



Major Research Themes in CS

A personal and partial view

Marc Snir

What is CS?

- Computer Science, Computer Engineering, Software Engineering, Information Science... what's all that?
- CS is a misnomer
 - science finds the truth about nature
 - engineering builds better (cheaper, more useful) gadgets
 - CS is an engineering discipline!
- The boundaries between CS, CE, SE, IS... are arbitrary
 - mostly a reflection of historical evolution and education foci
 - still evolving

What is Engineering Research?

Alternative View



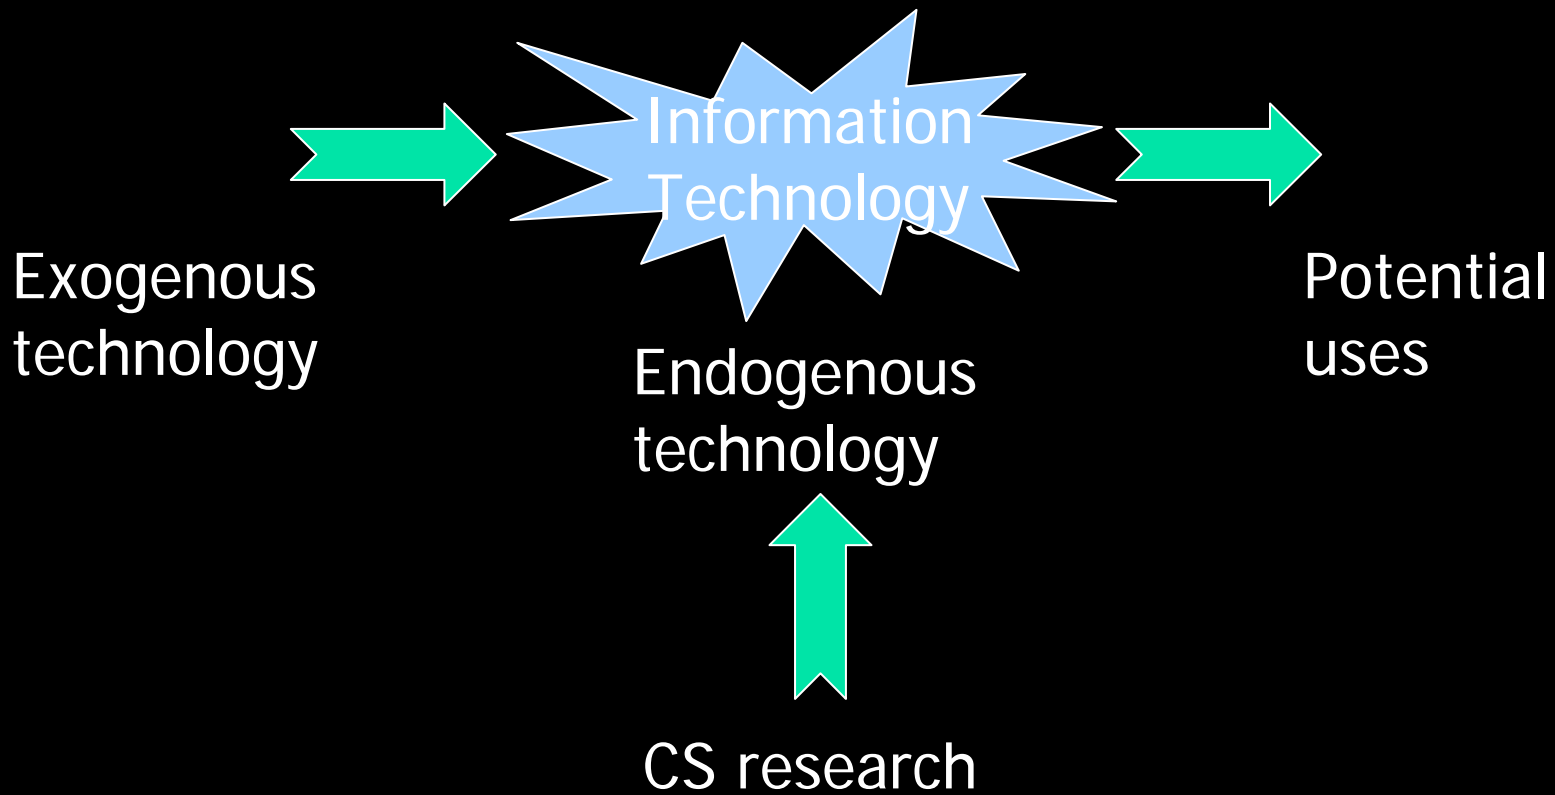
How is CS/CE different from other engineering disciplines?

