

Computing Community Consortium: Update

Ed Lazowska
Susan Graham
Andy Bernat

July 2008



Highest-level overview



- Hypothesis of the original goal
 - Provide air cover for GENI
- NSF solicitation and CRA proposal
 - Support the computing research community in creating compelling research visions and the mechanisms to realize these visions
 - GENI (\Rightarrow NetSE) is one such vision
 - Serving as the ongoing “voice of the community” for large-scale initiatives such as GENI is one key function
- Slow start due to the need for inclusiveness
- Moving well now
 - We'd give ourselves a B+
 - We feel we're having a significant, and growing, impact

What we say about ourselves

Computing

- Advancing learning
- Advancing other
- Advancing
- Network



YouT

Research foundations




- Time
- Com
- Net
- Pers
- Win
- RIS
- Mod
- RAI
- Para

Much

- Enter
- Data
- Port
- The
- Spee
- Broo

The future is full of opportunity

- Creating the future of networking
- Driving advances in all fields of science and engineering
- Wreckless driving
- Personalized education
- Predictive, preventive, personalized medicine
- Quantum computing
- Empowerment for the developing world
- Personalized health monitoring => quality of life
- Harnessing parallelism: many-core and DISC
- Neurobotics
- Synthetic biology
- The algorithmic lens: Cyber-enabled Discovery and Innovation



We m
artic

- The of t
- The and
- The brig

To th
the C

- To c
cons
- T
- T
- T
- T
- T
- A "c
- C

The s

- CCC
- T
- b
- The
- T
- In
- R
- The
- E
- 50
- The
- A

■ Those involved in shaping CRA's response to NSF's original challenge

- | | | |
|----------------|---------------|----------------|
| ■ Andy Bernat | ■ Dick Karp | ■ Dan Reed |
| ■ Randy Bryant | ■ Ken Kennedy | ■ Wim Sweldens |
| ■ Susan Graham | ■ Ed Lazowska | ■ Jeff Vitter |
| ■ Anita Jones | ■ Peter Lee | |

■ Inaugural CCC Council

- | | | |
|-----------------------|--------------------|---------------------|
| ■ Greg Andrews | ■ Dick Karp | ■ Fred Schneider |
| ■ Bill Feiereisen | ■ John King | ■ Bob Sproull |
| ■ Susan Graham (v ch) | ■ Ed Lazowska (ch) | ■ Karen Sutherland |
| ■ Anita Jones | ■ Peter Lee | ■ David Tennenhouse |
| ■ Dave Kaeli | ■ Andrew McCallum | ■ Dave Waltz |
| | ■ Beth Myratt | |

CCC Council coordination



- Frequent interaction among Ed, Susan, and Andy
- Full-council conference calls every other week
- 3 face-to-face meetings per year
 - December 2006, DC
 - April 2007, Seattle
 - July 2007, Seattle
 - November 2007, DC
 - March 2008, Bay Area
 - July 2008, DC

Agenda, Monday July 7

8:30 – 9:00 / Breakfast

9:00 – 9:15 / Review of CCC financials / Andy Bernat

This is a fiduciary responsibility of the CCC Council

9:15 – 10:15 / Positioning computing research with the next Administration / Ed Lazowska

How can we ensure that computing research is front-stage with the next Administration?

Guest:

Henry Kelly, Federation of American Scientists and Clinton White House

10:15 – 10:30 / Break

10:30 – 11:45 / Transitioning our most successful workshops into funded programs / Susan Graham

Using Cyber Physical Systems as an example, we will discuss in a roundtable with funding agency representatives the ways in which we might transition our most successful workshops into funded programs, and, more broadly, what appropriate strategies might be for initiating and sustaining research.

