Introducing the Computing Community Consortium

Susan Graham

Pehong Chen Distinguished Professor Emerita and Professor in the Graduate School, University of California, Berkeley

Vice-Chair, Computing Community Consortium

Jennifer Rexford

Professor of Computer Science, Princeton University

Member, GENI Science Council

Computing has changed the world

- Advances in computing change the way we live, work, learn, and communicate
- Advances in computing drive advances in nearly all other fields
- Advances in computing power our economy
  - Not just through the growth of the IT industry – through productivity growth across the entire economy



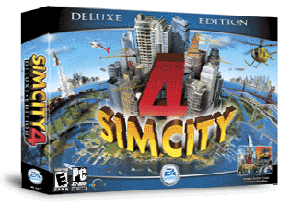

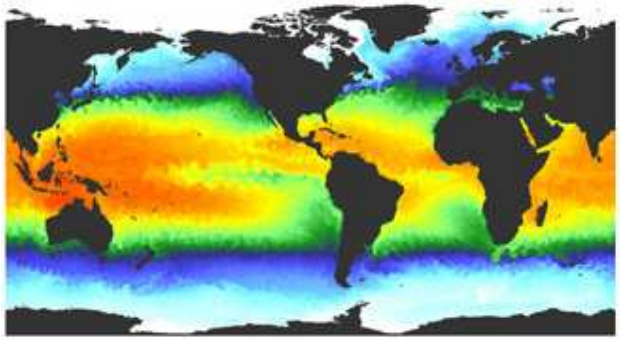
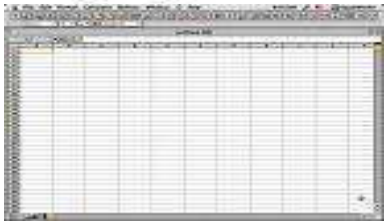
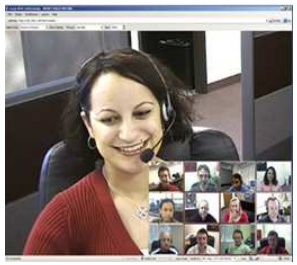


Research has built the foundation

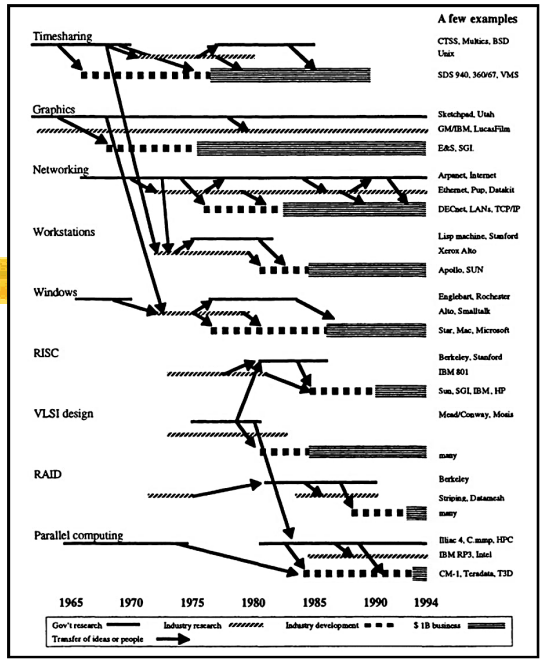

- Timesharing
- Computer graphics
- Networking (LANs and the Internet)
- Personal workstation computing
- Windows and the graphical user interface
- RISC architectures
- Modern integrated circuit design
- RAID storage
- Parallel computing

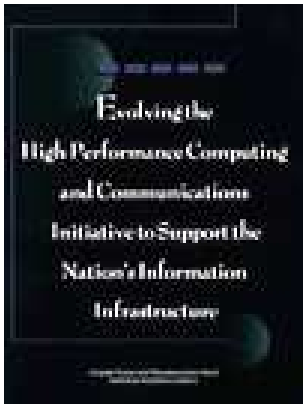

Much of the impact is recent

- Entertainment technology
- Data mining
- Portable communication
- The World Wide Web
- Speech recognition
- Broadband last mile

