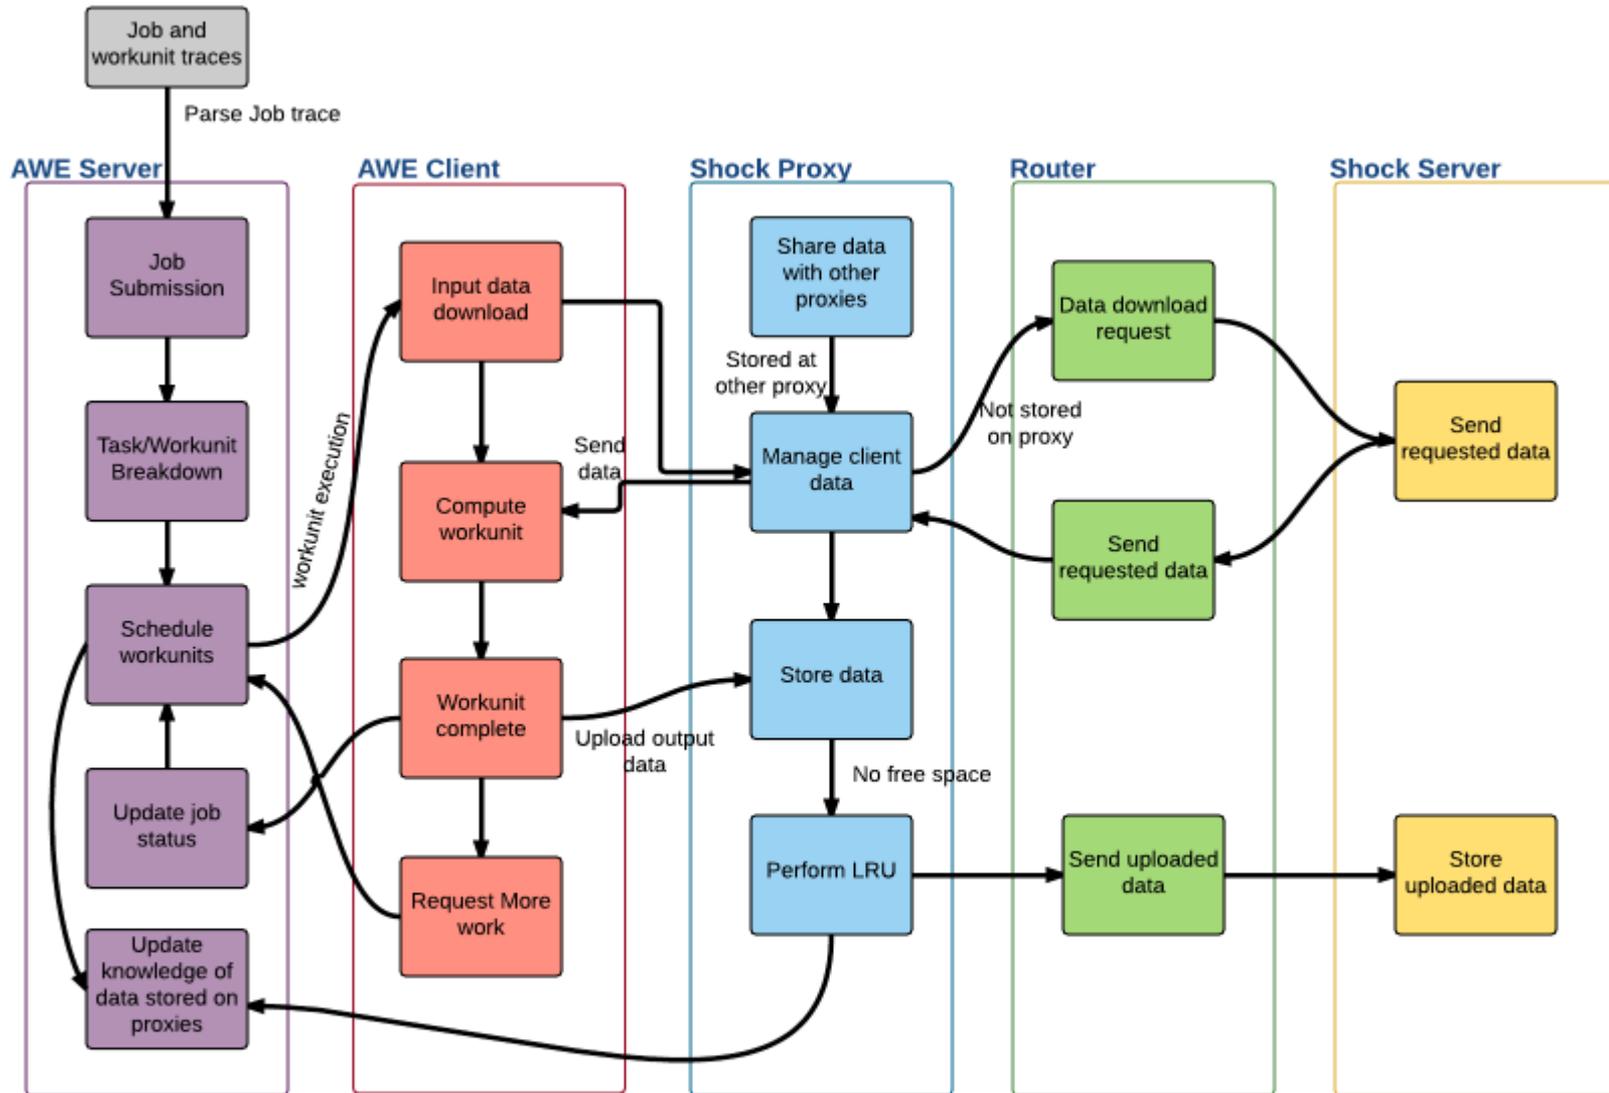
Shock Proxy

- Adding a Shock Proxy LP
- Has a configurable amount of storage
- Sim can have one or more proxies at each site
- AWE clients now communicate with proxies instead of directly with Shock
- AWE server knows where data is stored when scheduling workunits to clients

MG-RAST System with Proxies



AweSim Flow Chart



Preliminary Study - Configuration

