Digital gadgets demand ever more of our attention with their rude and thoughtless interruptions. Engineers are now testing computers, phones and cars that sense when you're busy and spare you from distraction

CONSIDERATE COMPUTING

By W. Wayt Gibbs

"YOUR BATTERY IS NOW FULLY CHARGED," ANNOUNCED THE LAPTOP COMPUTER to its owner, Donald A. Norman, with enthusiasm—perhaps even a hint of pride?—in its synthetic voice. Norman, a chief advocate of the notion that computers and appliances ought to be programmed with something akin to emotions, might normally have smiled at the statement. Instead he blushed—and no doubt wished that his computer could share his embarrassment. For at that moment, Norman was onstage at a dais, having addressed a conference room of cognitive scientists and computer researchers, and his Powerbook was still plugged into the public address system. Many in the audience chuckled at the automated faux pas and shook their heads. The moderator, flustered, shot Norman a less than sympathetic look.

And yet we've all been there. Our cell phones ring during movies. Telemarketers interrupt our dinners with friends. Our laptops throw up screensavers in the middle of presentations. "You've got mail!" derails our train of thought just as we get in the groove.

To be sure, distractions and multitasking are hardly new to the human condition. "A complicated life, continually interrupted by competing requests for attention, is as old as procreation," laughs Ted Selker of the Massachusetts Institute of Technology Media Lab. But increasingly, it is not just our kids pulling us three ways at once; it is also a relentless barrage of e-mail, alerts, alarms, calls, instant messages and automated notifications, none of them coordinated and all of them oblivious to whether we are busy—or even present. "It's ridiculous that my own computer can't figure out whether I'm in front of it, but a public toilet can," exclaims Roel Vertegaal of Queen's University in Ontario.

Humanity has connected itself through roughly three billion networked telephones, computers, traffic lights—even refrigerators and picture frames—because these things make life more convenient and keep us available to those we care about. So although we could simply turn off the phones, close the e-mail program, and shut the office door when it is time for a meeting or a stretch of concentrated work, we usually don't. We just endure the consequences.

"We take major productivity hits with each interruption," says Rosalind Picard, a cognitive scientist at the M.I.T. Media Lab. People juggle the myriad demands of work and daily life by maintaining a mental list of tasks to be done. An interruption of just 15 seconds causes most people to lose part of that to-do list, according to experiments by Gilles O. Einstein of Furman University.

Numerous studies have shown that when people are unexpectedly interrupt-

PETER HOEY

> **"It's ridiculous that my own computer can't figure out whether I'm in front of it, but a public toilet can."**

ed, they not only work less efficiently but also make more mistakes. "It seems to add cumulatively to a feeling of frustration," Picard reports, and that stress response makes it hard to regain focus. It isn't merely a matter of productivity and the pace of life. For pilots, drivers, soldiers and doctors, errors of inattention can be downright dangerous.

"If we could just give our computers and phones some understanding of the limits of human attention and memory, it would make them seem a lot more thoughtful and courteous," says Eric Horvitz of Microsoft Research. Horvitz, Vertegaal, Selker and Picard are among a small but growing number of researchers trying to teach computers, phones, cars and other gadgets to behave less like egocentric oafs and more like considerate colleagues.

To do this, the machines need new skills of three kinds: sensing, reasoning and communicating. First a system must sense or infer where its owner is and what he or she is doing. Next it must weigh the value of the messages it wants to convey against the cost of the disruption. Then it has to choose the best mode and time to interject.

Each of these pushes the limits of computer science and raises issues of privacy, complexity or reliability. Nevertheless, "attentive" computing systems have begun appearing in newer Volvos [*see box on opposite page*], and IBM has introduced Websphere communications software with a basic busyness sense. Microsoft has been running extensive in-house tests of a much more sophisticated system since 2003. Within a few years, companies may be able to offer every office worker a software version of the personal receptionist that only corner-suite executives enjoy today.

But if such an offer should land in your inbox, be sure to read the fine print before you sign. An attentive system, by definition, is one that is always watching. That considerate computer may come to know more about your work habits than you do.

