1. INTRODUCTION

We propose to develop an intelligent network for sensing and analyzing the local environment. The network will be composed of smart dust particles. Each dust particle will be a millimeter size system that contains sensors, a micro-controller, a transceiver and a power source. An illustration of a dust particle is shown in Figure. 1.1. The network can contain several to hundreds or even thousands of dust particles. The dust particles sense physical attributes of the environment, and then communicate these attributes to other dust particles. The communication network is wireless. Using distributed computing algorithms, the dust particles fuse their data, and make decisions based on the cumulative information of the overall network. The unique aspects of the network will be integration of unique communication algorithms with physics-based designs, including MEMS technology and 3D integration.

Developing such a network is a challenge. Dust particles will be small size, but must still house sufficient radio frequency and computational circuitry to operate effectively on low power. Dust particles will either be powered externally or by an internal power source that is commensurately sized with respect to the particle.

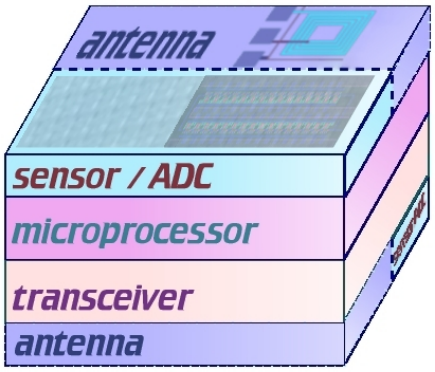

Figure 1.1. Dust Particle Illustation

To perform RF communication, computations and sensing while minimizing power use, and maximizing flexibility, we propose to develop an *ad-hoc* smart sensor network. The network will combine time division multiplexing (time-division multiple access or TDMA) with frequency division multiplexing (frequency-division multiple access or FDMA). We plan to have the network self-assemble into a set of overlapping, flexible cells. Each cell will have a small number of dust particles, a small number of operating frequencies, and a small number of time allocations. Information will be sensed and decisions made within each cell. Intracellular information will then be fused with that from other cells to make network-wide decisions. The information will be communicated digitally using frequency shift keying (FSK) modulation. To develop this network, we have divided the tasks into six major categories:

- Radio Frequency CMOS VLSI Hardware
- Communication Protocols, Networking and Algorithm Development: Energy-Efficient Smart-Dust Networks for Fusion
- Digital Design: Microcontrolling and Computational Components for Interfacing with RF Transceiver, Managing Network Protocols and Data Analysis.
- Antenna Design
- Power: MEMS, Batteries, Capacitor Structures and Solar Cells
- Analog Sensing

We plan to develop the smart dust network in several phases.

- **Phase 1**: Develop a smart dust prototype ad-hoc network using commercially off the shelf (COTS) parts. This network will be prototyped on PC boards, and be used to test basic communication algorithms; set up the frequency and time divisions for the FDMA-TDMA multiplexing schemes, and establish channel and source coding protocols. The test vehicle for this prototype network will be to first sense a local maximum (e.g. temperature or light, for example). Next, we will sense the location of the local maximum. Finally, we plan to track the maximum's trajectory in time and space.
- **Phase 2**: Develop smart dust prototype network using chips designed by ourselves, and fabricated through MOSIS. Use well-established process that is not economically prohibitive, and one that we already have experience with (either the 0.25µm or the 0.5 µm MOSIS). Start integrating IC's into 3D systems. Experiment with other sensors to detect chemical agents and vibrations. Continue to improve algorithms for power and noise tradeoffs.
- **Phase 3**: Develop state-of-the-art smart dust ad-hoc network using a 0.13 or 0.09 micron process. Scale the technology we developed in Phase 2 to lower dimensions that are commensurate with the higher frequency, low size, lower power 0.13µm or 0.09µm processes (use which ever is more appropriate 18 months from now).

2. RADIO FREQUENCY VLSI HARDWARE

Design Criteria:

The dust transceiver system will have to be low power, low noise, and have physical dimensions in the millimeter range. In addition, the transceiver must be relatively inexpensive to fabricate so that thousands can be deployed of without prohibitive costs. The frequency of operation must be in the gigahertz range, so that commensurately sized antennas may be utilized. The design will need to accommodate the TDMA-FDMA multiplexing scheme using FSK modulation in order to establish the ad-hoc network described in the Communication section of this white paper. Also, the design should be well suited for 3D integration, which is a key technology we are developing with the LPS. Incorporation of 3D integration transcends the current state of the art, which mainly relies on 2D integration, which should facilitate the development of unique compact smart dust systems. The size, frequency, modulation, power, 3D integration and cost criteria all argue for the use of CMOS technology to develop and fabricate smart dust particles. In our background research, we have developed numerous RF chips and IC's for 3D integration. In addition, we have developed numerical models for analyzing the parasitic effects of EM cross talk and heat production in 3D IC's[1-6]. This background work also points to the use of CMOS technology for achieving 3D smart dust particles. With the constant reduction of feature size, CMOS transistors are operating well into the GHz range. Currently, IBM is offering its 0.13 micron CMOS-RF process through the MOSIS foundry. This process boasts a unity gain frequency f_T = 90GHz. The next generation, which is the 0.09 micron process, has already been made commercially available from UMC and it claims to have an f_T = 160GHz. Over the next ten years, the Semiconductor Industry Association Roadmap indicates that gate lengths will continue to be reduced to below 0.035 microns, and operating frequencies should increase at least linearly with gate length reduction. Such high frequency operation should allow for continued scaling of dust particles to sub millimeter dimensions. In addition, reduction in size also corresponds to

commensurate reduction in supply voltages, and for various components of the system, a quadratic reduction in power.

Transceiver Topology:

The transceiver topology is shown in Fig. 2.1. The transceiver will be narrow band and will contain receiver and transmitter blocks. Initial designs will be centered to operate in the 10GHz range with a bandwidth of approximately 20MHz. Noise levels at the input should be small enough so that we can detect signals which are at least as low as -100dBm. Precise values of these numbers will depend on the technology, power, noise levels, and the details of the TDMA-FDMA communication network we develop. The receiver hardware will consist of a low-noise amplifier (LNA), a mixer, an IF amplifier, and a local oscillator (PLL-frequency synthesizer). The transmitter will require a frequency synthesizer, a power amplifier and will share the PLL with the receiver.

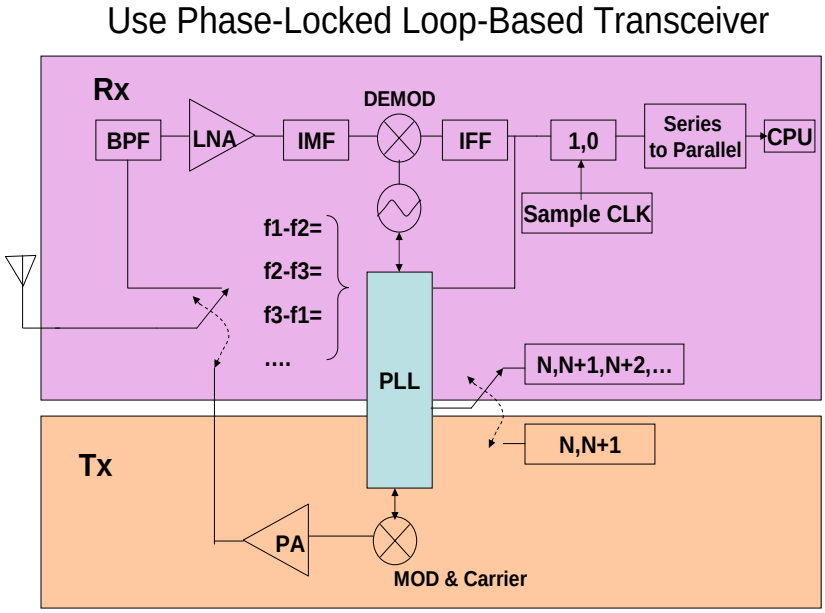

Figure 2.1. Transceiver Topology

- **Phase-Locked Loop**: The transceiver will be centered around a second order phase-locked loop topology. The PLL will be used for frequency synthesis required for both broadcasting and tuning. The PLL will be composed of a phase detector, filter, voltage-controlled oscillator (VCO), and a counter. The counter will be controlled by the microprocessor. Tuning in both the frequency (FDMA) and time (TDMA) domains will be electronic by interfacing the PLL with the microcontroller component of the dust system. The PLL will also serve as the demodulator for the FSK signal by reading corresponding voltage levels from the input of the PLL's VCO. The layout of a PLL chip we designed and had fabricated is shown in Figure 2.2.

