

Beyond Weiser: From Ubiquitous to Collective Computing

Gregory D. Abowd, Georgia Tech

Considering the technological changes across computing's first three generations, how might the next serve humanity? Three critical technologies—the cloud, the crowd, and the shroud of devices connecting the physical and digital worlds—define the fourth generation of collective computing.

More than 25 years ago, Mark Weiser identified ubiquitous computing as the third generation of computing, following the first generation of mainframe computing and the second generation of personal computing.^{1,2} His vision demonstrates how we can define current and future computing generations based on the technological innovations and applications that influence the human experience with computing.^{1,2}

This framework yields the initially surprising outcome that the fourth generation's vision has already happened. Three critical technologies have emerged that dramatically influence how humans connect to computing. Distributed computing has matured through cloud services that provide effectively infinite bits and cycles. Human computation has been reintroduced to solve problems, and social computing platforms offer quick ways to communicate with targeted audiences and curate knowledge. The Internet of Things (IoT) and wearable computing have re-emerged, creating a layer between potentially every physical object (including

humans) and the digital world. These technologies—the *cloud*; the *crowd*; and what I call the *shroud*, or the layer of digital technology that connects the physical properties of people, places, and things to the digital domain—define a new era of cooperation between humans and computing that enhances both computational capabilities and the human experience. I call this new computing generation *collective computing*, but the name is less important than the phenomenon it represents and the applications it enables.

Although the critical technologies of the cloud, crowd, and shroud have roots going back decades, recent advances have made them much more promising individually and, more importantly, in combination. Collective computing applications provide individuals with an unprecedented level of self-sufficiency, helping them to harness information in real time, regardless of their training, in fields as diverse as navigation, entertainment, health, and education. Collective computing also augments organizations' capabilities. The core fourth-generation technologies already have research

TABLE 1. A framework for comparing computing generations, inspired by Mark Weiser.

Generation	Time frame	Human–computer ratio	Canonical device	Application	
				Initial	Follow-on
1	Mid-1930s	Many–1	Mainframe	Scientific calculation	Data processing
2	Late 1960s	1–1	PC	Spreadsheet	Database management, document processing
3	Late 1980s	1–many	Inch/foot/yard	Calendar and contact management, human–human communication	Location-based services, social media, app ecosystem, education
4	Mid-2000s	Many–many	Cloud/crowd/shroud	Personal navigation and entertainment	Health advisors, educational assistants, supply chain logistics

communities; however, no research community has yet adopted the combination suggested by collective computing. The window of opportunity is open now.

A HISTORICAL PERSPECTIVE

One motivation for defining a fourth generation of computing is the time that has passed since Weiser defined the third generation of ubiquitous computing. Inspired by his framework, Table 1 summarizes the evolution of computing generations since the 1930s, associated changes in the human–computer relationship, canonical devices representing each generation, and, finally, the driving applications that encouraged and then leveraged wide-scale adoption of those technologies.² In progressing from one generation to the next, the previous generations’ devices and applications do not disappear but, rather, are augmented by those of the next generation.

Generation 1: The mainframe

Automated computing’s origins can be traced back many centuries. The first vision leading to practical implementation was Alan Turing’s 1936 formulation of a computational engine, the so-called Turing machine, which has since influenced the architecture of computational devices.³ Turing’s ideas and work inspired the creation of automated computing machines during World War II. The assumption of the human–computer relationship was that a

single “mainframe” device would support many individuals, initially one at a time but eventually in seemingly simultaneous fashion. Ironically, the initial “killer app” for this first generation of mainframe computing was to help military powers decrypt enemy messages and calculate ballistics to more accurately target their forces. Once the war ended, large corporations realized that mainframes could automate much of their data processing needs. They acquired their own mainframes to support business activities involving important but tedious calculations, such as accounting.