Minding Your Busyness

MOST PEOPLE AREN'T AS BUSY as they think they are, which is why we can usually tolerate interruptions from our inconsiderate electronic paraphernalia. James Fogarty and Scott E. Hudson of Carnegie Mellon University recently teamed up with Jennifer Lai of IBM Research to study 10 managers, researchers and interns at work. They videotaped the subjects and periodically had them rate their "interruptibility." The amount of time the workers spent in leave-me-alone mode varied from person to person and day to day, ranging from 10 to 51 percent. On average, the subjects wanted to work without interruption about one third of the time. In studies of Microsoft employees, Horvitz has similarly found that they typically spend more than 65 percent of their day in a state of low attention.

Today's phones and computers, which naively assume that the user is never too busy to take a call, read an e-mail, or click "OK" on an alert box, thus are probably correct about two thirds of time. (Hudson and Horvitz acknowledge, however, that it is not yet clear how well these figures generalize to other jobs.) To be useful, then, considerate systems will have to be more than 65 percent accurate in sensing when their users are near their cognitive limits.

Fortunately, this doesn't seem to require strapping someone into a heart monitor or a brain scanner. Fogarty and his collaborators have found that simply using a microphone to detect whether anyone is talking within earshot would raise accuracy to 76 percent. That is as good as the human judgment of coworkers who viewed videotapes of the subjects and guessed when they were uninterruptible. When Fogarty's group enhanced the software to detect not only conversations but also mouse movement, keyboard activity and the applications running on machines, the system's accuracy climbed to 87 percent for the two managers. Curiously, it rose only to 77 percent for the five scientists, perhaps because they are a chattier bunch.

Bestcom/Enhanced Telephony, a Microsoft prototype based on Horvitz's work, digs a little deeper into each user's computer to find clues about what they are up to. Microsoft launched an internal beta test of the system in mid-2003. By last October, Horvitz says, about 3,800 people were using the system to field their incoming phone calls.

Horvitz himself is one of those testers, and while we talk in his office in

Overview/*Sensing Attention*

- Computers continue to grow cheaper, more powerful and more pervasive. Human attention, in contrast, is a scarce and fixed resource. As we spend more of our time surrounded by "smart" devices, their productivity-sapping, stress-inducing interruptions increasingly detract from their value.
- Researchers at corporate and academic labs have developed several powerful techniques that enable computerized machines to estimate their user's cognitive load and focus of attention.
- Engineers are testing prototype systems that can automatically prioritize, reschedule or forward incoming phone calls and digital messages, much as a personal receptionist might. Similar attention-sensing technology has begun to appear in cars and may lead to more "considerate" everyday appliances.
- Monitoring a person's attention requires sophisticated reasoning based on in-depth surveillance. The work thus raises issues of privacy and reliability.

Attentive Autos

First the eyelids droop, then the head begins to bob. The car drifts out of its lane, then jerks back erratically. The signs of a drowsy or distracted driver are not hard to spot, but one must look for them. Many vehicles soon will.

The U.S. National Highway Traffic Safety Administration (NHTSA) figures that 20 to 30 percent of crashes reported to the police—about 1.5 million every year—result at least in part from inattentive drivers. Technology contributes to this problem, cluttering up the cabin with phones, DVD players, touch-screen maps and other doodads. Now technology of a different kind may help to solve it.

In 2003 Volvo added an attention management system to its S40 sedans. Sensors pick up steering actions, accelerator position and other vehicle dynamics. They feed into a computer, which looks for evidence of swerving, overtaking or hard braking. When it notices such demanding maneuvers, the system suppresses nonsafety-critical messages from the onboard phone, navigation system, warning lights, and so on. In June, Motorola and DaimlerChrysler demonstrated a minivan outfitted with a similar system. More recently, Volvo has been testing cameras that can detect drowsy eyelids and suspicious lane crossings [*see illustration at right*].

Last March the European Union launched a four-year, €12.5-million project to develop industrywide standards for adaptive driver-vehicle interfaces by 2008. Engineers must still solve many tricky issues.

