RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Autonomics, Cyberinfrastructure Federation, and Software-Defined Environments for Science

Manish Parashar
NSF Cloud and Autonomic Computing Center (CAC)
Rutgers Discovery Informatics Institute (RDI[2])
Rutgers, The State University of New Jersey
http://parashar.rutgers.edu/

With J. Diaz-Montes, M. AbdelBaky, M. Zou, I. Rodero R. Tolosana, I. Petri, O. Rana, …

RUTGERS

Pervasive Computation/Data Ecosystems

A data-driven and information-rich computational ecosystems at the network edges that seamlessly and symbiotically combine data and computing power to model, manage, control, adapt and optimize virtually any realizable sub-system of interest

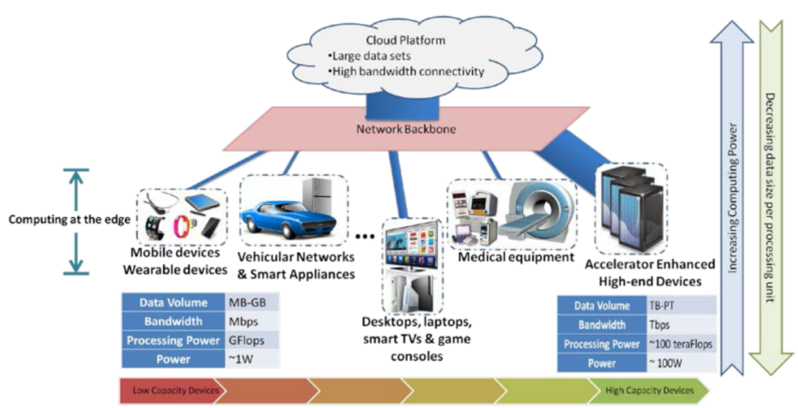

RUTGERS

Computing in Place / Computing at the Edges

- Leverage resources and services at the logical extreme of the network and along the data path to increase the semantic content of data while reducing its volume / transform data into knowledge and insights that are actionable
- Exploit the rich ecosystem of data and computation resources at the edge based on the principle that **data is not moved** from data generation
- Identify the **high level of concurrency** that is pervasive throughout the ecosystem as the key to realizing scalable data-centric applications
- Usecases span precision medicine, instrumented oilfields, disaster management, etc.

RUTGERS

Outline

- Autonomics, clouds/federated computing, software defined systems, …. and Science
- Initial explorations with autonomic federation using CometCloud
- Towards a software-defined federated infrastructure for science
- Conclusion

RUTGERS

CLOUDS, FEDERATED COMPUTING, SOFTWARE DEFINED SYSTEMS

RUTGERS

The Lure of Clouds

- Cloud services provide an attractive platform for supporting the computational and data needs of academic and business application workflows
- Cloud paradigm:
  - "Rent" resources as cloud services on-demand and pay for what you use
  - Potential for scaling-up, scaling-down and scaling-out, as well as for IT outsourcing and automation
- Landscape of heterogeneous cloud services spans private clouds, public clouds, data centers, etc.
  - Heterogeneous offering with different QoS, pricing models, availability, capabilities, and capacities
  - Hybrid cloud infrastructures could integrate private clouds, public clouds, and data centers
- Novel dynamic market-places where users can take advantage of different types of resources, quality of service (QoS), geographical locations, and pricing models
- Cloud federations extend as-a-service models to virtualized data-centers federations

RUTGERS

Cloud Federation

- Cloud Bursting (scaling out to a cloud when needed)
  - Extending local cluster to a cloud with different scheduling policies (M. D. de Assuncao et. al)
  - Extending Austrian Grid with a private cloud (S. Ostermann et. al)
  - Extending grid resources to a Nimbus cloud (C. Vazquez et. al)
- Hybrid Grid and Cloud
  - Creating a large-scale distributed virtual clusters using federated resources from FutureGrid and Grid'5000 (P. Riteau et. al)
  - Infrastructure to manage the execution of service workflows in a union of a grid and a cloud (L. F. Bittencourt et. al)
- Cloud of Clouds
  - Federation of Amazon EC2 and NERSC's Magellan cloud (I. Gorton et. al)
  - Using Pegasus and Condor to federate FutureGrid, NERSC's Magellan cloud and Amazon EC2 (J.-S. Vockler et. al)
- Federation Models
  - Composing cloud federation using a layered service model (D. Villegas et. al)
  - Cross-federation model using customized cloud managers (A. Celesti et. al)
  - A reservoir model that aims at contributing to best practices (B. Rochwerger et. al)

