Intel Research Berkeley Collaborating to Change the World

Introduction

On the 13th floor of the PowerBar building in downtown Berkeley, California, researchers are exploring new ways to deliver energy and meet performance challenges. But these scientists aren't focused on the latest protein bar or sports nutrition drink. They are designing a novel networking infrastructure that will enable computing and communications to reach remote regions of the developing world.

Welcome to Intel Research Berkeley—one of three labs in the Intel Research Network, an innovative approach to industry-university research pioneered by Intel to enhance and accelerate long-term research. Researchers flow freely between the lab and UC Berkeley, conducting joint research projects under a model designed to foster collaboration and generate breakthrough results. The lab is exploring a range of research projects in four areas: technology development for emerging regions of the world, systems and networking, programming languages, and the user experience. Since its opening in November 2001, Intel Research Berkeley has made an enormous impact.

> "It's always great to have a lab with first-rate researchers close by; it provides more opportunities for collaborations and for interesting work. We promote the fact that Intel Research Berkeley is close to campus when recruiting faculty and students. The lab is a big asset to us."
>
> *Ion Stoica*
> *Associate Professor of Computer Science, UC Berkeley*
> *Faculty Affiliate Researcher, Intel Research Berkeley*

The Open and Collaborative Research Model

The open and collaborative research (OCR) model practiced at Intel's university labs was designed to eliminate the conflicts over intellectual property (IP) rights that hinder or prevent many collaborations between companies and universities. "Our model provides a certain freedom that you can't get with a more proprietary approach," says Eric Brewer, Director of Intel Research Berkeley. "Research can take its natural course, since researchers don't have to worry about IP. A more closed model, even if it's well intentioned, slows down the research process. And it can produce a chilling effect: just worrying about potential violations of IP rights prevents work from getting done."

Under the OCR model, Intel's university labs operate on the principles of collaboration and non-exclusive IP rights. Intel owns and funds the labs, but much of their research is published and widely shared. Patents may be issued but are expected to be rare, because of the OCR model's emphasis on collaboration, not competition. "The fact that Intel stresses sharing of results makes it very easy to collaborate with our lab," says Brewer.

> "Initially, I didn't realize the extent to which Intel was willing to pursue open research. Because of its open and collaborative model, Intel was able to recruit very high-level academic researchers as lab directors. This has enabled the labs to capture a much larger mind share of academic departments than any of the other research groups that I've seen. That, to me, was the brilliant insight."
>
> *Scott Shenker*
> *Professor of Computer Science, UC Berkeley*
> *Group Leader of Networking, International Computer Science Institute (ICSI)*

The idea of foregoing IP rights is difficult for some companies to embrace, but the benefits are significant, according to Allan Knies, Associate Director of Intel Research Berkeley. "Intel had an intellectual breakthrough when it embraced the OCR model for our university research labs," he says. "Adopting an open IP model requires a strength of conviction that few companies can muster. The agreement has allowed us to work with visionary researchers in the UC Berkeley community. That would not be possible if it weren't for Intel's willingness to share IP in order to advance long-term, exploratory research."

Profiting from Proximity

Both Intel and the university benefit from the proximity of Intel Research Berkeley to the UC Berkeley campus. "We have a large number of students flowing back and forth between the lab and the university," says Brewer. "It's this flow of students that's the most critical form of sharing. The primary means of collaborating with the university is via students; they're the mechanism that really makes it work."

Working with researchers across the street from campus can be a richer experience than collaborating with colleagues across the country or even across town. The interchange of ideas is more fluid; student researchers can easily stop by the lab to ask questions, discuss a difficult research problem, or work on a project. Lab researchers can meet with their university colleagues on campus or at a local café to talk about their work. These casual interactions, which require close physical proximity, make the collaboration much more interactive, relaxed, informal, and productive.

Applying the Resources of Intel

For UC Berkeley researchers, another benefit of the Intel lab is the ease with which research can be initiated, according to Knies. "One of the key benefits of our lab is that we can quickly move budgetary resources into new areas," he says. "We have the ability to support students, provide research grants, and add people to projects without having to spend a year in planning. That allows for innovative ideas to ramp quickly, and with less overhead than researchers would otherwise be subject to if they had to go through a traditional planning or grant-writing process."