- More malleable: less constrained by underlying physical reality
- More universal: information systems are (or can be) everywhere
 - researchers have more freedom to use/abuse
 - technology changes at faster rate
 - faster change in research directions, but greater superficiality (ideas are forgotten, experiments are not validated)
 - “social” constraints (standards, conventions) play a greater role
 - larger “surface-to-volume” ratio: more interaction with other disciplines

How is academic research different from industrial research?

- education vs. product motivation
- long term vs. short term
- public good vs. private good
 - public goods are of greater importance in computer industry than in most other industries
 - defensive IP (MAD) vs. offensive IP

Pushpull



Push: the tyranny of the exponentials

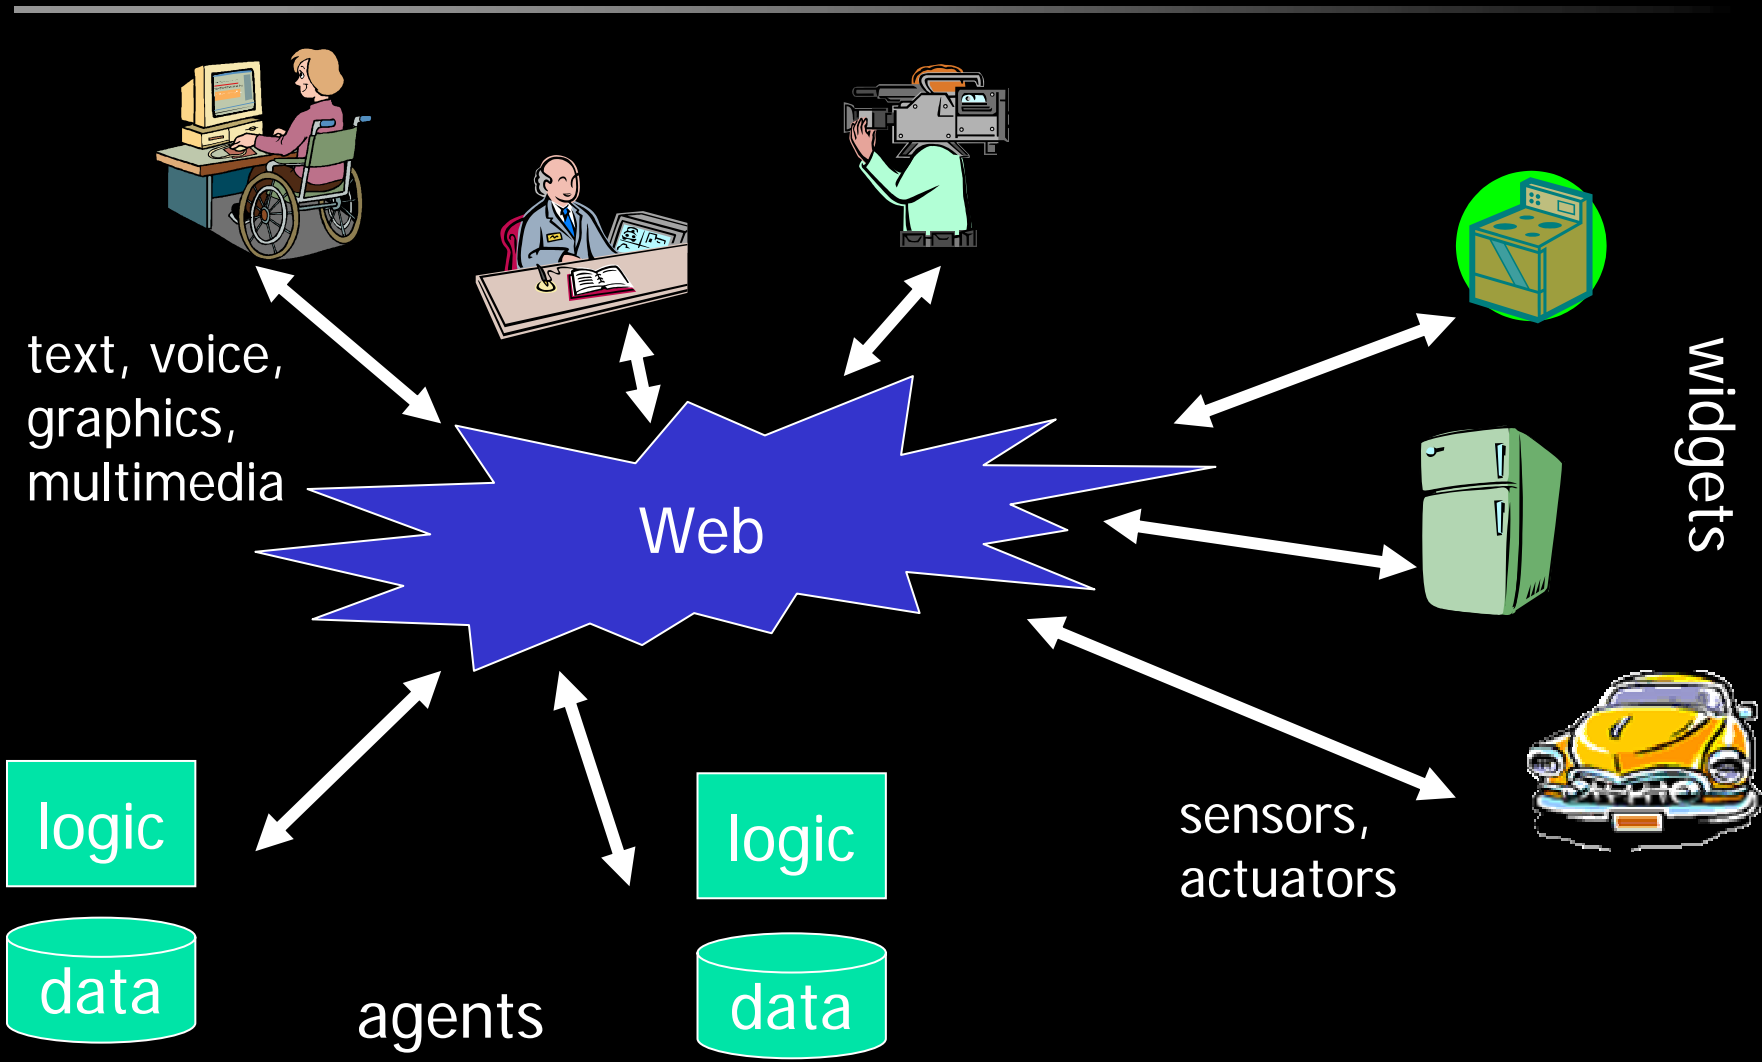
- Exponential growth in compute power
 - doubles every 2 years (at constant price)
- Exponential growth in storage (memory, disk)
 - doubles every (18, 12) months (at constant price)
- Exponential growth in bandwidth (LAN, WAN)

- Billion transistor chip and terabyte disk in 4-6 years

Can exponential growth continue forever?