"Liability is a major stumbling block in the U.S.," observes Trent Victor, who designs driver awareness systems for Volvo. Automakers must be certain the computer will not make a bad situation worse, even if the driver uses it improperly.

Indeed, it seems inevitable that as the system adapts to drivers, drivers will adapt right back. "There are people today who use the rumble strips [which buzz when a car approaches the edge of a road] as a way to help them watch TV," Trent says. —*W.W.G.*

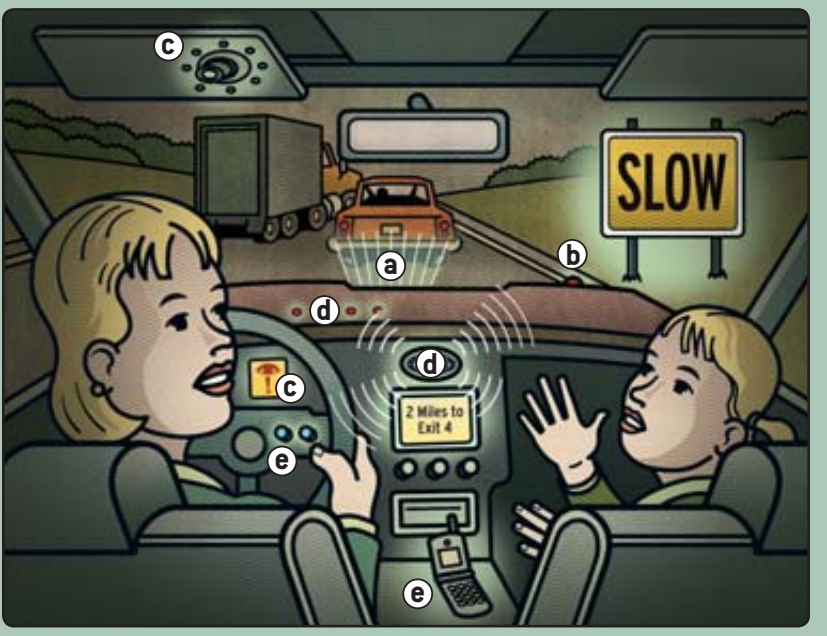

"CONSIDERATE CARS" of the future could combine a variety of systems to sense the attention level of the driver and prevent dangerous distractions. In this example, Alice is driving at high speed in heavy traffic, but her concentration is interrupted by a rambunctious daughter. The car ahead has begun to brake as it enters a construction zone. But at the same moment, Alice's mobile phone receives a call, and the onboard navigation system decides to alert her to an upcoming exit.

(a) **MIND THE GAP:** A short-range radar mounted on the front of the car notices that the vehicle ahead is getting closer. The adaptive cruise-control system automatically eases off the gas and gently applies the brake to maintain separation. Such systems are already available on certain luxury cars.

(b) **IT'S A SIGN:** A forward-looking video camera watches the road. An onboard computer continually scans the video for lane markers and road signs. It recognizes the construction warning ahead and raises the car's estimate of the current burden on the driver's attention. Alexander Zelinsky and his co-workers at the Australian National University have demonstrated such sign-reading systems.

(c) **WATCHING THE WATCHER:** LEDs on the visor shine invisible infrared light into Alice's eyes; a small infrared camera picks up the reflections from her pupils and deduces the direction of her gaze. If the driver takes her eyes off the road or shows signs of drowsiness, an icon in the instrument panel lights up, and the car goes on high alert. It vibrates the steering wheel if needed to rouse her. Zelinsky started a company, Seeing Machines, to commercialize such gaze monitors, and Volvo and others have tested them in cars and trucks.

(d) **EYES ON THE ROAD:** Because Alice is looking away at a moment when the demands on her attention are high, the system emits an alert sound through the car speakers and flashes a sequence of lights on the dash that draws her attention back to the car ahead. Tests by Volvo have shown this method to be effective.

