

计算机学者发表论文之目的与追求

张晓东

计算机科学与工程系

俄亥俄州立大学

(Ohio State University, USA)

Purposes of Scientific Research

- **Discoveries**
 - New understanding and creating new knowledge
 - Formulating new problems
- **Innovations**
 - Solutions (algorithms, methodologies, models, et. al)
 - Foundational theory and useful systems
- **Turing's computability describes "D" and "I" as**
 - (1) problem abstractions, (2) algorithms, (3) reasonable complexity.
- **Well documented papers for community and society**
 - Influential presentations with clear information flow
 - "Science is organized knowledge. Wisdom is organized life".
Immanuel Kant (1724-1804, 康德, 德国哲学家)


Major Computing Technology Advancements are Driven by Academic Research in Universities

- Brooks Sutherland Report, 1995, for US Academies of Sci/Eng shows
 - ✓ **Every aspect of computing** in our lives can trace back its root of basic research
 - ✓ **Basic research** is starting point for advancement without any patent protections
 - ✓ Claude Shannon, **Boolean logic, information theory**, MIT, 1937
 - ✓ Len Kleinrock , **packet communication**, UCLA, 1964
 - ✓ Alan Kay, **small talk** (original concept of object oriented programming), Utah, 1969
 - ✓ Bob Metcalfe, **Ethernet**, Harvard, 1973
 - ✓ Bill Joy, **BSD Unix**, Berkeley 1977
 - ✓ D. Patterson, **RISC**, Berkeley, 1981
 - ✓ T. Berners-Lee, **WWW**, CERN, 1989
 - ✓ L. Torvalds, **Linux**, U. of Helsinki, 1991
 - ✓ Mark Andreessen, **Mosaic browser**, UIUC, 1994
- Computer business could make some wrong predictions
 - **T. J. Watson, IBM Founder**: “There is a market for may five computers”, 1943.
 - **K. Olsen, DEC Founder**: “There is no reason anyone would want a computer in their home”, 1977.

Business Driven R&D are Common in Other Disciplines

- **Industrial Revolution** (18th to 19th century, 100 years)
 - James Watt's Steam Engine (1763-1775), it was patented.
 - many other patented innovations in coal mining, textile manufacture, metallurgy, chemicals, and gas lightening.
- **Wright Brothers** built the first airplane in the world in 1903
 - “Flying Machine” was patented
 - “Wind-tunnel” was patented
 - It took 40 years to enter massive productions (World War II)
- **All Edison's inventions** have been patented (1,093)
 - Light-bulb (1879), phonograph, motion pictures, film projector, ...
 - It took 60 years for US families to widely use light-bulbs

Documented Results and Concepts are Important



Von Neumann model

- a **memory** containing both data and instructions
- a **computing unit** for both arithmetic and logical operations
- a **control unit** to interpret instructions and make actions

- Before **Von Neumann's** computer project, several operational computers were built:
 - 1936, Zuse's **Z1** (1st binary digital computer) in Germany
 - 1937, Atanasoff and Berry's **ABC** (1st electronic digital computer) in Iowa State University
 - 1943, **ENIAC** in UPenn
- The most important milestone Von Neumann left is his paper: “**First Draft of a Report to the EDVAC**”, 1945. (a consulting report to US Army)

Computer Science Research Makes Direct Impact to Society

- **New York Times' top 20 best innovations in last 30 years** (judged by Wharton Business school, Upenn, March 2009)
 - (1) Internet, broadband (2) PC and laptop (3) Mobile phone
 - (4) e-mail (5) DNA sequencing (6) MRI (核磁共振成像)
 - (7) Microprocessors (8) Fiber optics (9) office software
 - (10) Robotic surgery (11) open source software (12) LED(发光二极管)
 - (13) Liquid crystal display (14) GPS (全球定位系統) (15) e-commerce
 - (16) data compression (17) microfinance(小额信贷)
 - (18) Solar energy (19) large wind turbines (20) Social networks
- **A unique and high expectation to Computer Science Research:** it must be relevant to the advancement of the technology to improve our society


追求与目的之一: 把握技术发展方向和应用需求, 做重要的问题

- **CS research is technology driven and fast changing**
 - New problems are caused by rapid technology updates, such as Moore's Law, application demands, and infrastructure cost change
 - Many research problems come and go, and some are persistent, such as performance, security, reliability
 - There are a limited number of classical problems in our field. Knuth says, we have less than 100 general-purpose algorithms
- **CS research is easy to lag behind, “bubble”, or irrelevant**
 - Many algorithms design are architecture dependent, e.g. $O(n)$ measures CPU cycles, but performance determines by data
 - Many real-world issues can be abstracted into math problems with unrealistic assumptions to publish papers
 - Many important development are engineering oriented, needs only a few papers.

