Testimony of
Robert F. Sproull

Before the House Subcommittee on Research and Science Education
Committee on Science, Space and Technology

Oversight of the Networking and Information Technology Research and Development Program and Priorities for the Future

September 21, 2011

Thank you Chairman Brooks, Ranking Member Lipinski, and the other members of the Subcommittee for this opportunity to discuss the Federal government's Networking and Information Technology Research and Development program. I am pleased to offer my perspective on your questions based on more than 40 years of experience doing or managing computing research in academia and industry and also advising high-technology venture capital investors. Among other roles, I currently serve as the chair of the National Research Council's Computer Science and Telecommunication Board (CSTB), and recently retired from Oracle as the Director of Oracle Labs. This is an applied research laboratory first started by Sun Microsystems in 1990 and retained by Oracle when they acquired Sun in 2010. I present today's testimony as an informed individual and not as a representative of any particular organization.

Introduction

Extraordinary economic and societal benefits have exploded from the U.S. NITRD ecosystem, which is a complex interplay of government, academia, and industry that dates back more than 40 years. Some of the technologies themselves have improved extraordinarily, such as the price/performance of microprocessors; equally, new markets have grown explosively as networking infrastructure and low-cost electronics have enabled innovative products and businesses. I will describe below some of the aspects of this ecosystem, especially the importance of fundamental research and the interplay of government, academic, and industrial roles.

I wish to stress at the outset, however, that this ecosystem would not have been born, nor would it be successful today, without a vigorous, thoughtful strategy of Federal investment in fundamental research in NIT. Especially important in the early days were programs of long-term research sponsored by NSF and ARPA. An important milestone was the High Performance Computing Act of 1991, which recognized the importance of high-performance computing to Federal missions, especially those of Defense and Energy. But as IT technology itself became more pervasive in the U.S., signaled most vividly by the blossoming of the World Wide Web in 1993, a wide class of NIT technologies became critical to short and long term requirements of many more Federal agencies. The Act and its research coordination role were appropriately extended to address the expanded set of challenges. This extension in scope must continue: today, NIT's role in national security, national competitiveness, and national priorities is far broader than high-performance computing alone.

NITRD Goals

The nation has identified advances in energy, transportation, health, and cybersecurity as important national priorities. I concur with the PCAST NITRD Working Group, on which I served, that these are important drivers where NIT research and innovation can make enormous contributions, and with the PCAST report recommendation to expand NITRD's purview as necessary to address these areas. H.R. 2020 is an excellent first step, identifying cyberphysical systems in particular for more attention. The recently announced National Robotics Initiative is a concrete example of investing in cyberphysical systems research. But there is an even wider need for cyberphysical systems in achieving national priorities, for example as part of controllers and systems that achieve efficiencies in energy and transportation, and for monitoring patient health. Indeed, the national priorities show a broad panorama of areas, including high-performance computing, in which NITRD investments will be essential.

Sun Microsystems' Research Lab, an Industrial Contributor to the NITRD Ecosystem

Sun Microsystems was founded in 1982 to build advanced computer workstations, based on results of research conducted primarily at Stanford, Berkeley, and Bell Laboratories. In 1990, Sun created a research laboratory. I was a founding member and eventually became its director. When Oracle acquired Sun in 2010, they retained the lab as a way to start Oracle Labs. I retired from Oracle earlier this year.

I characterize the lab as an "applied research lab," in that most of its research projects, though risky, have medium-term objectives (e.g., less than 3 years) that, if successful, would have a significant impact on a Sun product or product line. Our job is to selectively explore risky ideas and reduce their risk to a level that would be acceptable to an engineering team. Ideally, our research team would then transfer to the engineering organization, carrying its ideas and insights into a larger engineering team. We like to say that "technology transfer is a contact sport," meaning that the most effective transfers from research to engineering are those that transfer people.

