Intel Research Expands Moore's Law

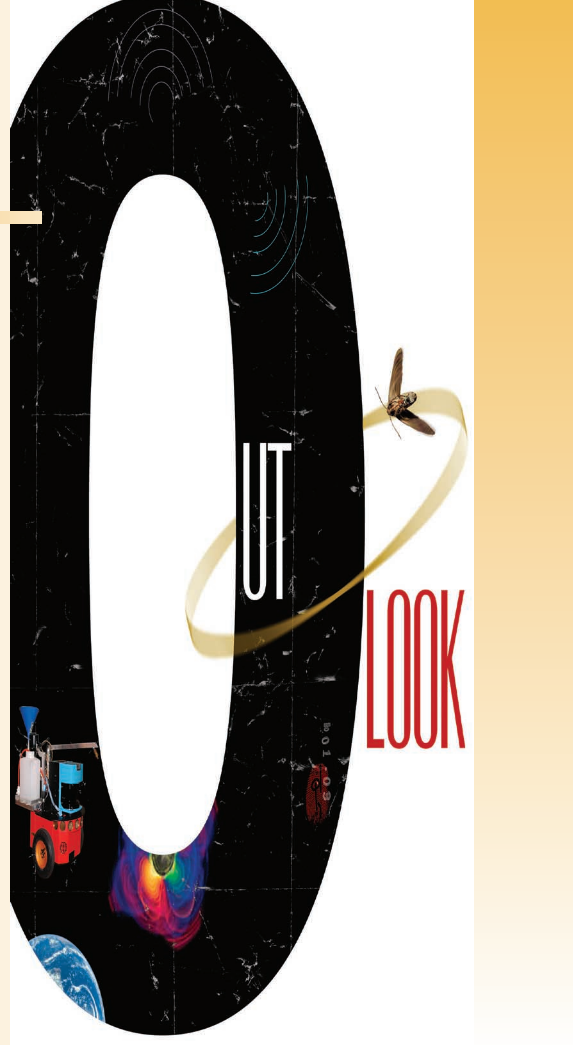

Intel's lablet network leverages industry and academic synergy to nurture off-the-roadmap ideas and technologies and provides a proving ground for testing their viability.

Scott Hamilton
Computer

A funny thing happened on the way to the forum—the Intel 2002 Developer Forum—this past September: The hoopla at the final keynote wasn't reserved for Itanium or Madison, but for Intel Research's vision of proactive computing.

During the keynote, Intel VP and CTO Pat Gelsinger shared center stage with Sunlin Chou, senior VP and general manager of Intel's Technology and Manufacturing Group. The pair introduced a research agenda that includes Intel's 90-nanometer fabrication processes, scheduled for a 2003 debut with Prescott, and cutting-edge research in extreme ultraviolet lithography that will help extend Moore's law. However, the presentation showcased Gelsinger's vision of an *expanding* Moore's law, as exemplified by *disruptive technologies* such as

- MEMS microradiators for cooling chips and MEMS-based smart antennas and RF components for switches, resonators, and filters that—for certain devices like cellular handsets—would be too high powered to implement in normal silicon;
- ad hoc sensor networks with wireless communications; and
- photonic devices such as optical switches and cheap tunable lasers intended ultimately for chip-to-chip communication.

Obviously, Intel has no plans to relinquish its lead in either microarchitecture and circuit design or silicon process technology and manufacturing, to which it devotes most of its $4 billion research budget. Indeed, Intel highlighted its new 90-nanometer manufacturing process technology using silicon germanium, which will be introduced when the retooled fabrication plant in Hillsborough, Oregon, goes into full commercial production this year.

Yet Gelsinger's keynote sent a strong message that Intel is seriously championing emerging technologies that could advance Intel's business by creating new markets and opportunities or potentially disrupt it if left unanticipated or ignored. According to Gelsinger, Intel's vision for computing's future foresees "a world where people no longer conform to computers but computers conform to people and how they operate so that computing is integrated into every aspect of our lives."

Intel's demonstrated mastery over Moore's law lets it explore and develop niche technologies at an affordable price, thereby increasing demand and ultimately making these technologies widely available to new markets, products, and user populations. Ad hoc networks with wireless sensors that sell for pennies will make possible novel services and applications for the home, the office, the factory floor, healthcare, and the environment. These net-

0018-9162/03/$17.00 © 2003 IEEE Published by the IEEE Computer Society

Corporate Technology Group

From the start, Intel's management—especially Gordon Moore, who had been head of R&D at Fairchild—insisted that R&D spending have clear commercial goals. Intel kept research decentralized in a distributed internal network of labs connected to, and funded by, the appropriate business group where their research was most likely to be transferred. Thus, the silicon and transistor research labs are in Sunlin Chou's Technology and Manufacturing Group, and the microarchitecture and processor labs are in Justin Rattner's Microprocessor Research Lab.

Moore's experience at Fairchild taught him that the hardest part of R&D was the transfer of new ideas from labs into the product groups. Over time, Intel's existing labs have become very good at turning research into profitable tools, processes, and products. One exploratory research lab delivers its best ideas into an advanced technology lab, which passes prototypes to a products lab, which then hands over the designs to manufacturing. However, this very efficiency makes it difficult for the business labs to devote the funds and long-term attention necessary to find and nurture new applications or truly disruptive, breakthrough technologies. This was why Intel created the Corporate Technology Group and, most recently, Intel Research.