Generation 2: The PC

By the late 1960s, visionaries like J.C.R. Licklider recognized opportunities for boosting human performance through enhanced connection to computation. Douglas Engelbart’s NLS/Augment project, famously demonstrated in 1968, showed for the first time how computing could augment human cognitive and communicative capabilities. Alan Kay and his Xerox PARC colleagues—inspired by the Ethernet, raster displays, and laser printing—created the first examples of a “personal” computer. This device transformed the human–computer relationship into one where each individual had his or her own computational device.

While these visions and prototypes explored applications for every individual, the PC industry did not take off

until the adoption of the spreadsheet for use in businesses. Using the metaphor of the accountant’s ledger, electronic spreadsheets became an essential tool for accounting and forecasting functions. Once businesses had invested in PCs for many of their employees, follow-on applications such as database management and document processing programs leveraged this investment and encouraged further purchases, with PCs eventually moving into homes.

Generation 3: Ubiquitous computing

By the late 1980s, personal computing had taken hold and new visionaries were dreaming of what was to come next. Weiser first articulated a computing revolution by claiming that the human–computer relationship would lure individuals to own and interact with multiple devices. Weiser, as well as Ken Sakamura (University of Tokyo), Andy Hopper (Olivetti Research Laboratory), and William Newman and Michael Lamming (RankXerox EuroPARC), also envisioned computational devices of different sizes and capabilities. Weiser used the analogy of inch-, foot-, and yard-scale devices that differed not only in size but also in mobility and ownership.

Two applications spurred ownership of inch-scale devices in the mid-1990s. First, simplified synchronization of PC-based calendar and contact information to pocket-sized PDAs pushed the sale of those devices to busy, highly mobile

professionals. However, those devices did not exploit connectivity beyond wired synchronization to a PC. The first driving application for connected inch-scale devices arrived a few years later as people across the globe began purchasing mobile phones to support more flexible human-human communication. Mobile phones (which have morphed into smartphones) became more ubiquitous than PCs, and text communication matched and surpassed voice communication. The next wave of follow-on applications—music being a major example—began to exploit these devices' sensing, storage, and networking capabilities.

Research in mobile and wireless computing also pointed toward similar opportunities for more portable forms of computation, adopting the tagline of “pervasive computing.”⁴ Foot-scale devices—desktops and luggable laptops—already existed in the PC generation. By the late 1990s, laptops had become smaller and easier to transport and provided reliable wireless networking capabilities. Pen computing also emerged but was not commercially successful for foot-scale devices. Touchscreen tablet devices' introduction in the mid-2000s was motivated by easier consumption of digital media through, for example, the Web and streaming media. At the yard scale, advertising has spurred the adoption of large public electronic displays while education, particularly for pre-university-aged students, has driven wide-scale adoption of interactive whiteboards.

THE FOURTH GENERATION'S CRITICAL TECHNOLOGIES

As inspiring as ubiquitous computing's vision has been, there are certainly detractors. Rather than focus on socio-technical criticisms,^{5,6} I draw attention

to missing elements in ubiquitous computing's technical definition that could be leveraged today. Weiser's vision did not really expose the opportunities to enhance interaction across individuals. Our research communities have long recognized computing's importance as a means of supporting human-human

on integration, albeit imperfect, is important going forward.

Greater opportunities for integration exist today because we have access to essentially limitless computation and storage through network-enabled services—initially in the form of the mid-1990s World Wide Web, and now

THESE TECHNOLOGIES—THE CLOUD, THE CROWD, AND THE SHROUD—DEFINE A NEW ERA OF COOPERATION BETWEEN HUMANS AND COMPUTERS.

interaction. Fourth-generation technologies directly address this gap, recognizing that many people interact with one another through many devices, and vice versa. Two and a half decades beyond the original conception of ubiquitous computing, technologies have emerged that Weiser and his contemporaries did not predict, such as the World Wide Web. The question now is which of these technologies will be the most influential in the coming years.

