The Poly-sensing Environment and Object-Based Emergent Intention Matrix:
Toward an integrated Physical / Augmented Reality Space

- What problem are you solving?
- What technology or usage issues motivate this problem?
- What is your research hypothesis?
- How will you verify that hypothesis?
- What is the potential impact on industry practice if the hypothesis is verified?

1. Overview

Architecture and applications

The Poly-sensing Environment (PSE) is a responsive informational space that promotes rich new forms of collaboration in scientific, cultural and artistic applications. Each PSE is heavily instrumented with a network of multi-modal sensor nodes. Through signal processing and statistical pattern recognition, data collected by the sensor network offers a kind of "machinic perception" of the activities taking place in the environment. The PSE is also a multimedia space equipped with projection displays and multi-channel audio output. Through a software construct known as the Emergent Intention Matrix (EIM), we can "program" the PSE, focusing the attention of the sensor network on specific activities and scripting responses on the room's display hardware.

The media objects available for presentation in the PSE are organized and stored in an "adjoining" virtual space. Occupants of the PSE can use the network of sensors to create new media, or sources can be pulled from the Internet. The coordination between the PSE and the virtual media space has tremendous potential for creating new modes of collaboration; the system enables any object, behavior or "sensed activity" to become an interface to the stored media objects or other sources available through the Internet. The PSE provides a rich new set of models for storing, structuring, accessing, presenting, browsing and searching complex media.

At the heart of the PSE is a network of poly-sensing chips. These specially designed devices contain several sensors as well as a digital signal processor (DSP). To support a broad range of collaborative projects, we plan to make many kinds of "senses" perceptible by the PSE including sound, temperature, pressure, motion, chemical, balance, voltage, and position. Along with the sensor bundle, each poly-sensing chip will also have a radio device that can communicate with other nodes in the network and with several centralized servers. This architecture allows us to coordinate data processing and collection within the PSE, dividing the effort between the DSPs on the nodes and the central servers.

Communication in the PSE extends beyond the sensor network and the accompanying virtual media space. We anticipate constructing several PSEs and coordinating activity sensing and data sharing between the environments. An extension of the EIM can be used to articulate models for security, privacy and overall attention control when we link PSEs.

In Figures 1 and 2 we map out the basic structure of a collection of PSEs. We first illustrate the network of nodes in a single PSE capturing a simple action (in this case, the occupant has moved a book), and then we describe the various representations of this event in both the virtual media space and a series of other PSEs. Through the EIM, we define interesting actions and script the accompanying reactions to be taken by the PSE. Before we explain the programming metaphor that went into designing the EIM, we present a number of collaborative scenarios that can take place in our instrumented environments.

- A scientist could use the sensor network to register any kind of sensed "difference over time" in the environment i.e. social interactions or, from a very different perspective, plant growth. They could use the visualization and sonification capabilities of the virtual media space to generate a time-based associated representation of this change using pre-programmed generative media processes that are chosen through use of the EIM. Networking capabilities would enable this generated environment to be observed anywhere else on the globe, enabling researchers to discuss developments in the space and to see a dynamic representation of the data gleaned from the sensors. As additional focused sensors are added, multi-modal information is recorded as a set of parallel streams – forming an n-dimensional virtual data space. A dynamic relation between the sensor nodes and the server is developed to optimize the system. Multiple representations of the space, each related to a different layer of sensed data, can be brought together to form a time-based "machinic perception" of the ongoing process. Consider the multi-modal collection of light, heat, weight, and pressure data, etc. all informing a cross-referenced model.

- The PSE can also be applied in an educational context. The sensors could be trained to pay attention to specific objects in the space, and the behavior of people relating to those objects. In this case the EIM could author a very literal "associated" relation to the physical environment — "if " this chosen object is picked up "then" get specific information from the database that is linked to that object and provide that information in the augmented reality viewing environment — on screen, "augmented reality glasses", new forms of viewing apparatus, or through a sound system, i.e. someone picks up a particular historical object. The EIM is used to program a focused search to bring up information about that object available in the system. It could be programmed to call up a picture of the object, run a historical video, show a 3D model of the object, play a sound file related to the object, display a text related to the object, etc. An Internet search query could also be facilitated with that information appearing in a window of the viewing environment.

- This also suggests potential application for those with disabilities through new adaptive interface design approaches. An environment could be facilitated where sensed physical behavior could be linked to specific media and/or computer-driven or robotic processes. One might navigate in a wheelchair, turn on and off different media, explore the Internet via various programmed "if" / "then" statements all through the focusing of the Poly-sensing environment.

- In the same way that the educational use of the system can be explored, Seaman is interested in using the system as an artistic/creative/expressive authoring system. Thus, sculptural objects in a room could have poetic texts "associated". These could be called up into the augmented viewing space based on particular interactions or behaviors that are sensed in the space and are focused through the EIM. Also sound files, musical sections, artistic video clips, 3D animations etc. could all be brought into the augmented reality space based on how the "associations" of "if"/"then" statements are authored. Dance could be sensed and dynamic media "associated". New forms of musical interaction/composition could also be explored.

- The networking capabilities could enable one to access the augmented reality aspect of the environment, as well as representations of the data streams in distant locations. This has ramifications for distance learning, where one could navigate on a virtual map of the space and trigger "associated" data streams that are delivered via the Internet or Internet 2; this might serve ethnic groups that might want to access particular historical or educational poly-sensing environments, thus transcending the difficulties of geographic proximity via Internet access;

- The sensor streams could be functionally focused to make the poly-sensing environment a surveillance-based space. Verbauwhede also has worked on security schemes for related sensing technologies, where a user could also insure privacy of sense data.

The Emergent Intention Matrix

With these applications in mind, we now briefly describe the central control or authoring tool for a PSE, the Emergent Intention Matrix or EIM. Essentially this is a modular, object-based encoding scheme that lets the average user of the system define a series of if/then statements connected by conjunctions "and/or"; "if" a particular behavior, action or activity is sensed by the system and some other action is sensed either simultaneously or within a specified window of time, "then" the server goes into action and either presents a particular media element or media process in the adjoining virtual media space. This is achieved through the use of an elaborate matrix of choices. The user of the system can spin container-wheels of a virtual rolodex-like system that enables them to focus the "attention" of the sensors as well as to "associate" a particular media element or process to the chosen behavior in the physical space (see a prototype at the web site http://www.design.ucla.edu/~fwinkler/Poly_Sensing/). One at a time a user can add object-based code potential to the complexity of the set of "observed" activities in the space and the kinds of media that are called up to augment activity in the space. The potential of the system is to enable any object, behavior or "sensed activity" in a given space to become an interface to an interactive augmented reality media space and/or the Internet.

Part of the system would help map the space and articulate different objects and/or behaviors to be sensed or made known to the EIM. Thus a tool to define a particular object used in tandem with embedded sensors (housed in or on those objects) would be developed and would function via linked GPS information in a multi-modal manner to specifically articulate an object or person in space from the perspective of multiple sensor streams. As we learn from data gleaned from the space we will also be able to look for specific kinds of behavior. The notion of sensing an object, moving it, re-scanning the object with various sensors, then subtracting out difference information also presents a possible way to encode objects such that the "system" can recognize them.

The output from the EIM could also drive computer-based affectors, robotic systems, architecturally responsive machine responses, atmospheric relations etc.

Broader impact

Central, is the functional ability of the system to enhance the infrastructure for research, education, the creation of new kinds of machinic sensing research facilities; new forms of instrumentation that arise through visualization and sonification of the sensed environment; the networked potential of sharing this information or cultural production, and the potential partnerships arising through the networking capabilities of the environment. Clearly the long-term broader impact of this IT system is immense in that it advances knowledge across a series of disciplines, and becomes a tool that can later be used by each of the researchers for furthering their individual inquiry and practice. Yet the potential benefits for society at large cannot be underestimated in relation to educational, training, learning, cultural and scientific production. Thus, a connected physical space and associated augmented reality environment can be networked, linking individuals from around the globe.

