


Remembering Ken Kennedy



With Ken Kennedy's passing in February 2007 we lost a great intellectual pillar in our field and a truly wonderful human being. Ken's work in compiler technology and parallel tools was well known and served as a foundation for successive generations of computational scientists to use advanced computer systems; but more than that he was a catalyst in building the community of computer scientist and computational scientists alike. He was a devoted educator and mentor. He was a Ph.D. adviser to 38 students, and he mentored countless others.

Ken Kennedy joined the Rice faculty in 1971, and in 1980 he became a professor of computer science. He was instrumental in forming several new research centers at the school, including the department of computer science in 1984, the Computer and Information Technology Institute in 1986, the Center for Research on Parallel Computation (CRPC) in 1989. In 2000, the CRPC became the Center for High Performance Software Research (HiPerSoft), which Kennedy directed from its inception. HiPerSoft is the Rice administrative home for several multi-institution projects, including the Virtual Grid Application Development Software (GrADS) Project, an NSF-sponsored effort involving seven universities, and the Los Alamos Computer Science Institute (LACSI), a consortium of five universities and the Los Alamos National Laboratory. In 1997, Ken was chosen to co-chair the President's Information Technology Advisory Committee (PITAC), a congressionally mandated committee charged with advising the president, Congress and other federal agencies on advanced information technology. The panel's 1999 report urged U.S. leaders to increase spending for computing research by more than \$1 billion, and it served as a catalyst for increased IT research support from numerous federal agencies.

Ken received many awards during his 36-year career, including the 1967 Hugh Cameron Award for Service to Rice; the 1995 W.W. McDowell Award—the IEEE Computer Society's highest research award—for his high-performance computer systems research; the 1999 Lifetime Achievement Award from the Association of Computing Machinery's Special Interest Group on Programming Languages; and the 2002 Distinguished Alumnus Award from the Association of Rice Alumni. In 2002, he was promoted to University Professor, only the fourth person in the history of the school to receive such an honor.

Ken Kennedy authored more than 200 papers and wrote two books. He was also a Fellow of the ACM, the American Academy of Arts and Sciences, and a member of the National Academy of Engineering.

Ken was a special person in many ways. His work and our memory of him will be with us for a long time and his presence will be sorely missed. Below are some remembrances of Ken's life and style from a few of his many friends, colleagues, and collaborators.

Jack Dongarra
Knoxville, Tennessee
May 2007

Fran Berman

University of California, San Diego

I first met Ken as a beginning Assistant Professor in Computer Science and have been fortunate enough to become his friend and collaborator on multiple projects through the years. Ken's fundamental contributions to Computer Science are well known. National Academy member and international leader, he is recognized as one of the premier Computer Scientists in the world.

Ken, his students, and collaborators provided intellectual leadership in the areas of compiler design and parallel programming environments for decades. In the 90's, as grid computing emerged as an important platform for computing, Ken applied his extraordinary curiosity and passion for scholarship to the problem of designing compilers and programming environments for grid platforms. The GrADS and VGrADS projects, led by Ken, brought together an A-list group of researchers, and it was through Ken's leadership, scholarship, and perseverance that this group tackled and began to address some of the most challenging problems in grid software.

Ken focused on quality and inclusiveness in all his activities, and created an environment in which colleagues, students, and other technical professionals could thrive, and were fully integrated as leaders and team members into projects and programs. Ken coupled his exceptional research contributions with a deep commitment to education, mentorship, and the broadening of participation within the Computer Science community. Through his extensive community leadership activities, he kept quality and inclusiveness on the radar, promoting these themes as co-chair of the Presidential Advisory Committee on High Performance Computing and Communications and the Next Generation Internet, as Director of the Center for Research in Parallel Computation (CRPC), within the PACI partnerships, and as advisor or participant of virtually every important

CS research policy-making organization from the NSF CISE Advisory Committee, to the National Research Council, to the US Congress, etc.

Ken was instrumental in developing one of the first Computational Science programs in the country at Rice, and always made sure that education and training were well-integrated with research. He spent countless hours with new generations of Computer Scientists feeling that this was an integral part of his role and responsibility as a professional Computer Scientist.

I miss Ken greatly. His personal commitment to supporting, encouraging, nurturing, and promoting excellence when he saw it - across a wide spectrum of people, programs, and environments, and at all different career levels - set a community and a leadership example for all of us. When he died, I told a friend that we would all have to try even harder to achieve the best and give back to the community to fill the tremendous gap that he left.

Vint Cerf

Google

My first and most vivid interactions with Ken Kennedy were in connection with President Clinton's Information Technology Advisory Committee. Ken was the co-chair, together with Bill Joy, during the 1997-1999 period and he served in that role with extraordinary panache and effectiveness. He was ever the gentleman but you knew when he wasn't happy with progress. His skillful guidance and patient energy led this group to make substantial inputs into the country's information technology research program. More than anything, I believe Ken wanted people to understand things deeply. This philosophical principle permeated his work with PITAC and my impression is that this extended to his academic work at Rice University. Although I did not have the privilege of working with him in that setting, all the reports I have heard underscore this theme. His commitment to the computing community, especially in high performance computing and programming, propelled Rice University to the forefront of this field. He will be deeply missed but his impact in professional and personal capacities will not diminish. I feel deeply privileged to have known and worked with Ken, albeit for far too brief a time.

Andrew A. Chien

VP Research, Intel

My first exposure to Ken came thru his seminal work in program analysis and transformation. As a graduate student at MIT in the 1980's working on parallel architecture and programming, I was an avid reader of Ken's work as it was one of the intellectual lighthouses for understanding program analysis for vectorization and later parallelization. Ken's work was visionary, joined the areas of program analysis, program transformation, vector and distributed memory machines, and led to data parallel languages such as Fortran D, High Performance Fortran, and a whole class of related languages. Over four decades, Ken was an intellectual force in the HPC research community.