CTG focuses on platform technologies, standards, and new applications and works to develop new core technologies to fuel industry growth: Ultrawideband, Universal Serial Bus 2, and Intel's new portable video player concept are examples of CTG programs.

CTG is the organization that works the most with the industry and other external industry players, and its various advanced development labs, such as the Emerging Platform Laboratory, are developing products in the three-to-five-year timeframe. Intel Research, on the other hand, pursues research that might bear fruit in three to ten years.

works will generate huge volumes of data and information for data centers and desktops—further fueling Intel's core businesses. Or at least Intel is betting that they will.

INTEL RESEARCH AND PROACTIVE COMPUTING

To help set and realize this new research agenda, Intel hired David Tennenhouse in 1999 to found Intel Research. Previously chief scientist and director of DARPA's Information Technology Office, Tennenhouse clearly sees proactive computing as the driver for IR's agenda: For 40 years, the IT and research communities have worked successfully to fulfill J.C.R. Lickleiter's vision of interactive computing. But, according to Tennenhouse, much has happened since we achieved one computer per person, and now that we have five or 10 per person, we still haven't moved away from the direct interactive computer toward proactive computing as the new personal computing paradigm.

How do you keep the individual empowered but make computing proactive? What do we need to develop now to prepare for the day when there are 1,000 smart devices per person so that the devices can all work together to proactively serve us? In short, how do you prevent computer science from becoming a mature field nearing the end of its developmental road—the analog of Detroit worrying about bells and whistles instead of improving the engine—and push it to the next level by an order of magnitude or two?

Intel Research is a small group of researchers and staff, described in the "Corporate Technology Group" sidebar. Of the current 90 Intel Research employees, only a few work full time as dedicated management and staff housed at the corporate headquarters in Santa Clara, Calif. Most work in lablets—Intel-owned and funded university laboratories—or on projects in labs strategically placed in the various business units distributed geographically across the company. IR's charter is to transform emerging and disruptive technologies into products within three to 10 years and to perform strategic exploratory research that looks five to 10 years down the road.

As Intel's VP and director of research, Tennenhouse performs two roles:

- he functions as champion of disruptive research throughout the company, identifying technology needs and options and lobbying senior management for the resources to realize these research capabilities—either internally within the product development groups or externally through university research grants; and
- he initiates exploratory strategic research projects (SRPs) within Intel Research and through the Intel Research Network of Intel-funded and wholly owned labs adjacent to universities.

The research agendas for Intel's Architecture Group and its Technology and Manufacturing Group are firmly in place for the next eight years and the next several generations of processors and processes. Tennenhouse, on the other hand, is more interested in research at the 10-nanometer scale and is investing in the future through Intel's approximately 250 university research grants, making sure Intel has a good set of viable options and the core talent needed to implement them in the next 10 years.

Similarly, when it identifies a technology area in which Intel must increase its investment, Intel Research will nurture a needed competency until it can find a home for it in one of Intel's business units. One example of this approach is Radio Free Intel. Tennenhouse and others determined that the company would at some point require an RF competency, which none of Intel's labs then supported. Intel Research started three SRPs and used them to build up an RF competency. A year and one-half later, the company reached the conclusion that it would need an RF facility, and Intel Research transferred those projects and resources to instantly establish a full-fledged RF lab.

Intel Research has primary responsibility for initiating and monitoring long-range strategic research through the university and industry grants the Intel Research Council awards, including long-term research in its core competencies. But IR is more closely focused on *off-the-roadmap* research intended to discover new research areas of interest that Intel is not competitively pursuing. To do so, IR invests in its own strategic research projects and expects them to deliver results in three to five years, and it also invests in the Intel lablets, which target delivery in five to seven years.

INTEL RESEARCH NETWORK: THE LABLET MODEL

In early 2001, Intel began establishing its innovative Intel Research Network (IRN) of wholly owned and funded university labs to address the challenges of a future world where billions of small, affordable, and connected computing devices, sensors, and actuators will be embedded within our physical environment and work proactively on our behalf.

Intel positioned each of these lablets adjacent to a university. Each facility employs up to 20 Intel researchers and an equal number of student interns, graduate students, faculty members, and visiting researchers. The lablet model's purpose is to foster long-term, open collaboration between industry and academia—in this case Intel and its university partners. This collaboration ensures that Intel has access to new ideas, new technologies, and well-trained researchers. But the concept also, and more altruistically, feeds knowledge back into the academic ecosystem by letting university researchers undertake projects on a scale heretofore unimaginable.

Defining the model

Intel selected the four current university partners—the University of California, Berkeley; the University of Washington; Carnegie Mellon University; and Cambridge University—for their academic reputation and long-standing collaborative research relationships with the company. Each lablet has a well-established, tenured faculty member on partial leave as an Intel employee who serves as lab director and largely defines the lab's research agenda. The director rotates out after two or three years to keep the research agenda fresh.

The Intel Research codirector's role is to provide for the lab's long-term stability and facilitate a strong upstream and downstream communications channel between the lablet, IR, and, ultimately, the various product groups. Intel technologies and marketing plans must flow into the lablets, and research advances must funnel back into Intel's advanced development labs and product groups.

