Computing at Extreme Scale –

The Challenge and Excitement of Modern Computing

Stuart Feldman
VP, Engineering
Google

Unicamp,
Campinas, Brasil
11 August 2008

Outline

Future of Computers

Expanding Frontiers of Computer Science

Demands of modern online information systems: search and ads

Google's unique approaches to infrastructure: hardware, software, middleware

Successes and Problems

The Future – Concerns and Opportunities

Who Am I

Vice President-Engineering for Google's "East Coast" engineering offices

- includes Canada, Brazil, Chicago, Texas

Also, Past President of the ACM

On a number of boards and committees

- US National Science Foundation (2)
- Business Schools
- Australia's NICTA

Past

- PhD in theory of galaxies
- Early member of the "Unix bunch" at Bell Labs
- Wrote Make, first Fortran 77 compiler, etc.
- Went to Bellcore Research – software engineering and a detour to large-scale development
- Went to IBM Research – Internet, E-commerce, long-term computer science research.

Future of Computers

Scale Measures

Number of users

- Ever
- Recent
- Simultaneous

Amount of Data

- Stored
- New
- Changed

Amount of Networking and Communication

- Messages per second
- Bytes per second
- Kilometer-bytes per second

Amount of computing

- Total operations
- Operations per byte of data
- Bytes of data analyzed per second

Processors

Hardware progress has driven field forever

- Range of possibility relevant to topics, details, constraints

Current state

- Stuck on practical thread speed
- Able to execute lots of independent threads- multicore and massive multicore
- Wirth’s Law as well as Moore’s
- Specialized hardware possibility

Other Hardware

Other relevant changes

- Interconnect
- Network
- Memory
- Storage hierarchy

Challenges

Long term bets

- Quantum computing
- Molecular, etc.

Return of Architecture, after years dormant

- Won’t be rescued by or run over by clock speed
- Large potential factors for specific approaches
- App/Software/Hardware co-evolution

The Multi-Core Challenge

Hardware management

Software Structure

Middleware

Where will the apps come from

Who will re-think and rethink them?