- **Low Noise Amplifer**: The RF receiver input will be a cascode low noise amplifier (LNA) designed to be impedance matched to the antenna, while at the same time providing approximately 20dB gain and simultaneously minimizing power consumption and noise. This constraint of minimizing power and noise simultaneously will be a research challenge. Generally, noise decreases when transconductance increases. However, supply current also increases with higher transconductance. We will optimize our smart dust system so that for the given communication algorithms as discussed in Section3, we will optimize the noise/power tradeoff. Use of minimum supply voltages will be explored which should help to minimize power without compromising noise performance.

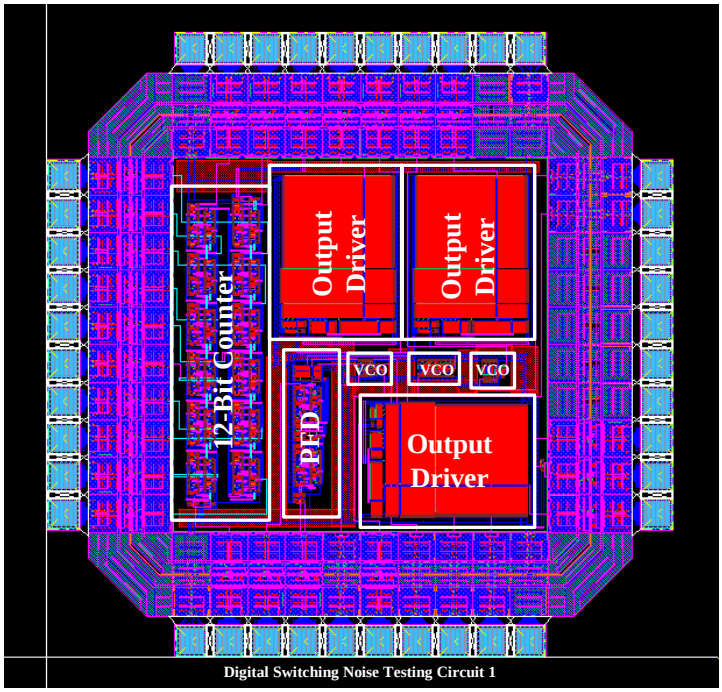

Figure 2.2: PLL Chip Designed for FSK Modulation

- **Mixer**: The initial mixer will be of the Gilbert-cell type, with accommodations made for maximizing swing at low DC bias levels. The PLL will serve as the local oscillator. The mixer will mainly serve as a down-converter to the intermediate frequency (IF). Other mixers topologies will be explored in an effort to minimize supply voltage levels.
- **Power Amplifier:** We will investigate the use of switch-mode type transmitter power amplifiers (PA) in order to maximize power efficiency. A switch mode amplifier theoretically can achieve 100% efficiency because the current and voltage are opposite in phase. Since the development of switch-mode amplifiers require use of passive components, 3D integration should facilitate their application. We plan to operate at the highest frequencies allowable by our active devices. This will facilitate minimizing the size of the passive components, especially antenna, inductor and capacitor structures.
- **Passive Structures (Inductors, Capacitors and Antennas):** Passive components will present a challenge with respect size constraints. Inductors will be in the nanohenry range, while capacitors will be in the tenth picofarad range. 3D integration should be a great asset in fabricating these passive components because the added volume will greatly facilitate fabrication of these elements while at the same time minimizing their resistive parasitic elements which reduce the quality factor. Specific antenna designs are described in Section 5 of the paper, however, the added layers of 3D integration should allow for more efficient antennas and more optimal radiation patterns. We also plan to expand our backgournd on carbon nanotubes[7], examine their use as inductors. Our background investigation

indicates that CNTs may exhibit very large kinetic inductance, and require a small volume, thereby saving considerable IC real estate.

3. Communication Protocols, Networking and Algorithm Development:Energy-Efficient Algorithms for Integrated Networking and Fusion

In this project we will develop implement and optimize energy-efficient algorithms for data processing, communication, and networking, that allow computation of fusion tasks in large networks of randomly distributed sensors. Although, our sensor prototypes will be initially limited to sensing temperature information, the fusion algorithms we will develop and implement, can be employed, in principle, to perform fusion tasks that arise in a broad range of applications, including weather forecasting, detection and monitoring of chemical and biological agents, and surveillance tasks, such as target detection, localization and tracking. In particular, the fusion objectives of interest in this project comprise:

- distributed algorithms for computing global statistics of the measurements collected across the network, such as maxima and averages;
- methods for event detection, classification, and localization, such as detection of measurements exceeding a threshold, and localization and tracking of threshold exceeding events;
- algorithms for localized post-detection processing and tracking, such as computation of localized statistics and localized event tracking.

Several important constraints invariably arise in designing resource-efficient physical- and network-layer communication algorithms for performing fusion tasks over large networks of randomly distributed sensors. Limitations in sensor battery power and computing resources place tight constraints in the rate and form of information that can be exchanged among sensors and the type of on-sensor processing that can be performed. Equally important, changes in the network topology and size, due to node mobility node failures and battery outage, make global knowledge of the changing network topology impractical. As a result, ad-hoc networking and decentralized approaches to fusion are becoming increasingly preferable over their hierarchical centralized counterparts for such large sensor networks.

3.1 Ad-hoc Networking and Fusion

The ad-hoc networking algorithms we propose to develop will provide short-range communication nodes the capability to communicate and fuse information across large networks. In particular, based on its transmit power, each dust particle will be able to establish direct bidirectional communication with a small subset of neighboring nodes (e.g., nodes within an 1-10 m radius) and will thus be aware of (the IDs of) a small subset of communicating nodes.

Performing local processing and fusion prior to relaying data across a smart-dust network can provide substantial data compression savings and as a result, savings in both power and bandwidth usage. A promising approach to performing distributed fusion involves viewing each node in the network as a local fusion host and developing integrated routing/fusion algorithms such that the

fusion objectives are eventually made available to every node in the network. In its most basic form, ad-hoc networking amounts to localized routing algorithms allowing data from any single node to reach a subset of designated receiver (or in principle, all) nodes for future interrogation. More generally, the integrated networking/fusion algorithms we plan to implement will allow for event detection, localization, and tracking and are described in the following sections.

Distributed Computation of Global Maxima and Averages

The localized processing and fusion algorithms we propose to implement and optimize will allow distributed computation of global and local statistics in the measurements, such as global maxima and weighted averages of the sensor data. These algorithms are locally constructed at each node, i.e., they only require that each node is aware of the subset of neighboring nodes with which it can establish bidirectional communication. They are also inherently distributed scalable and fault tolerant, and can allow substantial savings in communication bandwidth and power usage through local data processing and fusion prior to retransmitting information over the network.

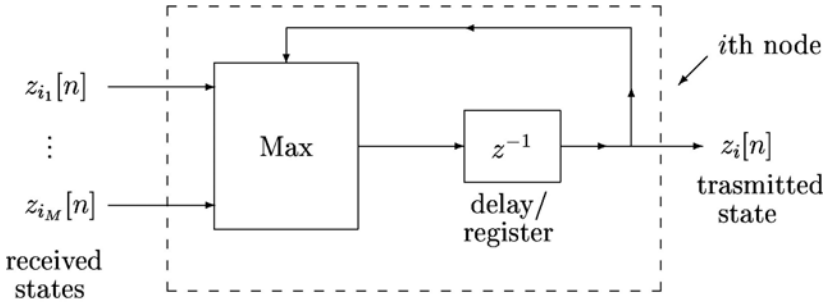

Figure 3.1: Local processing at the *i*th node yielding the global maximum

The simplest task that illustrates the benefits of local fusion involves the computation of the highest temperature (global maximum) among the sensed temperatures. The local fusion rule at a given node is shown in Fig. 3.1, and sets its new local-maximum estimate as the maximum of all the estimates transmitted by neighboring nodes and the node's own local estimate. Simulations, assuming uniformly distributed sensor data, in an ad-hoc network of 10,000 randomly distributed nodes with approx. 10 neighbors/node, and whereby each node broadcasts only improvements in its local estimate, reveal that on average less than 8 transmissions/node are required until every node in the network obtains the global maximum. Furthermore, the average number of transmissions per node (and thus average energy usage) required for global convergence can be significantly reduced at the expense of delays, if nodes broadcast improved local maxima only a fraction of the time.