Guests:

Helen Gill, Program Director, Embedded and Hybrid Systems Program, NSF
CISE CNS

Anita Jones, Lawrence R. Quarles Professor of Engineering and Applied Science,
University of Virginia, and former Director of Defense Research and
Engineering (DDR&E)

Carl Landwehr, Program Manager, Intelligence Advanced Research Projects
Activity (IARPA)

Peter Lyster, Program Director, Center for Bioinformatics and Computational
Biology, NIH National Institute of General Medical Sciences (NIGMS)

Jack Stankovic, BP America Professor of Computer Science, University of
Virginia, and Cyber Physical Systems co-initiator

11:45 – 12:15 / Updates on various topics / Ed Lazowska
Possible Sematech-like structure for CS research – Fred Schneider
Blog – Peter Lee
CISE Symposium – Dave Kaeli
CISE nuggets – Andy Bernat
Theory booklet – Dick Karp
CCC/CRA/NSF “Get Outside Your Comfort Zone” workshop – Ed Lazowska

12:15– 12:30 / Break and pick up working lunch

12:30 – 1:30 / Research Visions / Ed Lazowska

Guests:

Haym Hirsh, Rutgers University, Artificial Intelligence

Ed Felten, Princeton University, computer science and public policy

1:30 – 2:30 / NetSE Council Activities / Ed Lazowska

Presentation and discussion

Guest:

Ellen Zegura

2:30 – 2:45 / Break

2:45 – 3:30 / How are we doing? / Ed Lazowska, Susan Graham, Andy Bernat

Discussion, specifically including post-workshop transitioning activities, NetSE Council progress, value-add of CCC vs. uncoordinated smaller awards

3:30 / Meeting close

Ed Lazowska, Susan Graham, Andy Bernat travel to NSF

Activities to date



- Definition and execution of a bootstrapping procedure for the CCC
 - Not straightforward, because community ownership was essential
- Five plenary talks at the Federated Computing Research Conference (June 2007) to introduce CCC to the computing research community
 - Embracing and amplifying efforts that are already underway

- Countless additional talks (Snowbird forthcoming)

The Computing Community Consortium: Stimulating Bigger Thinking

Ed Lazowska

Bill & Melinda Gates Chair in
Computer Science & Engineering
University of Washington

Chair, Computing Community Consortium

Rice University

April 2008

<http://www.cra.org/ccc/>



Articles in CRN, CACM (forthcoming), ...

COMPUTING RESEARCH NEWS

A Publication of the Computing Research Association

January 2008

Vol. 20/No. 1

The Computing Community Consortium: Who, What, When, Where, Why, and How

Computing Research News interviews Ed Lazowska, Bill & Melinda Gates Chair of Computer Science & Engineering at the University of Washington, and Chair of the Computing Community Consortium. Short articles on the CCC in its formative stages appeared in the November 2006 and May 2007 issues of CRN.

Computing Research News: Begin by describing the Computing Community Consortium. What is its role?

Ed Lazowska: The National Science Foundation created the Computing Community Consortium with the goal of stimulating the computing research community to imagine, articulate, and pursue more audacious research visions—visions that will capture the imagination and change the world. The CCC is funded through an NSF award to the



Ed Lazowska

Computing Research Association; the CCC's Council operates as a committee of CRA.

CRN: Who is on the CCC Council? How were they chosen?

EL: The CCC Council comprises a Chair and 15 members on staggered 3-year terms. Members of the Council are listed on the CCC website: <http://www.cra.org/ccc/>. The Council was chosen through an open process led by Randy Bryant, Dean of the School of Computer Science at Carnegie Mellon University. I chair the Council (this selection, too, involved an open process led by Randy), and Susan Graham from UC Berkeley serves as Vice Chair.

CRN: What's the role of the Council?

EL: The Council serves as a facilitator. It is the computing research community as a whole that must imagine, articulate, and pursue more audacious research visions. The Council helps the process in various ways.

CRN: How?

