Active Tapes: Bus-Based Sensor Networks

Andrew Tjang, Michael Pagliorola, Hiral Patel, Xiaoyan Li, Richard P. Martin
{atjang,mtp42,hiralp,xili,rmartin}@cs.rutgers.edu
Department of Computer Science, Rutgers University
110 Frelinghuysen Rd., Pisacataway, NJ 08854

Department of Computer Science Technical Report DCS-TR-557

***Abstract*— In this work we present Active Tapes, a bus organization for sensor networks. We construct simple cost models to compare active tapes against more traditional wireless sensor networks. Our models show that density, lifetime, and power consumption play significant roles in determining overall deployment and maintenance cost. We then characterize regimes where tapes are always cost effective, sometimes cost effective, and never cost effective. We then describe three real-world applications using our models. Next, we present an initial implementation of an active tape for power distribution and networking. Finally, we conclude with a discussion of future directions for the active tape architecture.**

I. Introduction

Technology trends continue to make wireless sensors cheaper, smaller, and more power efficient without sacrificing computing power and sensing ability. As these trends continue, sensor networks will become more ubiquitous, ultimately becoming an integral part of the vision of ubiquitous computing. To date, there has been a tremendous amount of research on wireless sensor network nodes and their resulting networks and protocols. In all of these works, a key architectural assumption is that each sensor, or *node*, is a self-contained, generally programmable computing device with a processor, memory, storage, and wireless networking. A second assumption in these works is that each node uses batteries, greatly limiting its energy consumption. We call an organization of individual, wireless and power-scarce nodes the *standard sensor network model*, which is fully described in [1].

Although the standard model is often an attractive organization, there is a dearth of research on alternative architectures for sensor networks. In this work, we thus explore *Active Tapes*, a novel sensor network architecture. An active tape is a sequence of sensor nodes and related units (such as batteries) organized around a bus. A bus of programmable sensor nodes is, in effect, a programmable linear array, thus the term active tape. Buses are a primal mechanism in computer organization, as evidenced by the fact that every modern computer today contains several. A bus provides a simple mechanism to share resources. The same sharing principle applies to sensor networks as it does in more standard computing environments, such as PC's and servers. However, in a sensor network context, the primary sharing concerns center around energy, sensing and networking [10], rather than the traditional triumvirate of I/O, processing, and memory.

This paper serves as an introduction and initial exploration into the design space of active tapes, both as an architectural and organizational sub-unit. We hope our introduction will stimulate novel sensor network applications and spur further research into how tapes can ease power and data management in sensor networks. Our work also serves as a first step toward a formal quantification of the circumstances under which wiring sensors is both practical and economical.

An active tape organization may be desired in different scenarios for different benefits. Clearly, for stationary, dense, high-duty cycle, long-lived networks close to external power, an active tape is a more desirable organization than the standard model. A bus enables the nodes to use external power, and allows higher bandwidth and lower latency communication than wireless networks. The reduced communication cost in turn eases higher-level tasks, such as time synchronization, distributed task scheduling, and distributed sensing. In low-density networks, active tapes would allow sharing patterns that are impossible in the standard model. For example, several sensors can share power supplies, power storage, specialized sensors, and networking, where in the standard model each sensor is limited to its own units. In this context, active tapes function as the organizational unit in a larger sensor network. The unit of deployment in the sensor network thus becomes a collection of tapes, rather than a set of isolated sensors.

There are many interrelated factors determining when an active tape organization is more appropriate as compared to the standard model. These include sensor density, node mobility, network lifetime, proximity to external power, wiring costs, recurring human labor, security, and communication patterns. Our results show that for static deployments, the three key factors impacting the cost effectiveness of active tapes are density, network deployment lifetime, and sensor duty cycle; the desirability of the tape architecture changes dramatically in response to any one of the above factors. A research challenge is to map out the parameter space of where tapes will play a valuable role in sensor networks.