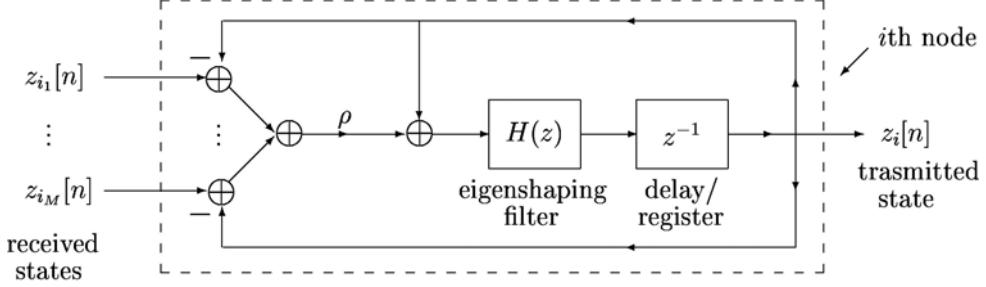

Figure 3.2: Local processing at the *i*th node yielding the global average

Distributed implementations of the majority of fusion rules (including linear fusion rules) are not, in general, as straightforward, as the global network topology is not locally available at individual nodes. For instance, local averaging of the available states at each node, yields local sequences of estimates that in general do not converge to the global average, but rather to a weighted average that depends on the (unknown) global network topology. In [Scherb03a] we developed methods for local processing and fusion that allow distributed computation of various computations, including weighted averages of functions of the node data, over ad-hoc networks with arbitrary network topology. These low-complexity local-processing rules are shown in Fig. 3.2, and can be locally constructed and optimized to provide at each node fast-converging approximations to the desired global computation. Fig. 3.3 shows the convergence rates to the global average for a typical network of 400 nodes for various such fusion rules. As the figure shows, the rule corresponding to the steepest curve requires less than 20 local communication and fusion steps to obtain an estimate of the average with 0.1% (–30dB) relative squared-error distortion. We plan to leverage our experience in this area and in developing optimized DSP implementations [ssb[2]], to implement, test, and optimize such fusion algorithms over the smart dust networks that we will deploy, and investigate the trade-offs between processing/communication power, and fusion quality.

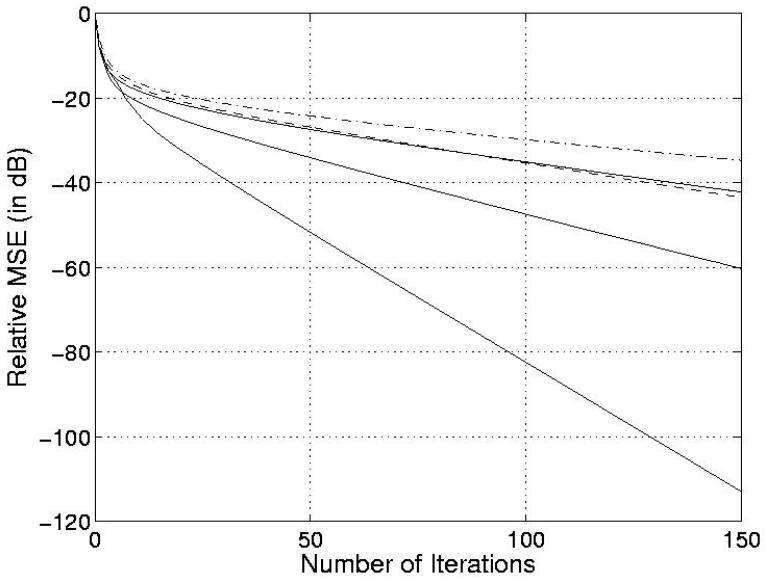

Figure 3.3: Convergence rates of local processing rules for computing averages

Sensor-Location Aided Event Detection and Localization

The algorithms in [Scherb03a] can be employed to perform more complex fusion tasks, which can be decomposed into sets of weighted average computations of local function of the sensor data. These include a number of surveillance tasks, such as target detection by majority voting, and target localization. For instance, in [Pham04], we employ these algorithms for distributed target localization via an ad-hoc network of acoustic sensors, whereby each node knows its own location and has measurements that provide relative-range information between the node and the target. In the early stages of this project involving networks with node placed in known predetermined locations, we plan to leverage of prior experience with distributed source localization from range-measurement information [Pham04], to develop implement and optimize distributed triangulation algorithms for sensor-location aided event localization and tracking.

Sensor Location Estimation via Networking

The minimum number of hops required to communicate between any two nodes in the network provides a coarse estimate of the relative distance between the two nodes. By combining such relative range information between a node and several properly distributed nodes with known locations it is possible to obtain an estimate of the position of the node. This approach to position estimation is attractive in dense networks, since the minimum number of hops required for communication between any two nodes can be obtained via localized broadcasting and routing. In its simplest form, it involves one of the nodes serving as the source that broadcasts a counter-type message, which will provide the minimum number of hops information when it reaches its destination. In particular, the message is sequentially propagated throughout the network by receiving nodes by locally updating (incrementing by 1) the minimum-number-of-hops counter message and rebroadcasting, until the message reaches the destination node. Fig. 4 depicts the localization performance of this method for a network of 5,000 nodes. The node in the center of the blue circle has detected an event, while the 8 circles in the periphery denote reference nodes with (locally) known positions. A single broadcasted message by the detecting node allows each of the reference nodes to obtain range information from the detecting node that can then be used to estimate of its location via distributed source localization [Pham04]. The blue circle in the figure denotes the localization error range of this method, which, for this network, is of the same order as the range of a single-hop transmission. We plan to investigate and optimize such coarse sensor location estimators that are based on the minimum number of hops and additional locally available

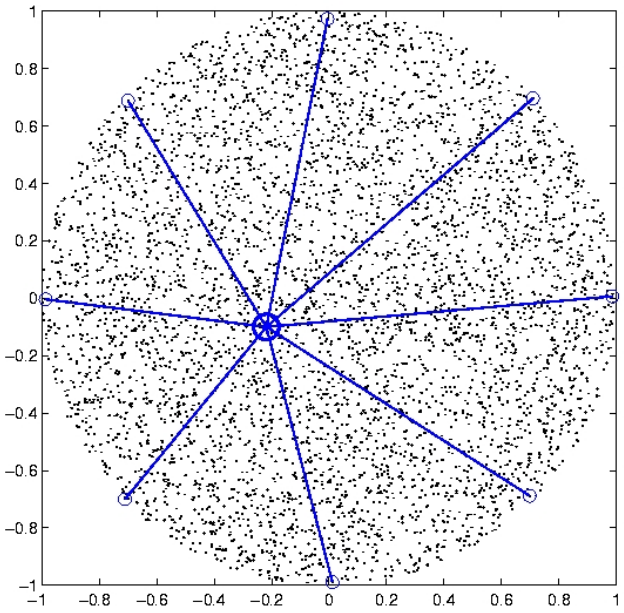

Figure 3.4: Ad-hoc network aided coarse node localization in dense smart dust network

information (e.g., received signal strength in single-hop transmissions), and analyze their localization performance characteristics as a function of the network node-density and size, the broadcasting range, and the number and location of reference nodes used for triangulation.

Localized Averages and Localized Tracking

At later stages of the project, we plan to focus on methods for performing localized fusion tasks, whereby only a (potentially time-varying) subset of nodes close to the event of interest have relevant measurements for fusion. Our approach consists of methods for distributed localized detection and localization, followed by localized tracking via a properly adapted subset of detecting nodes. In this context we plan to leverage our prior experience [Scherb03b] involving methods for distributed computation of fusion tasks based on data from only a subset of the nodes in the network, by exploiting the whole network topology. These algorithms are very attractive in the context of tracking and surveillance, as only a subset of the whole set of nodes in the network (those close to the source) have information-bearing measurements, yet, it is in general advantageous (both in terms of power usage and computation delays) to employ the whole network for the computation [Scherb03b].

