Peter J. Denning

Editor: David Walden

Photo courtesy of Louis Fabian Bachrach.

A leading scientist in computing since his graduation from Massachusetts Institute of Technology in 1968, Peter J. Denning is best known for his pioneering work in virtual memory, especially for inventing the working-set model for program behavior, which eliminated thrashing in operating systems and became the reference standard for all memory management policies. He is also known for his work on the principles of operating systems, operational analysis of queueing network systems, the design and implementation of the Computer Science Network (CSNET), ACM digital library, and codifying the principles of computing. A primary goal of Denning's career has always been promoting the science in computer science through education, research, and the general health of the field.[1]

***David Walden*: Please tell me a bit about your early life.**

Peter J. Denning: I had interests in math, science, and nature from a young age. At school I was too small to be any good at athletics, which were socially popular, so I devoted myself completely to academics, which were not.

By age 12 I developed an interest in magician performances, especially those that depended on mathematical tricks. By age 13 I had discovered a deep fascination with electricity and electronics, which seemed to have a magic all their own.

My parents sent me to Fairfield Prep in 1956 to get me into an intellectual community and out of the athletics-infatuated public school culture. Under the wing of a gifted science teacher, I entered three science fairs with computers made of pinball parts and vacuum tubes—one to compute sums, one to solve linear equations, and the last to solve cubic equations. The second computer won the science fair. The third computer worked perfectly but fared poorly at the fair because I paid no attention to marketing and presentation—a valuable life lesson.

From Fairfield Prep, I went to Manhattan College to study electrical engineering in 1960. Although short on computing, its curriculum gave me a solid grounding in practical engineering—the building and testing of things people could use.

I came out on top of my class at Manhattan in 1964 and got a National Science Foundation fellowship good at any graduate school. I applied to MIT in fulfillment of my father's advice (he had wanted me to attend MIT rather than Manhattan).

***Walden*: Say a bit about MIT.**

Denning: MIT had a completely different philosophy from Manhattan about EE principles and organization. To prepare for the PhD exams at the end of first year, I took all the MIT EE core courses in addition to my required master's courses. That intense preparation was barely enough. With the help of my master's thesis advisor, Jack Dennis, who took me under his wing, I passed the PhD qualifiers on the second try. He and I have had a long and productive friendship for almost 50 years.

My master's thesis was about scheduling requests for a rotating disk or drum memory so as to minimize mean access time, a critical issue for an experimental time-sharing system Jack Dennis had been developing. During that year, I worked closely with Allan Scherr, who taught me about systems programming, language design, compiling, data collection in an OS kernel,

Background of Peter J. Denning

Born: 6 January 1942, Queens, New York.

Education: Fairfield College Prep, 1960; Manhattan College, BEE, 1964; Massachusetts Institute of Technology, MS, 1965; MIT, PhD, 1968.

Professional Experience: Princeton University, 1968–1972; Purdue University, 1972–1983; NASA-Ames RIACS, 1983–1991; George Mason University, 1991–2002; Naval Postgraduate School, 2002 to present.

Honors and Awards: Southern Connecticut Science Fair Grand Award, 1959; Princeton Engineering Teaching, 1971; IEEE Fellow, 1982; AAAS Fellow, 1984; ACM Distinguished Service, 1989; ACM Fellow, 1994; ACM Karlstrom Outstanding Educator, 1996; ACM SIGCSE Outstanding Educator, 1999; Centennial Engineer, Manhattan College, 1999; George Mason University, School of Engineering, best teacher, 2002; George Mason University, best teacher, 2002; SIGOPS Hall of Fame, 2005; NSF CISE Education Fellow, 2007; Postel Award for CSNET, 2009; ACM SIGCSE Lifetime Achievement, 2009.

 Published by the IEEE Computer Society 1058-6180/12/$31.00

discrete simulation, and queueing theory. Through the thesis, Jack and I showed that shortest latency time disk scheduling was optimal for time-sharing systems.

On passing my PhD qualifiers in the spring of 1966, I decided to tackle a much tougher resource allocation problem, which was looming in the design of Multics. The problem was how to build a stable computing system from multiprocess computations, which could have large variations in their processor and memory demands. I had to learn how to measure the demands of multiprocess computations, configure a system with appropriate capacity for the demand, and manage the allocation of CPU and memory dynamically. Jerry Saltzer told me of thrashing, a major instability they were encountering with multiprogrammed virtual memory systems, and challenged me to find a solution. That solution turned out to be much harder than either of us imagined. My quest produced the theory of locality, the working set model for program behavior, and a method of system balance for optimal control.[2]

During my PhD years, I also helped Jack Dennis teach a course on computational models.

