# CAREER: Scalable Sensor Infrastructure for Sustainably Managing the Built Environment

**Overview:**

Federal sustainability targets mandate that 50% of U.S. commercial buildings become net-zero energy by 2050. A range of options exists to achieve this goal, but economic concerns require a measured and controlled approach--both figuratively and literally. Today, critical gaps exist in the energy and water measurement technology, and in the indoor climate control science, needed to benchmark options, operate buildings efficiently, and maintain occupant comfort.

This CAREER project proposes a long-term, integrated program of research, education, and outreach to create foundational measurement science and technology, increase awareness of resource demands and usage patterns, and promote more sustainable stewardship of resources through better management of the built environment. The transformative aspects of the research agenda include (i) fundamental understanding of pervasive, energy-harvesting sensors that infer resource usage indirectly through side-channel emissions, (ii) new protocols to network and manage the sensors, and the applications that use them, and (iii) new toolkits that combine the novel measurement technology with algorithms, systems, and methods to better control buildings. Collectively, these contributions enable a new measurement paradigm that offers unparalleled sampling granularity and scale, challenge empirically unfounded assumptions of resource use, and crucially, provide a measured and controlled pathway to a sustainable energy future.

The proposed ideas are evaluated by building, deploying, and operating a prototype system in a “living laboratory” and comparing it with prior methods. The core infrastructure supports several applications that share the underlying technology--including continuous energy and water audits segregated by end use types, demand-controlled daylight harvesting as an overlay on the existing lighting system, intelligent thermostats that use distributed temperature, humidity, and occupancy sensors to maintain comfort, and demand-controlled ventilation driven by occupancy and air quality levels--which in turn inform fundamental design choices.

**Intellectual Merit:** This proposal explores a new kind of distributed, energy-harvesting metering architecture, and its applications to sustainability. The key insight is that the transfer and use of energy usually emits energy, often in a different domain, and that this emitted energy is often enough to intermittently power simple, energy-harvesting sensors whose duty cycle is proportional to the energy being transferred or used. Hence, the mere activation rate of the sensors signals the underlying energy use. The power-proportional relationship between usage activity and side channel harvesting, when coupled with state-of-the art, millimeter-scale, nano-power chips and whole-house or panel-level meters, enables small and inexpensive sensor tags that are pervasively distributed with unbounded lifetimes. But, networking and tasking them, and making sense of their data, requires a fundamental rethinking of low-power communications, control, and data fusion to abstract the intermittent, unreliable, and noisy sensor infrastructure.

**Broader Impact:** A National Science and Technology Council (NSTC) report suggests that, “For the foreseeable future, the greatest national energy saving potential lies with improvements to existing buildings.” This project lowers the barrier to implementing these recommendations by delivering a scalable sensor infrastructure that forms the measurement-driven foundations needed for the intelligent, efficient, and responsive buildings at the heart of global sustainability efforts. This proposal’s broader impact stems from undergraduate and graduate courses that produce students who bridge disciplines, operating at the intersection of measurement science, information technology, and sustainability policy. The clear social impact of the research aids in recruiting female students. Collaborations with DOE’s [redacted] Lab ensures the research is timely and relevant. Industrial collaborations with [redacted] and [redacted] rapidly translate scientific discovery and technical knowledge into beneficial commercial products. Open source infrastructure enables third parties to readily use, improve, and extend the research artifacts. The research findings are disseminated through scholarly publications, industry presentations, and academic leadership to Federal agencies.

## D.1 Significance—A National Grand Challenge

Our unbridled consumption of fossil fuels is depleting the planet's resources [Hubbert56], changing its climate [IPCC07], and contributing to armed conflict [Woodward07]. To address these energy-related challenges, a National Science Board (NSB) study [NSB09] states that "Transformation of the U.S. fossil fuel-based energy economy to **a sustainable energy economy is a critical grand challenge** facing the Nation today." The report recommends that the U.S. Government should lead "a nationally coordinated research, development, demonstration, deployment, and education (RD3E) strategy to transform the U.S. energy system to a sustainable energy economy that is far less carbon intensive." Executive Order 13514—Federal Leadership in Environmental, Energy, and Economic Performance—has launched just such an RD3E program [Obama09]. The Order directs Federal agencies to "implement high performance sustainable Federal building design, construction, operation and management, maintenance, and deconstruction."


However, many agencies report poor performance on this goal. For example, the DOE scorecards for the past three years report RED on Green Buildings. This means that the DOE cannot demonstrate compliance with the guiding principles of the Executive Order [DOE11, DOE12, DOE13]. These reports are concerning because a National Science and Technology Council (NSTC) report [NSTC08] states that "for the foreseeable future, **the greatest national energy saving potential lies with improvements to existing buildings.**" If the DOE is having trouble, how will the rest of us fare?

### D.1.1 Why Focus on Buildings?

Buildings account for a significant share of U.S. resource use. They represent 39% of the energy, 73% of the electricity, 55% of the natural gas, and 12% of the water consumed in the United Stated, and powering them results in $400 Billion in annual expenditures. If current trends continue, by 2025 buildings worldwide will consume more energy than the transportation and industry sectors combined [NSTC11]. Buildings are also responsible for significant emissions and pollutants including 38% of the $CO_2$, 49% of the $SO_2$, 25% of the $NO_x$, and 60% of all nonindustrial waste. These facts are a major reason that the NSTC stressed the importance of *improvements to existing buildings* in its 2008 report.

The billion-dollar question is: *which improvements?* The answer varies from building to building but it is crucial for assessing the economic viability of competing projects and prioritizing their implementation. To help evaluate projects before *and* after deployment, one critical focus area is developing new measurement science and technology to enable deeper and broader monitoring of building resource usage [NSB09, NIST10]. A recent report recommends that the building community view energy and water submetering as a key tool to baseline performance, identify inefficiencies, and commission continuously, but the report also notes that "no widely available tools exist to help them understand their energy consumption or to compare it with peer groups" [NSTC11]. While several commercial products and research projects have attempted to address this gap, scalable, affordable, and easy-to-use systems for fine-grained submetering do not exist today.