I first met Ken at NCSA in the early 1990's, as the national supercomputing centers transitioned from vector to parallel architectures. At that time, I was a junior professor at the University of Illinois and Larry Smarr had asked me to lecture in NCSA/CRPC's education programs on parallel architecture, and Ken was lecturing on parallel programming systems. What I recall from those meetings is an immediate feeling of kinship and shared mission – in driving the frontiers of parallel computing. Ken was approachable, open, thoughtful, and friendly. What grew out of that was a relationship where over the years, Ken became an intellectual co-traveler, research collaborator, and mentor. That 20 year relationship has spanned technical areas of parallel programming and program analysis, parallel input/output, object-oriented programming, and most recently grid computing. We have gone “toe-to-toe” on difficult technical issues, and both emerged the better for it. We have worked together to build large complex systems, which advanced the state of the art and represented breakthroughs for the research community. And Ken has provided invaluable career guidance and advice for which I will always be in his debt. I feel privileged to have been able to work with Ken, and have benefited immeasurably from it.

As significant as Ken's research contributions are, Ken's community leadership was perhaps even more remarkable, in the formation of the Center for Research in Parallel Computing, HPCC Initiative, and of course later the PITAC. A tangible example of how he identified key problems for the field, and moved us forward is his catalytic role in the development of the Message Passing Interface (MPI), which unified a dozen vendor specific messaging interfaces. His energy and intellectual leadership was so great that he easily held the respect of his peers while moving the community forward in major ways. As I began to take on leadership and service roles, my appreciation for Ken, and what Ken achieved at Rice and in the community expanded greatly. Exposure to how difficult these things “really” are to achieve, made not only his accomplishments, but the way he carried them off – with kindness, charm, and unbelievable grace under adversity and pressure. In this, I believe Ken will always be a role model for us all.

Keith Cooper

Rice University

Ken Kennedy was a close friend, a long-time colleague, and my mentor. We started working together in about 1980, when he agreed to advise my Ph.D. Over the years, we wrote papers together, hatched political plots to take over the University (or, at least a major share of its budget), and worked together on countless matters. In reflecting on Ken's influence on me, on his students, on Rice, and on the larger community, several of Ken's traits stand out.

- Ken had, in his research, an overriding concern with practicality. He wanted the work to find application in real systems that helped solve important problems. In a purely technical sense, this trait surfaced as a deep concern for the complexity of any algorithm that he considered. For a technique to find widespread adoption, he believed, it must have small constants and a low asymptotic complexity; the former for speed on small problems and the latter for scaling.

One of his first significant insights in the area of dependence analysis was that we could compute transitive dependence using depth-first search, which has a cost that grows linearly in the number of dependences; prior practice used Warshall's algorithm, which has complexity that is cubic in the number of dependences. The underlying goal was to reduce the cost of vectorization so that it would become a practical technology. The effect, however, was to make his students acutely aware of algorithmic issues—an effect seen in much of their subsequent work to this day.

- Ken was at his happiest when completely engrossed in a challenging problem; on a plane, in a meeting, and at the beach, his mind would explore approaches, solutions, and their practical ramifications. When he made a breakthrough, he would rush to explain it to his nearest colleagues with a combination of child-like joy and a deep sense of accomplishment.

I recall one brief period where he became involved in a discussion with a horse breeder on the subject of inherited genetic flaws. He spent countless hours mastering the material, checking the breeder's records and insights. The problem was far outside his normal area of expertise; nonetheless, he attacked it with the same enthusiasm and determination that he brought to bear on problems from program analysis to federal funding for research.

- Ken was an intensely social man, who collected friends, collaborators, and students. He built an academic department, an NSF Science and Technology Center, and numerous collaborative ventures such as the GrADS, VGrADS, and LACSI projects. In each, he found new friends and collaborators. His sense of humor, and his deep belief in the problem solving ability of his colleagues elevated meetings and teleconferences from functional gatherings into meetings of friends. Many of us in the computational community experienced this side of Ken.

With Ken, however, neither age nor status governed the nature of his friendship. He was as serious talking about the prospects of his beloved Houston Astros with a middle-school boy as he was when discussing the PITAC report with Congress. He prized the conversations that he had with his students, both undergraduate and graduate. He would focus the same intensity on any problem and on any subject.

In the aftermath of his death, I have corresponded with hundreds of his friends and colleagues and heard countless stories of his warmth, encouragement, and concern. His laugh, loud and strong, echoes through the minds of many in our community. While I fondly remember his humor, I am more struck by his lifelong optimism. It was perhaps

best embodied in his love of gadgets. Ken had an absolute faith that technology (particularly if made by Apple) could cure any problem. My deep hope is that we can keep alive that unique combination of curiosity, optimism, playful fun, commitment to excellence, and faith in one another that so marked his career.

Jack Dongarra

University of Tennessee and Oak Ridge National Laboratory

My first memory of meeting Ken was at IBM Yorktown Heights back in 1980. I was spending part of the year working with George Paul and Alan Karp and others on the IBM vector processor. One day Ken was visiting and gave a talk about his compiler technology. John Cocke, Fran Allen and others were there to hear Ken's seminar. Someone asked Ken a question and he stopped and started drawing on the white board filling it up with details to explain the point. After Ken finished he realized that he was writing with a permanent marker on the white board. This caused John Cocke to remark in his distinctive southern drawl that he would be remembering that explanation for a long time. At which point Ken broke out into his loud infectious laughter. I came to know and hear that characteristic laugh of Ken's many times in our relationship over the next 27 years.

My serious interactions with Ken began in earnest at the forming of CRPC. I knew of him and interacted with him before this but I really came to know Ken during that 10 year effort plus the various projects that evolved until his death.

Ken had a commitment and loyalty to his friends that ran deep. He was skilled at building group consensus and that showed in his ability to bring together the computing community. His intellectual leadership and vision with PITAC and other efforts are a model for us to follow. His ability to bring together groups of researchers, prepare proposals and research papers is an experience I try to emulate in my own efforts.

