On the Feasibility of the UC Davis Metanet

Leo Szumel, John D. Owens

August 13, 2003

1 Problem Statement

Many of the design challenges for large scale wireless sensor networks are well known, among them:

Scalability: Large networks quickly saturate the radio spectrum available to them. Moreover, the algorithms to efficiently utilize spectrum are not well defined and, depending on the application, may be intractable. Part of scalability is the capability to easily add nodes to an existing network.

Geographic range: The scalability of a network is affected by the density of the nodes in the network. The tradeoff between processing and transmission cost changes as a function of network density.

Lack of infrastructure: Deploying sensor networks in real environments leads to the logistical challenge of communicating with that remote environment. Many remote environments do not contain "easy" long-haul communication options such as telephone lines.

Cost: In order to deploy and test large scale networks, large amounts of equipment need to be purchased. Because the technology is in its infancy, many of the parts are hand-built and thus quite expensive.

Longevity: One cannot replace batteries regularly in a large scale network. The only easily accessible infinite sources of energy, such as solar, place very real logistical limitations on network deployment.

Heterogeneity: Sensor networks, especially long-lived ones, should be able to incorporate multiple architectures. It is difficult to provide communication between heterogeneous nodes, in part due to implementation difficulty and in part due to efficiency concerns—distributed cooperative algorithms are difficult to design when nodes have heterogeneous capabilities and resources available.

Mobility: Mobile nodes require dynamic and flexible routing, which can be difficult to design in an energy efficient way.

These design challenges are well known; however, they are not necessarily well understood because they have largely been analyzed from a theoretical standpoint. Practical limitations, derived from actual experiments, have only begun to be explored.

We will propose the "U.C. Davis Meta Network," referred to herein as the DMN, as a method of simultaneously designing realistic answers to the aforementioned design challenges *and demonstrating through empirical studies the effectiveness of these answers.*

2 Introduction

The goal of this study is to examine the current state of the art of wireless sensor networks and then to apply that knowledge to analyze the feasibility of the Davis Meta Network proposal.

The existing design challenges, as well as the future expected progress, both weigh heavily on the feasibility of the DMN. Because the DMN effort will be relatively small scale in terms of manpower, we will rely heavily on the development of the larger WSN community; as a result, it is critical that we understand the limitations and future directions of that community.

In Section 10, we will apply this analysis to some proposed applications in order to determine their feasibility.

3 Approach

My approach to this feasibility study is to examine the existing technology and, as best as we can, estimate the future trends in WSN technologies. we will then examine the technological limits with respect to some example environmental monitoring scenarios.

We identify existing applications and try to map the results of those experiments to the example scenarios laid out in Section 10. For example, the experience of the Great Duck Island researchers has some, but not complete, overlap with the example scenarios. To date, no one has attempted such a large-scale project using WSN nodes of this cost scale, so projecting existing results into this larger application is critical.

4 Challenges

In this section we will focus on the primary challenges facing the DMN and how they will affect its feasibility.

4.1 Scaling

Scalability is a function of the network topology and the physical constraints of the nodes and infrastructure devices in the network. Gupta and Kumar point out that for random communication patterns in a constant-density network of n nodes, the per-node usable bandwidth is proportional to $1/\sqrt{n}$ [10]. Li et al. argue that locality of traffic is key to making networks scale [15], and we agree. The key result of this work is that any and all communication that can be focused locally, such as collaboration or data aggregation, must be done in order to obtain the maximum scalability of the network. If communication can be localized completely, the network can scale without bound.

Servetto showed that for estimation problems where the nodes' readings are correlated, the network can scale without bound if the per-node measurement rate decays at least as fast as the per-node capacity [20].

We have seen that local communication scales, and global communication can scale under certain idealized constraints. In real-world situations, the scalability will be limited but can certainly be maximized by using the results of Servetto's analysis and Li et al.'s advice recarding locality. The result is that real-world apps should attempt to utilize as much localized communication as possible, while also making use of correlations in data when deciding on measurement strategies.

See Section 10 for some examples of proposed applications and how they might be able to concentrate traffic locally in order to alleviate the scaling issue.

Our scalability will be primarily limited by the data throughput of the radio technology available to the node architecture that we pick. A second limit is the transmission density of the network, to be described in Section 4.1.3.

Longevity is a function of the routing scheme and the energy available in the network, and places another, more difficult to predict, constraint on scalability.

4.1.1 Effect of Distance on Throughput

To try to estimate the transmission range in which maximum throughput is achievable, we scanned the web for empirical data displaying throughput-vs-distance relationships. A study using MICA motes [27] found that 100% packet success rate could be achieved out to approximately 50 feet. This contrasts with a maximum achievable range of approximately 100 feet for the MICA radio [6].

4.1.2 Effect of Multiple Hops on Throughput

Even without data congestion, the throughput of a multihop chain is limited by interference when a single shared channel is used. This has been shown to result in a maximum throughput of $\frac{1}{4}$ of the maximum attainable single-hop throughput [15]. This effect becomes more difficult when there is interference in two dimensions (a net rather than a chain), but the specific results will depend on the application.

It is my opinion, though it has not been shown yet theoretically, that localized routing protocols, such as min-cost [26], will result in the least interference problems. While large-scale interference *can* be avoided via smart global scheduling, such scheduling is intractable [1]. When scheduling is done only locally, with TDMA for example, relatively complex scheduling algorithms can be used due to the small problem size.

The motes in the Great Duck Island experiment (see Section 7 only lasted for 2.5 months in the best case, when the expected lifetime was 6 or more months [23]. Perhaps more importantly, in that same experiment, the half-life of the system (time until one half of the network is unresponsive) was only about one month for a 4 month experiment. At the end of 4 months, all but 3 of the 30+ nodes were dead. This is almost surely a result of the fact that some modes have a higher cost-to-transmit than others and as a result were burdened unfairly.

The reasons behind these differing costs are generally:

Hop distance: The farther a node has to send data, the more expensive each transmitted bit becomes.

Traffic load: The more data that a node is responsible for forwarding, the more energy it must consume.

Friendly interference: Radios overhear the "chatter" of neighboring nodes, and just receiving data uses energy. In addition, packet collisions result in retransmission, which increases the per-good-bit cost of transmission. As a result, network capacity is often defined in terms of bit-meter/sec, indicating the number of bits that can be sent over a one meter distance each second, regardless of who is sending. This metric is used by Gupta and Kumar and accounts for multi-hop and interference effects.

Other interference: Radio interference causes packet loss. Additionally, multipath interference patterns can cause "nulls" in the received power at specific points in space [22]. These nulls cause either partial or complete degradation of the signal strength. The location of these nulls is dependent on the geometry of the environment and the frequency used for communication.

Simplistic multi hop algorithms that simply shuffle data towards the "sinks" in the network will result in the nodes nearest the sink dying first, because they are burdened with larger traffic load than other nodes in the network. This is undesirable because it would be better to have the network die off in a more random manner, so that the sensed area will degrade evenly.

