Some thoughts on IO/Storage Curricula/Centers

Greg Ganger

PDL Director

Outline/Plan

- What we've been doing at Carnegie Mellon
  - research (PhD education) and classes/program
- One perspective on challenges being faced
  - perception and critical mass
- A few seeds for brainstorming
  - dreams and possible steps

Parallel Data Lab (PDL)

- CS and CE students, staff and faculty exploring the future of storage systems and advancing the state of the art …
  - 14 contributing faculty and approx. 40 students & staff
- Significant industry involvement
  - Companies: APC, EMC, EqualLogic, HP, Hitachi, IBM, Intel, Microsoft, Network Appliance, Oracle, Panasas, Seagate, Sun, Symantec
  - PDL Retreat: 30+ technical leaders from these companies dedicate 3 days to us each Fall
    - another 20+ for one-day Visit Day each Spring

14 years of PDL Research and counting…

- Founded in 1992 by Garth Gibson
- Early-mid 90s: advanced RAID architectures
  - multi-disk failure mechanisms
- Mid-late 90s: object-based storage
  - foundation of OSD standards, Lustre, and Panasas
- Late 90s: IP-based network-attached storage
- Early 00s: increased storage device functionality
  - device-embedded security
  - more expressive storage interfaces
- Current focus: Self-* Storage
  - Goal: “lights out” administration for storage
    - by integrating automation and survivability into system design
  - Building/deploying a storage infrastructure for CMU researchers
    - allows for real experiences and measurements

PDL's Broad Research Portfolio

- Exploring the future of storage infrastructures and advancing the state of the art …
  - new storage architectures
  - storage management and self-*-ness
  - storage system security and reliability
  - home/consumer data organization and search
  - database and file systems
  - designing systems with new storage technologies
  - storage characterization and modeling
  - storage networking and decentralized caching

18-746 (Storage Systems)

- A storage systems course in its fifth year
  - target: seniors and first-year grad students
    - goal: deep, end-to-end understanding of storage systems
    - components, trends, systems, applications, management
  - also working on book (cough)
- Topics from first version (Spring 2001)
  - disks, firmware, and disk performance enhancement
  - interconnects and storage networking
  - file systems and database back-ends
  - multi-disk subsystems and disk arrays
  - distributed file systems and NAS
  - storage management

Towards MS/BS specialization

- Working to educate students on opportunity
  - explaining the job/technology landscape
    - e.g., advertising in 18-200 and OS class
  - explaining class sequences they should take
    - foundations in OS, networking, distributed systems
    - 18-746, of course
  - working to formalize it a little bit
    - something they can list on their resume
- Uphill battle
  - both with students and (to lesser extent) administration
  - see next topic…

Outline/Plan

- What we've been doing at Carnegie Mellon
  - research (PhD education) and classes/program
- One perspective on challenges being faced
  - perception and critical mass
- A few seeds for brainstorming
  - dreams and possible steps

Interrelated challenges

- Lack of critical mass
  - too few academics view storage/IO as their area
    - even some for which it obviously is
  - consequence: difficult to sustain & move forward
- Weakness on perception front
  - too few people recognize importance of storage
  - consequence: difficult to energize others to fund, help, join, etc.

One way to see our thin presence

	Acquisitions	Education
Networking	$61 B	Multi-class sequences at every University
Comp. Arch.	$55 B	Multi-class sequences at every University
Storage Sys.	$25 B	**One upper-level class at ~15 Universities worldwide**

- Note: very rough analysis here!
  - ignores administration costs (and backup)
    - another breakdown: 40% of data center costs is storage
  - ignores PCs/laptops/other such things
  - other assumptions and shortcomings

Vicious cycle for education

- Few faculty -> few classes -> few students -> little momentum -> few faculty…

- Many aspects
  - students aren't prepared for storage jobs
    - EMC Academy, for example
  - students don't apply to grad school for storage
    - primary exception: after few years in storage industry
  - who would teach a storage class (at most schools)?
  - most deans/chairs don't see the point of storage classes
    - after all, other schools don't bother
    - in fact, why even hire such people?

Outline/Plan

- What we've been doing at Carnegie Mellon
  - research (PhD education) and classes/program
- One perspective on challenges being faced
  - perception and critical mass
- A few seeds for brainstorming
  - dreams and possible steps

What are some things we might do?

- Has to be said: apply resources
  - change has to be induced
- Tell your friends
  - a bit of marketing (mottos, pitches, etc.)
- Grand challenges?
  - one traditional way of energizing

Thin presence not just in classes

	Acquisitions	Education	Gov't funding
Networking	$61 B	Multi-class sequences at every University	Dozens of CFPs
Comp. Arch.	$55 B	Multi-class sequences at every University	Dozens of CFPs
Storage Sys.	$25 B	One class at ~15 Universities worldwide	HECURA 2005

- Note: HECURA 2005 is a very very good thing
  - first "storage program" I've seen!
    - thank you, from us all, to everyone who made it happen!
  - if we want to grow/sustain energy focused on storage systems problems…

Two forms of inducing, perhaps

- Generally more storage-targetted resources
  - draws people in (examples: HECURA, security)
  - this is about growing toward critical mass
- Resources invested specifically
  - a small current community means that we have to lean on a limited set of people to lead change
    - some time has to be created for those people, even while they are pulled in many directions
  - this is about guiding the new community

Going to have to do some marketing

- Most of us shrink from it
  - but, can't do that if going to induce this change
- How many of us have explicitly encouraged non-obvious peers to submit their work to FAST?
- Some mottos heard yesterday
  - "storage is hard"
  - "storage is hard because mistakes last forever"
  - "can't just restart when storage fails"
  - "storage is the toughest form of distributed system"
- As one example, slides I've used to market 18-746

What makes storage systems so cool?

1. Combines so many topic areas:
   - hardware meets OS meets networking meets embedded systems meets distributed systems meets security…
   - so, it's just plain cool!

What makes storage systems so cool?

1. Combines so many topic areas
2. This is where great jobs are!
   - Designers and implementers still needed
     - not just testing
   - Continuing growth area for the future
     - The Internet is a network, but the web is a storage system
     - Strong existing companies: EMC, Veritas, NetApp, IBM, HP, …
     - and still support for start-ups: Panasas, Spinnaker, and many others

What makes storage systems so cool?

1. Combines so many topic areas
2. Great careers
3. Still so much room to contribute:
   - performance actually matters here
     - in fact, it dominates other parts of system performance in many cases
   - storage management wide open

Storage systems: fun quotes

"I/O certainly has been lagging in the last decade"

- Seymour Cray, 1976

"Also, I/O needs a lot of work"

- David Kuck, 1988

"Storage will dominate our business in a few years"

- Compaq VP, 1998

"In 3 to 5 years, we will start seeing servers as peripherals to storage"

- SUN Chief Technology Officer, 1998

"Scalable I/O is perhaps the most overlooked area of high-performance computing R&D"

- Suggested R&D topic report for 2005-2009

Grand Challenges

- Sufficient visibility may require this
  - hard to convince congress folk with plumbing
    - other option is FUD
- We all know today's coolest instance, yes?
- Needs demonstrations and tech transfer
- Some types:
  - A clearly storage-dominated cool application?
    - index/search of interesting corpus?
  - Something numerical?
    - XXX creates/sec, YYY files indexed/sec, ??