Dynamically linked physical and virtual interaction, will promote knowledge production, networked communication, the sharing of data, specific visualization and sonification related "transparently" to the physical space, and or expressive/poetic relations "associated" to sensed behavior in the environment. Thus, the augmented reality space is driven by the integration of the sensor nodes, focused by the EIM, enabling ongoing analysis both on the chip via a dynamically responsive DSP as well as on a higher order level analysis on the server. A rich array of hard scientific problems will be researched over the 4-year period. Our goal is to create a "poly-sensing" environment. Indeed, advances in sub-micron and nano-scale semiconductor technologies enable the integration of multiple heterogeneous sensors on one "system-on-a-chip." The unique aspect of our technology will be the collection of information from the parsing of an integrated "collaboration" between a diverse collection of micro-scale sensing devices and larger integrated sensing apparatus. Thus this IT system will integrate research and education by advancing discovery and understanding through the bottom up process of the development of the potentials of the system. A number of researchers are currently exploring the parsing of individual streams of sense data. Our new paradigm will enable every chip to define a "poly-sensing" "neighborhood". Another key feature of the research explores the following: the technology will be created with a flexible and adaptive means of focusing the "attention" of these multiple-sensing devices in recombinant groups of intercommunicating distributed fields. The "Poly-sensing" environment potentially enables the intelligent storage, triggering, and calling-up of media elements and media behaviors, images, relevant libraries, and databases, accessing both local and distributed (Internet based) memory systems as well as operative programs to be elicited by physical events in the "sensed" space. The system enables any object or behavior in the field of "attention" to become an interface through multi-modal sensing. This represents a significant paradigm shift in human/computer interface design.

The Poly sensing environment

Seaman and Hansen will be primarily responsible for the interface, visualization and sonification potentials of the Poly-Sensing Environment, placing special emphasis on the extensibility of the environment for use by other teams of researchers.

The long-term goal is the creation of an advanced tool for scientific, cultural and artistic production. Central, is the functional ability of the system to enhance the infrastructure for research, education, the creation of new kinds of internationally distributed machinic sensing research facilities; linked physical and data spaces; new forms of instrumentation that arise through visualization and sonification of various sensed environments; the networked potential of sharing this information for cultural production; and the potential partnerships arising through the networking capabilities of differing version of the environment. Clearly the long-term broader impact of this IT system is immense in that it advances knowledge across a series of disciplines and becomes a tool that can later be used by each of the researchers for furthering their individual inquiry and practice.

Reconfigurable – reprogrammable embedded system
The emergent intention matrix will receive multiple information streams coming from the distributed sensor field. This poly-sensing environment has several inter-functioning levels. It provides a very interesting driver from an embedded systems research perspective, which will be addressed in the next sections.

Level 1: Poly-sensing data collection

The first level is the collection of raw data by the individual sensing devices. These sensing devices will range from cheap infrared or temperature sensors to more expensive video camera's. The idea is to provide sensory support for each of the human senses: visual, hearing, smell, taste and touch extended with non-human "sensing": chemical (extension of smell and taste), infra-red (extension of vision), ultra-sonic (extension of hearing) and so on.
The sensor nodes don't have to be all the same. For instance, one could sprinkle cheap sensor nodes, as described in [9]or in the Piconet project [45], onto the walls of the room, and combine them with a fewer expensive video nodes. The idea is that the authoring tool selectively turns on or off areas of interest and thus areas of sensors.
In this way, video might be one set of nodes in the system (not on the chip but multiple cameras distributed in the room and possible the virtual rooms) -- or a new kind of "video" image might be constructed out of individual sensor data streams, given a particular perspective and series of single sources of light collection. Recognition may take place on a low level at the chip with high level processing of data at the server once data is collected from distributed chips.

Below is a list of sensors that are currently being investigated for use in our environment: sound (simple sound sensor with the potential of using sonar); heat; pressure; infra-red (and other light spectrum related data); color; ultra-sonic information; sub-sonic information; motion; chemical; balance; voltage; position, etc. Other sensors will be considered in an ongoing manner using a modular design scheme.

Level 2: Embedded Digital Signal processing

The raw data from the sensors is collected and partially processed on the sensor node by the embedded DSP processor. In a first phase this will consist of existing DSP processors and embedded micro-controllers. In a second phase, the goal is to develop flexible domain specific processor units adapted to the specific real-time tasks (see below "level 3").

The signal processing level provides an interesting *computation – communication trade-off*. Indeed, the simplest implementation would consist of forwarding all the raw collected data directly to the central server. However, this creates a huge communication bottleneck and will overload the general purpose processor in the central server. Hence the idea is to provide pre-processing on the nodes. The nodes have a limited processing capability (e.g. no floating point unit) and very limited memory. Thus this results in some interesting research trade-offs. We can identify several signal processing research components:

1. Preprocessing of the data from the individual sensors. These might be relatively simpler tasks such as correlating data between different consecutive measurements. For instance, the node could be instructed to check the variation in temperature or humidity.
2. Higher end signal processing tasks on one stream of data. They are still at the level of one type of sensing. Examples are: image recognition, text recognition, gesture recognition, video, data, sonar data etc. The speech recognition project described below is an example of this level of signal processing.
3. The third level of signal processing consists of signal processing between nodes and between different sensing devices. Here a clear distinction has to be made on what tasks can be executed within one node, and what tasks require either collaboration between nodes or interaction with the intention matrix. For instance, collaborative processing between nodes could indicate the movement of an object and is used to track objects. To save power, the sensors will only be woken up if there is activity in their neighborhood.

This can be accomplished as follows: slow, cheap sensors can monitor activity, and send their data on a slow periodic rate to the central server. This server will, based on the information it receives and the activity it wants to be programmed through the intention matrix, wake up other higher energy sensors. E.g. infra red sensors can be used to detect the presence and approximate location of a human in the room. Based on this information, other sensors will be activated.

Level 3: System on a chip

An iPAQ node is useful as a prototype development with only a few nodes. However, it is not scalable and does not have the low power operation required by a real system. Hence, the next step consists of developing our own embedded system on a chip. The sensor nodes of the poly sensing environment are an excellent driver for our RINGS architecture (Reconfigurable interconnect for next generation systems) [48][49].

Several programs exist where small sensor nodes are being built. Examples are the Smart Dust nodes [Culler] and the Pico radio project [45]. However, these early nodes only contain a small micro-controller and very limited memory. The next generation of nodes will however need to process more complicated real-time signal processing algorithms, e.g. for speech and image recognition, feature extraction, etc. This cannot be done on the small micro controllers, yet at the same time the power and energy constraints are the same.

Domain specific processing engines have the advantage that they combine flexibility (in the form or programmability or reconfigurability) with low energy consumptions. This is especially true for applications that perform poorly on general purpose programming platforms. Examples of such processors are DSP processors and processors for cryptography.

Within the RINGS architecture, as shown on figure 4, these processors are connected together with a reconfigurable interconnect. Together with a small embedded processor, they form a embedded system on a chip. This is an example of a large grain reconfigurable/reprogrammable system. In figure 4, this architecture is applied to en embedded security device [48]. It contains a processing engine for the network connection, a signal processing engine for the template matching and feature extraction (in our case for fingerprint recognition) and an crypto engine to implement all aspects of the security protocol.