Heinzelman et al. propose a clustering algorithm called LEACH [11] which attempts to spread energy usage evenly among the nodes by using a clustering scheme in which nodes take turns being "cluster heads" which are responsible for long-distance transmission of several nodes worth of data.

We have analyzed the behavior of LEACH and other clustering algorithms, and while they do improve the situation they do not fix it [24]. Moreover, the long-distance communications used by LEACH will cover a very large area of the network with interference, forcing a global scheduling scheme in order to avoid costly collisions. Global scheduling is known to be NP-hard, and thus is definitely not scalable.

4.1.3 Assumptions

For the purposes of this study, we assume that the per-node throughput will be on the order of tens of kbps, and that this throughput will only be achievable at roughly one-half of the specified transmission range. Further, we will assume that roughly 1/4 of the specified throughput is achievable in the coverage area.

We will use an unusual capacity metric of bits/sec/m^2, which we derive from the assumption that the nodes transmit with the same power and, hence, distance. We chose this metric to decouple the physical medium's capacity from a particular application's communication pattern. This metric does not account for multi-hop effects because doing so would make assumptions about the communication patterns of the network, a feature which is application-dependent. Gupta and Kumar, for instance, assumed that the nodes communicate with other nodes randomly and thus the mean number of hops was the mean message distance divided by the max transmit distance.

My metric simply provides the maximum number of bits that can be transmitted per unit area per unit time, regardless of the traffic patterns. Application multi-hop statistics can easily be used to find the reduced capacity; for example, if the mean number of hops is four, the capacity will be reduced to 1/4 of the maximum. The reason Gupta and Kumar's result shows the $1/\sqrt{N}$ relationship is that the mean message distance, and hence number of hops, is proportional to $\sqrt{N}$.

$$\begin{aligned} TransmissionArea &= 2\pi(TransmissionRange)^2 \\ TransmissionDensity &= \frac{\frac{1}{4}(TransmissionRate)}{TransmissionArea} \end{aligned} \tag{1}$$

To check the validity of my model at its bounds, let us consider a transmission range of nearly zero, and a transmission range that is very large. For the former, the transmission density is nearly infinite, which makes sense because one could pack many nodes into each unit area and they wouldn't interfere with eachother. The latter option results in a very low transmission density (approaching zero), which is reasonable because nodes must be placed very far apart to avoid interference.

For an example node architecture given in Section 5, the first-generation MICA, the transmission area is 5655m^2 and the transmission density is 1.77 bits/sec/m^2. That would mean that one sensor generating 177bps could be placed every 100m^2.

Using multiple frequencies (OFDM for example) should provide a linear increase in the transmission density.

4.1.4 Comments on Scaling

It is my opinion that we need to perform simulations, perhaps using the ns simulator, to study the scaling behavior of various routing algorithms. We have not yet come across a good paper which applies such analysis to many routing algorithms *with emphasis on scalability.* If such a paper is really missing in the literature, it would be a good place for us to make a contribution. We could provide a "benchmark" to fairly compare all the algorithms under equal circumstances (both realistic and synthetic) and then propose a better solution based on our observations.

4.2 Management

Because existing sensor nodes are not designed for large-scale applications (see Section 5), the logistics of network management may be a significant hurdle. For instance, in the Great Duck Island experiment, the design and maintenance of a reliable satellite uplink (no phone lines were available) was a significant proportion of the effort required [16]. Many sensor network applications, due to their remote nature, will incur these management issues.

In addition to the requirement of removing data from the WSN, it is also necessary to provide a reliable and convenient database for storing, retrieving, and/or analyzing the sensor data. For the size of the DMN physical experiments, this should not be an issue. For larger experiments, however, this will be a significant problem that needs to be addressed.

4.3 Cost

The cost of the initial DMN experiments will be dominated by sensor costs. This includes the cost of the sensor node, the sensors attached to the node, and any auxiliary equipment (antenna, solar panels, etc). Other costs associated with the network, such as base stations and administrative costs, will be amortized over the many nodes.

Because we lack the resources to design or build our own sensor nodes, we must buy nodes from an outside source. Currently, the MICA motes are the only commercially available microsensor nodes. It may also be possible to acquire (or purchase) nodes from another research institution provided they are willing to do so.

Crossbow Inc., the maker of the Mica brand motes, only publishes development kit prices. A kit containing 8 MICA motes sells for approximately $2000 [6], for a per-node cost of about $250.

We contacted Crossbow regarding volume purchases, and they furnished the information in Figure 1. This information indicates a per-node base cost of under $190 for the MICA2 mote. The standard sensor board would cost roughly $190, though the DMN would likely eschew that board because it contains sensors that are not needed.