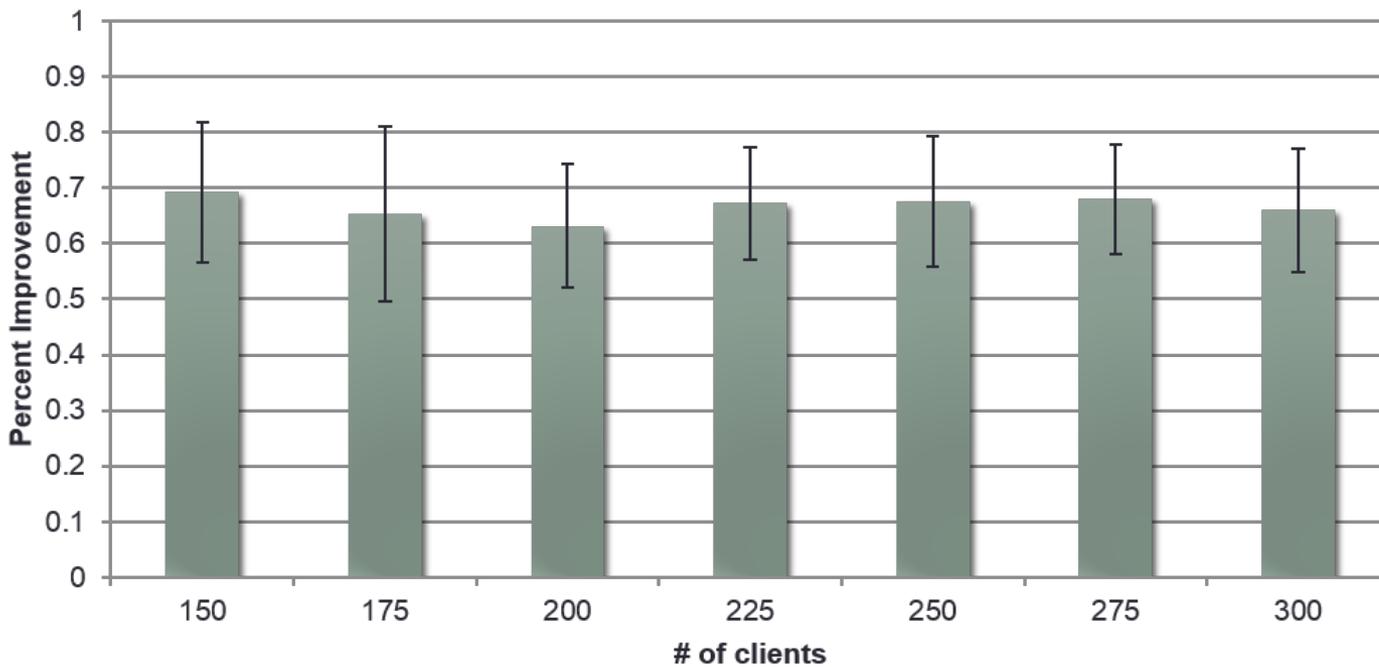
- # of clients: 150 to 300
- # of proxies: 1 to 300
- Storage size available on the proxies: 5 to 100 TB
- Collect stats on:
 - Data movement overhead
 - Response time for jobs
 - Client utilization
 - Accesses to Shock
 - Amount of data uploaded and downloaded
- Preliminary results are using one month of traces from MG-RAST
 - 2976 jobs
 - 29760 tasks
 - 57565 workunits

