

In Memoriam Kenneth Clem Sevcik

May 7, 1944 - October 4, 2005



March 26, 2006

Ken Sevcik passed away on October 4, 2005. Ken was an exceptional mentor, an exceptional academic, and a man of profound wisdom. This tribute to Ken and his achievements is written in recognition of his great intellect and his generous friendship.

Ken received his B.S. in 1966 from Stanford University and his Ph.D. in 1971 from the University of Chicago. Ken joined the faculty at the University of Toronto in 1971, and was Chair of the Department from 1990 to 1992. He also served as Director of the Computer Systems Research Institute (CSRI) and on the Governing Council of the University of Toronto.

Ken's work in performance evaluation has been called "an exemplar of experimental computer science." He was a founder of the field of computer system performance. He made influential contributions to both the theory and practice of computer system performance evaluation. His seminal work on queueing network models has had a lasting impact on industrial practice and provides a foundation for much of the performance research that has followed it. Database researchers know him for his work on grid files, an early multidimensional indexing technique that remains on the required reading list in most Ph.D. programs, and for his extensive work on database system performance. He was an early member of the ACM SIGMETRICS Board of Directors, the premier organization for the promotion of research in performance analysis. In 2004, he received the SIGMETRICS Achievement Award in recognition of a career devoted to sustained, long-lasting, and high-impact contributions to computer system performance evaluation in the technical, leadership, and mentoring dimensions.

Ken's research program has been devoted to answering the difficult performance questions that arise in the design, implementation, and operation of complex computer systems. He is credited with having brought some of the most effective tools in programming analysis to practicing engineers. Together with his graduate students and research collaborators, he devised new techniques for constructing mathematical models of the essential characteristics that govern the performance of complex systems. With three co-authors (including two of his former students), Ken wrote the 1984 book *Quantitative System Performance* which is widely used by researchers, educators, and practitioners responsible for designing and managing computer systems. It is still today the definitive work on analytic modeling of computer systems.

Ken is also responsible for some of the deepest insights in queueing network theory. Together with Isi

Mitrani, he established the equality of two distributions arising in quite different contexts. This result, now known as the Sevcik-Mitrani theorem, forms the basis of mean value analysis. This and other foundational results were central to enabling the widespread, practical use of queuing networks.

Effective management of very large databases requires making performance decisions beyond those that determine the underlying hardware of the database system. Logical data design, physical data design, index selection, and query optimization all influence the overall performance of a database system. Ken proposed a novel layered methodology for predicting database system performance. This approach has been used for performance prediction in several commercial database management systems. His recent work branched out beyond system performance prediction to apply robust information-theoretic techniques to the representation and management of large, high-dimensional data sets.

In large-scale scientific computing, there are major performance issues related to determining how different types of scientific applications will perform on diverse architectures (for example, on large Cray supercomputers versus hypercubes or clusters of slower, less expensive processors). Ken made fundamental contributions to formally describing the architecture of various machines, including a precise characterization of their capability to carry out multiple operations in parallel. Based on his models, he was able to characterize the amount of parallelism available to diverse applications and predict application performance. This work was complemented by his involvement in the Hector project at the University of Toronto, an early influential project focused on building a highly-parallel multiprocessor.

Ken's facility in mastering a diverse range of areas, spanning the breadth of computer systems research, is further illustrated by his work while on sabbatical at NASA Ames Research Center. He evaluated candidate mechanisms and protocols for the communication network that was being designed for the U.S. modules of the Space Station. One aspect of this work involved experimentation with, and assessment of, the priority mechanism in the then recently standardized FDDI protocol for optical fiber-based local area networks. This work was followed by a series of influential papers on the design of inter-connection networks to support multiprocessor systems with a large number of processors.

Ken is a rare breed of university researcher who has received international recognition from both the academic and industrial research communities. In addition to his SIGMETRICS Achievement Award, his contributions to industry were recognized by a prestigious Award for Excellence in Research from the Information Technology Association of Canada (ITAC) and the National Science and Engineering Research Council (NSERC). This award recognizes research excellence of benefit to academia and industry.

Throughout his career, Ken has been committed to professional service. His CV lists numerous influential positions he held in support of his department, his colleagues, his university, and his field. To quote Ed Coffman: "Ken owes much of his exceptional leadership skills to the simple fact that he combines a gentle and compassionate personality with his invariably correct, well-informed, and firmly held views of his profession and how it should be evolving. He is and always has been a role model among role models."

