Perspectives on Pervasive Sensing Informed Through Applications in the Internet of Things

Joseph A. Paradiso, Responsive Environments Group, MIT Media Lab

The Internet of Things assumes ubiquitous sensate environments. Without these, the cognitive engines of this everywhere-enabled world are deaf, dumb, and blind, and can't respond relevantly to the real world events that they aim to augment. Advances over the past decade have been rampant, as sensors tend to exploit Moore's Law. Accordingly, sensors of all sorts now seem to be increasingly in everything, indicating an approaching phase transition once they are properly networked, like we saw when web browsers appeared and our interaction with computers fundamentally changed. This shift will create a seamless electronic nervous system that covers the planet – and one of the main challenges for the computing community now is in how to merge the rapidly evolving "omniscient" electronic sensoria onto human perception. This article tours aspects of this coming revolution guided by several recent and ongoing projects in author's research group at the MIT Media Lab that approach this precept from different perspectives. We give examples that range from smart buildings to sports, exploiting technical trends ranging from wearable computing to wireless sensor networks.

Introduction

Just as Vannevar Bush foresaw multimedia, the personal computer and so much more in his seminal 1945 article [41], Mark Weiser predicted the Internet of Things in a similarly seminal article in 1991 [42]. Whereas Bush's article, rooted in extrapolations of WW II technology, was mainly speculative, Weiser illustrated the future that he predicted via actual working systems that he and his colleagues had developed (and were already living with) at Xerox PARC. Coming from an HCI perspective, Weiser and his team looked at how people would interact with networked computation distributed into the environments and artifacts around them. Now upstaged by the 'Internet of Things,' which much more recently gestated in industrial/academic quarters enamored with evolving RFID technology, Weiser's vision of 'Ubiquitous Computing' resonated with many of us back then working in the domains of computer science involved with human interaction. Even within our community and before the IoT moniker dominated, this vision could take many names and flavors as factions tried to establish their own brand (we at the Media Lab called it 'Things That Think', my colleagues at MIT's Lab for Computer Science labeled it 'Project Oxygen', others called it 'Pervasive Computing', 'Ambient Computing', 'Invisible Computing', 'Disappearing Computer', etc.), but it was all still rooted in Mark's early vision of what we often abbreviate as 'UbiComp'.

Appears in *The Next Step: Exponential Life*, Open Mind Collection, BBVA Press, 2016, pp. 47-75.

I first came to the MIT Media Lab in 1993 as a High-Energy Physicist working on elementary particle detectors (which can be considered to be extreme 'sensors') at the then-recently-cancelled Superconducting Supercollider and the nascent Large Hadron Collider at CERN. Since I had been facile with electronics since childhood and had been designing and building large music synthesizers since the early 70s [43], my earliest Media Lab research involved incorporating new sensor technologies into novel human-computer interface devices, many of which were used for musical control [44]. As wireless technologies became more easily embeddable, my research group's work evolved by the late 90s into various aspects of wireless sensors and sensor networks, again with a human-centric focus [45]. We now see most of our projects as probing ways in which human perception and intent will seamlessly connect to the electronic 'nervous system' that increasingly surrounds us. This article traces this theme from a sensor perspective, using projects from my research group to illustrate how revolutions in sensing capability are transforming how humans connect with their environs in many different ways.

Inertial Sensors and Smart Garments

Humans sense inertial and gravitational reaction via proprioception at the limbs and vestibular sensing at the inner ear. Inertial sensors, the electrical equivalent, once used mainly for high-end applications like aircraft and spacecraft navigation [46], are now commodity. MEMS (Micro Electro-Mechanical Systems) based accelerometers date back to academic prototypes in 1989 and widespread commercial chips coming from Analog Devices in the 1990s [47] that exploited their proprietary technique of fabricating both MEMS and electronics on the same chip. Still far too coarse for navigation applications beyond basic pedometry [48] or step-step dead reckoning [49], they are appearing in many simple consumer products, e.g., toys, phones, and exercise dosimeters and trackers, where they are mainly used to infer activity or for tilt-based interfaces, like flipping a display. Indeed, the accelerometer is perhaps now the most generically-used sensor on the market and will be embedded into nearly every device that is subject to motion or vibration, even if only to wake up more elaborate sensors at very low power (an area we explored very early on, e.g., [50,51]). MEMS-based gyros have advanced well beyond the original devices made by Draper Lab [52] back in the 1990s. They have also become established, but their higher costs and higher power drain (because they need a driven internal resonator, they haven't pushed below the order of 1 mA current draw and are harder to efficiently duty-cycle as they take significant time to start up and stabilize [51]) preclude quite as much market penetration at this point. Nonetheless, they're becoming common on smart phones, AR headsets, cameras, and even watches & high-end wristbands, for gesture detection, rotation sensing, motion compensation, activity detection, augmented reality applications, etc., generally paired with magnetometers that give a ground-truth orientation using Earth's magnetic Field. Triaxial accelerometers are now a default, and triaxial gyros have also been under manufacture. We've seen since 2010 commercial products that do both, providing an entire 6-axis inertial measurement unit on a single chip together with 3-axis magnetometer (e.g., starting from Invensense with the MPU-6000/6050 and with more recent devices from Invensense, ST Microelectronics, Honeywell, etc.) [53].

The MIT Media Lab have been pioneers in applications of inertial sensing to user interface and wearable sensing [54] and many of these ideas have gestated in my group going back 20 years [55]. Starting a decade ago, piggybacking on even older project exploring wearable wireless sensors for dancers [56], my team has employed wearable inertial sensors for measuring biomechanical parameters for professional baseball players in collaboration with physicians from the Massachusetts General Hospital [1]. As athletes are capable of enormous ranges of motion, we need to be able to measure peak activities of up to 120 G, and 13,000°/s, while still having significant resolution within the 10 G, 300°/s range of normal motion. Ultimately, this would entail an IMU with log sensitivity, but as these aren't yet available, we've integrated dual-range inertial components onto our device and mathematically blend both signals. We use the radio connection primarily to synchronize our devices to better than our 1 ms sampling rate, and write data continuously to removable flash memory, allowing us to gather data for an entire day of measurement and analyze the data to determine descriptive and predictive features of athletic performance [57]. A myriad of commercial products are now appearing for collecting athletic data for a variety of sports that range from baseball to tennis [58], and even the amateur athlete of the future will be instrumented with wearables that will aid in automatic/self/augmented coaching.

Accelerometer power has also dropped enormously, with a modern capacitive 3-axis accelerometers drawing well under a milliamp, and special purpose accelerometers (e.g., devices originally designed for implantables [2]) pulling under a microamp. Passive piezoelectric accelerometers have also found application in UbiComp venues [3].

Sensors of various sorts have also crept into fabric and clothing – although healthcare and exercise monitoring have been the main beacons calling researchers to these shores [4], fabric-based strain gauges, bend sensors, pressure sensors, bioelectric electrodes, capacitive touch sensors, and even RFID antennae have been developed and become increasingly reliable and robust [5], making clothing into a serious user interface endeavor. Recent high-profile projects, such as Google's Project Jacquard [59], led technically by Nan-Wei Gong, an alum of my research team, boldly indicate that smart clothing and weaving as electronic fabrication will evolve beyond narrow and boutique niches and into the mainstream as broad applications, perhaps centered on new kinds of interaction [6], begin to emerge. Researchers worldwide, such as my MIT colleague Yoel Fink [60] and others are re-thinking the nature of 'thread' as a silicon fiber upon which microelectronics are grown. Acknowledging the myriad hurdles that remain here in robustness, production, etc., the research community is mobilizing [61] to transform these and other ideas around flexible electronics into radically new functional substrates that can impact medicine, fashion, and apparel (from recreation to spacesuits and military or hazardous environments), and bring electronics into all things stretchable or malleable.

Going beyond wearable systems are electronics that are attached directly to or even painted onto the skin. A recent postdoc in my group, Katia Vega, has pioneered this in the form of what she calls 'Beauty Technology' [62]. Using functional makeup that construes conductors, sensors, and illuminating or color-changing actuators [63], Katia and her colleagues have made skin interfaces that react to facial or tactile gesture, sending wireless messages, illuminating, or changing apparent skin color with electro-active rouge. Katia has recently explored medical implementations, such as tattoos that can change appearance with blood-borne indicators [64].