Intel Research Berkeley also makes it easier for the university to scale its research by providing access to the resources of Intel. "There are certain projects that need scale if they are to achieve their goals," says Brewer. He points to the TIER project, which aims to bring technology to developing regions, as one example. "The work that we are doing at UC Berkeley and our lab on technology for developing regions is helped tremendously by the fact that Intel has an entire platform dedicated to emerging regions," he says. "We are also working with product groups within Intel that care about developing regions, to learn from them about the needs of people in those regions and provide them with new technology ideas. That kind of access to Intel is critical, and it's the key reason I was excited to accept the role of lab Director. Having the resources of Intel behind us enables us to go beyond ideas to really influence the ecosystem and products that are available in developing regions."

Access to the Best and Brightest

In exchange for providing resources to UC Berkeley, Intel gains access to some of the best and brightest faculty and students in the world. "We're working with talented young collaborators who might one day become our employees," says Brewer. "By having a lab near campus and collaborating with the university, Intel significantly enhances its visibility and prestige with young researchers. They start to think of Intel not as some company off on the horizon but as those folks down the street who are doing some leading-edge research. That's a huge factor in enhancing Intel's reputation and ability to recruit top talent."

The value to Intel goes beyond having access to top university researchers who may become employees one day. Having labs located near leading universities also helps the company stay up to date on the latest academic research developments. "When you have collaborations going on with three of the best universities in the world—and Intel has strategically located its labs close to three top schools—you will hear about every important new research concept," says Brewer. "Either someone at the university will be working on new research that's of interest to Intel, or some visiting professor will be giving a talk on it. So you have access to every big idea that comes along."

> "We are incredibly fortunate to be able to work with some of the best students in the world at UC Berkeley. Because we have spent years building great relationships with the faculty, we have a steady flow of top-notch student collaborators coming through the lab, creating excitement, and helping our own full-time staff to reach critical mass on their research."
>
> *Allan Knies*
> *Associate Director, Intel Research Berkeley*

That exposure to leading-edge academic research also helps to ensure that Intel's university labs will not become insular. "One of the real dangers of industrial research labs is that, despite everybody's best effort to the contrary, they become isolated and self-reinforcing," says Scott Shenker, Professor of Computer Science at UC Berkeley and Group Leader of Networking at the International Computer Science Institute (ICSI). "The Intel labs, by being small and so closely connected to the university, have a constant influx of new ideas and new talent. That 's one aspect of the labs that I think is really crucial."

> "The positive impact on the wider research community of the work done at the Intel lab in collaboration with UC Berkeley researchers has been dramatic, particularly considering (or perhaps because of) the lab's size. It's a small, close-knit team that operates through wide-ranging collaboration."
>
> *Timothy Roscoe*
> *Professor, Department of Computer Science, ETH Zürich*

Lab Leadership

In June 2005, after completing a two-year rotation as Director of Intel Research Berkeley, Joe Hellerstein returned to his tenured position as Professor of Computer Science at UC Berkeley and Eric Brewer took the reins of leadership. Under Brewer's direction, the lab has continued to build its reputation as a place where groundbreaking research is performed, according to Shenker. "Intel Research Berkeley, under the leadership of David Culler, then Joe Hellerstein, and now Eric Brewer, has provided a set of tremendously exciting and world-changing projects," he says. "The intellectual depth of the ideas involved, their potential to change the world, and the level of resources Intel has made available for these projects has been extraordinary."

About Eric Brewer

Eric Brewer became the Director of Intel Research Berkeley in June 2005, taking a leave of absence from his position as Professor of Computer Science at UC Berkeley to assume leadership of the lab. Brewer's research focuses on all aspects of Internet-based systems, including technology, strategy, and government. As a researcher, he has led projects on scalable servers, search engines, network infrastructure, sensor networks, and security. His current focus is (high) technology for developing regions, with projects in India, Cambodia, Mexico, Sri Lanka and Bangladesh (so far), and including communications, health, education, and e-government.

In 1996, Brewer co-founded Inktomi Corporation with a UC Berkeley graduate student based on their research prototype, and helped lead it onto the Nasdaq 100 before it was bought by Yahoo! in March 2003. In 2000, he founded the Federal Search Foundation, a 501-3(c) organization focused on improving consumer access to government information. Working with President Clinton, Brewer helped to create FirstGov.gov, the official portal of the Federal government, which launched in September 2000.