> **Intel Research has primary responsibility for initiating and monitoring long-range strategic research.**

The university side provides built-in momentum that can vitalize the industrial side, while the industrial side provides stability, attention span, and the ability to follow through by tapping into additional manufacturing or engineering resources.

IRN's novelty, however, stems from the open collaborative process in which Intel and its university partners can explore futuristic technologies, widely disseminate the research through publication, and release those concepts to the research community—most often in open source form. Both UC Berkeley and Intel Research Berkeley (IRB), for example, make all the hardware design specifications for their mote project freely available even while contracting with Crossbow Technology, a supplier of inertial sensor systems, to produce them. Even more surprising, Crossbow has followed suit by making its extensions to the Berkeley mote architecture openly available as well.

Berkeley has a long history of open source development, beginning with BSD Unix, Spice, RISC, and RAID. As David Culler, professor of computer science at Berkeley and director of IRB, states, "Historically, especially in the EECS, we pride ourselves on the multibillion-dollar industry that has grown up out of the research ideas—the workstations, the PCs, the enterprise systems—much more than on the individual startups. The department culture has been one of putting the ideas with potentially big impact out in the public domain."

Indeed, Berkeley's EECS culture anticipated recent studies by the University of Texas's Strother Moore[1] and Berkeley's former dean, Dave Hodges,[2] which have overturned the popular argument that computer science intellectual property generates substantial revenues for universities. These studies show that the majority of licensing income for universities resides in single, clearly definable patents for biotechnology, agriculture, and medicine.

Carving the pie. UC Berkeley's long-standing tradition of open development enabled Intel and the university to negotiate an Open Collaborative Research (OCR) agreement—with considerable input from the three lab directors—that was acceptable to both parties and, subsequently, to the University of Washington and Carnegie Mellon as well. The current IP agreement establishes four principles:[3]

The OCR defines what research will be conducted in the open and clearly specifies and retains the IP rights of both parties.

- collaboration—not just throwing money over the wall—should be the norm;
- the IP of this collaborative research should be nonexclusive;
- many areas of collaborative research favor not filing for IP protection; and
- proprietary advantage should be generated when the principals take ideas downstream through, for example, internal R&D programs.

The OCR defines what research will be conducted in the open, identifies which researchers will be involved, and clearly specifies and retains the IP rights of both parties—though for the most part granting nonexclusive access to IP both for participants and third parties.

The addendum deals with how credit and royalty revenues are apportioned to the partners for inventions that derive from their collaboration. It clearly spells out the conditions under which an invention is considered to belong to Intel, the university, or both partners jointly. The addendum also gives the university the right to charge royalties on inventions that derive primarily from its researchers' efforts, although such charges must be reasonable and nondiscriminatory and Intel may have already paid them fully and in advance by contributing facilities to the project.

The addendum can also restrict the direction that collaborative research can take. So, for example, even though Intel and a given campus may be involved in a sensor net project, Intel may not be allowed access to university research in other fields such as biotechnology. Thus, each specific project pursued jointly by Intel and a university works from a research project document.

This document clearly defines the project's scope of work and the open content boundaries for the use and publication of any knowledge gained through the project. The document also spells out what content is nonexclusive and royalty-free, not just to the two primary collaborators, but to third parties as well.

Everybody wins. What does Intel get out of this approach? It will compete on the downstream translation of the product: Rather than trying to retain the initial idea so that nobody else has it in an early form, Intel agrees to let it be open so that the knowledge thus released can foster a rich industry. Intel expects to keep its competitive edge because its people gain access to the core knowledge earlier, accumulate a deeper body of expertise about the technology, and develop good relationships with the labs at which these ideas germinate.

On the flip side, as IRB director Culler notes, "Ask yourself what you would do if building an industrial research organization as groundbreaking as Xerox PARC or Bell Labs. Today, it's not the little devices invented by an individual in a lab—transistors, core memory, and so forth—that make it happen.

"It's increasingly the large-scale systems that represent a confluence of diverse technologies, often contributed by industry. It's almost impossible to explore such large-scale systems at the academic level because of limited resources. Such exploration is equally impossible at the industrial level because the research demands a significant long-term investment without a clear return in the form of a product."

IRN's goal is to blend these two models in an open, collaborative, synergistic fashion in which its lablet program can explore ubiquitous computing scenarios at different levels of abstraction.

Intel Research Berkeley

IRB provided an ideal environment for lab director David Culler and codirector Hans Mulder to evolve, and continue to evolve, Intel's lablet and IP models. The real challenge, says Culler, was to build a world-class research lab from scratch in one year. Because UC Berkeley had already laid the groundwork for advanced ad hoc sensor networks through Kris Pister's smart dust project and Berkeley's subsequent mote prototypes, Culler decided to jumpstart IRB by bringing Intel's resources to bear on tiny, low-power, wireless embedded platforms.

This kind of research cuts across the traditional areas of architecture, operating systems, databases, networks, languages, and human-computer interaction. Moreover, IRB's research findings apply not only to small devices, but also to very large systems composed of these devices, and, ultimately, to planetary-scale services as exemplified by IRB's follow-on PlanetLab project. This research sharpened IRB's focus on scale in *extreme systems*, both large and small.