After Ken's death I was at Rice University for a project meeting. As we assembled in the room where a few months before Ken led us in discussion, we couldn't quite get the meeting to begin. It was as if we were waiting for Ken to come in with his Starbucks latté in one hand and his Mac computer in the other to get things going.

Ken continues to be an inspiration and his strong presence will be greatly missed. For me he was a mentor, colleague, and friend. I miss my friend.

Kelvin Droegemeier

University of Oklahoma

Because my graduate advisor in atmospheric science was a computer scientist, I became familiar with Ken Kennedy in the early 1980s as high performance computing, defined by the CDC 7600 and Cray-1, was in its infancy. As a new faculty member at the University of Oklahoma, I had the good fortune of working on an NSF Science and Technology Center proposal while, unknown to me, Ken was doing the same at Rice University. When we both won, to our amazement, I immediately contacted Ken to see if we could leverage one another's activities. I also wanted Ken's expertise on our external advisory panel given that the vision of our center was to take computer weather prediction down to the scale of individual thunderstorms – a computationally immense problem for which we were developing the world's first massively parallel storm-scale model. Because so many different architecture paradigms were being pursued at that time, we needed Ken's genius to help us straddle the fence in designing our code so that we could optimize it once the way forward was more clear. Needless to say, his guidance was superb and our code benefits from it to this day. His work in advanced languages, compilers and most recently, virtual grid applications environments was profoundly creative, practicable, and of broad impact.

Working with Ken was exhilarating, exhausting and humbling! His quick mind kept us on our toes, requiring that we did our homework. One always knew that getting a piece of Ken's time was a privilege, and that the payoff would be far greater than imaginable. I'll always remember two important facts about Ken. First, those working for him were intensely loyal – John Mellor-Crummey comes to mind as a great example – and many remained with him from the start of their careers until he left us. In addition, despite his tremendous stature, Ken always took a personal interest in one's ideas and aspirations. He was a great encourager and had a tremendous sense of discernment.

Many work hard to leave a legacy but for Ken, it came with ease. His influence on people, on technologies that improve and save lives, and on institutions will extend for decades. He made us think, he helped us do more than we thought we could, and he left the world a better place.

Thom H. Dunning, Jr., Danny Powell, and Rob Pennington

National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign

Professor Ken Kennedy had many talents. Everyone in scientific computing, whether you are a computer scientist, a computational scientist, or an applied mathematician, is familiar with Professor's Kennedy's technical contributions – in parallel computing, programming environments, and languages and compilers. In this tribute to the memory of Professor Kennedy, we will stress another aspect of Ken's talents – his ability to pull people together and establish a community that draws together the leaders of a discipline. He first showed this talent when he helped found the Department of Computer Science at Rice University. But fortunately for the scientific computing community, his interests extended far beyond Rice.

Back in the 1980's, "interdisciplinary and inter-institutional collaboration" was the buzzword, but it was actually a rarity in practice. Ken led several large, group projects where he was able to bring together many of the key leaders in computer and computational science (Paul Messina, Jack Dongarra, Geoffrey Fox, Andy White, Herb Keller, Rick Stevens, John Dennis, Richard Tapia, John Reynders, Fran Berman, Dan Meiron, Mary Wheeler, Dan Reed, Mani Chandy, Andrew Chien, Carl Kesselman, Ian Foster, Rich Wolski, Lennart Johnsson, and others) in true, coordinated team efforts. The work of these groups was acclaimed nationally and internationally and included projects such as the Center for Research on Parallel Computing, the Los Alamos Computer Science Institute, and the Center for Grid Application Development Software project. The partners in these projects planned and worked together on a scale that is still a rarity today. What drove the team integration was Ken's incessant desire to create something of real benefit to others, to be demonstrated by its use by others and to be found still in use (or built upon) long after the original team project was over. This practical bent was what endeared Ken to computational scientists.

What cemented these teams together was Ken's ability to listen to others, incorporate their ideas into a defined team goal, allocate resources appropriately, and give credit to the partners for their contributions. It was rarely a Rice project with other partners contributing – it was a team project with everyone having a critical role, and pulling together to make the whole result much greater than the sum of the individual parts. What we find, in retrospect, is the establishment of strong relationships between key, individual players across the country in computer and computational science, that were developed as part of Ken's team building abilities, and that have lasted to this day – driving many of the advances in scientific computing that are now coming to fruition. Ken will be sorely missed.

Bill Feiereisen

Los Alamos National Laboratory

In 1998 many of us gathered at Argonne National Laboratory to discuss the Grid. At that time Grid

computing was coalescing from individual research projects into a community for distributed scientific computing. My interest in distributed computing came from changes that I was planning for supercomputing at NASA Ames where I worked at the time. Having come from a largely computational community I was a newcomer.

By chance I sat behind a tall gentleman with a mustache, who spoke several times during the day. I was impressed by the calm, measured and elegant arguments about the importance of Grid computing to the country, and to the world. This was Ken Kennedy who I have come to know as a colleague and friend in the intervening years. Ken and others made me feel welcome as part of this community and helped me with the arguments and advocacy that were critical for success in founding the Information Power Grid, NASA's first distributed computing attempt at scale.

I moved to Los Alamos National Laboratory in 2002, where one of the pleasures was to jointly manage the Los Alamos Computer Science Institute (LACSI) along with Ken. LACSI was focused on the computer science of high performance computing in support of the nuclear weapons program. Ken had assembled a stellar team of collaborators from multiple schools, all interested in the success of these "mission" computations and the CS technologies behind them. The scientific areas ranged from algorithms and libraries, through performance profiling and tuning to new concepts in parallel languages and compilers that was Ken's life-long mission. LACSI was a significant seed that contributed to laboratory successes such as "Science Appliance," and our performance modeling capability.

The Computer Science Institute was not Ken's day job, but he put considerable effort into his technical work for LACSI while still managing to represent the work of all his colleagues and help us plan for future CS technologies in the weapons program. The success of LACSI was in no small measure due to Ken's efforts and his talents. He was a top level computer scientist and

yet at the same time he could formulate and articulate “the big picture.” I have remarked many times in the last five years that LACSI was the closest and most productive collaboration between University and Lab colleagues that I have ever seen.