RUTGERS

Clouds as Enablers of Science

- Clouds are rapidly joining traditional CI as viable platforms for scientific exploration and discovery
- Possible usage modes:
  - Clouds can simplify the deployment of applications and the management of their execution, improve their efficiency, effectiveness and/or productivity, and provide more attractive cost/performance ratios
  - Cloud support the democratization
  - Cloud abstractions can support new classes of algorithms and enable new applications formulations
  - Application driven by the science, not available resources -- Cloud abstractions for science?
- Many challenges
  - Application types and capabilities that can be supported by clouds?
  - Can the addition of clouds enable scientific applications and usage modes that are not possible otherwise?
  - What abstractions and systems are essential to support these advanced applications on different hybrid platforms?

RUTGERS

Cloud Usage Modes for Science

- ***HPC in the Cloud*** – outsource entire applications to current public and/or private Cloud platforms

- ***HPC plus Cloud*** – Clouds complement HPC/Grid resources with Cloud services to support science and engineering application workflows, for example, to support heterogeneous requirements, unexpected spikes in demand, etc.

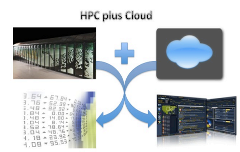

- ***HPC as a Cloud*** – expose HPC/Grid resources using elastic on-demand Cloud abstractions

RUTGERS

Federated Computing for Science (I/II)

- Emerging applications can have large and diverse compute and data requirements
- Federated computing is a viable model for effectively harnessing distributed resources
  - Combine capacity, capabilities
- HPC Grid Computing - monolithic access to powerful resources shared by a virtual organization
  - Lacks the flexibility of aggregating resources on demand (without complex infrastructure reconfiguration)
- Volunteer Computing - harvests donated, idle cycles from numerous distributed workstations
  - Best suited for lightweight independent tasks, rather than for traditional parallel computations

RUTGERS

Federated Computing for Science (II/II)

- Current/emerging application workflow exhibit heterogeneous and dynamic workloads, and highly dynamic demands for resources
  - Various and dynamic QoS requirements
    - Throughput, budget, time
  - Unprecedented amounts of data
    - Large size, heterogeneous nature, geographic location
- Such workloads are hard to be efficiently supported on classic federation models
  - Rigid infrastructure with fixed set of resources
- Can we combine the best features of each model to support varying application requirements and resources' dynamicity?
  - Provisioning and federating an appropriate mix of resources on-the-fly

RUTGERS

Software Defined ….

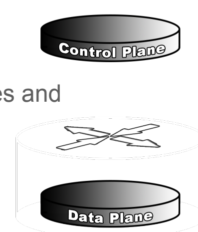

- Software Defined Networks
  - An approach to building computer networks that separates and abstracts elements of these systems (Wikipedia)
  - E.g., separation of control and data plane
- Software Defined Systems
  - Based on software defined networking (SDN) concepts
  - Allow business users to describe expectations from their IT in a systematic way to support automation
  - Enable the infrastructure to understand application's needs through defined policies that control the configuration of compute, storage, and networking, and it optimizes application execution
    - Open virtualization, Policy driven optimization and elasticity – autonomics, Application awareness
- See also software defined data centers, ….

RUTGERS

AUTONOMICS

RUTGERS

Motivations for Autonomic Computing

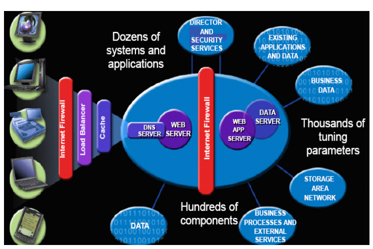

Source: http://www.almaden.ibm.com/almaden/talks/Morris_AC_10-02.pdf

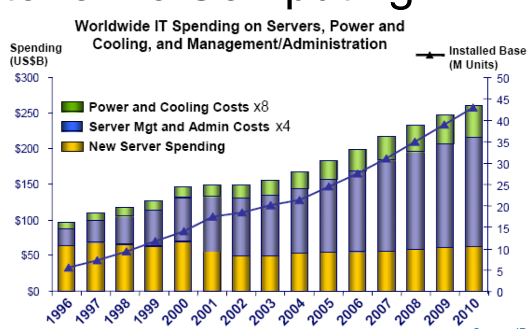

Source: http:idc 2006

2/27/07: Dow fell 546. Since worst plunge took place after 2:30 pm, trading limits were not activated

8/3/07: (EPA) datacenter energy use by 2011 will cost $7.4 B, 15 power plants, 15 Gwatts/hour peak

8/1/06: UK NHS hit with massive computer outage. 72 primary care + 8 acute hospital trusts affected.