Motes. Pister's early smart dust project had as its goal embedded computing's holy grail: a cubic millimeter piece of silicon with integrated computing, communication, sensing, and power.[4] Shortly thereafter, his team built the first Berkeley mote, which measured one and one-half inches in diameter, to show what could be done using today's off-the-shelf components.

The Berkeley motes have continued to evolve, culminating in the current Mica platform shown in Figure 1.[5] Mica's Atmel Atmega 8-bit microcontroller runs at 4 MHz and delivers 4 MIPS. It has a 128-Kbyte flash program memory, 4-Kbyte SRAM, an internal 8-channel 10-bit DAC, 48 I/O lines, an external universal asynchronous receiver/transmitter, and one serial peripheral interface port. The node interfaces to a fairly rich sensor platform that contains light and temperature sensors, 2D accelerometers and magnetometers, and a microphone and sounder—with other sensor configurations possible.

Although Pister viewed his early prototype as merely an engineering demonstration, Culler considered it a watershed because he wanted to go deeper and explore operating systems structures that would scale down in accordance with Moore's law. This interest led to Berkeley's TinyOS operating system (http://today.cs.berkeley.edu/tos/), a multithreaded, event-based OS designed precisely to explore new interfaces for this embedded realm. TinyOS uses a modular structure so that developers can replace any part of the operating system—for example, the network stack or even the scheduler—with their own version.[5,6]

Significantly, the Berkeley group handed the initial mote design over to Crossbow (www.xbow.com). In return, Crossbow manufactured the first few hundred motes for the Berkeley group at a reasonably low cost. Now effectively having a hardware platform, an OS, and a hardware manufacturer and distributor, the research community could begin building sensor networks on a large scale, developing the network layers, and experimenting with various system-level concepts and applications.

Today, more than 100 research groups—including commercial entities like Honeywell and Siemens—experiment with the platform. Intel Research also has internal SRP activities that deal with both the hardware and software aspects of ad hoc sensor network design.

Ad hoc sensor networks. IRB also collaborates with the College of the Atlantic in an ecological study that uses motes equipped with microweather stations to study the nesting habits of storm petrels, which each year build burrows underground on Great Duck Island, Maine. IRB fielded 40 to 50 nodes equipped with meteorological sensors that detect light, temperature, humidity, and barometric pressure. Researchers placed the sensors both above ground and in petrel burrows, then established a transit network to stream the data back to a base station and on to the college via a satellite uplink.

In addition to letting researchers study implementation issues in real-world conditions in which they are likely to be deployed, networks such as these present revolutionary opportunities for the life sciences because they let researchers monitor a habitat without disturbing it. They also provide new opportunities for disciplines such as environmental science and civil engineering.

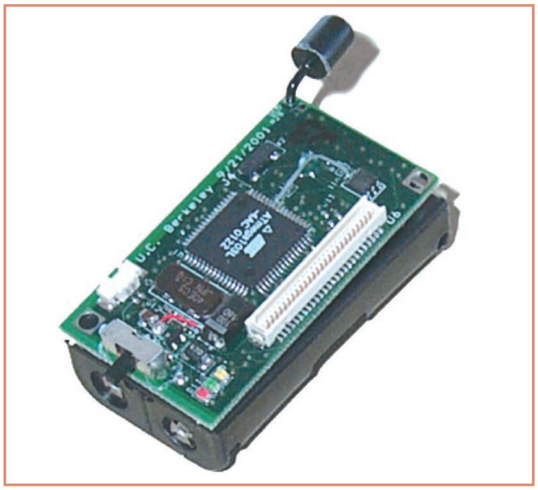

Figure 1. Mica wireless sensor platform. The node measures 1.25 x 2.25 inches, as dictated by the two-AA battery power supply, which supports applications that last for several years. Mica was developed as a low-cost, energy-efficient, experimental platform, and a single-chip version of the architecture is in development.

In conjunction with this effort, IRB, in collaboration with Berkeley, developed its TinyDB (www.intel-research.net/Berkeley/index.asp) to help process in real time the huge volume of data such sensor networks generate. As Culler puts it, in the world of ad hoc wireless sensor networks, a profound relationship exists between in-network query processing and content-based routing. In this example, because you don't really care about node 37—only the ones associated with inhabited nests—the database work and networking work fuse in interesting ways.

PlanetLab. At the other extreme, there is a significant synergy between the very-large-scale Internet and a sensor network that contains vast numbers of tiny processors. IRB recently launched its second major research project, PlanetLab, a test bed to explore a new kind of Internet that has begun to emerge, attracting high-quality researchers throughout the world. This new Internet—which consists of technologies such as content distribution, peer-to-peer file-sharing networks, and network-embedded storage—is beginning to blur the distinction between packet forwarding and application processing.

Although researchers seek to define the next Internet services beyond name translation and routing, they lack the means for trying out these ideas. To remedy this shortcoming, IRB has started building an overlay network to deploy and test the new service models that might comprise the next-generation Internet on a planetary scale.

The fully programmable nodes that comprise this overlay network will provide a diversity of link behavior and widespread geographic coverage, connecting a large client population to the overlay. PlanetLab is designed to be both an experimental research test bed with realistic network behavior and a deployment platform that can support users of a popular new service.