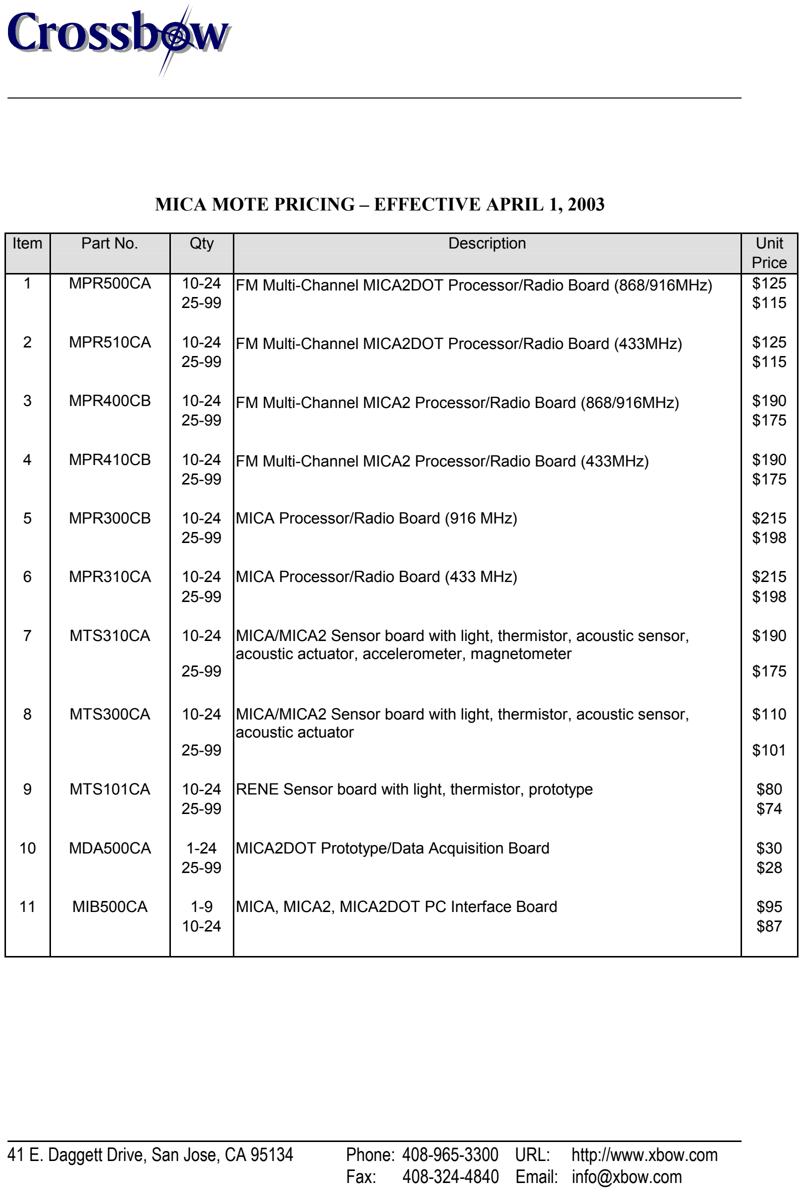

Crossbow

MICA MOTE PRICING – EFFECTIVE APRIL 1, 2003

Item	Part No.	Qty	Description	Unit Price
1	MPR500CA	10-24 25-99	FM Multi-Channel MICA2DOT Processor/Radio Board (868/916MHz)	$125 $115
2	MPR510CA	10-24 25-99	FM Multi-Channel MICA2DOT Processor/Radio Board (433MHz)	$125 $115
3	MPR400CB	10-24 25-99	FM Multi-Channel MICA2 Processor/Radio Board (868/916MHz)	$190 $175
4	MPR410CB	10-24 25-99	FM Multi-Channel MICA2 Processor/Radio Board (433MHz)	$190 $175
5	MPR300CB	10-24 25-99	MICA Processor/Radio Board (916 MHz)	$215 $198
6	MPR310CA	10-24 25-99	MICA Processor/Radio Board (433 MHz)	$215 $198
7	MTS310CA	10-24 25-99	MICA/MICA2 Sensor board with light, thermistor, acoustic sensor, acoustic actuator, accelerometer, magnetometer	$190 $175
8	MTS300CA	10-24 25-99	MICA/MICA2 Sensor board with light, thermistor, acoustic sensor, acoustic actuator	$110 $101
9	MTS101CA	10-24 25-99	RENE Sensor board with light, thermistor, prototype	$80 $74
10	MDA500CA	1-24 25-99	MICA2DOT Prototype/Data Acquisition Board	$30 $28
11	MIB500CA	1-9 10-24	MICA, MICA2, MICA2DOT PC Interface Board	$95 $87

41 E. Daggett Drive, San Jose, CA 95134
Phone: 408-965-3300 URL: http://www.xbow.com
Fax: 408-324-4840 Email: info@xbow.com

Figure 1: Volume pricing information obtained from Crossbow, Inc.

Dr. Kevin Delin, lead of the JPL Sensor Webs Project, has informally stated that the sensor pods cost "a few hundred dollars" in parts. Assuming that they could be obtained at-cost (a possible partnership with JPL), the price may be competitive with Crossbow's offerings.

4.4 Physical limitations

For the experimental DMN, deployment, recovery, and waste management issues will not be a problem. For the ultimate DMN vision to be realized, however, these issues must be addressed. Nodes should be deployable with little or no attention given to orientation, placement, etc. The sensor network should provide for easy recovery of dead nodes, so they can be replaced and to prevent "littering" the environment with dead nodes. Not only will the sensors need to know their own location, but some element of the network will need to reliably log the last known position of other nodes in the network.

Whether this log should be centralized or distributed is not clear; for mobile nodes, it would be prohibitively expensive to keep a centralized log because the location information from every node would have to be sent to a common place—the "many to one" communication that is inherently unscalable. If the nodes are immobile, however, such communication would be infrequent enough that a centralized log would be feasible.

The key paradigm in sensor network theory is that individual nodes do not matter. That is, with large networks, the user will not be interacting with individual nodes. As such, if a node dies, the user will not have any reason to notice. Thus, the network must provide a notification system that does notice when a node goes "off the radar," and makes that information available to the network user.

5 Existing Solutions

While there are only one or two architectures commercially available (MICA & possibly WINS), it is useful to examine the state of the art in the research community as well, as the commercial sector will inherit the fruits of research.

Campbell Scientific [3] has a well tested line of sensor network products that are currently in use all around the world. They use a system of "dataloggers," external sensors, and external radio to comprise a sensor node. The system requires manual set up and is quite expensive (on the order of thousands of dollars per node).

Rockwell [19] has a well-developed line of modular sensor nodes that are being used primarily by the military and somewhat by UCLA researchers. The main downside to WINS is that it is relatively high power (uses two 9V batteries and runs approximately 15 hours) and cost, though unknown, is probably high.

JPL has a project called Sensor Webs which is, in principle, similar to WSN. We believe the project is not getting enough press coverage because it exists outside academia and does not publish as often in the same journals as MIT et al. In many ways, the Sensor Webs project is a pioneering one in the realm of sensor networks. For instance, JPL deployed an environmental monitoring network at Huntington Botanical Gardens fully 9 months before Berkeley deployed its network at Great Duck Island (see Section 7). The project is fairly far along, having constructed at least three versions of the design and implemented several test networks, from California to a NASA launch site in Cape Canaveral [8].

The Sensor Webs project differentiates itself from data-gathering sensor networks by focusing on *in situ* data processing and decision making. In this scheme, sensor data is used in the network to analyze and react to the environment. Another differentiating factor is that the Sensor Web pods are designed to be robust to environmental and handling conditions, and to be installed by the end-user rather than a member of the Sensor Webs project. The pods have been demonstrated to handle temperature extremes and very long deployments—in excess of a year. Evidence has shown that the sensor pods can last on the order of months without any solar recharging, making them comparable to other non-rechargeable node architectures when energy harvesting is not possible.

Medusa is a UCLA offering that seems to have fallen through the cracks. There is very little news available about it, although UCLA's CENS claims it will be incorporating the MK-2 Medusa node in the future, likely as a supplement to Berkeley motes.

MICA is a product of the Berkeley mote architecture, and is currently produced and available from Crossbow, Inc. This architecture is very well tested, having been used by researchers throughout the world for various projects. MICA2 is an updated form that uses a programmable-frequency radio from ChipCon.

PicoRadio is another Berkeley effort, but it focuses more on the ultimate goal of "smart dust," microscopic nodes that have all the functions integrated onto a tiny piece of silicon. As such, the PicoNodes are not available for purchase and are basically just for laboratory use.

MIT's μAmps project is between the PicoNode and the mote. It is sized closer to a mote but is just at a research stage in terms of hardware development and production. The ultimate goal of the μAmps project, as stated at their website, is to develop energy aware microsensor network technology that targets the small distance, low data rate application space.

The MANTIS nymph, from University of Colorado, is a new entry and closely resembles the Berkeley MICA2 motes.