Brewer received an MS and PhD in EECS from the Massachusetts Institute of Technology, and a BS in EECS from UC Berkeley. He was named by the World Economic Forum as a "Global Leader for Tomorrow," by The Industry Standard as the "most influential person on the architecture of the Internet", by InfoWorld as a top ten innovator, by Technology Review as one of the 100 most influential people for the 21st century (the "TR100"), and by Forbes as one of the magazine's 12 "e-mavericks," for which he appeared on the cover. In February 2007, Brewer was elected a member of the National Academy of Engineering (NAE), among the highest professional distinctions accorded to engineers, in recognition of his design of Highly Scalable Internet Services.

About Allan Knies

Allan Knies is a Principal Engineer and Associate Director of Intel Research Berkeley. For his first nine years at Intel, Knies was a contributor to understanding the capabilities of the Itanium® Processor Family architecture, its performance, and associated compiler technology. He represented Intel on the joint architecture committee with Hewlett-Packard and prototyped a number of advanced technologies and tools for Itanium. Prior to joining Intel Research, Knies was leading a research team working on programming languages and compiler technology. His most recent research is aimed at producing an extremely high efficiency many-core IA architecture.

Knies has published several papers on the Itanium® architecture and performance evaluation and has given talks at numerous customer sites, academic institutions, and conferences. He has one patent and several other filings relating to architectural enhancements of IPF.

Knies earned his Bachelor of Science in Computer Science from Ohio University in 1988, then went on to Purdue University to earn his Masters in Computer Science in 1990 and a Ph.D. in Computer Engineering in 1995.

Research Agenda

Under Intel's model, when a university research lab transitions to a new director (every two or three years), roughly two-thirds of the research agenda remains the same, while the new director sets the remaining third of the agenda. The new component of the research agenda that Eric Brewer brought to the lab reflects his strong interest in technology for the emerging regions of the world.

Technology and Infrastructure for Emerging Regions (TIER)

The TIER (Technology and Infrastructure for Emerging Regions) project is a multidisciplinary effort that aims to make technology accessible to emerging regions by developing hardware, software and infrastructure that's explicitly designed to address the physical, political and economic realities of those regions. The primary funding for the project comes from the National Science Foundation.

Intel is an active participant in TIER, and one of several companies providing additional financial support. Following is an overview of TIER-related research underway at Intel's Berkeley lab. The focus of the Intel research is on developing novel wireless networking solutions.

Delay Tolerant Networking (DTN)

Today, most networking technology is designed to provide continuous connectivity and Internet access. In very remote areas and many portions of the developing world, this may be very expensive or simply impractical due to the lack of continuously available power and networking infrastructures.

The goal of delay tolerant networking (DTN) is to provide reliable, secure, asynchronous communication for difficult to reach places that may only have intermittent connectivity. DTN is being designed for application in a wide range of environments, including rural villages in developing countries.

Through the Delay Tolerant Networking Research Group (DTNRG), a research group associated with the Internet Research Task Force (IRTF), Intel researchers and other government and industry researchers are working together on an alternative to the standard Internet TCP/IP communications protocol that is better suited to the communications challenges of environments where connectivity may be intermittent and where a choice of communication technologies may be available. The alternative uses hop-by-hop storage and retransmission to form a delay-tolerant, heterogeneous overlay network, and provides a messaging service interface that is conceptually similar to email but is application-independent and supported by specialized reliability, routing, and security capabilities. DTN can work in conjunction with TCP/IP in environments where connectivity is good, but can also operate in situations where TCP/IP is not available.

Implementation and wide deployment of DTN would represent a significant step forward in communications capability for remote areas and developing regions. DTN enables a range of applications to be used in environments with poor connectivity-from email and voicemail to offline search engine queries, electronic form filling, and "instant-enough messaging"-at a reduced cost.

Rural Communications Platform

There's a tremendous need for affordable, robust wireless networking in rural areas of the developing world. Existing infrastructure that is useful in high-density areas, where costs can be spread over many users (notably, cellular telephony), is not appropriate for sparsely-populated rural areas.