### D.1.2 State-of-the-Art in Metering and End Use Disaggregation

Whole-building meters [TED13, Wattvision13] provide an overall view of energy or water use, but they do not *disaggregate* the data in a way that allows attribution to individual loads. Analytics can take meter data and disaggregate the readings using appliance signatures with a technique called non-intrusive load monitoring or NILM [Hart92]. This approach works when the loads are sufficiently few (e.g. 6-7), mostly large (e.g., air conditioners and stoves), and have distinctive signatures (e.g. refrigerators and ovens). NILM has difficulty with smaller loads (e.g. electronics) or multiple instances of a particular load (e.g. several 60 W light bulbs). Moreover, NILM approaches cannot pinpoint individual devices. Extensions of this general approach use higher frequency sampling (e.g. MHz) of the current and voltage waveforms, higher dimensional data (e.g. real and reactive power), and complex signatures (e.g. wideband spectra arising from the flick of a switch or a toggle of a faucet) to identity individual loads [CGF+10, FLC+09, GRP10, PRK+07], but these approaches still require training (the association of the loads to their signatures), are susceptible to small changes in the environment (e.g. moving a load from one outlet to another), and remain costly (due to the use of high-rate sampling and processing techniques).

At the other extreme, plug load meters allow individual loads to be measured [LFO+07, PKGZ08, JDDC09, WBD+11, iMeter13, P313, WattsUp13]. Standalone meters display usage data locally, which is useful for casual use but not for automated aggregation and analysis. Networked (typically wireless) meters send their data to servers for analysis and visualization. However, plug load metering faces some coverage and cost disadvantages. Some loads are built-in or hard-to-access, including ceiling lights, HVAC equipment, and some major appliances, making them poorly suited to such meters. Furthermore, at a price of \$25-\$50 for standalone meters and \$75-\$250 for networked meters, covering a home or office can cost thousands or tens of thousands of dollars.

Due to the drawbacks of both the NILM and direct metering approaches, some efforts have explored hybrid models in which additional sensors augment NILM and aid with disaggregation [KSC+08, KSCS09, MCS12, SSW13]. The extra sensors help NILM scale beyond a half-dozen loads by providing it an additional signal that reflects the state of an individual load. Sensors that detect the on-off states [JS10] or more finely quantized energy emissions of appliances including light, sound, and magnetism [RBR10], have been shown to aid greatly in disaggregation. However, the major scalability impediment that this approach faces is the cost of the sensors and the overhead of periodic battery replacement. Low-cost, mains-powered sensors that detect appliance state can address these problems, but they are currently restricted to electrical plug loads [WS12].

### D.1.3 Scalable Sensor Infrastructure

To address the cost, coverage, and scalability challenges of energy and water metering, this proposal explores a new kind of distributed, energy harvesting metering architecture. The key insight is that ***the transfer and use of energy usually emits energy, often in a different domain, and that this side-channel energy is often enough to intermittently power simple, energy-harvesting sensors whose activations rate is proportional to the energy transferred or used.*** Hence, the mere activation rate of the sensors signals the underlying energy use or flow, either directly on indirectly. We can use this principle to build inexpensive and easy-to-install sensors.

Figure 1: Monjolo (water mill) hammers at a rate that is proportional to the water flow rate.

**Metering Principle.** This model of metering is motivated by *monjolo*, Portuguese for water mill (Figure 1). The rate at which the monjolo pestle strikes is proportional to the water's flow rate. We call this the ***monjolo principle***, and note that in addition to the strike rate being proportional to the flow rate, each strike signals a fixed quanta of water.

In this example, monjolo directly measures water flow. But sometime such direct measurement is not possible, but a side-channel measurement is. For example, the magnetic field surrounding a conductor feeding a load is related to the current in that wire, the light intensity close to dimmable overhead light is related to the power dissipated by the bulbs, the air flow through ductwork is related to the power drawn by an air handling unit, and the heat intensity on a range top is related to the power draw of the stove.

**Distinction.** In contrast with prior work that attempts to directly measure these side-channel emissions (including ones that harvest the side-channel energy and use it to power active sensors [CLC+10, MCS12]), we propose to simply harvest the side-channel energy and transmit a radio packet when enough energy has been accumulated to do so. Harvesting just enough to send a packet requires less energy than revenue-grade sensing, and is easier to install, thus enabling smaller and less expensive, but also cruder sensors. But much of what the sensors lose in individual quality, they gain through sheer numbers, and in conjunction with data fusion algorithms. Like the NILM and hybrid disaggregation techniques, we also use a whole-building meter. In our model, data from the whole-building meter is combined with the packets from the multitude of inexpensive sensors that are delivered over a lossy radio channel, to infer the contributions of individual loads. Depending on the relationship between a sensor's activation rate and the underlying energy flow—whether linear or not—the fusion may use one of several different convex optimization algorithms—linear, non-linear, or geometric programming.

**Preliminary Results.** Our prior work has demonstrated the monjolo principle with mesoscale (cm/inch-size) sensors built from commercial components that respond to low levels of light [YCB+12] and magnetism [DCD13]. We hypothesize that the power-proportional relationship between usage activity and side channel harvesting, when coupled with our work on mm-scale, nano-power integrated circuits [LKB+12], whole-house metering [SCD10], and energy profiling [FDLS08], will enable small and inexpensive sensors that can disaggregate loads, be pervasively distributed, and run with unbounded lifetimes, thus solving the two scalability challenges—coverage and cost.

### D.1.4 Proposed and Enabled Research

Section 2 outlines our approach to realizing a new class of energy-harvesting energy and water meters based on the monjolo principle, and more conventional energy-harvesting climate and occupancy sensors. These sensors will support studies to better characterize indoor climate and energy use. They will provide a sensing tier to complement emerging third-party home and building operating systems [DMA+12, DKTC13] and middleware services [ABC+12]. They will support smart building applications like real-time occupant energy footprinting [CCZ+12] and demand-controlled ventilation [TKDC13].

## D.2 Research Vision and Plans

Our long-term research plan aims to make buildings more efficient and scalable using improved sensing technologies and control strategies. Over the past two years, we have been laying the groundwork for this long-term vision by creating a living laboratory in our building. These efforts have yielded many new platforms, protocols, and services to support our research, as Figure 2 shows. But this work has also demonstrated the poor scalability of many sensing techniques, most notable energy and water usage in buildings.