8/12/07: 20K people + 60 planes held at LAX after computer failure prevented customs from screening arrivals

Key Challenge

Current levels of scale, complexity and dynamism make it infeasible for humans to effectively manage and control systems and applications

RUTGERS

Integrating Biology and Information Technology: The Autonomic Computing Metaphor

- Current paradigms, mechanisms, management tools are inadequate to handle the scale, complexity, dynamism and heterogeneity of emerging systems and applications
- Nature has evolved to cope with scale, complexity, heterogeneity, dynamism and unpredictability, lack of guarantees
  - self configuring, self adapting, self optimizing, self healing, self protecting, highly decentralized, heterogeneous architectures that work !!!
- Goal of autonomic computing is to enable self-managing systems/ applications that addresses these challenges using high level guidance

"Autonomic Computing: An Overview," M. Parashar, and S. Hariri, Hot Topics, Lecture Notes in Computer Science, Springer Verlag, Vol. 3566, pp. 247-259, 2005.

RUTGERS

Adaptive Biological Systems

- The body's internal mechanisms continuously work together to maintain essential variables within physiological limits that define the viability zone
- Two important observations:
  - The goal of the adaptive behavior is directly linked with the survivability
  - If the external or internal environment pushes the system outside its physiological equilibrium state the system will always work towards coming back to the original equilibrium state

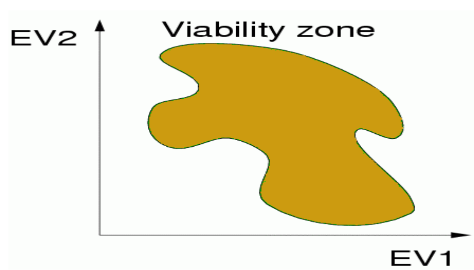

RUTGERS

Ashby's Ultrastable System

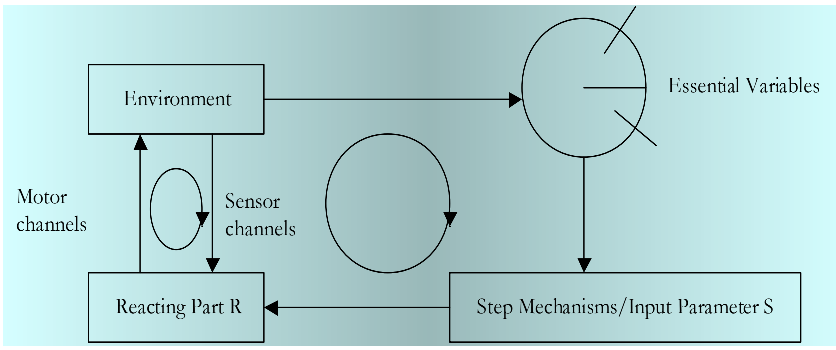

RUTGERS

Autonomic Computing – A Pragmatic Approach

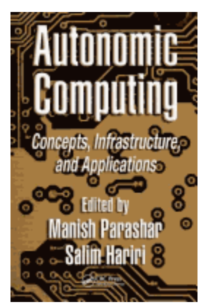

- Separation + Integration + Automation !
- Separation of knowledge, policies and mechanisms for adaptation
- The integration of self–configuration, – healing, – protection,–optimization, …
- Self-* behaviors build on automation concepts and mechanisms
  - Increased productivity, reduced operational costs, timely and effective response
- System/Applications self-management is more than the sum of the self-management of its individual components

M. Parashar and S. Hariri, ***Autonomic Computing: Concepts, Infrastructure, and Applications,*** CRC Press, Taylor & Francis Group, ISBN 0-8493-9367-1, 2007.

RUTGERS

Autonomic Computing Research

- Integrates and advances several fields
  - Distributed computing
    - Algorithms and architectures
  - Artificial intelligence
    - Models to characterize, predict and mine data and behaviors
  - Security and reliability
    - Designs and models of robust systems
  - Systems and software architecture
    - Designs and models of components at different IT layers
  - Control theory
    - Feedback-based control and estimation
  - Systems and signal processing theory
    - System and data models and optimization methods

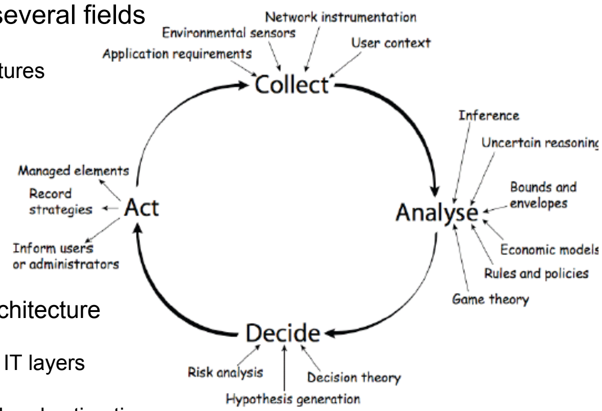

(From S. Dobson et al., ACM Tr. on Autonomous & Adaptive Systems, Vol. 1, No. 2, Dec. 2006.)