Ken always had time for the people around him and truly cared for them. He served on many committees and was motivated by the good that he could do for people and their institutions. Ken was the chair and a member of our “Division Review Committee” who regularly helped guide us through the labyrinth of life at a national laboratory. Although they convened annually with us and our management to make recommendations, Ken was here every couple of months to do the technical work of LACSI. He would meet with Lab management about the latest Lab controversies and provide a steady hand. He became a mentor to me, and many others here at Los Alamos and throughout the community.

Ken was a great communicator and would astonish me with almost instantaneous replies to email at rather odd hours. We had known that he was sick and fighting hard for his health, but in late January he stopped replying to email. Calls to Houston discovered that he had gone into the hospital again. This had happened several times before, but I was worried. I was sitting at my desk in Los Alamos when Linda Torczon called and told me of Ken’s death earlier that afternoon. I had to shut the door for a while, having lost a role model and a good friend.

Ian Foster

Argonne National Laboratory and

The University of Chicago

I first met Ken in 1989, when by a fortunate series of events I became involved in the Center for Research on Parallel Computing (CRPC). Over the next several years CRPC was a big part of my professional life. Looking

back, what I find most remarkable is how unremarkable it seemed to me at the time.

As a young postdoc just arrived in the US, I assumed that projects of this sort were the norm and that they always involved such outstanding people, were run as smoothly and collaboratively, were as productive scientifically, were as effective in producing software and influencing practice, had as strong an outreach and diversity program, and were as successful in sustained these accomplishments over many years. I now appreciate that these characteristics of CRPC were far from inevitable--were, in fact, highly unusual--and that the creation, longevity, and broad success of CRPC are all a testament to Ken's exceptional leadership skills.

Later, I watched as Ken led the President's Information Technology Advisory Committee (PITAC) that spurred the creation of (among other things) NSF's Information Technology Research program. That program supported a tremendous amount of good work across the US--another example of how Ken did not simply excel in research, but also worked to shape the computer science research landscape. Certainly for me, and I think for many others, Ken's example of what it means to be an engaged and successful researcher in applied computer science was incredibly influential. His many students and collaborators ensure that his legacy continues.

Geoffrey Fox

Indiana University

Ken Kennedy changed my career when we met in 1987 at an ICS meeting in Athens and I discussed the application requirement and he the compiler opportunities for parallel computing. These interactions grew into my participation over the next decade in CRPC, the Center for Research in Parallel Computation. Ken taught me a lot over this time giving me some credentials as a computer scientist so I was no longer just a physicist. I also should have learnt from him to keep calm and turn the other cheek and smile in aggravating situations.

Ken excelled as a scientist, leader and colleague. His skills in the compiler field are documented in his books and papers, the colleagues and students he trained and in new ideas. I am most familiar with his work on FortranD and HPF High Performance Fortran and although the latter was not broadly successful that more reflects the resources devoted to it rather than its merit. In fact the key ideas of HPF can be seen in the new HPCS languages from Darpa and I expect variants of these ideas to be very promising. Ken worked hard to nurture his colleagues and I and others benefited from his caring nature. His leadership in the field was first seen by me in his inspirational governance in CRPC and then on an international scale in PITAC. I had just started working with Ken again on

multicore systems over the last year and in spite of his illness, he approached these new challenges with vigor and enthusiasm. I have lost a friend and the field has lost a great scientist.

Chris Johnson

University of Utah

My first meeting with Ken came about in a way that emphasized how he stood out – above – in so many ways, and so I remember it vividly. In the mid 90s we were interviewing faculty candidates in systems. The interviews included several compiler candidates. Some of the presentations were given using transparencies created from Latex, and commenced with slide after slide of do loop optimizations. As you can imagine, many of the presentations were less than exciting. Then, Ken Kennedy came to give an invited talk. What a difference! Not only did Ken use multiple colors in his presentation (perhaps a first for a compiler talk at the time), he gave his work a meaningful context, providing a vision for high performance computing research and for how high performance compilers made up an integral part of such a vision. By the end of the presentation, I was starting to feel that high-performance Fortran was going to save the day for HPC.

Of course, I tell this story somewhat in jest. Still, I think it exemplifies Ken's many amazing qualities. He was a visionary in the field of high-performance computing. He was a great communicator and motivator. And few can compare with Ken's accomplishments both with respect to research leadership and national information technology policy leadership. Finally, Ken was a warm and generous person. His loss will be felt throughout the computing community and certainly by me personally.

Carl Kesselman

ISI

In writing about Ken Kennedy's impact, it is tempting to focus on his significant technical contributions. As these are well documented in his papers and the extensive collection of awards and recognitions, I would like to take this opportunity to focus on a facet of Ken's impact that may not be apparent to those who did not have the privilege to know him personally.

I finished my Ph.D. in 1991, a time when there was great excitement and progress in high-performance computing. Parallel processing at a massive

scale was showing great promise; computers with upwards of five hundred processors were being built. If only we could figure out how to program these awkward beasts. A problem of this magnitude required a comprehensive coordinated approach, and in recognition of the National Science Foundation funded the Center for Research of Parallel Computation (CRPC). As the CRPC PI, Ken created an organization that brought together many of the leading thinkers in parallel programming tools, languages and applications with a mission to significantly advance the state of the art and practice of parallel programming.

Somehow, I was lucky enough to participate in this wonderful organization and for the next ten years, I had the chance to collaborate with Ken and the stellar collection of researchers who participated in CRPC.

Calling CRPC an organization or a center does it a great injustice. In shaping CRPC, Ken created a true scientific community, which was as important to him as the technology. I look back with great fondness to the annual treks to Houston in the hot humid days of August to attend the annual project meetings and a chance to collaborate not only with my peers but the leadership of CRPC and to sit down at the end of the day with them, have a beer and perhaps some Texas BBQ.