- No major technology barrier prevents trends to continue for 10-15 years
- Practical barriers arise 15-20 years from now, for current technologies
 - physical limitations (quantum effects for charges and magnetism)
 - economic limitations (price of manufacturing)
 - disruptions due to differential growth rates
 - the “von Neumann” bottleneck
 - the disk head bottleneck
- Physics do not prohibit continued exponential growth for many generations
 - no inherent physical constraints on computation
 - no practical physical constraints on information

Pull: the ubiquitous information infrastructure



Pervasive digital infrastructure

- Most interactions (human-to-human, human-to-widget) are mediated by digital devices
- Most data generated by humans and sensors are stored
- Digital devices and stores are all connected
 - raise level of interaction with physical world
 - enhance human information processing capabilities
 - abolish distances in space, time, language, ability
 - integrate data to create actionable information
 - data \Rightarrow information \Rightarrow knowledge \Rightarrow action


The information revolution

- Information technology ushers a revolution more fundamental than any seen so far by the human species
 - the industrial revolution enhanced our muscles and our ability to change the physical world
 - danger: irreversible, undesirable changes in the physical world
 - the information revolution enhances our brain and our ability to sense and process information, as an individual, and as a community
 - danger: we loose contact with reality
- Revolutions take time...

Human-to-widget example: car electronics

- New capabilities, due to better control
 - emissions and fuel economy, safety, comfort
- Cheaper
 - fewer wires
- Advanced diagnostics
 - Preventive maintenance
 - Remote maintenance
- Malleable
 - late binding, customization, field upgrades
- Risky
 - bugs, security

Modern car has more than 50 microprocessors

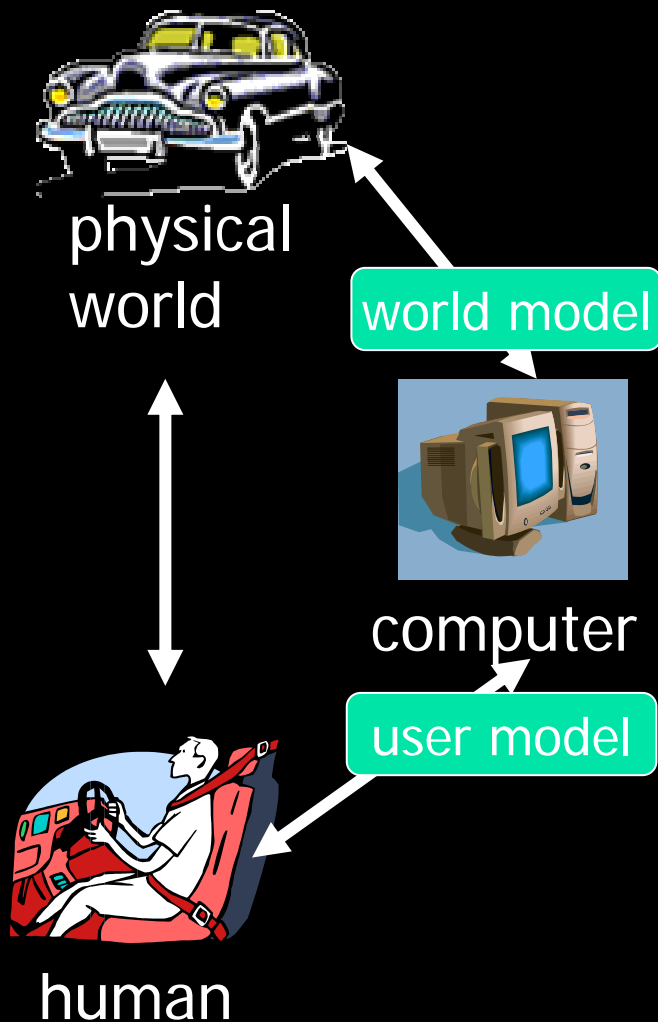


Same ideas applied soon in medicine (implanted sensors and artificial organs)