- Requires experimental validation

RUTGERS

Autonomic Cloud-based Federation

- Assemble a federated cloud on-the-fly integrating clouds, grids and HPC
  - Cloud-bursting: dynamic application scale-out/up to address dynamic workloads, spikes in demand, and other extreme requirements
  - Cloud-bridging: on-the-fly integration of different resource classes
- Provide policy-driven autonomic resource provisioning, scheduling and runtime adaptations
  - Policies encapsulate user's requirements (deadline, budget, etc.), resource constraints (failure, network, availability, etc.)
- Provide programming abstractions to support Science
  - Master/worker, MapReduce/Hadoop, Workflows

Autonomic Reservoir Management: "Closing the Loop" using Optimization

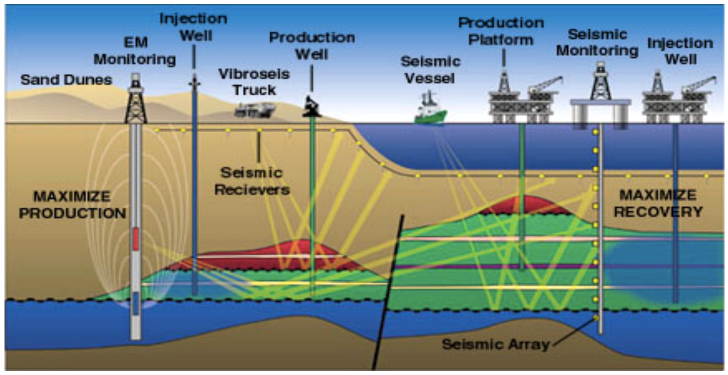

Detect and track changes in data during production.
Invert data for reservoir properties.
Detect and track reservoir changes.

Model Driven

Data Driven

Assimilate data & reservoir properties into the evolving reservoir model.
Use simulation and optimization to guide future production.

RUTGERS

Autonomic Data Streaming & In-Transit Processing

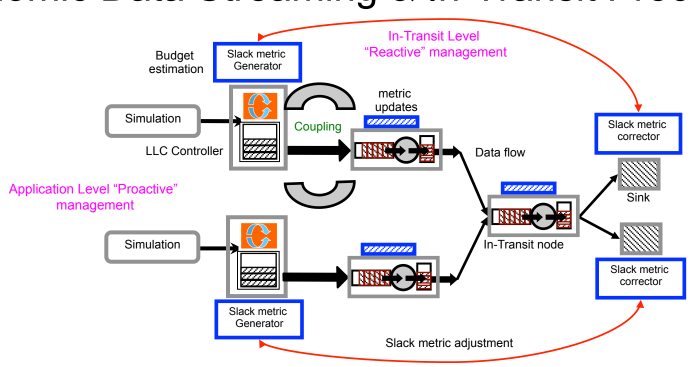

- Application level
  - Proactive QoS management strategies using model-based LLC controller
  - Capture constraints for in-transit processing using slack metric
- In-transit level
  - Opportunistic data processing using dynamic in-transit resource overlay
  - Adaptive run-time management at in-transit nodes based on slack metric generated at application level
    - Adaptive buffer management and forwarding

RUTGERS

EXPLORING FEDERATED INFRASTRUCTURE FOR SCIENCE USING COMETCLOUD

RUTGERS

CometCloud

- Enable applications on dynamically federated, hybrid infrastructure exposed using Cloud abstractions
  - **Services:** discovery, associative object store, messaging, coordination
  - **Cloud-bursting:** dynamic application scale-out/up to address dynamic workloads, spikes in demand, and extreme requirements
  - **Cloud-bridging:** on-the-fly integration of different resource classes (public & private clouds, data-centers and HPC Grids)
- High-level programming abstractions & autonomic mechanisms
  - Cross-layer Autonomics: Application layer; Service layer; Infrastructure layer
- Diverse applications
  - Business intelligence, financial analytics, oil reservoir simulations, medical informatics, document management, etc.