EL: We're in the early stages—the Council was appointed six months

ago, after an open process that engaged the entire computing research community. Here are some examples:

- The CCC sponsored a set of five plenary talks at the Federated Computing Research Conference in June—talks by Christos Papadimitriou, Bob Colwell, Randy Bryant, Scott Shenker and me that described specific research visions for the field. See <http://www.cra.org/ccc/frc/>.
- The CCC is providing support for "visioning workshops" organized by members of the computing research community. Sponsorship can be obtained through a lightweight proposal process; the first awards have already been made, and we were excited by the level of participation. See <http://www.cra.org/ccc/rfp/>.
- We will coordinate closely with funding agencies so we can help to transition visionary ideas into funded programs.


"Our field has accomplished so much, and there is so much more to do. The opportunities are extraordinary."

- We are preparing an inspirational website and booklet describing a wide range of research visions for the field. We are also initiating a blog where the entire research community can participate in real time.

CRN: How does the CCC relate to CRA and to NSF?

EL: CRA and CCC both are concerned with the health of the computing research community—CRA in a broad-based way, and CCC with a narrower focus on research visions. One way to view the relationship is that NSF, by funding the CCC through CRA, has provided CRA and the computing research community with the means to dramatically expand our efforts in this particular area. The

The CCC
Continued on Page 5



- Definition and execution of an RFP process to support visioning by the computing research community

- Quarterly deadlines, but a rolling process

- Five efforts launched thus far:

- "Big Data Computing Study Group"

- "Cyber-Physical Systems"

- "Visions for Theoretical Computer Science"

- "From Internet to Robotics: The Next Transformative Technology"

- "Network Science and Engineering"



■ *Big Data Computing Study Group*

■ Topic:

- "The Big Data Computing Study Group will undertake efforts to explore and enable opportunities on the research and application of high-performance computing over very large data sets."

■ Leadership:

- Randy Bryant, CMU
- Thomas Kwan, Yahoo! Research

■ Initial activities:

- Hadoop Summit, March 25, Sunnyvale CA
- Data-Intensive Scalable Computing Symposium, March 26, Sunnyvale CA



SEARCH

NSF Web Site



Computer & Information Sciences & Engineering



[CluE Program Overview](#)

[Usage Agreements](#)

[Cluster Allocation \(forthcoming\)](#)

[Available Data Sets \(forthcoming\)](#)

[CluE Program Page](#)

[CluE Program Solicitation Page](#)

CISE - Cluster Exploratory (CluE)

The Cluster Exploratory (CluE) program is providing NSF-funded researchers access to software and services running on a Google-IBM cluster to explore innovative research ideas in data-intensive computing. Funded proposals will cover a broad range of activities exploring the potential of this technology to contribute to science and engineering research as well as to applications that promise benefit to society as a whole.

This page provides access to information that may be of use to proposers as well as to PI's of awarded projects.

[Usage Agreements](#)

Awarded projects must execute a usage agreement before any access to the cluster will be provided.

Cluster Allocation

Details for acquiring a cluster allocation will be described here.

Available Data Sets

A catalog describing data sets on the cluster available to researchers will be provided here.



University of Washington

Computer Science & Engineering

Welcome to the 2008 NSF Data-Intensive Scalable Computing in Education Workshop

▷ CSE Home

▷ About Us ▷ Search ▷ Contact Info

Quick navigation:

- [Motivation](#)
- [Location](#)
 - [Campus Map](#)
- [Tentative Schedule](#)
- [Application](#)
- [Online Resources](#)
- [Contact Information](#)

Motivation

Data-intensive scalable computing (DISC) is becoming an increasingly relevant area of computer science education. Given the rapid rate of change in this field, existing curricular efforts need to be revisited to address the unique challenges for designing computer clusters, software platforms for large-scale data computing, and applications that effectively use them.

The goal of this workshop is to inspire the development of new coursework in large-scale data-intensive application design and cluster computing. Educators will be introduced to existing curriculum components for similar coursework, as well as provide in-depth hands-on experience using software platforms that make this manageable in an undergraduate setting. Time will be allocated for discussions between attendees and representatives from industry and the open-source community to help formulate new ideas to carry back to the academic institutions of the attendees.