***Walden*: I was a student of yours in that course; it was a great course.**

Denning: I loved teaching that material and developed a deep understanding of computation and the essential role of machines in doing it. Our class notes caught the attention of a Prentice-Hall editor, and we signed a contract for a book, *Machines, Languages, and Computation*, in 1967. Unfortunately, writing a book was more work than I ever imagined—we did not finish until 1978.

In January 1968 Jack told me I had plenty of material for my PhD thesis. I went on a crash program of writing and working with my committee. I graduated with my MIT PhD in May 1968.

As graduation approached, I pondered where to go next. I had offers from MIT and three other universities. I chose Princeton because it was more attractive to my family.

***Walden*: At Princeton you continued and expanded the scope of your research in the areas of operating systems, as well as teaching and beginning other research.**

Denning: Yes, my four years at Princeton were productive. I developed and taught new courses in the principles of operating systems and computer architecture. I took on two PhD students and worked with several others. I collaborated closely on several projects with computer scientists Ed Coffman, Jeff Ullman, and Al Aho and with electrical engineers Stuart Schwartz and Bruce Eisenstein. Those projects extended the working set theory and validated it with experiments. They also codified operating systems principles.

Prior to Princeton, I helped Jack Dennis organize the first ACM Symposium on Operating System Principles (SOSP), held in Gatlinburg, Tennessee, in 1967. At Princeton, Ed Coffman and I organized a follow-on, SOSP-2, in 1969. There was a huge interest in gaining a fundamental understanding of operating systems, which were the most complex computing systems then known. The SOSP has continued every two years since that time.

In 1969 and 1970, I chaired a task force for the NSF Cosine (Computer Science in Engineering) project, which was developing prototypes of new core courses for computer science programs. I invited Jack Dennis, Nico Habermann, Butler Lampson, and Dennis Tsichritzis to the team on OS principles. Our recommendations, released in 1971, were adopted nationally as many universities created their first systems-oriented core courses.

After the task force, Ed Coffman and I decided to write a book with the bold title *Operating System Theory*.[3] Published in 1973, it contained the best material we could find on the fundamental principles of operating systems.

***Walden*: You were then recruited to Purdue University?**

Denning: By my fourth year at Princeton, promotion was not looking good because of a cap on tenured faculty—no more than two promotions in the engineering school in the next five years and at most one in our EE department, where CS was a minority. Early in 1972 I encountered Sam Conte, the CS chair at Purdue, on an elevator at a conference. He said, "I hear you are looking around. I can make you an offer as tenured associate professor and pay you 50 percent more salary." Now that was a great elevator pitch!

I interviewed at Purdue in the dead of winter. The faculty members were warm and welcoming. I accepted an offer from Sam a few weeks later.

My Purdue years were also productive. With the help of several graduate students, I continued the working-set project. We showed that the working set model was very general; it could simulate any paging algorithm with memory contents that obeyed an inclusion property with increasing value of the control parameter. We validated the phase transition model of locality and used it to show that working-set memory management was within 5 percent of optimal for memory controllers that could not see the future. We answered the basic questions about locality and memory system management we set out in 1966. I published my final research paper on working sets in 1980. Working-set principles were widely adopted in operating systems.

In the meantime, I had taken up research into system performance modeling, working with Jeff Buzen, who had made queueing network models popular in 1971 with his fast algorithm for computing throughput and response time. He was interested in reformulating queueing theory around assumptions commonly met by computing systems. I helped him develop the "operational analysis" theory and showed that it applied to real computing systems.

During the Purdue years, I published another book, many more papers, and advised 10 PhD students.

***Walden*: Your wife Dorothy also was at Purdue, correct?**

Denning. I met Dorothy in late 1971 when I visited the University of Rochester to interview for chair of the new CS department, where she was an instructor. I saw her again at the 1972 Spring Joint Computer Conference. When she described her interest in getting a PhD, I suggested she apply to Purdue. Sam Conte was impressed with her credentials and hired her as an instructor. We soon fell in love and got married in early 1974. We have had a long and happy marriage, and she has become renowned in data security and cryptography. She was recently inducted into the Cyber Hall of Fame.

***Walden*: At MIT you were already beginning to be active in ACM activities, including editing and great concern for good writing. I remember visiting your office where you were reading Strunk and White's little book on writing style.[4]**

Denning: Let me address the writing topic first. From a young age, I had a knack for writing and liked doing it. In high school, my science teacher gave me my first teaching experience, for which I wrote a series of lectures about basic electricity for the science club. I also wrote articles and even drew cartoons for the school's magazine. In college, I won a couple of essay awards. At MIT I wrote extensive course notes, which as I mentioned earlier, resulted in a book-publishing offer. Jerry Saltzer completed his PhD thesis in 1966; it was a masterpiece of clear exposition. Jerry told me that he valued clear writing and used Strunk and White as a guide. I got a copy, loved it, and used it as my guide too. I was well prepared for an active publication and editing subcareer.