IRP will tackle the software challenges related to this *global ubiquitous storage* approach—including how to access large amounts of data from mobile devices.

IRB plans to build the initial tools and the basic infrastructure that will encourage new Internet services. It will base part of its design approach on planned obsolescence: PlanetLab will specify or create the initial software components—a *virtual machine monitor* (VMM) that runs on each node and specifies a controlled interface to which services are written and a management service that controls the overlay network. But IRB intends for the community at large to develop many of the scalable replacements for those tools.[7]

From the experimental test bed perspective, researchers don't need services on a single node to write research papers; they require a large slice through thousands of nodes over a six-month span to field a novel service and run a workload evaluation. The overlay must therefore be self-organizing and support distributed virtualization, with the VMM that runs on each node allocating and scheduling computational and networking resources to multiplex multiple services.[7] Providing this capability means drilling through the basic operating system substructure, the networking structure on top of it, and a whole set of security and privacy issues.

IRB began planning the PlanetLab project in March 2002, then started seeding the initial set of nodes with researchers from 15 or so universities around the US in July. Currently, more than 100 nodes are online around the world. These efforts are intended to lead to a self-supporting consortium (www.planet-lab.org)—including industrial researchers—as the network grows to the planned 1,000 nodes.

Overall, this project exemplifies the kind of research IRB and PlanetLab want to do and should do. No rational university would attempt to engineer and support a planetary-scale infrastructure, says Culler, and no sensible industry player would want a single company to tackle such a project with only its vested interest at stake. Significantly, this project illustrates the potential power and synergy of the lablet model, as it requires many new research contributions by the Intel Research Network.

Intel Research Pittsburgh

IRP lablet director Mahadev (Satya) Satyanarayanan focuses his lab's efforts on storage solutions for distributed environments. These projects will target widely distributed storage systems to support future mobile and ubiquitous computing systems. These systems span a broad range of technologies from wearable devices to server farms, from gigabit LANs to kilobit wireless links, and from mission-critical database applications to entertainment products. Satyanarayanan expects that work on integrating new and existing hardware storage technologies into innovative, higher-level software systems will lead to new storage paradigms and superior implementations of conventional storage systems.

Research into Internet-based storage models provides a major focus. Under these models, files and databases would be stored and maintained on the Internet and could easily be accessed from anywhere in the world. IRP will tackle the software challenges related to this *global ubiquitous storage* approach—including how to access large amounts of data from mobile devices with limited battery life and how to maintain users' privacy and control over data stored in cyberspace.

IrisNet. Ubiquitous Webcams collect vast amounts of potentially useful data. However, researchers lack effective tools for querying this data. Thus, the Internet-Scale Resource-Intensive Sensor Network (IrisNet) project seeks to develop a scalable software infrastructure that will let users with Internet access mine a wide range of sensors such as Webcams and microphones.

IrisNet will facilitate the deployment of sensor services that leverage both real-time and historical feeds from any number of sensors. Locally attached to a computer, each sensor can perform resource-intensive processing, such as image recognition or speech recognition. Whereas the Berkeley smart dust approach is exploring the capabilities of tiny devices, IRP seeks to seamlessly applying high-end processing to such data to provide "brilliant rocks."[8]

Internet Suspend/Resume. The Internet Suspend/Resume (ISR) project builds on the metaphor of notebook computers. When shut down, these systems preserve their execution environment intact. Then, when the user opens the computer again, no matter where, the system restores the identical execution environment.

The ISR project seeks to achieve the same effect without transporting the computer by letting a user seamlessly suspend work at one workstation and resume it at another. After moving to the second workstation, the user would find the identical execution environment. IRP plans to realize this OS-independent capability through a combination of virtual machine technology and distributed file systems.[9]

Data staging, an effort closely related to ISR, seeks to improve the performance of interactive applications running on small, storage-limited clients. Data staging speculatively prefetches data from distant file servers and stages them on nearby surrogate machines. Clients redirect file requests to surrogates, reducing access times. Strong end-to-end encryption ensures data privacy and authenticity, allowing users to implement surrogates on untrusted hosts such as commodity Web servers.[10,11]

Disk-Assisted Search. The goal of the Disk-Assisted Search project is to enable rapid interactive search of huge volumes—many terabytes or larger—of nonindexed, loosely-structured data such as images, video, x-rays, CT scans, and so on. This data typically is scattered in different parts of an organization and resides on disks that may be directly attached to a LAN or attached via a host.

A key problem that users face is "finding the needle in the haystack"—in other words, finding a vaguely specified piece of relevant and urgently needed data that may or may not be present somewhere in a huge mass of data. Processing embedded in disks enables the distribution and parallelization of such searches. Such "active disks" execute *searchlets*, pieces of code that a domain-specific front-end search application generates.

Intel Research Seattle

Seattle lab director and University of Washington computer science professor Gaetano Borriello defines his lab's role as making ubiquitous computing desirable to consumers who simply don't want to integrate another device into their systems. Unfortunately for these consumers, today's reality ensures that almost any device above a certain price—and soon almost every device at any price—will contain some intelligence.

To avoid the resistance most consumers have to even simple programming and configuration tasks, these devices must be seamlessly integrated into the consumer's environment. Without such integration, consumers who balk at the intricacies of managing their home entertainment center will not attempt to confront a device array that is orders of magnitude more complex.