We are living in an era that has started to technically mine the wrist, as seen in the now ubiquitous smart watch. These devices are generally touchscreen peripherals to the smartphone, however, or work as improved activity monitors such as have been in the market since 2009. I believe that the wrist will become key to the user interface of the future, but not implemented like in today's smart watches. User interfaces will exploit wrist-worn IMUs and precise indoor location to enable pointing gestures [65], and finger gesture will still be important in communicating (as we have so much sensory-motor capability in our fingers), but rather than using touch screens, we will be tracking finger gesture by sensors embedded in smart wristbands with our hands by our sides. In recent years, my team has developed several early working prototypes that do this, one using passive RFID sensors in rings worn at the upper joint of fingers interrogated by a wrist-worn reader [66], another using a short-range 3D camera looking across the palm from the wrist to the fingers [67], and another that uses a pressure-imaging wristband that can sense tendon displacement [68]. My student Artem Dementyev and his collaborator Cindy Kao have also made a wireless touchpad into a fake fingernail, complete with a battery that lasts a few hours of continuous use [69], enabling digital interaction at the extremes of our digits, so to speak. Hence, I envision the commuter of the future not staring down at their smartphone in their crowded self-driving bus, but rather looking straight ahead into their head-mounted wearable, nudging their heavily context-leveraged communication via finger movements with hand resting innocuously at their side.

My former student Mat Laibowitz and I introduced the concept of 'Parasitic Mobility' over a decade ago [70]. We proposed that low power sensor nodes could harvest 'mobility' as opposed to just garnering energy, hopping on and off moving agents in their environment. In this era of drones deployed from delivery vans, such a concept is now far from radical. Our examples used people as the mobile agents, and the nodes were little robots that could choose to be 'worn' by attaching in various ways to their hosts, inspired by ticks, fleas, etc. in the natural world, then detaching at the location of their choice. My current student Artem Dementyev, together with Sean Follmer at Stanford and other collaborators, have recently pioneered a different approach to wearable robots in their 'Rovables' [71]. These are very small (e.g., 2 cm) robots that navigate across a user's clothing, choosing the best location to perform a task, ranging from a medical measurement to a dynamic display. This project displays an extreme form of human-robot closely proximate collaboration that could bear profound implication.

Everywhere Cameras and Ubiquitous Sensing

In Orwell's 1984[72], it was the totalitarian Big Brother government who put the surveillance cameras on every television – but in the reality of 2016, it's consumer electronics companies who build cameras into the common set-top box and every mobile handheld. Indeed, cameras are becoming commodity, and as video feature extraction gets to lower power levels via dedicated hardware, and other micropower sensors determine the necessity of grabbing an image frame, cameras will become even more common as generically embedded sensors. The first commercial, fully-integrated CMOS camera chips came from VVL in Edinburgh (now part of ST Microelectronics) back in the early 90s [7]. At the time, pixel density was low (e.g., the VVL "Peach" with 312 x 287 pixels), and the main commercial application of their devices was the "BarbieCam," a toy video camera sold by Mattel. I was an early adopter of these digital cameras myself, using them in 1994 for a multi-camera precision alignment system at the Superconducting Supercollider [73] that evolved into the hardware used to continually align the 40- meter muon system at micron-level precision for the ATLAS detector at CERN's Large Hadron Collider (LHC). This technology was poised for rapid growth – now, integrated cameras peek at us everywhere, from laptops to cellphones, with typical resolutions of scores of megapixels and bringing computational photography increasingly to the masses. ASICs for basic image processing are commonly embedded with or integrated into cameras, giving increasing video processing capability for ever-decreasing power. The mobile phone market has been driving this effort, but increasingly static situated installations (e.g., video-driven motion/context/gesture sensors in smart homes) and augmented reality will be an important consumer application, and the requisite on-device image processing will drop in power and become more agile. We already see this happening at extreme levels, such as with the recently-released Microsoft HoloLens, which features 6 cameras, most of which are used for rapid environment mapping, position tracking, and image registration in a lightweight battery-powered head-mounted self-contained AR unit. 3D cameras area also becoming ubiquitous, breaking into the mass market via the original structured-light-based Microsoft Kinect a half-decade ago. Time-of-flight 3D cameras (pioneered in CMOS in the early 2000s by researchers at Canesta [74]) have evolved to recently displace structured light approaches, and developers worldwide race to bring the power and footprint of these devices down sufficiently to integrate into common mobile devices (a very small version of such a device is already embedded in the HoloLens). As pixel timing measurements become more precise, photon-counting applications in computational photography, as pursued by my Media Lab colleague Ramesh Raskar, promise to usher in revolutionary new applications that can do things like reduce diffusion and see around corners [75].

My research group began exploring this penetration of ubiquitous cameras over a decade ago, especially applications that ground the video information with simultaneous data from wearable sensors. Our early studies were based around a platform called the "Portals" [8] – using an embedded camera feeding a TI DaVinci DSP/ARM hybrid processor, surrounded by a core of basic sensors (motion, audio, temperature/humidity, IR proximity) and coupled with a Zigbee RF transceiver, we scattered 45 of these devices all over the Media Lab complex, interconnected through the wired building network. One application that we built atop them was "SPINNER" [9], which labeled video from each camera with data from any wearable sensors in the vicinity. The SPINNER framework was based on the idea of being able to query the video database with higher-level parameters, lifting sensor data up into a social/affective space [10], then trying to effectively script a sequential query as a simple narrative involving human subjects adorned with the wearables. Video clips from large databases sporting hundreds of hours of video would then be automatically selected to best fit given timeslots in the query, producing edited videos that observers deemed coherent [9]. Naively pointing to the future of reality television, this work aims further, looking to enable people to engage sensor systems via human-relevant query and interaction.

Rather than try to extract stories from passive ambient activity, a related project from our team devised an interactive camera with a goal of extracting structured stories from people [11]. Taking the form factor of a small mobile robot, "Boxie" featured an HD camera in one of its eyes – it would rove our building and get stuck, then plea for help when people came nearby. It would then ask people successive questions and request that they fulfill various tasks (e.g., bring it to another part of the building, or show it what they do in the area where it was found), making an indexed video that can be easily edited to produce something of a documentary about the people in the robot's abode.

In the next years, as large video surfaces cost less (potentially being roll-roll printed) and are better integrated with responsive networks, we'll see the common deployment of pervasive interactive displays. Information coming to us will manifest in the most appropriate fashion (e.g., in your smart eyeglasses or on a nearby display) – the days of pulling your phone out of your pocket and running an app are severely limited. To explore this, we ran a project in my team called 'Gestures Everywhere' [76] that exploited the large monitors placed all over the public areas of our building complex [77]. Already equipped with RFID to identify people wearing tagged badges, we added a sensor suite and a Kinect 3D camera to each display site. As an occupant approached a display and were identified via RFID or video recognition, information most relevant to them would appear on the display. We developed a recognition framework for the Kinect that parsed a small set of generic hand gestures (e.g., signifying 'next', 'more detail', 'go-away', etc.), allowing users to interact with their own data at a basic level without touching the screen or pulling out a mobile device. Indeed, proxemic interactions [30] around ubiquitous smart displays will be common within the next decade.

The plethora of cameras that we sprinkled throughout our building during our SPINNER project produced concerns about privacy (interestingly enough, the Kinects for Gestures Everywhere didn't evoke the same response – occupants either didn't see them as 'cameras' or were becoming used to the idea of ubiquitous vision). Accordingly, we put an obvious power switch on each portal that enabled them to be easily switched off. This is a very artificial solution, however – in the near future, there will just be too many cameras and other invasive sensors in the environment to switch off. These devices must answer verifiable and secure protocols to dynamically and appropriately throttle streaming sensor data to answer user privacy demands. We have designed a small, wireless token that controlled our portals in order to study solutions to such concerns [12]. It broadcast a beacon to the vicinity that dynamically deactivates the transmission of proximate audio, video, and other derived features according to the user's stated privacy preferences – this device also featured a large 'panic' button that can be pushed at any time when immediate privacy is desired, blocking audio and video from emanating from nearby Portals.