http://cometcloud.org

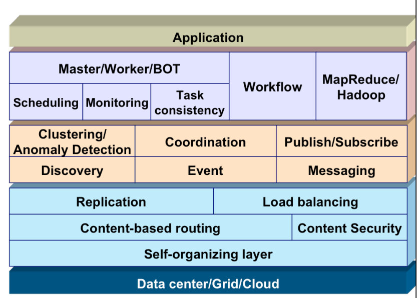

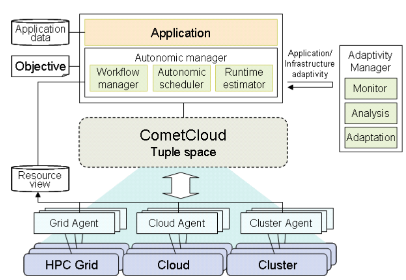

Federated (hybrid) computing infrastructure

RUTGERS

Autonomics in CometCloud

- **Autonomic manager** manages workflows, benchmarks application and provision resources.
- **Adaptivity manager** monitors application performance and adjusts resource provisioning.
- **Resource agent** manages local cloud resources, accesses task tuples from CometCloud and gathers results from local workers so as to send them to the workflow (or application) manager.

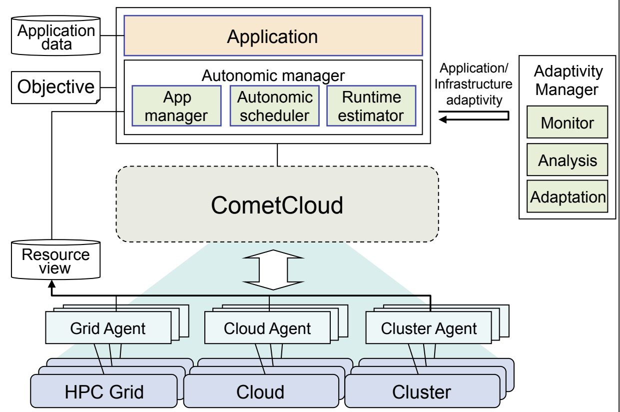

RUTGERS

On-Demand Elastic Federation using CometCloud

- Autonomic cross-layer federation management
  - Resources specified based on availability, capabilities, cost/performance constraints, etc.
  - Dynamically assimilated (or removed)
  - Resources coordinate to:
    - Identify themselves / verify identity
    - Advertise their resources capabilities, availabilities, constraints
    - Discover available resources
- Federation coordinated using Comet spaces
- Autonomic resource provisioning, scheduling and runtime adaptations
- Business/social models for resource sharing

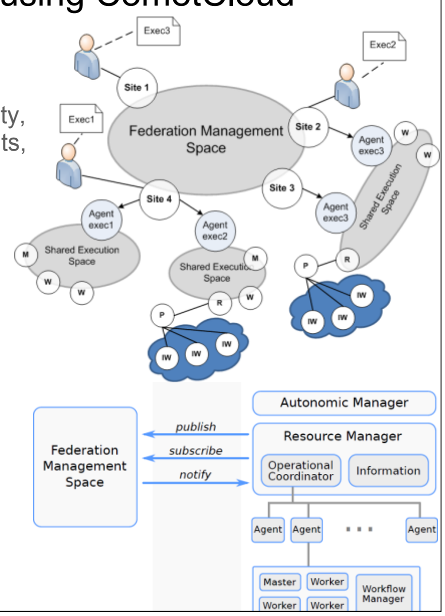

RUTGERS

An Initial Experiment: Fluid Flow in Microchannel

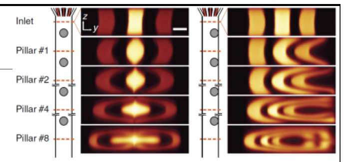

- Controlling fluid streams at microscale is of great importance for biological processing, creating structured materials, etc.
- Placing pillars of different dimensions, and at different offsets, allows "sculpting" the fluid flow in microchannels
- Four parameters affect the flow:
  - Microchannel height
  - Pillar location
  - Pillar diameter
  - Reynolds number

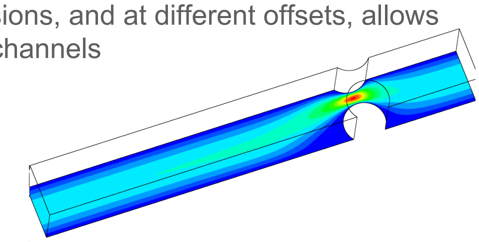

- Each point in the parameter space represents simulation using the Navier-Stokes equation (MPI-based software)
- Highly heterogeneous and computational cost is hard to predict a priori
- Global view of the parameter space requires 12,400 simulations (three categories)