R&D of Applications Have Been Driven by Moore's Law

R&D of RDBMS to effectively utilize multi-cores

After 10,000 time increase in clock rate (400KHz in 1971 to 4GHz, in 2005), we have entered multi-core era by increasing # cores with low clock rate.



1970
Relational Data Model

1976:
System R & Ingres

Parallel RDBMS
DIRECT, Gamma,
Paradise, ...

Cache-Optimized RDBMS
MonetDB, StageDB, EaseDB ...


dropping DRAM price
from 400 \$/MB
to 0.09 ¢/MB


Multi-cores

Quad
core

Dual
core

Single-core Architecture has been Inefficient for 10 years





of transistors on chip increases following Moore's Law

386

Pentium 4

Dual-Core Itanium 2

- Transistors (000)
- Clock Speed (MHz)
- Power (W)
- Perf/Clock (ILP)

>10

What Matter most in Single-cores and Multi-core?

- Single-core performance mainly depends on **hardware advancement**
 - ✓ **Clock speed**: stop increase in 2005 due to power consumption limit
 - ✓ **Execution optimization**: pipelining, out-of-order, et. al. (little gain after 2000)
 - ✓ **Cache**: large and high associativity, major source for performance gain
 - single cores are inefficient** (**1,000+** times less from its ideal model now, and reach to **30,000** times in 2020)
- Multi-core performance mainly depends on **software advancement**
 - **Number of cores**: continue to increase (inevitable hardware advancement)
 - **Concurrency**: parallel/concurrent execution is scheduled by OS/software
 - **Shared cache**: data sharing and conflict avoidance is managed by OS/software
 - **Throughput = concurrency / latency**
 - Major efforts**: high concurrency (scheduling) and low misses (exploiting locality)

Architecture Teaching Lags behind Research

- ILP is still a dominant portion in architecture classroom (H&P Book)
 - ✓ **Instruction Set Principles and Examples:** Chapter 2, 72 pages
 - ✓ **Instruction-Level Parallelism and its Dynamic Exploiting:** Chapter 3, 117 pages
 - ✓ **Exploiting Instruction-Level Parallelism by Software:** Chapter 4, 175 pages

A total of 361 pages
- Only one chapter for memory performance
 - **Memory Hierarchy Design:** Chapter 5, 124 pages.
- **Other Chapters are non-microprocessor topics**
 - **Multiprocessor and Thread-Level Parallelism:** Chapter 4, 138 pages (SMP)
 - **Storage Systems:** Chapter 7, 101 pages (I/O is an independent topic)
 - **Interconnection Networks and Clusters:** Chapter 8, 90 pages (a part of parallel computing class along with Chapter 4, SMP).

追求与目的之二: 争取让更多的同行读我们发表的论文

- **Why publish in top venues?**
 - Each field has its own flagship and leading conferences and journals, the recognitions come from reputation and influence
 - Most researchers only read papers from these venues
 - Number of citations is an indicator of the amount of readerships
- **Conferences versus Journals**
 - In certain technology based fields, such as systems, architecture, networking, and databases, top conference papers are prestigious, highly visible, with 4-5 more reviews in depth, and low acceptance rate.
- **SCI and EI driven publications are not desirable**
 - An SCI/EI entry may not ensure the quality of the journal/conference
 - The quality should be judged by the peers in the field
 - The impact factor sometimes is field size dependent

追求与目的之三:发表有影响力的论文

- **Seeking useful theory to layout foundations**
 - As basic theories are missing or lagging behind, technology developers can only have ad-hoc and sub-optimal solutions
 - A true theory guides us to develop best solutions
- **Understanding intrinsic nature of complex problems**
 - A complex problem is a problem that we do not understand
 - “The search for simplicity is the search for a structure within which the complex become transparent” (R. Gallager, a pioneer of computer networks)
- **Examples of influential papers**
 - Shortest path algorithm, (Dijkstra’s notes, 1959)
 - Relational data model, (CACM’70)
 - Co-scheduling in distributed systems (ICDCS’82)
 - HITS algorithm lays out a foundation for search engine (SODA’98)