The Seattle lab is exploring three major research areas that the staff hopes will contribute to the popularity of ubiquitously embedded computing devices:

- association of devices,
- Zero3 applications, and
- proactive applications.

Exploring an example of each area reveals much about the Seattle lab's plans and processes.

Applications by association. As an example of device association, take a Bluetooth-enabled phone specifically configured to interact with the hands-free speakerphone in your car. Suppose, however, you're traveling to a different city and rent a car once you arrive. Later, your phone rings while you're driving on a highway, and you nearly crash looking for the device that usually responds to a button on your dashboard.

Ideally, both systems should have discovery mechanisms that register and authenticate each other, perhaps through sensors that let them both know they are traveling at the same speed and experiencing the same phenomena, therefore they are in the same car—not in the car in the next lane and still within radio range.

The Seattle lab's role is to make ubiquitous computing desirable to consumers who simply don't want to integrate another device into their systems.

As this example shows, ideally, ubiquitous devices should interact with their immediate environment to determine what communication relationships to establish. Embedded sensors can provide an environmental signature that facilitates this discovery or association of devices.

Location sensing—the tracking of people and objects using current technologies like the Global Positioning System, RFID tags, or beacons—augments our ability to associate devices, but with varying degrees of precision or applicability. GPS, for example, may need a large antenna, or even a second antenna, for land-based differential signals, and it doesn't generally work indoors. To provide increased precision and environmental coverage, we can use *fused location sensing* with multiple location-sensing technologies and *ad hoc location sensing* with a cluster of sensors sharing data with one another to converge on an accurate position.[12]

Researchers face the challenge of building a system that combines these multiple location and sensing technologies and lets application developers write to a single API that abstracts the various underlying technologies. Providing this capability will let developers focus on what really interests them—coordinate systems.

The Seattle lab is currently developing Location Stack, middleware that combines multiple location technologies so that applications can take advantage of an environment's available information. Modeled on the Open Systems Interconnection model, the Location Stack partitions the sensor network into seven increasingly abstract layers that only need to know how to communicate with their nearest neighbors. The lowest Sensor Layer, for example, contains hardware and software drivers that detect sensor data and export it to the Measurements Layer residing above it, which in turn transcribes the raw data into canonical measurement types—

Figure 2. The PlantCare robot hardware platform consists of a Pioneer 2-DX mobile robot with custom hardware for watering plants, inductive charge coils for recharging both the robot and the deployed sensors, a laser scanner used for navigation, and an IEEE 802.11-enabled laptop that runs the robot's control and navigation software.

distance, angle, proximity, asserted position, and so forth—for export to the Fusion Layer, which combines data from multiple sensors to create a more precise representation of the object's position and orientation.[13]

Location Stack provides a modular framework for encapsulating program behavior and permits the introduction of new sensor technologies and services without altering the entire application. Intel Research has already transferred elements of this sensor fusion technology to the Emerging Platform Lab within CTG and has begun to incorporate the technology into product.

Zero3 applications. To be practical, associated devices should require zero configuration, maintenance, and downtime. Consider being able to stop at a hardware store, for example, where you buy a suite of motion, glass-breaking, door, and temperature sensors. You take them home and plant them around the house, then simply log on to a Web site where you can activate the sensor network, thus enabling your own home security system.

How does the security system determine what components you have and what policies it should enforce? Can the system actually observe patterns within the house and infer human activities? Can it determine that windows open and close at predictable times and therefore don't pose a threat? Can it determine that hot water starts running at 7:00 a.m. for showers and doesn't run again until after dinner, so why not lower the water heater temperature in the interim to save energy? If you move a motion sensor, can the sensor recognize its new location? If you unplug your PC to move it, or even to replace it with a new one, will the system seamlessly transfer control to another machine and automatically provide software upgrades as needed?

Seattle's PlantCare project takes this technology a step further. Figure 2 shows this autonomous robotic system, which sports custom hardware for watering plants and for recharging both the robot and the sensor network.[14] This system tackles problems associated with both wireless sensor networks and mobile robots. More importantly, it also explores what it means to develop an application that users only set up once—even if they rearrange the furniture or move the plants.

Seattle lab researchers started by adding wireless light, temperature, and soil humidity sensors to plants. The sensors periodically transmit readings to the Rain middleware, an XML-based software framework that lets applications discover one another and interact as a collection of cooperating services that communicate via asynchronous events.

The 15 Rain services that make up PlantCare include a low-level robot control for navigation and for sensor and actuator activation. A high-level robot manager provides functionality for performing tasks such as watering plants. A gardener service examines the sensor data and consults specific plant-care instructions in a plant encyclopedia service, then notifies the task server if specific plants need care. Similarly, a mechanic service monitors power levels in the sensors and robots. The task server coordinates all tasks and sends them to the mobile robot.

Although it is by no means intended as the next killer app, PlantCare provides a rich experimental test bed for exploring Zero3 applications and, more particularly, the software infrastructure needed to support them.

Proactive applications. The Seattle lab is exploring proactive applications that can, for example, monitor the home environment of an Alzheimer's patient, use learning algorithms to identify normal or preferred behavior, and assist caregivers and the patient in performing activities of daily living. To succeed, such applications must provide a balance between system autonomy and human decision making.