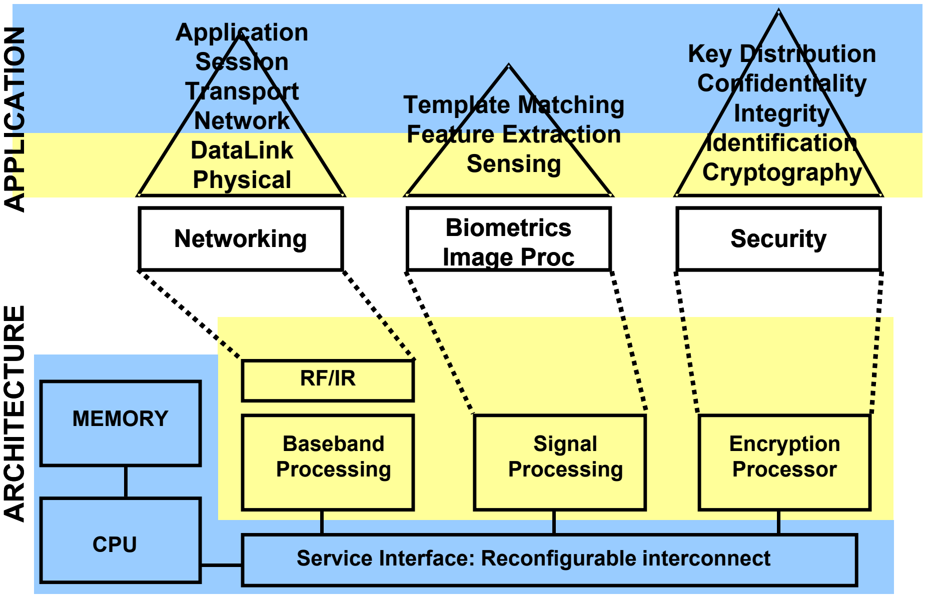

Figure 4: RINGS architecture

The architecture of figure 4 is a generic architecture. In this proposal, we plain to build a flexible low power RINGS architecture for the distributed sensor nodes. Programmability is extremely important because the emergent intention matrix will be allowed to reconfigure and train the nodes.

At this level, it is not the intention to develop new radio protocols or new ad-hoc wireless communication protocols. These have been extensively studied within sensor research projects and commercial products, such as Bluetooth or 802.11 wireless LAN's. For our prototype, the idea is to start from existing radios and integrate them in the poly sensing node prototype.

The speech recognition experiment:

In one experiment, shown in the figure 5 below, we have implemented a speech recognition system on the iPAQ node. A video clip, documenting this experiment can be seen at the following web link http://www.desgin.ucla.edu/~fwinkler/PSE/speech.html . The raw speech data is processed by the iPAQ. A small vocabulary of useful words for the system are trained and stored on the node. The recognized words are then forwarded to the intention matrix on the server which will use it in the "if – then –else" statements of the intention matrix.

For this experiment, the Compaq iPAQ H3835 running Familiar Linux is chosen as our sensor node platform because it has reasonable battery life and it provides a few peripherals that can be used by our project. The iPAQ is a commercial handheld using a StrongARM SA-1110 processor with a memory of 64MB RAM. It includes a built-in microphone. The serial port can be used to connect other sensor devices. A PC running WindowsXP is used as the server. To transmit data we equipped the iPAQ with Orinoco Silver Wireless PC Card, choosing 802.11b as our communication protocol.

This experiment was chosen to study the trade-off between the computation requirement on the embedded device and the bandwidth requirement on the wireless communication channel. The speech recognition task is divided into two parts: front-end feature extraction and recognition. On the sensor part, the ETSI standard is used for extracting a MFCC feature vector of the speech. Then the complex recognition step is realized on server.

The strongARM SA-1110 processor, which is used in the iPAQ, does not have an on-chip floating-point processor, so all the floating-point operations must be emulated in software. The algorithm described in the ETSI distributed speech recognition standard contains a lot of floating-point math functions (e.g. triangle function, logarithm, etc.) and it is usually computation intensive. Research showed that 90% of the energy is consumed by floating-point emulation. On top of this, the execution time is intolerantly long. To meet the time and energy constraints of the embedded system, fixed-point arithmetic is introduced into the source code optimization.

The algorithm portions for extracting the speech features, pre-emphasis, framing, windowing, and computing the power spectrum can be ported into fixed-point in a rather straight forward way. However, when computing Mel spectrum and Mel ceptrum, the multiplier of filter coefficient and the power spectrum needs to be implemented. The product often overflows the 32bit register. To solve this problem, C source code is substituted by assembly code. Another challenge is to take natural logarithm. In our project, a very efficient fixed-point algorithm with low complexity is used to solve this problem.

The analysis of the code shows that the total number of cycles of the code is about 218M to process 2 seconds of speech data and it needs less than 2 sec to finish it on the iPAQ.

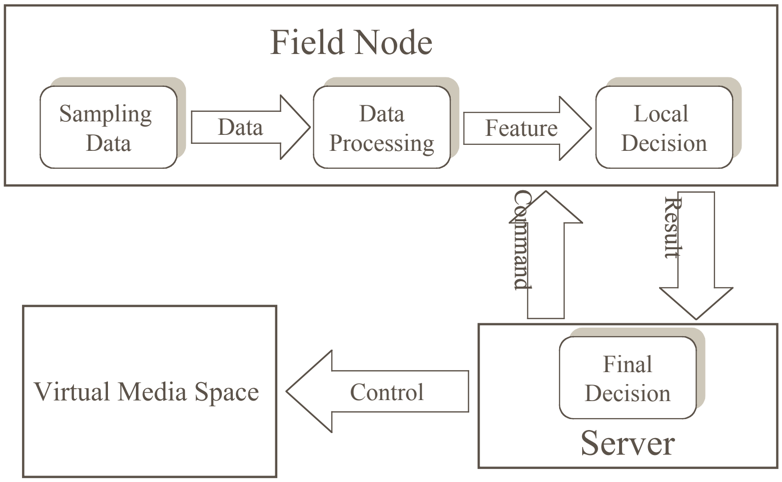

Figure 5: The speech recognition experiment.

The server recognizes the words and will take some actions, as directed by a small emergent intention matrix. For instance, the "open poem" will display a poem on the screen.

Statistical Modeling and the Poly-Sensing Environment

At the core of the poly-sensing environment is a network of heterogeneous sensor nodes. As mentioned previously, each node can be thought of as an embedded system on a chip consisting of one or more sensors together with specialized processing engines that implement specific real-time data analyses. Through the EIM, virtually any user activity can become an interface to the augmented media space and other informational resources. Therefore, the sensor network must be equipped to recognize a variety of events. Nodes will be designed to measure different physical aspects of the environment, allowing us to assemble a rich characterization of users' actions. One component of the proposed research involves developing a framework for expressing high-level events in terms of low-level sensor measurements. Statistical characterizations of user behaviors will be assembled from "information" residing at different nodes in the network. Here, "information" refers to both raw and processed (through the embedded specialty engines) sensor measurements as well as dynamic statistical models designed to capture some aspect of the state of the environment as seen through one or more nodes. To facilitate coordination between nodes, we will develop an event model through which we can transmit both information held at a node and control scripts that direct computations and can guide further propagation of the event through the network. As we will see, this approach extends several threads of current research in communication algorithms for distributed sensor networks, and incorporates an emphasis on the basic tasks associated with statistical analysis and modeling.

Computation and Communication

In early implementations of sensor networks, it was common to have each of the nodes communicate directly to a central processing unit that was responsible for aggregating the various measurements and performing any computations on the data. Statistical models for event detection, object tracking and pattern recognition can all benefit from having access to the "complete" set of measurements. The scale of recent network proposals, however, necessitates other kinds of communication. In the poly-sensing environment, for example, many of the nodes will share information wirelessly. Data transmission to and from these nodes will require much more power, making the existence of a single analysis point impractical. With wider distribution and greater diversity in the underlying sensors, emphasis is now placed on local computation and cooperation between the nodes to determine when and where events have occurred.

As the poly-sensing environment is essentially a very general computer interface device, we must also be concerned with techniques for specifying what actions should be taken once events are recognized. Such higher-level functioning will be articulated through the EIM and "authored" by programmers of the system. This open-ended, extensible set of events and actions makes the poly-sensing environment unique among many sensor network applications, and poses some very real statistical challenges. We propose a flexible event-driven model for communication and computation that is rich enough to support statistical functions like pattern recognition. This emphasis is different from the so-called data-centric motivation of the directed diffusion [36] scheme and even the application-specific approach used in LEACH [50].