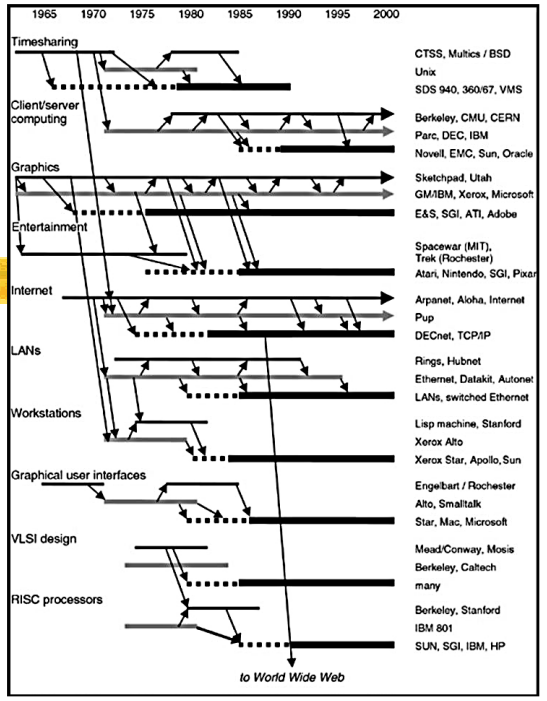

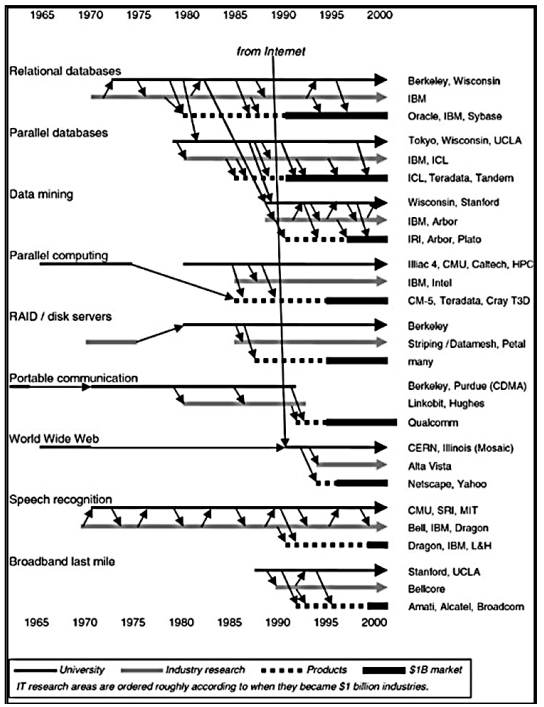

The future is full of opportunity

- Designing a new Internet – FIND + GENI
- Driving advances in all fields of science and engineering
- Wreckless driving
- Personalized education
- Predictive, preventive, personalized medicine
- Quantum computing
- Transforming the developing world
- Personalized health monitoring => quality of life
- Data-intensive supercomputing
- Neurobotics
- Synthetic biology
- The algorithmic lens: Cyber-enabled Discovery and Innovation

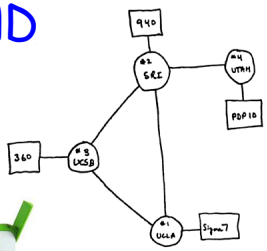

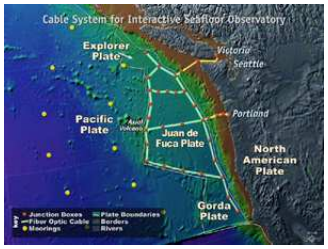

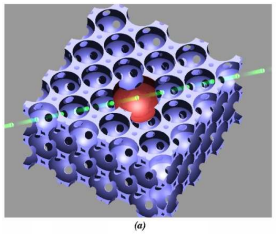

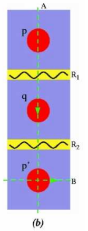

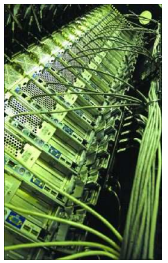

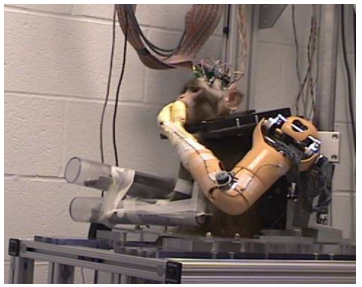

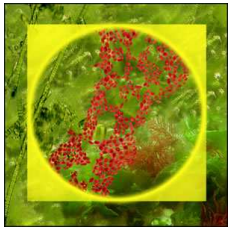

We must work together to establish, articulate, and pursue visions for the field

- The challenges that will shape the intellectual future of the field
- The challenges that will catalyze research investment and public support
- The challenges that will attract the best and brightest minds of a new generation

