CAREER: Abstractions and Middleware for D3 Science on NSF DCI

Shantenu Jha

http://radical.rutgers.edu

A brief (self-) introduction

- Educational/Professional Preparation:
  - PhD 2004, Syracuse University.
  - Post-doc, Dept. of Chemistry, University College London
  - Research Scientist & Director of CI: LSU (2007-10)
  - Assistant Professor: Rutgers (2011-14)
  - Associate Professor: Rutgers (2014-)

Research at the triple point of Cyberinfrastructure R&D, Computational Science & Applied Computing.

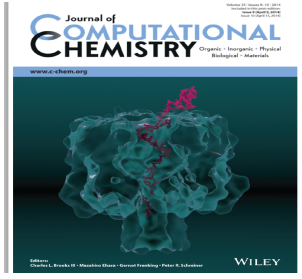

Distributed Scientific Applications

Scalable Algorithms & Methods

Distributed Computing Infrastructure

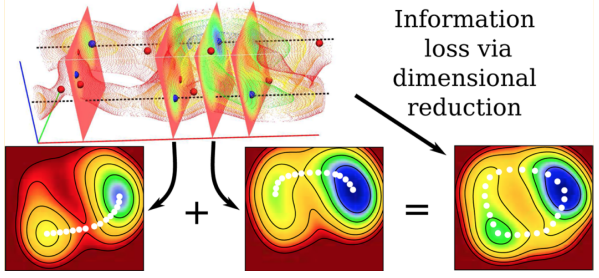

Open Science Grid

XSEDE
Extreme Science and Engineering Discovery Environment

Context

High-Performance Distributed Computing Infrastructure (DCI) as a fundamental enabler of both big project science and long-tail science.

- Need to Scale-up, Scale-out, and Scale-across heterogenous DCI
  - Access Tier 0/1/2 resources, clouds, data repositories, ...
  - Exploit uniqueness and geo-locality of diverse resource platforms.
- ***Need for collective utilization, i.e. whole is > sum of the parts.***

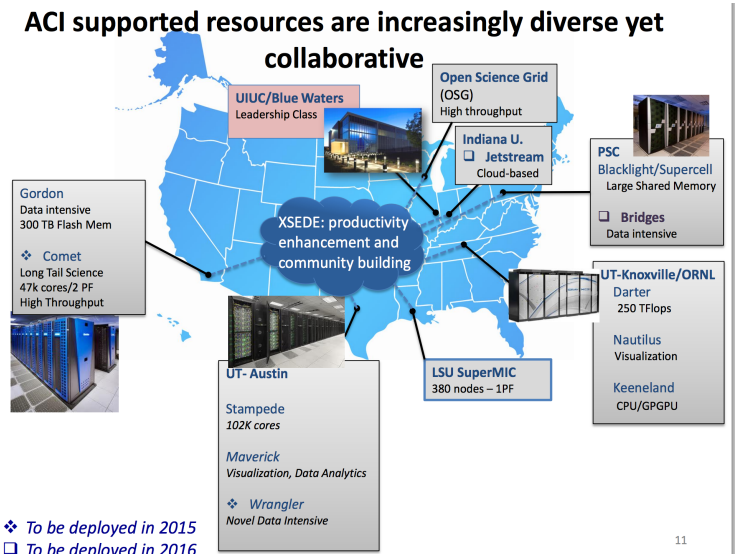

Image/Slide Courtesy Irene Qualters (NSF/ACI).

DCI Status Quo

We are still learning how to architect large-scale production DCI. Missing design principles and abstractions.

