Scaleability and Immortality

C Gordon Bell
Microsoft Research, Silicon Valley
835 Market Street, Suite 700
San Francisco, CA 94111
(01) 415 640 8255
gbell@microsoft.com

ABSTRACT

James Nicholas Gray's understanding and experimentation gave him a special perspective. From 1995 his commitment was building indefinitely scalable tools by working on really hard data-intensive application problems with other scientific disciplines. His attention to research for both understanding and use made him a unique researcher. Jim pioneered a new kind of 21st century science based on data analytics, requiring computer scientists to collaborate as an equal with scientists in other fields.

Gray was an advocate and principal supporter of the MyLifeBits project aimed at extending our memory, memex. The by-product includes digital immortality that we herein speculate and explore.

1. BARC, Scalability, Computer Systems, and the Laws that Govern Computing

I don't recall how Jim and I decided to first meet in 1994, but I feel either flattered and honored, or smart and smug, for what is a wonderful friendship. Jim had just spent the four years since 1991 heading DEC's San Francisco Lab at 455 Market Street and had turned consultant; since 1989, I had been a consultant aka unemployed, Silicon Valley angel investor, being called a crusty hardware guy, and consultant to Microsoft Research.

Our first meeting at my Los Altos home was wide-ranging: scalable architectures, the importance of industry standards, and building influential systems in advanced development. I had just abandoned a 30 year, o(30) system quest of building multiprocessors (scale-up) for clusters (scale-out). We consulted with several companies, but Jim felt that he needed the confines of an organization, and he convinced me that I needed structure, too. He had been talking to Microsoft. We believed Microsoft was the place for setting standards. I jumped the gun and emailed the Redmond folks to hurry up and start a Microsoft SF Lab for Jim:

> Sun Jan 08 15:41:55 1995
> To:rashid@microsoft.com,nathanm@microsoft.com,... Davec@microsoft.com,
> From: gbell@mojave.stanford.edu
> **Subject: Approaches to Servers and Scalability... and an AD Lab here!**
> >Folks, Here's how Jim Gray and I see the next decade or two: A Scalable Network and Platforms (SNAP) architecture (Figure 1) is predicated on one set of standards: a ubiquitous ATM network **and PC-sized platforms**.
>
> SNAP allows upsizing i.e. building world-scale computers from a single platform in a scalable fashion. SNAP will encourage further industry de-stratification. It eliminates the traditional computer price class distinctions (PCs, minis, mainframes) and goes a long way to eliminate the stratified business models that supports traditional, high overhead computer suppliers. SNAP will cause an computer industry upheaval greater than the early 1990s client-server downsizing wave. That wave created a large UNIX market displacing IBM mainframes and proprietary minis. But the UNIX market is fragmented and small when compared to Compaq and NT. UNIX would have to consolidate around one or two dialects in order to get the volumes required to compete with NT. This seems improbable, so Microsoft's NT is likely to become the dominant server standard for all hardware platforms, just as Windows garnered the desktop or client side.

Jim, wrote the recipients, pointing out that he had not "put me up" to write the email, and also enumerating the difficulties of operating a remote lab, validated the importance of scalability. The Bay Area Research Center (BARC) opened in the summer of 1995. Tom Barclay came back from Redmond to join Jim. I was honored and delighted to join in August to work on Telepresence. In October, Jim penned a brief "NTclusters Research Agenda", outlining a BARC project that could be built to test the concept and the need to work with the Redmond NTclusters group.

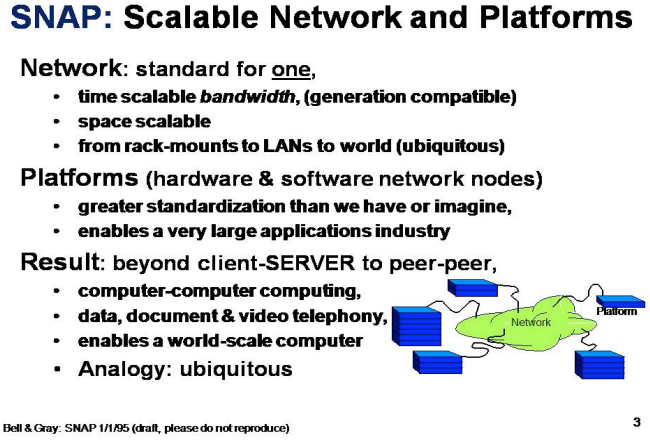

Figure 1 Gray SNAP PowerPoint c1994 positing building everything from a single platform for indefinite scalability.