To this end, NSF asked CRA to create the Computing Community Consortium

- To catalyze the computing research community to consider such questions
  - To envision long-range, more audacious research challenges
  - To build momentum around such visions
  - To state them in compelling ways
  - To move them towards funded initiatives
  - To ensure "science oversight" of "at scale" initiatives

The structure

- CCC is all of us!
  - This process *must* succeed, and it *can't* succeed without broad community engagement
- There is a CCC Council to guide the effort
  - The Council *stimulates* and *facilitates* – it doesn't "own"
  - The initial Council was created through an open process led by Randy Bryant
- The Council is led by a Chair
  - Ed Lazowska, University of Washington
  - 50% effort – not titular
- The CCC is staffed by CRA
  - Andy Bernat serves as Executive Director

The CCC process

Nucleation

The germ of a vision. in the minds of a small number of people

CCC can encourage through exemplars

Crystalization and broadening

Broadening of involvement and crystallization of the vision

CCC can support study groups

Program formation

Work with agency staff to formulate a program

CCC can provide guidance and create relationships with federal agency staffs

Program realization

Agency places the program into its budget request

CCC can work with initiators to ensure inclusion in budget

Execution

Do it

Two motivating examples

- CDI - Cyber-enabled Discovery and Innovation
  - Started as a white paper from theory community - the algorithmic lens on science
  - Now a recently announced funded program at NSF
- GENI – Global Environment for Network Innovations
  - Proposed MREFC instrument for computing research
  - CCC Council formed the GENI Science Council to create and guide the research and education plan and to interact with the GENI Project Office funded by NSF
  - Jen Rexford will describe GENI

GENI: Global Environment for Network Innovations

- What is GENI?
  - Shared, wide-area experimental facility
  - …to evaluate clean-slate network architectures
  - …that are visions for the future Internet

- Natural questions you might have…
  - Why worry about Internet's future?
  - Why should we think from a "clean slate"?
  - Why do we need an experimental facility?
  - What kind of experimental facility?

Why Worry About the Future Internet?

- The Internet is great at what it does.
  - Everyone should be proud of this.
  - All sorts of things can be built on top of it.
- But…
  - Security is weak and not getting better.
  - Availability continues to be a challenge.
  - It is hard to manage and getting harder.
  - It does not handle mobility well.
  - A long list, once you start…

Why Think From a "Clean Slate"?

- Clean Slate is a *means*, not an end
  - No one expects direct adoption of radical designs
- New *insights* can impact the Internet's evolution
  - Clean-slate designs ➔ insights ➔ Better Internet
- Intellectual foundation for network architecture
  - Understanding trade-offs between many design goals
- NSF's FIND (Future Internet Design) program
  - See www.nets-find.net

Why Do We Need an Experimental Facility?

- Need to build and try out ideas
  - Paper designs are just idle speculation
  - Simulation is only occasionally a substitute
- We need:
  - Real implementation
  - Real experience
  - Real network conditions
  - Real users
  - To live in the future
- But this is hard to do today….

What Kind of Experimental Facility?

- **Shared**: many experiments in parallel
  - Amortize the cost of the facility
  - Long-running deployment studies
- Programmable: new network designs
  - Experiments with radical designs
  - Revisit the divisions between the layers
- Real: attracting real user traffic
  - Users stress a system, and vote with their feet
  - Experiments grappling with scale and failures