Event sponsors:



CCA-08:

Cloud Computing and Its Applications

HOME

More details will be provided shortly.



E.D.S.' service management center in Plano, Tex. (Photo: Electronic Data Systems)

Organizing Committee

Charlie Catlett, Argonne National Laboratory

Ian Foster, Argonne and University of Chicago

Joe Hellerstein, University of California Berkeley

October 22 and 23, 2008

Gleacher Center

Chicago, IL

Dramatic growth in data and equally rapid decline in the cost of highly integrated clusters has spurred the emergence of the data center as the platform of choice for a growing class of data-intensive applications. To encourage conversations between those developing applications, algorithms, software, and hardware for such "cloud" platforms, we are convening the first workshop on Cloud Computing and its Applications (CCA'08).

This workshop will include a mixture of invited and contributed talks on cloud computing, data intensive scalable computing, and related topics.

Topics of interest include:

- compute and storage cloud architectures and implementations
- map-reduce and its generalizations
- programming models and tools
- novel data-intensive computing applications
- data intensive scalable computing
- distributed data intensive computing



| *Cyber-Physical Systems*

| Topic:

- “The integration of physical systems and processes with networked computing has led to the emergence of a new generation of engineered systems: Cyber-Physical Systems (CPS). Such systems use computations and communication deeply embedded in and interacting with physical processes to add new capabilities to physical systems. CPS range from miniscule (pace makers) to large-scale (the national power-grid). This effort will identify the science and technology challenges facing CPS.”

| Leadership:

- Bruce Krogh, CMU
- Jack Stankovic, University of Virginia
- 12 others

| Initial activities:

- Multiple preliminary workshops
- Cyber-Physical Systems Summit, April 24-25, St. Louis MO



■ *Visions for Theoretical Computer Science*

■ Topic:

- “The purpose of the visioning workshop will be to identify and distill broad research themes within TCS that have potential for major impact in the future ... The workshop will aim to produce compelling “nuggets” that can quickly convey the importance of a research direction to a layperson [and] could be used by the CCC or anyone else making the case for a sustained investment in long-term, foundational computing research.”

■ Leadership:

- Richard Ladner, Washington
- Bernard Chazelle, Anna Karlin, Dick Lipton, Salil Vadhan

■ Initial activities:

- Workshop prior to STOC, May 17, Seattle WA



■ *From Internet to Robotics: The Next Transformative Technology*

■ Topic:

- "This study will generate a roadmap of applications for robotics across users, producers and researchers. The objective is to provide a comprehensive view of use of robotics, the main obstacles to deployment, and the key competencies required to facilitate the transformation."

■ Leadership:

- Henrik Christensen, Georgia Tech, and 10 others

■ Initial activities:

- Workshop on manufacturing robotics, June 17, Arlington
- Workshop on medical/healthcare robotics, June 18-19, Arlington
- Workshop on emerging technologies and trends in robotics, August 14-15, Snowbird
- Workshop on domestic and professional service robotics, August 7-8, San Francisco



■ In the pipeline:

■ Approved after review and resubmission:

- | One Teacher per Student: Global Resources for Online Education (GROE)
 - Beverly Park Woolf, University of Massachusetts, and 14 others

■ Reviewed and recently resubmitted; awaiting re-review:

- | Envisioning National and International Research on the Multidisciplinary Empirical Science of Free/Open Source Software (FOSS)
 - Walt Scacchi, UC Irvine, and others
- | ...



- | Reviewed and awaiting resubmission:

- | ...

