CRA Snowbird'08 Conference

Workshop on the Instrumentation Needs of CISE Research

Preliminary Report

Azer Bestavros

Boston University

July 13, 2008

Workshop Goals

- Assess the nature and needs for instrumentation development, acquisition, utilization, and sharing for purposes of ongoing and anticipated research in different CISE areas

- Discuss the opportunities and limitations of existing funding mechanisms available to the CISE community and provide feedback on possible improvements

Instrumentation Evolution

~ 1980s: instrumentation was mostly concerned with acquisition and providing access to computers and networks

~ 1990s: instrumentation evolved to acquisition of specialized hardware in support of collaborative research projects

~ 2000s: instrumentation extended to development and support of community resources, testbeds, and infrastructures

Key Factors

- Declining cost of hardware and increasing cost of maintenance and support
- Increasingly critical role that software plays in instrumentation
- Emergence of utility computing resources, grids, and clouds
- Broadening of the CISE constituents in need of instrumentation beyond HPC

Funding Mechanisms

- Research grants from NSF, DOE, DOD, industry and other sources
- NSF grants for instrumentation acquisition or development
  - Major Research Instrumentation (MRI) program
  - CISE Computing Research Infrastructure (CRI) program
- Grants from other agencies such as DOE, DOD & NIH (e.g. DURIP)
- Industry funding for equipment acquisition or infrastructure use (e.g. the IBM SUR program)

NSF CISE Instrumentation

- Evolved over the years to match the CISE community needs – RI, MII, CCLI, CRI

- Current CRI solicitation distinguishes between two types of infrastructures:
  - Institutional Infrastructure (II)
  - Community Infrastructure (CI)

NSF/CISE CRI Instrumentation

- CRI II proposals must not be for infrastructure or instruments available through existing CI awards

- CRI CI proposals must show community need and buy-in
  - CI Planning grants of up to $100K for 1 year
  - CI Acquisition, Development, Deployment and/or Operations (CI-ADDO) grants of up to $4 million for durations of up to 4 years

NSF MRI Instrumentation

- Seeks to “increase access to scientific and engineering equipment for research and research training” – awards up to $4M
- Managed under the Office of Integrative Activities (OIA)
- Separate budget that supports various directorates at NSF – an MRI award increases the CISE funding base
- Allocation of funds to a directorate is dependent on “proposal/budget pressure” from that division