Slices

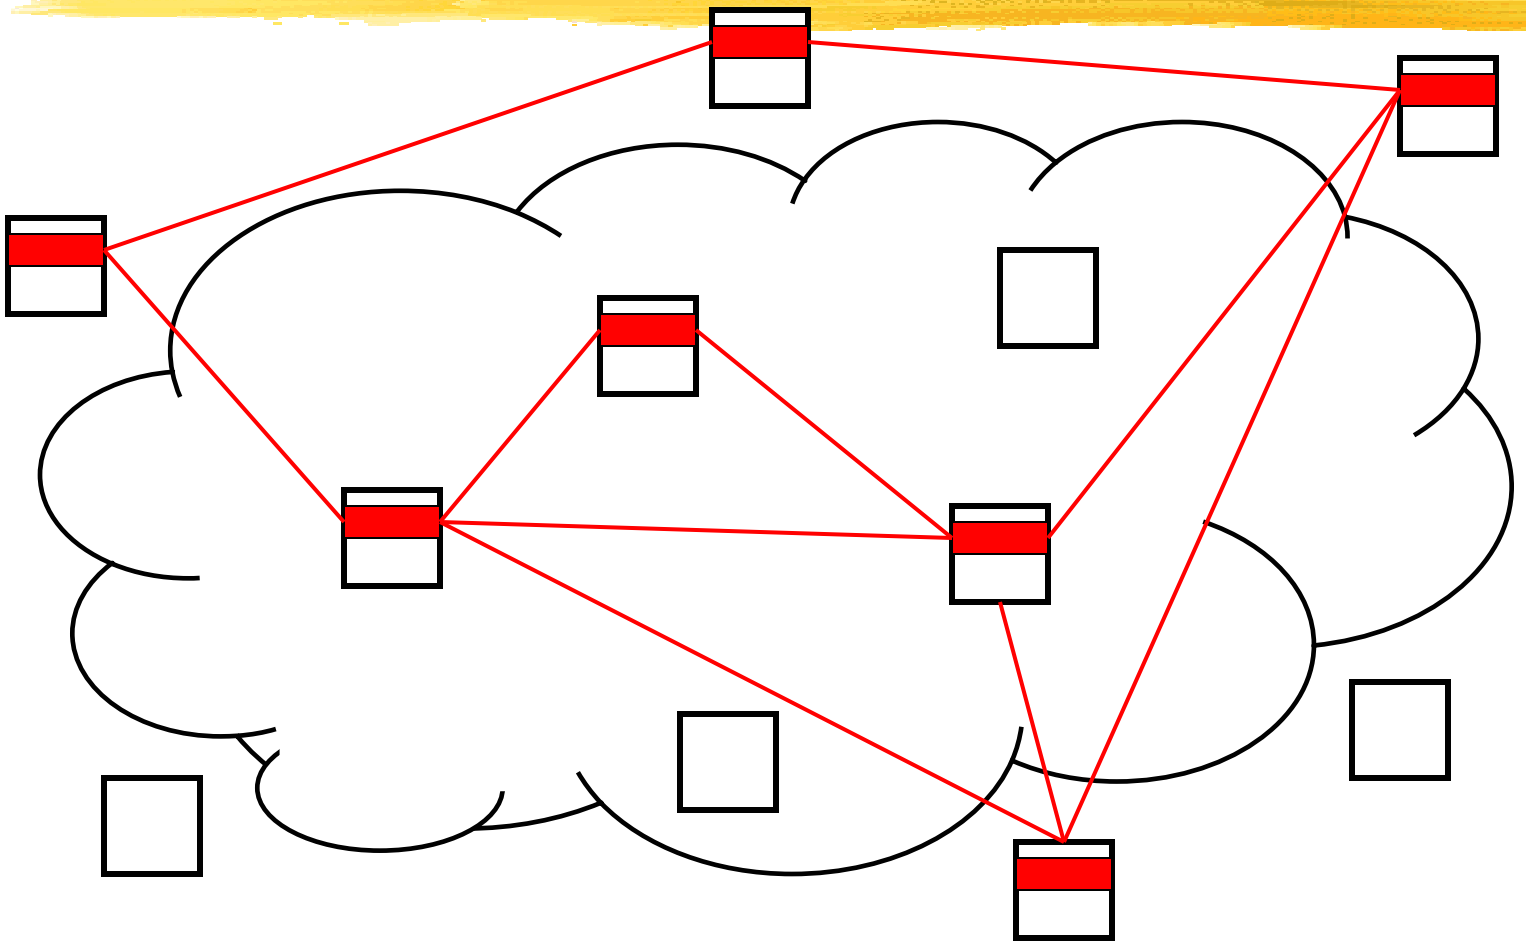

Slices

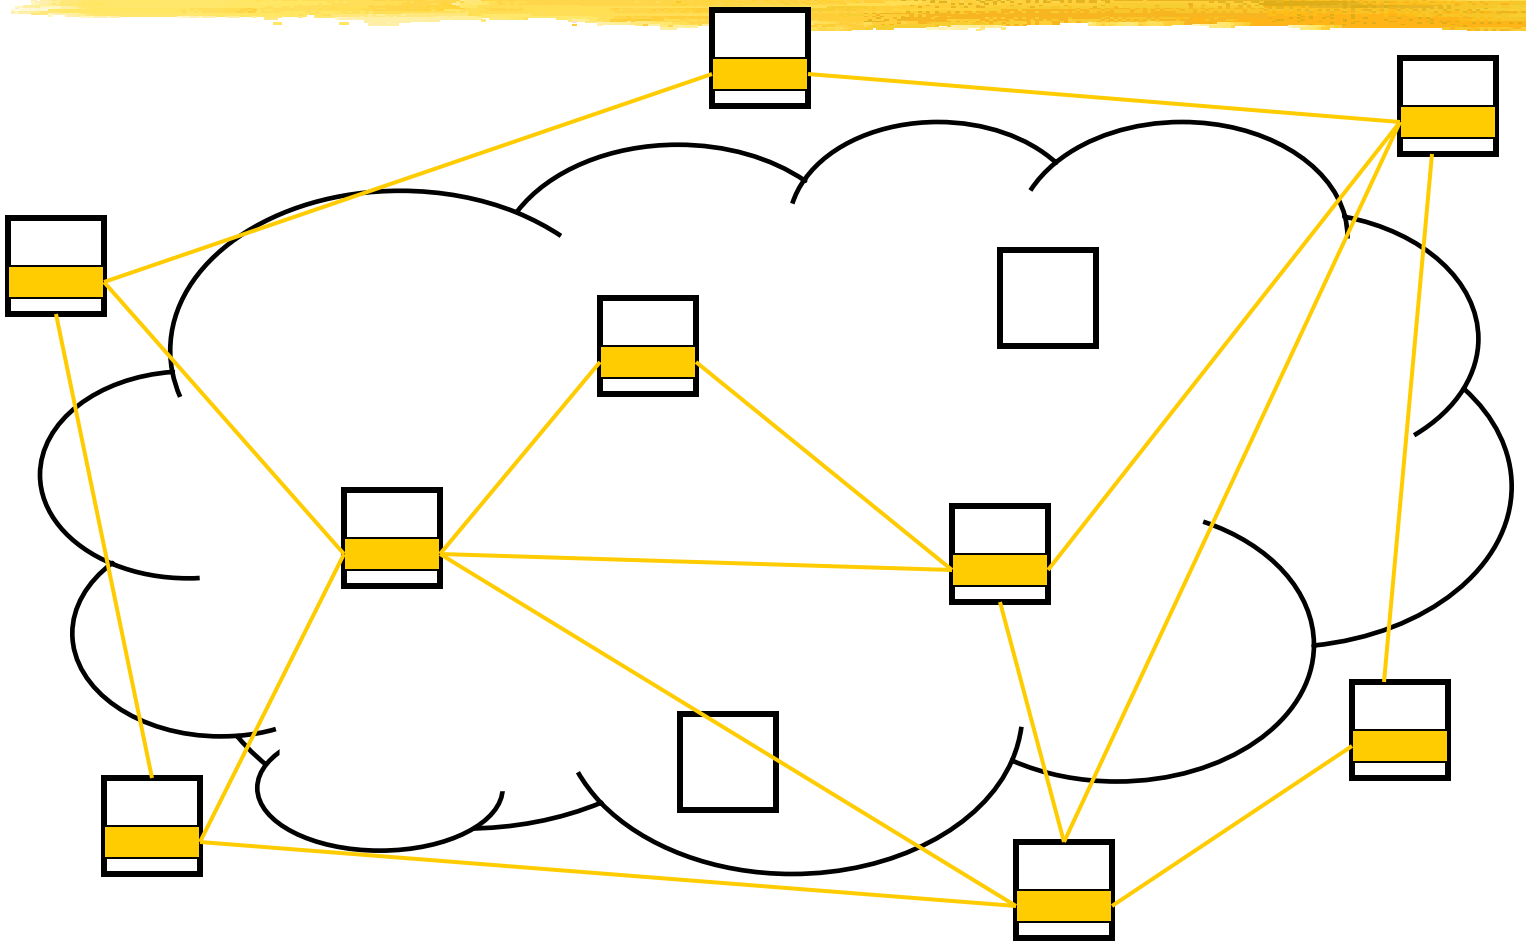

User Opt-in

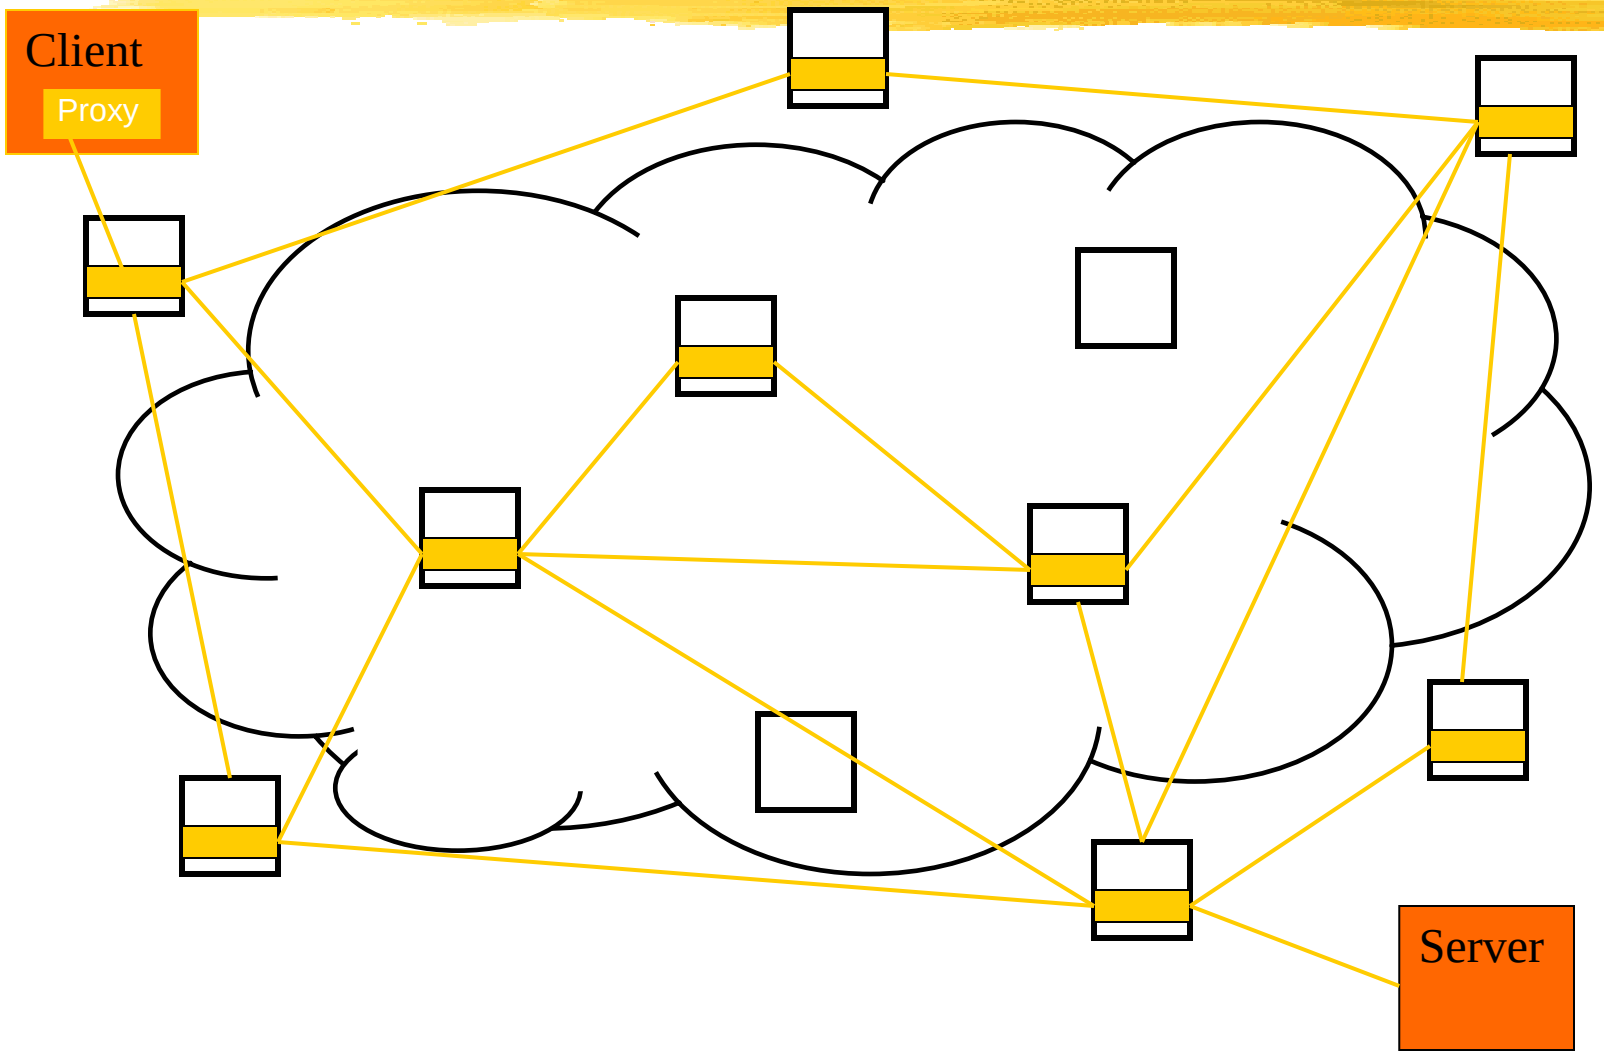

Realizing the Ideas

- Slices embedded in a substrate of resources
  - Physical network substrate
    - Expandable collection of building block components
    - Optical switches, routers, servers, wireless, sensors…
  - Software management framework
    - Knits building blocks together into a coherent facility
    - Embeds slices in the physical substrate
- Builds on ideas in past systems
  - PlanetLab, Emulab, ORBIT, X-Bone, …

GENI: Current Status

- An initial design and science plan
  - Created by a large group of volunteers
  - See http://www.geni.net/documents.html