Architecture	Cost	Range	Thpt	bits/sec/m^2	Power	MIPS	Regulation
Campbell	\$?	Miles	9.6kbps	1x10^{-8}	High	N/A	FCC
WINS	\$?	100m	100kbps	4x10^{-1}	Medium	133	ISM
JPL SWEBS	~\$300	150m	50kbps	9x10^{-2}	Medium	?	ISM
MANTIS Nymph	N/A	200m	77kbps	7x10^{-2}	Medium	4	ISM
Medusa MK-2	N/A	10m	115kbps	45	Low	44	ISM
MICA mote	\$200	30m	40kbps	1.8	Low	4	ISM
MICA2 mote	\$200	160-320m	38kbps	6x10^{-2}	Low	4	ISM
PicoNode	N/A	10m	1Mbps	400	Low	25	ISM
μAmps I	N/A	10m	1Mbps	400	Low	133	ISM
μAmps II	N/A	10m	1Mbps	400	Low	?	ISM

Figure 2: Comparison of some existing sensor node architectures. ISM is "Industrial, Scientific, Medical" is defined by the ITU, and includes the common 433MHz, 915MHz, and 2.45GHz bands. Individual countries may place additional limits in ISM bands, but generally under a certain power limit the frequencies are unregulated.

The architectures presented in Figure 2 can be loosely grouped by power. It is interesting to note that while power generally goes down as communication range is reduced, the bandwidth goes up. This is due to the fact that SNR often becomes much better for short distance communications. We also see that the network capacity falls markedly with longer range transmissions, as depicted in Figure 3, due to the shared communications medium. As a result, technologies like BlueTooth can achieve 1MBps communication over distances up to 10 meters, using very low power transmissions (1mW).

As a result, higher density networks can actually be the best case in terms of scalability (provided they are using short range communications), because they have the highest data rate to work with and lowest power requirements to satisfy. The downside, of course, is that a high density network is currently costly (in terms of money and time) to deploy, because many more nodes are required to cover a given area.

Even though architectures like μAmps may provide the best technical advantage, real-world constraints will steer us towards MICA or JPL SensorWebs scale devices in the near term. The smaller architectures will become viable once they are produced in volume.

5.1 Limitations

The main limitation of existing node solutions is the radio technology. None of the currently used radios are especially power efficient. The groups making the most progress in that area are MIT μAmps [17] and Berkeley PicoRadio [2] projects. These teams are working on highly integrated and optimized radios that may provide an order of magnitude or more decrease in the power required per bit transmitted.

The radio in the original MICA mote is single channel and omnidirectional. This means that two motes within interference distance of one another (a distance significantly longer than the attainable transmission distance [15]) cannot successfully transmit simultaneously. The MICA2 uses a programmable-frequency

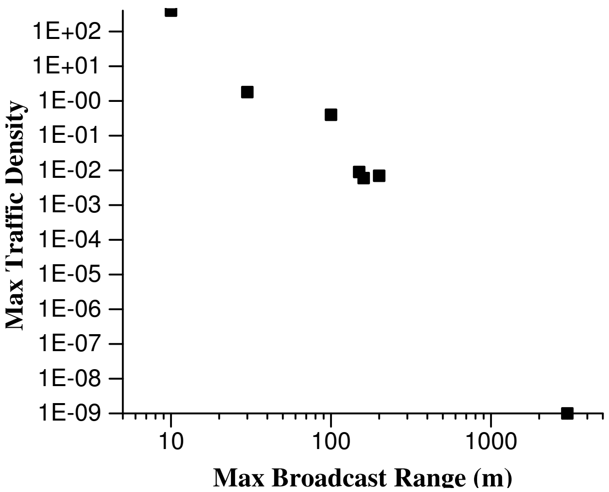

Figure 3: Network transmission capacity, as measured by the bits/sec/m^2 metric. Long-range communications significantly limit the transmission capacity of the network.

radio from ChipCon, allowing multiple channels to be used simultaneously in the same geographic space with fewer interference problems. Though the ChipCon radio used (model CC1000) supports four center frequencies, the MICA2 platform only appears to support one center frequency per model, with between 8 and 100 channels possible in each center frequency. The MANTIS nymph allows all the center frequencies (315, 433MHz, 868, and 915MHz) via software reprogramming.

Slightly more complicated devices, such as the Rockwell WINS and the Campbell Scientific nodes, use spread spectrum technology. While this alleviates the scheduling problem for interfering nodes (bandwidth is simply degraded gracefully—i.e., it is implicitly shared), it comes with a cost: power. Spread spectrum technology costs in terms of data processing, but there is an additional cost which is that the fairly complicated hardware being used is COTS and thus is “tacked on” to sensor nodes rather than being part of the integral design.

BlueTooth radios use frequency hopping spread spectrum (FHSS) to mitigate interference. 802.11b uses direct sequence SS (DSSS) while 802.11a orthogonal frequency division multiplexing (OFDM).

Another limitation of the mote radios is that they are raw data devices that must interrupt the processor to call a software handler for each byte transmitted or received; this includes the preamble and checksum. Future designs will incorporate some smartness into the radio, but for now the per-bit cost is artificially high on these devices. However, the cost is still way below that for larger nodes like the WINS nodes.

Some of these radio limitations may be overcome by configurable radio technologies. These technologies range from simply providing multiple radio channels to full blown “software defined radios,” or SDRs, such as the Vanu SDR System [4], or VSR. Systems such as VSR do not currently fit the low power paradigm of the WSN world—VSR is targeted primarily at high power applications such as wireless base stations and vehicular communications—but some level of configurability will almost definitely find its way into sensor nodes in the future. The most likely first step would be so-called “reconfigurable SDR” systems, where software is able to reconfigure a portion of the baseband hardware (e.g., in an FPGA) to provide dynamic radio services.

6 Future Trends

We have seen an interesting trend wherein some nodes have been getting bigger rather than smaller. To some extent, this is to facilitate testing and experimentation. However, we believe it is partly due to a divergence in WSN goals: some architectures are moving towards “smart dust” design space while others are moving towards macroscopic applications.

Microscopic nodes that can collaborate and provide very fine granularity have two significant drawbacks:

extremely low power budgets and no real durability. We believe there will always be a place for larger, more durable nodes. As technology advances, those nodes will necessarily become more capable and long-lived, but they will not necessarily get significantly smaller. Physical constraints such as antennas, solar panels, and weather-proof enclosures will necessitate a certain minimum form factor in the near future. Eventually, advanced energy harvesting techniques and laser communications may allow a major shift in node form factor.

6.1 Power Sources

Ultimately, energy harvesting is expected to be a huge part of the energy strategy for WSN architectures. In the near term, however, disposable batteries will continue to be popular because of the small scale of the majority of experiments going on. They're easier to analyze in terms of energy capacity and utilization, and it removes the complexity of adding and managing a renewable energy source (like solar).

The JPL nodes seem to be the only ones that routinely provide support for solar charging. Though nodes like the Berkeley mote do allow for external power supplies, one would have to add a custom charging circuit to provide rechargeability.

6.2 Radio

Radio enhancements will be key to the success of all WSN architectures. For instance, for small transmissions the radio "set-up" cost is higher than the transmission cost. Groups like MIT and Berkeley are working on better, more power-aware and system-integrated single-chip radio designs. We can expect to see results from that work in the next 1-2 years. PicoRadio III, a single chip SoC node, is supposed to be operational by June 2003 [2]. The MIT μAmps-II final design (also single chip SoC) is in progress and does not have a published completion date.