The lab was deliberately kept small, with a budget of about 2% of Sun's total R&D budget. SunLabs hires mostly PhDs in computer science and engineering fields, but also high caliber college graduates in those fields. When the lab started, our CEO Scott McNealy explicitly asked us to be "eyes and ears" for Sun, to participate in the global IT research community, to learn from it, and to contribute to it. Our researchers are nationally and internationally known, attending and presenting papers at international conferences.

SunLabs does very little fundamental or long-term research. An applied research project might develop broadly applicable results, but that is not its principal objective. In order to import a broad range of fundamental new ideas, we pay careful attention to academic researchers and their results, as McNealy requested.

Sun evolved a system of "collaborative research" with academic partners. We would contribute money or equipment to an established university research project that we judged might be able to contribute to Sun's technologies. Then our researchers would interact closely with those in the university. We encouraged academic researchers and graduate students to work with us at Sun, as consultants or student interns, to learn from their ideas--again through people. For example, Sun's embracing of Reduced Instruction Set (RISC) processor technology – a technology behind most computer processors in use today -- was accelerated by collaborating with the RISC research group at U.C. Berkeley and by consulting help from its principal investigator, Prof. David Patterson. This model served Sun well, and helped us sustain innovation at a time of rapid technological change. These collaborative interactions with academia also allowed us to present challenging Sun problems to academics and thus influence academic research agendas.

Though Sun Labs managed almost all the *research* projects at Sun, it was responsible for only a fraction of Sun's *innovations*. The product engineering organizations, developing both hardware and software, routinely innovate. For example, Sun is famous for introducing in 1984 the "network file system" (NFS), which allowed computers to share files over a computer network. Though innovative, its development was not the direct result of research.

Incidentally, I dislike the word "breakthrough," because it is too often assumed that breakthroughs are the only objective of research and stem only from research, especially fundamental research. To the contrary, high-impact innovations can emerge in many ways, and sometimes the principal reason for the high impact--and thus perhaps the perception of a "breakthrough"--may simply be a sharply lower price or rapid market penetration. But these dramatic advances usually depend on much varied research, much incremental, perhaps some revolutionary, and often far earlier than the apparent "breakthrough."

The NITRD Ecosystem -- The Big Picture

As part of an early assessment of the High-Performance Computing Act of 1991, a study by the National Research Council developed a graphic presentation known as "the tire tracks diagram" to illustrate some of the features of the complex interactions among government, academic, and industrial players that lead from early research to several billion-dollar subsectors of the IT economy. The graphic is attached below.

The graphic charts the development of technologies from their origins in industrial and Federally-supported university R&D, to the introduction of the first commercial products, through to the creation of billion-dollar industries and markets. The principal features of the NITRD ecosystem that this diagram illustrates are:

- Contributions are made by universities (usually Federally-funded) and industry, in varied orders and magnitudes. Ideas and people often contribute to different paths; there are frequent flows from academia to industry and vice-versa. There is no direct path from research to impact.

- Initial research often takes a long time to pay off; 15 years is typical.

- Research often pays off in unanticipated ways: developments in one sector often enable advances in another, often serendipitously.

- Innovations occur at all points along technology trajectories, not only in research settings.

- University and industry research are different: university research favors long-term fundamental problems, while industry generally focuses on the next product cycle or two (at most a few years). Results of university research are public and available to all, creating a challenge for industry uptake.

The original diagram produced by the NRC in 1995 identifies 9 billion-dollar sectors. The updated diagram produced in 2005 shows 19 billion-dollar subsectors of the IT economy, each of which bears the clear stamp of Federal investment, usually in high-risk research with uncertain commercial application or payoff. The Council is at work now producing the next version of the chart, and they are likely to identify several new billion-dollar sub-sectors – search and social networking, for example -- that have emerged just since 2003.

The NITRD Ecosystem -- A Java Example

In the late 1990's Sun Microsystems introduced the new Java programming language. Although new programming languages are rarely adopted widely, Java became popular because of its ability to run robustly on many different computer types and because of its modern design, especially features that reduced some of the tedious chores of programming; that is, it increased programmer productivity. Many IT staffs and product developers embraced Java to program their products and services. Today, Java is often taught to high school students as their first programming language. One of the reasons Oracle acquired Sun is that much of Oracle's product suite had come to depend on Java.