Perhaps Ken's most remarkable intellectual legacy are his students. We reproduce here a tribute written by four of Ken's students, on the occasion of his 60th birthday. We then conclude with some memories from Ken's friends.

Ken Sevcik as an Advisor and Mentor

Ed Lazowska, University of Washington
Satish Tripathi, University at Buffalo
John Zahorjan, University of Washington
Derek Eager, University of Saskatchewan
August 2004

Ken Sevcik joined the Computer Science faculty at the University of Toronto in 1971, following his doctoral studies at the University of Chicago and his bachelors studies at Stanford. Like many of Toronto's early hires in Computer Science, Ken was persuaded by then-chair Tom Hull to "give it a try for a couple of years," which became a career.

The four of us are privileged to have been among Ken's early Ph.D. students numbers 2, 3, 5, and 8 of his 22 Ph.D. students as of 2004 (see Table 1 at the end of this document). Ken also has graduated 50 M.Sc. students as of 2004 (see Table 2), and the 7 of his 22 Ph.D. students who pursued academic careers have given him, thus far, 61 Ph.D. grandchildren, 86 great grandchildren, 32 great great grandchildren, and 2 great great great grandchildren (see Figure 1) a total of 203 Ph.D. descendants in all!

As faculty members ourselves, we often reflect on the characteristics that made Ken such a successful advisor and mentor, and try as best we can to emulate them.

Ken is open and approachable. When you meet Ken, whether for the first time or the thousandth, you can count on being received with a smile. No matter how busy Ken is, while you're talking with him that conversation has his full attention. No matter how naïve what you have to say is, he gives it full consideration (before gently pointing out what might be wrong with it).

Ken is incredibly conscientious. As a student, you need many things from your advisor (including many things you don't know you need, but you do). You can count on Ken to provide these, both the good (e.g., introducing you to the community at conferences) and the bad (e.g., insisting that your thesis really does need another chapter's worth of work before it's complete). You can count on him looking out for your welfare, even when the effort to do so benefits him at most extremely indirectly.

Ken leads by example. He engages in all the aspects of the work that a student has to learn to do, as well as those carried only by the advisor. He sets high standards for his students, and even higher ones for himself.

Ken is always open to new ideas and new approaches, and eager to learn new things. His willingness to tackle new problem areas is evident from the titles of the Ph.D. and M.Sc. theses that he supervised: scheduling theory, queueing network modeling techniques, distributed algorithms, database management systems, applications of performance models, multiprocessor architecture, multiprocessor operating systems, communication network protocols, computer system performance measurement, file structures, spatial databases, and many more.

Finally, Ken managed to be our friend as well as our advisor, without letting either of those relationships interfere with the other. This is at the heart of the immense respect we have for him, both professionally and personally. In the final tally, how well one has done as an advisor is reflected in how much one's students have learned. We owe Ken a terrific debt for the lessons he gave us about performing and leading research, and about how to lead our lives.

I first met Ken when I was visiting the U. Chicago campus in 1967. He was admitted to the Committee on Information Sciences a year before me. The Committee was a small, friendly, interdisciplinary group, and it was Ken's welcoming enthusiasm, along with the charm of his classmates Rick Weiland and John Pomeranz, that made me feel at home immediately. I entered with a class that included Adele Goldberg, who later became quite well known through Smalltalk and her leadership in the ACM, and Alice Peters, now of A.K. Peters Publishing.

We were classmates and house-mates for about four years: Ken, Rick, and I lived with a few other grad students in faculty houses vacated during their sabbaticals. How we convinced homeowners that this was a good idea I still don't know, but I do know we sometimes did not internalize what home ownership and stewardship really entailed. An example of my responsible attitude to other's possessions: I borrowed Ken's car to drive 300 miles south over Christmas vacation. Borrowed, that is, without his permission or knowledge. And of course he came back early while I was still out of town. I should have emerged from that caper either plus an auto-theft conviction or minus some teeth, but Ken's grace and good nature transcended my awful judgement and I learned some lessons.

House-sitting allowed for a very comfortable life and extras like our wonderful year with Bonnie the Black Labrador: the four of us could almost keep up with her appetite for exercise. We had, even for that dissipated era, some really serious parties (I remember our department chair at the time chasing – as in running – after my date all over the house with serious intent, for example.)