The cloud

Despite the current proliferation of various inch-, foot-, and yard-scale devices, their seamless integration was not sufficiently emphasized in the third generation's vision and execution. Researchers have presented interesting scenarios for integrating experiences across devices, but commercial tools still do not facilitate the development of applications that work smoothly across heterogeneous devices. Some argue that such seamless integration will never be achieved;⁵ I argue that an emphasis

the cloud.⁷ The notion of distributed computing can be traced to Philip Enslow's seminal paper defining a distributed data processing system.⁸ However, the relationship between people and information changed dramatically when the National Center for Supercomputing Applications introduced a simple way to view and author content across various computing platforms. When [Salesforce.com](https://www.salesforce.com) first demonstrated the delivery of enterprise applications as a Web service in 1999, the network's role in facilitating the delivery of computation changed. Amazon released Elastic Compute Cloud (<https://aws.amazon.com/e2>) in 2008, demonstrating the network's value in supporting both computing and storage services.

Today's cloud, which encompasses the Web, has revolutionized the delivery and access of computation and data. With networked devices all around us, information is truly at our fingertips. The cloud could eventually enable immediate and cross-device transmission of ever-rich data

types (for example, traces of human experiences through video or other physiological signals), allowing us to communicate past and present phenomena to anyone. While this capability's security and privacy implications are daunting, its enormous potential motivates consideration of appropriate social and legal protection measures.

The crowd

One benefit of the Web is the provision of easier publication mechanisms, which let groups of otherwise disconnected individuals cooperatively produce and curate large, valuable corpora such as Wikipedia. Additionally, as Luis von Ahn and Laura Dabbish's ESP game demonstrated,⁹ thousands of individuals can be playfully induced to label image content through an online, competitive game, providing information that current image-recognition

is here to stay. Computing with the crowd uses cloud computing services to connect people to problems that need solutions. Toolkits are emerging to support more programmatic use of crowd platforms, such as TurKit for Amazon Mechanical Turk.¹¹ Michael Bernstein and his colleagues have shown real-time results for complex human tasks such as copyediting or searching.^{12,13} Even tasks well beyond the scope of algorithmic solution, such as for supporting social problem-solving¹⁴ or food journaling,¹⁵ are feasible with the crowd.

Investigating the crowd as an isolated computing platform is worthwhile, but the combined algorithm-and-crowd platform should also be considered. Theoretically, what does it mean to add the crowd as a computational operator? When is the crowd a temporary replacement for an eventual algorithmic solution, and when is

experience and what machines on the periphery communicated or computed. Collective computing blurs the distinction between what is human and what is computational. For any single computational activity, will we know or care whether the result's source is the algorithm and/or the crowd?

The shroud

Collective computing also amplifies the connection between the physical and digital worlds through two technology trends that have emerged in the past 20 years. The first is wearable computing. Notions of the cyborg have existed in science fiction and futuristic nonfiction since the 1960s, but it was not until wearable computing devotees arose in the late 1990s (primarily out of MIT) that wearable computing became reality, albeit initially for a fringe community. The second trend is the IoT—connected devices that communicate information to other devices and the cloud—a term that Kevin Ashton coined for a 1998 presentation to Proctor & Gamble.¹⁶

Merging wearable devices with the IoT results in the Internet of Nouns (IoN)—a digital universe of people, places, and things that mirror their counterparts in the physical world. The IoN constitutes a shroud enveloping the physical world that continuously updates and reacts to changes in that world, facilitated by the cloud.

Fitness tracking has produced various small devices that, when attached to the body, monitor physical activity. Body-mounted cameras were initially driven by the desire to capture first-person-perspective video easily and continuously. Head- and wrist-mounted displays are fed by a need for quicker access to information. Easy-to-use devices (for example, the Nest

**COLLECTIVE COMPUTING BLURS THE
DISTINCTION BETWEEN WHAT IS HUMAN
AND WHAT IS COMPUTATIONAL.**

algorithms are unable to produce. Von Ahn coined the term *human computation* to refer to systematically harnessing human intelligence as a complement to or substitute for algorithmic computation on devices.¹⁰

Whether through persuasive or playful means of generating and curating content, or through paid services like Amazon Mechanical Turk (www.mturk.com), the crowd's power as a "new" computational element

it an essential element of any solution? When creativity is involved, the crowd might always be needed.