Java was designed by James Gosling in 1991 as part of a research project exploring ways to use graphical point-and-click user interfaces to control televisions, set-top boxes, kitchen appliances, and other consumer gear. This product objective did not succeed, but Java found a foothold in the mid '90s as a way to program Web browsers to create animated and interactive experiences. Early releases for this purpose reached a large number of programmers, the language became quite popular, and Sun went on to develop versions for conventional computer systems (as opposed to browsers).

Java's design and implementations draw heavily on preceding research in many areas. Object-oriented programming languages had long been studied by industry and academia. Especially important was the SmallTalk language, developed by Xerox researchers in the 1970s, inspired by a language named Simula, developed by Norwegian researchers in the 1960s. Research to speed up execution of SmallTalk programs became a popular focus of university research on a wide range of fundamental language implementation problems. For example, a graduate student at Stanford, Urs Hölzle, developed a revolutionary way to generate fast code for the Self language, a close kin of SmallTalk. He and others founded a startup, Animorphic, to exploit this technology in a commercial SmallTalk system to compete with other SmallTalk offerings from a small group of startups (none of which survive today).

When Java became popular, the Animorphic team quickly retargeted their work to a Java implementation, judging that it would have greater commercial value than SmallTalk. Sun, looking for ways to speed up its own Java implementations, bought Animorphic, and the team incorporated their technology into Sun offerings, where it became known as "HotSpot technology."

This is but one of many threads from research to product that contributed essential components to Java technology.

This detailed glimpse of one of Java's technology paths shows the NITRD ecosystem at work. The players are global; there are complex interactions among industry and academic researchers; people and ideas flow rapidly; startups play an important role whether or not they ultimately succeed as standalone businesses; fundamental innovations may take a long time to reach mainstream products; a commercial success will track back to countless research projects and results, many of them funded by the Federal government. The ecosystem collapses without Federal support of fundamental research.

Characteristics of the Ecosystem

Using the term "ecosystem" to describe the complex interactions among participants in NITRD activities may seem a stretch, but the term is apt. There are many distinct players, with varied but blurry roles, and complex dependencies. As we've seen, an IT product depends on other NITRD activities in complex ways akin to the dependencies in a biological system. Different players perform complementary roles. Long-term academic research provides new results whose impact cannot be predicted at the time. Industry amplifies these results through its own applied research and product development processes. The overall health of the system depends both on funding from government and from revenues received for products and services offered by healthy IT businesses.

Like a biological ecosystem, the NITRD ecosystem could be disrupted or damaged inadvertently. The NRC report *Assessing the Impacts of Changes in the Information Technology R&D Ecosystem* addressed exactly this concern in 2009. It concludes that Federal investment in research is essential and dangerously thin. It points to the importance of venture funding. It also points out that the ecosystem includes customers: "The most dynamic IT sector is likely to be in the country with the most demanding IT customers and consumers." Thus, for example, improving U.S. broadband networking is essential to creating the demand to develop world-class innovative services.

The most dangerous and least visible threat to the ecosystem is that we all focus on short-term research and payoffs, thus underinvesting in the long-term research that may lead to extraordinary technical advances and returns.

Investing in fundamental research is risky, and the amount and character of payoff cannot be predicted. But Federal sponsors have an excellent record of directing fundamental research, in concert with the research community itself. DARPA, for example, pursues military needs, and its long-term vision and investments have resulted in fundamental and high-payoff results such as interactive computing, networking, and RISC microprocessors. NSF's recognition that digital libraries would become important led to high payoff in search engines, which can be seen in today's search offerings from Google, Yahoo, and Microsoft. The wisdom of long-term Federal research investments is evident in the productive ecosystem they have spawned.

As I remarked in my introduction, the national goals in energy, transportation, health, and cybersecurity are excellent guides for today's NITRD research investments. Who knows what billion-dollar NIT industries may emerge from research toward these goals?