Jim Gray Tribute, May 31, 2008, Berkeley, CA, USA..

The memo posited Commodity Clusters and How Clusters Differ from a Distributed System. Thus, he had outlined a rough plan for the next decade. He never wavered in the approach to finding challenging science as a way to test the tools so as to refine them!

The Terra Server with 120 disks holding a Terabyte became live in 1996. In the May 1997 Scalability Day, Jim gave an on-stage demo with Bill Gates using over a hundred processors in a room of 20 plus racks to achieve one billion transactions per day.

I also recall the glee when Jim had just finished measuring a half billion transactions per day on April Fool's Day, 2005 using his relatively old laptop and observing: *A $2k computer can execute about 8K transactions per second. This is 80x more than one of the largest US bank's 1970's traffic – it approximates the total US 1970's financial transaction volume. Very modest modern computers can easily solve yesterday's problems[1].*

The Terra Server and Scalability Day results illustrate Gray's fondness for understanding through constant building and experimentation. The extensive publications on Jim's web site shows his drive for understanding through experimentation – standard database and transaction processing metrics, component level performance, and most important, the use of "rules of thumb" backed with constant re-calibration. These metrics and measures have become the basis of measuring progress.

Research is inspired by:	Considerations of use?	
	No	Yes
Quest for fundamental understanding? Yes	Pure, basic research (Bohr)	Use inspired, basic research (Pasteur)
No		Pure applied research (Edison)

Figure 2 In Pasteur's Quadrant characterizes Gray's mode of operation and collaboration with other scientist.

Pasteur's Quadrant [2] (Figure 2) describes Jim's position on research. It characterizes scientists and engineers in research and development. Jim is one of those rare persons in the quadrant where research is inspired by both fundamental understanding and utility. The Sky Server work with Alex Szalay is archetypical.

Through various paths, Jim infected me with the importance of data – it's "all about the data". In one of our more playful times, while discussing how to get the concern for data into the national computing agenda, we bumped into John Markoff, a friend and columnist at the New York Times. We proceeded to posit our (Jim's) view that science is missing the point by just thinking about computation speed aka FLOPS (floating point operations per second). John took our picture in the lab on Friday and the article and our picture (Figure 3) appeared the Sunday New York times, 1 June 2003. Our friends in Washington were not especially happy, but they acknowledged we were right and after years are slowly changing their way of doing research--eScience.

Figure 3 Jim Gray and author at Microsoft's Bay Area Research Lab photo by John Markoff, New York Times, that appeared in his 1 June 2003 article "On Weighing Sheer Power Against Pools of Data"

Two other papers with Jim[3][4] were about the "emperor's lack of clothes" with a singular focus on peak computation speed and lack of concern or balancing storage and database performance. Jim's work on storage and transaction processing performance has been a cornerstone to move the focus to data and the need for databases. Many of today's scientists have yet to understand and embrace databases, other than to use them to track their computational experiment files. The tendency is to "grep" files for an answer versus indexing results using a database properly in the first place in order to get the answers directly. Databases need to be a cornerstone computational science tool.

Ironically, the last talk[4] that we have a record of on his web site posits a Fourth Paradigm for Science based on data. This wide-ranging talk covers the philosophical change in data based science, the need for databases and collaboration with computer scientists, the whole area of scientific publication based on openness, peer review, the use of Wikis, and ends with the need for data provenance.

We went out on a limb at the ACM 50th and speculated about the next 50 years[6]. For years we intended to write a book: *The Laws governing computing—standards, rules of thumb for production, market*, and *industry formation* based on our decades of observations and especially all the of the experimental data. Jim was always too busy to start it.

2. Immortality

After joining I continued to work on telepresence with my BARC colleague, Jim Gemmell. In 1997, our live capture of the ACM 50th Anniversary conference was attended by over 20X the viewers in cyberspace, and convinced us that Telepresentations were the "killer app". But by 1999 I *rightly* felt that more telepresence applications were still a decade away—and I am fond of short term results.

In early 1999, motivated by Bush's Memex and with Jim Gray's suggestions and encouragement I started to move into cyberspace by collecting and digitizing my past lives (articles, books, correspondence, documents, music, and videos), rationalizing that teleworkers required efficient and paperless work at home environments.

In 2001 Jim Gray and I co-authored a paper on Digital Immortality,[7] marking the real beginning of the MyLifeBits project and tools. Roger Lueder joined Gemmell and I and the work began in earnest, with Jim's encouragement and support to understand the use of and essential need for a database (I had initially told Jim: "Why do we need a stink'n database?") for holding all of a person's information. Jim had posited the clear research goal (one of a dozen) in his 1997 Turing Lecture: "Personal Memex: Record everything a person sees and hears, and quickly retrieve any item on request." MyLifeBits has remained a quest for Gemmell and me.