7 Existing Applications

The following is a brief summary of some existing applications that are similar in scope to the proposed DMN, with an eye on their accomplishments and weaknesses. we will contrast with the DMN's goals.

Figure 4: JPL SensorWeb node, left, and Crossbow MICA family, right.

Great Duck Island: The GDI project uses approximately 30 MICA motes. The motes are organized into "sensor patches," each of which has a single dedicated "gateway node" which has infinite energy and implements a long-haul link to the island base station using a Yagi directional antenna. The communication range in the patch is on the order of tens of feet; between patch and base station is about a thousand feet. GDI is one of the first physical WSN's set up and run for a considerable period of time. They have published a considerable amount of results and insights that are of use to anyone designing a WSN. While the test was successful, the researchers admit that no new scientific knowledge was produced by the experiment [23], though they are aware of this and one of their goals remains useful application of WSN technology.

Likewise, the DMN's goal is to design the network in collaboration with non-EE researchers with the goal of producing valuable scientific information that is not otherwise attainable.

Huntington Botanical Gardens: JPL's experiment at HBG is largely a proof of concept, like GDI, but it seems to be focused more on the application goals than simply testing out the hardware. For instance, the JPL sensors are helping researchers at HBG understand the life cycle of a rare tropical plant [13]. JPL has similar webs at other locations, such as Cal Poly Pomona, Kennedy Space Center, and a farm in Virginia. The JPL nodes have been shown to work for over a year in these applications, presumably unattended, due to their solar energy harvesting ability.

James Reserve: UCLA's CENS group has set up a habitat monitoring project at the U.C. James Reserve. The project has involved huge infrastructure investments, including massive solar arrays, battery systems, AC inversion circuits, and network connectivity. A combination of Berkeley motes and HP iPAQ PDAs are used in the sensor network to provide general sensor readings, while video is provided by more expensive, and wired, Axis video cameras and servers. The project is making significant progress in terms of providing useful data to researchers and providing infrastructure to support multiple future experiments.

8 Cost Estimates

The hardware costs, at least per-node, have already been estimated in the low hundreds of dollars range. The remaining costs will be incurred in deploying and maintaining the network and its infrastructure. Additionally, the cost of manpower for research is significant.

The ultimate per-node cost is an open question. Berkeley's BWRC desires nodes that cost about $.50 each. The only thing that is sure is that node prices will come down as SoC designs, which incorporate everything on a chip (except perhaps antenna and battery), become commercially available.

8.1 Infrastructure

To successfully deploy one DMN test network, we will need at least one base station and computing resources available to capture and log all the sensor data. As the GDI researchers discovered, this is no small feat. Their system ended up incorporating two laptops for redundancy, as well as a satellite uplink for data transmission. Assuming $6000 for laptops and about $600 for the satellite uplink, equipment costs come to about $6600. There is also a $60 per month charge for the satellite uplink [12].

The DirecWay satellite uplink used by at GDI provides uplink rates similar to a modem. Other services intended to be available soon, such as Hughes SpaceWay and WildBlue [25] will offer better uplink speeds, up to 16Mbps. The main benefit of these systems is that they offer consumer-priced satellite uplinks. If a base station is serving a suitably large WSN, the relatively small monthly access fee is amortized over a very large network investment.

In the GDI experiment, there was a secure building available to house the base station resources. Some sort of secure building would be advisable because of the valuable equipment at the base station, and the DMN should be able to provide this due to the presence of reserve buildings. If the base station must be exposed to the outdoor environment, this will add the expense of an enclosure and, more importantly, a power source for the base station. A suitable enclosure could probably be constructed from Rubbermaid totes or even a small shed, and a 100W solar panel with integrated battery charger is commercially available from Fry's Electronics for $700 [9]. Thus it seems reasonable to assume an additional $1000 cost for a base station enclosure.

The good news is that while up-front costs are high, the continuing infrastructure cost is quite low.

8.2 Maintenance

The maintenance cost of the network (finding and replacing/repairing failed nodes, replacing batteries, etc) is very difficult to estimate. We will assume that batteries only need to be replaced in a failure mode situation, because for a large network it is impracticable to use a disposable battery system. Thus the cost comes from having to visit the field to do replacements or repairs. The GDI researchers note that, for this very reason, it would be valuable to have an experiment that is close to the research center. The DMN proposed applications (see Section 10) are all near U.C. Davis, which should significantly reduce the cost of maintaining the networks.

As a quick and dirty estimate, let us assume that nodes have a hardware failure rate of 10% per year, a battery failure rate of 10% per year, and that it requires one man-day per failure. If a man-day of labor costs $200, a node costs $300, and batteries are (essentially) free, we can expect a per-node yearly cost of about $70.

It becomes apparent in these calculations that it is imperative to avoid site visits if at all possible. This makes the design of the DMN's remote management and diagnostics abilities critical to cost management of the project.

Ultimately, the continuing cost to us will be lower because another entity will be responsible for maintaining the network. This requires the DMN designs to be fully "plug and play" so that people who do not understand the underlying technology can still deploy and recover nodes. This implies that the network is fully self-configuring, self-healing when possible, and provides easily accessible information to the network user.

8.3 Summary of Rough Cost Estimate

Failure	Failure Rate	Per-Node Yearly Cost
Node HW	10%	$50
Battery	10%	$20
Total	20%	$70

9 Critical Challenges

The critical challenges for our proposed applications are cost and scalability. Existing node architectures are very simple, and the available bandwidth is quite limited. Some questions to answer: can we use existing algorithms? How effective will they be? Is there evidence to show that those algorithms have attained their stated efficiency, under real-world conditions?

Considerable research has been done on ad-hoc communications, though not always in the realm of WSN's. Networks such as 802.11 and cellular have spawned many significant research studies into scalability, routing, and power-aware techniques. However, in the WSN arena, no routing protocols have been implemented at large scales due to the cost and logistical problems of deploying large scale WSN's.

The largest WSN deployment that we have been able to determine was by Berkeley for an Intel Developer's Forum demo in August 2001 and involved 800 MICA-like motes [7]. Though the demo mostly worked, the difficulty of managing such a large network became apparent, and the seeming lack of any other similar scale experiments in the following 2 years indicates that the state of WSN protocols is not sufficient to make such large networks feasible.

Node architectures provide another significant challenge because they do not currently provide the level of robustness, customization, cost, or efficiency required for large scale WSN deployment. However, we do not believe we are positioned to develop sensor node architectures at this stage, so this challenge will be met by other research institutions such as Berkeley and UCLA, and companies such as Crossbow and SensiCast.

We believe the primary challenge is in developing, simulating, and testing realistic routing and aggregation methods, since these can and must be layered on top of any sensor node architecture. See Section 4.

9.1 Roadmap

The following is a short term roadmap that outlines some of our immediate needs.