Ken was incredibly generous with his time, his ideas and with credit. During the duration of the CRPC, Ian Foster and I had introduced ideas about Grid computing into the research agenda of the CRPC. Ken encouraged and embraced this work, bringing Ian and myself into leadership roles in the center. Without question, collaborating with Ken and being part of the CRPC and the community that Ken created within the CRPC played a significant role in shaping my career.

When the funding for CRPC drew to an end, it was important to Ken to find a way to find funding to not only keep elements of the research program continuing, but to keep the community together. This was typical of Ken's world view. He was committed to his collaborators, committed to his colleagues, committed to his University and committed to the computer science community as a whole. Because of this, Ken's impact goes far beyond people who have read his papers, used his software, or taken one of his classes. Ken's sense of scholarship and community were unique, and because of this, he will be greatly missed.

Rusty Lusk

Argonne National Laboratory

Ken is well known for his leadership role in compilation research and the creation of the HPF parallel Fortran language. Less well known is the key role he played in creating the MPI standard. It was through this latter activity that he had the greatest influence on my own life and career.

In April of 1992 Ken organized the Workshop on Standardization of Message-Passing Libraries, supported by his own Center for Research in Parallel Computation and held in Williamsburg, Virginia. All the extant message-passing libraries were represented, from hardware manufacturers, independent software vendors, and research groups. Near the end of the meeting, after listening to the presentations, I thought to myself, "There is too much diversity here; it is just too early to standardize message passing." At the very end, Ken said with little preamble, "OK, who wants to participate in an effort to standardize the message passing interface?" It was a brilliant stroke. Obviously no one wanted to be left out, and I put up my hand along with everyone else. Thus MPI was born.

Although Ken did not attend the MPI Forum meetings, his influence was strong. The HPF Forum, which he led, had just concluded its efforts successfully, and the MPI Forum confidently adopted Ken's rules on membership, committee structure, voting processes, and so forth. We even decided to meet in the same hotel where the HPF Forum had met. All of these decisions helped the MPI process succeed, and most of my own work since that time has been related to MPI.

I most recently worked with Ken on a successful SciDAC proposal for an Institute to hold a series of summer workshops to further all aspects of high performance computing. Typically for Ken, who was still cheerfully working on this project with us the last week of his life, its purpose was not just to advance his own work, but everyone's.

John Mellor-Crummey

Rice University

I joined Rice University in 1989 as a fresh Ph.D. to work with Ken Kennedy as a member of the Center for Research on Parallel Computation (CRPC) - a multi-institutional NSF Science and Technology Center that Ken led. Ken's vision for the center was to "make parallel computing truly usable." This was a broad, compelling, and ambitious vision. As a young researcher, working with Ken in the CRPC was a unique privilege. The CRPC's broad research agenda enabled me to pursue a range of interests in my early career, including synchronization algorithms for shared memory programs and detecting data races in parallel programs. The early years of the CRPC were an idyllic time for me: the center had long-term funding and it enabled me to focus almost exclusively on research. As a young researcher, the CRPC introduced me to many of the national leaders in the area of high performance computing. Many of these individuals are among my active collaborators today.

As part of the CRPC, Ken led the project in Fortran Parallel Programming Systems. The project's early work focused on the ParaScope Parallel Programming environment for shared-memory parallel systems. Work soon shifted to explore programming systems for distributed memory platforms as they became the dominant scalable parallel architecture. At Rice, work in this area began with the Fortran D research project; other similar efforts

were underway at other institutions and companies. Ken saw the opportunity to organize these independent efforts and in 1992 established the High Performance Fortran Forum. Under Ken's leadership, this group produced the High Performance Fortran language standard in 1993. The standard was only the first piece of the puzzle; the work of compiling a broad spectrum of HPF programs into efficient code had only begun. Through collaboration with Ken, I became deeply involved in compiler technology for parallel systems. Over the better part of the next decade, my research included work on dHPF - a second generation compiler for HPF. Even today, part of my current research portfolio includes compiler technology for Co-array Fortran - a single-program multiple data language for parallel programming on distributed memory architectures. When I joined Rice to work with Ken in 1989, I hadn't expected that nearly twenty years in the future, I would still be involved in compiler technology for parallel programming models based on Fortran!

Many of those who worked with him regularly will remember Ken for his love of sharing funny stories in meetings and how he punctuated these stories with his distinctive laugh.

He will be greatly missed.

Paul Messina

From my perspective, Ken's activities had significant impacts on the following areas of computer science and computational science:

Compilers and programming languages. Ken's work on Fortran compilers for generating efficient code for scalar, vector, and parallel computer architectures not only yielded many important advances in the fields but also legitimized work in those fields. Ken undertook research in those areas despite the disdain many computer scientists had (and still have) for Fortran and the comparatively limited funding available for research on compilers. By recognizing that most scientific applications have an extreme need for runtime efficiency and developing ways to achieve it, his work helped demonstrate that research on compiler technology can provide enormous leverage; an excellent optimizing compiler can yield much higher performance on a given platform, thus reducing how much money is spent on hardware.

Multidisciplinary research. Ken led and engaged in true multidisciplinary research, working closely with scientists who were the end users of the technology, as well as fellow computer scientists and mathematicians. He did so without compromising his computer science research output and provided a stellar example of the benefits of such research for all participants. I worked closely with Ken for over ten years as a participant in the NSF Science and Technology Center he directed, the Center for Research on Parallel Computing (CRPC). It was in the early years of CRPC that I observed how skillfully Ken

brought together a large and diverse collection of independent researchers from various fields and molded them into a cohesive group with common goals and values. That example guided many of my subsequent projects, such as the Intel Touchstone Delta consortium, the Scalable I/O Initiative, and the National Virtual Observatory. I also applied the lessons learned in CRPC in my role of Chief Architect for the NPACI and in leading the DOE-NNSA Advanced Scientific Computing Initiative (ASCI).