追求与目的之四: 更新和创造核心技术, 推动社会进步

- **Systems research has a high potential to directly advance technology**
 - Solving real problems in core components of system infrastructure
 - Creating new system frameworks for low costs and high returns
 - Building new systems that responds rapid demands in society
- **One more step after research prototypes can make a big difference**
 - Open source software, Linux, BSD, MySQL, Hadoop, BitTorrent, ...
 - Commercial products, RISC CPU, GPU, Yahoo!, VMWare, ...
- **Examples of influential papers changing the world**
 - Ingres/PostgreSQL (TODS,76, most advanced open source database)
 - 汉字激光照排技术, (计算机学报, 81, 90+%的国内外华文报业市场)
 - RAID, (SIGMOD'88, \$10 billion storage market)
 - WWW, (A proposal of Information management, CERN, 1989)
 - Google, (WWW'98, the World dependent search engine)
 - Computer architecture advancement, (by influential papers)
 - OS advancement, such as Linux and BSD, (by influential papers)

A List of Research Contributions by UC Berkeley (1957-2007)

- **Theoretical Foundations and Influential Algorithms**
 - Fuzzy logic (Zadeh)
 - Theory of NP-Completeness (Cook and Karp)
 - Karmarkar's Algorithm for Linear Programming (Karmarkar /Karp)
 - Complexity of Cryptography (Blum, Micali and Goldwasser)
- **Advancing Technology and Improving Society Productivity**
 - Design automation: SPICE and others (Pederson, Rohrer, Sangiovanni, and Newton; -> Cadence and Synopsys)
 - INGRES (Stonebraker and Wong; -> Oracle and PostgreSQL)
 - Berkeley Unix (Fabry/Joy; -> Sun, FreeBSD, and NetBSD)
 - IEEE 754 Standard for Floating-point Arithmetic (Kahan)
 - RISC (Paterson/Sequin; -> chip designs in Sun, Fujitsu and others)
 - RAID (Paterson/Katz/Gibson; -> more than 15 vendors)

Requirements to Computer Science Ph.D. Students

- **Rigorous training with a strong foundation**
 - Both systems and analytical skills
 - merge one into a focused area
- **Identify significant problems**
 - Literature search and reading, attending seminars
 - Interns in industries and research labs
- **Completing a Ph.D. dissertation**
 - Working hard (100 hours per week)
 - Articulating ideas and results well in papers and presentations
 - Achieving goals #1 and #2
 - Aiming at goals #3 and 4

Results of Computer Science Ph.D. Research

- 读博士，就是挑选一个狭窄并重要的领域作研究，毕业的时候交出一篇世界一流的毕业论文，成为这个领域里世界首屈一指的专家。任何人提到这个领域的时候，都会想起你的名字。
- 你从学校带走的、最有价值的、不是这份论文，而是你分析和独立思考的能力、研究和发现真理的经验，还有科学家的胸怀。当你某一天不再研究这个领域的时候，你依然能在任何一个新的领域做到最好。

王选定义的一流研究

我认为在国内，达到以下四个目标之一的项目均属第一流成果。

1. 文章发表在权威杂志上，有国内外学者的引用。
2. 在原理性样机或系统中，提出了新思想、新技术或新方法，被后来的研究者采纳，并最终影响了工业界。
3. 把新技术首次应用于某一领域，并在应用中效果显著，从而得到了一定的推广，随之而来的是该领域的商品热潮。
4. 风靡市场的商品

王选, “科研选题和制定目标时要注意的几个问题”, 《计算机世界》1994年

- Aiming high to **technology transfer** (goal #4), and he himself is a role model
- **Influential scholarships** also include theoretical foundations and leading new directions

Summary

- CS Research advances technology and impacts the society
 - Computing technology in the society **deeply roots to research**
 - Patent-free, open source, and conference driven research make our field **fast moving and highly competitive**
- How to measure achievements of Goals #1 and #2
 - **Citations**
 - High citations reflect high readerships and follow-ups
 - High citations means the area is active
 - However, high citations can also mean the area is “bubbled”
 - **Persistency and classics** of the problems
 - **Potentials** to achieving goals #3 or/and #4
- Measuring Goals #3 and #4 needs hard evidences and time
 - **Quality citations:** laying out foundations (5-10 years)
 - **Updating technologies:** new standard, new model, new products
 - **Systems are widely used:** supporting many applications

Last Words: Contribution-driven Research

- Einstein said:
 - “Try not to become a man of **success** rather try to become a man of **value**”
- **Success** can be measured by
 - A large number of publications, visibility, and more grants
 - Awards, promotions, and high salary
 - Upgrading official levels and respects
- **Value** is measured by impact and contributions
 - New findings and foundations for **advancing knowledge**
 - Research results directly **updating technology**
 - **How many people** use it, and **how large scope** it affects
 - **Service** by widely used software, textbooks, and inspiring lectures.

Thank You !