We begin our exploration of active tapes in Section II by defining the factors, such as scale, density and deployment lifetime that impact the decision to use active tapes or the standard sensor network model. We then derive some simple linear-models to mathematically describe these trade-offs in terms of cost to deploy and maintain the network. We focus on power rather than communication trade offs because energy consumption has been the primary focus of much sensor network research [2], [4], [6].

Our models show that in many situations, the network deployment lifetime is a dominant factor for deciding when to use active tapes. Since the standard model is quite sensitive to the cost of changing batteries, often comparing active tapes against the standard model depends on whether the cost of running wires will eventually be offset by the battery replacement cost. For example, our model explains why many sensor network applications use wires during the development stage: the high-duty cycles due to debugging, and proximity to external power make running wires cheaper. Of course, after deployment these factors may change to favor the standard model. In general, the longer the long network life time, especially with high application power consumption, the cost of running wires between nodes will be less than using batteries. In contrast, the standard model can become very attractive under very low power consumption conditions. In particular, when the lifetime of a single sensor's battery exceeds the network lifetime, under even moderate densities, an active tape organization is more expensive than the standard model.

The two additional important factors contributing to the cost of active tapes are node density and application power-consumption. For many applications, however, it is often easier to reason about power-consumption in terms of duty cycle, or fraction of time the sensor node is active. Because each sensing operation requires a fixed amount of power, duty cycle effectively represents power-consumption requirement of the application. We explore the dependency of these three parameters in Section II-B.

In Section III, we describe three different implementations to demonstrate the practicality and simplicity of active tapes. In Section IV we apply our models from Section III to compare three applications' costs when each is implemented in a pure active tape as well as a pure wireless network. We show that real applications vary enough in density, power and lifetime that both the active tape and standard model are appropriate. Section V describes related work. In Section VI we summarize our experiences with the platform and speculate on future research directions. Finally, in Section VII, we conclude.

II. Cost Models and Comparisons

Our approach to comparing the two organizational approaches is to build simple cost models of each organization and compare the Deployment and Maintenance Cost (DMC), as a function of different factors. Because we are most interested in the ratios of these costs, and we assume the node cost without batteries as well as the disposal costs are same in both organizations, we do not consider the more typical total cost of ownership (TCO), which would include these additional costs. Our model also assumes all sensors fall into one of the two classes; that is, all sensors are either wired or wireless. We leave modeling mixed organizations as future work.

We begin by first enumerating the important factors influencing the desirability of active tapes: energy consumption, density, wiring cost, battery cost, deployment lifetime, duty cycle, mobility, resource sharing, proximity to external power, and communication rate. In many cases the factors are not independent and trade-offs exist. For example, duty cycle, density, network lifetime, and communication rates are often interdependent, thus leading to countless trade-offs among these factors alone. We found that although the direction of influence is clear, accounting for mobility, resource sharing, proximity to external power and communication patterns in the model increased its complexity beyond what was practical.

Factor	Factor Variable	Cost
Battery Cost	$C_{battery}$	\$3.00
Recurring Labor	C_{labor}	10.00 \$/hr.
Node Replacement Time	α_T	0.20 hr.
Downtime Cost	$C_{downtime}$	\$0.00
Physical Wire	C_{wire}	0.05 \$/ft.
Node Setup Time	α_{setup}	0.20 hr.
Electricity Cost	$C_{electricity}$	0.05 \$/kw-hr

TABLE I

Actual costs of factors used in the model.

In the remainder of this section we first develop a wireless and a wired model. We then explore the design space varying the deployment lifetime, duty cycle, and number of nodes.

A. Cost Models

Once we succeeded in enumerating the associated factors, we needed to define their relationships. As previously mentioned, the various dependent and independent variables prohibits making generalized statements as to the cost effectiveness of either wireless or wired implementations. However, the following cost models do give us a starting point for analysis.