Collective computing emphasizes two forms of computation: traditional machine-based (the algorithm) and human-generated (the crowd). Weiser's ubiquitous computing vision blurred the distinction between the physical and the digital, which tagging technologies made possible, but always distinguished between the human

thermometer or Philips Hue lighting) are bringing home automation to the masses. The last decade's human body augmented by a powerful pocket-size smartphone will face competition from a physically disaggregated collection of devices screaming for an aggregated interaction experience.

The Weiserian vision of a physical universe calmly enhanced by peripheral services must be updated to include a shroud that spontaneously reconfigures itself to deliver the right interaction experience for one or more humans. Roy Want and his colleagues' Personal Server was one approach to this challenge, although it was limited to a single individual's perspective and did not address situations in which places or things need to co-opt resources.¹⁷

Currently, the shroud has many holes, which compelling use cases will fuel the drive to patch. This brings us to the next defining feature of the fourth generation of collective computing: the applications.

THE FOURTH GENERATION'S DRIVING APPLICATIONS

Numerous applications already drive adoption of the core technologies of collective computing: software as a service, commoditized computing cycles, and storage services exploit the cloud, while social networks, crowdsourced information repositories, a growing economy of microwork epitomized by Amazon Mechanical Turk, and microfunding leverage the power of the crowd. Tools have emerged for cloud creation and management as well as various social and crowd computing tasks. Although less pronounced, health, convenience, and curiosity are driving wearable technologies, and environmental automation

is motivating place- and thing-based adoption of shroud technologies. Some complex and simplified tools exist for managing things and places; people-based wearable technology is a clear opportunity to improve aggregated behaviors.

Beyond these isolated advances, we should focus on applications that combine fourth-generation technologies to demonstrate the true power of collective computing; that is, to harness information and expertise in real time, allowing people to solve hard problems for which they have no previous training.

Initial drivers: personal navigation and entertainment applications

Personal navigation devices have moved from being a luxury to a near necessity. Without them, we are often literally lost. But with them, we travel confidently, often alone, in places otherwise foreign to us. Similarly, when we want to discover potential leisure activities, we use online resources such as Yelp or Fandango to help us select options and then navigate to them. Vehicles can even sense when our usual route to work is congested and suggest an alternative, indicating how much time will be saved. All of this is made possible by combining cloud, crowd, and shroud technologies.

Navigational devices rely on location technologies, such as GPS and cellular telephony, but supplementing these with other technologies—for example, services like Waze that provide real-time traffic conditions (crowd + cloud) and environmental and bus-mounted proximity beacons to continuously update bus schedules (shroud)—often dramatically improves the driving or public transportation experience.

Similarly, crowdsourced reviews and online social networks help us decide what to experience and when, as well as inform us who might be there to share the experience. These services can even determine the best time to take a break during an event, information that would be best presented on an immediately available display (shroud).

Collective computing's initial driving application is personal navigation and entertainment, but there is room for improvement, particularly regarding shroud technologies. VizWiz shows how people with visual impairments can use a smartphone to garner immediate information from the crowd about environmental surroundings.¹⁸ Wearable devices such as smartwatches and head-mounted displays allow faster access to information through visual and haptic interaction modes.

Follow-on applications: harnessing real-time expertise

Opportunities for collective computing exist in other domains as well, including health, education, and commerce. Each example demonstrates a different take on a more general idea of allowing individuals and organizations to harness real-time expertise.

Health. Lee Rainie's recent Pew Foundation for Internet Studies report reveals some intriguing facts about how we use technology to address our health concerns:¹⁹

- › 59 percent of US adults looked online for health information in 2012.
- › 35 percent of US adults are "online diagnosers," using the Internet to figure out what medical condition they or another might have.

- › 53 percent of online diagnosers talked with a clinician about what they found online.
- › 41 percent of online diagnosers had their condition confirmed by a clinician.