The Research Workforce

I want to offer a few comments on the workforce available to industrial research groups. Note that this is a small subset of the overall IT workforce. I offer these comments to emphasize the varied nature of skills and training in the workforce.

At SunLabs, we hired mostly PhDs, many fresh out of graduate school. Candidates come from all over the world. In most cases we know of students finishing their degrees because we have ongoing collaborations with their professors or the students themselves. In all cases, we seek candidates who have demonstrated research skills in areas aligned with the research project we are staffing. For example, a project to explore new ideas in building Java "virtual machines" seeks candidates who have built virtual machines, garbage collectors, or other programming-language artifacts as part of their academic research. Consistent with our objective of transferring technology by transferring people, we seek researchers adept at building systems and willing to join engineering teams.

Although we expect staff to work from one of our two lab locations in the U.S., this is not always possible. Some candidates have family constraints that prevent a move. Some foreign nationals cannot obtain visas, or must work from abroad until a visa can be obtained. The international Internet makes remote work ("dispersed R&D") possible, but not preferable. Location still matters, but as networking improves, it matters less.

Understanding Federal Research Investments

The PCAST Working Group that examined the NITRD program had trouble determining the levels of research investment in different areas because of difficulties in labeling and measuring expenditures. In industry, we make clear distinctions between different kinds of investment in IT, in part so that the investments can be balanced appropriately.

First, support of fundamental and applied *research*. The goals of this work are too risky to depend on results to meet customer or market needs.

Second, investments to *develop* new IT products and services, some for sale and some for internal use. These developments may be routine or highly innovative, but the development itself is not very risky: schedules, milestones, tests, and periodic releases characterize the work.

Finally, investments in NIT *infrastructure* that support all parts of the business, including the two items mentioned above, NIT research and NIT development. These investments are usually the least risky and innovative of all, and are usually driven by estimates of computing and networking capacity needed. As NIT infrastructure becomes necessary to support almost all business activities, these expenses are similar to those for space and utilities, and are accounted as an overhead for the activities they support.

Federal budget reporting makes it difficult to distinguish these three classes of investment. Infrastructure, in particular, should not be characterized as an *NIT* R&D investment unless it supports NIT R&D itself. For example, a Web server that provides citizen access to an agency's database is not an NITRD investment, though it is a use of NIT. While distinctions between *research* and *development* (the first two categories) are sometimes blurry, the appropriate measure is one of risk and reward. It is the risky but potentially broadly valuable investments that should be classified as research.

NITRD program coordination would be improved if the participating NITRD agencies were required to report their R&D expenditures more clearly. To coordinate research activities, actual *research* investments must be reported. Either better categories such as the ones I've outlined or more thorough line-item reporting would help. This is an area where a bill such as H.R. 2020 could contribute.

Conclusion

The NITRD program has demonstrated an ability to coordinate Federal investments in essential research, starting with high-performance computing and now extending to a broader set of national goals. The challenge now, for sponsors and researchers alike, is to make the case to an increasingly broad set of NITRD mission agencies that long-term investments in fundamental NITRD research lead to large rewards for their missions and for the nation.

Interplay of university research, industrial research, and development for IT in the US