Rather than block the video stream entirely, we have explored just removing the privacy-desiring person from the video image. By using information from wearable sensors, we can more easily identify the appropriate person in the image [13], and blend them into the background. We're also looking at the opposite issue – using wearable sensors to detect environmental parameters that hint at potentially hazardous conditions for construction workers and rendering that data in different ways atop realtime video, highlighting workers in situation of particular concern [78].

Energy Management

The energy needed to sense and process has steadily declined – sensors and embedded sensor systems have taken full advantage of low power electronics and smart power management [14]. Similarly, energy harvesting, once an edgy curiosity, has become a mainstream drumbeat that's resonating throughout the embedded sensor community [15,79]. The appropriate manner of harvesting is heavily dependent on the nature of the ambient energy reserve. In general, for inhabited environments, solar cells tend to be the best choice (providing over 100 $\mu A/cm^2$ indoors, depending on light levels, etc.). Otherwise, vibrational and thermoelectric harvesters can be exploited where there is sufficient vibration or heat transfer (e.g., in vehicles or on the body) [16, 17], and implementations of ambient RF harvesting [18] have not become uncommon either. In most practical implementations that constrain volume and surface area, however, an embedded battery can be a better solution over an anticipated device's lifetime, as the limited power available from the harvester can already put a strong constraint on the current consumed by the electronics. Another solution is to beam RF or magnetic energy to batteryless sensors, such as with RFID systems. Popularized in our community starting with Intel's WISP [19] (used by my team for the afore-mentioned finger-tracking rings

[20], commercial RF-powered systems begin to appear (e.g., Powercast, WiTricity, etc.), although each has it's caveats (e.g., limited power or high magnetic flux density). Indeed, it's common now to see inductive 'pucks' available at popular café's to charge cellphones when placed above the primary coils embedded under the table, and certain phones have inductive charging coils embedded behind their case, but these seem to be seldom used as it's more convenient to use a ubiquitous USB cable/connector, and the wire eliminates position constraints that proximate inductive charging can impose. In my opinion, due to its less-than-optimal convenience, high lossiness (especially relevant in our era of conservation), and potential health concerns (largely overstated [80], but more important if you have a pacemaker, etc.), wireless power transfer will be restricted mainly to niche and low-power applications like micro-power sensors that have no access to light for photovoltaics, and perhaps special closets or shelves at home where our wearable devices can just be hung up to be charged.

The boom in energy harvesting interest has precipitated a variety of ICs dedicated to managing & conditioning the minute amounts of power these things typically produce (e.g., by major IC providers like Linear Technology and Maxim), and the dream of integrating harvester, power conditioning, sensor, processing, and perhaps wireless on a single chip nears reality (at least for extremely low duty cycle operation or environments with large reservoirs of ambient energy). Indeed, the dream of such Smart Dust [111] has been realized in early prototypes, such as the University of Michigan wafer stack that uses its camera as an optical energy harvester [21] and the recently announced initiative to produce laser-accelerated, gram-level 'Star Wisp' swarm sensors for fast interstellar exploration [81] will emphasize this area even more.

Although energy harvesting is often confused with sustainable energy sources by the general public (e.g., [22]), the amount of energy harvestable in standard living environments is far too small to make any real contribution to society's power needs. On the other hand, very low power sensor systems, perhaps augmented by energy harvesting, minimize or eliminate the need for changing batteries, thereby decreasing deployment overhead, hence increasing the penetration of ubiquitous sensing. The information derived from these embedded sensors can then be used to intelligently minimize energy consumption in built environments, promising a better degree of conservation and utility than attainable by today's discrete thermostats and lighting controls.

Dating back to the Smart Home of Michael Moser [23], work of this sort has recently become something of a lightning rod in Ubiquitous Computing circles, as several research groups have launched projects in smart, adaptive energy management through pervasive sensing [24]. For example, in my group, my then-student Mark Feldmeier has used inexpensive chip-based piezo cantilevers to make a micropower integrating wrist-worn activity detector (before the dawn of commercial exercise-tracker wristbands) & environmental monitor for a wearable HVAC controller [3]. Going beyond subsequently-introduced commercial systems such as the well-known

Nest Thermostat, the wearable provides a first-person perspective that more directly relates to comfort than a static wall-mounted sensor, plus is intrinsically mobile, exerting HVAC control in any suitably equipped room or office. Exploiting analog processing, this continuous activity integrator runs at under 2 microamperes, enabling our device (together with its other components, which also measure and broadcast temperature, humidity, and light level each minute) to run for over 2 years on a small coin cell battery, updating once per minute when within range of our network. We have used the data from this wristworn device to essentially graft onto the user's "sense of comfort", and customize our building's HVAC system according to the comfort of the local users, as inferred from the wearable sensor data discriminated by "hot" & "cold" labels, obtained by pushing buttons on the device. We estimated that the HVAC system running under our framework used roughly 25% less energy, and inhabitants were significantly more comfortable [3]. Similar wireless micropower sensors for HVAC control have recently been developed by others, such as Schneider Electric in Grenoble [25], but are used in a fixed location. Measuring temperature, humidity, and CO_2 levels every minute and uploading the readings via a Zigbee network, the Schneider unit is powered by a small solar cell – one day of standard indoor light is sufficient to power this device for up to 4 weeks (the duration of a typical French vacation).

Lighting control is also another application that can benefit from Ubiquitous Computing. This is not only for energy management purposes – we also need to answer the user interface challenge that modern lighting systems pose. Although the actuation possibilities of contemporary commonly-networked digitally-controlled fluorescent and solid-state lights are myriad, the human interface suffers greatly, as especially commercial lighting is generally managed by cryptic control panels that make getting the lighting that you want onerous at best. As we accordingly lament the loss of the simple light switch, my group has launched a set of projects to bring it back in virtual form, led by my students Nan Zhao, and alums Matt Aldrich and Asaf Axaria. One route that we have pursued involves using feedback from a small incident light sensor that's able to isolate the specific contribution of each nearby light fixture to the overall lighting at the sensor, as well as to estimate external, uncontrolled light [26]. Our control algorithm, based around a simple linear program, is able to dynamically calculate the energy-optimal lighting distribution that delivers the desired illumination at the sensor location, effectively bringing back the local light switch as a small wireless device.

Our ongoing work in this area expands on this idea by also incorporating cameras as distributed reflected-light sensors and as a source of features that can be used for interactive lighting control. We have derived a set of continuous control axes for distributed lighting via principle-component analysis that relevant to human perception [82], enabling easy and intuitive adjustment of networked lighting. We are also driving this system from the wearable sensors in Google Glass, automatically adjusting the lighting to optimize illumination of surfaces and objects you are looking at as well as deriving context to set overall lighting, for instance seamlessly transitioning between lighting appropriate for a working or a casual/social function [83]. Our most recent work in this area also incorporates large, soon-to-be ubiquitous displays, changing both images/video, sound, and lighting in order to servo the user's emotive/attentive state as detected by wearable sensors and nearby cameras [84] to healthy and productive levels.

Radio, Location, Electronic "Scent", and Sensate Media

The last decade has seen a huge expansion in wireless sensing, which is having a deep impact in ubiquitous computing. Mixed signal IC layout techniques have enabled the common embedding of silicon radios on the same die as capable microprocessors, which have ushered in a host of easy to use smart radio ICs, most of which support common standards like Zigbee running atop 802.15.4. Companies like Nordic, TI/Chipcon, and more recently Atmel, for example, make commonly used devices in this family. Although a plethora of sensor-relevant wireless standards are in place that optimize for different applications (e.g., ANT for low duty-cycle operation, Savi's DASH-7 for supply chain implementation, Wireless Hart for industrial control, Bluetooth Low Energy for consumer devices, Lorawan for longer-range IoT devices, and low power variants of WiFi), Zigbee has had a long run as a well-known standard for low-power RF communication, and Zigbee-equipped sensor modules are easily available that can run operating systems derived from sensor net research (e.g., dating to TinyOS [85]), or custom-developed application code. Multihop operation is standard in Zigbee's routing protocol, but battery limitations generally dictate that routing nodes need to be externally powered. Although it could be on the verge of being replaced by more recently-developed protocols (as listed above), my team has been basing most of its environmental and static sensor installations on modified Zigbee software stacks.