Clearly, the Internet allows individuals to harness medical information. A variety of online forums also support knowledge gathering for wellness. There are valid concerns about the veracity and safety of such independent medical “research.” But rather than cautioning against this independence, we should consider how collective computing can support an individual’s desire and even society’s economic necessity to spontaneously gather actionable medical advice. While some might exclaim that this is the wrong kind of “killer app” (pun intended), I argue that we will want and need this capability, provided that we can trust the sources providing us answers and advice and that privacy and liability safeguards are established. I am not suggesting we replace trained health professionals, but rather that we effectively and efficiently extend their expertise and services to patients.

We are increasingly able to collect clinically meaningful data outside of clinical spaces such as the doctor’s office and hospital. Soon, such data will dwarf that received from clinical spaces because devices can be placed on or near the body (shroud) to continuously gather relevant physiological evidence. Place-based shroud technologies collect other relevant environmental information, and human-human communication in various modalities such as speech or text can reveal emotional and social states. These data can be collected, reviewed,

and shared via the cloud. Given the right set of data, an expert should be able to provide a diagnosis remotely, and we can certainly imagine a worldwide network of experts (a specialized doctoring crowd) that provides speedy for-fee responses to data and patients in search of a diagnosis. All of this is in addition to the empathetic healing and advice that can be garnered from other “patients” through social networks.

Education. Education has been a major force in the adoption of ubiquitous computing’s yard-scale technology. Massively open online courses, or MOOCs (cloud), are a potentially democratizing and disruptive technology application for delivering appropriately vetted educational material (although there is debate about just how democratizing MOOCs are in practice). They are often combined with social networking services so that students across the globe can learn collaboratively. These technologies could be complemented with shroud technologies that help identify teachable moments or alert students to an expert’s presence in their physical or online vicinity. Learning in the moment is similar in spirit to searching for and finding entertainment advice. And output-oriented shroud technologies can deliver small pieces of information to support microlearning.

Commerce. The previous applications were examples of humans and computing working together to assist human needs. In our collective view, it makes sense to consider opportunities for humans to improve computing. The world is already driven by IT infrastructure, and an organization’s most important and powerful people are often those who operate and support


that infrastructure. With increased data analytics capabilities to support supply-chain management and logistics, companies are using shroud technologies from the ubiquitous computing generation—for example, barcodes, QR codes, and RFID and near-field communication tags—to track inventory from production and distribution to purchase and beyond. Large, sophisticated retailers such as Walmart now regularly use crowds to help refine data to make supply chains more efficient and profitable. Collective computing could grant other retailers access to those valuable services through the cloud to aggregate relevant product information.

A fourth generation of computing has emerged over the past decade, one that is distinctly different from the third generation of ubiquitous computing. The revolutionizing technologies of the cloud, the crowd, and the shroud—developed and exploited mostly independently—together present major new opportunities for research and commercial activity in an era of collective computing. A key feature of collective computing is the blurring of human and computational elements; we no longer need concern ourselves with whether answers come from a collection of computational elements, humans, or both. In fact, harnessing the crowd could help machines communicate with human intelligence as efficiently as they do with other machines. A new class of applications that enables individuals and organizations to spontaneously garner expertise awaits exploration in many domains.

Individual cloud, crowd, and shroud research communities already

ABOUT THE AUTHOR

GREGORY D. ABOWD is a Regents' and Distinguished Professor in the School of Interactive Computing at Georgia Tech. His research focuses on the human-centered application of computing technologies to everyday challenges in domains as wide-ranging as education, health, and developmental disabilities. Abowd received a D.Phil. in computation from the University of Oxford. He is a member of IEEE and a Fellow of ACM. Contact him at abowd@gatech.edu.

exist. However, their isolation might inhibit collective computing, whose goal is to combine constituent technologies to improve both computing and the human experience. Research and industry must work together to apply the lessons of the three previous computing generations to the fourth generation.²⁰ But we must hurry—the fifth generation is probably only a decade away! 

ACKNOWLEDGMENTS

While the concept of fourth-generation computing is my own, it has been greatly influenced by two decades of pursuing Mark Weiser's original ubiquitous computing vision. Mark's writings have deeply influenced me, as have my collaborations with many graduate students and colleagues over the years; limited space prevents me from giving everyone the explicit acknowledgment they deserve.