RUTGERS

Experiment Summary [IEEE CiSE 2014]

- 10 different HPC resources from 3 countries federated using CometCloud
- 16 days, 12 hours, 59 minutes and 28 seconds of continuous execution (in spite of failures, etc.)
- 12,845 tasks processed, 2,897,390 CPU-hours consumed, 400 GB of data generated

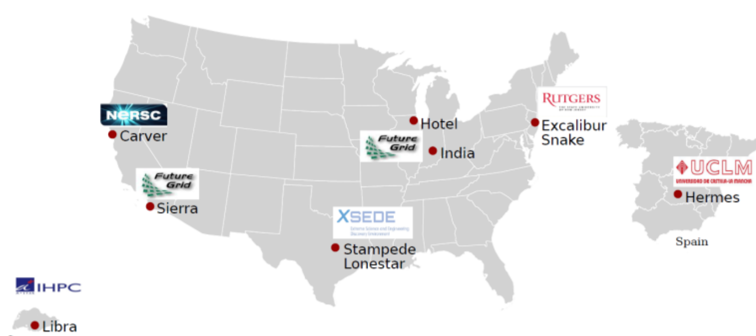

RUTGERS

Experiment Summary (II)

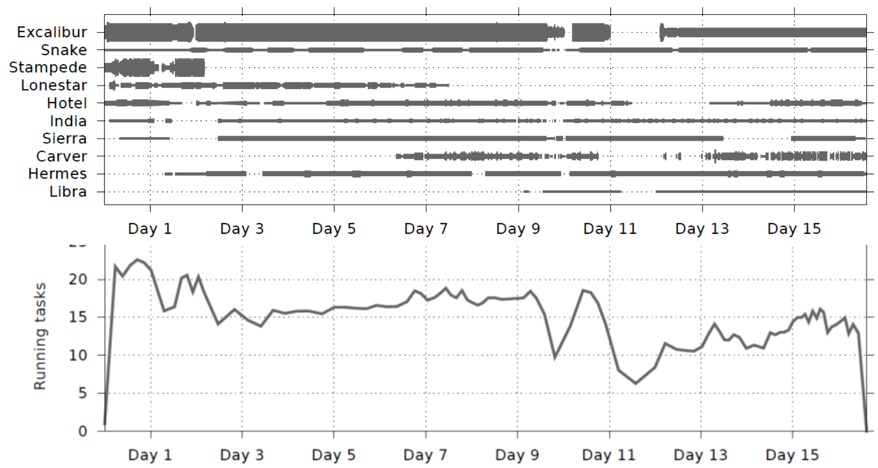

RUTGERS

Throughput and Queue Waiting Time

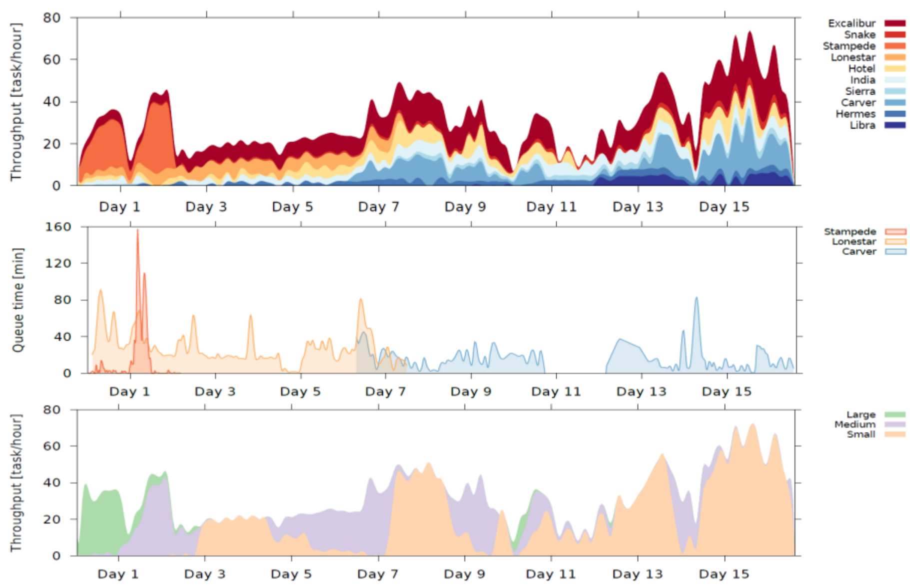

RUTGERS

Science Outcomes*

- The most comprehensive data on the effect of pillars on microfluid channel flow
- Library of flow transformations
- Arranging pillars is possible to perform basic flow transformation
- What is the optimal pillar arrangement to achieve a desired flow output?
- Useful for medical diagnostics, smart materials engineering, and guiding chemical reactions