(e) **NOT NOW, PLEASE:** The onboard computer has a model of the driver's abilities and is always monitoring the cognitive load, so it knows that now is a bad time for a phone call or a navigational alert. The computer turns off the ringer on the cell phone and instead lights an unobtrusive "incoming message" button on the steering wheel that Alice can press to take the call when it is safe to do so. The car likewise lights another button to let her know that the navigation system has an instruction for her. Chrysler has built such gentle notification devices into some of its prototype cars.

PETER HOEY

AN AUTOMATED PERSONAL RECEPTIONIST

Notification Platform, a prototype system developed by Horvitz and his co-workers at Microsoft, performs triage on incoming communications much as a receptionist would. The system, which runs on a central server, analyzes messages to decide whether, when and how to notify the recipient of their arrival. It would handle the same e-mail message to Alice from her boss differently if she were conducting a job interview (*scenario 1*), reading e-mail (*2*) or packing for a business trip (*3*). One part of the system estimates the value of the message and how that value decays with time. A second part uses sensors and a statistical model to guess the focus and intensity of Alice's attention, what devices she is using, how likely she is to see the message without an alert, and how these variables will change in the future. A third part decides what kind of alert to issue.

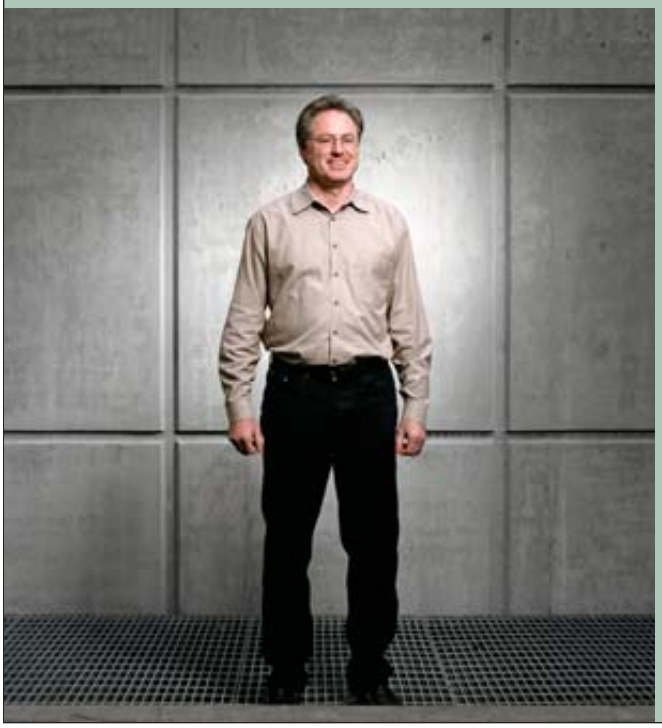

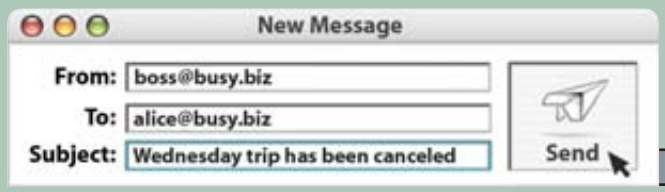

ERIC HORVITZ (*left*) pioneered the use of sensors and statistical models (*above*) to build attention-aware communications systems.

Redmond, Wash., Bestcom silently handles one call after another. First it checks whether the caller is listed in his address book, the company directory, or its log of people he has called recently. Triangulating these sources, it tries to deduce their relationship. Family members, supervisors and people he called earlier today ring through. Others see a message on their computer that he is in a meeting and won't be available until 3 P.M. The system scans Horvitz's and the caller's calendar and offers to reschedule the call at a time that is open for both. Some callers choose that option; others leave voice mail. E-mail messages get a similar screening. When Horvitz is out of the office, Bestcom automatically offers to forward selected callers to his cell phone—unless his calendar and other evidence suggest that he is in a meeting.

Most large companies already use computerized phone systems and standard calendar and contact management software, so tapping into those "sensors" should be straightforward. Not all employees will like the idea of having a microphone on all the time in their office, however, nor will everyone want to expose their datebook to some program they do not ultimately control. Moreover, some managers might be tempted to equate a "state of low attention" with "goofing off" and punish those who seem insufficiently busy.