Experimental computer science. Ken believed in using real machines to test the research results with real applications in realistic computing environments. As Director of CRPC he ensured that CRPC researchers had access to a spectrum of parallel computer architectures, in many cases the largest configurations available in the world. He was also an early and vigorous supporter of the Intel Touchstone Delta project that I led in the early 1990s and helped make that project a success.

Ken was able to make an impact in these areas, all of which were unpopular in the computer science community; he had the courage to engage in them and earned stature and great respect from both the computer science and physical and mathematical science communities.

Ken also had direct impact on my own activities.

He was a key participant in a weekend-long brainstorming session to identify the focus of a follow-on project to the Delta project. Ken was among those who noted that input-output issues for parallel computers badly needed attention; that planning session resulted in the Scalable Input-Output Initiative, in which he was a major participant as well.

In the early days of CRPC, Ken asked me to terminate a research project I was carrying out with CRPC funding, in order to generate funds for a more worthwhile project. Ken's direct yet sensitive style when he met with me convinced me his judgment was right, the other project was more deserving of funding, and I was glad to end my project. That's style!

Juan Meza

Lawrence Berkeley National Laboratory

It was 30 years ago and I had just walked into the Advanced Programming course. Ken taught the course and he had just gone around the room and asked us why we were in the class and what we expected to learn from it. Most answers were of the form, "I want to be a better programmer", or "I want to learn more about computers". What I most remember about the class though was Ken's emphasis on "programming style". I recall

him saying that some of the best programmers that he knew were former English majors. As you might imagine that came as quite a shock to most of us students. As most of us (perhaps all) were science and engineering majors, we would have taken an English course only if it was absolutely required. Still Ken persisted and explained that programming was about organizing your thoughts and structuring what you wanted to do – exactly the same characteristics that you needed to be able to write well.

Looking back over the years I now realize that he was absolutely right. His emphasis on “style” and on first figuring out what you wanted to do has stayed with me and guided me on many a programming project. I’ve also witnessed the reverse, programmers that started coding almost immediately and then floundered when it came time to debug and test their codes. It comes as no surprise to me these days to hear that so many software projects come in past their deadlines and go over budget. If only they had had the advantage of taking a course from Ken!

I also remember many discussions with Ken in and out of class. In those discussions Ken was always thinking about what would be possible in computing 5, 10 years down the road. But to say he was a visionary isn’t enough. He had an incredible knack for generating excitement and getting you to believe that not only was it possible but that it would be *fun*. I’m reminded of the old saying, “Dream no small dreams for they have no power to move the hearts of men.” Ken had that ability to dream the big dream and we’ll all miss him.

Dan Reed

Renaissance Computing Institute

Of all the character traits I prized in my friend Ken Kennedy, three stand out above all the others. The first was his ability to listen quietly to a complex technical debate, with strong voices and opinions on all sides, and then ask the critical question, the one that made everyone stop and think deeply. Inevitably, such questions identified either the flaw in an approach or, more often, suggested a better way that combined the best ideas in a more general approach.

These were the transformative questions that shaped the research trajectories of careers, of programs and of agencies.

The second was Ken's unwavering commitment to a long-term, integrated computing and computational science research program. His tenure as co-chair of the President's Information Technology Advisory Committee (PITAC) reshaped the national computing and computational science research portfolio, elevated computing in the federal R&D agenda and secured substantial new funding. I know how difficult this was, both from watching the 1999 PITAC report unfold and from my own (later) experiences on PITAC and PCAST.

The third and most important was Ken's insatiable curiosity and love of research. The morning coffees, the afternoon discussions and the long dinners (always with wine) were filled with questions and possibilities. To spend a day with Ken was to emerge energized and reminded of why research matters, why each of us became computer or computational scientists.

Larry Smarr

University of California, San Diego

Ken Kennedy left a major impact on both computer science research and on computational science. Ken worked closely with me in setting up the National Computational Science Alliance, where he served as team leader for the Alliance Enabling Technologies Parallel Computing Team, and as a member of the Alliance Executive Committee. He hosted the critical NSF site visit at Rice in 1997. All of us there remember the "Spirit of Houston," in which Ken's leadership helped pull the broad set of researchers together into what became the successful Alliance program. Ken was specifically quite involved in developing parallel compilers for Steve Wallach's Convex mini-supercomputer line, which NCSA made available to the national computational science community. I remember Steve saying about Convex—"I get my compilers from Ken at Rice and I get my parallel applications from Larry at NCSA."

Shortly after Ken and I were asked to become members of the President's Information Technology Advisory Committee (PITAC), with Ken serving as co-chair. The committee provided guidance to the Clinton administration's efforts to accelerate development and adoption of information technologies. However, it was Ken's skillful leadership of PITAC that eventually led to a recommendation, later enacted by Congress, that added roughly a billion dollars to basic computer science research at NSF. This funded the Information Technology Research program. It was that program that enabled my NSF OptIPuter grant to be funded, which has formed the core of my computer science research program for the last five years.

Rick Stevens

Argonne National Laboratory

Dear Ken:

Just a short note to tell you what's been happening since you had to go. BTW I keep waiting for you to pop up on iChat, I suppose I should stop that.

I keep thinking about the multi-core discussions we had last summer, since all issues multi-core are really heating up and it looks like that (as you predicted) will be a major direction/opportunity for the software community for the future.

I have been missing you in the multi-core debates at dinner. I think some progress will be made on the idea of a better set of compiler abstractions for multi-core, but it will be slower than we hoped.

On the DOE front we've been organizing these big Town Hall meetings to make some progress on directions for the next decade in things like exascale architectures, software and applications in energy, the environment and global security (what some have started calling the Long War).

I think you would be enjoying these meetings and we could really use your help with the software discussions for exascale computing.

I really miss your ability to bring people together and get them to lower their shields long enough to see some common ground, we could use some of that magic right now.

I also wanted to say we could still use your help with the petascale software project, as a coherent community wide effort is needed even more now that we are trying to chart a course for the next ten years.

Well, I guess I better run as this note is getting harder to write the longer I try.