I learned about ACM from MIT faculty and joined as a student member in 1965. At the SOSP in 1967, Walter Kosinski, the co-organizer with Jack Dennis, suggested that I could broaden my interests and link up with more people in operating systems by joining the Special Interest Committee on Time Sharing (SICTIME). The principals of SICTIME not only welcomed me as a member, but they recruited me to be the SICTIME newsletter editor. I was honored to be able to contribute and meet new operating systems people at the same time. I wrote and distributed SICTIME newsletters for the next two years.

Next I was involved in transitioning SICTIME to a Special Interest Group (SIGOPS). I was put in charge of writing up the transition plan and the new bylaws. We got approval from the ACM Council in 1969, and President Bernie Galler appointed me the chair of the new SIGOPS.

In early 1970, the ACM Council approved a bylaw change that created a SIG board to oversee the 32 SIGs and added its chair to council. I got a call from the nominating committee that spring asking if I would be a candidate for SIG board chair. I was flattered but reluctant; I didn't want to be so involved with ACM that I would jeopardize my reputation as an academic. I decided I would run, reckoning I would not be elected since the other candidate was an old and respected hand in the SIGs. To my surprise, I won the SIG board election and became a member of the council.

That was indeed a turning point. My four years as the SIG board chair began a four-decade long series of high-level ACM positions (including president). Here is a partial list:

- editor in chief of *Computing Surveys* (1976–1978);

- president (1980–1982);
- editor in chief of the ''new *CACM*,'' to help the *Communications of the ACM* become a flagship magazine to further everyone's understanding of computing (1983–1992);
- chair of the Publications Board, where I led the team that developed, implemented, and delivered the ACM Digital Library to help establish an electronic community and expand the number of people who had access to CS knowledge (1992–1998);
- chair of the Education Board, where we initiated a new curriculum recommendation in 2001 and a new development online professional development center (1998–2004);
- leader of the ''IT Profession Initiative'' with the ACM Council to have ACM become the society for computing professionals (1999–2002).
- member of the Education Council (still today).
- editor in chief of *ACM Ubiquity*, an online peer-reviewed magazine devoted to the future of computing and the people creating it (since 2008).

I learned how to arrange my daily schedules so that I would have time for both my professorial and ACM work. It is possible I could have been even more productive with academics and research, but even with the weight of a ''second career,'' I think I was plenty productive enough to solidify my position as a genuine academic.

***Walden*: I read your articles in *American Scientist* for a number of years. How did that come about?**

Denning: The *American Scientist* project began in 1985. The AmSci editor, Michelle Press, had seen my ACM president columns in 1980–1982 and some of my other writings on behalf of computational science. She thought I could do a column for them. The goal was to reveal the science of computing so that scientists in all the fields served by AmSci could come to appreciate the field. In those days, many people thought computer science was a guild of programmers. The idea that computing is science was novel. From 1985 to 1993, I wrote 47 columns for the magazine. One of my undone projects is to scan those columns and put them up on my website—the science is unchanged!

***Walden*: Despite a successful tenure at Purdue, including becoming department head, you left and went to RIACS and from there to George Mason, where you were involved in the teachings of Fernando Flores. Eventually, you ended up at the Naval Postgraduate School. It seems like you were seeking something.**

Denning: Indeed, a lot of dots to connect! The best way to understand my transitions is to understand my quests.

My lifelong professional quest has been to understand and communicate the fundamental principles of computing. This has been driven by my ongoing sense of the magic, beauty, and joy of computing and by my ongoing conviction that computing can make big contributions across the board in the physical, life, and social sciences.

Another quest, which developed in the 1980s, is to help computer scientists successfully contribute their computing expertise to people in other fields seeking to solve problems. I never liked the caricature that computer scientists were antisocial. I wanted to help my students overcome that image.

Dorothy and I wound up moving to new locations and organizations about every 10 years in order to immerse in environments where we could most effectively pursue our quests.

While at Purdue, I eventually and reluctantly concluded that our computer science research was mostly not interesting to people in other departments. Thus, I decided to move to NASA-Ames, a different environment where I could learn more about how to make computer science more interesting. (Besides, Dorothy and I had been dreaming of getting to California.)

I arrived at NASA-Ames in Mountain View in 1983 to found the Research Institute for Advanced Computer Science (RIACS). Our mission was to help NASA's research groups move forward through computational science, a combination of computation and computer science. We were quite successful. Still, we had problems with customer relations because the CS notion that ''breakthroughs cannot be scheduled'' often clashed with the NASA notion of regular deliverables. This was not good in a research environment funded only by customers expecting operationally useful results. I became interested in more professional development. Terry Winograd introduced me to his colleague Fernando Flores, who had a new theory of management based on communication

Computer science has always been my first and foremost career.

rather than decision making. I studied management and leadership with Flores and began to apply his philosophy to my work at RIACS. I was amazed at how much more productive we became and how much more trust we earned from our customers.