Future car evolution

- Enhanced user interface
 - multi channel
 - context aware
- Enhanced reality display
- Automatic driving
 - vision

Future information systems

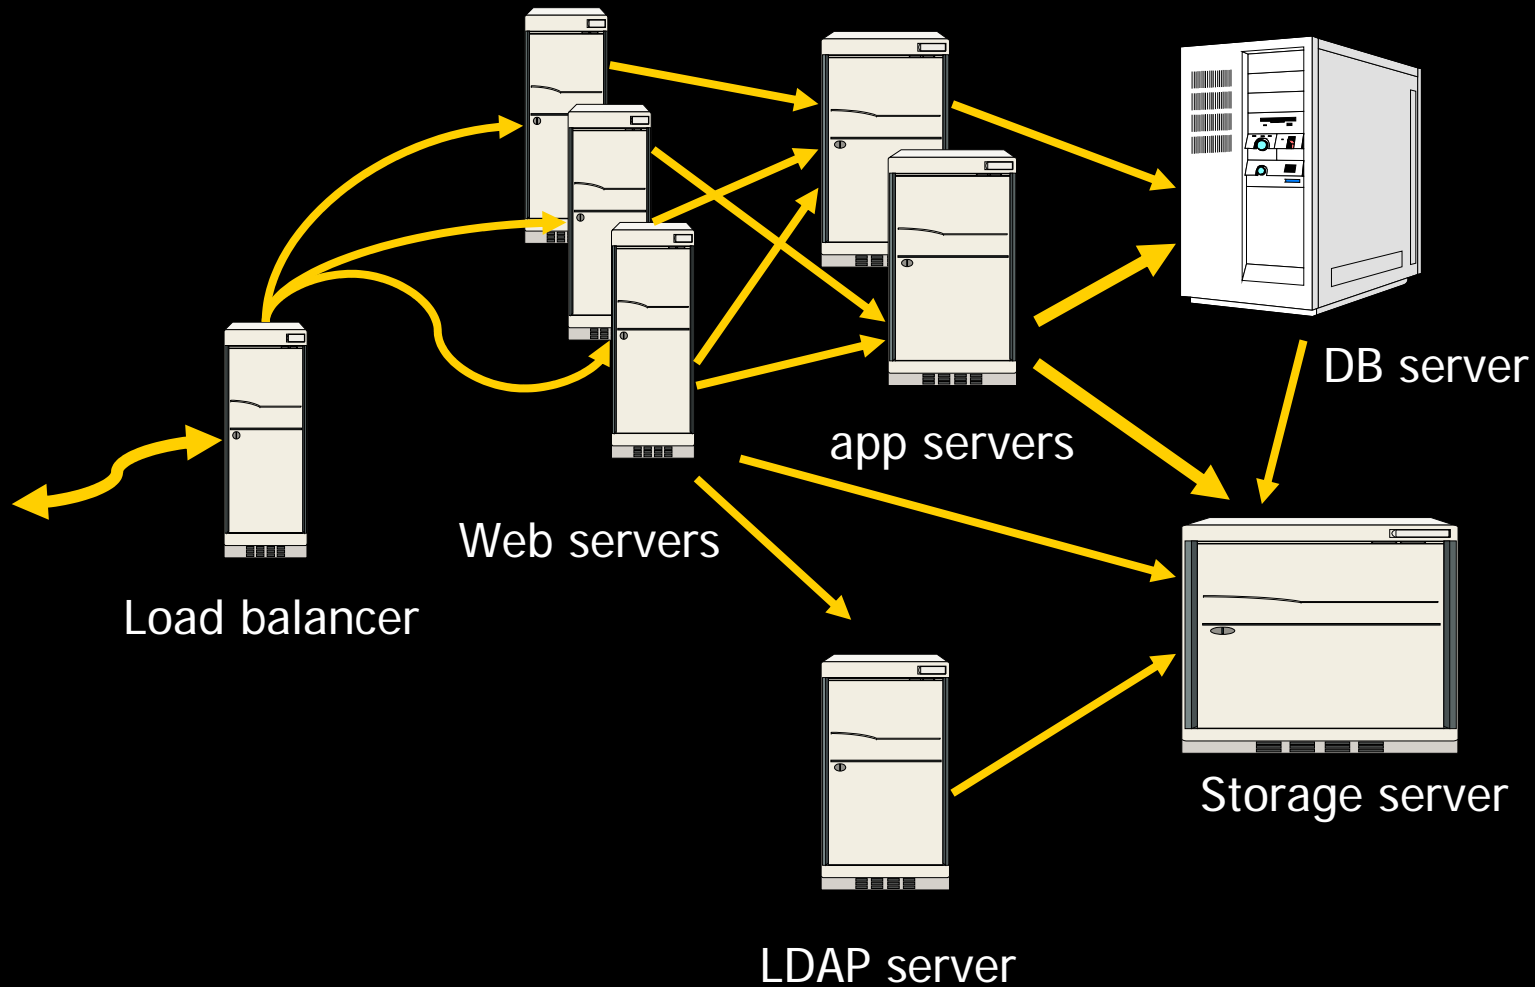


- Tight coupling: system needs to be studied as an integrated whole
- AI is a necessary technology for the computer node
 - the ability of simulating human cognition so as to maintain an adequate model of user state
- Not machines that think, but machines that know what we think