- More formal structure guiding the next phase
  - GENI Project Office (BBN, Chip Elliott)
    - Working groups to complete the GENI design
    - Support prototyping efforts on parts of the facility
  - GENI Science Council (Ellen Zegura, Scott Shenker)
    - Research requirements for the facility
    - Internationalization, education, …
- Still very much a work in progress…

Success Scenarios for GENI

- Expand the research pipeline
  - Sound foundation for future network architectures
  - Experimental evaluation, rather than paper designs
- Create new services
  - Demonstrate new services at scale
  - Attract real users
- Aid the evolution of the Internet
  - Demonstrate ideas that ultimately see real deployment
  - Provide architectural clarity for evolutionary path
- Lead to a future global network
  - Purist: synthesis of a single new architecture
  - Pluralist: virtualization supporting many architectures

Why Should You Care About GENI?

- The Internet belongs to all of us
  - We should play a lead role in its evolution
- Many areas of computing are crucial here
  - Networking, distributed systems, algorithms, …
- We can make our ideas come to fruition
  - Great ideas + experiments == insights and change
- Our community should "think big"
  - We can be more than the sum of our parts
  - If we can come together to address big intellectual and practical challenges

The desired outcome for CCC

- Broad community engagement in establishing more audacious and inspiring research visions for our field
  - Some will require significant research infrastructure (e.g., GENI); some will be new programs (e.g., CDI)
- Better public appreciation of the potential of the field
- Attraction of a new generation of students
- Greater impact!

Discussion

- Questions?
- Comments?
- Suggestions?

CCC is all of us!

- http://www.cra.org/ccc/

Extra slides

Initial CCC Council

- Greg Andrews
- Bill Feiereisen
- Susan Graham
- Anita Jones
- David Kaeli
- Dick Karp
- John King
- Ed Lazowska
- Peter Lee
- Andrew McCallum
- Beth Mynatt
- Fred Schneider
- Bob Sproull
- Karen Sutherland
- David Tennenhouse
- Dave Waltz