To organize local computations and route messages between networked nodes, we envision an event model that builds on the network-based tracking applications of [60] and [61] and the mobile agent approach to data sharing introduced in [58], [56] and [57]. In the former references, the location $\boldsymbol{x}$ of a moving object is described by a dynamic Bayesian model. At each stage in the process, a representation of the model itself is communicated from one node to another. Predictions for $\boldsymbol{x}$ are based on a formal posterior calculation given the set of measurements $z_1, \ldots, z_n$, taken from n previously visited nodes, $s_1, \ldots, s_n$. At the start of the $(n+1)st$ step in the process, the model resides at node s_n. We then attempt to determine which node in the network would have information that would provide the greatest improvement to the estimation of $\boldsymbol{x}$ and is obtainable at the lowest cost. We label this node s_{n+1} and send an approximation of the posterior from s_n to s_{n+1}. Upon arrival at s_{n+1} an observation z_{n+1} is made, and the posterior distribution is updated, conditioning now on the data $z_1, \ldots, z_{n+1}$. At a technical level, we will explore how this approach can be generalized and improved upon using the spline-based, nonparametric descriptions of [71] and [73], possibly in combination with so-called particle-filters [76]. More importantly, however, this approach illustrates how dynamic statistical models can be used to both encode an estimate for some state of the environment under observation, as well as to control communication among the sensor nodes.

In [56] and [58] we find a more explicit proposal for directing activities within the network. Here, mobile agents are transmitted in an attempt to identify the location of the maximum value observed in cluster of sensors that all attend to the same physical measurement. In this application, agents move from node to node, collecting data, processing it in some way on the node and then deciding which node to visit next. While the underlying data analysis is much less involved than that of the tracking application mentioned above, these authors explicitly refer to programmatic specifications (code) that travel with the agents rather than residing on the sensor nodes. Because the poly-sensing environment involves many different kinds of sensors, each with different costs associated with their operation, we can envision an agent's instruction set (code) directing a node to activate one or more of its sensors based on information carried with the agent. By exporting code in this way, we simplify the tasks like reprogramming the environment or generating models to recognize new kinds of behaviors specified in the EIM. It is unclear whether the processing unit on each node should act like a virtual machine, executing code that arrives with each new event, or whether a special kind of reprogramming agent should be developed that (somewhat infrequently) prepares the nodes to receive new data, reducing the overhead associated with each transmission. In either case, a kind of scripting language will be developed to author event-based controls for communication and computation within the network that are tailored to the fundamental statistical operations required by the environment.

Statistical Goals

At a basic level, we need to understand the operating characteristics of the different sensors in the network. These need to be matched, in some sense, to the pattern recognition tasks implied by the EIM. For example, if the location of a person in the environment is an important trigger for some action, then we should consider designing nodes with cameras, proximity sensors, and even microphones. These sensors have different power requirements, and register information about a person's location in very different ways. Combining or fusing measurements from different sensors appears in the literature on robot navigation, and indeed can be thought of as deriving a kind of "mechanic perception" of the environment.

Before any pattern recognition or other more advanced processing can take place, we need to understand the basic output from each kind of sensor, the degree of spatial correlation to be expected from neighboring sensors of the same kind, and neighboring sensors of different kinds. This work is similar to the analysis leading to [63]and [64], where in the context of a larger classification project we considered color measurements (spectral reflectance curves), video output from a CCD array; 1-dimensional time series output from a commercial barcode scanner (NCR); and multivariate time-series output from an artificial olfactory system. To support this exploratory phase, we will develop visualization tools and basic analysis methods to understand patterns evident in the assembled data streams. Special emphasis will be placed on sensor costs and, to the extent possible, effectiveness for identifying behaviors in the room. We will also begin to evaluate which low-level features can be extracted from the stream by one or more of the specialty processing systems that could be embedded along with the sensor on the node. While we are thinking mainly about simple DSP-type operations, we will experiment with more complex functions like speech recognition (mentioned above). In this case, sensor nodes could produce streams of words, with the correlation structure dictated by the position of the speaker in the room. Even at this basic level, we can imagine a simple thresholding operation being applied to sensor output to generate interesting events expressible in the EIM.

In many cases, however, we will need to consider combining information (again, either raw or processed sensor measurements). To represent data of this kind, we will explore multi-modal mixture models. This general framework has been used by [66], [69], [70], [67] and [65] to link text, audio and visual observations. These authors introduce a hierarchical specification of the mixture that allows for browsing and searching activities and use the standard EM algorithm to train the model. We will evaluate these mixtures in the context of dynamic data. Previously, in the context of [72], we applied similar schemes, proposing a variant of the EM computations (see [74] and [75]) that could be implemented directly in a web server. In so doing, the models could be fit (the system would learn) in an online fashion as data arrived at the server. Given the overall power constraints in the poly-sensing environment, we will augment model selection procedures, like the feature weighting ideas of [68] or information criterion [77], and specifically include the cost associated with activating (and communicating with) each sensor. New methods will have to be developed to handle the option that a measurement could be made available if the gain in classification rate is sufficiently large.

To make use of the flexible event structure outlined above, new techniques will have to be developed for expressing the underlying statistical models in a possibly distributed fashion. This will require new methods for training and evaluation. An open question in our design is how much data should be stored at the nodes (be they special processing islands in the network with enhanced storage and computing power or generic nodes) and how much knowledge should reside at the server. We will explore data handling architectures like DIMENSION [78] in relation to the statistical tasks outlined above.

Through statistical analysis of the data streams, we will be able to articulate a particular behavior, recognizing its occurrence in the space. We embody these behaviors in an object-oriented way through the EIM. So far, we have considered only top-down specification of events. We also anticipate searching and browsing capabilities that will allow us to examine different streams of sensor data with the goal of suggesting enhancements to the EIM. The exploratory visualization tools developed in the early phases of the research program will become an interface to these organized streams. The structure of the mixture models will allow efficient interrogation of the raw or partially processed poly-sensing output.

Approach and workplan

Seaman will work with 4 graduate student researchers and an undergraduate intern to:

1) Continue to model and define the functionality of the Emergent Intention Matrix (EIM);
2) Explore the means of articulating multiple visualizations/sonifications related to a chosen physical environment based on different forms of sense data, as that data is parsed and focused, generating a real-time “associated” augmented reality media space.
3) Enable the user to view singular streams as well as to toggle through different visualizations /images of sensed data;
4) Explore the means of superimposing the different representations to begin to formulate a machinic perception of cross-referenced data correlations;
5) Generate and illustrate the potential of “associational” links between the physical environment and the augmented reality space.
6) Communicate with the other researchers and work in a modular manner to incorporate other aspects of the research in an ongoing manner;
7) Make the generative virtual media space inter-operative with the sensing space, and the data processing spaces both on the server and via the DSP, as focused through the user friendly employment of the EIM.
8) Build tools to help articulate objects, behaviors and processes to be entered into the EIM for object-based exploration and recombinant use.
9) Design the system so any user can up-load relevant associated media databases as well as to potentially link in additional sensor streams.

Verbauwhede will work with two graduate students. One graduate student will be responsible for the "software" related issues. The other graduate student will work on the platform (hardware) related research issues.
The first graduate student will in a first stage work on the research challenges associated with implementing signal processing functions onto resource constrained, low power domain specific processing engines. In a first phase, this will be done on existing platforms, such as the StrongArm processor. The results of this work can be used to define optimized signal processing platforms for the RINGS architecture. In the past, Verbauwhede has done similar work for the design of programmable DSP processors for cellular phone systems and for the design of cryptographic co-processors.

The second graduate student will work on the RINGS hardware platform. It consists of building the co-processor units and the reconfigurable interconnect paradigm. In a first phase, a prototype of this architecture will be built on FPGA boards. In a second phase, a VLSI implementation through MOSIS is planned.

Together, these two graduate students will work on the embedded software research issues on the RINGS platform. Indeed, "low power" software development associated with the distributed RINGS architecture is not trivial. More specifically, an intelligent "distribution" of the application has to be made between the different processing engines and the embedded processor.