3.2 Resource-efficient Point-to-Point Communication

The point-to-point communication algorithms will allow measurements and data to be reliably communicated over a given link while making efficient use of energy (and bandwidth) resources. Transmitter and receiver block diagrams are shown in Figs. 3.5 and 3.6 and comprise the following subsystems:

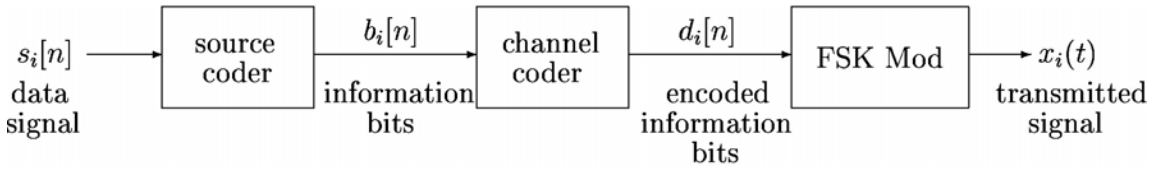

Figure 3.5: Transmitter for communicating a sequence of data or measurements

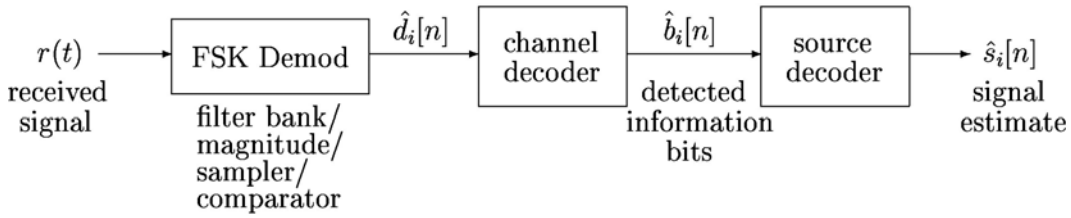

Figure 3.6: Receiver comprising FSK demodulation followed by channel and source decoding

- **Source coding (transmitter) and source decoding (receiver) algorithms**: The source coder compresses the data/measurements that are to be communicated into a sequence of information bits that is instead communicated over the channel. These algorithms perform signal compression in amplitude and time in order to eliminate redundancy from the data to be communicated. We will leverage our prior experience with low-complexity source encoding algorithms [Papa01]-[Abda01], in order to implement energy-efficient systems for encoding and communicating the analog measurements in digital form over the wireless channel.
- **Channel coding (transmitter) and decoding (receiver) algorithms**: The error-correcting encoder inserts controlled redundancy into the information bit stream in order to communicate it over the wireless channel to a receiving node at desired fidelity levels. The error-correcting codes we plan to employ are known as rate-compatible punctured convolutional (RCPC) codes [Hage88]. Their advantages are their low encoding and decoding complexity, and the fact that by design they can readily provide different levels of error protection to different information bits, thereby exploiting the fact that the signal distortion due to an erroneous detection at the receiver strongly depends on the information carried by the associated information bit.
- **Modulation algorithms**: As described in the transceiver design section, the channel-coded information bits are relayed from the microcontroller to the RF circuitry and are communicated over the wireless channel via frequency-shift keying (FSK), at a preassigned carrier frequency. The FSK carrier frequency range will be in the order of 10GHz in our original designs, and may reach 50 GHz at later stages.

3.3 Multiple-Access Protocols

Taking into account the inherent limitations in the RF hardware and the need for ad-hoc networking, we will employ TDMA-FDMA protocols, according to which user transmissions are multiplexed in both time and frequency. According to the protocol, each dust particle periodically communicates with closely located nodes only at preassigned time slots via FSK at a preassigned carrier frequency, and can receive on a single carrier frequency during any non-transmitting time slot. The use of FDMA (as opposed to CDMA) is motivated by the presence of sufficient bandwidth, allowing orthogonal (multiuser-interference free) communication among neighboring nodes. The need for multiple-access in time (TDMA) is motivated by limitations in the transceiver circuitry (dictating that a node can only receive at a single carrier frequency at any given time) and allows each node to receive transmissions from multiple neighboring nodes during any given time frame.

Protocol Assignment

Due to the nature of ad-hoc networking, there is a need for methods for spatial reuse and reallocation of time-frequency slots so as to avoid interference between transmitting nodes. In the initial network prototype, involving fixed sensor placement, the time-frequency slots will be preallocated (and will remain fixed) so as to avoid interference. In the context of random node placement, we will leverage our prior experience with evolutionary< algorithms and directed graph theory [ssb[1]] to generate time-frequency slot assignments so as to avoid interference. Congestions that may inevitably eventually arise due to the time-varying aspects of the network topology can be

alleviated by having sensors with interfering transmissions locally readjust their time-frequency slots in order to avoid interference. As an example, in a network of randomly placed sensors, whereby each node communicates on average with 10 other neighboring nodes within its range (e.g., an area around the node of 10m radius), and a frequency reuse factor of 3 (corresponding to no interference from all nodes outside an area of 60m radius), a 10×20 FDMA-TDMA slot allocation protocol more than suffices to obtain non-interfering transmissions. We remark that, as bandwidth efficiency is not as important as efficient battery-power use, congestions can be always alleviated at the expense of network computation rates (and thus delays in fusion), by increasing the number of TDMA slots in a frame. This approach is also attractive for generating an initial operational network configuration in networks of randomly dispersed dust particles. Indeed, using an FDMA-TDMA allocation table with a large enough number of time-access slots allows the nodes to self-organize and obtain an initial locally constructed multiple-access slot allocation, which can be used to obtain the network topology. The techniques described in Sec. xxx [ssb[1]] can then exploit such network topology information to obtain new more efficient multiple access allocations, involving FDMA-TDMA tables with shorter time frames, thereby obtaining networks with optimized computation rates.

4. Digital System Design: Microcontrolling and Computational Components for Interfacing with RF Transceiver, Managing Network Protocols and Data Analysis

The digital subsystem is responsible for control and data processing of a smart dust node, as illustrated in Figure 4.1. This requires close monitoring and control of transceiver circuitry. For example, the digital subsystem will set the counter value of the PLL based on the next channel on which reception will be monitored, carry out reception on this frequency for a pre-specified length of time, extract and examine the quality of the received data, and extend the duration of reception as necessary until the data on the channel is adequately received. The digital controller will repeat this sequence of steps for each active channel during each period of data collection. During transmission, the digital controller will determine what, if any, data needs to be transmitted; carry out the low complexity source and channel coding algorithms described in Section xx [reference to Babis's section on "Resource-efficient Point-to-Point Communication"]; and activate the transmitter with transmission frequency and duration based on the TDMA-FDMA protocols of the network.

Figure 4.1. An illustration of the structure of core digital system functionality.

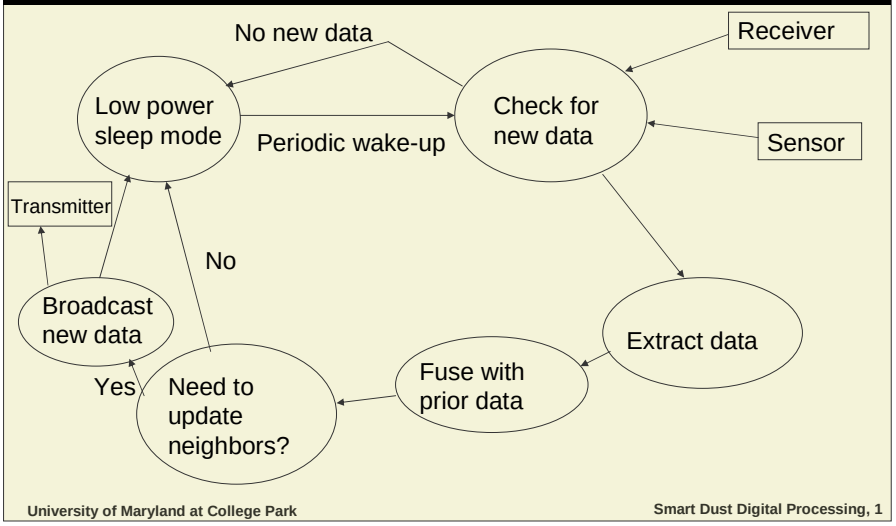

Digital controller design

In our initial prototype, the digital processing engine will be a low power microcontroller, which will integrate lightweight operating system functionality for monitoring and adaptation of node behavior with the data processing tasks to be performed at each node. Use of a low power programmable processor will enable exploration of more complex algorithms for sensor fusion, experimentation with different algorithms and protocols, and run-time adaptation of node behavior, and will provide access to important energy consumption features such as processor shutdown modes and dynamic voltage scaling. Research issues in the design and implementation of the digital processing platform include efficient integration with the radio subsystem, and the mapping of the desired processing algorithms and protocols into energy efficient and highly memory constrained implementations. For this latter task, we will build on our extensive experience in developing transformations for efficient and compact embedded software implementations (e.g., see [2][3]). Our work will involve modeling of the desired smart dust processing in terms of formal models of computation, such as dataflow graphs and finite state machines, and graph theoretic analysis of these models to map them into the constraints imposed by the restricted processing architecture.