The researchers seem to appreciate these risks. Hudson argues that an attentive system should not record audio, keystrokes or the like but simply analyze the data streams and discard them after logging "conversation in progress," "typing detected," and so on. "We built a privacy tool into Bestcom from the beginning," Horvitz emphasizes, "so users can control who is allowed to see the various kinds of information it collects about them."

Watching the Watcher

AS DIGITAL CAMERAS fall in price, that information may come to include video. With a simple $20 webcam, Horvitz's software can tell when a person is in view and whether she is alone or in a meeting. Fancier cameras can use the eyes as a window to the mind and perhaps extend the reach of considerate computers into the home.

Vertegaal has filled the Human Media Lab at Queen's University with everyday appliances that know when you are looking at them. "When I say 'on,' the lamp over there doesn't do any-

KC ARMSTRONG (*photographs*); PETER HOEY (*scenarios*); JEN CHRISTIANSEN (*diagram*)

SCENARIO 1

Monday, 10:07 A.M.

SCENARIO 2

Tuesday, 9:00 A.M.

SCENARIO 3

Wednesday, 6:15 A.M.

URGENCY CLASSIFIER consults its database to determine the relationship between the sender and recipient and their communication history. A linguistic analyzer scans the text for times, dates and key words. The classifier "stamps" the message with a dollar value that can fall as the communication grows stale.

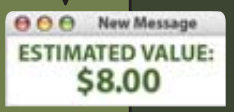

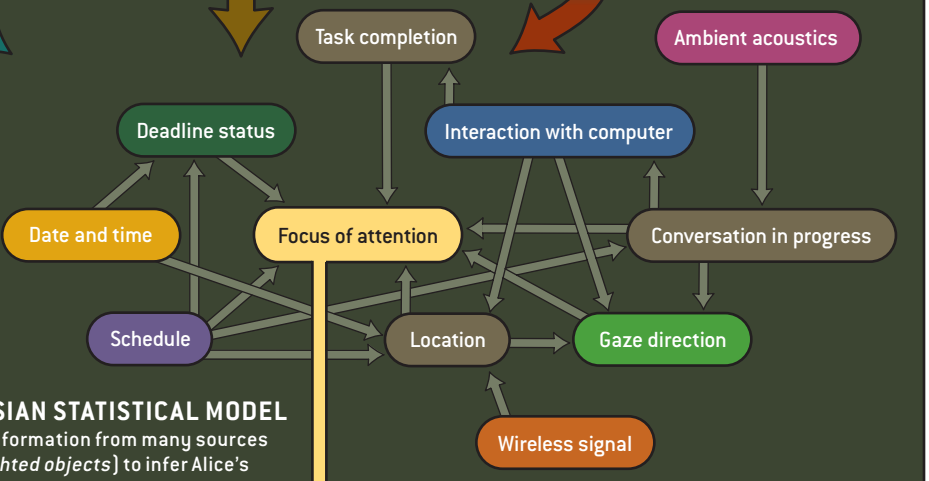

BAYESIAN STATISTICAL MODEL fuses information from many sources (*highlighted objects*) to infer Alice's current focus of attention.

The system also estimates the cost of interrupting Alice now and in the near future. It alerts her to the message only when the cost of a distraction falls below the value of a notification (*arrows*).