Still, we were a rather structured, rule-bound group. We had mandatory sit-down-dinners at a given time five or six nights a week, being cooked in rotation by the house-mates. The idea was no dinner was ever supposed to be less spectacular than the last, and so our standards were high. For example, I remember helping to drip egg yolk into boiling honey for some Ethiopian dessert or other. We invited our seminar speakers over to dinner (van Wijngarten comes to mind). By house rules, you cleaned up the kitchen after the guy who cooked before you (so you couldn't complain about its state). So we were not discouraged from making a huge mess, as I often did with flour in my capacity as puff-pastry expert. We had another strange rule: the beer bill was simply divvied up evenly, so if you didn't drink more than average you subsidized those who did.

Ken had done some opera in college, and he got us involved in vocal projects, most notably singing Beethoven's Ninth with the Chicago Orchestra. His renditions from Don Giovanni inspired me at long last (which Oberlin College and a summer of doing musical theatre on Cape Cod hadn't done) finally to get interested in real opera.

He was a hard worker, regularly at his desk thinking and writing (we used pads and paper back then). I was always impressed with his powers of concentration and of course with his productivity. He was very organized and was a seriously get-things-done guy.

Ken got a job offer at Toronto and stayed there. We stayed in touch and I followed his rise to international fame and also the rise in his local administrative responsibilities and attendant crises, like a devastating fire during his tenure as chair.

He had always had a natural athletic talent and was able to take full advantage of the superb opportunities for playing hockey offered at Toronto. Our own CS basically-pickup hockey team narrowly defeated the third team of Cornell's CS department (we were officiated by players from the Cornell CS women's team, who were their first- or second-best, I forget). Anyway after this victory I raised with Ken the possibility of us playing the Toronto CS team. Ken didn't actually laugh out loud, and somehow managed quite diplomatically to let me know that the projected score of such a mismatch was about 75-0. It was a triumph of tact. We did skate together on at least one visit, and he and Carmen were gracious hosts whose cosmopolitan and intelligent company was always a highlight in our all-too infrequent trips to Toronto.

I last talked to Ken when he called me from his hospital room in August 2005. Despite the drugs and discomfort, he still was the same person, in a way that many of my contemporaries, buffeted about by Life, are not. He was the same selfless, sincere, good-hearted, optimistic, wise, good-humored, smart guy I met some 40 years ago. Neither the demands and realities of academic leadership nor his years fighting cancer had eroded his essential self.

I never heard one mean-spirited, cynical, or manipulative utterance from him, or heard him say anything negative about anyone else. His enormous, durable integrity and his rare generosity of spirit never failed him and will always be an unattainable example for me.

Chris Brown

I'm honored to have been invited to speak today. I'm the second of Ken's 22 Ph.D. students. I'm the root of more than half of his nearly 200 Ph.D. descendants. And I was his friend for more than 30 years.

Ken was an absolutely extraordinary person, and mentor, and friend.

What was particularly remarkable about Ken is that he was extraordinary in a collection of pretty ordinary ways. This made him a superb role model.

Although Ken was plenty smart, he didn't achieve his success by being smarter than everyone else. That's really important if you're a student, because you look to your advisor for traits you can emulate, and you can't say "I think I'll try to be smarter, just like my advisor."

And although Ken was plenty big, he didn't achieve his success by dominating people, either physically or intellectually.

Instead, Ken led by positive personal example:

- He was incredibly conscientious.
- He was always receptive to new ideas and new approaches, and eager to learn new things.
- He was open, and approachable, and non-judgmental.
- He treated everyone as a peer.
- He was a loyal friend, who welcomed all of us into his life and home, and who we welcomed into our lives and homes.
- He was always there for us personally as well as professionally.

Ken was extraordinary because he combined these traits better than anyone else I've known.

Ken was also a lot of fun:

- There was the softball, and the basketball, and the hockey. (Ken never really learned how to skate, but he was such a phenomenal athlete that he would always fall in the right direction to help the play.)
- There was the pub in the basement of his condo development that we used to refer to as "Corporate Headquarters."
- There was the way he calmed my nerves before my first conference presentation by focusing for weeks in advance on what an experience it would be to travel to the Western US with him and have my first Coors beer.
- There was visiting his mom and dad after a conference in Colorado, along with fellow students John Zahorjan, Satish Tripathi, and Derek Eager, and seeing all his old high school baseball trophies lined up on shelves in his room – a room that looked as if nothing had been touched since the day he left.
- There was the sabbatical we spent together – most of it in the pool where he tossed around our 2-year-old and 4-year-old for hours on end.