Outdoor location within a several meters has long been dominated by GPS, and more precise location is provided by differential GPS, soon to be much more common via recent low-cost devices like the RTK (Real Time Kinematic) GPS embedded radios [86]. Indoor location is another story. It is much more difficult, as GPS is shielded away, and indoor environments present limited lines-of-sight and excessive multipath. There seems to be no one clear technology winner for all indoor situations (as witnessed by the results of the Microsoft Indoor Location Competition run yearly at recent IPSN Conferences [87]), but the increasing development being thrown at it promises that precise indoor location will soon be a commonplace feature in wearable, mobile, and consumer devices.

Networks running 802.15.4 (or 802.11 for that matter) generally offer a limited location capability that utilizes RSSI (received signal strength) and radio link quality. These essentially amplitude-based techniques are subject to considerable error due to dynamic and complex RF behavior in indoor environments. With sufficient information from multiple base stations and exploitation of prior state and other constraints, these "fingerprinting" systems claim to be able to localize mobile nodes indoors to within 3 to 5 meters [27]. Time-of-flight RF systems, however, promise to make this much better. Clever means of utilizing the differential phase of RF beats in simultaneously transmitting outdoor radios has been shown to be able to localize a grid of wireless nodes to within 3 cm across a football field [28], and directional antennas achieve circa 30 cm or better accuracy with Nokia's HAIP system [29], for example. However, the emerging technology of choice in lightweight precision RF location is low power, impulse-based UltraWideBand, where short radio impulses at around 5-8 GHz are precisely timed, and cm-level accuracy is promised. High-end commercial products have been available for a while now (e.g., UbiSense for indoor location, and Zebra Technology radios for stadium-scale (sports) location), but emerging chipsets from companies like Decawave and Qualcomm (via their evolving "Peanut" radio) indicate that this technology might soon be in everything – and as soon as every device can be located to within a few cm, profound applications will arise (e.g., geofencing and proximate interfaces [30] are obvious, but this capability will impact everything).

One example that Brian Mayton in my team has produced leveraging this technology was the 'Wristque' wristband [65]. Equipped with precise indoor location (via the UbiSense system) as well as a full 9-axis IMU and sensors to implement the smart lighting and HVAC systems discussed above, the Wristque also enables control of devices by simple pointing and arm gesture. Pointing is an intuitive way by which humans communicate, and camera-based systems have explored this ever since the legacy 'Put That There' demo was developed at the MIT Media Lab in the 80s [88]. But as the Wristque knows the position of the wrist to within cm via soon-to-be-ubiquitous radio location and the wrist angle via the IMU (calibrated for local magnetic distortion), we can easily extrapolate the arm's vector to determine what a user is pointing to without the use of cameras.

Different locations also exhibit specific electromagnetic signatures or "scents," as can be detected via simple capacitive, inductive, RF, and optical sensors. These arise from parasitic emissions from power lines and various local emanations from commonly modulated lighting and other electrical equipment – these signals can also be affected by building wiring and structure. Such background signals carry information relevant to energy usage and context (e.g., what appliance is doing what and how much current it is drawing) and also location (as these near-field or limited-radius emanations attenuate quickly). Systems that exploit these passive background signals, generally through a machine learning process, have been a recent fashion in UbiComp research, and have seen application in areas like non-intrusive load management [31], tracking hands across walls [32], and inferring gesture around fluorescent lamps [33]. DC magnetic sensors have also been used to precisely locate people in buildings, as the spatial orientation of the Earth's magnetic field changes considerably with position in the building [34].

My team has leveraged using ambient electronic noise in wearables, analyzing characteristics of capacitively-coupled pickup to identify devices touched and infer their operating modes [89]. My then-student NanWei Gong and I (in collaboration with Microsoft Research in the UK) have developed a multimodal electromagnetic sensor floor with pickups and basic circuitry printed via an inexpensive roll-roll process [35]. The networked sensing cells along each strip can track people walking above via capacitive induction of ambient 60 Hz hum, as well as track the location of an active GSM phone or Near-Field transmitter via their emanations. This is a recent example of what we have termed 'Sensate Media' [90] – essentially the integration of low-power sensor networks into common materials to make them sensorially capable – essentially a scalable 'electronic skin'. Our most recent work in this area has produced Sensor Tape [91], a roll of networked embedded sensors on a flexible substrate in the form factor of common adhesive tape. Even though our sensor tape is many meters long it can be spooled off, cut, and rejoined as needed and attached to an object or surface as desired to give it sensing capability (our current tape embeds devices like IMUs and proximity sensors at a circa 1" pitch for distributed orientation and ranging). We have also explored making versions of Sensor Tape that can exploit printed sensors and be passively powered [92], enabling it to be attached to inaccessible surfaces and building materials, then wirelessly interrogated by, for example, NFC readers. In cases where sensors can't be printed, they can also be implemented as stickers with conductive adhesive, such as recently pioneered by my student Jie Qi and Media Lab alum Bunny Huang for craft and educational applications [93] in their commercially-available Chibitronics kits.

Aggregating and Visualizing Diverse Sensor Data

We've witnessed an explosion of realtime sensor data flowing through our networks – and the rapid expansion of smart sensors at all scales should soon cause real-time sensor information to dominate network traffic. At the moment, this sensor information is still fairly siloed – e.g., video is for video conferencing, telephony, and webcams, traffic data shows up on traffic websites, etc. A grand challenge for our community is to break down the walls between these niches, and develop systems and protocols that seamlessly bridge these categories to develop a virtual sensor environment, where all relevant data is dynamically collated and fused. Several initiatives have begun to explore these ideas of unifying data. Frameworks based on SensorML, Home Plug & Play, and DLNA, for example, have been aimed at enabling online sensors to easily provide data across applications. Patchube (now Xively [94]) is a proprietary framework that enables subscribers to upload their data to a broad common pool that can then be provided for many different applications. These protocols, however, have had limited penetration for various reasons, such as the desire to own all of your data and charge revenue or limited buy-in from manufacturers. Recent industrial initiatives, such as the Qualcomm-led AllJoyn [95] or the Intel-launched Open Interconnect Alliance with IoTivity [96] are more closely aimed at the Internet of Things and are less dominated by consumer electronics. Rather than wait for these somewhat heavy protocols to mature, my student Spencer Russell has developed our own, called CHAIN-API [97,98]. CHAIN is a RESTful (REpresentational State Transfer), web-literate framework, based on JSON data structures, hence is easily parsed by tools commonly available to web applications. Once it hits the internet, sensor data is posted, described, and linked to other sensor data in CHAIN. Exploiting CHAIN, applications can 'crawl' sensor data just as they now crawl linked static posts, hence related data can be decentralized and live under diverse servers (I believe that in the future no entity will 'own' your data, hence it will be distributed). Crawlers will constantly traverse data linked in protocols like CHAIN, calculating state and estimating other parameters derived from the data, which in turn will be re-posted as other 'virtual' sensors in CHAIN. As an example of this, my student David Ramsay has implemented a system called 'learnAir' [99], in which he has used CHAIN to develop a protocol for posting data from air quality sensors, as well as building a machine-learning framework that corrects, calibrates, and extrapolates this data by properly combining ubiquitous data from inexpensive sensors with globally-measured parameters that can affect its state (e.g., weather information) as well as grounding it with data from high-quality sensors (e.g., EPA certified) depending on proximity. With the rise of Citizen Science in recent years, sensors of varying quality will collect data of all sorts everywhere. Rather than drawing potentially erroneous conclusions from misleading data, learnAir points to a future where data of all sorts and of varying quality will be combined dynamically and properly conditioned, allowing individuals to contribute to a productive data commons that incorporates all data most appropriately.