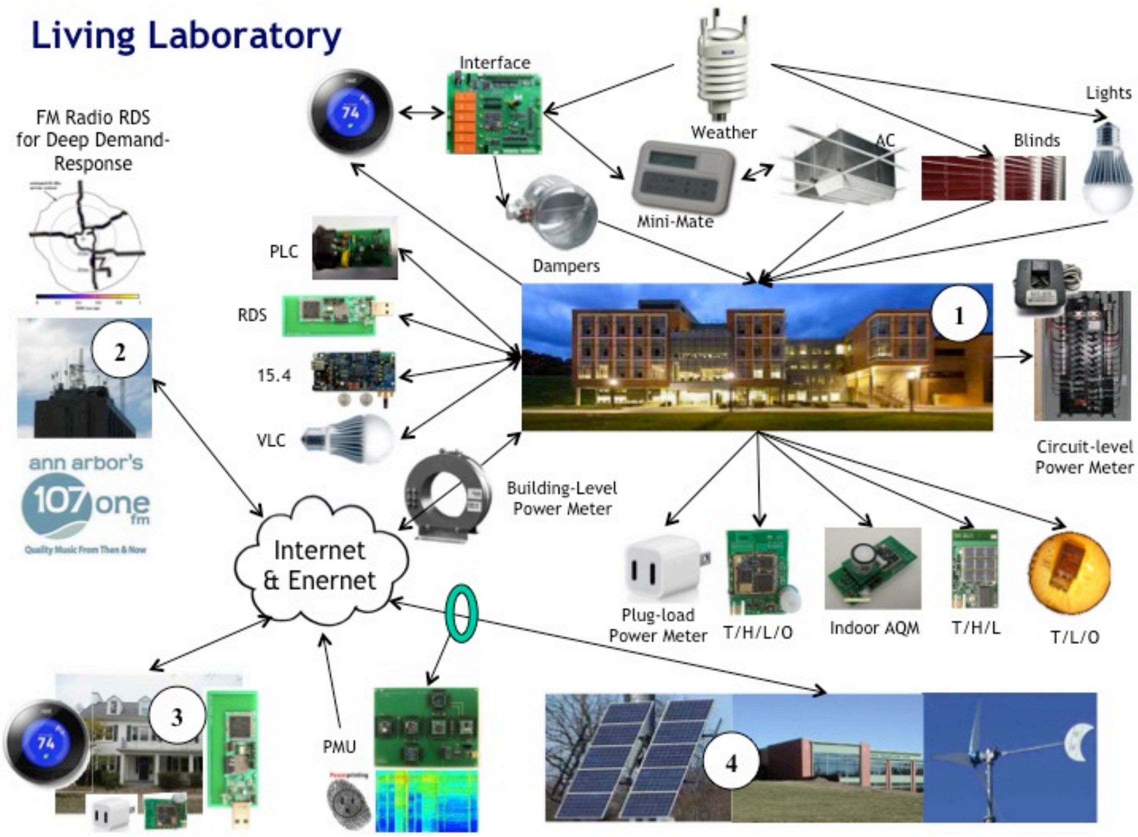


Figure 2: Living laboratories: (1) CSE Building instrumented with locally-designed platforms and software for (a) monitoring plug loads, panel meters, and building climate, including temperature, humidity, lighting, occupancy, $CO_2$ level, and air flow (in development), (b) communications over power line, RDS, 802.15.4, and visual light, and (c) programmatically controlling heating, lighting, and air conditioning; (2) Local radio station with RDS transmitter that allows us to transmit (digital) data on FM [redacted] to the [redacted] metro area, allowing wide-area broadcasts for demand response; (3) Instrumented home; (4) Instrumented local facility with renewables.

Some studies have shown significant potential for improved efficiency: reductions in lighting power by 51-72% [Kanellos09] and energy for heating, cooling, ventilation, and air conditioning by 75% [IEA08]. To achieve the "450 Scenario"—limiting $CO_2$ emissions so global temperature increases by no more that 2 degrees Celcius—requires end-uses to provide half the $CO_2$ reductions [Heffner11]. A key challenge in achieving greater end-use efficiency—that is, better managing buildings and their contents—lies with improving our understanding of the relationships between resources, services, occupants, and externalities. This requires us to measure not only energy and water but *all* critical parameters in a *scalable* way, and log them for real-time or deferred analyses.

### D.2.1 Model Side-Channel Energy and Characterize Harvester Performance

Our hypothesis rests on the assumption that there is sufficient side-channel energy from various loads to intermittently activate sensors [PS05]. A preliminary investigation of this critical assumption has led to promising results. In one set of experiments [YCB+12], we explored powering sensors from various light sources (e.g. fluorescent, incandescent, halogen, LED-2500K, LED-3000K, and natural), irradiance levels (e.g. 10-10,000 uW/cm$^2$), and solar cells (e.g. amorphous, crystalline, thin-film). Figure 3 shows the results—that message interval is closely related to irradiance. This figure also shows that sensors can transmit roughly every two minutes even under very low lighting levels.

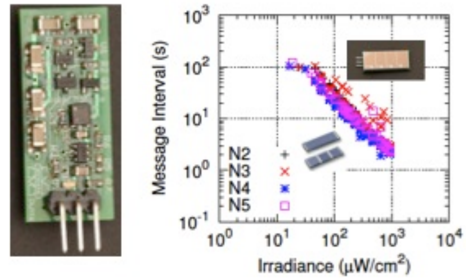

Figure 3: A small indoor photovoltaic harvester and the message interval under various lighting conditions. The system is roughly power monotonic over a 100x irradiance range.


Figure 4: Split core current transformer.

In other experiments [DCD13], we explored the viability of an energy-harvester that operates from the magnetic field emissions of an AC current-carrying conductor. The prototype, seen in Figure 5, shows a monotonic (but not linear) relationship between wakeup frequency and load power over a 25 W to 500 W range. We can adjust the higher and lower ends of the range. This basic design could be integrated into a split core current transformer (CT), aiding panel metering.

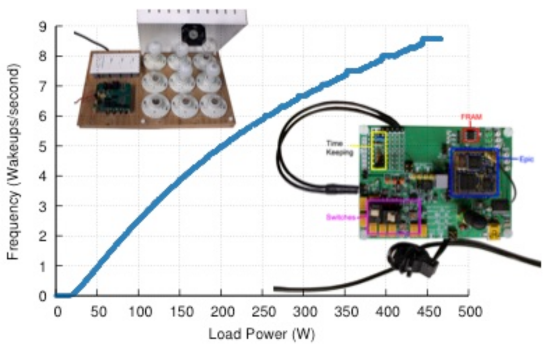

Figure 5: A current transformer based magnetic harvester shows a monotonic relation between activations and load power.