On the way downtown today, I asked my wife, Lyndsay, what she most remembered about Ken. She said that he was the world's best listener. And that he was also the world's best hugger. It just goes to show that there is no measuring up to Ken. She was kind enough not to also mention that he was the world's best looking man, but here's a story I remember clearly. At the end of my first year of graduate school, I was hit by a car while bicycling in to UofT, and spent a few months in Sunnybrook Hospital. One afternoon when Ken was visiting me, a cute young nurse came into my room – she was a UofT nursing student working there for the summer. She took one look at Ken, and dropped a cup full of thermometers on the floor, breaking them all. While Ken and I continued to work on my thesis, she scabbled around on her hands and knees gathering up shards of glass and pools of mercury.

Ken Sevcik was an extraordinary person. Each of us can pay tribute to Ken by trying to emulate his traits, making ourselves better people and the world a better place.

Ed Lazowska

My memories of Ken always include the hockey games that members of the department use to play every Friday afternoon. Some were good, and though Ken had never played before, he did become a good skater, and some such as me were not so good. I had never skated before and was made into a goal tender. But the games were fun, the beer afterward was fun, and all of us were much younger. My sympathies to his family.

Dick Swenson

I will remember Ken Sevcik with affection as a brilliant colleague and my co-author on two papers, as a devoted teacher and mentor, as my friend of close to thirty years and as my warm and charming visitor when I was in Paris. He is one of a small number of individuals who made the department of Computer Science at the University of Toronto stand out internationally.

With my warmest good wishes to his wife, students, colleagues and friends who stand proud in having known and admired him.

Erol Gelenbe

hink I

On the lighter side, lest we forget, there was a memorable basketball gang consisting of John (Skyhook) Zahorjan, Ken (The Rock) Sevcik, Fast Eddy Coffman, and Larry (The Franchise) Dowdy. (I hope I'm not missing anyone.) Ken's street name came from what it felt like when you ran into him or tried to out-rebound him. Had the [ACM] SIGs put together a basketball league, we would have been champs, easily. And Ken, Larry, and I played a few rounds of golf as well. Those were good old days; we need to revive that mix of sports and professional meetings that's gone missing.

Ed Coffman

Ken was the first database researcher (I always claimed him as a database researcher although his interests were more general) that I had met when I moved to Canada in 1984 after completing my PhD. He had come to give a talk at University of Alberta and I was among the dinner party. He had read a paper and a technical report that I had iend for more than 30 years.

Ken was an absolutely extraordinary person, and mentor, and friend.

What was particularly remarkable about Ken is that he was extraordinary in a collection of pretty ordinary ways. This made him a superb role model.

Although Ken was plenty smart, he didn't published as a graduate student (which surprised me) and had kind words about them (which made me feel like a king). Our interactions continued sporadically from then on. He was always very kind, supportive and humble despite his extremely high research profile. Someone once referred to him as an institution by himself. He certainly had done ground-breaking work, but it was his quiet leadership that was most important as far as I am concerned. He held leadership positions within Canadian research community and contributed greatly to Canadian computer science community. I will miss his common sense approach to issues, his kindness, and his generosity.

M. Tamer Özsu

I have a memory or two of Ken Sevcik to share, for your collection. My name is Robin Cohen. I was a graduate student in the Department of Computer Science at the University of Toronto from September 1975 to May 1983.

My first memory is of Ken participating in the infamous "Star Trek" skit that I co-wrote with Jim Hoover and Maria Klawe, for our yearly Christmas party. Ken played the role of the alien, and spoke the famous line: "I be M, the Monitor". He eagerly agreed to participate in the skit when we approached him – just another sign of his willingness to be a good sport.

My second memory is of watching Ken walk by very early one morning when I was playing intramural field hockey on a field on campus, and seeing him stop to wave hello and to actually watch the game being played for a while. Ken and I had played softball together in several of the games that were organized for grads and faculty together. He truly enjoyed sports and was there to support others who were into sports as well. It was very encouraging to see that he cared.

Although I am not likely to be able to attend the memorial service for Ken, my thoughts are with everyone else who will be remembering him that day. I will always have fond memories of his enthusiasm for all activities, both within and outside of academia.