***Walden*: Didn't CSNET happen about this time?**

Denning: The CSNET project overlapped Purdue and RIACS. In 1979, I teamed up with Larry Landweber, Dave Farber, and Tony Hearn to propose a computer science network that would make the Arpanet technologies available to all CS researchers and their students. The Arpanet-connected universities were rapidly moving ahead of the others, and we did not want to be left behind. We got the CS community behind the CSNET proposal and NSF funded it for five years. By 1986, CSNET included 120 CS departments and industry labs, about 50,000 faculty and students, and was self-supporting.[5]

***Walden*: How did RIACS wind down?**

Denning: By 1990, I had concluded that the customer-relations issues we experienced at RIACS were widespread in computing research. This pointed to a defect in the way we educated computer scientists, and I wrote a manifesto about needed reforms in computing education. I needed to return to academia to pursue this agenda. Dorothy and I found two positions in the Washington, DC, area in 1991—she became CS department chair at Georgetown, and I did the same at George Mason.

At George Mason, I concluded that I still needed more professional development with leadership to help interested faculty to change and help students learn how to get their ideas adopted. I undertook further leadership studies with Richard Strozzi Heckler in my spare time. I learned how to coach students and was certified as a master somatic coach. I initiated a design course called Sense 21 that aimed to put Flores-Heckler principles to work to help students be more successful as designers and innovators. Sense 21 was a surprising success. Its graduates formed an alumni club so that they could continue learning together. (None of the graduates of my operating systems classes ever wanted to form an alumni group!)

Also at George Mason I continued work on my quest to understand the fundamental principles of computing. I designed a new capstone course on the core of information technology.

***Walden*: And that wound down too?**

Denning: Yes, Dorothy and I wanted someday to return to California. Suddenly in 2002 an opportunity appeared when two professor positions opened at Naval Postgraduate School in Monterey. We took them and returned.

At NPS I focused on two main projects: great principles of computing and innovation. My NPS colleagues and I designed a course on the great principles that became one of our department's best. In 2004, our faculty reorganized the curriculum using the principles framework as a guide. Our website (greatprinciples.org) has all the materials.

The other project was about innovation. After all my years leading research, teaching students key leadership practices in Sense 21, and working with my colleague Bob Dunham to teach leadership practices to business people, I realized I had learned a lot about innovation. Bob and I collaborated on *The Innovator's Way*,[6] which is about eight essential practices for successful innovation.

In 2009 I organized a conference called "Rebooting Computing: The Magic and Beauty of Computer Science." Many people in the field were in anguish over the lack of acceptance of computer scientists as peers in other sciences and the decline of students choosing CS as their college major. The purpose of the conference was to stimulate conversations between segments of the field that had little to say to each other and, in so doing, to recover the sense of magic and beauty that had brought so many others into the field. We succeeded!

To sum up, computer science has always been my first and foremost career. ACM became a sort of second career that never eclipsed the first, but provided many rich channels to pursue my quests. My pursuit of innovation practice in the last two decades has become a third career. And my quests continue.

I have always been drawn to Lao Tzu's statement, ''A leader is best when people barely know he exists. Of a good leader, who talks little, when his work is done, his aim fulfilled, they will say, 'We did this ourselves.'''

***Walden*: You have had, are still having, an exceptionally broad career—technical innovations, education and professional society activities, and methods of innovation in business.[7] Thank you for this interview.**

Denning: My pleasure.

References and Notes

1. This interview was conducted by email in July and August 2012. The original 6,400 word interview was edited down to the length of this column by the interviewer with the approval of the interviewee.
2. P.J. Denning, ''Working Set Publications,'' http://denninginstitute.com/pjd/PUBS/Workingsets.html.
3. E.G. Coffman and P.J. Denning, *Operating Systems Theory,* Prentice-Hall, 1973.
4. W. Strunk and E.B. White, *The Elements of Style,* 4th ed., Longman, 1999.
5. P.J. Denning, A. Hearn, and C.W. Kern, ''History and Overview of CSNET,'' *Proc. ACM Symp. Comm. Architectures and Protocols* (SIGCOMM 83), ACM Press, 1983, pp. 138–145; http://www.internetsociety.org/sites/default/files/pdf/Comm83.pdf.
6. P.J. Denning and R. Dunham, *The Innovator's Way Essential Practices for Successful Innovation,* MIT Press, 2010.
7. See http://denninginstitute.com/denning for a full biography.