**Task Milestone:** *Record and model the side-channel energy emissions from various home and office loads and build harvesters to empirically validate available energy and determine the transfer function between activation rate and source energy* – We will build multiple versions of our mesoscale energy harvesters using custom and commercial components and will use them to characterize the harvestable energy in multiple domains from a range of everyday objects, including lamps, stoves, panel meters, pipes, air vents, etc. Principally, this will focus our attention on energy harvesting operation from indoor lighting, magnetic fields, acoustic sources, temperature gradients, air flows, and movement (e.g. door open/close). We propose to use off-the-shelf harvesters where possible (e.g. thermoelectric generation, or TEG) and explore the design of new harvesters (e.g. air flow using impellers), as Figure 6 shows. We propose to augment our work by creating empirical models that parameterize the design space. This will allow us to evolve our designs as the energy harvesters, power supply electronics, processor, and radios evolve, for example using piezo-electromagnetic (PEM) devices [PXC+12, XPS+13].

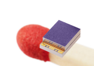

Figure 6: TEG and impeller.

### D.2.2 Integrated Sensors: Harvesters + Millimeter-Scale, Nano-Power Electronics

We have integrated photovoltaic and magnetic energy harvester designs with a processor, radio, and storage [DJT+08] to build two complete “Monjolo” devices. Our key findings were the importance of extremely low-leakage operation (less than 1uA, by power gating everything but the harvester), the need for a richer power supply control interface for system software (to allow triggered wakeup and software-controlled capacitor discharge and reset), and the value of an unconventional energy storage hierarchy (capacitor only). In addition, we found that protocols must not only be low-power, but they must also become low-energy (to minimize the size of the energy reservoir). Software, and especially bootup code and communications, must become low-latency, by eliminating most conservative timeouts and startup dependencies, and embracing high-concurrency (we improved the startup energy and time by 14x and 69x, respectively, over prior art). We also discovered the need for ultralow write-energy memory storage across power cycles (satisfied by ferro-electric or FRAM; necessary to keep counter state to detect and correct for duplicate/dropped packets over a lossy radio link). Applying these principles, we were able to show that the Monjolo principle is viable, at least using commercial, mesoscale devices.

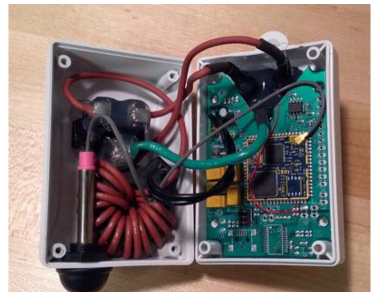

Figure 7: A "Monjolo" meter that works by harvesting the magnetic field surrounding an AC-carrying conductor.

**Task Milestone:** *Develop and evaluate microscale “peel-and-stick” sensors that can be affixed to everyday objects and can signal those objects’ internal state through their side-channel emissions* – We envision a range of small and inexpensive sensor tags that have a QR code printed on them so that can be read by a mobile phone and electronics embedded in them so that they can implement the monjolo principle. This task builds on our current work on millimeter scale computing (CSR-1111541), in which we have developed or obtained all of the essential building blocks—processor, radio, memory, power management, timers, battery, capacitor, solar cell, imager, interconnection bus, and optical programmer—to allow us to construct modular sensors.


Figure 8: Proposed "peel-and-stick" sensor tag. A 0.5" x 0.5" QR code sits next to our processor (ARM CM-0), radio, and solar cell.

One goal is to understand the scaling limits of this approach (area, volume, power) to support energy-harvesting operation. Our prototype devices draw approximate 10 nW to transmit a packet every 10 minutes with a range of a few meters. Given the output power of existing harvesters—expected to be in the µW range—we expect to have a tradeoff space of second- or subsecond-level temporal granularity or shorter or longer range.

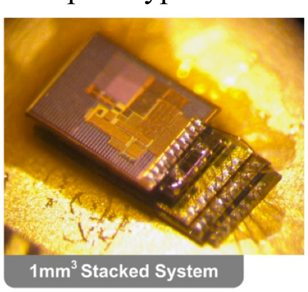


Figure 9: Die-stacked, cubic-mm scale, nano-power system (1 pkt/10 min/10 nW).

A second goal is to replicate our mesoscale prototypes and build completely new “Monjolo” sensors that are activated by the new harvester front ends we explore, including temperature gradients, mechanical vibrations, piezo-electromagnetic, or air flow. The third goal is to create a reusable family of sensor tags.

### D.2.3 Networking and Analytics

Our current approach to metering uses the architecture show in Figure 10. In the model, a Monjolo sensor transmits a packet per activation. Each packet identifies the source, packet number, and activation number (there may be multiple activations per packet, to minimize radio channel contention).

A data aggregator listens for packets from one or more sensors. Upon reception, it determines the interval between the packets and uses that value to estimate the power of the load attached to the sensor. The data aggregator then relays that information for further processing or display [ONS+12].

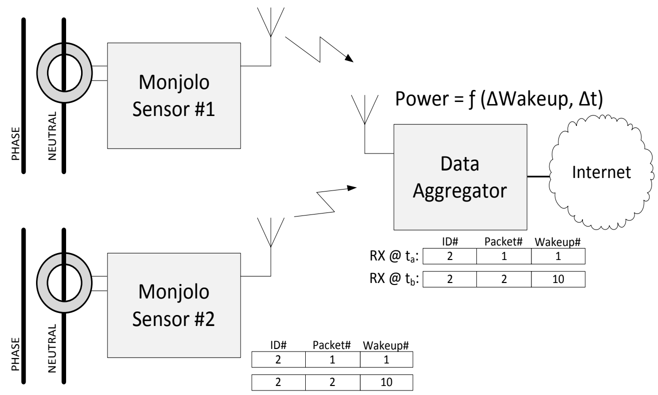


Figure 10: System and network architecture. Monjolo sensors transmit packets. A data aggregator tracks transmissions and converts packets to energy use for subsequent processing or display.

**Task Milestone:** *Develop the networking protocols and energy analytics necessary to combine the monjolo sensor transmissions with data from whole-building or panel-level meters to estimate the consumption of individual loads and evaluate the accuracy and precision of the approach at modest scales* – We propose an online data-processing architecture that accepts external pulse data from a conventional electricity, water, or gas meter (e.g. each pulse equals 1 W-hr, for a typical energy meter) and couples them with the packet transmission from the sensors to present a virtual sensor view of each load.