- Applications are characterized by “gluing” it to a platform.
  - Brittle, bespoke & local solutions
  - Lack of end-to-end solutions
- Minor difference in starting points of DCI, very different end points.
  - OSG vs XSEDE
- Inability to reason about spatial-temporal execution
  - *“Where/how can I best adapt my application to a DCI?” vice versa?*
  - *“Why did the system allocate this DCI to my application?*
  - *“How will my application perform?”*

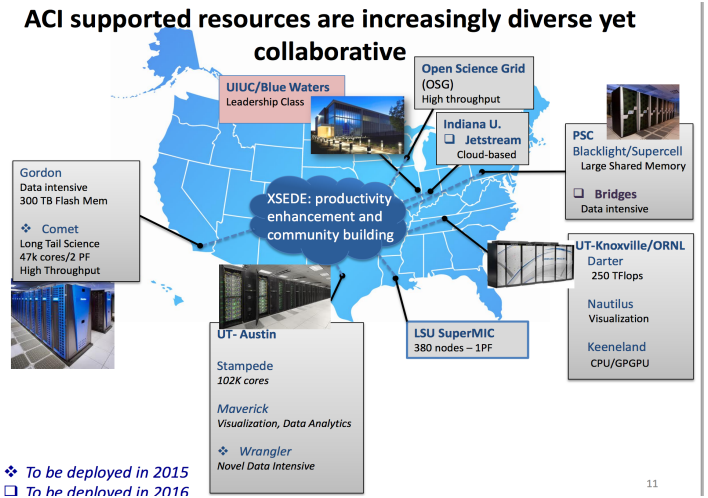

Distributed computing practice for large-scale science and engineering applications, Shantenu Jha et al, Concurrency and Computation: Practice and Experience Volume 25, Issue 11, pages 1559–1585, 10 August 2013 (doi: 10.1002/cpe.2897)

DCI Status Quo

We are still learning how to architect large-scale production DCI. Missing design principles and abstractions.

- Applications are characterized by “gluing” it to a platform.
  - Brittle, bespoke & local solutions
  - Lack of end-to-end solutions
- Minor difference in starting points of DCI, very different end points.
  - OSG vs XSEDE
- Inability to reason about spatial-temporal execution
  - *“Where/how can I best adapt my application to a DCI?” vice versa?*
  - *“Why did the system allocate this DCI to my application?*
  - *“How will my application perform?”*

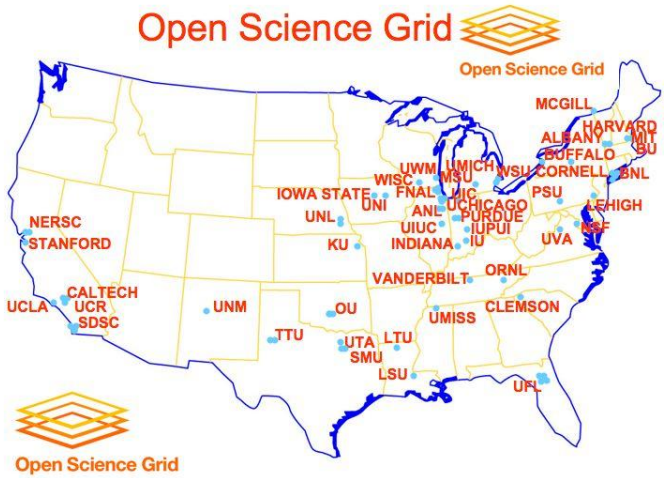

Distributed computing practice for large-scale science and engineering applications, Shantenu Jha et al, Concurrency and Computation: Practice and Experience Volume 25, Issue 11, pages 1559–1585, 10 August 2013 (doi: 10.1002/cpe.2897)

A Pore Man's View of the TeraGrid/XSEDE

- Developed algorithm that would exploit task-level parallelism.
- **2005-09**: Tried running O (1000) simulations on many supercomputers. *Did not work!*
- Reverted to running hundreds of simulations on sequential and single resources.