- | Recently submitted and awaiting review:

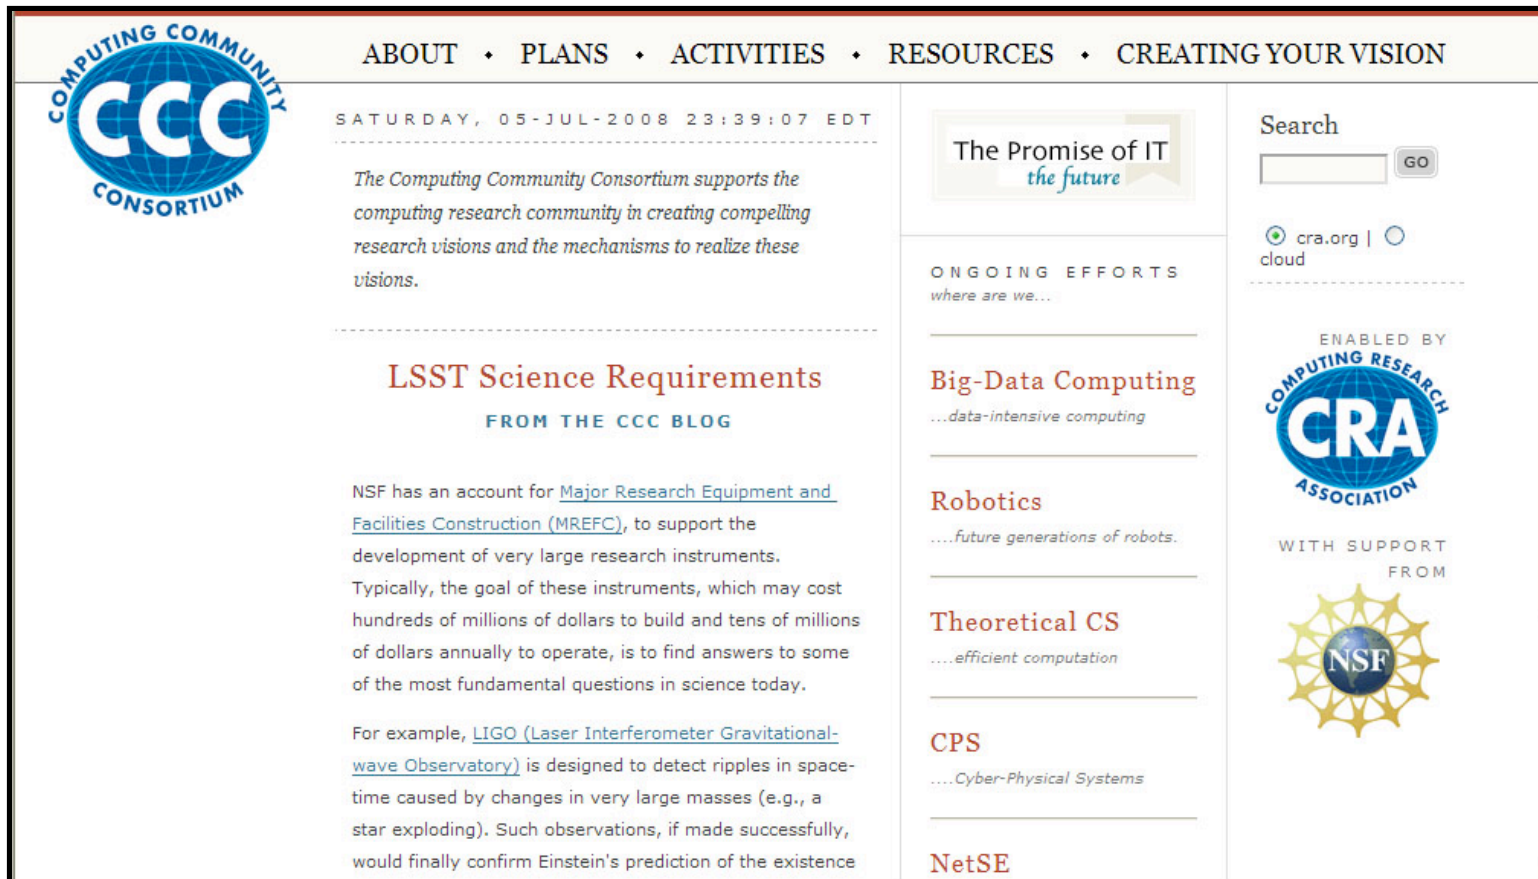
- | Predictable Systems from Unpredictable Components