In our model, a user will affix a tag to a device, then use a mobile phone to read the tag's QR code and take a picture of the device, and finally affix the tag to the device of interest. Once attached and activated, the sensor will transmit periodically, its rate potentially limited to prevent high channel utilization, and the data aggregator will begin to keep track of the packets and the pules from external meters. The aggregator will run some kind of convex optimization—linear, non-linear, or geometric programming—to breakdown the aggregate into the individual components. This will allow us to attribute to each load its share of the energy consumption. As a simple example with linear sensors, we could simply use weighted multivariate least squares to estimate the usage of each energy load. The input to this regression process would be a set of records, for each interval during which the load "states" are same (as indicated by a roughly fixed packet transmission interval), the aggregate resource usage (e.g. energy) during that interval ($\Delta E$), the length of the interval ($\Delta t$), and the load states of all devices during the interval ($\alpha_1, \ldots, \alpha_n$) provide all of the data need to isolate the *i*-th load [FDLS08]. We propose to perform these calculation online to provide a real-time estimate of load consumption. To evaluate the accuracy of our approach, we propose to employ calibrated plug load meters [JDDC09] or commercial current-transformer based solutions [TED13].

Another area that requires research is the underlying network stack that permits the data to be reliably forwarded. Sensors might not have enough energy to retransmit packets in case of loss. Packets might be transmitted multiple times, coded across transmissions, or synchronized to an external signal (e.g. ambient 60 Hz) to ensure greater reliability and channel efficiency, but the specifics depends on activation rate.

## D.3 Educational Plan

Embedded sensing, computing, and control systems are becoming increasingly important to many engineering and scientific disciplines, and these systems are finding applications on bridges [Lynch07], in buildings [JLT+09], on bodies [SKL+10], and in biology [TPS+05]. Many new crowd-funded startups are offering hardware products around an "Internet of Things" (IoT) theme. A number of traditionally software companies—like Microsoft—are redefining themselves around a "devices and services" model. In short, there's strong and growing demand for embedded systems engineers who can operate across the embedded + mobile + cloud space. This groundswell of interest in embedded systems has created strong demand for courses, tutorials, and hackerspaces.

### D.3.1 Need for a Smart Objects-Focused Embedded Curriculum

A recent report [CCW11] from the NSF-sponsored workshop on Pervasive Computing at Scale (PeCS) concluded that, "We need an interdisciplinary approach to educate PhD students about developing smart object ecosystems. As many participants observed, this is a cross-disciplinary area, including computer systems and software, communications (both lower and the upper network layers), and circuit design. At present, a typical faculty member encourages their PhD students to take courses from different disciplines with the expectation that the students will integrate knowledge across classes."

Unfortunately, such bottom-up integration of knowledge—trying to integrate an unrelated set of courses into a cohesive whole—is difficult for most students without a global frame of reference [Witkin81]. Today's electrical engineering and computer science curricula rarely provide the holistic approach that is needed, yet the need for engineers capable of designing and deploying embedded systems both within academia and in industry grows. In conversation after conversation, I hear from academic colleagues and industrial collaborators, "We need students who understand systems, can build things, and can transcend hardware and software layers!"

### D.3.2 Undergraduate Course Development and Research Opportunities

To meet the growing need for embedded systems training, I have been offering a completely revamped embedded systems course for juniors and seniors (EECS [redacted]), and I have worked with a colleague in my Department to design a senior-level embedded systems capstone course (EECS [redacted]). Both of these courses take a holistic approach to embedded systems design, grounded in fun weekly labs for the first half of the course, and culminate with an integrative, open-ended design project in the second half.

Students who take these courses and show an aptitude for embedded systems and possible interest research are recruited. In three years of teaching EECS [redacted], I have worked with nine students on undergraduate research. Three have gone on to obtain masters degrees and three others have joined my research group as Ph.D. students. This approach has been so successful that many of the preliminary results discussed in this proposal were first explored as class projects and subsequently as undergraduate research.

This mirrors the NSF report [CCW11]: "We need to design inter-disciplinary courses, both at the graduate and undergraduate levels, which are broad in terms of scope, and which focus on integration across the different disciplines. Holistic courses, for example, could require each student to build a smart object, including hardware platforms, software libraries, and web service, thus clarifying the need for integrative understanding."

**From Course Project to Undergraduate Research:** Students who take EECS [redacted] and [redacted] are encouraged to continue independent research projects after their course classroom work, culminating in a first-authored research paper. Students are mentored through the research process—from conceptualizing research to giving a conference talk—and they are supported to travel to present their work. Sometimes, the course projects and subsequent research projects influence the graduate research program as well. I plan to continue this approach to introducing research into the classroom.

**Teaching Undergraduates to Do Research:** Most undergraduate students (and first year graduate students) are not familiar with the process of research or the tools and techniques used to effectively communicate research results. To bridge this gap in the recent past, I offered a "research bootcamp" lunch series for students who took EECS [redacted] and wanted to turn their class projects into research publications. Each week, the lunch meetings discussed a different topic, including articulating hypotheses, scripting experiments; communicating through figures, tables, and captions; typesetting with LaTeX; version control; planning and writing papers; writing titles and abstracts; reading papers; doing literature searches; and choosing a publication venue.

Students found this experience so useful that I started presenting the material during the discussion section of the graduate operating systems course (EECS [redacted]) that I teach. In the future, I plan to offer this course as an undergraduate research lunch seminar. This class will expose students to research tools, tactics, topics, and taste. Its course notes will attempt to capture many little lessons that take years to learn. It will exist to help directed study undergrad and grad students establish their presence, execute their research plans, and eventually disseminate their findings via top-tier publications. Every student will be expected to work diligently on a research project and complete the assignments in a timely manner, leading to a publication. The lecture notes for this seminar will be publicly available, and hopefully broadly used.

**Dissemination of Teaching Materials:** The materials developed for EECS [redacted], including lecture slides, labs, online videos, hardware designs, software designs, software tools, sample problems, and sample exams are freely available and readily downloadable from the course homepages. From time to time, instructors at other universities contact me, asking to adopt them for their own use, which of course I support. This had led to the course being adopted at other institutions, including the University of [redacted].

**Integration of Research and Education – A Virtuous Cycle:** While undergraduates benefit from the research experience that follows the EECS [redacted] and [redacted] courses, the process also works in reverse—and faculty can learn from and be greatly inspired by the work that the students do. Indeed, much of the preliminary work for this proposal emerged from three EECS [redacted] projects carried out by undergraduate students. The original motivation for this work was three projects from two different teachings of EECS [redacted]. What these successful projects demonstrate is a virtuous cycle of motivating education with research, and research with education, while training undergraduates about graduate school and research careers.

### D.3.3 Graduate Education and Training

I teach the graduate operating systems course (EECS [redacted]). This discussion-heavy, project-oriented course introduces students to advanced topics in computers systems. To prepare, students read one or two papers per class, write reviews, and enter them in a conference management system. This allows the students to read each other's reviews before class (but only after they submit their own), fostering better interactions.