One specific family of candidates for use as the microcontroller in our initial prototype is the Texas Instruments MSP430 family of "ultra-low-power" microcontrollers. Components in this family offer high programmability with active-mode power consumption below 500 microwatts when operating near the lower limit of their voltage range. In addition, these components provide integrated A/D to and D/A converters, and processor shutdown modes in which power consumption is reduced to amounts as low as a few tenths of a microwatt during idle periods. Experimentation with microcontrollers such as these in our initial prototype will focus our digital design techniques on the performance and memory constraints that are critical to achieving the desired levels of miniaturization and low power consumption. It will also provide baseline implementations that provide reference points for our future efforts at size/power reduction and 3D integration.

In parallel with development of our initial prototype, we will explore three directions for achieving the high miniaturization and 3-D integration capabilities that will be targeted by our subsequent prototype designs. After a preliminary exploration in which the potential costs and benefits of all three directions are carefully estimated, we will select one of these directions to focus on for the development of our second prototype. Our first direction of exploration will involve developing our own streamlined implementation of the off-the-shelf microcontroller used in the initial prototype. The objective of this implementation will be to maintain at least some of the programmability of the original microcontroller component, but achieve 3D integration capability and significant reductions in power consumption by designing and implementing the integrated circuit ourselves. Mechanisms for power reduction in the implementation will center around removal of unnecessary or underutilized processor features (e.g., instructions, memory capacity, and addressing modes), and careful application of hardware/software co-design [4] to migrate critical functionality into custom hardware implementation. By designing for compatibility with the original microcontroller specification (or a carefully-selected subset of it),

we will be able to leverage the compilers, assemblers, and other software development tools that accompany it.

Our second direction of exploration for digital system design will be replacing the microcontroller component with an application-specific integrated circuit (ASIC) that employs dedicated circuitry for all functions. The ASIC will be carefully designed to support the complete digital system functionality for a specific smart dust application, with a specific set of communication protocols. Unlike a microcontroller-based design, this approach will not provide flexibility in accommodating changes in application or communication functionality (without re-fabrication). However, due to its specialization, we can expect significant streamlining of size and power consumption. Extensive use of low power digital system and logic design techniques will be employed. Such techniques include clock gating, controlling input signals to functional modules so that they change only when the module results are needed, arithmetic strength reduction (e.g., replacing multipliers with shift or shift-and-add modules), and reduction of switching activity at the integrated circuit pins. A modular architectural and hardware description language (HDL) design style will be employed so that supported features (e.g., data processing functions or communication protocols) can be easily reconfigured from one integrated circuit implementation to the next, thereby mitigating some of the decreased flexibility as compared to the microcontroller-based design.

Third, we will consider the possibility of licensing the commercial intellectual property (IP) core that underlies the original microcontroller, and using that core for 3D integration in our second-generation prototype. This approach will allow us to leverage state of the art, commercial low power microcontroller technology, and allow us to focus more of our digital system design effort on system-level configuration optimization, as described in more detail below. This possibility is highly speculative at this point, however, because the manufacturer of the microcontroller may not permit such licensing (especially for a low volume research prototype), and because, even if licensing is permitted, the cost is anticipated to be extremely high.

Network configuration

Several issues must also be addressed in configuring the smart dust network. These include setting up handshaking protocols, source coding and decoding, integrating with the transceiver to establish PLL timing, establishing error-correction coding and low-complexity decoding, assigning transmission power, and assigning processing tasks to network nodes. The parameters underlying these setup issues form a complex design space, whose efficient exploration will lead to significant power savings.

For example, streamlining the energy efficiency of processing in the smart dust network will require careful balancing of computation and communication across the network nodes. For this purpose, we will explore techniques for strategically assigning sensing and processing tasks to network nodes so that effective trade-offs are achieved across computational accuracy, processing energy consumption, and communication energy consumption. This will involve developing quantitative models of communication and processing energy consumption throughout the network, and modeling the overall network's processing algorithms as task graphs, where graph vertices represent sensing/processing tasks and connections between vertices represent data dependencies (communication requirements) between tasks. We will explore algorithms to

systematically embed these task graphs into the sensor network based on the overall energy consumption predicted by the energy consumption models. This embedding will determine which task graph computations are assigned to which network nodes, and which subset of network nodes is kept idle (turned off).

Similarly, for the assumed TDMA/FDMA communications protocol, the assignment of time/frequency slots to network nodes can be performed strategically to optimize the computation rate of the network, and this assignment may be adapted periodically based on the network state and channel conditions. This problem is discussed further in Section xx [reference to Babis's section on multiple access protocols]). To address such system-level optimization issues, we will explore the use of evolutionary algorithms, which provide a probabilistic search methodology for complex, non-linear optimization problems. In this work, we will leverage our previous experience in architecting evolutionary algorithms for embedded software optimization [1]. An illustration of the general framework of evolutionary algorithms is shown in Figure 2. Candidate solutions to an optimization problem are encoded as binary strings based on a "genetic representation" function. This "population" of binary strings is then adapted probabilistically through operations such as "recombination" (combining attributes from two existing binary strings to form a new binary string) and "mutation" (making small perturbations to an existing binary string). As these adaptations occur, the modified population is filtered with a strong bias toward retaining those binary strings that correspond to the most effective candidate solutions. This process of probabilistic adaptation and filtering continues through numerous iterations, and the best candidate solution in the resulting final population is chosen as the output. When the key evolutionary algorithm building blocks, such as the genetic representation, recombination and mutation operations, and filtering mechanism, are designed well, evolutionary algorithms have been shown to be effective across a wide variety of complex optimization problems (e.g., see [5]).

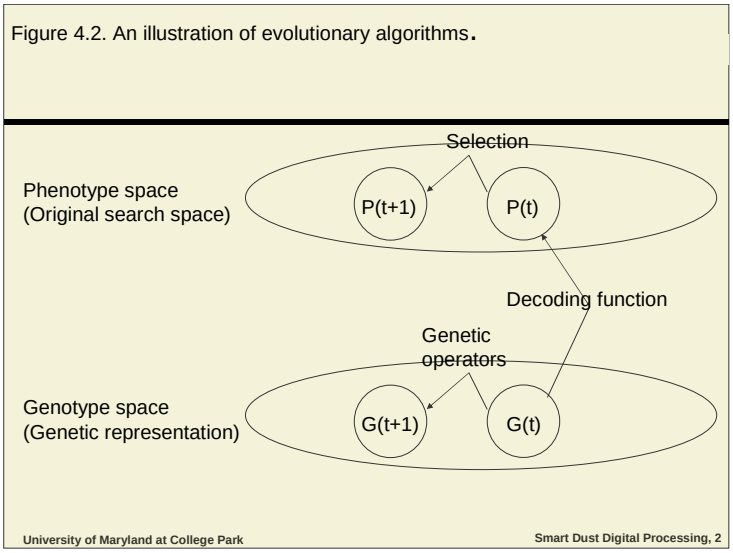

5. Milli and Micro ANTENNAS

Establishing an effective communication link between the dust particles require electrically and physically small antennas that should be capable of meeting the power requirements of the system. Providing that the antenna size profile (size profile is defined as the radius of the smallest sphere that completely encloses the antenna,) should be smaller than 5mm, we propose several antenna topologies that address two primary concerns: The antenna radiation resistance, which is a direct measure of the power gain of the antenna, and the directive gain, or radiation pattern. We assume that the dust particles will be distributed in a plane and therefore the radiation pattern of the identical antennas will be uniform in the polar angle and negligible radiation in the azimuth. Since

the dust particles are positioned randomly, maximizing the directive gain in a certain direction will not be useful. Initially, we will assume vertical (with respect to the ground) polarization.

The antennas proposed here are all mounted on a perfectly conducting platform. Antennas mounted on finite size planes provide alternative designs, however, as a proof of concept, we start with infinite perfectly conducting planes to minimize antenna simulation time. We assume that the frequency of operation is f=10GHz. If the frequency is increased, say to 50GHz, the antenna electrical and radiation characteristics will remain unchanged providing that all antenna dimensions are scaled uniformly.