Validation of Simulation

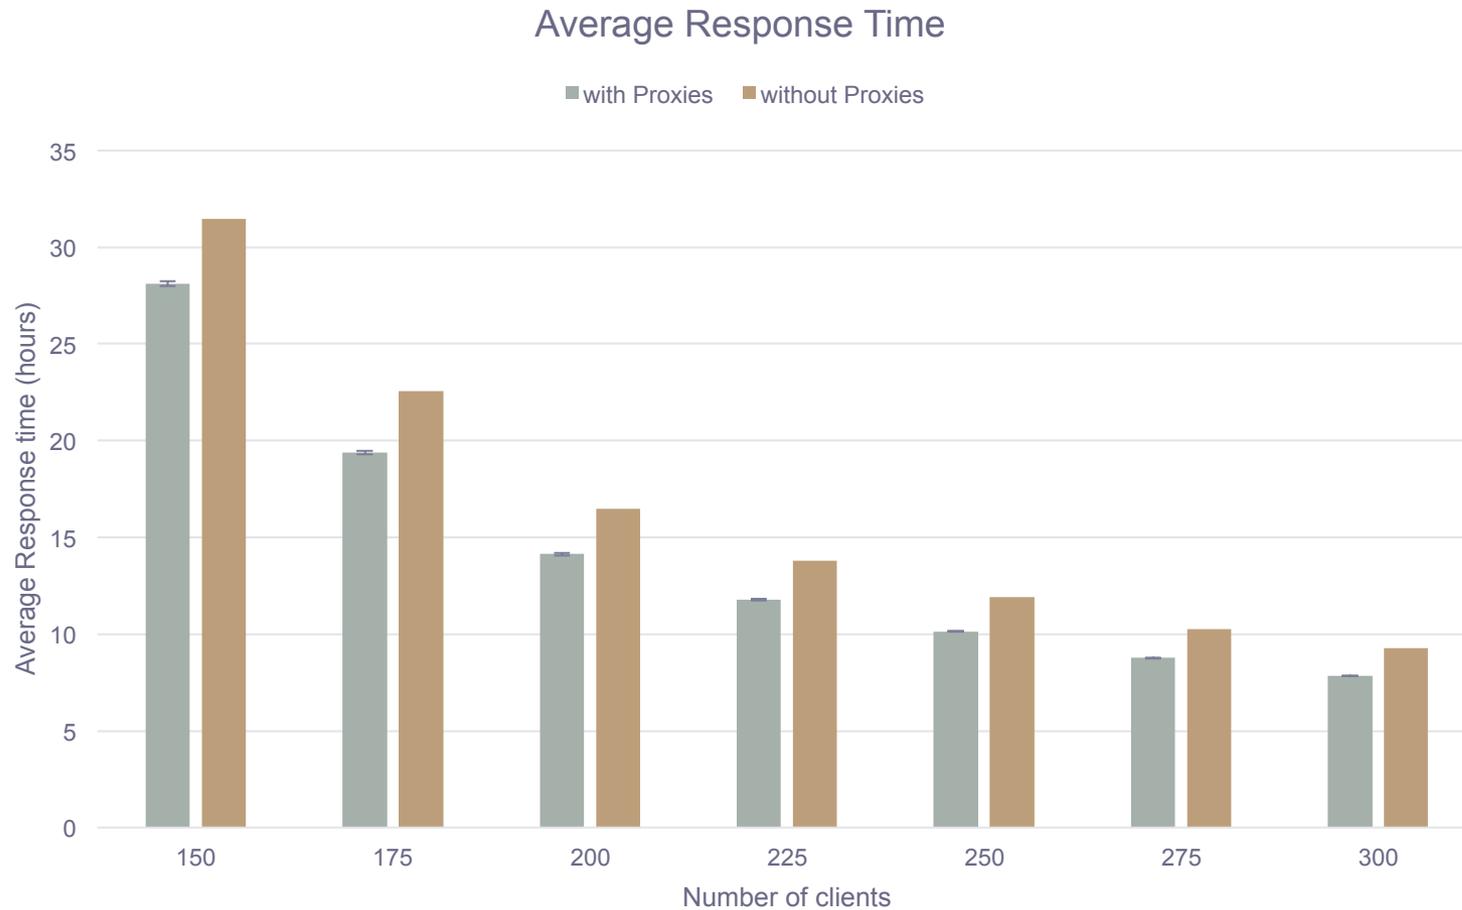
- Compared results to previous version of AweSim without proxies
- Number of accesses to central Shock server decreased by about 70%
 - 115130 accesses w/o proxies vs 34380 accesses with proxies
- Amount of data downloaded from Shock has substantially decreased with the addition of proxies
 - 26 TB downloaded w/o proxies vs 4 TB with proxies
- Average busy rate of clients stays approx. the same for same number of clients, regardless of # of proxies and storage size on proxies

Preliminary Results

Average relative data movement overhead improvement



Preliminary Results



Next steps

- Expand current simulation to two or more sites with their own set of clients and proxies
- Scale up to larger number of clients, larger input traces
- Add in more functionality to Shock proxies
 - Data-aware protocol for proxy selection

Questions?