Date	Goal
June 2003	Determine 1-3 possible node architectures
June	Create and/or acquire “benchmark” data sets
June-July	Investigate collaboration with other institutions
July	Obtain and begin testing / characterizing the nodes
July-Aug	Perform simulation analysis using benchmarks
Aug	Decide on a node architecture to use in round 1 deployment

10 Proposed Applications

We are collaborating with Dr. Gregory Pasternack, a Hydrologist at the UCD Land Air Water Resources Department, in designing a sensor network for monitoring some hydrologically interesting sites which are presented in Section 10.1 and Section 10.2. He already has a tethered sensor network in place, but it has several limitations, some of which are listed below:

- The nodes and sensors are expensive
- Acquiring data requires a field visit with a laptop
- The monitored area is limited due to the expense and manual labor required
- The data produced is delayed significantly by the manual monitoring method

Pasternack’s goal is, in addition to generating valuable research data, to provide data to classrooms for instructional use.

In Section 10.3 we describe another similar project proposal that involves ecological monitoring at three U.C. reserves.

We believe that an intelligent wireless sensor network, such as the DMN, could overcome the cost and logistical barriers currently facing such projects. Below is discussion of three projects with a focus on the feasibility with today’s WSN technology.

Besides addressing the cost and logistical problems, a WSN would also provide the benefit of *timely* data reports. When something of interest is going on, scientists could get notified in minutes and have a chance to go investigate while the event is still in progress.

10.1 Consumnes River Preserve

The CRP is a flood plain adjacent to the Consumnes river. During flood stage, the river dumps water into the flood plain. Further downstream, some of that water re-enters the river.

Currently, Pasternack measures the water velocity at a few key points along the boundary of the system. What he would really like is to measure the boundary as well as the interior of the system. Additionally, it would be invaluable for him to get real-time updates about the system.

As a first-order attempt at solving this problem, we have decided to analyze a network of 500 nodes, more or less evenly distributed throughout the CRP. The specific quantities that he would like to see reported are analyzed below.

10.1.1 Solution Space

We will do first-order analysis assuming no data compression is used, because it is difficult to estimate the achievable compression. Thus the estimates below are worst-case estimates.

At one sample every 15-minute interval:

- Water pressure

- Water temperature
- Water velocity, surface and subsurface
- pH
- Conductivity
- Turbidity

These quantities can be reported with about 18 bytes of uncompressed data. Sampled at 15 minute intervals, the entire network would produce a net of 80bps of data. This is clearly feasible with any of the WSN options available today.

At 100Hz burst every 15-minute intervals:

- Water velocity, surface and subsurface

The net data produced for this measurement is 104kbps, which stretches the limit of today's WSN technologies. Besides pushing the data bandwidth limit, it would also require a lot of power since every bit transmitted has a cost associated with it.

The key to wireless sensor networks is that the network cannot report all raw data and still be scalable. Clearly, at 500 nodes, we are reaching, if not exceeding, the scalability limit of today's WSN hardware. In order to push through this barrier, we must do some in-network processing and only report the resulting subset of data.

Pasternack has explained that the rapid 100Hz velocity measurements are required so that accurate rates, in the form of partial derivatives, can be calculated offline. By performing this operation in-network, the net data produced could be reduced by approximately three orders of magnitude (one report from 6000 samples). Careful analysis of the internal data flow in the network must be done, but the worst case will be well under the 104kbps of the naive approach. Discussions with Pasternack regarding the details of this in-network processing are ongoing and present an opportunity for both innovation in WSN design (to support the in-network processing) and experience in designing intelligent distributed algorithms for this type of application.

Once every 10 minutes at 6 locations:

- Webcam resolution snapshot

Perhaps surprisingly, the network load for supporting this feature is not infeasible. Assuming 640x480 resolution, 16bit color depth, and 6:1 JPEG compression, the net data produced would be 8kbps.

Once every 15 minutes at 1 location:

- Air temperature
- Relative humidity
- Air pressure
- Precipitation
- Solar radiation
- Air velocity

These measurements produce a negligible trickle of information, less than 8bps net, and can be reported directly to the researchers at no significant cost delta.

10.1.2 Feasibility with Today's Technology

The CRP system would be feasible with today's technology, though it would push the envelope significantly in terms of WSN implementations. No project described in the literature has scaled beyond dozens of nodes in size.

Using the transmission density metric, we see that each node in the CRP would produce roughly 200bps of data. If we assume the CRP covers an area of 1km^2, evenly distributed nodes would be placed approximately every 2000m^2. Using the estimates from Section 4.1.3, each node should have available to it approximately 3.5kbps, indicating that transmission density is not a limiting factor for the CRP.

The prime challenge for the CRP system would be the cost of the sensors and power sources for the 500 node system. The technical challenge of implementing the network would be significant but feasible, and would provide fertile ground for innovation in intelligent network routing and data aggregation techniques.

The CRP would not ultimately scale when used as a data-collection mechanism; however, it could scale as an event-alert mechanism. As a data-collection mechanism, the scalability limit would be on the order of 500 nodes. If in-network processing were used as discussed above, the limit would be extended because each node would produce dramatically less data.

In order to calculate the upper bound for scalability under this condition, we must re-evaluate the transmission density metric for a higher density network; i.e., where the transmission distances are lower than the maximum that the radio is capable of.

Let $R_m in$ be the minimum transmission distance for a CRP network with N nodes:

$$R_{\min} \quad = \quad \sqrt{2}\sqrt{\frac{1000^2\text{m}^2}{N\text{nodes}}} \tag{2}$$

The max link rate is 38kbps, of which 1/4 or 9.5kbps should be achievable. If we assume no webcams, each node generates 0.231 bps of data. The link limit on scalability is roughly $9500/0.231 \approx 41125$ nodes. Using Equation 2 we calculate the transmission radius as 6.97m, which should be achievable via the ChipCon radio on the MICA2 mote using its -20dBm (10 μWatt) transmit mode.

10.2 Trinity River System

In the TRS, 500 nodes would be placed at regular intervals along the river's path. The nodes would perform the same low-rate water measurements as in the CRP (single sample at 15 minute intervals).

10.2.1 Feasibility with Today's Technology

The net data produced by the system would be approximately 80bps, making this system highly feasible in terms of bandwidth.

Using the transmission density metric, we see that each node in the TRS would produce less than 1bps of data. Because of the linear topology, transmission density will not be a limiting factor for the TRS.

The primary challenge for the TRS would be fault tolerance, due to a lack of redundant communication paths. Preventing such a single point of failure (SPOF) may entail placing the nodes close enough that any one or more points of failure can be tolerated so long as they are not adjacent in the chain.

For instance, if the communications range is nominally 100m, the nodes could be placed at 50m intervals to increase reliability. If the probability of a single node failure (during the lifetime of the system) is 1%, and nodes are independent, 100m placement should result in a system failure probability of essentially 100%:

$$1 - (1 - 0.01)^{500} = 99.24\% \tag{3}$$

50m placement, on the other hand, should result in a probability of only 4.8% for an equal number of nodes or 9.4% for an equal covered distance.[1]

The TRS would not ultimately scale because of the data flooding towards the end of the network; however it is feasible at the proposed scale and could be extended to roughly 120 times its proposed size.