Robin Cohen

I met Ken more than 20 years when we were both young and full of . Thatdreams. Our first meetings were in data base conferences. I was greatly impressed with Ken's demeanor - polite and respectful. He was

willing to listen to other ideas and had no pretensions. He knew his strengths and weaknesses, always giving his best to you.

In 1989, when I was asked by IBM to create a new centre for Advanced Studies, Ken was an obvious person to approach to ask for advice. The idea was to create a centre which would bridge between industry and universities, creating a win-win situation combining talent and resources from both. Ken's advice was to make sure to keep the best people, irrelevant of where they came from - Canada, internationally, systems or theory. He definitely had a vision that IBM would gain the maximum if they allowed researchers to work directly with developers, rather than other researchers at IBM. His other great advice was to cut the paperwork as much as possible. His vision was that if you requested people to do research, then you can't expect them to sign contracts to deliver something, the role of the developer. As a result, we built a model based on PhD and Master students, with very little paperwork or contractual arrangements. This was a breakthrough for IBM which was used, at that point, to sign very complicated contracts.

At the beginning of the 1990's the Centre was created, with Ken as a key advisor to both myself and others at the Centre. Within a few years, academia became more involved with the Centre, as Ken helped me to bring theory people from U of T, such as Professor Allan Borodin, Professor Stephen Cook, and Professor Charles Rackoff. System people, such as Alberto Mendelzon, John Mylopoulos and Rick Holt from U of T were also brought in.

The Centre, by this time, had many projects and close to 50 PhD students during the summers. All had a great time doing interesting research that was recognized both by IBM and the academic community.

When I left for academia in 1997, again Ken was there to give me advice and prepare me for such a change in life. I started my job at Dalhousie University for the newly created Faculty of Computer Science. As you can all guess, Ken continued to have an advising role. With his great advice, we developed and improved an interesting faculty.

Ken was always very modest, brilliant in what he did, but his social skills were the ones that impressed many that knew him. He managed to have everyone feel good about themselves, was optimistic and always had words of encouragement when you needed it. At the same time, he could gently criticize when necessary. He was very tactful.

I was one of those fortunate people, who over time, became very good friends with Ken. As a friend, he was extremely loyal and honest. He gave his honest opinion, whether or not you agreed with it, putting an optimistic spin on it. This continued even though he became sick with cancer, and continued through the illness. We kept up our correspondence through his illness. He continued to be optimistic and upbeat, not letting the illness become a major focus of his life. You couldn't help but admire him for being so full of life, knowing there was no cure. In didnturn, I also tried to keep upbeat, to show my support, and I have no doubt that his wife, Carmen, was a huge pillar of support to Ken during this difficult time. She would see what Ken wouldn't show the rest of us.

In June of last year, I received a disturbing email from Ken. Reading between the lines, I could tell that the cancer was winning the battle and Ken's time with us was growing shorter. Even at this time, he had a lot of dignity and left his door open, offering to have dinner whenever I came to town or wanted to talk with him about my own health problems. He always kept my spirits and my hopes high. I used him as my model on how to life my life. Following his example, I have learned how to life a full life, not letting problems get me down.

He left a lot with me. Although he is no longer with me in body, he is in spirit. I will definitely miss Ken - not only as a colleague but as a very good friend.

Jacob Slonim

Ken was an inspirational figure for everyone who knew him. I consider him as an academic father, as a person who inspired and inspires my work everyday. I first met Ken in the summer of 2000, when I discussed with him the topic of my PhD thesis. Since then, my work took on a different flavor. Ken was able to understand my problems in a matter of minutes giving me the proper advice. Even though, I was not officially his student he got excited with my topic and was actively participating in my research. Not to mention that I would call him the "proof-reading machine", with useful and constructive comments. Apart from a good advisor, Ken was an excellent friend. Nothing would replace his witty jokes, his joyful laughter and the long conversations we had about his trips all over the world. In recent years, he would teach me, as well as everybody else, what it means to be strong and enjoy every moment in life. He was and he is a

positive influence for me, and he is one of my examples when I am stuck or when I need to give advice to others. He will be surely missed as he was irreplaceable.

Periklis Andritsos

There are so many memories of Ken and not a single one I can think of that isn't kind.