CISE's Share of MRI

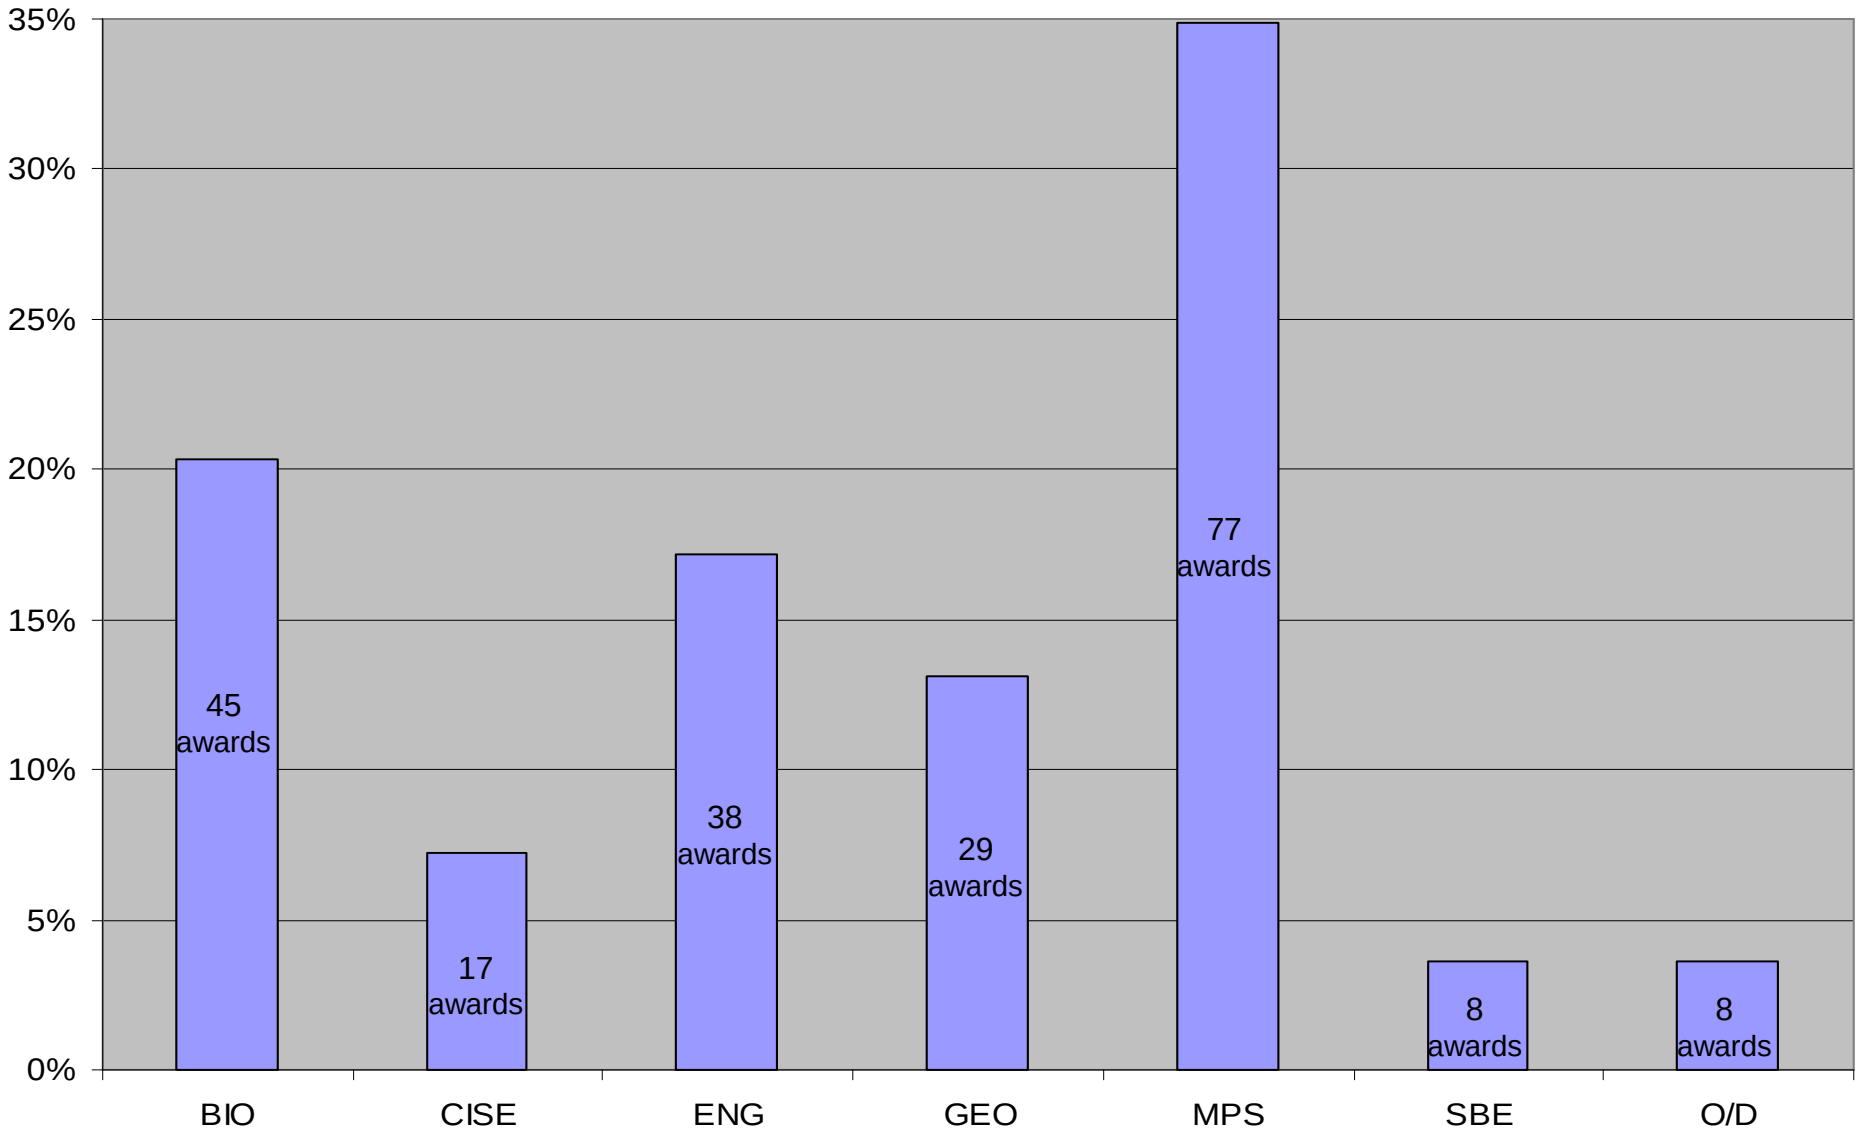

Distribution of MRI Awards

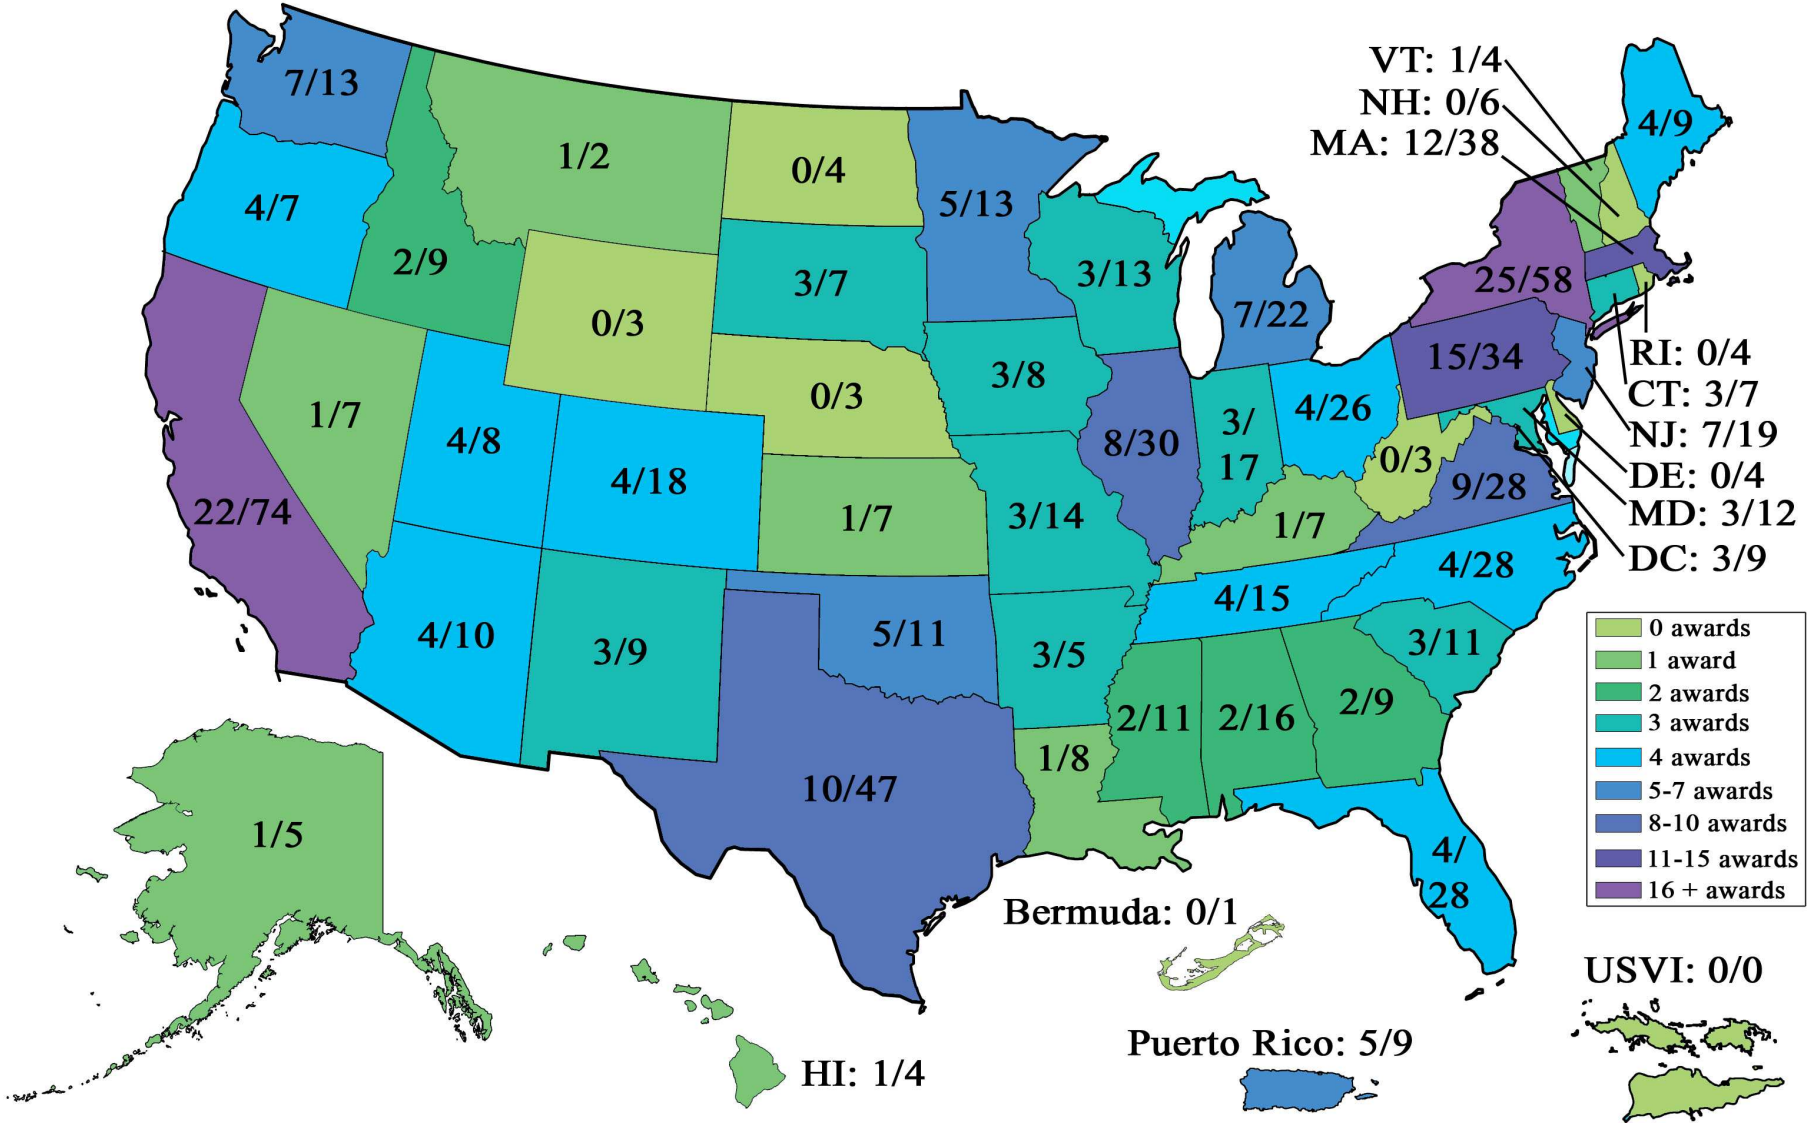

NSF MRI Instrumentation

- Limit of three proposals per institution
  - Need advocacy by chairperson
- Requires 30% institutional cost sharing
  - Need advocacy by chairperson
- Must adhere to scope & use the right lingo
  - Need to inform the community of MRI caveats