Power calculations for Simple Dipole Antenna

Simple calculations can be made to determine the minimum power requirements for a communication link between two antennas separated by a distance *r*. For simplicity, we will assume that the two antennas are positioned such that the transfer of power between them is maximum. This includes polarization and directivity (maximum directive gain). Let P_T be the total power transmitted, P_R be the total received power, then according to Friis Transmission Formula

$$P_R = P_T \frac{A_{en}^T A_{en}^R}{r^2 \lambda^2}$$

where r is the distance between the two antennas, A_{em}^T and A_{em}^R are the maximum aperture area of the transmitting and receiving antennas respectively. The aperture area is expressed as

$$A_{en}^T = \frac{V_{rms}^2}{4 R_r S_{av}}$$

Where V_{rms} is the voltage at the input terminals of the antenna, S_{av} is the average power density, and R_r is the radiation resistance of the antenna. What is important to note here is that the aperture area is inversely proportional to the radiation resistance.

In light of the severe size constraints placed on the antenna, an electrically small antenna (ESA) must be used. An ESA typically has dimensions $kr < 1$, where *k* is the wavelength. There are several antenna shapes types that can be used. The simplest is the dipole antenna (or the monopole antenna).

For a short (non-ideal) dipole, the radiation resistance, R_r is given by

$$R_r = 20\pi^2 \left(\frac{l}{\lambda} \right)^2 ,$$

and the maximum effective aperture is given by $0.48\lambda^2$. Substituting in Friis formula, and assuming identical transmit and receive antennas placed to maximize power coupling, we have

$$P_T = P_R \frac{r^2}{0.23\lambda^2}$$

As an example, under optimal matching and polarization conditions, and assuming that the received power should be greater or equal to -100 dBm (0.1 pico Watts), then according to the formula above, if an electrically short dipole antenna is used, we can have the following minimum power requirements:
For antenna separation of 1m, $(P_R)_{\min}$ = 0.482 nano Watts
For antenna separation of 10m, $(P_R)_{\min}$ = 48.2 nano Watts

In light of the fact that the aperture area is inversely proportional to the radiation resistance, it is expected that any increase in the radiation resistance (as offered by the alternate designs discussed below) would lead to lower power transmission requirements to achieve -100dBm.

It should be kept in mind, however, that the short dipole antenna has low radiation resistance which needs to be matched to the feed line, and a capacitance which needs to be tuned either by an inductor or a transmission line stub that is connected in parallel with the feed. Notice that the transmission line stub need not be stretched but can be meandered to reduce profile.

Inductively-Coupled Small Antennas

Aside from the dipole antenna, the most desirable type of antenna is the patch antenna which offers low cost, high reliability, and minimal physical profile. The microstrip patch antenna becomes resonant when one of its dimensions is approximately equal to one half the wavelength. For an operating frequency of 10GHz, the wavelength is 3cm which requires the smallest dimension of the patch antenna to be approximately 1.5cm. Since the microstrip antenna is inherently very inefficient radiator, the antenna cannot be used effectively when its length is decreased below the half wavelength mark.

Different antenna configurations can be considered, however, all of these configurations are electrically small in the sense that $kr < 1$. For a dipole operating at 10 GHz, and restricting the height of the antenna and radius to 5mm and 0.5mm respectively, the input impedance of the antenna is approximately

$$Z_{in} = R_{in} + jX_{in} = 30 - j$$

The short dipole antenna presents the simplest possible design configuration. More efficient, yet non-intuitive designs can be adopted that are based on using genetic type optimization algorithms that search for specific optimal antenna characteristics under specified constraints. Two antennas that will be discussed here have been obtained based on previous designs that were optimized.

The first antenna considered will be referred to as Antenna A, shown in Fig. 5.1. This antenna consists of a shorted feed line inductively coupled to a meander antenna. The inductive feed provides counter balances the inherent capacitive reactance of small antennas and the meander geometry extends the length of the antenna while maintaining small overall electric volume. The antenna radiation pattern is shown in Fig. 5.2. (Ansoft HFSS was used for all the simulation results presented here.) Note that the power is skewed towards one direction due primarily to the asymmetric feed geometry. This imbalance in the power distribution can be corrected by providing

a symmetric inductive feed topology. The antenna return loss is shown in Fig. 5.3, and the input impedance, presented in Fig. 5.4, shows that this antenna does need tuning and its radiation resistance is very close to 50 Ohms, thus providing excellent power match potential.

The second antenna, antenna B, shown in Fig. 5.5, builds upon the same concept of inductive coupling discussed above. However, unlike antenna A, the shorted element in antenna B is pushed further from the ground plane. This change provides a better directionality as shown in Fig. 5.6.

Since antennas A and B provide a much improved radiation resistance than the dipole antenna while eliminating the need for tuning which can reduce the overall efficiency of the antenna-feed system.

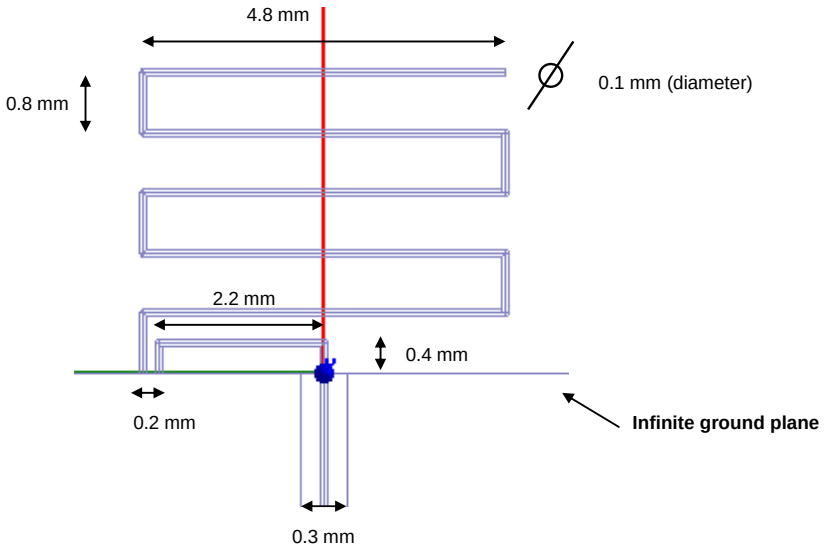

Fig. 5.1. Antenna A.