I miss you and I hope you are doing okay wherever you are.

Cheers,

--Rick

Richard Tapia

Rice University

I came to Rice in 1970, Ken came to Rice in 1971, and Mary Wheeler came in 1972. There we were --all fresh assistant professors out to conquer the world. We learned together and from each other. I built a world- class program in optimization, Mary built a world - class program in the numerical solution of partial differential equations, and Ken built a world -class program in computer science. We were young and bold and did not know that we were supposed to have senior leadership to build these fine programs. No, instead of senior leadership we depended on each other. We were like brothers and sister. We argued and disagreed often, but we respected, loved, and blindly supported each other. While Ken was the youngest, he was the most professionally mature. He showed

Mary and me how to deal with administration and how to disagree in a professional manner. He taught us much. I feel that Mary and I also gave him much. Ken was my best example of true graciousness. He unselfishly promoted others and received great satisfaction over their accomplishments.

Ken Kennedy was a visionary leader. As others were focusing simply on moving ahead technologically, Ken was also thinking of what the development of computation could represent in individual human lives, the societal changes that broad access to information could bring about, the social implications of Internet use, and how vital it was for the computing community itself to be inclusive. Because of Ken, computer science became far more open to women, Hispanics, and African Americans, and computer science is the richer for it. Ken was a mentor of mentors. I could not have achieved what I have without Ken. Others had told me, "Go ahead and do it," but Ken is the one who made it possible. He provided the resources and allowed me to do the things I had been dreaming about. He helped everybody become better. I carry his teachings with me today. For more than three decades I bounced things off of Ken and he bounced things off of me. Each day I realize more how much I miss him.

Linda Torczon

Rice University

Ken Kennedy is well known in the computer science community for his research, his intelligence, his warmth, and his vision. To me, one of the most impressive aspects of Ken's personality was his ability to forge consensus from a group of colleagues with disparate ideas and opinions. That ability, coupled with his stellar research record, lay at the heart of his successful interdisciplinary and inter-institutional research collaborations. Several of these collaborations (notably CRPC and the GrADS/VGrADS projects) were decade-long inter-institutional efforts. On that time-scale, technology changes, investigators come and go, new research directions arise, and old research directions become solved problems. For such collaborations to succeed, the group must have the ability to change directions, to reallocate resources, and to reorganize—all without damaging the spirit of collaboration among the team members. Ken ensured that the individuals associated with his research projects shared a central vision for the project; he encouraged participants to re-direct their research efforts, as necessary, to contribute toward the central vision as the goals for the project evolved with changes in technology. Ken also fostered communication and collaboration between project participants and members of the computer science community with the combined goals of disseminating project research results and contributing toward a community vision for the area.

One of Ken's lasting legacies is the Computer Science Department at Rice. He led the effort to create the department and served as its founding chair. His skills at collaboration and consensus building played a critical role in shaping the department. After all, he was, at 38 years old, the department's senior member. Yet, Ken was able to

forge consensus, lead communal projects, and foster both a sense of community and a shared vision that continues to this day.

Ken was also dedicated to increasing the number of minority students and women receiving degrees in science and engineering in general, with a particular emphasis on computer science. As director of the CRPC, he helped develop outreach programs targeting underrepresented students, which continue today under other funding and leadership. For decades, he both championed and was involved in a variety of student outreach programs as a mentor, speaker, or organizer. He was a strong proponent of increasing the diversity of the computer science graduate student population at Rice. During his career, he produced eight female Ph.D. students, including me. Ken also mentored and encouraged many of his colleagues and was generous in nominating them for awards and honors.

Ken and I were collaborators and friends for 25 years. We ran research centers and developed outreach programs together; we wrote myriad proposals and papers together. I miss his vision, his energy, his sense of humor, and his friendship.

Steven J. Wallach

Advisor CenterPoint Ventures

Ken and I go back to 1982.

At that time, Ken was a professor in the mathematical science department at Rice University (there was no computer science department at that time), and I was starting Convex Computers. I had read some of his papers on Vectran (a research effort at IBM). The content of the papers impressed me. One of my people called Ken up to ask some questions. He asked if he could meet with someone from Convex. I flew down to Rice and thus began our professional and personal relationship.

When we first met, I immediately sensed that this person was special and extremely gifted. I described what we wanted to do and my vision for supercomputing. I saw no reason why programming a supercomputer couldn't be as easy as programming a scalar computer. He shared my vision (not many other people did). He agreed to help us. The help coming in the form of design reviews and suggestions. Our compiler team would present problems and Ken would suggest solutions. Over the course of 2 years, we produced the world's first production compiler using dependency analysis and a data flow intermediate language for vectorization and global optimization.

Ken provided the intellectual guidance for the group. In his current office, is a picture of the Convex compiler group with Ken at the head. The background is the block structure of the compiler structure. This compiler had an immediate impact on the industry; both from a research and commercial sense. One could take standard scalar code and achieve

high levels of automatic vectorization and optimizations without resorting to special language extensions. Thus began the era of mini-supercomputers in the commercial world, but this also served as a catalyst for an every expanding academic emphasis on automatic compiler optimizations focused on vectorizations and parallelizations.

I remember when at one early conference, one of Ken's graduate students, came over to me prior to a panel session that had me as a participant. He told me that Ken wanted me to make sure that the audience knew that the Convex compiler was based on Rice University technology and not from someplace else. He and Rice were somewhat overshadowed by other academic organizations. But that quickly changed in subsequent years. In retrospect this incident was somewhat amusing.

In somewhat rapid fashion, a computer science department was created at Rice, with Ken as its first department head. Convex was the first industrial affiliate. The commitment was made before a process had been set up to accept affiliates. Once the department was set up, and with me being on the Rice engineering advisory board, a group of us began advocating the construction of a new building to house the computer science department and CRPC (Center for Research on Parallel Computing). Finally with the construction of this building and with the results of the past 10 years or so, Rice's Computer Science Department was recognized as one of the top departments in the United States.

