Short Biography of Bryant W. York

Dr. York holds an A.B in Mathematics from Brandeis University, the M.S. in Management from the Sloan School at MIT, and the M.S. and Ph.D. in computer Science from the University of Massachusetts at Amherst. He has held industrial research positions at the IBM Research Labs in San Jose and at Digital Equipment Corporation's Artificial Intelligence Technology Center in Massachusetts. He is currently professor and research director in the Computer Science Department at Portland State University.

Dr. York is a member of the Board of Directors of the Computing Research Association and is beginning his second term as a member of the Advisory Committee for the Computer Information Science and Engineering Directorate of the National Science Foundation (1992-97, 2002-present). He is a former member of the User Advisory Council for the National Computational Science Alliance (1999-2001), a former member (1991-94) of the Advisory Panel for the Ethics, Values, and Society program within the Social Behavioral and Economic Sciences directorate of the National Science Foundation. He is a member of several professional organizations including ACM, SIAM, IEEE-CS, AAAI, and AAAS. He was also a member of the ACM Education Board (1991-1996) and a member of the ACM U. S. Public Policy Committee (1992-2000). He was the 1994 chair of the ACM committee on minorities and co-chair of the Education Committee for the Supercomputing 97 Conference. Dr. York received outstanding service awards from ADMI in 1991 and 1997 and he is the 1998 recipient of the Computing Research Association's A. Nico Habermann award for service to underrepresented groups in computing. In 2001 he was the first recipient of the Richard A. Tapia Award for Scientific Scholarship, Civic Science, and Diversifying Science.

Dr. York has research publications in several areas of computer science, including computer vision, expert systems, software engineering, computer assistance for persons with disabilities, and parallel computation. His current research interests are in the development of parallel algorithms for image processing and advanced scientific computations, and in the development of educational software for teaching K-12 mathematics. He is working with protein crystallographers on parallel computations for crystallography, with mathematicians on applications of Clifford Algebra to image processing, and with bioscientists on proteomics. On the educational front Professor York is working with a large multidisciplinary group of scientists to produce tools to support culture specific interactive learning systems. Specifically he is working on a project entitled the Multicultural Mathium, a collection of knowledge-based tools to assist in the development of customized mathematics curricula for individual children, aged 6 to18. The goal of the work is to develop tools that help cognitive scientists, psychologists, cultural anthropologists, and others understand the kinds of curricula and technology that enhance learning in children.

Over the years the NSF, NASA, ARPA, AFOSR and Digital Equipment Corporation have supported his research.

DFL: A New Paradigm for Computing

Bryant W. York
Portland State University

At present the earth supports about 6×10^9 humans, each with about 10^{10} bits of storage and with processors that execute on a millisecond scale. Human sensory bandwidths are approximately 10^6 bits/s for the eye, 10^4 bits/s for the ear, and 10^2 bits/s for the fingertip. It is projected that the earth can support only about 15×10^9 such entities. I speculate that over the past fifty years individual humans have increased the average number of relationships they have with other human beings during a lifetime from about 10^3 to 10^5 primarily as a result of improvements in telecommunications.[1] At a coarse level of granularity we have a dynamically expanding graph of approximately 6×10^9 nodes each of in/out degree 10^5 (a non-SIMD massively parallel machine of sorts). I would call the exotic computation being executed by this machine, *species survival*. My challenge question is "How can we keep such a machine up and running?" The answer is not DoD support! More specifically, I should ask, "What can we as computer scientists contribute both architecturally and in terms of the *computation*?" Curiously, I am reminded of the nickname for the old CM5 network state-saving instruction, "All Fall Down," when I ponder the consequences of system failure.

Permit me to extend this metaphor a bit further. For the past generation the prevailing assumption has been that we could improve the reliability of the system and hence more effectively perform our exotic computation by adding progressively faster non-organic processors and memory to support the "slow" nodes. The number of fast processors has grown from a few tens in the late 1940s to probably on the order of 10^{12} currently with a concomitant growth in non-organic storage and bandwidth; yet the problem of keeping the slow node network up seems to be getting more difficult. Eventually we began to think that the source of the problem was that information did not flow from one site to another quickly enough, so we added high bandwidth networks to the mix. Slow nodes now have to interact with an ever-expanding high bandwidth network of fast nodes. These interactions are often less than satisfying and frequently disruptive. Clearly the quality of human life (a measure of the effectiveness of the exotic computation) has become highly dependent upon an individual's ability to manage these interactions. More specifically it is dependent upon the individual's ability to manage large numbers of transactions and large amounts of information in a timely and effective manner. This ability is not innate and training in this skill is differentially available to the earth's population. As a non-systems researcher, I would like to see the development of new computing and communications technologies that enhance the quality of life for all of the slow nodes.

Rather than pursue our current asymptotically bankrupt path of adding complexity and performance to the network without understanding the needs of its primary stakeholders, I think that new paradigms are required. In particular, I would like to see a new paradigm that includes not just human constraints in designing the I/O of a system, but also

[1] My estimate includes casual relationships such as calling your telephone service provider to complain about your bill.

includes human needs, aspirations and goals. I would call this paradigm, *Design for Livability.* Let me give some historical examples in support of my case. For years hardware systems were designed in a nearly random fashion, thus making professional life difficult for hardware test engineers. Eventually some brave souls put forth the idea of Design for Testability (DFT). It caught on because it made systems more reliable and more reliable systems were easier to sell. DFT was an acceptable solution only when the three primary stakeholders (hardware designers, hardware test engineers, and marketing people) agreed that locally optimal behavior by one of the stakeholders (hardware design engineers) did not lead to a global optimum. Although DFT improved the quality of life (an aspiration of hardware test engineers) for only a small segment of the population, it is representative of the DFL paradigm on a small scale. In the mid 1980s hardware and software support for persons with disabilities were afterthoughts. Through the efforts of a small but determined group of people, hardware and software designers were forced to grapple with accessibility issues at the initial stages of system design. Design for Accessibility (DFA) is a second example of DFL that addressed the needs and aspirations of a small population. Now what is the inductive step? We cannot address the individual needs of every person on earth, or can we? Before I address this question, let me note that slow nodes cluster into families, tribes, nations, religions, cultures, professional societies, trade associations, political parties, etc. to lobby for common interests. Problems arise when different clusters have conflicting goals and vie for a non-sharable or consumable resource. So it appears that DFL depends on DFCR (Design for Conflict Resolution) and might benefit from existing computer science models for resource allocation.

Now, let me return to the question of whether we can address the individual needs of every person on earth through advanced computing technologies. Amazon.com currently has data warehouses that store information on the purchasing preferences of millions of people and fulfillment centers to distribute the products. It regularly mines this data to develop individualized marketing programs. If Amazon is able to accomplish this with current technology, is it inconceivable that we could develop individualized learning programs for every individual on the earth? Why do I suggest individualized "learning programs?" I am betting that "smarter" slow nodes are really the key to the non-termination of the exotic computation. So, my grand challenge application is to develop all of the human resource on this planet through customized education/learning programs. Although this requires significant research in human cognition and in the ethical uses of technology, I will confine myself to listing three of the computing systems research challenges that I think are necessary to support this application.

- Research in large-scale data fusion and modulation for human bandwidths.
- Research in parallel reconfigurable architectures for data mining.
- Research in large-scale, high-dimensional data representation and manipulation.