MRI Caveats

- Acquisition or development must be for a single instrument (which itself could be made up of many components) with an identifiable location
- MRI does not support general purpose laboratory equipment that does not have a common or specific research focus
- MRI does not support instrumentation for medical research or education, or research with disease-related goals

Workshop Outcomes

- Resulted from reports by 7 workgroups

- 5 workgroups organized around research topics
  - Computer Artifacts
  - Intelligent Systems
  - Distributed Systems
  - Formal and Software Systems
  - HPC and Data-Intensive Computing

- 2 workgroups organized around nature of collaboration
  - Collaboration across multiple CISE disciplines
  - Collaboration involving CISE and non-CISE disciplines

Workshop Outcomes

- The necessity of educating the various constituents on the evolving nature of CISE research and instrumentation

- The importance of balancing inter-disciplinary efforts so that CISE research is leveraged by other disciplines as much as CISE leverages other disciplines

Nature of Research

Issues impacting instrumentation needs:

- Dealing with emerging behavior of large-scale software systems
- Software systems are embedded, and increasingly safety-critical
- Integration of computing systems with the human in the loop – man-machine composition
- Needed instruments cannot be readily acquired; they need to be developed (possibly stitching together many acquired pieces)

Spectrums

- User of instrument
  - A single PI in a “cave”
  - A community of PIs, scientists, students, …
- Role of instrument
  - Enabling new research
  - Sustaining successful research
- Lifecycle of instrument
  - Short-term – prove a concept and create a community
  - Long-term – nurture and transform a community
- Nature of instrument
  - Classical – e.g., simulators and visualizers
  - Emerging – e.g., web-scale auctions, SN games

Examples of Classical Instruments

- A simulator of the interplay between abstractions and computing fabrics at very large scales
- An instrument that enables visualization of emergent behaviors at large scales
- Acquisition of electrical source imaging to help with neuroscience for brain research
- Intelligent spaces, e.g., in museums, that enable new research involving social science topics

Examples of Emerging Instruments

- A software system for testing mechanism design on a web-scale auction
- A programming workbench that allows the composition of various verification theories
- An echo system for certification / quality control of open-source software
- An internet-scale virtual machine – think about building a VM out of cloud resources
- A data collection & associated tools that enable multi-disciplinary experimentation at scale

Instrumentation Impact

- Needed instruments provide higher abstractions that enable advancement in
  - CS Research
  - CS Education
  - K-12 Education
- Large instrumentation projects enhance the visibility within the university – a good strategy to improve a department's standing
  - Builds a community within a department
  - Facilitates acquisition of resources from administration
  - Effective for recruitment of graduate and undergraduate students

Observations

- A limitation of current MRI funding is that software development is not viewed favorably – yet it is critical

- Evolving nature of what constitutes a CISE instrument is hard for other disciplines to accept now – only a matter of time

Observations

- Good “science” is key to success – must argue that science cannot advance without the instrument

- CS community must bring advances in other disciplines to bear on CS research – to make allies and change perceptions

Observations

- Need to train the CS community on how to develop successful MRI proposals
  - Focus on development as opposed to acquisition proposals
- Need to train the CS community on how to evaluate impact and potential impact
  - What may be incremental within a community may be transformative for another or for industry and society – e.g., SLAM

Observations

- On the role of industry
  - Reaching out to industry to underwrite the development of instruments adds legitimacy
  - But academia's role is crucial in providing a neutral "echo system" for instrumentation and to ensure scientific trustworthiness
  - Talking points: The Haskell story at MS, industrial involvement in EU and Brazil

Take-Home Messages

- Importance of educating CS faculty about funding opportunities for instrumentation
- Importance of increasing the MRI proposal and budget pressure from CISE
- MRI proposals need chairs' support to push them through the institution
- Importance of recognizing/rewarding good science – not if you build they will come!