Expanding Frontiers of Computer Science

Landscape of Advanced Computing

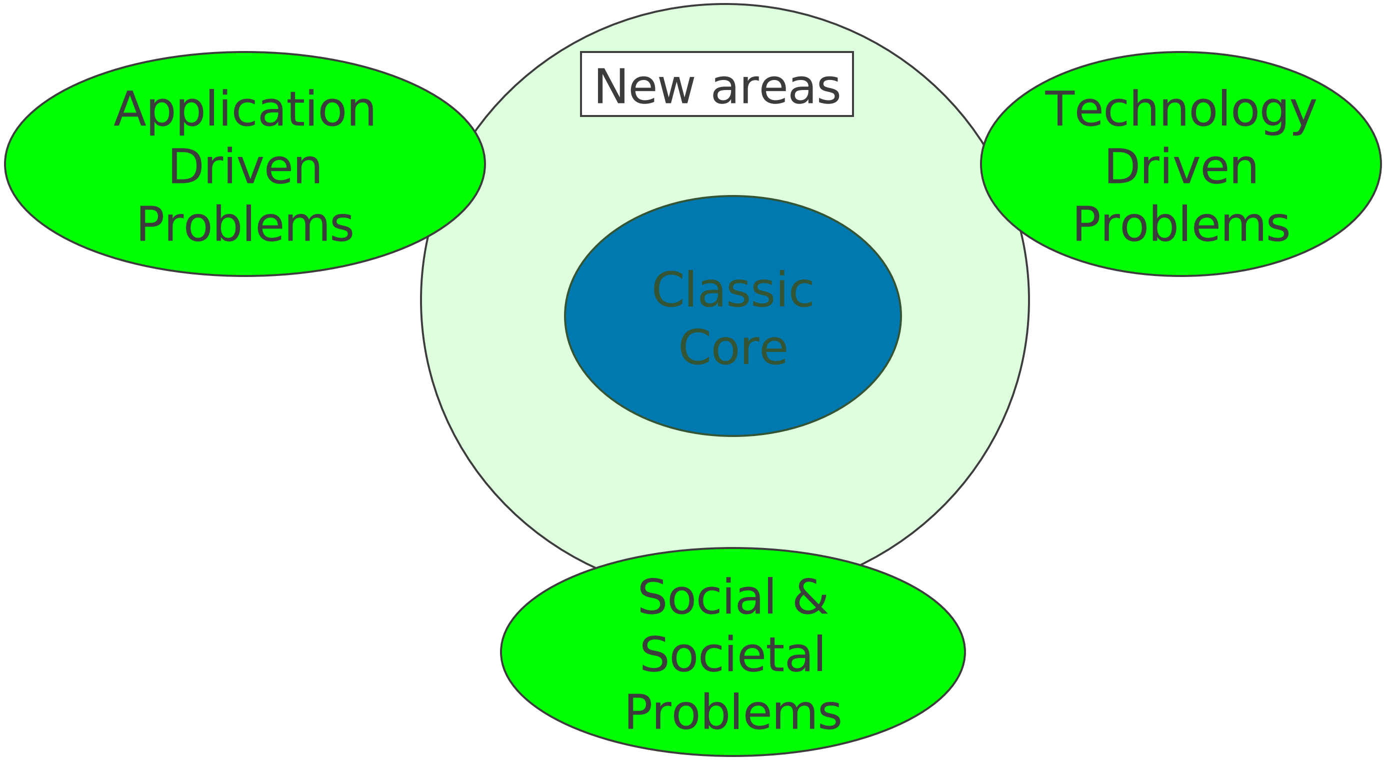

Research Fields

Typically focused on major insights, newly accessible problems

Driven by key theorems and demonstrator systems

Result in funding streams and publication outlets

The Expanding Core

Theory and basic models

Algorithms and complexity

Numerical Analysis

System and Hardware Architecture

Programming Languages

Operating Systems

Artificial Intelligence

The Expanding Core

Pre-1980 curriculum fields	Newer fields
Theory and basic models	Software Engineering
Algorithms and complexity	Databases
Numerical Analysis	HCI
System and Hardware Architecture	Information Retrieval
Programming Languages	Networking
Operating Systems	Distributed Systems
Artificial Intelligence	Storage
	Security

Pre-1980 curriculum fields

Newer fields

Punctuated Equilibrium

Fields seem to move in lurches, based on new insight, fashion, opportunity – cf Darwinian evolution

Examples

- Object-Oriented programming
- Hypervisors
- Relational Databases
- Web search

Future of the Computing Sciences

Fields Outside the Core

Real-world drivers

Very hard problems indeed

Intersections of multiple fields

Hybrid vigor

Are these technology and application driven areas the true future of CS?

Technology-Driven

Datamining

High-end graphics and interfaces

Super-Computing

Web technologies

Application Domain-Driven

Enterprise computing

Personal Computing

Bioinformatics

Healthcare

Services Sciences and Management

Auto-informatics

e-Government

Video, Gaming, 3D, …

Computing at the Extremes – And How it Expands CS

Classic High Performance Computing

Google-style massive data intensive computing

Highly distributed systems (the Internet, sensor/effector nets)

Enormously complex systems (enterprise computing)

Adaptive systems (service systems, embedded systems)

High Performance Computing

Classic high end scientific computing on supercomputers

- Whatever a supercomputer is today

Petaflops are the new (soon to reached) frontier)

- And people are already talking about exaflops.

How to harness huge parallelism for problems that aren’t perfectly symmetric

Performance tuning ranges from algorithm re-design to cache-sensitive coding.

Enterprise Computing

Not just “boring old MIS”

Genuinely complex problems

- Interacting rule sets
- Changing goals and measures
- Legal constraints
- Real-world requirements

Legacy issues exacerbate

Social & Societal Issues

Privacy

Employment

Economic shifts

Defense and implications

Networking and collaboration

And many many more

MODERN INFORMATION SYSTEMS ARE HARD TO BUILD

Google™

Demands of Modern On-Line Information Systems

Web Search is hard because of the structure of the web and the increasing expectations of users

Making a business out of search through advertising requires enormous computing resources

Supporting millions of users simultaneously is hard – whether it's games or spreadsheets

Structure of the Web

Key features of the Web – from the earliest days

- Easy but one-way links
- Freedom to create formats and content
- Extensibility
- Open protocols

All good choices

- Resulted in the explosive growth of the WWW
- Contribute to computational difficulty

Search

Earliest search engines did simple indexing

- Find as many pages as you can (or can afford)
- Find pages that satisfy query
- Figure out the most interesting

Didn’t work terribly well

Next generation based on structure of the web

- Properties of the graph, especially its link structure
- “Hubs and Authorities”
- Jon Kleinberg (Nevanlinna prize 2006), etc.