Ken and I (and Prof. Chris Brown of Rochester) were all housemates on the south side of Chicago and department mates in the Committee on Information Science at the University of Chicago. We all shared a graduate assistants' office whose walls were plastered with our current projects and our current whims. It was a fledgling computer science department, in the 60s, filled with new ideas and high energy. No PCs or supercomputers in those days – we worked mostly on a hand-built 2nd generation computer named Maniac.

One particular series of memories I have of Ken relates to one of the as incredibly conscientious. young faculty members in the department, who was also filled with new ideas, although many of them a little half-baked. The following interchange occurred a number of times. The young faculty member would run into the graduate assistants' office, announce he had devised a new theorem (usually relating to Formal Languages), write furiously on the blackboard, and then stand back for our admiration. Most of us were a little overwhelmed by this and struggled to follow what the new theorem was driving at. Ken would lean back in his chair, scrutinize the blackboard for a minute, and say "Isn't $k = 3$ a counter-example?" which, of course, it was. The faculty member would look at the board, scratch his head, say "Oh, yes," and run off to work on his next "theorem". Ken always managed to see through the fog, and share his clearer vision without ever being anything other than thoughtful and gracious.

It's an old cliché to refer to someone as A Gentleman and A Scholar, but that's what Ken was, superbly on both counts. I miss him.

Richard J. Weiland

Ken Sevcik was my friend and my mentor for more than 30 years.

We first met in 1971. Ken joined the Department while I was completing my Ph.D. program. Although he had no formal role on my advisory committee at the time he took an immediate interest in the research I was doing and gave me some excellent advice in a critical area. This advice helped shape the direction of my thesis and, in many ways, the direction of my research career.

Ken's unfailing enthusiasm for his work and for his students was an inspiration to everyone who knew him. Widely regarded as one of Canada's preeminent computer scientists, his contributions as one of the founders of the field of performance analysis were outstanding by any measure. Ken was never one to stand on his laurels, and his research took him in many directions. His students, including Derek Eager, Ed Lazowska, Satish Tripathi and John Zahorjan, went on to make major contributions of their own. These in turn mentored accomplished researchers such as Tom Anderson, Brian Bershad, and Sivarama Dandamudi. The "family tree" is broad and deep, and I know this was a source of considerable pride for Ken.

An accomplished scholar and a true gentleman, Ken earned the respect and admiration of his colleagues around the world. But those of us whose lives he touched in a special way will remember him as a quiet, caring friend – one who was always there when we needed him, whether receptive to new ideas and new approaches, and eager to learn new things. on the ball field or on a conference committee. I'll treasure the memories I have of Ken but I'll sure miss him.

Rick Bunt

We all know that Ken Sevcik was a very special person: big in stature, big in heart. He was sometimes more quiet and polite than lots of us, and at the same time had very definite opinions, talking with humour and conviction about those who are the "good guys" and those who are the "bad guys". He played the game of life with gusto and consideration, loving to win (yes we played hockey) but luckily for some of us, taking care not to run over us with his size and skill. Ken, we miss you and we wish the rest of us could play the game as well as you have.

Ric Holt

First time I met Ken was during my academic interview, here in the department. I entered his office humbled. I was just about to meet one of my academic grandparents. He was after-all the advisor of two of my academic mentors in graduate school. My confidence level suddenly dropped as I stepped into his office.

There are many things I remember about those 30 minutes. First thing I remember is Ken's warmth and grace. It was like meeting an old friend, somebody you have known for a long time. I still feel his kindness and elegance whenever I think of Ken.

I also remember the conversation. Not its content, but its rhythm. It was quick, invigorating, and

youthful. It was different. Different than I expected and in deep contrast with Ken's frail appearance.

I remember his smile as he went to his bookshelf and carefully extracted a green-colored, thick Master's thesis. From 1974. It said "Scheduling multiple resource computer systems". I cannot describe the feeling I had as I was holding it in my hand. It was from 1974, one year older than myself. Finally, I remember Ken's humor. As he was handing the thesis to me, he said with a light flickering in his eyes: "Here's Ed's Master's thesis. I'm sure he wants nobody to see this anymore."

Good-bye Ken.