1) Wireless Model: The cost of a wireless sensor network can be modeled as:

$$\lceil \frac{\alpha_T}{\beta} \rceil \times n \times (C_{batt} + \alpha_{replace} \times (C_{labor} + C_{downtime})) \tag{1}$$

where

$\alpha_{\mathbf{T}}$ is defined as the network deployment lifetime. That is, the total time the nodes will be in service in their current environment. This is perhaps the single most important factor when comparing DMC for ad-hoc networked system.

β is defined as the lifetime of a single battery pack, i.e., how long will a node be functioning before its batteries run out of energy. More information about how to compute β can be found in Section II-A.3.

$\mathbf{n}$ is defined as the total number of nodes in the system. This factor is the second most important factor for determining the absolute (as opposed to relative) DMC.

$\mathbf{C}_{batt}$ is defined as the cost of batteries per node. In this study, each node requires two batteries.

$\alpha_{replace}$ is defined as the amount of time required to replace each node's batteries when they have been depleted.

$\mathbf{C}_{labor}$ is defined as the cost of labor per time unit. In order to simplify the model, we have assumed that all labor costs are the same per time unit.

$\mathbf{C}_{downtime}$ is defined as the cost of downtime of each node per time unit. In this model, we assume that the network is homogeneous in that the downtime cost of each node is equal. We also assume that all nodes will fail within a small time difference, and replacement occurs en masse, as each node is taken down, batteries replaced, and restarted.

Notice the ceiling function applied over $\frac{\alpha_T}{\beta}$. The term represents the number of times the battery set for a node needs to be replaced during the application lifetime. The ceiling function is used to emphasize that the wireless option does not allow amortizing power cost. Even if a set of batteries can last longer than the network lifetime, multiple nodes cannot share them.

2) Active Tape Model: In this section we build models to quantify the advantages of not replacing batteries, and describe the conditions under which these advantages outweigh the disadvantages of wiring each node.

We model the cost of a wired-only sensor network as:

$$C_{upfront} + n \times P \times C_{electricity} \times \alpha_T \tag{2}$$

with

$$C_{upfront} \equiv n \times (C_{wire} \times d + \alpha_{setup} \times C_{labor}) \tag{3}$$

and where

$\mathbf{C}_{wire}$ is defined as the cost of wire per length unit.

d is defined as the distance between nodes. Again, we assume a homogeneous environment in which nodes are evenly spaced.

α_{setup} is defined as the time required to wire each node, and includes the time required to make modifications to the nodes. Additional wiring costs (such as digging trenches) could be included in this term.

$\mathbf{C}_{electricity}$ is defined as the cost of electricity per kilowatt hour.

P is defined as the power consumed by each node per time unit, as calculated by $P = I_{total}V$, which will be further discussed in Section II-A.3.

3) Application Current Draw: Application current draw, I_{total}, is a variable that encapsulates all the inner workings of the node. It represents an amalgamation of all the current draws of every component. It varies according to the specific applications, including the percentage of time each node spends in a sleeping, transmitting and receiving, blinking LEDS, and sensing. With the exception of sleeping, none of these activities are mutually exclusive; each has a unique power cost. I_{total} is also dependent upon hardware implementations. For this study we used the power consumption figures for the Mica motes, as charted in Table II, which are the sample values for all the current draws from CrossBow motes as provided by [7].

We define I_{total} as:

$$I_{total} \equiv \sum (\gamma_{component:state} \times I_{component:state}) \tag{4}$$

Component	Current (Amperes)
Processor (full op)	6.0×10^{-3}
Processor sleep	8.0×10^{-6}
LEDs (each)	20.0×10^{-3}
Radio Rx	8.0×10^{-3}
Radio Tx	12.0×10^{-3}
Radio sleep	2.0×10^{-6}
Sensor (full op)	5.0×10^{-3}
Sensor sleep	5.0×10^{-6}
Flash write	15.0×10^{-3}
Flash read	4.0×10^{-3}
Flash sleep	2.0×10^{-6}
Weather Board (full op)	2.48×10^{-3}
Weather Board sleep	14.5×10^{-6}

TABLE II

The factors that contribute to the overall current draw of a node as well as their corresponding values.