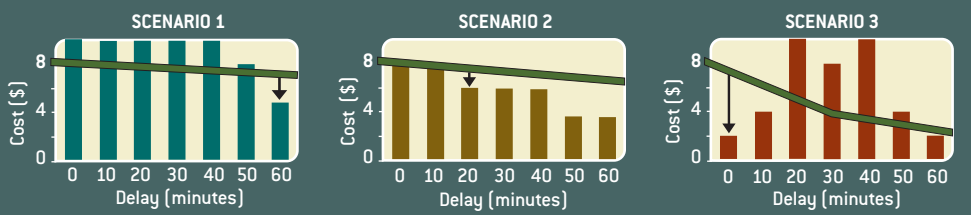

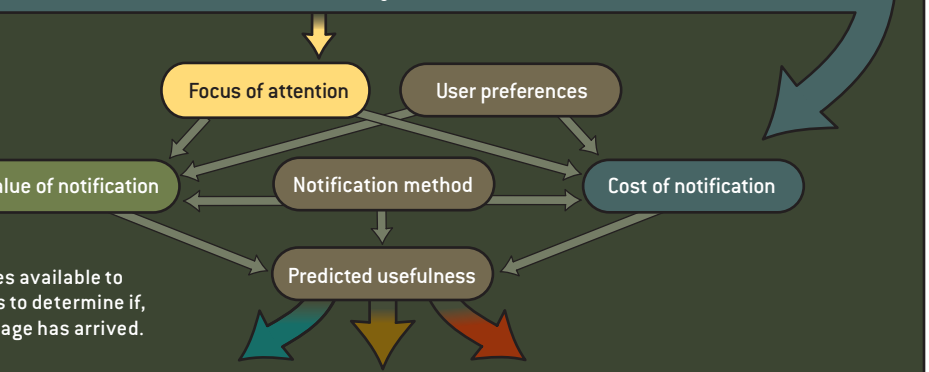

DECISION MODEL surveys the devices available to Alice and performs a cost-benefit analysis to determine if, when and how to signal her that this message has arrived.

SCENARIO 1

The Notification Platform adds the e-mail from Alice's boss to her inbox but does not notify her about the message until one hour later, after her appointment with the job candidate has ended.

SCENARIO 2

The system senses that Alice is at her PC and is composing an e-mail. It waits for her to finish writing, then chimes and displays a small alert box in the corner of the screen, notifying her about the message from her boss.

SCENARIO 3

The system matches the subject of the e-mail to an entry in Alice's schedule for today and assigns it highest priority. The e-mail is forwarded to her "considerate" cell phone, which signals her with an urgent ring tone.

WATCHING THE EYES TO FOLLOW THE MIND

Roel Vertegaal (*left*) and his students at Queen's University in Kingston, Ontario, have been enhancing televisions, phones, computers and video-conferencing systems with the ability to sense human eye contact. They have also put the technology to artistic use in AuraMirror (*right*). Hidden infrared lights, cameras and a computer transform a large monitor into a virtual magic mirror that superimposes nebulous "auras" over the images of people in front of it. The blobs extend outward and merge when two observers look at each other, giving visible form to an intangible human connection (*far right*).

> "Eye contact is the most accurate measure of attention that we have. But it's not perfect by any means."

thing," Vertegaal says, pointing over his shoulder. He turns to face the object.

"On," he says. LEDs mounted on a small circuit board stuck to the lamp shoot invisible infrared light into his pupils. The light reflects off his retinas, and an infrared camera on the board picks up two bright spots in the image, one from each eye. A processor does some quick pattern and speech recognition, and the lamp switches on.

Gaze detection can endow quotidian machines with seemingly magical behavior. Vertegaal answers a ringing telephone by looking at it and saying "Hello." When he stops talking and turns away from the phone, it hangs up. The TV in the lab pauses a DVD or mutes the sound on a broadcast show whenever it notices that there are no longer any eyes watching it. Some of Vertegaal's students walk around with eye-contact sensors on their hat or glasses. When the wearer enters a conversation, the sensor passes that information via a wireless link to the cell phone in his pocket, which then switches from ring mode to vibrate.

Although the technology is steadily improving, gaze detectors are still too expensive, bulky, ugly and unreliable for everyday use. "Eye contact is the most accurate measure of attention that we have—about 80 percent accurate in conversational settings," Vertegaal says. "But it's not perfect by any means."

Attentive appliances are mere parlor tricks, moreover, when they act independently. The real payoff will only come from larger, smarter systems that can both divine the focus of our attention and moderate our conversation with all our personal machines. Doing that reliably will require a nice bit of reasoning.