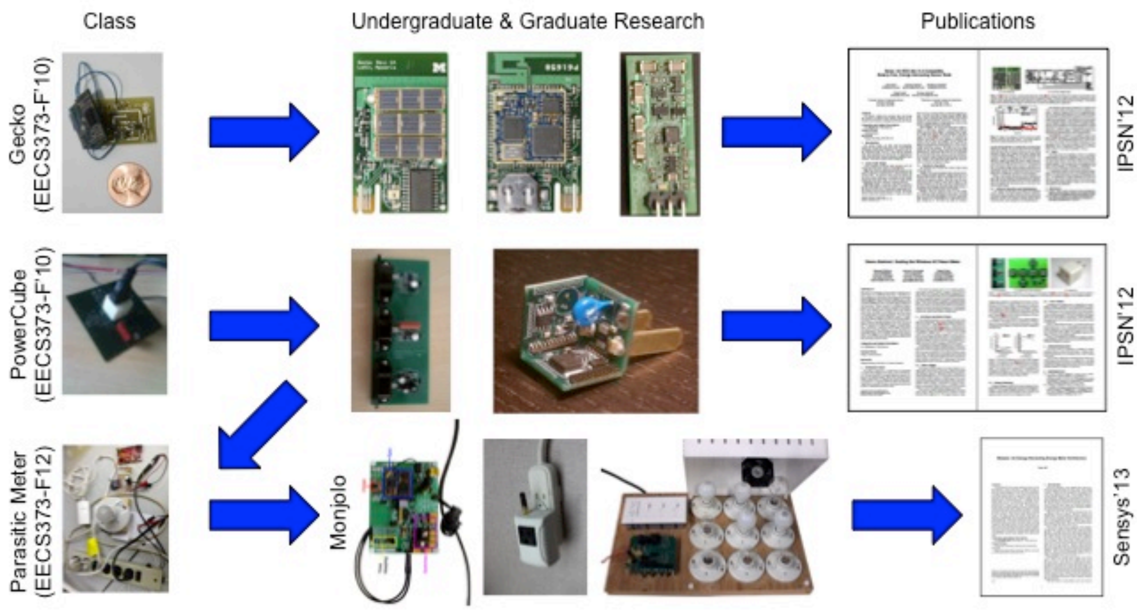


Figure 11: Three undergraduate course projects that were successfully transitioned to published research. Two students completed the class project and carried out the subsequent research without any graduate supervision. I met with the undergraduate students from time to time and advised them on the work. Two projects culminated with the undergraduate students writing a first-authored paper and a third paper motivated a follow on graduate student project. Several papers were published this way, leading to three students entering a Ph.D, program. The skills they learned in the process enabled them to win multiple [redacted] Graduate Fellowships, an [redacted] Graduate Fellowship, and [redacted] of a [redacted] [redacted] Fellowship. The projects inspired subsequent graduate research (and this project proposal).

As a graduate student, I participated in a mock SOSP program committee meeting. This experience was invaluable as it taught me how a reviewer reads a paper and how papers get discussed at a TPC meeting. Few students have such an opportunity, and since I found it to be so beneficial to my own development, I substituted a mock TPC meeting for the final exam. Students review each other's project papers, authors leave the room when their paper is discussed, and the reviewers choose the fate of each paper—accept or reject. This exercise has received such positive feedback and so improved students' ability to self-assess, that I am planning to offer two TPC meetings—one for the proposal (midway) and one for the paper (end) in future versions of this course.

Our Department offers an undergraduate hackerspace called the [redacted] [redacted] [redacted] [redacted] ([redacted]). Since I learned so much from tutoring others and helping them debug their systems, I encourage my Ph.D. students to teach hands-on tutorials and offer office hours. This helps undergraduates with schematic design, board layout, mechanical design, and programming, and helps graduate students acquire critical mentoring skills. I plan to expand this program as more students engage in embedded systems research.

## D.4 Broader Impacts

This project will have several broader impacts, including broadening participation of women in computing, enhancing the infrastructure for research and education, enhancing scientific and technological understanding, and benefitting society.

### D.4.1 Broaden Participation of Women in Computing

I have advised and mentored [redacted] female graduate students in my three and a half-years as faculty. I encouraged them to pursue projects that interested them, notified them of fellowships opportunities, and encouraged them to apply. Indeed, one was awarded an [redacted] Fellowship and another was awarded a [redacted] Scholarship. The third, [redacted], chose me as an adviser after her first semester working with a colleague. I was surprised, since my research is anything but [redacted]. But as I began to look at the pattern, I realized that all of these students had chosen to work with me to pursue *applications* of computer science—crowd-sourcing the hearing aid, healthcare for developing regions, and interactive systems for controlling the building—rather than computer science for the sake of computer science. They viewed my focus on research-enabling research in the service of society as being more aligned with their own motivations. My own experience parallels Fisher and Margolis' findings that, "the context of computing is often very important for many women students. [FM02]." The clear social impact of the proposed research should therefore aid in recruiting female students to this project and the discipline, and avert—at least at small scale—the gender imbalance reported more broadly [CN07, Rosenman10].

### D.4.2 Enhance Infrastructure for Research and Education

In prior work, over fifty Epic-based modules and platforms have been created by nearly two-dozen organizations, dramatically enhancing the infrastructure for research [Epic]. Over 500 Epic-based ACme power meters were deployed at [redacted] for a study of miscellaneous and electrical loads (see collaboration letter). This project will similarly open source key designs and encourage third parties to adopt, use, and improve them.

### D.4.3 Enhance Scientific and Technological Understanding

The outcome of this research program will benefit industry through the creating of energy harvesting data sets, information about the different ways in which resources are consumed in a building, and as demonstrations of future use cases or technology roadmap. For examples, as a diversified semiconductor companies that serve the wireless sensing and communications market, [redacted] and [redacted] have an interest in this work.

### D.4.4 Benefits to Society

The proposed work will have broader societal impact through collaborations that build upon, or leverage, the basic technology developed under this proposal. Several potential collaborations between disciplines and institutions, among U.S. academic institutions, and with industry and government have been identified and are being pursued. The greatest impact potential of these efforts surround environmental and energy monitoring in collaboration with [redacted] (see collaboration letter) and climate and water monitoring with the [redacted] ([redacted] for 2014-2016).