I could go on for many pages listing the technology achievements of Ken. His resume lists the papers and awards. I would like to discuss one more compiler tour de force. After dealing with vectorization and parallelization, at Convex we decided to delve into interprocedural analysis. Ken convinced me that he and his team had made sufficient strides in this area to have someone make an attempt to productize them. In essence, Rice had become Convex's research laboratory. Other companies could also benefit, but we were true believers. Consequently, based primarily on Ken's work, we produced the world's first interprocedural production compiler. There was doubter's initially, but over time, the benefits became clear.

While I have described how Ken's work directly impacted commercial compilers, the working relationships that Ken set up with Convex and other companies became the model for technology transfer between academia and industry. There is always lots of talk about this, but the method's and relationship became a model for other efforts.

Ken and I also developed a personal relationship. There was a HPC conference in Taiwan. Ken was the keynote speaker and I was one of the plenary speakers. On a ride from Taipei to Hsinchu, we came up with the following. Recently Berkeley was touting NOW (network of workstations). We decided that RICE would announce NOW2 (Network of Watches). We would get various HPC luminaries to voice support for this project (including CASIO), and John Markoff of the NYTIMES would publish an article on this project on April 1 (April fools day). April 1 came and went and no article. Markoff subsequently told me that he decided NOT to submit the article for publication. The reason being that no one would believe that the NYTIMES would do this and thus would interpret the article as real. And he would probably be fired.

Recently Ken has been the co-chair of PITAC (Presidential Information Technical Advisory Committee). I am fortunate to also be a member of this committee. The work of this committee and the conclusions reached (as presented in the final report of February 1999), are already beginning to have far a far-reaching effect on the US research community. NSF budgets have been increased, and a new focus on fundamental research as contrasted to applied research has occurred. As one would expect, the conclusions of this committee has also had effects in the international community. In particular, Germany, Taiwan, and Japan (their respective government bodies) have asked for presentations and follow up discussions on the findings of PITAC.

As co-chair, Ken was instrumental in getting the resources necessary to: obtain and analyze the budgets of various US government research efforts, to qualify them, and to get 20 or so high level executives and academics to work together as a cohesive team.

Gil Weigand ORNL and Alex Larzelere

Ken had the foresight to understand that the world of computing was quickly transitioning from gigaFLOP/s to teraFLOP/s to petaFLOP/s and beyond. He was always there for ASCI to help guide the way and never shied away from taking on the really hard problems. He made an important contribution to making sure that the hardware systems would be usable and he served as an important resource for ASCI to review proposals, to assess progress and to guide the program towards the future.

Andy White

Los Alamos National Laboratory

Ken Kennedy was the consummate leader, in clear air and turbulent. At the beginning, the Center for Research on Parallel Computation was a fragile being. The site visit at Caltech was postponed the very morning it was due to begin. When the review actually began at Rice, our performance the first day prompted many of us to consider early flights home. When finally underway, it was clear that a winning proposal and a successful center were two distinct entities. Not without anguish, Ken led us through the doubt and turmoil to eleven years as an extraordinarily successful Science and Technology Center. And he relished the unexpected challenge. Thus, despite our collective protestations, CRPC was reviewed and awarded the highest accolades by the National Academy of Public Administrators.

These risk-taking endeavors and the fellowship that grew up around them were inspirational and interesting things continued to happen. The High Performance Computing and Communications program itself followed on the heels of the CRPC.

Closer to home, collaborative efforts funded by the Office of Science (the High Performance Computing Research Center), The National Science Foundation (Supercomputer Centers) and the Accelerated Strategic Computing Initiative (the Computer Science Institute) built upon the pioneering work of the CRPC. On December 4, 1988 – one day after our second daughter was born – Rice held a press conference announcing the CRPC award. I sat next to a local TV cameraman who was collecting footage for the evening news. About 5 minutes into the festivities, he muttered, “boring,” under his breath and turned off the camera. Turned out it was everything but.

Katherine Yelick

University of California at Berkeley and Lawrence Berkeley National Laboratory

Ken Kennedy’s work on automatic parallelization, compilation of Fortran-D, definition of HPF, memory hierarchy optimizations, and scheduling all had a tremendous impact on the field of parallel computing. Ken’s work was typified by a clear vision of the future, but using a flexible rather than dogmatic approach; his research adapted over his career to the changes in hardware technology, inclusion of ideas from others in the field, and the preferences of the user community. Relative to my own work on the definition and compilation of explicitly parallel languages, his vision of telescoping languages was one of the most influential. Telescoping languages address the problems of productivity and performance in an approach that is both user-centric and practical. The idea is to allow languages to be nested within one another, with a high-level language tailored to the needs of the application scientists at the top, and languages with increasing control over performance features below. Examples of the high level languages in his work (joint with Bradley Broom and others) include Matlab and the statistical language S-PLUS, both of which leverage highly optimized libraries written in lower level languages like C and Fortran. In the telescoping language approach, highly optimized component libraries are annotated with specifications of optimization information, which is used by a compiler for the higher level language. Rather than imposing a new universal language on programmers, this approach simultaneously addresses concerns of language familiarity with the need for high performance and moderate compilation time.

The telescoping approach allows for significant time to analyze and optimize critical component libraries, because libraries are processed once at installation time rather than during application program compilation. This fits well with the idea of automatic search-based optimizations, in which multiple candidate programs are generated, measured, and then one is selected. While the compiler research community (with significant contributions from Ken) made tremendous strides in the analysis and transformation of programs to address the growing gap between processors and memory performance, the increasing complexity of the memory hierarchies made it nearly impossible to generate the optimal code from a conventional compiler. Faced with evidence that, for many machines, sophisticated compilers failed to select appropriate optimization parameters

even for the canonical problem of tiling dense matrix multiplication, Ken shifted his research agenda to include search-based optimization in compilers and organized multiple workshops on the topic. While prior work on autotuning (including Atlas, PHiPAC, and FFTW) used search for library-specific code generators, his work and others extended this idea to general-purpose compilers.