Stefan Saroiu

"It's nice to be important, but it is more important to be nice". Ken was both, to an extreme degree. But let me start from the beginning: I first met Ken in August 1981, as a fresh phd student, lost, looking for an office in the CS building. A tall, smiling gentleman stopped what he was doing, walked me to the right place, and welcomed me to U. of T. "Nice person", I thought to myself. Little did I know that this person was the first name in the Sevcik-Mitrani theorem, the most important theorem in computer system performance analysis; that he was a super-star teacher; that his "File Structure" course would be the single most influential course that I would ever take. Ken was an intellectual giant: I still remember my knees shaking, every time we had a technical discussion, out of fear that I might say something dumb in front of Ken. Of course, the fear proachable, and non-judgmental. was unjustified: Ken was too much of a gentleman and too much of a teacher to highlight the mistake, although it was clear in my mind that he had spotted it; instead, he would give a compliment, and he would tactfully correct it and suggest a similar, but much better idea.

Ken also had a great sense of humor: he would frequently participate in the hilarious Christmas sketches of the CS department, occasionally improvising his lines, when he would forget the original ones. Either way, the result was big laughs!

I could go on for a long time, but the point is that Ken is a role model for all of us that had the fortune to meet him. Ken, your body left us too soon, but your kind heart still lives in our hearts. Good-bye, Teacher!

Christos Faloutsos

For those looking for a reminder of Ken's unique voice and intellect, we recommend a look at a recent ACM interview which can be found at http://www.acm.org/ubiquity/interviews/v6i3_sevcik.html.

Table 1: Ken Sevcik's Ph.D. students

1. Su, Zaw-Sing, Dynamic Scheduling with Preemption: A Deterministic Approach (1975)
2. Lazowska, Edward D., Characterizing Service Time and Response Time Distributions in Queueing Network Models of Computer Systems (1977)
3. Tripathi, Satish K., On Approximate Solution Techniques for Queueing Network Models of Computer Systems (1979)
4. Chang, Ernie, Decentralized Algorithms in Distributed Systems (1979)
5. Zahorjan, John, The Approximate Solution of Very Large Queueing Network Models (1980)
6. Christodoulakis, Stavros, Estimating Selectivities in Data Bases (1981)
7. Galler, Bruce, Concurrency Control Performance Issues (1982)
8. Eager, Derek, Bounding Algorithms for Queueing Network Models of Computer Systems (1984)
9. Serry, Alaa, An Analytical Approach to Modelling IMS Systems (1984)
10. Walstra, B. Robbe J., Iterative Analysis of Networks of Queues (1985)
11. Casas-Raposo, Ignacio, PROPHET: A Layered Analytical Model for Performance Prediction of Database Systems (1986)
12. Hyslop, William, Performance Prediction of Relational Database Management Systems (1991)
13. Rolia, Jerome, Predicting the Performance of Software Systems (1991)
14. Li, Hui, Software Approaches to Memory Latency Reduction on NUMA Multiprocessors (1995)
15. Harzallah, Karim, Impact of Memory Contention in Large Scale Multiprocessors (1996)
16. Parsons, Eric, Coordinated Allocation of Processors and Memory in Multiprocessor Operating Systems (1997)
17. Zilio, Daniel, Physical Database Design Decision Algorithms and Concurrent Reorganization for Parallel Database Systems (1998)
18. Koudas, Nikolaos, Fast Algorithms for Spatial and Multidimensional Joins (1998)
19. Faloutsos, Michalis, The Greedy, The Naive, and The Optimal Multicast Routing: From Theory to Internet Protocols (1998) [co-supervised with A. Banerjea]
20. Anastasiadis, Stergios, Supporting Variable Bit-Rate Streams in a Scalable Continuous Media Server (2001) [co-supervised with M. Stumm]
21. Nguyen, Uyen Trang, Congestion Control for Multipoint Communications in ATM Networks (2003) [co-supervised with I. Katzela]
22. Wang, Hai, Concise and Accurate Data Summaries for Fast Approximate Query Answering (2004)

Table 2: Ken Sevcik's M.Sc. students

1. Tran, Tuan Khan, The Response Time Distributions of Some Computer Scheduling Disciplines (1972)
2. Lester, Eleanor A., The Investigation of Service Time Distributions (1973)
3. Lazowska, Edward D., Scheduling Multiple Resource Computer Systems (1974)
4. Lam, Lettice Hung, Comparison of Three Logical Views of Data in Hospital Information Systems (1974)
5. Schumacher, Helmut, The Synthesis of Optimal Decision Trees from Decision Tables (1974)
6. Tripathi, Satish, K., Approximate Models of Multi-Programmed Computer Systems with Paging (1975)
7. Lum, Wade C., Data Collection, Reduction and Analysis of Computer System Measurement (1975)
8. Cheung, Stephen Chun-Lap, A Computer Method for Cyclical Scheduling of Shift Workers (1975)
9. Shen, Helen, The Effect of Interarrival Times in Scheduling (1975)
10. Spirk, Franz, An Analysis of Different Deadlock Control Schemes (1975)
11. Lin, Hing-Lung, Closed Queueing Network Models: Computational Algorithms with Application to Computer System Performance Evaluation (1976)