## D.5 Work Plan

**Responsibilities:** The proposed work is to be carried out by the faculty member (PI), one graduate student research assistant (GSRA), and one undergraduate research assistant (URA) per year. The PI will be responsible for (i) curriculum development, (ii) community outreach, (iii) external collaboration, and (iv) project supervision. The GSRA will be responsible for (i) designing, implementing, and validating the network protocols, (ii) synthesizing the work of the URA into a cohesive system, (iii) translating the URA's mesoscale prototypes into microscale systems, (iv) deploying and evaluating the system at medium scale, and (v) supporting the use of the technology by collaborating parties and the research community. The URAs will be responsible for (i) characterizing the available energy and prototyping various sensor devices, (ii) designing several complete energy-harvesting sensor platforms, (iii) developing web infrastructure, services, and applications to process and visualize the sensor data, and (iv) developing mobile phone software to provide a friendly user interface to embedded sensors. Students in the PI's graduate course will explore applications that exercise the underlying infrastructure.

**Timeline:** The following table shows the detailed work plan timeline.

| Year | Electricity Sensing | Climate Sensing | Building Control | Water Sensing |
|---|---|---|---|---|
| Year 1 | Model Emissions | Design Sensors | Lighting Control | |
| | Capture Data | Deploy Sensors | Lighting SW | |
| Year 2 | Design Sensors | Analyze Data | Temp Control | Model Emissions |
| | Deploy Sensors | Publish Results | Temp SW | Capture Data |
| Year 3 | Analyze Data | | Ventil. Control | Design Sensors |
| | Publish Results | | Ventil. SW | Deploy Sensors |
| Year 4 | **Demonstrate "Measured and Controlled" Smart Building** | | | |
| | | | | Analyze Data |
| Year 5 | | | | Publish Results |
| | **Publish Datasets and Compare with Other Instrumented Buildings** | | | |

**Collaborations:** In addition to the work directly carried out by the faculty and students supported by this project, several external parties will collaborate as outlined below and detailed in the supplemental materials. [redacted] ([redacted]) has donated hardware to support the education and research efforts, and provides early access to products and people (see collaboration letter). [redacted] provides mentoring and technical feedback to students through its [redacted] Fellowship program (see collaboration letter). University of [redacted]'s Computer Science Division hosts a small-scale, battery-powered, sensor network to monitor conference room activity but will allow extending this to a medium-scale network to cover other spaces. [redacted] ([redacted]) has deployed a network of 500+ wireless AC power meters that I co-developed previously and they will evaluate the new sensors in a pilot test, validate their performance, and (if satisfactory) seek grant support for a large-scale deployment.

## D.6 PI Qualifications and Results from Prior NSF Support

My research sits at the interface between the cyber and physical worlds, and strives to bridge them through novel hardware/software systems with application to health, energy, and the environment. My research is supported through five (four current and one prior) NSF awards. Intellectual merit and broader impacts of the awards are described below.

**CSR-[redacted] ($2,533,000, 8/15/2011 – 7/31/2016):** “CSR: [redacted].” Applications of wireless sensor nodes are evolving at a previously unimaginable rate but enabling technology limits progress because devices are bulky—measuring one cubic centimeter or more—and hampered by short lifetimes. This project has produced replicable components for a 1 $mm^3$ sensor node. The ultra-miniature device is a complete wireless sensing platform, including transducers (imaging, temperature, pressure, etc.), wireless communications, timer, processor, memory, battery, and energy harvesting.

The central challenge in reducing the form factor lies in reducing power draw and more densely packing the discrete components (crystals, inductors, etc.) of all elements of the system including a near-threshold voltage processor, an ultra-low-leakage memory system, an ultrawide-band (UWB) transmitter and receiver that can communicate with other nodes over a distance of three meters with an integrated antenna, a low-jitter 100 pW timer that is temperature compensated, a new CMOS imaging approach capable of ultra-low power motion detection and image-acquisition, power management circuits for indoor photovoltaic and thermoelectric energy harvesting, and an energy-aware software development environment and operating system to control the nodes. This research has produced dozens of first- and second-generation systems; 100 complete systems will be disseminated to the broader community to enable research in a range of areas.

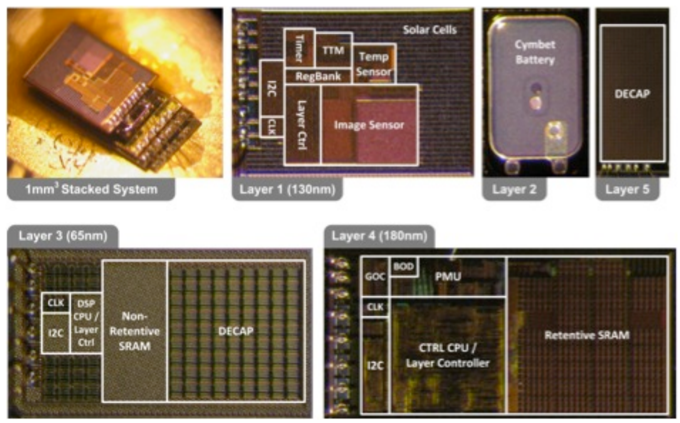


The development of cubic-millimeter sensor nodes enables applications that have long been envisioned but were unachievable—including the applications proposed in this project. Other examples include “sensory skins” that cover surfaces with a dense deployment of nodes to monitor the properties of the manifold itself or its surroundings. Meanwhile, “implantable intelligence” can enable deeply embedded physical and biological processes, e.g., malignant tumor growth monitoring or intra-ocular pressure sensing to reduce risk of retinal detachment. Such applications, and a myriad of other “thinking and linking” applications, can give everyday objects sensing, computing, communication, and tracking ability, enabling, for example, research ranging from the social network patterns of small insects to asset tracking in dynamic environments like hospitals. By shrinking sensor node size to one cubic millimeter, with potentially perpetual lifetime, the concept of “smart dust” is taken from science fiction to reality.

By disseminating the first generation of these sensors to members of the research community, this project dramatically accelerates the adoption of cubic-millimeter-class computing devices. This has broad and immediate impact on a wide array of research programs for intelligently sensing, tracking, measuring, and optimizing physical processes. This research promises a fundamental and long-term impact on a diverse set of societally important applications, including energy conservation, environmental quality management, and health care—the first of which the proposed research attacks.

The project has produced a set of reusable hardware “layers” that will be leveraged in the proposed research. The project has also produced a number of publications that demonstrate an integrated system [LBL+13, LKB+12], mesoscale systems for software development [YCB+12, DCD13], networking [KPSD12], and applications [VRD12].