Hansen will also work with two graduate students. Initially, this team will analyze data coming from prototype poly-sensing nodes as well as other data coming from CENS, UCLA's STC studying sensor networks (discussions are already underway with D. Estrin of CENS). In collaboration with Verbauwhede's group, Hansen will evaluate node-level DSP functionality and other potentially embeddable, on-board specialty processing engines for feature extraction and low-level event detection. Correlation within sensors of different kinds on the same node, as well as sensors on nearby nodes will be studied. Pattern recognition to identify preliminary gestures articulated in the EIM will be conducted using a complete set of sensor data. Visualization and other tools in support of exploratory data analysis will be created.

After these early evaluation tasks are complete, Hansen will begin to develop distributed statistical models for recognition within the poly-sensing environment. Specific studies of power constraints versus classification rates will be performed, possibly augmenting traditional methods for model selection in this context. Next, agent-based representations for communication and control will be considered. A set of primitive actions will be encapsulated in a scripting language. Tradeoffs between the expressiveness of the control language and communication overhead will be studied. Statistical methods will be cast into this framework, and new strategies will be developed for model fitting (training). Particular attention will be paid to the degree of localization achievable in implementing the computations. Interesting scenarios include complete distribution (each node can perform a variety of computations) versus limited distribution (islands of processing units or servers that have more computing resources).

In the final stages of the project, Hansen will develop search engine functionality. This will allow browsing and searching of the data streams emerging from the poly-sensing environment. Integrating new activities into the EIM will depend on a usable interface to these data.

Expected significance

The deepest intellectual questions concerning the project deal with our ability to learn, share knowledge, do research, share expressive media-based experience and communicate/share these process-oriented experiences via technological systems. The goal of developing a form of a multi-modal "machinic perception" that is linked to a memory augmenting space, has a deep history. Our senses present to us a lifelong set of elaborate patterns. Often our senses work in tandem to put forward a "multi-dimensional" perception of the world. The role of memory can not be underestimated in terms of making meaning out of these perceptions. An interest in augmenting memory and meaning production has in the past been embodied in memory techniques and memory theatres. [1] I will here briefly present a set of quotes related to some of these different memory/perception augmentation approaches, as well as relevant historical quotes related to human/computer interaction. In addition the poly-sensing research will explore a number of contemporary ideas surrounding the practice of "modeling" and/ or "abstracting" thought processes for application in computer-based environments. Search engines are one means of augmentation, enabling us to access vast stores of information distributed across the Internet or in more local databases. Here, human memory and computer capabilities work in tandem in the service of thought and learning. As we move to more complex media environments that include dynamic spatial patterns of image, sound and text forming computer-based contexts, it becomes important to articulate new multi-modal search techniques and strategies for advanced distributed connectivity. As we begin to collect information related to parallel data streams, this also becomes an interesting area of research — defining multi-modal search methodologies for such databases. Significantly different than the capacity of human perception, machinic sensing of environments can give us an abstracted and/or augmented perception of chosen spaces. As an artist, one begins to explore a poetics of machinic sensing, as well as abstract association between physical environment and media space as driven by multi-modal sensing and searching. Here are a series of relevant cultural, historical, and scientific passages related to the project:

Guilio Camillo's Teatro Del Mundo was designed as a "memory theatre" working with associational connections between symbolic images and memory. A spectator would sit at a central location inside a portable wooden structure, which contained seven groupings of information, each accessible from seven different levels. The viewer would engage with an environment designed to reveal secrets about the structure of the universe, from the microcosmic to the macrocosmic. [1]

Turing's description of the *ACE (Automatic Computing Engine)*, the first digital computer, saw the potential for a machine with programmed responsive, "operative" input and output "organs." [2] He described this system as being analogous to the mind, suggesting the machine would have "A finite set of states of mind," with the possibility of exploring "groups." We can think of this idea as the initial enabling concept behind the manipulation of constructed modules (or groupings) of particular symbolic entities; in particular media-elements as they relate to parallel streams of gleaned machinic sense data.

Licklider, in "Man-Computer Symbiosis" states "Man-computer symbiosis is an expected development in cooperative interaction between men and electronic computers. It will involve very close coupling between the human and the electronic members of the partnership. The main aims are, 1) to let computers facilitate formulative thinking as they now facilitate the solution of formulated problems, and 2) to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs. In the anticipated symbiotic partnership, men will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work that must be done to prepare the way for insights and decisions in technical and scientific thinking. Preliminary analyses indicate that the symbiotic partnership will perform intellectual operations much more effectively than man along can perform them." [3]

In their study of Vannevar Bush's *As We May Think*, Nyce and Kahn comment:

> Computer and information scientists today recognize Bush's article as containing the earliest description of a machine designed to support the building of trails of association through vast stores of information. Bush's writings on the Memex can be viewed as a proposal for an actual machine and as a body of essays that explore the potential utility and application of new kinds of machines for managing information and representing knowledge. [4]

The American researcher and visionary Ted Nelson, coined the term "hypertext" in 1965. Hypertext is a database style system that makes intelligent links between various linguistic components. It could be likened to an electronic card filing system that gives the user instantaneous access to specific kinds of information. For example, a reader could ask the machine to list all entries related to a particular topic of interest, and be provided with instant access to all related entries. Nelson came up with the idea for a system (called Xanadu) that would connect all of the world's literature by computer within an environment the viewer could interact with. [5] Here in the PSE we are developing physical/data/media linking structures.

It was not until 1987 that Apple introduced Hypercard (authored by Bill Atkinson) which quickly became an ideal hypermedia program for the average computer user. Hypercard is a system that allows a user to make an elaborate set of links between elements of language, image and sound. It provides people who do not write computer code with a means to create elaborate interconnected structures in a non-linear fashion.[5]

It wasn't until a sophisticated editing system was created that the principles for hypertext could be implemented for a real-world purpose. In 1967 the Hypertext Editing System and FRESS (File Retrieval and Editing System) were built at Brown University, Rhode Island, under the leadership of Andries Van Dam, a contemporary of Nelson. FRESS was the first step to create a functioning user controlled language system. [5]

Gordon Pask, in 'The Architectural Relevance of Cybernetics' discusses "Symbolic environments in architecture". He states "Many human activities are symbolic in character. Using visual, verbal or tactile symbols, man 'talks with' his surroundings. These consist in other men, information systems such as libraries, computers or works of art and also, of course, the structures around him." [6]

Daniel C. Dennett, in 'An Empirical Theory of the Mind: The Evolution of Consciousness', states "Since any computing machine at all can be imitated by a virtual machine on a von Neumann machine, it follows that if the brain is a massive parallel processing machine, it too can be perfectly imitated by a von Neumann machine. And from the very beginning of the computer age, theorists used this chameleonic power of von Neumann machines to create virtual parallel architectures that were supposed to model brainlike structures."[7]

Nicholas Negroponte, in an essay on 'Iconographics' from Being digital, states "In 1976 Craig Fields, a program director in the Cybernetics Technology Office a ARPA (and later the director of ARPA itself), commissioned a New York computer animation company to produce a movie of a fictitious desert town called Dar El Marar. The animated movie depicted a cockpit view from a helicopter flying around Dar El Marar, swooping into its streets, pulling back to see the whole townscape, visiting neighborhoods, and moving in close to see into buildings. The movie simulated being Peter Pan, not for the purpose of experiencing the townscape and a world of buildings but for exploring a world of information. The concept assumed that you had designed the town; you had built neighborhoods of information by storing data in particular buildings, like a squirrel storing nuts. Later, you would retrieve the information on your magic carpet by going to where you had stored it." [8] The Poly-sensing environment seeks to link physical space and movement with data space.