---

[1] Determined via simulation with 500 and 1000 nodes, respectively.

10.3 U.C. Reserve System

Cathy Koehler of U.C. Davis is leading a grant initiative to install sensor networks at three U.C. reserves (McLaughlin, Sedgwick, and Hastings). The goal of the URS project is to bring truly useful sensor data into the classroom to aid instruction and research.

The project would be rolled out over a three year period. Initially, basic environmental data such as temperature, humidity, solar radiation, and precipitation would be deployed at the three reserves in similar environments.

10.3.1 Feasibility with Today's Technology

There is nothing infeasible about this project. The proposed sensor package could be interfaced to an existing sensor node, though we may have to multiplex the signals to the A/D. The data rates generated are far within acceptable ranges, being considerably less than the CRP or TRS projects.

Assuming the reserve experiments cover a total of 1km^2, and sensors produce data at a similar rate as the TRS, we can reuse the analysis from Section 10.1 and state that the number of supportable nodes is roughly 41125.

The sensors would cost approximately $7000 per site, making cost the primary limitation of the scalability of this project. Because of the low data rates, even a periodic modem connection would be capable of relaying the data, though some reserves have high speed Internet connections.

The URS would not ultimately scale; however, the limit is quite high and not likely to be a problem.

10.4 Notes on compression

While aggregation is the most important tool in a sensor network, compression will play a very important role as well. It has been determined that with today's technology, one can run about a thousand instructions for the same energy cost as transmitting a bit of data. So, compressing data makes a lot of sense. It is likely that future node architectures will have dedicated hardware for performing compression and de/encryption, which will further reduce the energy cost of compressing data.

With inter-frame compression (likely to be high due to the immovable cameras) the compression rates achieved for the webcams would likely be much higher. The system could even be configured to report the images only if a significant change in the field of view had been detected.

11 Additional Topics

11.1 Compression

Compression of WSN data is valuable because transmitting bits is usually more expensive than processing. If a bit can be eliminated by a few hundred CPU cycles of work, it is likely worth it.

Compression only works when data has some correlation, i.e. it is not completely random. This correlation can be within a given node's readings or between the readings of nodes in the network.

Pradhan et al. show that correlation between nodes can be leveraged for network compression without the nodes having to collaborate. They give a small example [21] in which a 33% reduction in transmission cost is achieved using their DISCUS algorithm.

Petrovic et al. show that during aggregation of packets, the order in which packets are aggregated can be used to represent useful information [18]. Each ordering represents a symbol, which can be used to represent a bit pattern. This method can only be used in cases where the IDs of each packet are preserved in the aggregated packet. The method presented, called Data Funneling, helps to reclaim some of the overhead of transmitting the packet IDs.

Any compression method will depend on the entropy (randomness) of the data in the network. As long as there is some correlation between pieces of data, compression should be considered.

Let us consider the CRP example and analyze what might be achievable. First, we assume that the packet header overhead is ammortized by using large packets. Second, we assume that methods such as those proposed by Petrovic et al. are used to reduce the overhead of transmitting IDs for each sensor reading. What's left is to estimate the compression achievable on the actual data produced by the network.

To do so, we will assume that the data values, which are collected in 15-minute intervals, only rarely vary by more than ± 128 counts, or 8-bits of fixed point precision. This might correspond to a half-degree of temperature change. We assume that only 10% of all readings exceed this threshold and require a full 16-bit sensor report. Thus for every 10 readings generated, 9 will be 8-bit and 1 will be 16-bit. An additional bit is required to indicate whether the reading is a 8-bit delta or a 16-bit reference value. The compression rate is then given by:

$$1 - \frac{(9 \times 9 + 1 \times 16)}{10 \times 16} = 39.4\% \tag{4}$$

The simple algorithm above only attempts to compress using the autocorrelation of readings from each single sensor. The work of Pradhan et al. can be used when more is known about the correlation between nodes in the CRP network.

11.2 Estimated Transmission Statistics for CRP

Because CRP measurements are produced every 15 minutes, it is desirable to report the measurements equally as often. This precludes the option of buffering multiple measurements and transmitting only when a suitable number have been amassed in order to ammortize the transmission cost.

TinyOS, used by the MICA motes, uses 36 byte packets. A typical CRP measurement might be 18 bytes. Thus it would be ideal to aggregate two measurements per packet, but this would increase the maximum report delay from 15 minutes to 30 minutes.

An alternative approach is to have every sensor transmit a measurement at each 15 minute interval, and use the multi-hop nature of the communication to aggregate packets. In such a setup, the network load would approach one half, because every two packets produced in the network could be combined into a single packet.

Regardless of the aggregation method used, it is useful to analyze the multi-hope characteristics of the CRP project. We will perform this analysis for the case of zero compression. The results should extend to the case with compression, because all that changes is the rate of data produced.

To simplify the analysis, we make the following assumptions:

- The network is deployed in a circular region, with the base station at the center
- The coverage area is 1km^2
- The transmit range is 160m

Based on these assumptions and calculations, we see that the network is broken into three "zones." In zone 0, nodes can communicate directly with the base station. In zone 1, nodes require one additional hop to reach the base station. In zone 2, two additional hops are required. Area calculations show that the zones hold approximately 80, 321, and 179 nodes, respectively (for a total of 500 nodes).

Based on the MICA2 specifications, we calculate the energy to transmit a bit, receive a bit, and process an instruction:

$$\begin{aligned}
(\frac{1sec}{384000bits})(27mA)(3V) &= 211\frac{nJ}{bit_{tx}} \qquad (5)\\
(\frac{1sec}{384000bits})(8mA)(3V) &= 78.1\frac{nJ}{bit_{rx}} \qquad (6)\\
\frac{(8mA)(3V)}{4MHz} &= 6\frac{nJ}{instr} \qquad (7)
\end{aligned}$$

Note that the cost per instruction calculated in Equation 7 is based on the full MICA2 processing current and should incorporate the MCU plus all additional circuitry. The MCU itself, according to Atmel's specifications, should require no more than 3.75nJ per instruction [5].

A more efficient processor could achieve significantly better processing power efficiency than the Crossbow mote. The Imagine stream processor, for instance, achieves 4.7 billion 16-bit operations per Watt [14]. With its partitioned subword arithmetic, this translates to 9.4 billion 8-bit operations per Watt, 35 times better efficiency (with considerably higher performance) than the Atmel microcontroller.

To calculate the sum energy required for the network, we add up the transmissions from the zones and account for the additional receives and transmits required for mult-hop:

$$\begin{aligned} E_{zone2} &= 179[E_{tx} + 2(E_{tx} + E_{rx})] \\ E_{zone1} &= 241[E_{tx} + (E_{tx} + E_{rx})] \\ E_{zone2} &= 80E_{tx} \\ E_{total} &= 1099E_{tx} + 599E_{rx} \\ &= 279\mu J \\ E_{avg} &= 557\frac{nJ}{node \cdot bit} \end{aligned} \tag{8}$$

This figure indicates that there is an overhead of $(557 - 211)/211 = 163\%$ induced by the multi hop statistics that are a result of the CRP parameters in this analysis. This overhead would not decrease as a result of in-network processing.