- Nicholas P. Carter, Intel; Andre DeHon, Penn; Heather M. Quinn, LANL

- | Information and Communication Technologies for Development (ICTD): A New Grand Challenge for Computer Science Research

- Tapan S. Parikh, UC Berkeley, and 9 others

Creation of a website



The screenshot shows the homepage of the Computing Community Consortium (CCC) website as of July 5, 2008. The page features a navigation menu, a search bar, and several content sections. The main content area highlights a blog post about LSST Science Requirements, supported by NSF. A sidebar on the right lists ongoing efforts in Big-Data Computing, Robotics, Theoretical CS, CPS, and NetSE. Logos for the Computing Research Association (CRA) and NSF are also visible.

COMPUTING COMMUNITY CONSORTIUM

ABOUT • PLANS • ACTIVITIES • RESOURCES • CREATING YOUR VISION

SATURDAY, 05-JUL-2008 23:39:07 EDT

The Computing Community Consortium supports the computing research community in creating compelling research visions and the mechanisms to realize these visions.

LSST Science Requirements

FROM THE CCC BLOG

NSF has an account for [Major Research Equipment and Facilities Construction \(MREFC\)](#), to support the development of very large research instruments. Typically, the goal of these instruments, which may cost hundreds of millions of dollars to build and tens of millions of dollars annually to operate, is to find answers to some of the most fundamental questions in science today.

For example, [LIGO \(Laser Interferometer Gravitational-wave Observatory\)](#) is designed to detect ripples in space-time caused by changes in very large masses (e.g., a star exploding). Such observations, if made successfully, would finally confirm Einstein's prediction of the existence

The Promise of IT
the future

ONGOING EFFORTS
where are we...

Big-Data Computing
...data-intensive computing

Robotics
...future generations of robots.