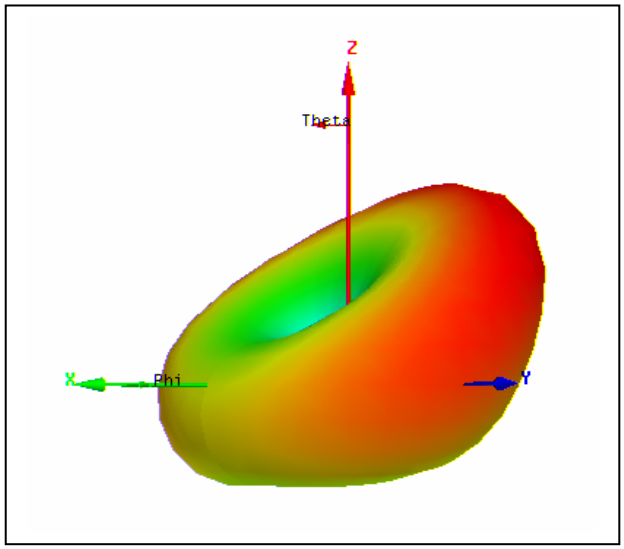

Figure 5.2: Radiation pattern of antenna A

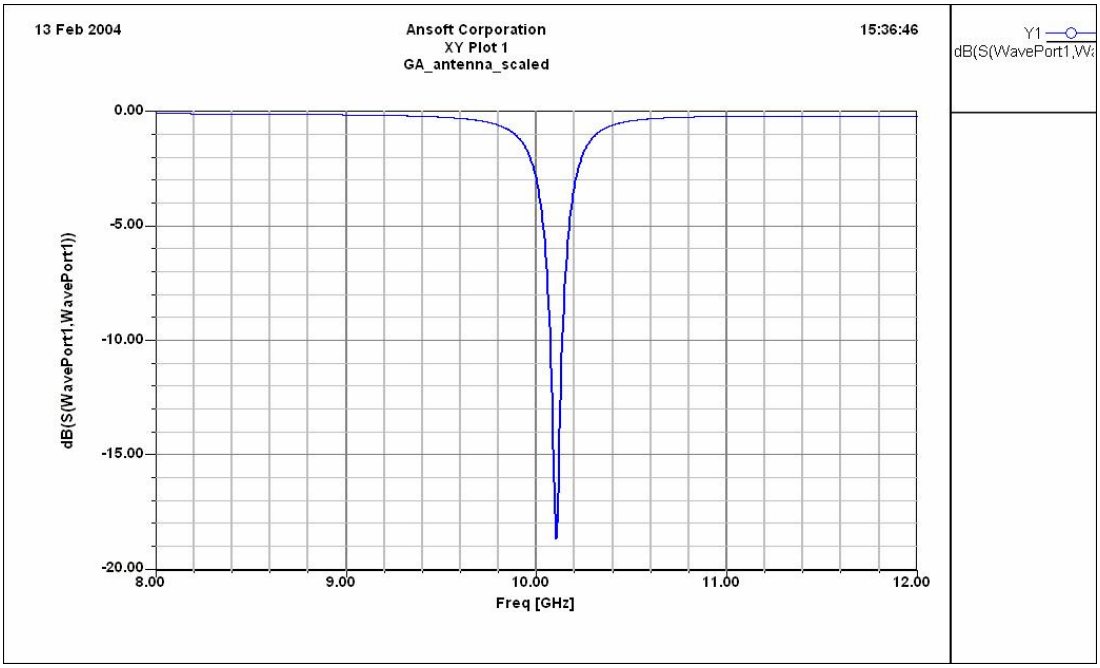

Fig. 5.3. Return loss (S_{11}) of antenna A.

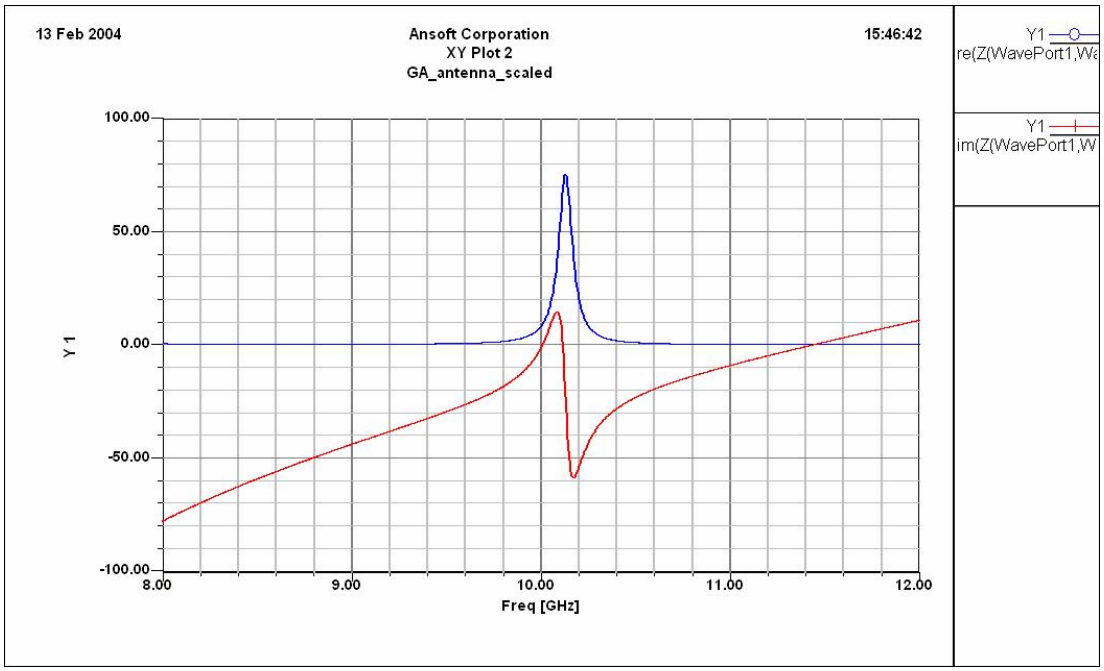

Fig. 5.4. Input impedance (resistance: Blue, Reactance: Red) of antenna A.

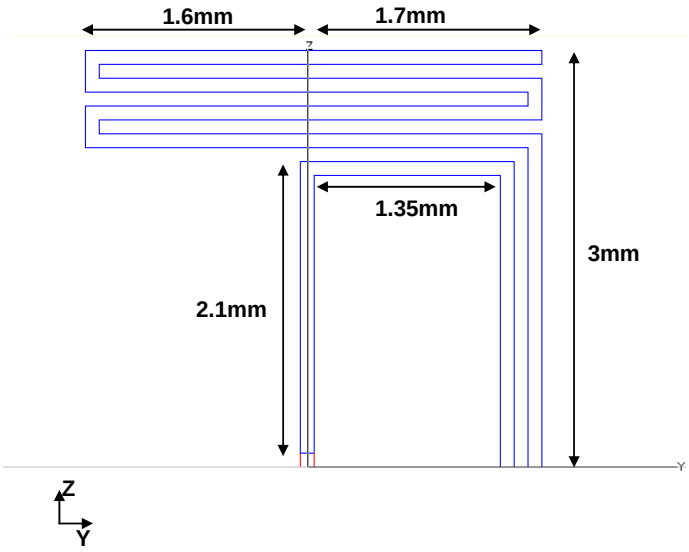

Fig. 5.5. Antenna B.

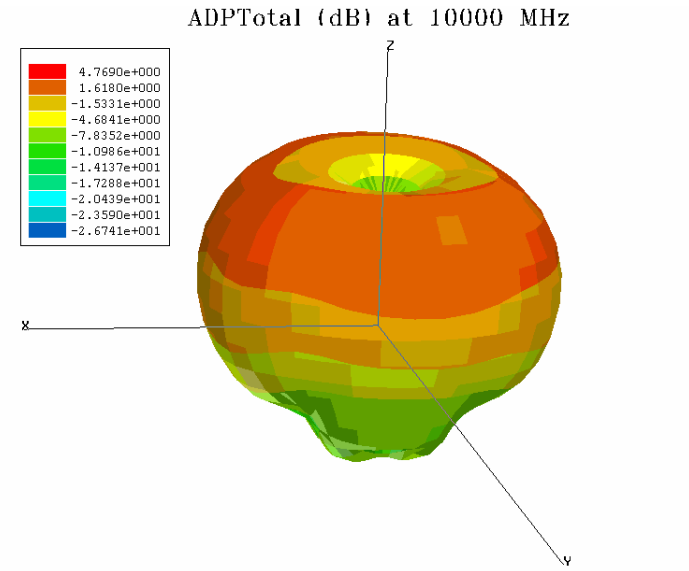

Fig. 5.6. Radiation pattern of antenna B.

Future Work

We emphasize here that the antennas discussed above were not optimized for directivity, radiation resistance and bandwidth. We propose to use genetic algorithms coupled with the Alternate-Direct Implicit Finite-Difference Time-Domain (ADI-FDTD) method to search for more effective topologies. The ADI-FDTD is an excellent algorithm for simulating structures with wide variance in resolution and can predict effectively Ohmic losses in the conductors and in the substrate. These losses are expected to increase with the frequency, leading to changes in the overall efficiency of the antenna.

We plan to plan to investigate volume-based curve-filling antennas and the hemispherical monopole above perfectly conducting ground plane. A detailed listing of the objectives is given below.

Objectives

Development of a full-wave three-dimensional ADI-FDTD algorithm that comprehensively simulate the antennas and the platform. The algorithm will be capable to address structures with wide variation in geometrical resolution while maintaining very efficient run time.
Adaptation of our genetic optimization algorithm to the antenna problem.
Application of genetic-based and non-genetic-based optimization algorithm to the antenna problem.
Feed design for maximum power transfer
Matching network design for maximum power transfer
Development of space-filling antennas
Incorporating electromagnetic band gap material to maximize gain for antennas operating in the 30-50GHz region
Antenna fabrication using micromachining
Antenna testing for power gain and radiation characteristics

6. Power: MEMS, Batteries, Capacitor Structures and Solar Cells

7. Analog Sensing

Sensors will be at the input of each dust particle. We plan to use temperature and light sensors in our initial prototypes. However, other sensors should be readily integrable, including MEMS and bio-sensors. We expect the output of the sensor to be an analog signal which we will then convert to a digital equivalent for operation with the microcontroller and the transceiver.