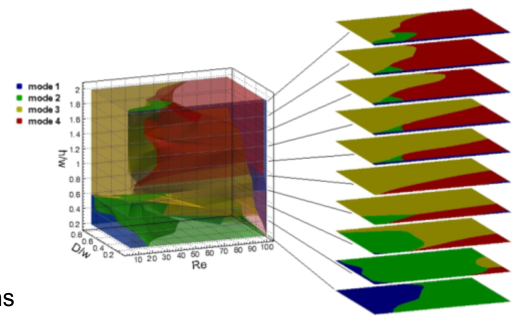

*Published in Nature Communications

Accelerating Protein Folding using Advanced Computational Infrastructure (Rutgers + BMS)

Individual trajectories
- Parallel NAMD trajectories
- Asynchronous communication in cometCloud

Science
- Be smart about using resources
- Commodity hardware versus high end resources
- Terminate or restart resources

Infrastructure
- Federated clouds

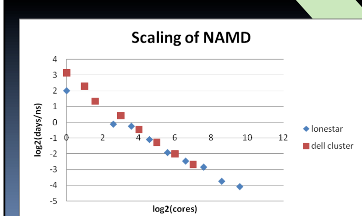

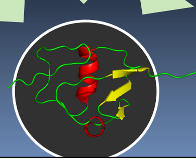

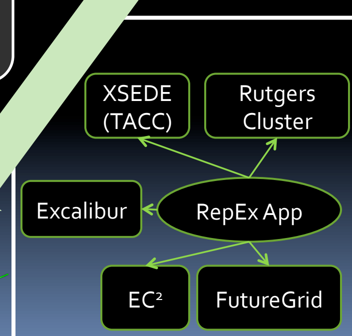

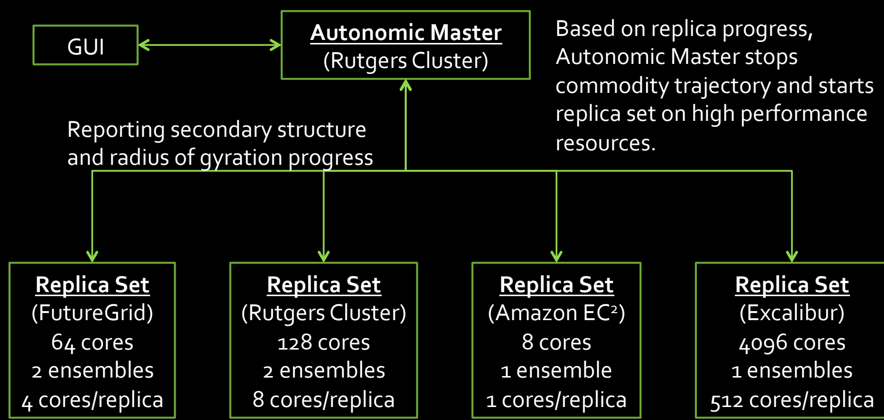

*Could run multiple replicas per temperature to improve likelihood of asynchronous exchange on heterogeneous hardware.

*8 temperatures = 1 ensemble

http://youtu.be/sg2C7N7g5CU

RUTGERS

Other experiments

- Data-Driven Workflows on Federated Clouds [Cloud'14]
- Federating Resources using Market Models [IC2E'14]
- Elastic Federations for Large-scale Scientific Workflows [MTAGS'13]
- HPC plus Cloud Federations [e-Science'10]
- …. [See cometcloud.org]

- Testbed using resource in US (RU, FutureGrid, XSEDE, IBM), UK (Cardiff), Amazon EC2

- Experiments successful…. but can the model be generalized?

RUTGERS

Software Defined Cyberinfrastructure Federations for Science ?