Our short paper on digital immortality went beyond bit preservation of legacy data, and the increasing number of live data sources including continuous, real time capture devices for immortality to the notion of two-way immortality. We went out on a limb and predicted: "we believe that two-way immortality, where one's experiences are digitally preserved and which then take on a life of their own, will be possible within this century." This still seems as far away as in 2001 even though avatars are in use! However, I do advise, the MyCyberTwin company.

Ideas embedded in algorithms, art, articles, books, and "code" including music that are the foundation of a community and used by future generations are the number one form of immortality. For personal immortality, unless some effort is made, e.g. Richard Feynman's book "What do you care what other people think?" no one outside of a circle of personal friends understand "what Jim is like". But friends are not immortal. The one hour video interview in "Behind the Code"[8], with blogs and Jim's responses, is the only item that gives a tiny personal glimpse of a Jim's attitudes and personality. *I have observed that videos of computer pioneers e.g. Atanasoff, Mauchly, Noyce, Olsen give a glimpse of their personalities, far better than audio or transcript interviews and their classic or cited articles that are their official "records".*

2.1 "Profession lives" immortality

Someone's work can be immortalized, as in the paintings of the great master's, buildings by a notable architect, or, for us scientists and engineers, some formula or approach. Going deeper, the *way* they worked may be immortalized: their techniques, their approaches, their professional relationships, and the stories of them at work. For instance, we know a fair bit about the work of Isaac Newton, including the story of Newton, in his early twenties, going to the countryside to avoid an outbreak of the plague and, like any typical young man with too much time on his hands, whiling away his time – inventing calculus and discovering the law of gravitation.

The World Wide Web is already being used to make this kind of knowledge more accessible and in-depth, so providing immortality of scientists and engineers on the web is not new. NIH's National Library of Medicine chronicled several of its notable scientists by digitizing their paper and computer files. These web archives usually have 50 kinds of items e.g. article, bibliography, brief, brochure, certificate, code, coding standard, correspondence, drawing, exam, financial record, interview, note, notebook, poem, resume, schedule, transcript, video recording. Web file archives give an impression of a professional's life output and their classic books and papers, but rarely capture or quantify a person's attitudes. Sites are typically devoid of personality –e.g. there are no personal recommendations or feelings about persons or mankind generally. Depending on the individual, it is essential to have a structure e.g. timeline, interpersonal relationships of one's life work. A few sites have videos, but interaction to easily browse is generally nil except for rare productions e.g. Leonardo's life. As the web has evolved, it is customary for an individual to have their own web page to display their professional lives, with web searches revealing other bits and pieces of professional lives e.g. photos, reports, from others in a community. To have a site that characterizes an individual's life for immortality most likely will require an effort similar to the professional effort of living the life itself! As a minimum, stories are essential to hold students, scholars, or even casual observers.

We have only other's reminiscences to describe, for example, why Jim is so unique. The Tributes are an excellent way to incorporate these views that hold the stories.

Based on MyLifeBits, I believe it is useful and perhaps essential to introduce the notion of lifelines or independent "lives" to characterize a person, over some period of time, in a community, operating toward a common goal or constraints such as a company, faculty, or professional group. It is furthermore usually necessary to divide these "lives" into personal and professional. Lives usually have minimal overlap, although with Jim, all of us count him on both out professional- and friend-life lines.

Jim's web site is a partial record of his many major "professional lives" including an extensive vitae starting in 1962 with a co-op job and 10 other academic and commercial organizations. The website has: 185 articles and books since 1966; and 212 Microsoft and system related talks since 1994. It lists 50 events in 2006 that Jim was part of; a dozen web sites Jim operated; education and honorary degrees; memberships in over 20 pub boards, advisory committees, program committees, and five societies including all three US Academies and the European Academy of Science. In 2006 Jim had awarded and was responsible for eight eScience group grants, so given Jim's nature, these likely amount to eight more lives as an unpaid consultant, contributor and helpful critic within these group.