Browsers for this ubiquitous, multimodal sensor data will play a major role, for example, in the setup/debugging [36] of these systems and building/facility management. An agile multimodal sensor browser capability, however, promises to usher in revolutionary new applications that we can barely speculate about at this point. A proper interface to this artificial sensoria promises to produce something of a digital omniscience, where one's awareness can be richly augmented across temporal and spatial domains – a rich realization of McLuhan's characterization of electronic media extending the human nervous system [100]. We call this process "Cross Reality" [8] – a pervasive, everywhere augmented reality environment, where sensor data fluidly manifests in virtual worlds that can be intuitively browsed. Although some researchers have been adding their sensor data into structures viewable through Google Earth (e.g., [37]), we have been using a 3D game engine to browse the data coming from our own building – as game engines are built for efficient graphics, animation, and interaction, they are perfectly suited to developing such architecturally-referenced sensor manifestations (our prior work used the communal shared-reality environment SecondLife [112], which proved much too restrictive for our aims). Our present system, called "DoppelLab" [38] and led by my student Gershon Dublon, makes hundreds of diverse sensor and information feeds related to our building visible via animations appearing in a common virtual world. Browsing in DoppelLab, we can see the building HVAC system work, see people moving through the building via RFID badges and motion sensors, see their public Tweets emanating from their offices, and even see real-time ping pong games played out on a virtual table [101]. We have also incorporated auditory browsing in DoppelLab by creating spatialized audio feeds coming from microphones placed all around our building, referenced to the position of the user in the virtual building. Random audio grains are reversed or deleted right at the sensor node, making conversations impossible to interpret for privacy's sake (it always sounds like people are speaking a foreign language), but vocal affect, laughing, crowd noise, and audio transients (e.g., elevators ringing, people opening doors, etc.) come through fairly intact [102]. Running DoppelLab with spatialized audio of this sort gives the user the feeling of being a ghost, loosely coupled to reality in a tantalizing way, getting the gist and feeling associated with being there, but removed from the details. Growing up in the late-60's era of McLuhan, this concept fascinated me, and as an elementary-school child, I wired arrays of microphones hundreds of meters away to mixers and amplifiers in my room. At the time, despite the heavy influence of then-in-vogue spy movies, I had no desire to snoop on people. I was instead just fascinated with the idea of generalizing presence and bringing remote outdoor sound into dynamic ambient stereo mixes that I'd experience for hours on end. DoppelLab and the DoppelMarsh environment described below have elevated this concept to a level I could have never conceived of in my analog youth.

We had planned to evolve DoppelLab into many tangent endeavors, including two-way interaction that allows virtual visitors to also manifest into our real space in different ways via distributed displays and actuators. Most of our current efforts in this line of research, however, have moved outdoors, to a 600-acre retired cranberry bog called 'Tidmarsh' located an hour south of MIT in Plymouth Massachusetts [103]. This property is being restored to its natural setting as a wetland, and to document this process, we have installed hundreds of wireless sensors measuring a variety of parameters, such as temperature, humidity, soil moisture and conductivity, nearby motion, atmospheric and water quality, wind, sound, light quality, etc. Using our low-power Zigbee-derived wireless protocol and posting data in CHAIN, these sensors, designed by my student Brian Mayton, last for 2 years on AA batteries (posting data every 20 seconds) or live perpetually when driven by a small solar cell. We also have 30 real-time audio feeds coming from an array of microphones distributed across Tidmarsh, with which we are exploring new applications in spatialized audio, real-time wildlife recognition with deep learning [104], etc.

In addition to using this rich, dense data archive to analyze the progress of restoration in collaboration with environmental scientists, we are also using the real-time CHAIN-posted data stream in a game-engine visualization. Within this 'Doppelmarsh' framework, you can float across the re-synthesized virtual landscape (automatically updated by information gleaned by cameras so the virtual appearance tracks real-world seasonal and climatic changes), and see the sensor information realized as situated graphics, animations, and music. The sounds in Doppelmarsh come from both the embedded microphones (spatialized relative to the user's virtual position) and music driven by sensor data (e.g., you can hear the temperature, humidity, activity, etc. coming from nearby devices). We have built a framework for composers to author musical compositions atop the Tidmarsh data [105] and have so far hosted four musical mappings, with more to come shortly.

Through our virtual Tidmarsh, we have created a new way to remotely experience the real landscape while keeping something of the real aesthetic. In Doppelmarsh, you can float across boggy land that's now very difficult to traverse (free of the guaranteed insect bites), hearing through nearby microphones and seeing metaphorical animations from real-time sensors (including flora and fauna that 'feed' upon particular sensor data, hence their appearance and behavior reflects the history of temperature, moisture, etc. in that area). We are also developing a wearable 'sensory prosthetic' to augment the experience of visitors to the Tidmarsh property. Based on a head-mounted wireless position/orientation-tracking bone-conduction headphone system we've developed called HearThere[106], we will estimate where or on what users are focusing attention via an array of wearable sensors, and manifest sound from related/proximate microphones and sensors appropriately. Users can look across a stream, for example, and hear audio from microphones and sensors there, then look down into the water and hear sound from hydrophones and sensors there – if they focus on a log equipped with sensitive accelerometers, they can hear the insects boring inside. In the future, all user interfaces will all leverage attention in this fashion – relevant information will manifest itself to us in the most appropriate way and leverage what we're paying attention to in order to enhance perception and not distract.

I see artists as playing a crucial role in shaping how the augmented world appears to us. As large data sources emanating from everyday life are now accessible everywhere, and tools to extract structure from them increasingly available, artists, composers, and designers will sculpt the environments and realizations that present this information to us in ways that are humanly relevant. Indeed, a common mantra in my research group now states that 'Big Data is the Canvas for Future Artists' and we have illustrated this via the creative musical frameworks we've made available for both DoppelMarsh [105] and realtime physics data coming from the ATLAS detector at CERN's Large Hadron Collider [107]. I can see a near future where a composer can do a piece for a smart city, where the traffic, weather, air quality, pedestrian flow, etc. are rendered in music that changes with conditions and virtual user location. These are creative constructs that never end, always change, and can be perceived to relate to real things happening in real places, giving them imminence and importance.

Conclusions

Moore's Law has democratized sensor technology enormously over the past decade - this paper has touched on a few of the very many modes that exist for sensing humans and their activities [39]. Ever more sensors are now integrated into common products (witness mobile phones, which have become the Swiss Army Knives of the sensor/RF world), and the DIY movement has also enabled custom sensor modules to be easily purchased or fabricated through many online and crowd-sourced outlets [40]. As a result, this decade has seen a huge rush of diverse data flowing into the network. This will surely continue in the following years, leaving us the grand challenge of synthesizing this information into many forms – e.g., grand cloud-based context engines, virtual sensors, and augmentation of human perception. These advances not only promise to usher in true UbiComp – they also hint at radical redefinition of how we experience reality that will make today's common attention-splitting between mobile phones and the real world look quaint and archaic.

We are entering a world where ubiquitous sensor information from our proximity will propagate up into various levels of what is now termed the 'cloud' then project back down into our physical and logical vicinity as context to guide processes and applications manifesting around us. We won't be pulling our phones out of our pocket and diverting attention into a touch UI, but rather will encounter information brokered between wearables and proximate ambient displays. The computer will become more of a partner than experienced as a discrete application that we run. My colleague and founding Media Lab director Nicholas Negroponte described this relationship as a 'digital butler' in his writings from decades back [108] – I see this now as more of an extension of ourselves rather than the embodiment of an 'other'. Our relationship with computation will be much more intimate as we enter the age of wearables. Right now, all information is available on many devices around me with at the touch of a finger or the enunciation of a phrase. Soon it will stream directly into our eyes and ears once we enter the age of wearables (already forseen by then Media Lab students like Steve Mann [109] and Thad Starner [110] who were living in an early version of this world during the mid-90s). This information will be driven by context and attention, not direct query, and much of it will be pre-cognitive, happening before we formulate direct questions. Indeed, the boundaries of the individual will be very blurry in this future. Humanity has pushed these edges since the dawn of society. Since sharing information with each other in oral history, the boundary of our mind expanded with writing and later the printing press, eliminating the need to mentally retain verbatim information and keep instead pointers into larger archives. In the future, where we will live and learn in a world deeply networked by wearables and eventually implantables, how our essence and individuality is brokered between organic neurons and whatever the information ecosystem becomes is a fascinating frontier that promises to redefine humanity.