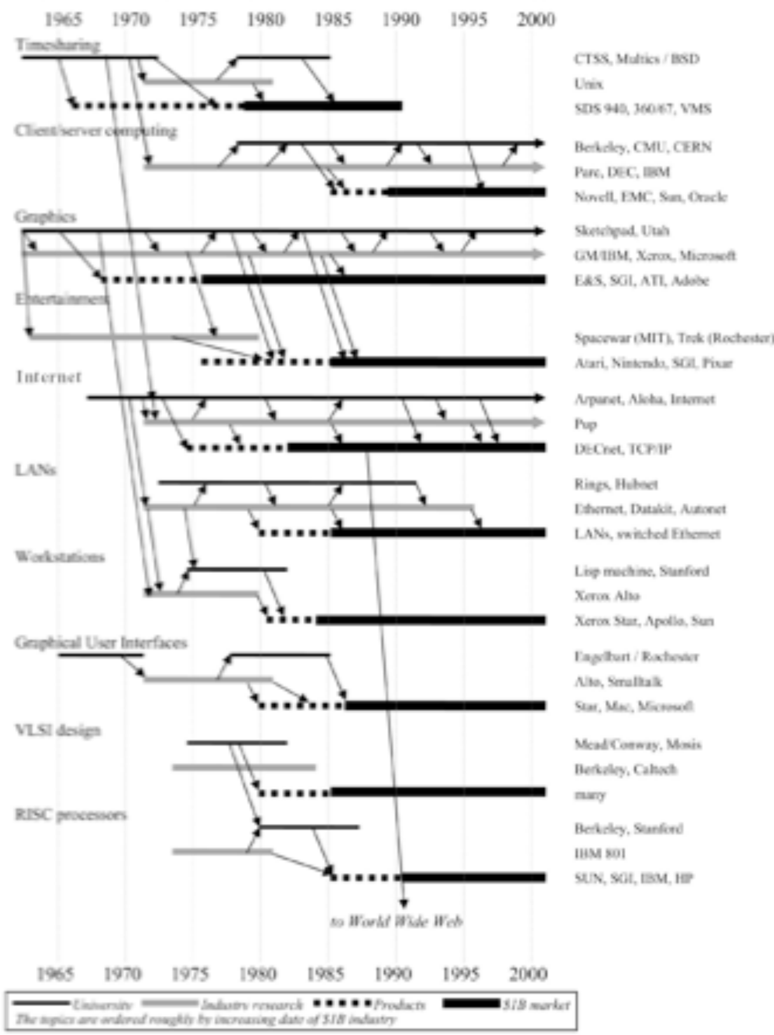

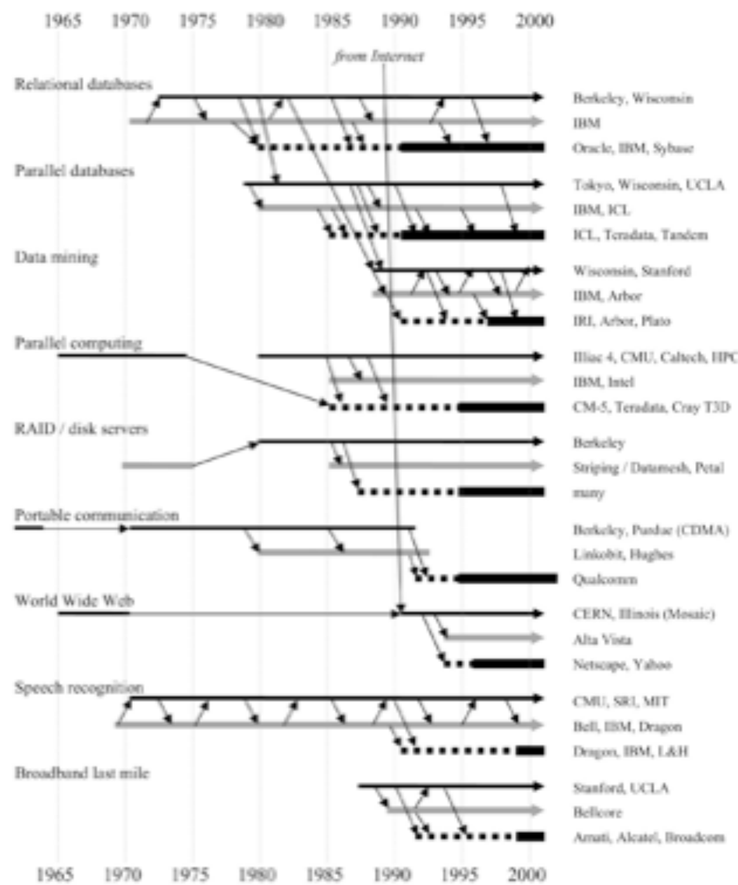

National Research Council, 2003