Intel researchers are addressing the challenge. Point-to-point long-distance wireless links between villages, augmented with a number of broadband satellite or wired Internet connections, seem to be one of the best connectivity choices.

A key component of this solution is a variation on IEEE 802.11 (Wi-Fi) technology. Low-cost Wi-Fi technology could form the backbone network in rural regions. While the 802.11 Media Access Control (MAC) protocol was not designed for long-distance communications, researchers believe that modifying the MAC layer should resolve the problem, without the need for hardware modifications.

Researchers are assessing the performance of point-to-point 802.11 long-distance links in a wireless testbed in and around Berkeley, California. They also have deployed several test links in India and more recently, in Ghana, where they have demonstrated a bandwidth of 5-7 Mbps over a relatively long distance: 10.5 kilometers. Researchers will continue testing and modifying the technology in additional deployments in India in 2007.

Creating a platform to demonstrate the viability of this approach to long-distance communications in rural regions requires inputs from multiple disciplines. As a result, researchers at Intel's Berkeley lab will collaborate with other researchers from Intel, UC Berkeley, the University of Washington, and UC Riverside.

Steerable Antennas

Long-distance wireless applications typically use directional antennas that provide the high gains necessary for communicating over great distances. However, the deployment of conventional directional antennas, such as parabolic reflectors, presents practical problems with initial installation and maintaining optimal alignment with remote counterparts over time.

Intel researchers are exploring the design and implementation of low-cost, electrically steerable antennas that provide high gain and directionality combined with the ability to steer the antenna pattern in both azimuth and elevation. In effect, these antennas combine the best aspects of omni-directional and directional antennas, and provide capabilities not previously available.

The researchers currently are focusing on the construction of controllable impedance reflecting surfaces that may deliver in excess of 24dBi of gain with steering across more than 45 degrees of azimuth and elevation. Early prototypes have demonstrated the radiophysical principles of operation with Wi-Fi transmissions and the feasibility of integrating these antennas with devices such as the Rural Communications Platform. Researchers will continue developing prototype systems for measurement and pilot deployments in the field while simultaneously seeking the lowest possible cost designs and implementations.

The multidisciplinary aspects of this project, combining elements of antennas design, high performance mathematical optimization, and systems integration require a diverse team. As a result, the research is being carried out in collaboration with Intel Russia and the Radiophysics Department of the Nizhny Novgorod State University.

P2: Declarative Networking

Any widely distributed system must track its participating nodes and must be able to send messages among those nodes. This capability can be thought of as an overlay network, since it provides an application with customized networking functionality (naming, topology, and routing) that runs as a layer over traditional IP networking. Such overlays are in widespread use, including in commercial mail/directory servers, application-level multicast systems, and distributed hash tables (DHTs).

It is tricky to design, build, and deploy an overlay network for a particular application and environment. To ease the process, researchers from Intel Research Berkeley and UC Berkeley have collaborated to implement P2, a system that uses a high-level declarative language to express overlay networks in a highly compact and reusable form. P2 has been used to specify and execute working, detailed overlays in a tiny fraction of amount of code used in traditional implementations. P2 automatically compiles high-level specifications to a dataflow-oriented runtime system, which can itself be used by expert programmers to optimize overlay efficiency and performance.

P2 is part of a broader effort to revisit networking technology through the lens of distributed database query processing. The P2 research could provide simpler, safer specifications for network protocols and-within the same framework-the ability to query, monitor and control all aspects of the network's distributed state.

Enterprise Network Security

In a typical enterprise network, all end hosts are configured identically. While this makes it easier to manage large enterprises, it is not optimal from a security standpoint. That’s because anomaly detection systems in enterprise networks often use pattern recognition techniques to identify abnormal activities that could signal a network attack, but what constitutes “normal” or “abnormal” varies from one user to the next.

Intel researchers are exploring the use of personalized security configurations as an alternative approach. They are developing profiling techniques that capture a broad range of user behavior in a compact form. The resulting profiles will be used to customize security configurations for each individual. The researchers’ hypothesis is that personalizing anomaly detectors will make it harder to implement a widespread attack on a network, since attackers would have to know the unique security configurations of each networked machine.