- Combine cloud abstractions with ideas from software-defined environments
- Independent control over application and resources
- Living federations that autonomously adapt to:
  - Changes in the environment
  - Application requirements

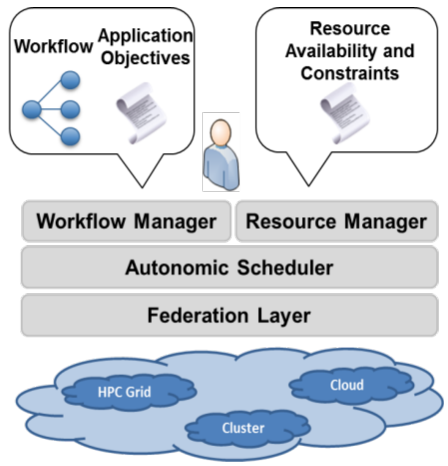

RUTGERS

Software Defined Cyberinfrastructure Federations for Science

User/Provider

Application: Provides performance, resource, and data requirements e.g. number of resources, memory, I/O rate, data dependency, workflow stages (varies in real-time)

Autonomic Manager: Adapt and reconfigure resources and network on the fly to meet objectives and constraints while executing workflow

User: Provide budget, deadline, resource usage policy (can be modified in real-time) , workflow requirement and constraints
Resource Provider: Provide resource constraints e.g. availability (varies in real-time), capacity, capability, cost, access policy

Scientific Applications & Workflows

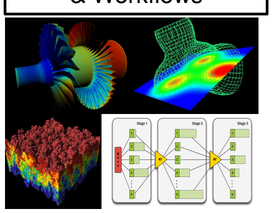

Synthesis a space-time federated ACI

Exposed as a cloud to the application/ workflow

Advanced Cyber-infrastructure Resources

RUTGERS

Software-defined ACI: ACI-as-a-Cloud

- Software defined ACI federations exposed using elastic on-demand Cloud abstractions
- Declaratively specified to define availability as well as policies and constraints to regulate their use
  - Use of a resources may only be allowed at certain times of the day, or when they are lightly loaded, or when they have sufficient connectivity, etc.
  - Prefer certain type of resources over others (e.g., HPC versus clouds or "free" HPC systems versus the allocation-based ones)
  - Specify how to react to unexpected changes in the resource availability or performance
  - Use resources only within the US or Europe due to the laws regulating data movement across borders
- Evolve in time and space -- the evaluation of these constraints provides a set of available resources at evaluation time
- Leverage software-defined networks to customize and optimize the communication channels or software-defined storage to improve data access

RUTGERS

Software-defined ACI: Platform as a Service

- Platform as a Service to decouple applications from the underlying ACI Cloud
- Key components
  1. An API for building new applications or application workflows
  2. Mechanisms for specifying and synthesizing a customized views of the ACI federation that satisfies users' preferences and resource constraints
  3. Scalable middleware services that expose resources using Cloud abstractions
  4. Elasticity exposed in a semantically meaningful way
  5. Autonomics management is critical
- CometCloud provides some of these; currently focusing on 2

RUTGERS

Some Related Efforts

- FED4FIRE (European Union FP7)
  - A common federation framework for developing, adapting or adopting tools that support experiment lifecycle management, monitoring and trustworthiness
- InterCloud (Univ. of Melbourne, Australia)
  - Utility-oriented federation of cloud computing environments for scaling of application services
- Business Oriented Cloud Federation (Univ. of South Hampton, UK)
  - Cloud federation model via computation migration for real time applications; targets real-time online interactive applications, online games
- ….

RUTGERS

Summary

- Emerging CDS&E workflows have dynamic and non-trivial computational/data requirements
  - Necessitate dynamically federated platforms that integrate heterogeneous resources / services
  - Provisioning and federating an appropriate mix of resources on-the-fly is essential and non-trivial
- Software-defined Advanced Cyber-Infrastructure for Science
  - Software defined ACI federations exposed using elastic on-demand Cloud abstractions
  - Application access using established programming abstraction/platforms for science
  - Autonomic management is critical
- Many challenges at multiple layers
  - Application formulation, programming systems, middleware services, standardization & interoperability, autonomic engines, etc.

RUTGERS

The CometCloud Team

Moustafa AbdelBaky
PhD student, Dept. of Electrical & Computer Engr.
Rutgers University
Email: moustafa@cac.rutgers.edu

Mengsong Zou
PhD Student, Dept. of Computer Science
Rutgers University
Email: mz228@cac.rutgers.edu

Javier Diaz-Montes, Ph.D.
Assistant Research Professor,
Dept. of Electrical & Computer Engr.
Rutgers University
Email: javidiaz@rdi2.rutgers.edu

Manish Parashar, Ph.D.
Rutgers Discovery Informatics Institute (RDI^2)
Rutgers University
Email: parashar@rutgers.edu

And many collaborators….

CometCloud: http://cometcloud.org

50

RUTGERS

Thank You!

Manish Parashar
Rutgers Discovery Informatics Institute (RDI^2)
Rutgers, The State University of New Jersey

Email: parashar@rutgers.edu
WWW: http://parashar.rutgers.edu/
CometCloud: http://cometcloud.org