Conclusions

The team of Verbauwhede, Hansen and Seaman, working in conjunction with a large group of graduate students over the next four years, has here presented an exciting interdisciplinary research methodology to a series of diverse, yet deeply related scientific problems. These researchers have carefully articulated specific scientific methodologies related to each of their disciplines. Yet, as a team, they bring together a synergistic energy and interdisciplinary team-based methodology that no single member could undertake. The team presents a rich top down approach to the research yet understands that the only way to achieve their common goal is through a series of heavily collaborative, process-oriented, bottom up hypotheses and tests. Meeting the larger goals of the research depends critically on the ongoing intercommunication of the three distinct teams. The difficulty in writing such a grant, to some degree, deals with the differences in language and vocabulary that must be bridged in order to take on such a project. This is a major problem for all interdisciplinary research. Historically each member of the team has been successful in large-scale projects in their own field. They have also each worked on large-scale interdisciplinary projects. Thus, the successful outcome of the research is highly probable. The researchers have clearly articulated a series of benefits to a wide ranging set of communities meeting the criteria set out by the NSF in a number of exciting and groundbreaking ways. As noted in the introduction, there are many exciting potentials related to the development of integrated sensing, networking and perception augmenting technologies at this time. In particular this research seeks to augment thought, perception, research, education, learning, human adaptability, and creative expression by developing integrated sensing/augmented reality spaces that through networking link multiple environments for shared knowledge; allied physical/data-oriented information schemes and processes; and dynamic engagement with search engines that call forth digital content "associated" with sensed behavior in a dynamically conjoined physical and data space. In particular the three interrelated teams are exploring the continuum between data space and physical space in an entirely novel manner. This research clearly impacts educational, scientific, artistic, and cultural realms. Perhaps the most beneficial to society is the fact that this research also suggests potential application for those people with disabilities, through new adaptive interface design approaches.

References:

[1] YATES, F. 1966. *The Art of Memory.* Chicago: University of Chicago Press.

[2] TURING, A.M. 1986. A.M. Turing's ACE Report of 1946 and Other Papers. Volume 10, p36. In: B.E. CARPENTER and R.W. DORAN, eds. The Charles Babbage Institute Reprint Series for The History of Computing. Cambridge/London: MIT Press, pp.21-124.

[3] Licklider, JRC, "Man-Computer Symbiosis" IRE Transactions on Human Factors in Electronics, Volume HFE-1, March 196.

[4] NYCE, J. and KAHN, P. 1991. *From Memex to Hypertext, Vannevar Bush and the Mind's Machine*. Boston: Harcourt Brace Jovanovich, Inc.

[5] Danny Goodman, The Complete Hypercard Handbook, Bantam Books, New York, 1987.

[6] Pask, Gordon, 'the Architectural Relevance of cybernetics', *Architectural Design,* 1969

[7] Dennett, Daniel C., Consciousness Explained, London: Penguin books Ltd., 1993.

[8] Negroponte, Nicholas, Being Digital, London: Hodder & Stoughton, 1995.

[9] D. Estrin, L.Girod, G.Pottie, M.Srivastava, "Instrumenting the world with wireless sensor networks" ICASSP 2001. Proceedings, 2001 IEEE international conference on Vol 4, 2001. P.2033-2036.

[10] Estrin, D.; Govindan, R.; Heidemann, J.; Kumar, S. Next century challenges: scalable coordination in sensor networks. MobiCom'99. Proceedings of Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking, New York, NY, USA: ACM, 1999. p.263-70. x+279 pp.

[11] S. Kumar, D. Shepherd, F. Zhao, "Collaborative Signal and Information Processing in Micro-Sensor Networks." IEEE Signal Processing Magazine, March 2002, to appear.

[12] Blatt, S.R. Collaborative signal processing for sensor networks. Proceedings of the SPIE - The International Society for Optical Engineering, vol.4126, (Integrated Command Environments, San Diego, CA, USA, 31 July 2000.) SPIE-Int. Soc. Opt. Eng, 2000. p.183-5.

[13] Beck, S.D.; Reynolds, J. Collaborative signal processing for a cluster of inexpensive distributed sensor nodes. Proceedings of the SPIE - The International Society for Optical Engineering, vol.4126, (Integrated Command Environments, San Diego, CA, USA, 31 July 2000.) SPIE-Int. Soc. Opt. Eng, 2000. p.97-108.

[14] Kokar, M.; Kim, K.H. Review of multisensor data fusion architectures and techniques. Proceedings. IEEE International Symposium on Intelligent Control (Cat. No.93CH3278-9), Proceedings of 8th IEEE International Symposium on Intelligent Control, Chicago, IL, USA, 25-27 Aug. 1993.) New York, NY, USA: IEEE, 1993. p.261-6. 620 pp.

[15] Luo, R.C.; Kay, M.G. Multisensor integration and fusion in intelligent systems. IEEE Transactions on Systems, Man and Cybernetics, vol.19, (no.5), Sept.-Oct. 1989. p.901-31.

[16] Bloch, I. Information combination operators for data fusion: a comparative review with classification. IEEE Transactions on Systems, Man & Cybernetics, Part A (Systems & Humans), vol.26, (no.1), IEEE, Jan. 1996. p.52-67.

[17] Hairong Qi; Iyengar, S.; Chakrabarty, K. Multiresolution data integration using mobile agents in distributed sensor networks. IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews), vol.31, (no.3), IEEE, Aug. 2001. p.383-91.

[18] Stephanie Lindsey, C. S. Raghavendra, K. Sivalingam, "Data Gathering in Sensor Networks using the Energy*Delay Metric", Proceedings of the IPDPS Workshop on Issues in Wireless Networks and Mobile Computing, April 2001

[19] Bass, T. Intrusion detection systems and multisensor data fusion. Communications of the ACM, vol.43, (no.4), ACM, April 2000. p.99-105.

[20] Martinerie, F. Data fusion and tracking using HMMs in a distributed sensor network. IEEE Transactions on Aerospace and Electronic Systems, vol.33, (no.1), IEEE, Jan. 1997. p.11-28.

[21] Viswanathan, R.; Varshney, P.K. Distributed detection with multiple sensors I. Fundamentals. Proceedings of the IEEE, vol.85, (no.1), IEEE, Jan. 1997. p.54-63.

[22] Blum, R.S.; Kassam, S.A.; Poor, H.V. Distributed detection with multiple sensors I. Advanced topics. Proceedings of the IEEE, vol.85, (no.1), IEEE, Jan. 1997. p.64-79.

[23] S. Meguerdichian, S. Slijepcevic, V. Karayan, M. Potkonjak, "Localized Algorithms In Wireless Ad-Hoc Networks: Location Discovery and Sensor Exposure" MobiHOC 2001, pp. 106-116, October 2001.

[24] N. B. Priyantha, A. Chakraborty, H. Balakrishnan, "The Cricket Location-Support System," Proc. of the Sixth Annual ACM International Conference on Mobile Computing and Networking (MOBICOM), August 2000.

[25] A. Ward, A. Jones, A. Hopper, "A New Location Technique for the Active Office," IEEE Personal Communications, Vol. 4, No. 5, October 1997, pp 42-47

[26] Goel, P.; Roumeliotis, S.I.; Sukhatme, G.S. Robust localization using relative and absolute position estimates. Proceedings 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human and Environment Friendly Robots with High Intelligence and Emotional Quotients (Cat. No.99CH36289), (vol.2), Piscataway, NJ, USA: IEEE, 1999. p.1134-40 vol.2. 3 vol. xxxviii+1925 pp.

[27] Nirupama Bulusu, Deborah Estrin, Lewis Girod and John Heidemann Scalable Coordination for Wireless Sensor Networks: Self-configuring Localization Systems. Proc. of the 6th International Symposiumon Communication Theory and Applications (ISCTA'01),Ambleside, UK, July 2001.1

[28] F. Zhao, J. Shin, J. Reich, "Information-Driven Dynamic Sensor Collaboration for Tracking Applications." IEEE Signal Processing Magazine, March 2002.

[29] L. Guibas, "Sensing, tracking, and reasoning with relations." IEEE Signal Processing Magazine, March 2002.

[30] Wang, A.; Heinzelman, W.B.; Sinha, A.; Chandrakasan, A.P. Energy-scalable protocols for battery-operated microsensor networks. Journal of VLSI Signal Processing Systems for Signal, Image, and Video Technology, vol.29, (no.3), (1999 IEEE Workshop on Signal Processing Systems. SiPS 99. Design and Implementation, Taipei, Taiwan, 20-22 Oct. 1999.) Kluwer Academic Publishers, Nov. 2001. p.223-37.