Theoretical CS
...efficient computation

CPS
...Cyber-Physical Systems

NetSE

Search

cra.org | cloud

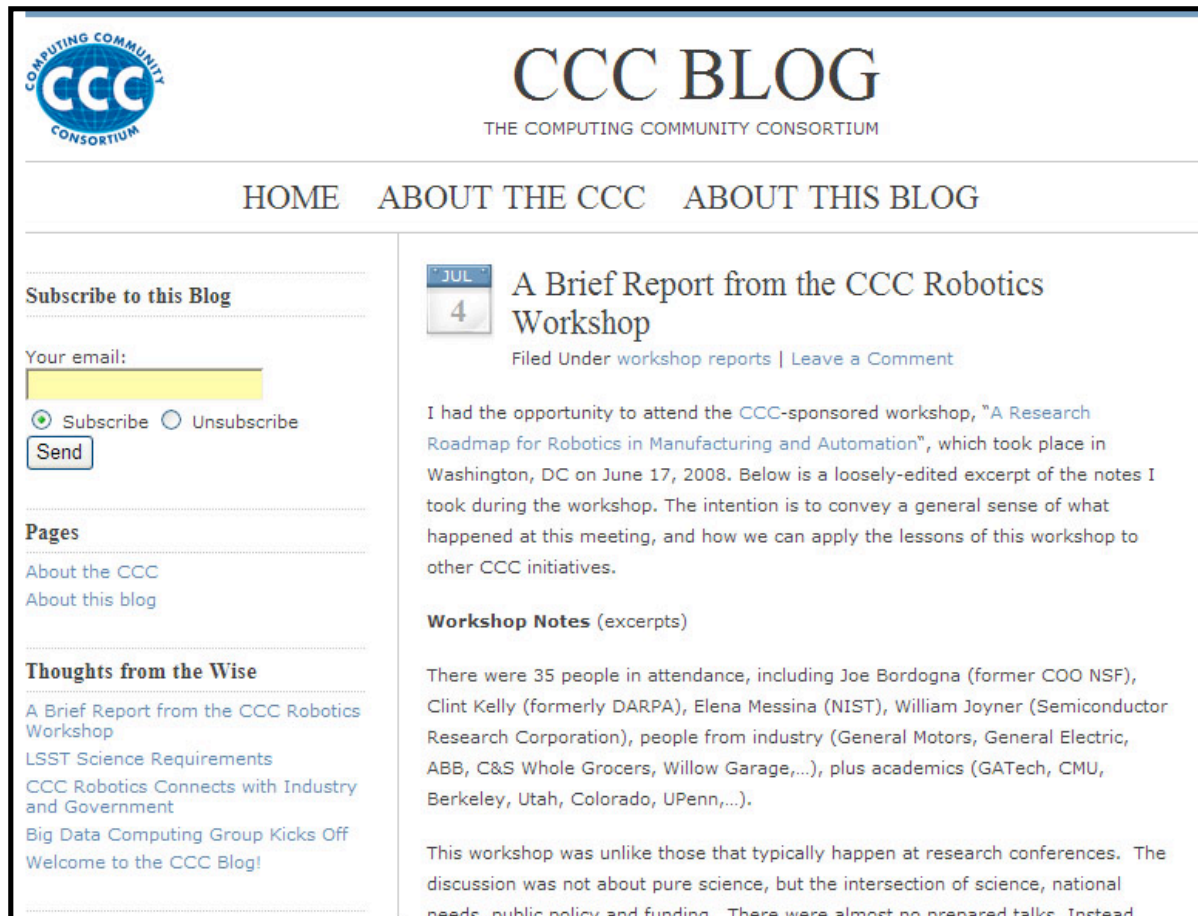
ENABLED BY

COMPUTING RESEARCH ASSOCIATION

WITH SUPPORT FROM

NSF

Creation of a research visions blog



The screenshot shows the homepage of the CCC Blog. At the top left is the CCC logo, which consists of a blue circle with 'CCC' in white and 'COMPUTING COMMUNITY CONSORTIUM' around it. To the right of the logo, the text 'CCC BLOG' is displayed in a large, bold, serif font, with 'THE COMPUTING COMMUNITY CONSORTIUM' in a smaller, sans-serif font below it. A horizontal navigation bar contains the links 'HOME', 'ABOUT THE CCC', and 'ABOUT THIS BLOG'. Below the navigation bar, the page is divided into three columns. The left column contains a 'Subscribe to this Blog' form with a text input field, 'Subscribe' and 'Unsubscribe' radio buttons, and a 'Send' button. Below the form are sections for 'Pages' (with links to 'About the CCC' and 'About this blog') and 'Thoughts from the Wise' (with links to 'A Brief Report from the CCC Robotics Workshop', 'LSST Science Requirements', 'CCC Robotics Connects with Industry and Government', and 'Big Data Computing Group Kicks Off Welcome to the CCC Blog!'). The middle column features a date widget for 'JUL 4' and a post title 'A Brief Report from the CCC Robotics Workshop' with a 'Leave a Comment' link. The right column contains the main text of the post, starting with 'I had the opportunity to attend the CCC-sponsored workshop, "A Research Roadmap for Robotics in Manufacturing and Automation", which took place in Washington, DC on June 17, 2008. Below is a loosely-edited excerpt of the notes I took during the workshop. The intention is to convey a general sense of what happened at this meeting, and how we can apply the lessons of this workshop to other CCC initiatives.' This is followed by a section titled 'Workshop Notes (excerpts)' containing a paragraph about the workshop attendees and a final paragraph starting with 'This workshop was unlike those that typically happen at research conferences. The discussion was not about pure science, but the intersection of science, national needs, public policy and funding. There were almost no prepared talks. Instead