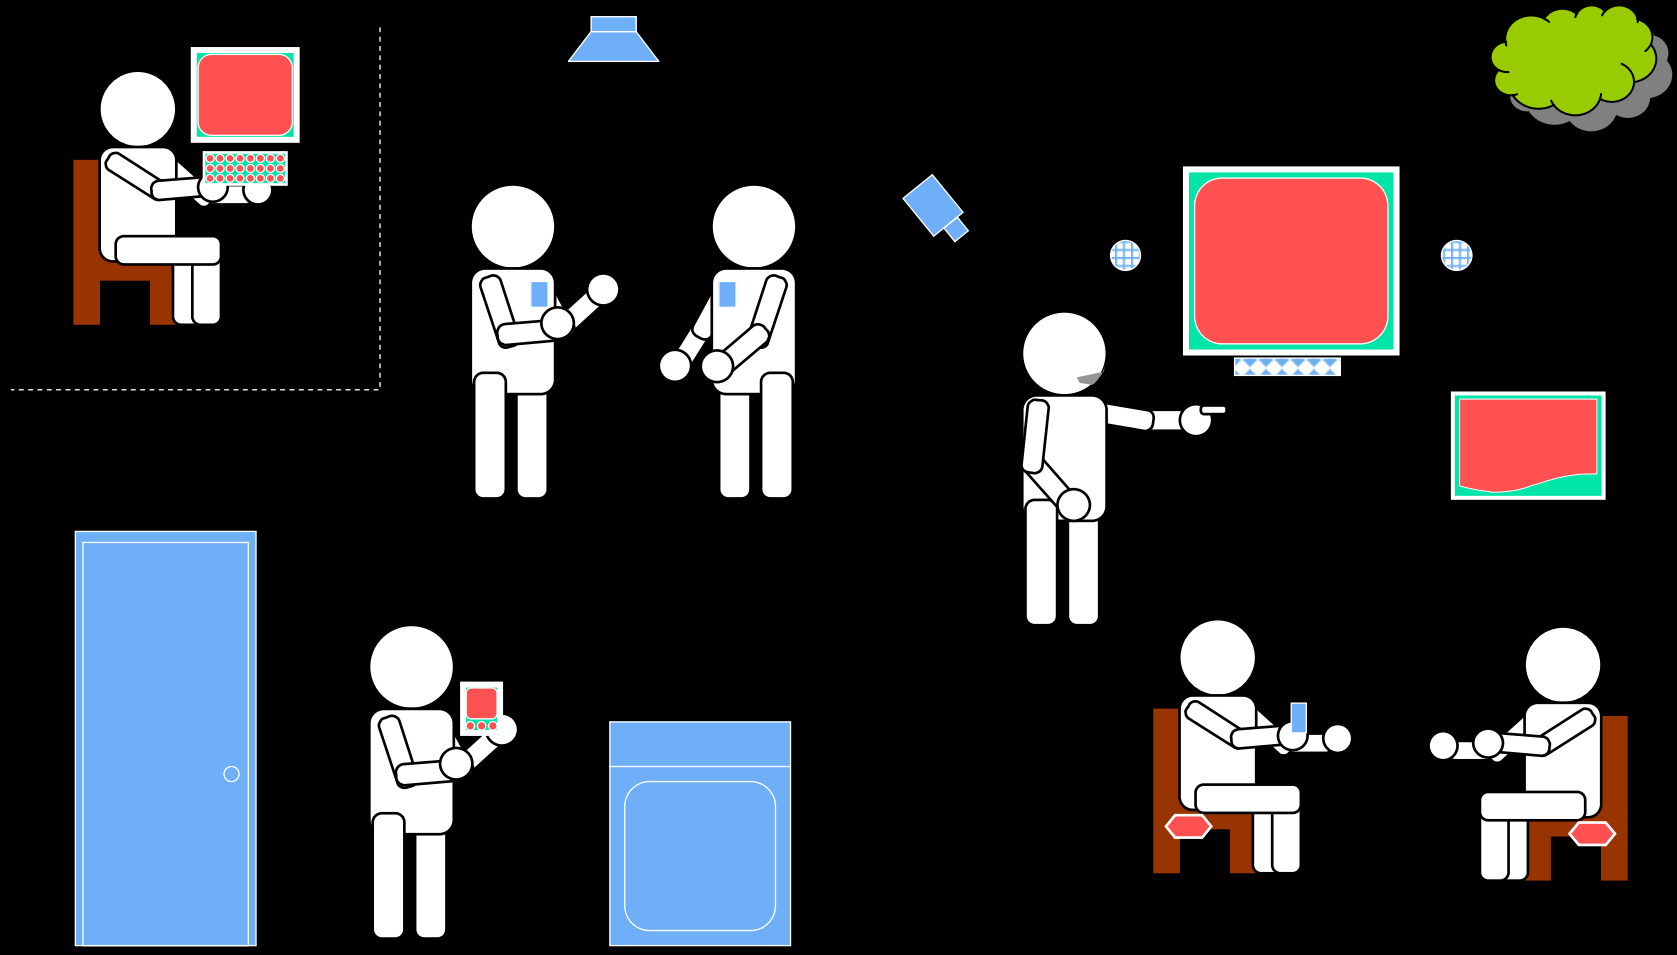
What is a computer?