Acknowledgements

The MIT Media Lab work described in this article has been conducted by my graduate students, postdocs, and visiting affiliates over the past decade, many of whom were explicitly named in this article and others represented in the cited papers and theses – readers are encouraged to consult the included references for more information on these projects along with a more thorough discussion of related work. I thank the sponsors of the MIT Media Lab for their support of these efforts.

References

[1] Michael Lapinski, Mark Feldmeier, and Joseph A. Paradiso, "Wearable Wireless Sensing for Sports and Ubiquitous Interactivity," IEEE Sensors 2011, Dublin Ireland, Oct. 28-31. 2011.

[2] For an early example, see the data sheet for the Analog Devices AD22365.

[3] M. Feldmeier and J.A. Paradiso, "Personalized HVAC Control System," in Proc. of IoT 2010, Tokyo Japan, Nov. 29 - Dec. 1, 2010.

[4] D. De Rossi, A. Lymberis, (eds.) IEEE Transactions on Information Technology in Biomedicine, Special Section on New Generation of Smart Wearable Health Systems and Applications, September 2005, volume 9, number 3.

[5] Buechley, L. "Material Computing: Integrating Technology into the Material World." In Price, S., Jewitt, C., and Brown, B. (eds). The SAGE Handbook of Digital Technology Research, Sage Publications Ltd., to appear 2012.

[6] Vinge, V., *Rainbow's End*, Tor Books, 2006.

[7] Eric R. Fossum, "CMOS Image Sensors: Electronic Camera-On-A-Chip," IEEE Transactions On Electron Devices, Vol. 44, No. 10, October 1997, pp. 1689-1698.

[8] Lifton, J., Laibowitz, M., Harry, D., Gong, N.W., Mittal, M. and Paradiso, J.A., "Metaphor and Manifestation – Cross Reality with Ubiquitous Sensor/Actuator Networks," IEEE Pervasive Computing Magazine, July-September 2009 (vol. 8 no. 3), pp. 24-33.

[9] Laibowitz, M., Gong, N-W., and Paradiso, J.A., "Multimedia Content Creation using Societal-Scale Ubiquitous Camera Networks and Human-Centric Wearable Sensing," in Proc. of ACM Multimedia 2010, Florence Italy, 25-29 October 2010, pp. 571-580.

[10] Paradiso, J.A., Gips, J., Laibowitz, M., Sadi, S., Merrill, D., Aylward, R., Maes, P., and Pentland, A., "Identifying and Facilitating Social Interaction with a Wearable Wireless Sensor Network," Springer Journal of Personal & Ubiquitous Computing, Volume 14, Number 2, February 2010, pp. 137-152.

[11] Reben, A. and Paradiso, J.A., "A Mobile Interactive Robot for Gathering Structured Social Video," in Proc. of ACM Multimedia 2011, October 2011, pp 917-920.

[12] Gong, N-W., Laibowitz, M., and Paradiso, J.A., "Dynamic Privacy Management in Pervasive Sensor Networks," in Proc. of Ambient Intelligence (AmI) 2010, Malaga, Spain, 25-29 October 2010, pp. 96-106.

[13] Dublon, G., "Beyond the Lens: Communicating Context through Sensing, Video, and Visualization." MS Thesis, MIT Media Lab, September 2011.

[14] Sinha, A. and Chandrakasan, A., "Dynamic Power Management in Wireless

Sensor Networks," IEEE Design & Test of Computers, March-April 2001, pp. 62-74.

[15] Paradiso, J.A. and Starner, T., "Energy Scavenging for Mobile and Wireless Electronics," IEEE Pervasive Computing, Vol. 4, No. 1, February 2005, pp. 18-27.

[16] P. D. Mitcheson, T. C. Green, E. M. Yeatman, A. S. Holmes, "Architectures for Vibration-Driven Micropower Generators," Journal Of Microelectromechanical Systems, Vol. 13, No. 3, June 2004, pp. 429-440.

[17] Yogesh K. Ramadass, Anantha P. Chandrakasan, "A Batteryless Thermoelectric Energy-Harvesting Interface Circuit with 35mV Startup Voltage," in Proc. of ISSCC 2010, pp. 486-488.

[18] Hiroshi Nishimoto, Yoshihiro Kawahara, and Tohru Asami, "Prototype Implementation of Wireless Sensor Network Using TV Broadcast RF Energy Harvesting," UbiComp'10, September 26–29, 2010, Copenhagen, Denmark, pp. 373-374.

[19] J. R. Smith, A. Sample, P. Powledge, A.Mamishev, and S. Roy. "A wirelessly powered platform for sensing and computation." In UbiComp 2006: Eighth International Conference on Ubiquitous Computing, pages 495–506, Orange County, Calif., 17–21 September 2006.

[20] Rachel Bainbridge and, Joseph A. Paradiso, "Wireless Hand Gesture Capture Through Wearable Passive Tag Sensing," in Proceedings of BSN2011Dallas, TX, May 23-25, 2011, pp. 200 - 204.

[21] Lee, Y., "Ultra-Low Power Circuit Design for Cubic-Millimeter Wireless Sensor Platform," PhD Thesis, University of Michigan (EE Dept), 2012.

[22] Michael Graham Richard, "Japan: Producing Electricity from Train Station Ticket Gates," http://www.treehugger.com/files/2006/08/japan_ticket_gates.php, August 4, 2006.

[23] M. Mozer, "The neural network house: An environment that adapts to its inhabitants," in Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligent Environments. Menlo Park, CA, 1998, pp. 110–114.

[24] Joseph Paradiso, Prabal Dutta, Hans Gellerson, Eve Schooler, (eds), Special Issue on Smart Energy Systems, IEEE Pervasive Computing Magazine, January-March 2011 (vol. 10 no. 1).

[25] Chabanis, G., "Self-Powered Wireless Sensors in Buildings: an Ideal Solution for Improving Energy Efficiency while Ensuring Comfort to Occupant," in Proc. Of Energy Harvesting & Storage Europe 2011, IDTechEx.

[26] Aldrich, M., Zhao, N., and Paradiso, J.A., “Energy Efficient Control of Polychromatic Solid-State Lighting Using a Sensor Network,” in Proc. SPIE (OP10) 7784 , San Diego CA, August 1-5, 2010.

[27] E. Elnahrawy, X. Li, and R. Martin, “The limits of localization using signal strength: A comparative study,” In Proceedings of the 1st IEEE International Conference on Sensor and Ad Hoc Communications and Networks, Santa Clara, CA, Oct. 2004.

[28] Miklós Maróti, et al, “Radio interferometric geolocation,” in Proc. of SENSYS 2005, pp. 1-12.

[29] Fabio Belloni, Ville Ranki, Ani Kainulainen, and Andreas Richter, “Angle-based Indoor Positioning System for Open Indoor Environments,” in Proc. of WPNC 2009.

[30] Saul Greenberg, Nicolai Marquardt, Till Ballendat, Rob Diaz-Marino, and Miaosen Wang. 2011. “Proxemic interactions: the new UbiComp?”, ACM Interactions 18, 1 (January 2011), 42-50.

[31] Jon Froehlich, et al, “Disaggregated End-Use Energy Sensing for the Smart Grid,” IEEE Pervasive Computing, Jan.-March 2010, pp. 28-39.

[32] Gabe Cohn, Dan Morris, Shwetak Patel, Desney S Tan, “Your Noise is My Command: Sensing Gestures Using the Body as an Antenna,” in Proc. CHI 2011, pp. 791-800.

[33] Sidhant Gupta, Ke-Yu Chen, Matthew S. Reynolds, Shwetak N. Patel, “LightWave: Using Compact Fluorescent Lights as Sensors,” in Proc. of UbiComp 2011, pp. 65-70.

[34] Chung, J., et al, “Indoor Location Sensing Using Geo-Magnetism,” in Proc. of MobiSys’11, June 28–July 1, 2011, Bethesda, Maryland, pp. 141-154.