[31] Alice Wang; Chandrakasan, A. Energy efficient system partitioning for distributed wireless sensor networks. 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221), (vol.2), (2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221), 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings, Salt Lake City, UT, USA, 7-11 May 2001.) Piscataway, NJ, USA: IEEE, 2001. p.905-8 vol.2. 6 vol. xci+2688 pp.

[32] Tsiatsis, V.; Zimbeck, S.A.; Srivastava, M.B. Architecture strategies for energy-efficient packet forwarding in wireless sensor networks. ISLPED'01: Proceedings of the 2001 International Symposium on Low Power Electronics and Design (IEEE Cat. No.01TH8581), New York, NY, USA: ACM, 2001. p.92-5. xii+395 pp.

[33] Meguerdichian, S.; Koushanfar, F.; Potkonjak, M.; Srivastava, M.B. Coverage problems in wireless ad-hoc sensor networks. Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213), (vol.3), Piscataway, NJ, USA: IEEE, 2001. p.1380-7 vol.3. 3 vol. xxvi+1810 pp.

[34] Savvides, A.; Park, S.; Srivastava, M.B. On modeling networks of wireless microsensors. Performance Evaluation Review, vol.29, (no.1), (SIGMETRICS 2001/PERFORMANCE 2001. Joint International Conference on Measurement and Modeling of Computer Systems, Cambridge, MA, USA, 16-20 June 2001.) ACM, June 2001. p.318-19.

[35] Stephanie Lindsey, C. S. Raghavendra, "PEGASIS: Power Efficient Gathering in Sensor Information Systems," IEEE Aerospace Conference, March 2002

[36] Intanagonwiwat, C.; Govindan, R.; Estrin, D. Directed diffusion: a scalable and robust communication paradigm for sensor networks. MobiCom 2000. Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking, New York, NY, USA: ACM, 2000. p.56-67. xii+300 pp.

[37] Coffin, D.A.; Van Hook, D.J.; McGarry, S.M.; Kolek, S.R. Declarative ad-hoc sensor networking. Proceedings of the SPIE - The International Society for Optical Engineering, vol.4126, (Integrated Command Environments, San Diego, CA, USA, 31 July 2000.) SPIE-Int. Soc. Opt. Eng, 2000. p.109-20.

[38] Yao, K.; Reed, C.; Hudson, R.; Lorenzelli, F. Random array processing for micromachined sensor network. Proceedings of the SPIE - The International Society for Optical Engineering, vol.3162, (Advanced Signal Processing: Algorithms, Architectures, and Implementations VII, San Diego, CA, USA, 28-30 July 1997.) SPIE-Int. Soc. Opt. Eng, 1997. p.263-74.

[39] Papadopoulos, H.C.; Wornell, G.W.; Oppenheim, A.V. Low-complexity digital encoding strategies for wireless sensor networks. Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98 (Cat. No.98CH36181), (vol.6), New York, NY, USA: IEEE, 1998. p.3273-6 vol.6. 6 vol. lxiii+3816 pp.

[40] Pradhan, S.S.; Ramchandran, K. (Edited by: Storer, J.A.; Cohn, M.) Distributed source coding: symmetric rates and applications to sensor networks. Proceedings DCC 2000. Data Compression Conference, Los Alamitos, CA, USA: IEEE Comput. Soc, 2000. p.363-72. xv+586 pp.

[41] Heidemann, J.; Silva, F.; Intanagonwiwat, C.; Govindan, R.; Estrin, D.; Ganesan, D. Building efficient wireless sensor networks with low-level naming. Operating Systems Review, vol.35, (no.5), (18th ACM Symposium on Operating Systems Principles (SOSP'01), Banff, Alta., Canada, 21-24 Oct. 2001.) ACM, Dec. 2001. p.146-59.

[42] M. Chu, H. Haussecker, F. Zhao, "Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks." Int'l J. High Performance Computing Applications, to appear, 2002. Also, Xerox Palo Alto Research Center Technical Report P2001-10113, May 2001

[43] F. Zhao, X. Koutsoukos, H. Haussecker, J. Reich, P. Cheung, C. Picardi, "Distributed monitoring of hybrid systems: a model-directed approach." Proc. of Int'l Joint Conf. on Artificial Intelligence, Seattle, August 2001

[44] Yuntao Cui; Samarasckera, S.; Qian Huang; Greiffenhagen, M. Indoor monitoring via the collaboration between a peripheral sensor and a foveal sensor. Proceedings. 1998 IEEE Workshop on Visual Surveillance (Cat. No.98EX129), Los Alamitos, CA, USA: IEEE Comput. Soc, 1997. p.2-9. vii+94 pp.

[45] J. Rabaey, J. Ammer, J. da Silva, D. Patel, S. Roundy, “Picoradio supports Ad Hoc Ultra-Low Power wireless networking,” IEEE Computer Magazine, July 2000.

[46] P. Schaumont, I. Verbauwhede, K. Keutzer, M. Sarrafzadeh, “A quick safari through the reconfiguration jungle,” Proceedings 38th Design Automation Conference, DAC 2001, pg. 172-177, Las Vegas, June 2001.

[47] P. Schaumont, I. Verbauwhede, “A reconfiguration hierarchy for Elliptic curve cryptography,” Proceedings 35th Asilomar conference on signals systems and computers,” Monterey, CA, Nov. 2001.

[48] P. Schaumont, I. Verbauwhede, “Domain Specific co-design for Embedded security,” IEEE Computer Magazine, April 2003, special issue on HW/SW co-design.

[49] I. Verbauwhede, M.F. Chang, “Reconfigurable Interconnect for next generation systems,” 2002 ACM/Sigda International workshop on System Level Interconnect Prediction (SLIP), April 2002, San Diego, CA.

[50] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "An Application-Specific Protocol Architecture for Wireless Microsensor Networks," IEEE Transactions on Wireless Communications, Vol. 1, No. 4, October 2002, pp. 660-670.

[51] C. Intanagonwiwat, R. Govindan, and D. Estrin, “Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks,” in Proc. ACM MobiCom ’00, pp. 56–67, August 2000.

[52] S. Tilak, N. Abu-Ghazaleh, and W. Heinzelman, "A Taxonomy of Wireless Micro-Sensor Network Models," ACM Mobile Computing and Communications Review (MC2R), Volume 6, Number 2, April 2002.

[53] A. Singhal and C. Brown. Dynamic Bayes net approach to multimodal sensor fusion. Proceedings of the SPIE - The International Society for Optical Engineering, 3209:2--10, October, 1997.

[54] D. Estrin, R. Govindan, and J. Heidemann. Scalable coordination in sensor networks. Technical report 99-692, USC Information Sciences Institute, May 1999.

[55] Deepak Ganesan, Deborah Estrin In Proceedings of First Workshop on Hot Topics in Networks (Hotnets-I), October 2002

[56] H. Qi, S. S. Iyengar, K. Chakrabarty, "Multi-resolution data integration using mobile agents in distributed sensor networks," IEEE Transactions on Systems, Man, and Cybernetics Part C: Applications and Reviews, vol. 31, no. 3, pp383-391, August, 2001.

[57] H. Qi, S. S. Iyengar, K. Chakrabarty, "Distributed sensor networks - a review of recent research," Journal of the Franklin Institute, vol. 338, pp. 655-668, 2001.

[58] H. Qi, X. Wang, S. S. Iyengar, K. Chakrabarty, "Multisensor data fusion in distributed sensor networks using mobile agents," Information Fusion, TuC2-11-16, Canada, August, 2001.

[59] M. Chu, H. Haussecker, and F. Zhao. “Scalable Information-Driven Sensor Querying and Routing for ad hoc Heterogeneous Sensor Netoworks”. Int’l J. of High-Performance Computing Applications, vol. 16, no. 3, Fall 2002.

