



THE UNIVERSITY of
NEW MEXICO

Center for Advanced Research Computing

*UNM Center for Advanced Research Computing:
A Strategic Research Emphasis Center of the University of New Mexico*

[Super]computing on the Edge: Challenges and Opportunities

Susan R. Atlas

Department of Physics and Astronomy

Co-Director, UNM Cancer Center Shared Resource for Genomics & Bioinformatics



Director, Center for Advanced Research Computing

Celebrating 20 years of academic supercomputing at the University of New Mexico , 1994-2014

What we do

CARC's mission is to enable **excellence in research** in science, engineering, biomedicine, humanities, and the arts, through support for parallel supercomputing, large-scale informatics, high-throughput computing, and advanced visualization, and by providing leadership to enhance computing-based **interdisciplinary research and education** at the University.

Research Areas

 Arts Technology • Computational Biology & Bioinformatics • Biophysics and Nanoscale Systems • Chemistry and Chemical Biology • Civil Engineering • Climate and Weather Modeling • Cyberinfrastructure • Digital Arts & Humanities • Electromagnetics • Energy Grid Modeling • Finite Element and Computational Fluid Dynamics Modeling • High Performance Computing and Scalable Systems • Image Processing • Materials Physics • Mechanical Engineering • Molecular Genomics • Nanoscience & Microsystems • Neuroscience & Genetics • Observational Astronomy • Open Science Grid • Particle Physics • Political and Social Science Modeling • Quantum Materials & Devices • Systems Biology & Biomedicine • Visualization and Virtual Environments 

How we do it: Leverage, leverage, leverage (repeat)

- All compute cycles and (reasonable) storage **free of charge** to UNM faculty, students, and staff (no cost recovery)
- Base funding (staff salaries + operations) from **OVPR F&A funds; \$560,000**. Director reports to Provost via OVPR; close ties with IT (regular 1-1 meetings with CIO; IT Cabinet)
- Staff at full complement: 2.5 FTE admin/operations; 4.6 technical; .2 Director
- Center ~100% over budget FY '09; on-budget past 4 years
- Operate as a scrappy, entrepreneurial 'start-up'
- Shared staff positions with UNM Cancer Center, IT, Psychology, MechEng, ARTS Lab
- 'Resident research groups' lead to new research areas, collaborations, grants
- \$1.3M renovation/network upgrade (2010 and 2013) via partnerships with UNM IT, Interdisciplinary Film and Digital Media Program; \$200k HVAC upgrade (2013) funded by Provost, OVPR, and University Libraries
- Acquisitions of supercomputers, clusters, storage via consortia, grants, partnerships: community + hosted
- "Hand-me-down" supercomputers/PCs-into-clusters
- Industrial Partnership Program (cost center) – in process



What keeps me up at night: “solvable” issues

- **Flat F&A budget; insufficient staff to grow/develop (mostly) novice user base**
 - **Work with Provost’s office to obtain state lines** (staff “tenure”)
 - **Partner with faculty to write technical staff into proposals** (up to .5 FTE senior staff, .7 FTE junior). Affordable growth with increased diversity/expertise (4.75 FTE = 7 people)
 - **Outsource or barter to obtain specialized expertise and resources** (IT fee-for-service ⇔ Center networking and firewalls; free cluster hosting ⇔ student-led accelerator workshops; student mentoring/internships ⇔ peer-to-peer user support, new technologies (Hadoop))
- **Shortage of qualified (advanced) systems staff**
 - **Hire smart kids and let ‘em loose** (also helps with salaries :-)
- **Repercussions from failed \$14M NM state supercomputer center**
 - **Ignore it. Work with legislators, governor’s office to highlight *bona fide* success stories**
 - **Avoid forced tie-ins to economic development. Let the science speak for itself**
- **Successful university center = takeover target and/or internal competition**
 - **“All computing is engineering”; cool apps are good PR [SOE]**
 - **Supporting infrastructure is boring: supercomputing is fun and exciting [IT, HSC IT]**
 - **“Not invented here” syndrome: anything you can do, we can do better; trust issues [SOM]**

Keep Calm and Carry On

“The rule of seven touches” and “the long tail of science”

"To create new norms, you have to understand people's existing norms and barriers to change. You have to understand what's getting in their way... [by] working with [them]... one by one... In the era of the iPhone, Facebook, and Twitter, we've become enamored of ideas that spread as effortlessly as ether. We want frictionless "turnkey" solutions... we prefer instructional videos to teachers... People and institutions can feel messy and anachronistic. They introduce, as the engineers put it, uncontrolled reliability. But technology... [is] not enough.... **Every change requires effort, and the decision to make that effort is a social process.** This is something that salespeople understand well... Evidence is not remotely enough... however a strong case you may have. You must also apply "the rule of seven touches." Personally "touch" the [person] seven times, and they will come to know you; if they know you, they might trust you; and, if they trust you, they will change... **[H]uman interaction is the key force in... speeding change.**"

– Atul Gawande, “Slow ideas” (The New Yorker, July 29, 2013)

What *really* keeps me up at night



- “I had some dreams they were clouds in my coffee...”
The cloud can solve all advanced computing problems.
It’s cheap and it’s easy. Press a button, run a workflow, publish a Science article.
- (Related) ***The University does not need an advanced computing center.***
 - This amounts to a subsidy of a few professors/users. There are other priorities.
 - Computers are fast and cheap; I can put some in a closet in my lab, pay \$5k for 100 TB of storage, and let my postdoc manage the system.
 - Computational science isn’t real science. It is not the province of ‘true’ research.
 - Computational science is a low-level service activity, akin to technician work.
 - Corollary 1. Anyone can do it; specialists/experts are not required
 - Corollary 2. Research grants should pay for analysis by the hour, rather than as collaborations
- (Related) ***GUIs, apps, and “software as service” (e.g. MATLAB) make life (too?) ‘easy.’***
Students are graduating without fundamental programming and numerical skills. Old knowledge is ‘rediscovered’: MapReduce as master/slave; ‘data parallelism’ for GPUs.
Are we (collectively) losing advanced technical and analytic skills?
- Big data/fast networks as disruptive technologies: (clinical) genomics = ‘killer app’?
 - Research Storage Consortium (4 large NSF projects + UL)
 - Data librarians complement applications scientists: curate, annotate \Rightarrow merge, correlate, analyze
 - Decade-long collaboration with UNM Cancer Center
 - “The gulf of Lomas”: ***bridging the NSF/NIH cultural divide in a time of amazing possibilities***