[35] Gong, N-W., “Nan-Wei Gong, Steve Hodges, Joseph A. Paradiso, “Leveraging Conductive Inkjet Technology to Build a Scalable and Versatile Surface for Ubiquitous Sensing,” in Proc. of UbiComp 2011, pp. 45-54.

[36] Mittal, M., and Paradiso, J.A., “Ubicorder: A Mobile Device for Situated Interactions with Sensor Networks,” in IEEE Sensors Journal 11(3), Special issue on Cognitive Sensor Networks, pp. 818-828, 2011.

[37] Nicolai Marquardt, Tom Gross, Sheelagh Carpendale, Saul Greenberg, “Revealing the invisible: visualizing the location and event flow of distributed physical devices,” in Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction (TEI 2010), pp. 41-48.

[38] Gershon Dublon, Laurel S. Pardue, Brian Mayton, Noah Swartz, Nicholas Joliat, Patrick Hurst and Joseph A. Paradiso, "DoppelLab: Tools for Exploring and Harnessing Multimodal Sensor Network Data," in Proc. of IEEE Sensors 2011, Limerick Ireland, October 2011.

[39] Teixeira, T., Dublon, G., and Savvides, A., "A Survey of Human-Sensing: Methods for Detecting Presence, Count, Location, Track, and Identity, ACM Computing Surveys, Vol. 5, Issue 1, pp. 59–69, 2010.

[40] Paradiso, J.A., Heidemann, J., Zimmerman, T.G., "Hacking is Pervasive," IEEE Pervasive Computing, Vol. 7, No. 3, pp. 13-15, July-Sept. 2008.

[41] Bush, V., "As We May Think," Atlantic Monthly, No. 176 July 1945, pp. 101-108.

[42] Weiser, M., "The Computer for the 21st Century," Vol. 265, No. 3, September 1991, pp. 66-75.

[43] Paradiso, J., "Modular Synthesizer," in Stocker, G., and Schopf, C. (eds), *Timeshift* (Proc. of Ars Electronica 2004), Hatje Cantz Verlag, Ostfildern-Ruit, Germany, 2004, pp. 364-370.

[44] Paradiso, J.A., "Electronic Music Interfaces: New Ways to Play," *IEEE Spectrum,* Vol. 34, No. 12, December 1997, pp. 18-30.

[45] Paradiso, J.A., Sensor Architectures for Interactive Environments, in Deleaney, K. (ed), **Augmented Materials & Smart Objects: Building Ambient Intelligence Through Microsystems Technologies**, Chapter 16, Springer, 2008, pp. 345-362.

[46] MacKenzie, D., **Inventing accuracy: a historical sociology of nuclear missile guidance**, MIT Press, Cambridge MA, 1990.

[47] Walter, P.L., "The History of the Accelerometer," Revised Edition, Sound and Vibration, January 2007, pp. 84-92.

[48] Weinberg, H., "Using the ADXL202 in Pedometer and Personal Navigation Applications," Analog Devices Application Note AN-602, Analog Devices Inc. 2002.

[49] Renaudin, V., Yalak, O., and Tomé, P., "Hybridization of MEMS and Assisted GPS for Pedestrian Navigation", Inside GNSS, January/February 2007, pp. 34-42.

[50] Malinowski, M., Moskwa, M., Feldmeier, M., Laibowitz, M., Paradiso, J.A., "CargoNet: A Low-Cost MicroPower Sensor Node Exploiting Quasi-Passive Wakeup for Adaptive Asychronous Monitoring of Exceptional Events," in *Proceedings of the 5th ACM Conference on Embedded Networked Sensor Systems (SenSys'07)*, November 6–9, 2007, Sydney, Australia, pp. 145-159.

[51] Benbasat, A.Y. and Paradiso. J.A., “A Framework for the Automated Generation of Power-Efficient Classifiers for Embedded Sensor Nodes,” in *Proceedings of the 5th ACM Conference on Embedded Networked Sensor Systems (SenSys’07),* November 6–9, 2007, Sydney, Australia, pp. 219 - 232.

[52] Bernstein, J., “An Overview of MEMS Inertial Sensing Technology,” Sensors Magazine, Feb. 1, 2003.

[53] Ahmad, N., Ghazilla, A.R., and Khain, N.M., Kasi, V., “Reviews on Various Inertial Measurement (IMU) Sensor Applications,” International Journal of Signal Processing Systems, Vol. 1, No. 2, December 2013, pp. 256-262.

[54] Verplaetse, C., “Inertial proprioceptive devices: self-motion-sensing toys and tools,” IBM Systems Journal, Vol. 35, Issue 3-4, 1996, pp. 639-650.

[55] Paradiso, J.A., “Some Novel Applications for Wireless Inertial Sensors,” in *Proc. of NSTI Nanotech 2006*, Vol. 3, Boston, MA, May 7-11, 2006, pp. 431-434.

[56] Aylward, R. and Paradiso, J.A., “A Compact, High-Speed, Wearable Sensor Network for Biomotion Capture and Interactive Media,” in the *Proc. of the Sixth International IEEE/ACM Conference on Information Processing in Sensor Networks (IPSN 07)*, Cambridge, MA, April 25-27, 2007, pp. 380 - 389.

[57] Lapinski, M.T., “A Platform for High-Speed Biomechanical Data Analysis using Wearable Wireless Sensors,” PhD Thesis, MIT Media Lab, August 2013.

[58] Lightman, K., “Silicon Gets Sporty: Next-Gen Sensors Make Golf Clubs, Tennis Rackets, and Baseball Bats Smarter Than Ever,” IEEE Spectrum, March 2016, p. 44-49.

[59] I. Poupyrev, N.-W. Gong, S, Fukuhara, M. E. Karagozler, C. Schwesig, K. Robinson, “Project Jacquard: Interactive Digital Textiles at Scale” In Proc of CHI 2016, PP. 4216-4227.

[60] Hou, C., Jia, X., Wei, L., Tan, S-C, Zhao, X, Joanopoulos, and Fink, Y., “Crystalline silicon core fibres from aluminium core performs,” Nature Communications, Vol. 6, February 2015.

[61] Chandler, D.L., “New institute will accelerate innovations in fibers and fabrics,” MIT News Office, April 1, 2016.

[62] Vega, K., **Beauty Technology: Designing Seamless Interfaces for Wearable Computing**, Springer, 2016.

[63] Liu, X., Vega, K., Maes, P., and Paradiso, J.A., “Wearability Factors for Skin Interfaces,” in Proceedings of the 7th Augmented Human International Conference 2016 (AH '16). ACM, New York, NY.

[64] Vega, K., Jing, N., Yetisen, A., Kao, V., Barry, N., and Paradiso, J.A., “The Dermal Abyss: Interfacing with the Skin by Tattooing Biosensors,” submitted to TEI 2017.

[65] Mayton, Brian D., Zhao, N., Aldrich, M., Gillian, N., Paradiso, Joseph A., “WristQue: A personal sensor wristband,” in Proc. of the IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN ‘13), 2013, pp.1-6.

[66] Bainbridge, R. and Paradiso, J.A., “Wireless Hand Gesture Capture Through Wearable Passive Tag Sensing,” in *International Conference on Body Sensor Networks (BSN ‘11)*, Dallas, TX, May 23-25, 2011, pp. 200-204.

[67] Way, D. and Paradiso, J., “A Usability User Study Concerning Free-Hand Microgesture and Wrist-Worn Sensors,” In *Proceedings of the 2014 11th International Conference on Wearable and Implantable Body Sensor Networks (BSN '14)*. IEEE Computer Society, June 2104, pp. 138-142.

[68] Dementyev, A. and Paradiso, J., “WristFlex: Low-Power Gesture Input with Wrist-Worn Pressure Sensors,” in the *Proceedings of UIST 2014, the 27th annual ACM symposium on User interface software and technology*, Honolulu Hawaii, October 2014, pp. 161-166.

[69] Kao, H-L., Dementyev, A., Paradiso, J.A., and Schmandt, C., “NailO: Fingernails as Input Surfaces,” to appear in the Proceedings of *the International Conference on Human Factors in Computing Systems (CHI'15)*, Seoul Korea, April 2015.