12. Dubien, Ronald J., Extending a Commercial Database System to Provide Data Independence (1977)
13. Kienzle, Martin G., Measurements of Computer Systems for Queueing Network Models (1977)
14. Kidd, Robert John, A Common Model for Multi-Key Access Performance Evaluation (1978)
15. Lai, Chi Sang, Patient Identification and Record Linkage (1979)
16. Au, Matthew, Some Properties of Queueing Network Models (1981)
17. Bell, Barbara, Database System Performance Prediction: The First Two Levels of a Multi-level Modelling Framework (1981)
18. Casas-Raposo, Ignacio, Analytic Modelling of Database Systems: The Design of a System 2000 Performance Predictor (1981)
19. Eager, Derek, Robust Concurrency Control in Distributed Databases (1981)
20. Bobrowski, Christoph, The Principle of Maximum Entropy in Some Computer System Modelling Problems (1983)
21. Vopalensky, Jan, The Scheduling Problem in Distributed Computer Systems: Its Definition and Partial Solution (1983)
22. Martin, Vickie, SPIRAL ONE - Dynamic Hashing Scheme (1983)
23. Aldwinckle, John, A Scheduler for UNIX (1984)
24. Gelblum, Michael, Some Results of the File Placement Problem (1984)
25. Rolia, Jerome, Performance Estimates for Multi-tasking Software Systems (1987)
26. Coatta, Terry, Queueing Networks With State-Dependent Service Rates (1987)
27. Clark, Graeme, Performance Properties of the FDDI Token Ring (1987)
28. MacLean, Richard, Performance Characteristics of the FDDI Token Ring Priority Mechanism (1989)
29. Lee, Jei-min, Concurrency Control Through Hierarchical Synchronization (1989)
30. LeBel, Edgar, Sparse Undistributed Memory: A Local Storage Model for Pattern Matching (1990)
31. Cupit, Brian, Parse Tree Based Revision Control and Program File Compression (1990)
32. Verma, Raj, A Metric Approach to Isolated Word Recognition (1991)
33. Srikantiah, Nandini, Processor Scheduling in Hierarchical NUMA Multiprocessors (1991)
34. Bacque, Ben, SUPERMON: Flexible Hardware for Performance Monitoring (1991)
[co-supervisor with K. C. Smith]
35. Marwood, Simon, Performance Comparison of Concurrency Control Techniques in Partitioned Systems (1993)
36. Wu, Chee Shong, Processor Scheduling in Multiprogrammed Shared Memory Computer Systems (1993)
37. Larson, Johan, Domain-Partitioned Parallel Sort-Merge Join (1995)
38. Wang, Corina, An Adaptive Rendering and Display Model for Networked Applications (1996) [co-supervisor with M. van de Panne]
39. Tam, Anita, Performance Prediction for Parallel Applications (1996)
40. Anastasiadis, Stergios, Parallel Application Scheduling on Networks of Workstations (1996)
41. Soreanu, Gabriel, A Comparative Study of Interconnection Networks in Large-Scale Multiprocessors (1996)
42. Wang, Hai, Approximate MVA Algorithms for Solving Queueing Network Models (1997)
43. Gibbons, Richard, A Historical Application Profiler for Use by Parallel Schedulers (1997)
44. Kennedy, Brett, Declustering Spatial Data for Range and Join Queries (1999)
45. Piegaze, Peeter, Incorporating Time into Computer Cartography: A Digital Atlas of History (2001)
46. Lu, Jingjing, Fast Algorithms for Distance-Based Spatial Queries (2001)
47. Garcia Arellano, Christian, Quantization Techniques for Similarity Search in High-Dimensional Data Spaces (2002)
48. Tjioe, Wei, Finding Structure in Data (2003)
49. Pham, Hang, Accurate Two-Dimensional Histograms for Fast Approximate Answers To Queries on Real Data (2004)
50. Dong, Junfeng, Indexing High-Dimensional Data for Main Memory (2004)