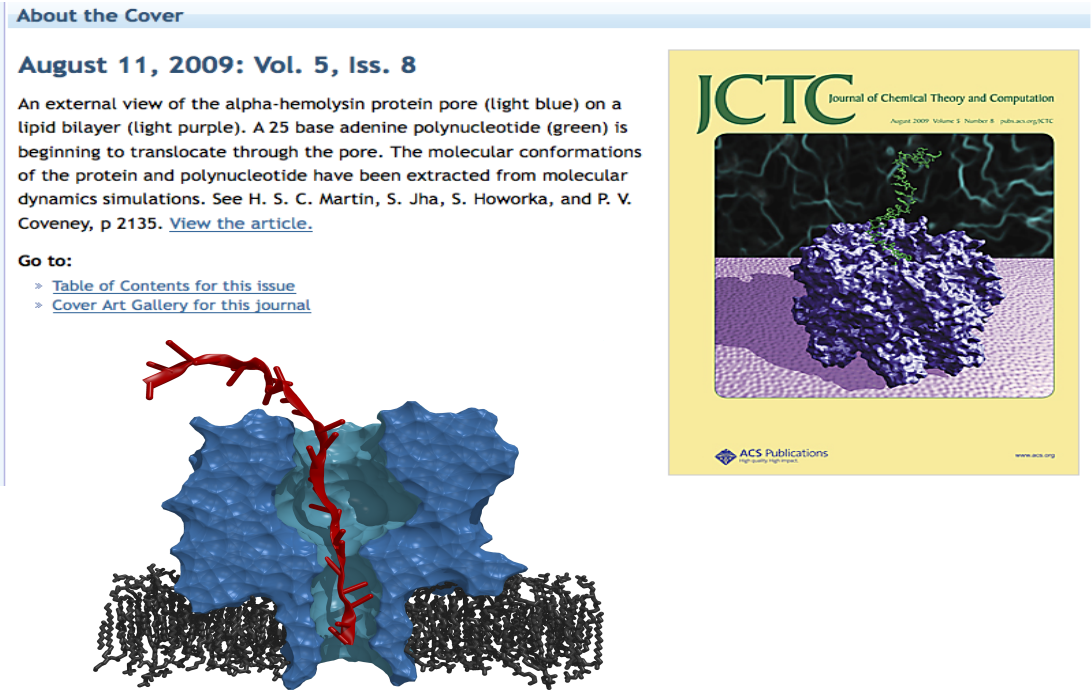

Project Objectives

Grand Challenge: How to federate heterogeneous dynamic DCI? How to use multiple sites effectively? Design and Develop Scalable Applications?

- Primary Research Objective:
  - Identify, develop abstractions and approaches for DCI collective utilization
  - Address the three considerations.
- Development Objective:
  - Prototype implementation that hides complexity, yet exploits diversity
    - “Hiding complexity is not the same as exposing simplicity”
- Validation and Assessment:
  - Useful for application types/characteristics.
  - Generalizable to diverse infrastructure.
  - Enable spatio-temporal reasoning about executing distributed workloads
    - Abstractions → Models → Reasoning about execution on DCI

Pilot-Data: Abstractions for D3 Applications

Pilot-Jobs: A system that generalizes a placeholder job, allows application-level control and supports multi-level scheduling via a scheduling overlay.

- Feature of Pilot-Jobs:
  - Decouples workload from resource management, thereby flexibility
- Pilot-Data: Decoupling of the specification of dynamic data from storage location/type.
  - Implement Pilot-Data (C-D placement) in conjunction with Pilot-job capability
  - Mathematical model of pilot-data to support spatio-temporal reasoning.