**CNS-[redacted] ($1,500,000, 2/15/2011 – 1/31/2015): “[redacted]”** The mobile phone's audio jack, which until recently was used to connect a microphone, headphones, and a remote control unit, has become the interface of choice for many new phone peripherals. What was an academic curiosity just a few years ago—whether it was possible to draw enough power from the headset port to power a microcontroller and communicate bidirectionally with it—is now central to a growing set of research projects and commercial products. HiJack demonstrated that it was indeed possible to use the headset port for power transfer and data communications [KVSD10]. This project has demonstrated its generality across a range of mobile phone platform [VRD12] and has supported work on mobile software radios [KPSD12]. This project has also had significant broader impacts. Hundreds of projects across six continents use the HiJack development kit [Seeed13], and HiJack’s basic principles have been used in several research projects including Dr. Chi [Fu12], CO-GPS [LPHR12], Diabeats [Nagesh12], and MusicalHeart [NDL+12]. A network of 12,000 radiation detectors that use the HiJack principle to connect with mobile phones are in use in Japan [bGeigie13]. The devices are used by the concerned public to independently monitor and visualize radiation levels in Japan in the aftermath of Fukushima [SafeCast13]. Commercial headset-based peripherals also exist, most notably for mobile payments—like the Square Reader [Square11], PayPal Here [Paypal12], and Intuit GoPayment [Intuit12]—and television remotes—like the RedEye Mini [ThinkFlood12] and My TV Remote [RYZM11].

**CNS-[redacted] ($315,000, 5/1/2010 – 4/30/2014): “[redacted].”**

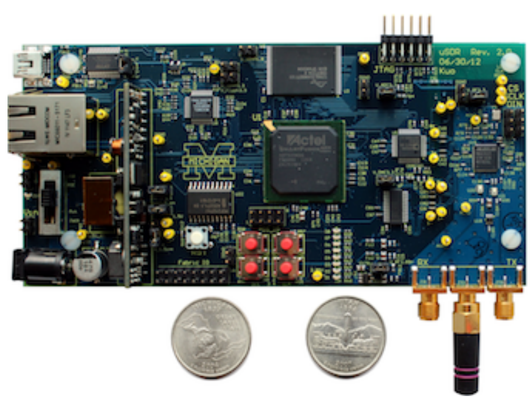

The conventional architecture for short-range, low-power, wireless hardware integrates a radio transceiver and a general-purpose microcontroller. However, microcontrollers have limited computational power and hence, restrict the kinds of algorithms that can be implemented. Moreover, closed radio architectures prevent the direct access to the physical layer needed for many novel applications. The emergence of low-power FPGAs with their efficient duty cycling support presents an opportunity to create a novel wireless platform to overcome these limitations. The project explores a flexible hardware/software architecture for mobile wireless networking based on low-power

FPGA devices. High performance DSP and other algorithms can be implemented directly in hardware, while the rest of the code runs on a soft processor core inside the FPGA. This flexible hardware/software boundary increases the complexity of software development that is eased via sophisticated development tool support and an extensive software component library that the project develops. The, small, inexpensive platform offers two orders of magnitude higher performance than current microcontroller-based hardware at equivalent power for many applications. It has stimulated research in novel protocols and enables new evaluations and applications previously not possible.

This award has supported graduate student's training and the overhaul of an undergraduate embedded systems course that uses a variant of the hardware/software platform. The award has supported the design of a high-resolution, low-power time keeping and synchronization system ([redacted] Best Paper Award [SDS10]) and service [SCA+10], the networking challenges of intermittently powered wireless sensors ([redacted] Best Paper Award [SSSD10]), data over an audio channel ([redacted] Design Contest First Place Award [KSD10]), the design and evaluation of a receiver-initiated link layer protocol ([redacted] Best Paper [DDC+10]), the architecture [DLK+10] and evaluation of a low-power software-defined radio [VSS+10, KPSD12], new programming methods [VSL+10], and rapidly prototyping applications in energy metering [SCD10] and custom ICT platform for developing regions [FDS+10].

**NSF-GRFP ($120,000, 9/1/2005 – 8/31/2007):** This award supported the XSM [DGA+05], Trio [DHJ+06], Radar [DAB06], and Epic [DTJ+08] platforms and their applications [ADB+04, ARE+05, JDD+09]. A 1,000+ node network of XSMs and 557 node network of Trios were deployed, at the time largest and longest-running research sensor networks, respectively. Over fifty Epic-based modules and platforms have been created by nearly two-dozen organizations, dramatically enhancing the infrastructure for research. ***Thirty nine such modules are show inset for illustrative purposes*** [Epic]. Aginova, Arch Rock (now Cisco), Crossbow (now Memsic), Moteiv (now Sentilla), Moteware, SonnOnet, and Vectare have commercialized some of these platforms. The Award supported research into energy metering [DFPC08, JDDC09, JDCS07, JLT+09] ([redacted] Best Paper [DFPC08] and [redacted] Design Contest Winner [DFT+08]), energy-efficient and low-power architectures and operation [CDE+05, DC05, DCS07, JTO+07, TDJ+07], mobile sensing [DC08, DC09], energy tracking [FDL+08], secure code dissemination [DHCC06], multihop data transport (48-hop) wireless networks [KFD+07], link layer primitives [DMST08], time synchronization [KDL+06, SKLD06, SDS10], and link estimation dynamics [SDTL06, SDTL10]. The Award also supported the creation of a course at [redacted] in Fall'05 ([redacted]: Graduate Seminar on Sensor Actuator Networks).

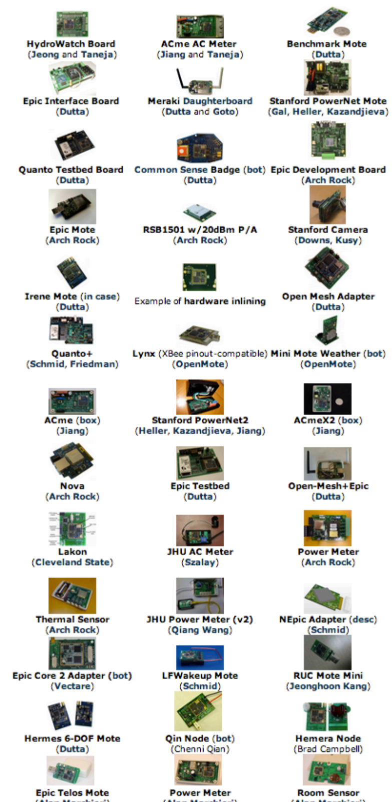

Figure 12: Some of the research platforms enabled by Epic.