The tuple {component,state} indicates that a given component, in a given state draws $I_{component:state}$ amps for γ percentage of the time. For example, the radio component of a mote that is in receive mode 50% of the time and in transmit mode 50% of the time, would contribute to I_{total} in the following way:

$$\begin{aligned} I_{RadioTotal} = 0.50 \times \left(8.0 \times 10^{-3}\right) \\ +0.50 \times \left(12.0 \times 10^{-3}\right) \\ +0 \times \left(2.0 \times 10^{-6}\right) \end{aligned} \tag{5}$$

And the values of current drawn from the other components are calculated similarly.

In our cost models, the resulting I_{total} is ultimately used to calculate β in the wireless models (with $\beta = \frac{BatteryCapacity}{I_{total}}$) and P (power) in the wired models (with $P = I_{total}V$).

4) Comparison to traditional cost models: In traditional network costs models, the fixed cost is defined as the cost of building the network, and the variable cost is defined as the cost of providing an additional unit of service to an existing customer, e.g. sending a packet or setting up a circuit [8], [16].

Our models do follow this general pattern, capturing the variable costs in the current draw, duty cycle and lifetime parameters. However, our model does not account for fixed costs independent of the number of sensor nodes. i.e., all costs are a function of n. We found it difficult to generalize any difference between wireless and wired sensor networks in relation to fixed costs. For example, it is likely the fixed costs of hiring design engineers would be the same in both cases. We thus leave an exploration of the types and impact of fixed costs as future work. However, the lack of a fixed cost in our model leads to the counter-intuitive result that the cost ratio (as opposed to the absolute difference) between the two types of networks is independent of the number of nodes in the network.

B. Comparing Costs

In this section we characterize the shapes of the costs curves to give a better understanding of the general trends predicted by our cost model. Figure 1 shows the general trend of how

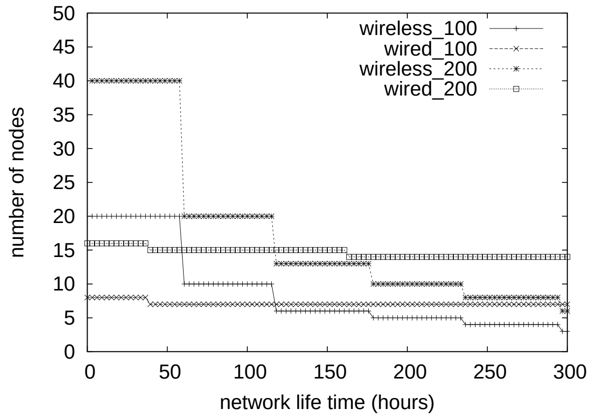

Fig. 1. Overview of supported number of nodes (n) versus network lifetime (α_T) at fixed cost. Lines represent points of constant cost.

networks scale, i.e. the supported number of nodes changes over different network lifetime at a fixed cost budget (although the actual numbers may not always be representative). Since we are limited to a discrete number of nodes and batteries, the curve is a stepwise function. Normally, the cost for setting up a wired network is expected to be higher than the cost of setting up a wireless network. So, at relatively short lifetimes, especially when the first set of batteries for wireless network nodes could last for the whole network lifetime (below 50 hours in Figure 1), the wireless network is more cost efficient than the wired network. However, when the lifetime becomes relatively long, the maintenance cost for wireless networks, i.e., the cost for replacing batteries, becomes dominant. When it reaches the point (around 120 hours in Figure 1) that the accumulated maintenance cost exceeds the initial setup cost for wired networks (the maintenance cost for wired networks, the electricity cost, is negligible), wired networks will catch up to become the more cost efficient choice for a given life time. In summary, choice of wired or wireless networks is basically the comparison of initial setup cost vs. long-term maintenance cost. Since maintenance cost directly relates to network lifetime, the choice over wired or wireless depends on how long the application needs the network.