■ Other activities

- CISE Symposium
 - | Planning is actively ongoing
- CISE Nuggets
 - | A mining exercise
- Prior Grand Challenge Efforts
 - | Another mining exercise
- "Get Outside Your Comfort Zone" workshop
 - | CCC, CRA, NSF - planning is actively ongoing
- Undergraduate institutions
 - | Karen Sutherland is driving
- Enhanced communications
 - | CRA has selected Xenophon Strategies to work on communicating the research message and Gimga Group for design issues

NetSE: As a visioning exercise

■ *Network Science and Engineering (NetSE)*

■ Topic:

- Our evolving networks are extraordinarily complex. Is there a **science** for understanding the complexity of our networks such that we can **engineer** them to have predictable behavior? We must develop a compelling and broad-based research agenda for the science and engineering of our evolving, complex networks.

■ Leadership:

- Ellen Zegura, Georgia Tech, chair of NetSE Council
 - 19 members
- Chip Elliott, BBN, director of GENI Project Office

■ Initial activities:

- Workshops going back several years, and continuing
- GENI Engineering Conferences, ongoing
- Research workshops and meetings, Summer/Fall 2008
- Delivery of V1.0 NetSE research plan, December 2008

NetSE: Current activities



■ Re-populating NetSE Council

- Email to current members sent June 20
 - | Politely requested signup for specific duties, or resignation
- Additional members targeted from individuals active in GPO activities
 - | The various NetSE research theme areas are already adequately represented
- Ellen will close this by July 21



■ Research agenda workshops

■ *Network Science and Network Design*

- | John Wroclawski, USC/ISI, and John Doyle, Caltech
- | July 29 - 30, Southern CA

■ *Behavior, Computation, and Networks in Human Subject Experimentation (BCN)*

- | Michael Kearns, Penn, and Colin Camerer, Caltech
- | July 31 - August 1, La Jolla CA

■ *Network Design and Societal Values*

- | David Clark, MIT, and Helen Nissenbaum, NYU
- | September 24-26, Washington, DC

- These are on-track, have good co-chairs, have deliverables specified, and have individuals assigned to extract facilities requirements from the discussions



■ Meetings

- Integrating the Theory of Networked Computation
 - | Ellen Zegura, Joan Feigenbaum (Yale), John Byers (BU), June 11
- Network Design
 - | Engaging members of the Planning Group and others active in GPO activities to extract network design research topics from the existing research plan and other sources
 - | Ellen Zegura is coordinating
 - | Likely held just before SIGCOMM (August)



■ Committees in formation

- NetSE Facilities Requirements Committee
 - | Construct a coherent set of facilities requirements from the outputs of the workshops and meetings
- NetSE Communications Committee
 - | Liaison with the computing research community
- NetSE Agenda Writing Committee
 - | Responsible for delivery of the V1.0 science plan by December 2008

NetSE: Major issues




- Continue current momentum
- Utilize GPO
 - BBN has allocated outstanding talent
 - Likely to be repurposed if not utilized
 - Some frustration with grant award process
 - Can GPO be used to engage industry interest/effort?

CCC value-add



- Overall: CCC is a catalyst serving the computing research community and NSF
 - Strong, diverse group of community members
 - | We want *somebody* to be thinking about and doing these thingsb. Who better?
 - Speed of action
 - | Less than a week in two cases
 - Shepherding
 - | Many of the visioning proposals
 - Stimulus for community
 - | Karp says theory workshop simply would not have happened

- 
- Re-focus existing subfields as well as explore new ones
 - Robotics effort
 - A vehicle/agent
 - New NSF/CRA/CCC "out of the box" workshop
 - Frank discussions - meetings not public

Idea for reverse site visit



- We conduct a self-assessment
 - Write the questions we think are appropriate for evaluation
 - Write the answers
- Reverse site visit to explore these questions and answers (and bring up additional questions)
- Results (or at least our self-assessment) are public