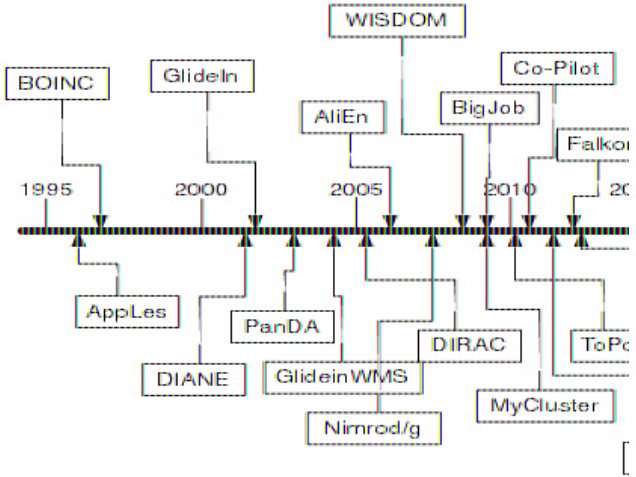

- P* Model (2010-12) of Pilots addresses inconsistent definition and terminology.

Picking the Right Problem (1)

One needs to find the right scope – too ambitious vs. not ambitious enough, how did you pick the problem?

- Manageable slice of the “Grand Challenge” Pie.
  - Necessary and minimally complete.
- Achievable (given effort and time) AND interesting AND impactful.
  - Arguably a challenge with every proposal(!)
  - Tension between achievable (~5 yrs) and impactful?
- Several layers of the Pie.
  - Theory, modeling, systems & experiments “in the wild”.
  - Development/implementation en route to validation.
  - Integration with some application drivers.

Picking the Right Problem (2)

Integration of research and education is a critical aspect of the proposal. How did you plan your education component as well as your plan to integrate this with your research component?

- Rutgers Internal:
  - Develop a scientific computing course that goes beyond numerical analysis and algorithms, to discuss types of infrastructure.
    - Sapir-Whorf Hypothesis: “Language shapes thoughts”
    - How does infrastructure influence how you do research?
  - Work with Byrne, Douglass College and Aresty Programs to bring project into undergraduate research and training curriculum
- External:
  - NSF/ACI Futuregrid (→ NSF Cloud Chamelon) Testbed education program. Similarly with XSEDE.

Process

The deadline is in July, when did you start?

- Starting to think about problem >1 year in advance.
- Attended Rutgers internal CAREER workshop (April-May)
    - General proposal and informational
    - Helpful resources
- Visited NSF to meet Program Director (Barry Schneider) May
    - Developed Education Plan soon after
    - Developed Collaboration with FutureGrid June
- Developed Research Collaboration (International) also in June
    - Unsure whether I wanted XSEDE letter of support
- Writing the proposal: 10-15 days!

Hofstadter's Law: It always takes longer than you expect, even when you take into account Hofstadter's Law.

Process

Did you talk to one or more NSF PDs, how did you pick the PD, your experience – how difficult was it finding the right match

- Natural choice: XD Program Director (Barry Schnieder)
  - Had interacted with Dr. Schneider on another project.
  - Did explore the possibility of Chemistry “co-submission”.
- Discussed education and research plan with PD.
  - PD engaged (surprisingly intensely!)
- In general, the meeting and discussion was informative.
  - Note: Dr Schneider wasn’t on the panel, as there were changes within ACI by time of submission/review!

Other Points

- *Did you have a mentor, how much did s/he help? Did you receive support from your department, do you have any suggestion on how to go about it etc. …*

  No mentor for the proposal.

- *Did you have some example projects if so how did you get them?*

  Not project, but discussion with a previous ACI CAREER awardee

- *...talk about your proposal writing/preparation guidance ..*

  Shared draft of graduation education plan with graduate students

Other Points

What additional advice, from your personal experience, would you give to PIs as they plan and write their CAREER proposals?

"The only good thing to do with good advice is pass it on; it is never of any use to oneself." -- Oscar Wilde

- You define the problem; the problem is likely to define you too
  - Its the defining problem of your pre-tenure phase and likely your tenure package!

- So be a dreamer, be a romantic, think blue skies or yonder, but whatever you do, define a problem that you feel passionate about!

CAREER: Abstractions and Middleware for D3 Science on NSF DCI

NSF ACI 1253644

shantenu.jha@rutgers.edu
http://radical.rutgers.edu