Realistic searches depend on much more

- Dynamic and static information

Search Quality

Practice is enormously data dependent

Observe user reactions (click through etc.) on many millions of queries to improve algorithms

Behavior not stationary

- Interests shift
- World responds to Google changes
- Spam and gaming

In real life, >200 factors are considered

Search – Computing Challenges

Google innovation – PageRank

- Still a key input to decisions
- Requires global knowledge of graph

Computing Effort

- Huge memory
- Web-scale analysis
- BIG indexes

User expectations

- Low latency and consistency
- Reliable responses

Fast response to many customers requires

- Huge networking
- Huge computing
- . . .

Algorithmic challenges

System challenges

Advertising

New generation of advertising

- Democratic (anyone can place an ad)
- Fair (auctions and quality rule)
- Valuable ($ billions)
- Trackable (can monitor actions, pay for results)

Approach designed for web

- But may translate to other media

Requirements

- Need to deliver ads fast
- Need to make decisions fast
- And LOTS of them

Advertising – Computing Challenges

How to identify promising ads to show

- Based on customer activity – e.g., queries
- Based on customer interest – e.g., content
- Based on advertiser preferences – e.g., demographics
- Decisions based on limited information
- Running myriad simultaneous auctions

Off-Line computations

- Clustering and machine learning with thousands of computers

How to do everything at high speed and accuracy

Applications and Hosting

Google supports a variety of applications

- Data-driven (e.g., Finance)
- User-driven (e.g., spreadsheet, document)

Requirements

- Availability
- Data protection

GOOGLE'S APPROACH TO INFRASTRUCTURE

Infrastructure Challenges

How to scale to meet expectations (billions of pages, billions of potential users, …)

How to do fault-tolerant computing in real life?

How to distribute the load?

How to make these problems manageable?

How to make these problems programmable?

How to maintain security and confidentiality without hobbling use?

Google™

Prefixes – Scientific Notation

Just a little review

- Mega = million = 10^6
- Giga = billion = 10^9
- Tera = trillion = 10^{12}
- Peta = quadrillion = 10^{15}
- Exa = quintillion = 10^{18}

Peta for Google is like Tera for most other companies

Platforms – Servers and Hardware

Google designs our own hardware,

- tailored to our specific needs

Cannot always get what we need in the market

- we take risks on the "bleeding edge"

Numbers – very very large

Infrastructure – Platform Software

Operating systems

- we make our own modifications to Linux

Management systems

- Network management
- server monitoring
- cluster management

Infrastructure - Middleware

Distributed processing

- MapReduce

Storage

- Classic data management systems (file systems, database systems) are not a good fit to Google's computing scale and approaches
- We have built many systems
  - Google File System
  - BigTable,
  - Many more to provide different tradeoffs of transaction, availability

Distributed and Parallel Computing

Distributed computing is the default

- If it doesn’t scale, we don’t do it
- Cluster “cloud” computing is what we do

Fault tolerance is essential

- Hardware breaks
- Data problems arise
- Networks partition
- Systems stall

Performance always matters

- People demand fast and predictable response
- Loss of customer data is unacceptable

Summary

Google operates at impressive scale

- Transactions per second
- Pages crawled and analyzed – number of pages, amount of data growing
- Size of indexes – sizes will grow by factors
- Number of advertisers and advertisements
- Complexity of engineering
- Size of data centers

Life will only get more challenging and exciting

- More sources and types of information
- More users in more situations
- More applications
- Whole new areas of activity – to be explored