Obviously, such systems for machine learning will require a lot of AI capabilities, as well as compute cycles, and Intel is working closely with a number of universities to provide these sophisticated machine-learning techniques.

One lab project that has progressed further than most, the Labscape ubiquitous computing system, simplifies laboratory work for cell biologists by making procedural information instantly available and organizing experimental data into a formal representation they can share with colleagues.

The Seattle lab's researchers carefully observed the work habits and procedures within an authentic user community—the Cell Systems Initiative in the University of Washington's Department of Bioengineering and Immunex Corp. They then abstracted common biology lab procedures into six simple operations—combination, incubation, dispensing, separation, detection, and storage and retrieval—that they could combine in the flow graph structures shown in Figure 3. Biologists can step through their experiments and enter procedural information on a terminal at the lab bench. They can also transparently migrate the user interface to other terminals in the lab using an infrared-based active-badge system.[15]

The project's developers initially intended for Labscape to be a full-on ubiquitous computing system that would use wireless sensors and recognition-based modalities to automatically construct an experiment while minimizing explicit interaction. However, the observed recalcitrance of biologists to carry around or use any tools other than their lab notebooks persuaded Labscape's developers to adopt a design that did not intrude on the laboratory environment and that relied on basic computing equipment and networking infrastructure.[15] This willingness to forego the smart environment as a platform for technology evaluation, focusing instead on user-centered design, demonstrates the Seattle lab's continued adherence to its charter of making ubiquitous computing desirable.

Although Labscape does not preclude the introduction of a rich sensor base, deep integration with existing instrumentation, and new recognition-based modalities—and will incorporate them if researchers find them useful—the baseline system exemplifies functionality in its purest form.

Intel Research Cambridge

Cofounder of Microsoft's Cambridge research center and former head of Cambridge's Marconi Labs, Derek McAuley now directs Intel's newest lablet. IRC will include fundamental networking research, from mathematical modeling of network traffic to emerging technologies such as optical switching.

IRC's Scope project will measure and model network performance—providing monitoring technology capable of packet capture and online compression at 10 and 40 Gbps, as well as offline analysis tools that can relate network packets to the behavior of the application on the user's computer. Cambridge is also looking into novel uses of virtual machines running simultaneously on many different operating systems on a single computer system, as the PlanetLab project requires.

THE ROAD AHEAD

This vision of a world containing billions of inexpensive embedded sensors begs an interesting question: How does it fit into Intel's business model of high-end, high-margin processors? Suppose these disruptive projects really take root, and Intel develops some early market products that begin to sell. Will the corporation invest $100 million to launch a new business, in addition to building a new fab to produce them at $1.00 per device?

As Pat Gelsinger admits, "Personally, I think we need to deliver the end devices. But at the end of the day, there will be an intense conversation, and Intel may indeed decide that it is too much of a low-margin commodity to actively manufacture."

Even so, the excitement about wireless sensor networks seems justified in that, if successful, they will unleash a wealth of information destined to be delivered to the desktop in real time. From Intel's perspective, this alone will drive the demand for net- work access points, personal computers, and high-end servers. And if planetary-scale networking takes off, Intel will have a highly developed understanding of the next-generation Internet that will let it supply the silicon building blocks in that space as well.

This research effort's most exciting aspect, however, is that it may succeed in realizing the vision of proactive computing and provide a new paradigm for computer science in general. As David Culler states, "Intel didn't buy a brain trust; instead, it is interested in amplifying the brain trust that is out there through collaboration with academia. Intel has provided us with some capabilities we've never had before."

What is unique about Intel's lablets, says Edward Lazowska, former chair of the University of Washington Computer Science and Engineering Department, is that they are created with the explicit purpose of collaboration, and the specific projects are picked from the intersection

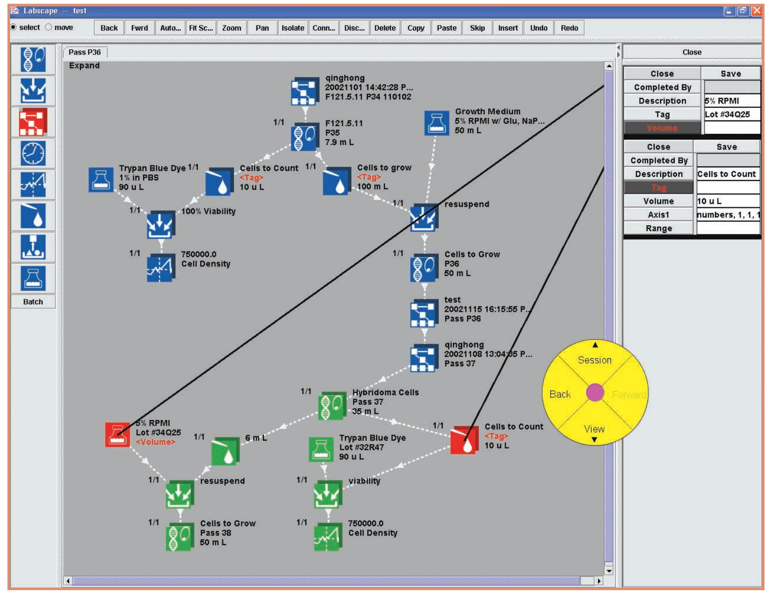

Figure 3. Biology lab flow graph. Labscape keeps track of materials as they flow through the laboratory. The buttons along the top and left of the screen provide interactive editing capabilities via touch panel computers distributed throughout the laboratory. Dialog boxes in the right column allow viewing and editing of properties in each step in the procedure.

of what's interesting to both Intel and the neighboring university.