Trusting the Black Box

BROADLY SPEAKING, computers can use two techniques—rules or models—to decide when and how to transmit a particular piece of information. Both approaches must face the bugbear of complexity.

If the system is limited to following a few rules, users can predict exactly how it will treat a given message. Many e-mail programs, for example, manage spam by maintaining lists of known spammers and of legitimate contacts. When each e-mail arrives, its sender is compared against both lists and either deleted or delivered. Such systems are simple and clear—but infamously inaccurate.

Spam filters and network firewalls improved significantly when they began to rely on statistical models, called Bayesian networks, that are built by machine-learning algorithms. The user gives the algorithm many examples of desirable messages and also some counterexamples of undesired traffic. "The software identifies all the variables that influence the property that you are interested in [for example, not spam], then searches over all feasible relationships among those variables to find the model that is most predictive," Horvitz explains.

Bayesian networks can be eerily accurate. "They use probabilities, so they are wise in the sense that they know that they can't know everything," Horvitz

> "Artificial intelligence couldn't deliver the personal secretary. I'm pretty sure we can deliver a personal receptionist."

KC ARMSTRONG

elaborates. "That allows them to capture subtle behaviors that would require thousands of strict rules." In January he plans to present the results of a field trial of a model trained on 559 past appointments taken from a manager's datebook. When challenged with 100 calendar entries it had never seen, the model correctly predicted whether the manager would attend the meeting 92 percent of the time. And in four out of every five cases, the model matched the manager's own estimate of the cost of interruption during the meeting.

That sounds impressive, but some experts in the field remain skeptical. Users may have a very low tolerance for a system that erroneously suppresses one out of every 10 important calls. "The more 'attentive' things become, the more unpredictable they are," warns Ben Shneiderman of the University of Maryland. "We have a history in this community of creating 'smart' devices that people don't use because they can't understand how they operate."

Indeed, Vertegaal reflects, "artificial intelligence couldn't deliver the personal secretary, because it was too complicated." Nevertheless, he adds, "I'm pretty sure we can deliver a receptionist."

That would be welcome, but will considerate computing really reduce interruptions and boost productivity? At least for certain specialized tasks, the answer is: unquestionably.

Consider Lockheed Martin's HAIL-SS (Human Alerting and Interruption Logistics-Surface Ship) system. In much the way that Bestcom interposes itself between the phone system and an office worker, HAIL-SS keeps an eye on the sailors operating an Aegis naval weapons system and mediates the many alerts that Aegis produces. In combat simulations, HAIL-SS cut the number of interruptions by 50 to 80 percent, allowing sailors to handle critical alerts up to twice as quickly. The software lowered the perceived difficulty and stressfulness of the job by one quarter. The U.S. Navy now plans to deploy HAIL-SS throughout the fleet.

No comparable studies have yet been done in the office environment, however. Even with Bestcom diverting callers to voice mail and squelching e-mail alerts, Horvitz was interrupted 14 times in the course of our five-hour interview. Two fire alarms, a FedEx deliveryman and numerous colleagues poking their head into the office were merely examples of a large class of disruptions that will never disappear, because they benefit the interrupter.

Vertegaal is optimistic nonetheless. "By opening up these new sources of information about how available someone is, people will naturally adapt and use them to apply existing social rules of etiquette," he predicts. "So just by virtue of letting people know when you're busy, you'll get fewer interruptions." SA

W. Wayt Gibbs is senior writer.

MORE TO EXPLORE

Attentive User Interfaces. Special section. Edited by Roel Vertegaal in *Communications of the ACM,* Vol. 46, No. 3, pages 30–72; March 2003.

Learning and Reasoning about Interruption. Eric Horvitz and Johnson Apacible. Proceedings of the Fifth International Conference on Multimodal Interfaces, November 2003.

Proceedings of the International Workshop on Progress and Future Directions of Adaptive Driver Assistance Research. Washington, D.C., May 2004. Available online at **www.volpe.dot.gov/opsad/saveit/**

A variety of relevant technical articles are available online at **http://interruptions.net**

KC ARMSTRONG