REFERENCES

1. M. Weiser, "The Computer for the 21st Century," *Scientific American*, vol. 265, no. 3, 1991, pp. 78–89.
2. M. Weiser, "Some Computer Science Issues in Ubiquitous Computing," *Comm. ACM*, vol. 36, no. 7, 1993, pp. 75–84.
3. A.M. Turing, "On Computable Numbers, with an Application to the Entscheidungsproblem," *Proc. London Math. Soc.*, ser. 2, vol. 42, 1936, pp. 230–265.
4. M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Comm.*, vol. 8, no. 4, 2001, pp. 10–17.
5. G. Bell and P. Dourish, "Yesterday's Tomorrows: Notes on Ubiquitous Computing's Dominant Vision," *Personal and Ubiquitous Computing*, vol. 11, no. 2, 2007, pp. 133–143.
6. Y. Rogers, "Moving on from Weiser's Vision of Calm Computing: Engaging Ubicomp Experiences," *Proc. 8th Int'l Conf. Ubiquitous Computing (UbiComp 06)*, 2006, pp. 404–421.
7. M. Armbrust et al., "A View of Cloud Computing," *Comm. ACM*, vol. 53, no. 4, 2010, pp. 50–58.
8. P.H. Enslow, "What is a 'Distributed' Data Processing System?," *Computer*, vol. 11, no. 1, 1978, pp. 13–21.
9. L. von Ahn and L. Dabbish, "Labeling Images with a Computer Game," *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI 04)*, 2004, pp. 319–326.
10. E. Law and L. von Ahn, "Human Computation," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 5, no. 3, 2011, pp. 1–121.
11. G. Little et al., "TurKit: Human Computation Algorithms on Mechanical Turk," *Proc. 23rd Ann. ACM Symp. User Interface Software and Technology (UIST 10)*, 2010, pp. 57–66.
12. M.S. Bernstein et al., "Crowds in Two Seconds: Enabling Realtime Crowd-Powered Interfaces," *Proc. 24th Ann. Symp. User Interface Software and Technology (UIST 11)*, 2011, pp. 33–42.
13. M.S. Bernstein et al., "Soylent: A Word Processor with a Crowd Inside," *Proc. 23rd Ann. ACM Symp. User Interface Software and Technology (UIST 10)*, 2010, pp. 313–322.
14. F. Boujarwah, G. Abowd, and R. Arriaga, "Socially Computed Scripts to Support Social Problem Solving Skills," *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI 12)*, 2012, pp. 1987–1996.
15. E. Thomaz et al., "Feasibility of Identifying Eating Moments from First-Person Images Leveraging Human Computation," *Proc. 4th Int'l SenseCam and Pervasive Imaging Conf. (SenseCam 13)*, 2013, pp. 26–33.
16. G. Santucci, "From Internet of Data to Internet of Things," *Proc. 4th Int'l Conf. Future Trends of the Internet*, 2009; http://cordis.europa.eu/pub/fp7/ict/docs/enet/20090128-speech-iot-conference-lux_en.pdf.
17. R. Want et al., "The Personal Server: Changing the Way We Think about Ubiquitous Computing," *Proc. 4th Int'l Conf. Ubiquitous Computing (UbiComp 02)*, 2002, pp. 194–209.
18. J.P. Bigham et al., "VizWiz: Nearly Real-Time Answers to Visual Questions," *Proc. 23rd Ann. ACM Symp. User Interface Software and Technology (UIST 10)*, 2010, pp. 333–342.
19. L. Rainie, "E-patients and Their Hunt for Health Information," Pew Research Center, 26 July 2013; www.pewinternet.org/2013/07/26/e-patients-and-their-hunt-for-health-information-2.
20. G.D. Abowd, "What Next, UbiComp? Celebrating an Intellectual Disappearing Act," *Proc. ACM Conf. Ubiquitous Computing (UbiComp 12)*, 2012, pp. 31–40.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.