[70] Laibowitz, M. and Paradiso, J.A., “Parasitic Mobility for Pervasive Sensor Networks,” in H. W. Gellersen, R. Want and A. Schmidt (eds): *Pervasive Computing*, Proceedings of the Third International Conference, Pervasive 2005, Munich, Germany, May 2005, Springer-Verlag, Berlin, pp. 255-278.

[71] Dementyev, A., Kao C-L., Choi, I., Ajilo, D., Xu, M., Paradiso, J., Schmandt, C., and Follmer, S. , “Rovables: Miniature On-Body Robots as Mobile Wearables,” in Proc. of ACM UIST 2016.

[72] Orwell, G., 1984, London: Secker and Warburg, June 8, 1949.

[73] Paradiso, J.A., "New Technologies for Monitoring the Precision Alignment of Large Detector Systems," *Nuclear Instruments and Methods in Physics Research* A386, 1997, pp. 409-420.

[74] Gokturk, S.B., Yalcin, H., and Bamji, C., "A Time-Of-Flight Depth Sensor – System Description, Issues and Solutions," CVPRW '04 Proceedings of the 2004 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'04) Volume 3, p. 35-45.

[75] Velten, A., Wilwacher, T., Gupta, O., Veeraraghavan, A., Bawendi, M.G., and Raskar, R., "Recovering Three-Dimensional Shape around a Corner using Ultra-Fast Time-of-Flight Imaging," Nature Communications, Vol. 3, 2012.

[76] Gillian, N., Pfenninger, S., Russell, S., and Paradiso. J.A., "Gestures Everywhere: A Multimodal Sensor Fusion and Analysis Framework for Pervasive Displays," In *Proceedings of The International Symposium on Pervasive Displays (PerDis '14),* Sven Gehring (Ed.). ACM, New York, NY, June 2014, pp. 98-103.

[77] Bletsas, M., "The MIT Media Lab's Glass Infrastructure: An Interactive Information System," in IEEE Pervasive Computing, Feb. 2012, pp. 46-49.

[78] Mayton, B., Dublon, G., Palacios, S., and Paradiso, J.A., "TRUSS: Tracking Risk with Ubiquitous Smart Sensing," in *IEEE SENSORS 2012*, Taipei, October 2012.

[79] Boyle, E., Kiziroglou, M., Mitcheson, P., and Yeatman, E., "Energy Provision and Storage for Pervasive Computing", in IEEE Pervasive Computing, October-December 2016.

[80] Park. R.L., **Voodoo Science: The Road from Foolishness to Fraud**, Oxford University Press, 2001.

[81] Overbyte, D., "Reaching for the Stars, Across 4.37 Light-Years," New York Times, April 12, 2016.

[82] Aldrich, M., "Experiential Lighting - Development and Validation of Perception-based Lighting Controls," Ph.D. Thesis, MIT Media Lab, August 2014.

[83] Zhao, N., Aldrich, M., Reinhart, C.F., and Paradiso, J.A., "A multidimensional continuous contextual lighting control system using Google Glass," to be presented at the *2nd ACM International Conference on Embedded Systems For Energy-Efficient Built Environments (BuildSys 2015),* Nov. 2015.

[84] Axaria, A., "INTELLIGENT AMBIANCE: Digitally mediated workspace atmosphere, augmenting experiences and sup- porting wellbeing," MS Thesis, MIT Media Lab, August 2016.

[85] P. Levis, et al, TinyOS: An operating system for sensor networks," in **Ambient Intelligence**, Springer Berlin, 2005, pp. 115-148.

[86] Gakstatter, E., Finally, A List of Public RTK Base Stations in the U.S.," GPS World, Jan. 7, 2014.

[87] See: https://www.microsoft.com/en-us/research/event/microsoft-indoor-localization-competition-ipsn-2016/

[88] Bolt, R., ""Put-that-there": Voice and gesture at the graphics interface," in ACM SIGGRAPH Computer Graphics, Vol. 14(3), July 1980, pp. 262-270.

[89] Zhao, N., Dublon, G., Gillian, N., Dementyev, A., Paradiso, J., "EMI Spy: Harnessing Electromagnetic Interference for Low-Cost, Rapid Prototyping of Proxemic Interaction," presented at the 2015 International Conference on Wearable and Implantable Body Sensor Networks (BSN '15), June 2015.

[90] Paradiso, J.A., Lifton. J., and Broxton, M., "Sensate Media - Multimodal Electronic Skins as Dense Sensor Networks," *BT Technology Journal*, 22(4), October 2004, pp. 32-44.

[91] Dementyev, A., Kao, H-L, and Paradiso, J.A., "SensorTape: Modular and Programable 3D-Aware Dense Sensor Network on a Tape," in Proc. of *ACM UIST* 2015.

[92] Gong, N-W., Wang, C-Y., and Paradiso, J.A., "Low-cost Sensor Tape for Environmental Sensing Based on Roll-to-roll Manufacturing Process," in *IEEE SENSORS 2012*, Taipei, October 2012.

[93] Qi, J., Hong A., and Paradiso J., "Crafting technology with circuit stickers," in Proc. of ACM IDC '15 (14th International Conference on Interaction Design and Children), 2015, pp. 438-441.

[94] https://www.xively.com/xively-iot-platform/connected-product-management

[95] https://allseenalliance.org/framework

[96] https://www.iotivity.org/

[97] Russell, S. and Paradiso, J.A., "Hypermedia APIs for sensor data: A pragmatic approach to the Web of Things," in the *Proceedings of ACM Mobiquitous 2014*, London UK, December 2014, pp. 30-39.

[98] https://github.com/ResEnv/chain-api

[99] Ramsay, D., "LearnAir: toward Intelligent, Personal Air Quality Monitoring," MS Thesis, MIT Media Lab, August 2016.

[100] McLuhan, M., **Understanding Media : The Extensions of Man**, McGraw-Hill, 1964.

[101] Dublon, G., Paradiso, J. A., "How a Sensor-Filled World Will Change Human Consciousness," *Scientific American*, July 2014, pp. 36-41.

[102] Joliat, N., Mayton, B., Paradiso, J.A., "Spatialized Anonymous Audio for Browsing Sensor Networks Via Virtual Worlds," in *The19th International Conference on Auditory Display (ICAD)*, Lodz, Poland, July 2013, pp. 67-75.

[103] http://tidmarsh.media.mit.edu

[104] Duhart, C., Chapter 8, "Ambient Sound Recognition" in *Toward Organic Ambient Intelligences – EMMA*, PhD Thesis, l'Université du Havre, France, June 2016.

[105] Lynch, E. and Paradiso, J.A., "SensorChimes: Musical Mapping for Sensor Networks," in Proc. of the NIME 2016 Conference, Brisbane Australia, July 11-15, 2016.

[106] Russell, S., Dublon, G., and Paradiso, J.A., 'HearThere: Networked Sensory Prosthetics Through Auditory Augmented Reality,' in Proc. of the ACM Augmented Human Conference, Geneva Switzerland, February 2016.

[107] Cherston, J., Hill, E., Goldfarb, S. and Paradiso, J., "Musician and Mega-Machine: Compositions Driven by Real-Time Particle Collision Data from the ATLAS Detector," in Proc. of the NIME 2016 Conference, Brisbane Australia, July 11-15, 2016.

[108] Negroponte, N., **Being Digital**, Knopf, NY, 1995.

[109] Mann, S., "Wearable Computing: A First Step Toward Personal Imaging," IEEE Computer, February 1997, pp. 25-32.

[110] Starner, T., "The Challenges of Wearable Computing, Parts 1 & 2", IEEE Micro 21(4), July 2001, pp. 44-67.

[111] Warneke, B., Last, M., Liebowitz, B., and Pister, K.S.J., "Smart Dust: Communicating With a Cubic-Millimeter Computer," IEEE Computer, 34(1), Jan. 2001, pp. 44-51.

[112] Lifton, J. and Paradiso, J.A., "Dual Reality: Merging the Real and Virtual," in *Proceedings of the First International ICST Conference on Facets of Virtual Environments (FaVE)*, Berlin, Germany, 27-29 July 2009, Springer LNICST 33, pp. 12-28.