Another aspect of the research focuses on designing anomaly detection solutions for large-scale distributed monitoring systems. Under the current approach, network monitors located throughout the enterprise network generate continuous streams of data from every node in the network. The data is sent to a centralized location for analysis and detection of anomalies that could suggest that a worm or virus has infected the network.

This approach suffers from two key limitations. First, it can take several minutes to detect an anomaly. Secondly, the hundreds of monitors in a large enterprise network generate a massive amount of data that is forwarded to a central location for analysis, clogging the network and wasting valuable resources.

Intel researchers are exploring an alternative approach that will not require monitors to send such a high volume of data to a centralized location for analysis. Different algorithms are used to detect different types of security breaches. By understanding the network conditions that are indicative of each type of security breach, monitors could be programmed to send data only when the network conditions approach a certain “danger zone.” The researchers’ approach seeks to detect attacks on sub-second time scales via continuous event tracking.

IVY

Software has become embedded in our everyday lives, used in arenas ranging from shopping and banking to healthcare and government. Unfortunately, much of this software contains bugs and vulnerabilities, exposing users to new risks and deterring people from using software to improve their lives further. Moreover, much of the software we rely on today are legacy programs written in “C”—a language that is 30 years old and that makes it difficult to create secure software applications.

The Ivy project is striving to make it practical for developers to improve the security of their existing legacy C code. Ivy is a set of extensions and modifications to the C language that provide features such as type safety, verified memory management, and safe concurrency.

The primary goal of the project is to provide practical tools that can be used by any C developer and project. To that end, Ivy is simple, can be implemented incrementally, and without affecting the integrity of the original C program: Projects can be ported to Ivy one line at a time. Porting requires no skills beyond those expected of typical C programmers. And programs do not become dependent on Ivy; they continue to be valid C programs independent of Ivy.

To demonstrate the practicality of applying Ivy to large legacy code-bases, the Ivy team has recently been working to guarantee type correctness and correct memory management for the Linux kernel.

Mobile Times

The goal of the Mobile Times Project is to identify new opportunities for mobile technologies, based upon a deeper understanding of time as an aspect of everyday life. The project researchers will use comparative ethnology to develop an understanding of time, exploring diverse sites and populations such as bike messengers, people living green lifestyles, busy families, and freelancers to develop meaningful cultural models of time. They also will use “ethno-mining,” a novel mixed-methods approach that combines ethnographic and data mining techniques. This will enable the project to bridge radically different kinds of data, information, knowledge and epistemologies to inform an understanding of time.

The resulting models of time will deliver insights on how time impacts consumer perception and adoption patterns of mobile devices, as well as suggest opportunities for technological innovation. Mobile Times is a collaboration between Intel Research Berkeley and the People and Practices Research Group within Intel Research.

Continuous Monitoring (CoMo)

Developing new tools to analyze network data often is a complex and error-prone process. Current practices require developers to possess an in-depth understanding of the original data sets and to develop ad-hoc software tools to extract the relevant information from the data, then implement the internals of the new algorithm. This development process results in long delays during the analysis of the data and in the production of software, which is often hard to reuse, debug or optimize.

Intel researchers have developed the Continuous Monitoring (CoMo) system, which provides an abstraction layer both for the network data and for the hardware architecture used to collect and process the data. This allows developers to focus on the correctness of the implementation of their analysis tools while the system makes the tool amenable to optimization when running on different hardware architectures.

CoMo aims to become the building block for a network monitoring infrastructure that will allow researchers and network operators to easily process present and past network traffic data over multiple sites. Individual systems with CoMo software installed would operate on backbone high speed links, edge access links and wireless access points, and would provide a unified interface for network queries. The systems would then cooperate to rapidly disseminate queries throughout the network, allowing network operators to "drill down" to relevant data locations, to rapidly identify and respond to attacks or investigate past anomalies.

Beyond Research: Advancing Education

The close connection between Intel Research Berkeley and the university next door extends beyond research to education. "We have tried to create an environment such that the lab feels like it's an extension of the Berkeley campus," says Brewer.

To that end, the lab regularly hosts seminars, at the lab or on campus, that are attended by Berkeley faculty and students. In addition, some of the lab's researchers teach courses at UC Berkeley. "Graduate students often move into research as a result of taking advanced level courses, so having Intel lab researchers teach at UC Berkeley complements the lab’s research mission,” says Edward Lee, Chair of the Department of Electrical Engineering and Computer Sciences at UC Berkeley. “The lab offers our students and faculty exposure to cutting-edge industry research, and an opportunity to collaborate on that research within an open environment.”