Since 2000, Jim had practiced another kind of interaction with science as typified by his work with Alez Szalay to create the Sky Server. Gray and Szalay pioneered an approach which can be summed up as: "*Find a real project within a scientific discipline that is limited by data understanding, find a post-doc and just do it*". This methodology is transformational to science at a time when science is moving from simulation to data-based science enabled by tera-, peta-, and even exa-scale data storage. Jim's heuristics for collaborating with scientists who work with real world data, may be the key to success. At the same time it transforms both the science and computer science. It doesn't take too much reading through the extensive publications on Jim's web site to glean his drive for understanding through the measurement. However, while you will find the term "eScience" and the relevant projects on Gray's web site, the approach he pioneered with Szalay needs to be more explicit and then institutionalized in a fashion that computational science never achieved.

But the website is just a glimpse of Jim's prodigious output. By his vita and the Berkeley Tribute, we see o(20) of the major active parallel lives or timelines within his lives at IBM (11 years), Tandem Computer (10 Years), and DEC (4 years) including lives within the computer science, database, fault tolerance, performance, transaction processing, and systems communities. Microsoft "lives" since 1995 include eSciences and systems generally, with contributions to: astronomy, bioinformatics, databases, oceanography, servers, and terrestrial data. One's

professional society and educational institutions extend to more "lives". Another view is how all of our lives intersect. I have: 13K emails, 1600 web pages, 100 presentations, 50 photos, and 600 documents that are either joint or reference Jim, plus one phone call, and the penciled in calendar entry of our first 1994 meeting and a December Saturday sail.

Of course, Jim's computer and all the corporate archives, have several million items including: awards, bibliographies, budgets, calendars, code, correspondence and email, design notes and notebooks, interoffice memos, personnel evaluations, plans, recommendation, reports, reviews, specifications, technical reports, transcripts. His computer could reveal music and numerous photos we'd like to see shared, for example even if they don't always include us with Jim. All the stories of the meetings and social interactions we have but can rarely recall to tell are sadly missing –these stories are essential for immortality beyond each of us. It is unclear how to ever capture, produce, share, and find these. Recorded tributes help understand these lives through all the stories we may use to explain our respect and love for Jim.

2.2 "Personal lives" immortality

Personal immortality adds many more lives in roles within families as child, sibling, spouse, parent, grandparent the area of personal immortality. Few of us maintain the "family tree" or relations beyond the immediate family (parents…great grandchildren, plus aunts, uncles, and cousins). With "personal lives" immorality may only be about stories that are triggered by communication (now email, but formerly letters), photos, videos, and wetware aka personal memory.

2.3 Toward a Digital Immortality

Rarely do any of us create and archive the stories that can be transmitted beyond each of the shared lives and save them permanently for others. Immortality for an engineer or scientist is in their artifacts, ideas and methods carried by artifacts (e.g. article, book, code) and especially colleagues.

The public tribute with videos that capture our stories of Jim's many lives help amplify and enlighten a few of the many *personally-professional* lives beyond a formal stark web or other archive.

Can we use this tribute as a challenge to the creation, enhancement, and technology for preservation and immortality? As a scientist, I suspect Jim would insist on this challenge – another chance to create understanding and something of value, not just a chance to see old friends, express our respect, and hardest of all—test our feelings!

3. ACKNOWLEDGMENTS

Jim Gray for the teaching, mentoring, and all the encouraging during our brief collaboration since 1994.

4. REFERENCES

[1] Gray, J. "Thousands of Debit-Credit Transactions per Second: Easy and Inexpensive", MSR Technical Report MSR-TR-2005-39

[2] Stokes, D.E. Pasteur's Quadrant: Basic Science and Technological Innovation, Brookings Institution, 1997.

[3] Bell, C.G., Gray, J. "What's next in high-performance computing?", CACM, V.45(2) Feb. 2002, pp. 91-95. Also see MS-TR-2001-76.

[4] Bell, C. G., Gray, J., Szalay, A, "Petascale Computational Systems: Balanced Cyber-Infrastructure in a Data-Centric World ," Letter to NSF Cyberinfrastructure Directorate, IEEE Computer, V. 39(1), , January, 2006 pp 110-112.

[5] eScience Talk at NRC-CSTB meeting Mountain View CA, 11 January 2007. Available on research.microsoft.com/~gray

[6] Bell, C.G., Gray , J., "The Revolution Yet to Happen", in Beyond Calculation, The Next Fifty Years of Computing, PJ Denning, RM Metcalfe editors, Springer Verlag, NY, 1997. Also see MSR-TR-98-44

[7] Bell, C.G., J. Gray, "Digital Immortality" CACM V.44(3) Mar. 2001, pp. 28-31. Also see MSR-TR-2000-101

[8] Interview of Jim Gray, Behind the Code, March 3, 2006 MSDN Channel 9