Joseph M. Hellerstein is an Associate Professor of Computer Science at the University of California, Berkeley. He is an Alfred P. Sloan Research Fellow, and a recipient of the NSF CAREER, NASA New Investigator, and Okawa Foundation Fellowship awards. In 1999, MIT's Technology Review named him one of the top 100 young technology innovators (TR100).

Hellerstein's research focuses on database and networked information systems, including adaptive dataflow systems, sensor and peer-to-peer networks, interactive data analysis and transformation, query optimization and execution, federated systems, and indexing of complex data and queries.

Hellerstein received his Ph.D. at the University of Wisconsin, Madison, a Masters degree from UC Berkeley, and a Bachelor's degree from Harvard College. He spent a pre-doctoral internship at IBM Almaden Research Center, and a post-doctoral internship at the Hebrew University in Jerusalem.

Hellerstein is a member of the advisory boards of ACM SIGMOD and Ars Digita University, and has served on the technical advisory boards of a number of software companies. He was a co-founder of Cohera Corporation (now part of PeopleSoft), where he served as Chief Scientist from 1998-2001. He currently serves as a Faculty Research Associate of the Intel Berkeley Research Lab.

Health Care: "Killer" App Challenges for Systems Research

Joseph M. Hellerstein, UC Berkeley

Computer systems research is notoriously self-involved. Many influential systems (e.g. UNIX) were developed to enable computer scientists to do more computer science. Commercial descendants of these system infrastructures (e.g. recent Windows, MacOS) are in many respects extensions of the inwardly focused original systems. One hardly needs to remind the CS community that the most widely used computing "system" – the World Wide Web – was the brainchild of physicists, who developed both the successful metaphor and the initial software to truly bring IT to the people.

Of course the sense of focus that characterizes the CS systems community has led to enormous innovation and maturation in a relatively short period of time. But as traditional computing becomes commoditized, the CS systems community must do its own work reaching out and justifying its existence in broader spheres – with applications that speak to society as a whole.

Choosing a Killer Societal App

The US Federal Budget serves as a coarse but useful metric of the value – in social and economic terms – of various national projects. Therefore an overarching direction for grand challenges is to reduce the cost and/or raise the quality of major items in the Federal Budget.

A quick glance at the *Citizen's Guide to the Federal Budget* (http://w3.access.gpo.gov/usbudget/fy2001/guidetoc.html) shows a number of obvious candidates for attention:

- The largest fraction of the budget (23%) goes to Social Security, which by definition should not be amenable to reduction via clever engineering. Hence we can ignore this category.
- Discretionary defense spending represents 16% of the budget. We already have significant input in the defense community, and can hardly revolutionize our existing level of impact. Hence we can ignore this area as a potential new grand challenge.
- When combined, Medicare and Medicaid form 19% of the 2001 Federal Budget. Computer systems research is notably absent in this arena, even though health care is an environment highly receptive to new technologies.

Hence an obvious target for "grand challenges" is in health care in general, and in applications for elderly and low-income recipients in particular. In addition to having obvious budgetary value, it has the serendipitous advantage of being a social – one might say moral – imperative. Perhaps the "killer" app (so to speak) will be found by trying to improve both the quality and efficiency of health care.

Computer Systems for Systemic Health Care Improvements

Many of the building blocks needed for making fundamental improvements in health care are hot topics in computer systems today. Hence we need not necessarily derail the community's

interests to address the next killer app. At the risk of leaving out broad swaths of research, we list three selected example areas here:

- **Embedded, networked, autonomous devices:** Embedded biosensors and actuators are an excellent example of challenges and opportunities in tiny, wireless, networked devices. Autonomous biosystems are well-suited to elderly and low-income patients, who typically have a lower rate of "compliance" with typical medical protocols (e.g. regularly taking drugs in the right doses) than younger, affluent populations. These devices have systems research challenges in MEMS design, ad-hoc networking, power-sensitive protocols and algorithms, zero-administration adaptive systems, and so on.
- **Federated information systems:** Health care information is spread across numerous governmental and private entities, with complex privacy and other legal obligations. While the medical world loves to spend money on new treatments, they tend to lag in IT investment. Hence many of these systems are legacy systems. The development of federated middleware to link these systems is a major issue. Challenges here include protocols for allocating resources and scheduling computation across administrative domains, integration of heterogeneous data without unified schemata, and reliability and fault-tolerance under distributed administration. Peer-to-peer systems represented a simplistic but achievable toehold on the mountain of problems in federated systems.
- **Security and Privacy:** There is good reason for embedded biosystems to participate in wireless communication: e.g., monitoring, connectivity to server-side decision-making software, and even peer-to-peer networking among biosensors within a single body, or across bodies (e.g. for collaborative sensing of public health threats). Network security is critical here, especially in bio-actuators: nobody should hack into pacemakers or drug delivery systems (especially *en masse*, a frightening terrorism threat). Similarly, privacy concerns require security in biosensors as well. A wrinkle here is that the security must not seriously compromise power consumption. Federated issues arise again, e.g. in managing access control across disconnected administrative domains under various (possibly life-threatening) circumstances.

Confessions and Conclusions

The author should be clear that by the standards here, he has been working in the old-fashioned style of a traditional CS systems researcher, doing "plain old" systems research without obvious direct societal benefits. He does not claim to have thought deeply about how to use his computer systems skill set to make a tangible difference to Medicare or Medicaid recipients, or to the Federal budget. That said, a call for "Grand Challenges" would seem to require descriptions of as-yet-unmet goals … hence a very speculative paper seemed more appropriate than a sales pitch for current work.

Perhaps the biggest challenge in making a real difference in this area is the broader *systemic* challenge itself: understanding the interlocking and complex relationships between computer systems, users, the administrative and governmental realities of modern health care, the political and ethical aspects of embedded computing in health care, and other issues that are not readily apparent in a speculative essay. This requires more than one person's clever ideas. It needs a coalition of people across computer science and other aspects of society to be energized and working together – a real systems challenge.