The distributed server



Smartspaces – the disembodied computer

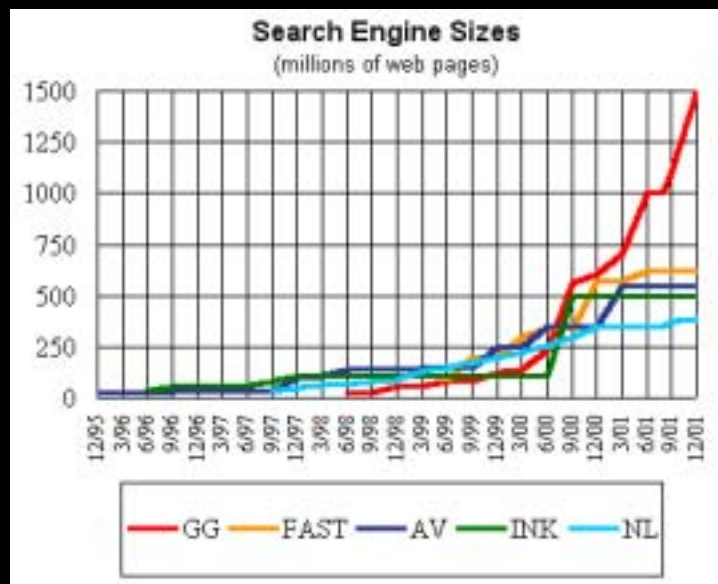


From the closed world to the infinite universe

- The closed computer is replaced by a dynamically changing assemblage of hw and sw components, within a connected universe of billions of devices
 - discovery
 - dynamic binding (hw, OS, application)
 - adaptation
 - migration
 - performance management
 - security...

The world of bountiful data

Web



>2B pages indexed, as of now

Commercial and scientific databases

TB	Database
40	Sloan Digital Sky Survey Raw Data Volume
24	Walmart's item information
20	Contents of Library of Congress
12	Two Micron All-Sky Survey (2MASS) Raw Data Volume
6	Entire indexable WWW in February 1999

From data to information, to knowledge, to action

- How do find the needle in the haystack?
- How do manage the tower of Babel?
 - the semantic web utopia
 - semantic integration
 - common language is natural language
 - learning: deep reasoning from few examples vs. shallow reasoning from many examples

The main impediments to the information revolution

- Can we program it? (productivity)
- Will it work? (correctness)
- Is it safe? (trustworthiness)
- The grand scandal of CS: no, no, no
- The CS community has a key ethical responsibility to ensure that software used in critical systems works safely

Why does software fail?

- Complexity of software systems
 - some inherent, some superfluous
- Interactions across large social systems
- Brittleness of software languages, systems, methodologies
 - due to the use of too powerful languages and tools?
- Inherent complexity of mesoscale systems
- Need new ideas!
 - autonomic, adaptive, reflective, learning subsystems
- Need new ethos for software engineers

CRA conference on grand research challenges in CSE

- safety.net: Ubiquitous Computing for Disaster Mitigation, Response and Recovery
- A Teacher For Every Learner: Scalable Learner-Centered Education
- Systems You Can Count On
- Enhancing Individual Productivity and Capability Through Cognitive Assistance
- Conquering Complexity: Building Systems with Billions of Parts

Systems you can count on

- To create the infrastructure for *Systems You Can Count On*, making today's applications reliable and secure, and enabling whole new classes of critical services
- Today, information technology is *the weakest link* in many of our most critical applications and services. Digital computing and communications increasingly pervade our lives, our economy, and our nation's critical infrastructure. In many applications this technology simply can't be counted on; it creates problems that range from chronic aggravation to unacceptable vulnerability
- Our challenge is to design a new generation of *Systems You Can Count On*:
 - Systems that are reliable, secure, available, predictable, and trouble-free
 - Systems that can be used with confidence in all elements of our global critical infrastructure
 - Systems that enable new services ranging from a currency-less society to 24x7 location--independent personalized healthcare
 - These goals require fundamentally re-thinking the way we design, deploy and support our global infrastructure— an essential technical grand challenge

Conquering Complexity

- To reformulate computing systems architectures at all levels (from circuits to global-scale distributed systems) that break through the complexity wall to deliver robust, scalable, long-lasting, systems.

The social dimension of computing

- IT is driven to a large extent by standards and conventions
- As IT becomes a critical infrastructure, govt. and society will increasingly involve itself in the definition of those
 - privacy (and snooping) laws
 - IP laws
 - export controls
- CS community need to become more involved in public policy issues

Summary

- Information technology is likely continue its rapid evolution for 10-20 years, driven by hardware evolution and insatiable appetite for new applications
- The main obstacle to such a continued evolution is us (lack of adequate software technologies)
- CS will continue to be a quick changing discipline, where fast innovation is essential
- As long as the rate of change keeps up, predictions, including these ones, are likely to be proven false