[60] F. Zhao, J. Shin, J. Reich, "Information-Driven Dynamic Sensor Collaboration for Tracking Applications." IEEE Signal Processing Magazine, March 2002.

[61] J. Liu, J. Reich, F. Zhao, "Collaborative In-Network Processing for Target Tracking." J. on Applied Signal Processing, 2002.

[62] L. Guibas, "Sensing, tracking, and reasoning with relations." IEEE Signal Processing Magazine, March 2002.

[63] M. H. Hansen, H. Hess, P. Mitra, and G. Thomas, U.S. Patent No. 6,424,745, "Method and Apparatus for Object Recognition," granted July 2002.

[64] G. Thomas, M. H. Hansen and D. X. Sun, U.S. Patent No. 6,075,594, "System and Method for Spectroscopic Product Recognition and Identification," granted June 2000.

[65] Barnard, K. and Forsyth, D. (2001). Learning the semantics of words and pictures, International Conference on Computer Vision, Vol. 2, pp. 408– 415.

[66] Barnard, K., Duygulu, P. and Forsyth, D. (2001). Clustering art, Computer Vision and Pattern Recognition, Vol. 2, pp. 434–439.

[67] Eric Brochu, Nando de Freitas and Kejie Bao. The Sound of an Album Cover: Probabilistic Multimedia and IR. AI-STATS.

[68] Peter Carbonetto, Nando de Freitas, Paul Gustafson and Natalie Thompson. Bayesian Feature Weighting for Unsupervised Learning, with Application to Object Recognition. AI-STATS.

[69] Eric Brochu and Nando de Freitas. "Name that Song!": A Probabilistic Approach to Querying on Music and Text. NIPS 2002.

[70] Pinary Duygulu, Kobus Barnard, Nando de Freitas and David Forsyth. I. Jordan. Object Recognition as Machine Translation: Learning a Lexicon for a Fixed Image Vocabulary. ECCV 2002.

[71] M. H. Hansen and C. Kooperberg (2002), Spline Adaptation in Extended Linear Models, with discussion, *Statistical Science*, Vol. 17, No. 1, pp. 2-51.

[72] M. H. Hansen and E. Shriver (2001), Using navigation data to improve IR functions in the context of Web search, *Proceedings of CIKM 2001*, pp. 135-142.

[73] C. Stone, M. H. Hansen, C. Kooperberg, Y. Truong (1997), Polynomial Splines and Their Tensor Products in Extended Linear Modeling, (with discussion), *The Annals of Statistics*, Vol. 25, No. 4, pp. 1371—1470

[74] M. Sato and S. Ishii (2000), "On-line EM algorithm for the normalized gaussian network, Neural computation, 12 (2), pp 407-432.

[75] Neal, R. and Hinton, G. E. (1998), "A view of the EM algorithm that justifies incremental, sparse and other variants," in M. I. Jordan (editor), Learning in Graphical Models, pp 355-368.

[76] Arnaud Doucet, Nando de Freitas & Neil Gordon (eds). Springer-Verlag. 2001. ISBN 0-387-95146-6.

[77] M. H. Hansen and B. Yu (2001), Model Selection and the Principle of Minimum Description Length, *Journal of the American Statistical Association*, Vol. 96, No. 454, pp. 746-774

[78] Deepak Ganesan, Deborah Estrin, and John Heidemann. "Dimensions: Why do we need a new data handling architecture for sensor networks?" 152 In First Workshop on Hot Topics in Networks (Hotnets-I), Princeton, New Jersey, USA, October 2002.

Figure 1. Top: The poly-sensing environment – exploring poetic/informational grammars of attention and functionality. Bottom: Interacting with the poly-sensing environment.

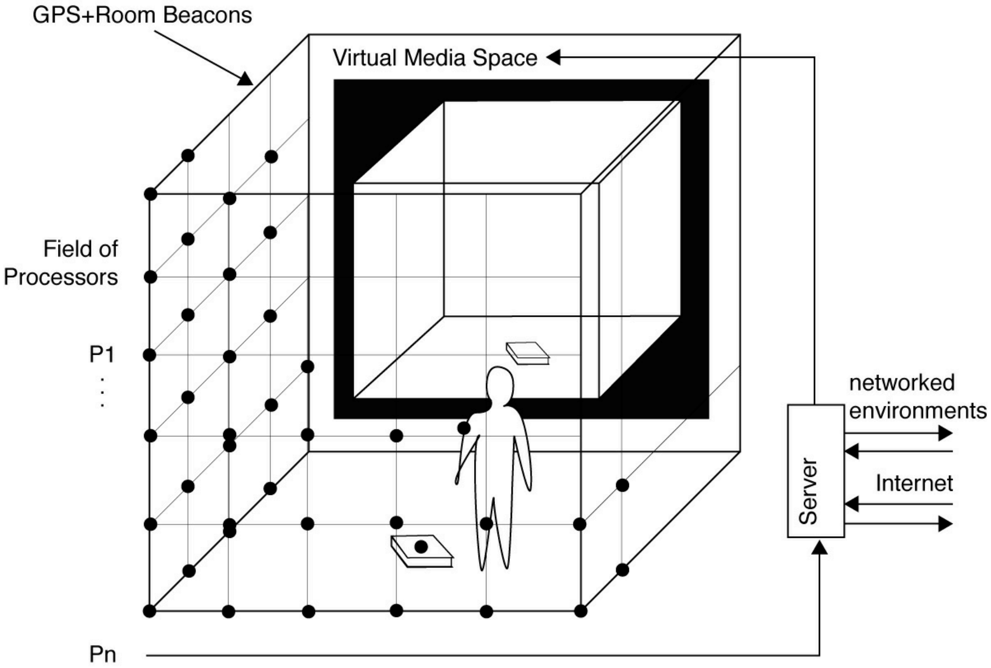

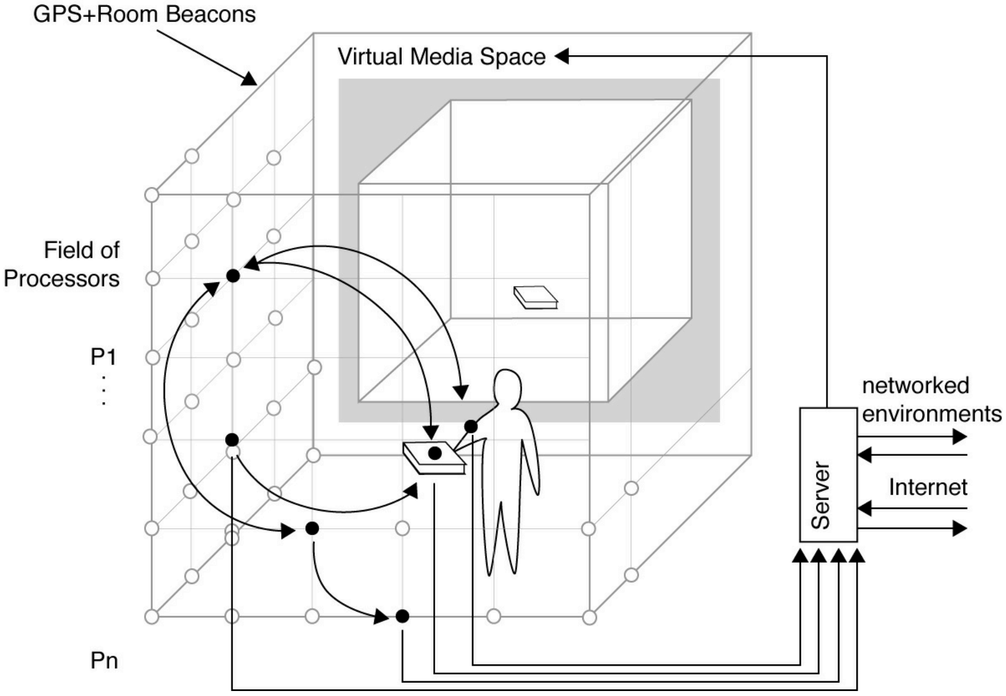

Figure 2. Communicating with the virtual media space and with a collection of other PSEs.