An initial prototype will be implemented with off-the shelf components, that measures temperature within a geographical. Of course, powering the dust is also an issue. Our first prototype will be internally powered with batteries. Later, we plan to use external power obtained from light or an EM radiation source. Once we demonstrate network viability, we will develop our own components using CMOS VLSI. We will design our analog circuits with Cadence, and digital systems with Verilog. Layout files (GDSII) will then be shipped to MOSIS for fabrication using state of the art process technology. We also plan to utilize our existing program in 3D integration. This will allow for unique design and implementation, especially for passive components such as antennas and inductors, which tend to be larger. Having 3D integration capability is a unique attribute which we plan to exploit to facilitate robust smart dust development.

4. Milestones

6-month milestones:

- Simulation and implementation of smart dust network that computes highest sensed measurement (global maximum) and provides elementary operations, such as node-mode control (turn on and sleep mode) and node querying. Testing of hardware components, communication algorithms and protocols over a small network with predetermined node placement. The algorithms that will be implemented include: measurement quantization and compression, low-memory RCPC coding for unequal error protection, and FSK-based modulation with preset TDMA/FDMA slot assignments. Computer-simulation based testing and optimization of these computations on larger networks involving random node placement.

- Simulation and preliminary implementation of digital processing algorithms using microcontroller development platform.

12-month milestones:

- Computations of weighted measurement averages and optimization of the source-coding and channel-coding algorithms for energy-efficient communication and local fusion. Power-optimized designs with integrated digital/analog processing communication and networking for global maxima and averages. Study of the trade offs between energy efficiency sensor density and computation quality.
- Complete implementation, testing, and integration of digital processing engine and embedded software for first prototype.
- Optimize design, fabricate and test basic antenna structures. Upgrade ADI-FDTD code to model antenna details. Link antenna structures with RF circuits.

24-month milestones:

- Implementation of sensor-location aided localized routing algorithms for energy-efficient event-detection and event localization in networks with predetermined node placement. Coarse event-range estimation in networks of randomly placed nodes based on the minimum number of hops from the event location.
- System architecture design and simulation for second-generation (3D integrated) prototype of digital processing engine. Evolutionary algorithm formulation of optimizing network configuration.
- Begin looking at improved antenna structures based on genetic algorithms. Scale antennas for higher frequencies.

36-month milestones:

- Dust particle position estimation based on routing length and distributed triangulation in large networks of randomly placed sensors. Event detection, and event localization in large networks

of randomly placed nodes. Proof-of-concept implementation and testing over reduced-scale network of randomly placed nodes.

- Complete low power implementation, testing and 3D integration of digital processing system for second-generation prototype. Complete implementation of evolutionary algorithm and integrate optimization results into system prototype.
- Fabricate and measure antennas that are based on genetic and non-genetic based algorithms. Scale antennas for higher frequency operation.

48 month milestones

- Extend digital subsystem with libraries and reconfiguration capabilities for design flexibility and multi-mode operation. Power/performance modeling of system prototype and integration of modeling into evolutionary optimization framework.

60-month milestones

- Integrated optimized networks for (RF circuitry and microcontroller) energy-efficient event detection, localization, and tracking in large networks of randomly distributed sensors. Distributed network sectoring, sector-based event detection, localization, and tracking based on a dynamically updated detecting subnetwork.

References (Neil and Omar)

1. Akturk, N. Goldsman and G. Metze, ``Faster CMOS Inverter Switching Obtained withChannel Engineered Asymmetrical Halo Implanted MOSFETs, *Solid State Electronics*, vol. 47, pp.~185--192, 2003.
2. X. Shao, N. Goldsman, O. M. Ramahi, P. N. Guzdar, A New Method for Simulation of On-Chip Interconnects and Substrate Currents with 3D Alternating-Direction-Implicit (ADI) Maxwell Equation Solver. *International Conference on Simulation of Semiconductor Processes and Devices,* pp. 315-318, 2003.
3. X. Shao, N. Goldsman, and O. M Ramahi, The Alternating-Direction Implicit Finite-Difference Time-Domain (ADI-FDTD) Method and its Application to Simulation of Scattering from Highly Conductive Material, *IEEE International Antennas and Propagation Symposium and USNC/CNC/URSI North American Radio Science Meeting: URSI, Digest*, p. 358, 2003.
4. Y. Bai, Z. Dilli, N. Goldsman and G. Metze, Frequency-Dependent Modeling of On-Chip Inductors on Lossy Substrate, *International Semiconductor Device Research Symposium*, pp. *292-293,* 2003.
5. X. Shao, N. Goldsman, O. M. Ramahi,and P. N. Guzdar, Modeling RF Effects in Integrated Circuits with a New 3D Alternating-Direction-Implicit Maxwell Equation Solver, *International Semiconductor Device Research Symposium*, pp. 532-533, 2003.
6. Akturk, L. Parker, N. Goldsman, G. Metze, Mixed Mode Simulation of Non-Isothermal Quantum Device Operation and Full-Chip Heating, *International Semiconductor Device Research Symposium*, pp. *508-509,* 2003.

7. G. Pennington and N. Goldsman, Semiclassical transport and phonon scattering of electrons in semiconducting carbon nanotubes, *Phys. Rev. B*, 68 (4): Art. No. pp. 045426-1 -- 11, 2003.

(Babis)

[Papa01] H. C. Papadopoulos, G. W. Wornell, and A. V. Oppenheim, "Sequential signal encoding from noisy measurements using quantizers with dynamic bias control," IEEE Trans. Inform. Theory, vol. 47, no. 3, pp. 978--1002, March 2001.

[Abda01] M. M. Abdallah and H. C. Papadopoulos, "Sequential signal encoding and estimation for distributed sensor networks," in Proc. IEEE Int. Conf. Acoust. Speech, Signal Processing, pp. 2577--2580, vol. 4, May 2001.

[Scherb03a] D. S. Scherber and H. C. Papadopoulos, "Locally constructed algorithms for distributed computations in ad-hoc networks," in Inform. Processing Sensor Networks, April 2003.

[Pham04] T. Pham, D. S. Scherber, and H. C. Papadopoulos, "Distributed source localization algorithms for acoustic ad-hoc sensor networks," submitted to 2004 IEEE Sensor Array and Multichannel Signal Processing Workshop.

[Scherb03b] D. S. Scherber and H. C. Papadopoulos, "Distributed computation of averages over ad-hoc networks," submitted to IEEE J. Select. Areas Commun.

[Hage88] J. Hagenauer, "Rate-compatible punctured convolutional codes (RCPC codes) and their applications," IEEE Trans. Commun., vol. 36, no. 4, April 1988.

[ssb[1]] N. K. Bambha, S. S. Bhattacharyya, J. Teich, and E. Zeitzler, "Systematic integration of parameterized local search in evolutionary algorithms," accepted for publication in IEEE Trans.

(Shuvra)

[1] N. K. Bambha, S. S. Bhattacharyya, J. Teich, and E. Zitzler. Systematic integration of parameterized local search in evolutionary algorithms. IEEE Transactions on Evolutionary Computation. To appear.

[2] S. S. Bhattacharyya, R. Leupers, and P. Marwedel. Software synthesis and code generation for DSP. IEEE Transactions on Circuits and Systems --- II: Analog and Digital Signal Processing, 47(9):849-875, September 2000.

[3] P. K. Murthy and S. S. Bhattacharyya. Shared buffer implementations of signal processing systems using lifetime analysis techniques. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 20(2):177-198, February 2001.

[4] S. S. Bhattacharyya. Hardware/software co-synthesis of DSP systems. In Y. H. Hu, editor, Programmable Digital Signal Processors: Architecture, Programming, and Applications, pages 333-378. Marcel Dekker, Inc., 2002.

[5] T. Back, U. Hammel, and H-P Schwefel. Evolutionary computation: Comments on the history and current state. IEEE Transactions on Evolutionary Computation, 1(1):3-17, April 1997.

Facilities

Anechoic chamber for electromagnetic testing (chamber size: 25 ft x 10 ft)
Low-to-High frequency source generator (45MHz to 27GHz)
Electromagnetic spectrum analyzer (45MHz to 26GHz
State-of-the-art Vector Network Analyzer (45MHz to 50GHz) with highly accurate 2.4mm test set
High-frequency precision spherical dipole for shielding and calibration
600MHz oscilloscope
Dielectric probe
Antennas and assorted high-precision probes
Access to micro-fabrication facilities at the University of Maryland at College Park