SUCCESSES

OPPORTUNITIES

CHALLENGES

Google- Successes

We handle a remarkable number of queries per second

Our information is fresher and more up to date than ever before

Our ad targeting gets better

Our data are kept secure

Google - Opportunities

More types of information are available and interesting

- And we want to organize it all and make it accessible worldwide

Computation costs are continuing to fall, as are networking and storage

Numbers of users and numbers of devices increasing

Mobile access provides new possibilities and needs

Google Challenges - Technical

- Processors are not speeding up, just becoming more numerous
  - the industry challenge of multi-core
  - energy consumption – power is expensive and hard to get
- Networking is not uniformly available
  - inconsistent networking and bandwidth
  - intermittent access
  - different capabilities in different regions
    - affects latency, optimal locations of caches and computing
- Mobile access presents numerous difficulties
  - nonuniform devices
  - user expectations ahead of device capabilities
- Handling new data types is difficult and compute-intensive
  - Images, video, audio

QUESTIONS AND

(maybe) ANSWERS ?

ADDITIONAL TOPICS

Future of Computing Engineering and Software

Engineering Approaches to Computing Problems

Performance modeling and analysis

System Engineering

Software Engineering

Computing Approaches to Engineering Problems

Measurement and analysis of engineered systems

- and natural ones too

Process analysis and management

Business analysis

Simulation

In silico science

System Engineering

System Engineering

- Quantitative models
- Deep analysis of performance, distributions, security, etc.
- Whole-system understanding, not just local optimization
- Control theory and instabilities

Humans in the loop

- Systems are for people
- Systems depend on people

Software Engineering

Wouldn’t that be a nice change

Software Engineering

Wouldn’t that be a nice change

How many “software engineers” are comfortably quantitative?

How many really understand “real” engineering

What techniques carry over, which do not?

Special techniques, capabilities

Software Evolution - Methodologies

Some development models and methodologies

- Waterfall
- Concurrent
- Individual
- Chaotic
- Distributed
- Agile (choose your brand)

They differ in rigidity of organization, definition of roles, time scales, product attributes, applicability

All have religious supporters and condemners

Some are demonstrably more appropriate for particular types of product and schedule than others

All can be misused

Finding the right combination is difficult but essential

Software Evolution - Languages

- Assembly language/low-level languages
  - Close to hardware model, every little bit of performance matter
- Higher level languages
  - fairly direct translation but more problem-oriented syntax
- Popular paradigms and abstractions
  - structured programming
  - object-oriented programming
  - functional programming
  - service-oriented computing

Software Evolution – Delivery Models

Single purpose software

- Single purpose, single installation
- Programmers often stay with the app

Packaged/general purpose software

- Must address wide range of potential uses
- Products are delivered to a customer, who
  - May (or may not) use it
  - At some unknown time
  - In some unexpected way
  - For a surprisingly long time
  - Expects upward compatibility, perpetual back-level support, etc.

Services

- Created with the customer, delivered at time of production
- Can change from interaction to interaction
- Provide results not mechanism
- Results can be measured and service improved

Future of Computing Research

Wide Open Frontiers

The classic core continues essential

- Algorithms, analysis, etc. are the mark of a computer scientist
- Need to support and extend

New fields being added

- Much of growth and excitement naturally comes from new topics and ideas that should be addressed

Driven by new uses and applications

- Which introduce new questions

A Truly Global Activity

Computing research was once narrowly practiced

- Small number of countries funding
- Small number of universities participating
- Small set of companies contributing
- Hardware expensive and rare

Now a truly international activity

- Collaborative research, funding
- e.g., ACM conferences, publication, and membership

Funding and Support Trends

Government funding fairly strong worldwide

- Governments see correlations if not necessarily causation
- Rigorous objective selections create halo effect

Classic industrial research has been decreasing

- Decline of old near-monopolies, increasing competition

But new entrants in specific arenas

- Own high technology
- Cooperation with other researchers

Future of Computing Jobs

The Good and Bad News

Market is strong for the very best

- Silicon Valley
- Google
- Many other sites

The Good and Bad News

New needs

- Deep modern computing knowledge
- **T** and **Π** shaped people
  - Broad knowledge (top part of letter) plus depth in one or two specifics
- Locality still matters – many positions are not exportable or fungible
- Occupations are changing
  - Partly as a result of progress in computing research – more automation, more networking, more formalization

The Good and Bad News

Real concerns

- Globalization, offshoring, etc. are real and present concerns, both for national competitiveness and for individual career choices
- The field continues to change – must continue to learn and adapt
- ACM Job Migration Task Force report

Landscape of Advanced Computing