A Crossroads for Computer Science Research

Eric Brewer is pleased with how Intel Research Berkeley is evolving and excited about its prospects for the future. "The lab has become a crossroads for computer science research," he says. One indication of the lab's strong reputation for research is the fact that two leading postdoctoral students postponed accepting faculty positions to spend more time working at Intel Research Berkeley. Sam Madden spent six months at the lab before joining the faculty of MIT and Carlos Guestrin decided to work for a year at the lab before accepting an Assistant Professor position at Carnegie Mellon.

"The Intel lab is one of the best places in the world—if not the best—for research in sensor networks," says Guestrin. "When I was offered the opportunity to work there, I was thrilled. The lab has given me the opportunity to explore the interaction between sensor networks and machine learning. This unique learning opportunity, along with the possibility of interaction with Intel researchers and UC-Berkeley students and faculty, was irresistible and has been extremely productive."

Many lab researchers who move on to new positions continue to collaborate with Intel Research Berkeley. "I loved the four years I spent as a Principal Researcher at Intel Research Berkeley, and I'm proud to retain an association with the lab,” says Timothy Roscoe, who is now a Professor in the Department of Computer Science at ETH Zürich. “The combination of industry resources, the openness of the university relationship, and the energy brought by the students lead to a ‘friction-free’ research atmosphere that it is unparalleled in my experience—fun, stimulating, and highly effective.”

Each year, Intel Research Berkeley showcases the results of its collaborative approach through an open house. The 2006 event provided more evidence that Intel Research Berkeley has become a computer science crossroads. A mix of visitors that included academic researchers and members of the Bay Area’s technical and venture capital communities mingled with researchers and viewed approximately 20 demonstrations of research projects underway at the lab. “I think the most important thing about the open house is that it makes people realize how broad the research agenda is at Intel,” says Brewer. “It shows that we are exploring not only classic technologies like networking and systems, but also innovative areas such as urban computing and technology for developing regions. That’s what makes Intel such a great place to do research; there are always some exciting projects underway.”

Looking Ahead

From the viewpoint of UC Berkeley, Intel's experiment in establishing a new model of industry-university collaboration has thus far been a success. "Intel's open and collaborative research model has set a new standard for university-industry research," says Scott Shenker. "Being very open and public about the agreement between Intel Research Berkeley and UC Berkeley not only helped the relationship between Intel and the campus; it also gave the campus a template that it can show to other companies who want to collaborate, as a starting point for the discussion. I think that openness helps to preserve the integrity of university research."

Under Brewer's direction, Intel Research Berkeley is expanding its relationship with the university next door, Allan Knies notes. "As our lab has become a fixture near campus, our set of collaborators has expanded dramatically," he says. "Of course, we have many ongoing projects with engineering and computer science faculty, but we are also a working with the Architecture Department and campus research centers such as CITRIS. All of this relationship building provides a foundation for even greater collaboration with the university in the coming years."

> "When I returned to Berkeley after more than 30 years in industrial research, I was struck with one of the most important attributes of the Berkeley Engineering faculty: their commitment, through teaching and research, to having a profound impact on science, industry, and society. What a wonderful thing for Intel to do to make that process so much easier through its laboratory in Berkeley! The relationships among the Intel researchers, the faculty, and especially the students have exactly the kind of 'enhanced intellectual transfer' effect that benefits all of those involved and society in general."
>
> *Gary L. Baldwin*
> *Executive Director, CITRIS*
> *The Center for Information Technology Research in the Interest of Society*

Brewer's goal is to continue building the relationship with UC Berkeley, and to sustain the lab's record of success while evolving his research agenda for the developing world. "By the end of my rotation as Director of Intel Research Berkeley, I hope that Intel will have a strong focus on technology for developing regions, and maybe even have products in the pipeline that came out of our lab's research," he says. "Our ultimate goal is to change the world with our research, and in my view, Intel is a big lever to help make that happen. If I can use that lever effectively while I'm the lab director, if our research can help to bring affordable computing and communications technology to the developing world, that would be incredibly rewarding."