The processing cost for the CRP routing is difficult to estimate as doing so would require knowledge of the instruction stream produced by the network. However we can pick a contrived example to get a feeling of the relationship between processing cost and transmission cost. Suppose that the data flow produced at each sensor in the network could be compressed to 75% of its original value by issuing 10 instructions per bit (80 instructions per byte). The communication cost saved would be $(0.5)(279nJ) = 139nJ$. The computation cost would be $(\frac{6nJ}{instr})(10instr)(500nodes) = 30\mu J$. The net savings in energy would be $109nJ$, or 20%.

12 Conclusion

Based on my calculations, we are sure that the **CRP project is feasible** but the critical limitations will be power, material cost, and developmental cost, in descending order of criticality. With some basic data compression and simple data fusion techniques, the bandwidth limitation is solvable. More challenging are the logistical problems of node cost, sensor cost, and deployment. The predominant power paradigm, AA batteries, will work for dozens of nodes. For the proposed 500-node network, we must develop and interface a renewable energy source, such as solar power.

The TRS project is simpler than the CRP project on several levels. Firstly, the geography of the system is essentially one-dimensional, making topology decisions very straightforward. Secondly, this simple topology will make the network easier to debug and maintain. Thirdly, the bandwidth requirements are lower and easier to analyze due to the topology. We rate this project as **very feasible**. The drawback of TRS is that the topology may be so simple that analysis of the network will not be very interesting. However, we feel that the sheer scale of the project may be enough to warrant its exploration.

The URS project is the most immediately feasible of the three proposed applications, primarily because it will not involve many sensors. We rate this project as **highly feasible**. The primary point of interest is that the URS exhibits one of the key attributes of sensor networks that we wish to design for: future growth. While the CRP and TRS could be grown in the future, they would likely grow in very predictable ways, i.e. by increasing the density of the sensors. The URS network will start small, and then grow in a somewhat unpredictable way as patches of sensors with different sensing capabilities and tasks are added in regions of interest. In addition, the URS involves three geographically diverse sites, which would require robustness in our design of the DMN. One of The key challenges will be designing the network architecture so it can gracefully handle the increased network size and usage.

None of the three proposed applications appear to violate either the transmission density or per-node capacity.

References

[1] E. Arikan. Some complexity results about packet radio networks (corresp.). *IEEE Transactions on Information Theory*, 30(4):681–685, July 1984.

[2] Berkeley Wireless Research Center. BWRC homepage. `http://bwrc.eecs.berkeley.edu/`.

[3] Campbell Scientific. Homepage. `http://www.campbellsci.com/`.

[4] J. Chapin and V. Bose. The vanu software radio system. In *Software Defined Radio Technical Conference*, 2002.

[5] Atmel Corporation. ATmega128L. `http://www.atmel.com/dyn/products/product_card.asp?part_id=2019`.

[6] Crossbow, Inc. MICA wireless measurement system. http://www.xbow.com/.

[7] David Culler. Largest tiny network yet. `http://webs.cs.berkeley.edu/800demo/`.

[8] Kevin A. Delin, Shannon P. Jackson, Scott C. Burleigh, and David W. Johnson. The JPL Sensor Webs project: Fielded technology. In *Proceedings of Space Mission Challenges for Information Technology*, July 2003.

[9] Fry's Electronics. SolarPRO Plug'n' Play 100 Watt solar powered battery charger, 2003. `http://www.outpost.com/product/3321592`.

[10] Piyush Gupta and P.R. Kumar. Capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[11] Wendi Rabiner Heinzelman, Anantha Chandrakasan, and Hari Balakrishnan. Energy-efficient communication protocols for wireless microsensor networks. In *Proceedings of the Hawaiian International Conference on Systems Science*, pages 3005–3014, January 2000.

[12] HUGHES. American satellite and entertainment DIRECWAY 2-way satellite internet. `http://hughes.americansatellite.com`.

[13] Jet Propulsion Laboratory. JPL Sensor Webs project homepage. `http://sensorwebs.jpl.nasa.gov`.

[14] Ujval J. Kapasi, Scott Rixner, William J. Dally, Brucek Khailany, Jung Ho Ahn, Peter Mattson, and John D. Owens. Programmable stream processors. *IEEE Computer*, pages 54–62, August 2003.

[15] Jinyang Li, Charles Blake, Douglas S. J. De Couto, Hu Imm Lee, and Robert Morris. Capacity of ad hoc wireless networks. In *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking*, pages 61–69, 2001.

[16] Alan Mainwaring, Joseph Polastre, Robert Szewczyk, David Culler, and John Anderson. Wireless sensor networks for habitat monitoring. In *ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02)*, pages 88–97, Atlanta, GA, September 2002.

[17] MIT. Micro adaptive multi-domain power aware sensors (μAmps). `http://www-mtl.mit.edu/research/icsystems/uamps/`.

[18] D. Petrovic, R.C. Shah, K. Ramchandran, and I. Rabaey. Data funneling: Routing with aggregation and compression for wireless sensor networks. In *Proceedings of the First IEEE International Workshop on Sensor Network Protocols and Applications*, pages 156–162, 2003.

[19] Rockwell Scientific. WSN: Wireless Sensor Networks. http://wins.rockwellscientific.com/.

[20] S. D. Servetto. On the feasibility of large-scale wireless sensor networks. In *40th Annual Allerton Conference on Communication, Control and Computing*, Urbana, IL, October 2002.

[21] Pradhan S.S., Kusuma J., and Ramchandran K. Distributed compression in a dense microsensor network. *IEEE Signal Processing Magazine*, 19(2):51–60, March 2002.

[22] Reto Stutz. Performance analysis and optimization of a 2.4ghz, multi-hop, wireless self-configurable network. (Thesis), 2003.

[23] Robert Szewczyk, Joe Polastre, Alan Mainwaring, and David Culler. GDI environmental monitoring application data and lessons learned. `http://webs.cs.berkeley.edu/retreat-1-03/slides/szewczyk-GDILessons.pdf%`.

[24] Leo Szumel and Kathryn Ewell. Wireless sensor network scalability. For EEC289Q, Introduction to Computer Networks, 2002.

[25] WildBlue Broadband. Homepage. `http://www.wildblue.com/`.

[26] Fan Ye, Alvin Chen, Songwu Lu, and Lixia Zhang. A scalable solution to minimum cost forwarding in large sensor networks. In *Proceedings of the 10th IEEE International Conferenceon Computer Communications and Networks (ICCCN 2001)*. IEEE Computer Society Press, 2001.

[27] Wei Ye, John Heidemann, and Deborah Estrin. A flexible and reliable radio communication stack on motes. Technical Report ISI-TR-565, University of Southern California, 2002.