Yes, Intel ultimately wants to sell more processors, but these projects are tackling important computer science and engineering research problems. Moreover, Intel's decision to locate the lab in Seattle proper is an enormous boon to the region as, along with Microsoft Research, it contributes to IT research by bringing in more people doing interesting things—as do the many Japanese research labs located in New Jersey, Boston, and Northern California. In fact, says Lazowska, the region needs even more of these labs, and he hopes that the IR Seattle lab is a model for industry-university collaborative research that can be cloned in this and other disciplines.

As current University of Washington CSE chair David Notkin states, "In truth, this is an ecosystem that the region is extremely fortunate to have. With Microsoft Research and now with Intel Research, we can work really hard to enhance this wonderful situation by further improving our department, establishing new industry-university partnerships, and working with venture capital to create startups. With luck and hard work, if all goes according to plan, great people will come to us—they'll benefit and we'll benefit—and when the economy turns around, it will be absolutely phenomenal for the region." ■

Acknowledgments

I thank Gaetano Borriello, David Culler, Pat Gelsinger, Edward Lazowska, Hans Mulder, David Notkin, David Tennenhouse, and Roy Want for devoting their time to the interviews that provided the background for this article. I also thank Kevin Teixeira and Judy Aydt for facilitating these interviews.

References

1. S. Moore, "Industrially Sponsored University Research in Information Technology: Some Recommendations Regarding Intellectual Property Agreements," *CRA Chairs Conference*, 2002; http://www.cs.utexas.edu/users/moore/publications/ip-memo.ps.
2. D. Hodges, "Electrical Engineering and Computer Science Intellectual Property Pilot Program," Office of Technology Transfer UCB, memo no. 2000-02, 30 Aug., 2000; http://patron.ucop.edu/ottmemos/docs/ott00-02.html.
3. H. Mulder, "The Intel Research Network: An Innovative Model of Industry-University Collaboration," *Computing Research News*, Nov. 2002, pp. 1, 9.
4. B. Warneke et al., "Smart Dust: Communicating with a Cubic-Millimeter Computer," *Computer*, Jan. 2001, pp. 44-51.
5. J. Hill and D. Culler, "Mica: A Wireless Platform for Deeply Embedded Networks," *IEEE Micro*, Nov./Dec. 2002, pp. 12-24.
6. D. Culler et al., "A Network-Centric Approach to Embedded Software for Tiny Devices," *Proc. First Int'l Workshop Embedded Software* (EmSoft 2001), LNCS 2211, Springer Verlag, 2001; tech. report IRB-TR-01-001; www.intel-research.net/Publications/Berkeley/120520021014_7.pdf.
7. L. Peterson et al., "A Blueprint for Introducing Disruptive Technology into the Internet," www.planet-lab.org/pdn/pdn02-001.pdf.
8. S. Nath et al., "Mining a World of Smart Sensors," tech. report IRP-TR-02-05, Aug. 2002; www.intel-research.net/Publications/Pittsburgh/081620021315_95.pdf.
9. M. Kozuch and M. Satyanarayanan, "Internet Suspend/Resume," *Proc. IEEE Workshop Mobile Computing Systems and Applications*, IEEE CS Press, 2002; tech. report IRP-TR-02-01; www.intel-research.net/Publications/Pittsburgh/081620021258_88.pdf.
10. J. Flinn, S. Sinnamohideen, and M. Satyanarayanan, "Data Staging on Untrusted Surrogates," tech. report IRP-TR-02-02, May 2002; www.intel-research.net/Publications/Pittsburgh/081620021310_92.pdf.
11. R. Balan et al., "The Case for Cyber Foraging," *Proc. ACM SIGOPS European Workshop*, ACM Press, 2002; tech. report IRP-TR-02-04; www.intel-research.net/Publications/Pittsburgh/081620021313_ 94.pdf.
12. J. Hightower, "Location Systems for Ubiquitous Computing," *Computer*, Aug. 2001, pp. 57-66.
13. J. Hightower, B.Brumitt, and G. Borriello, "The Location Stack: A Layered Model for Location in Ubiquitous Computing," *Proc. IEEE Workshop Mobile Computing Systems and Applications*, IEEE CS Press, 2002, pp. 22-28.
14. A. LaMarca et al., "PlantCare: An Investigation in Practical Ubiquitous Systems," *Proc. Int'l Conf. Ubiquitous Computing*, LNCS 2498, Springer-Verlag, 2002; tech. report IRS-TR-02-007; www.intel-research.net/Publications/Seattle/072320021545_76. pdf.
15. L. Arnstein et al., "Labscape: A Smart Environment for the Cell Biology Laboratory," *IEEE Pervasive Computing*, July-Sept. 2002, pp. 13-21.

Scott Hamilton *is senior acquisitions editor for* Computer. *Contact him at s.hamilton@computer.org.*