Figure 1 summarized the general costs when comparing wired networks and wireless networks. However, the specific network parameters still depend on application level knowledge, as slight variances in any of the aforementioned variables can tip the scales in favor of either paradigm.

III. Implementations

In this section, we provide three example scenarios utilizing active tapes. The first is a power-sharing bus, where a user can mix and match sensor nodes along with power generation and storage modules on a single bus. The second is an active Christmas tree, where sensor nodes along an active tape implement a tree that can sense and respond to the environment. The third is a localization application, where a series of tapes implement a simple localization infrastructure in an office building environment.

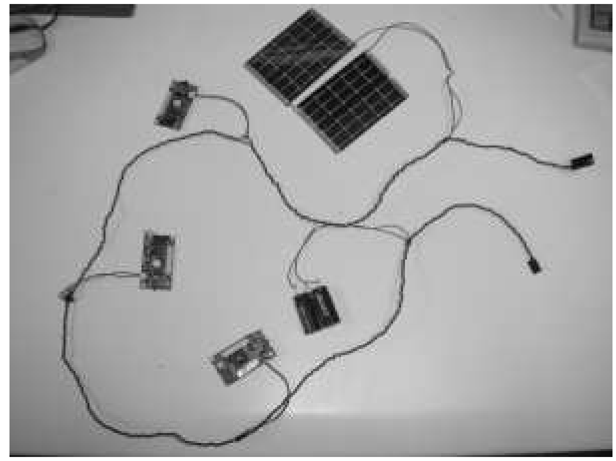

Fig. 2. The image shows a prototype active-tape. The tape is populated with a generator unit, consisting of two solar panels, a capacitor based storage unit, and 3 sensor nodes.

A. Power Sharing Bus

The idea of the power-sharing bus is the core idea of active tapes. When a sensor network utilizes a power-sharing bus, forming an active tape, power costs can be divided among the cluster of nodes attached to the bus. A bus not only reduces the cost per node, but also allows for cost effective energy harvesting, and better power availability via redundant storage modules. By using a bus, a tape can comprise a wide combination of sensor, power, and storage modules. Currently, we have implemented two types of power modules and a storage module for our prototype tape.

Our example implementation of the active tape concept is a general purpose power-sharing bus. Figure 2(a) shows a photograph of the implementation of the active tape. The power bus has been designed with plug-in interfaces that can be customized to meet the power, sensing, computational and storage requirements for arbitrary sensor tasks. Each six foot buses has four pin snap connectors at the ends which allow the buses to be linked end-to-end to create longer buses, as well as allowing for the addition of a data bus at a later date. A tap drops off the bus every foot that can be populated by a sensor module, power module or a storage module.

The first power module we constructed is a standard 12-volt power supply attached to a DC-to-DC converter that provides a constant 3.3 volts to the bus. This module was designed to attach to one end of the power bus. However, with a connector change from four pins to two pins, the module could be attached at any tap on the bus. The other module is a flexible solar cell module, rated at 3 volts and 100 mA when in direct sunlight. Testing has shown two of these power modules, in parallel, are sufficient to power a mote acting as an active localization anchor node on a sunny day. We also found that connecting three of these solar cells in series is sufficient to power the same localization mote in an office lighting environment, or a cloudy day.

The storage module, as implemented, consists of two 2.5-volt capacitors rated at fifty farads linked in series to provide suffi-

Fig. 3. The figure shown an image of the Christmas tree application. The tree outfitted with an active tape that changes the blinking pattern and makes sounds in response to changes in the environment.

cient voltage to power the motes. When linked to a power bus, the storage module can either store excess power generated from energy harvesting, or power the system. Attaching both the storage and solar modules creates a cost effective solution for a low power sensor network. With the ability to share power along the bus, all nodes in the network do not need to be attached to individual power and storage modules. All nodes are no longer required to be in the optimal position for power generation, as only the solar cells need access to light. Excess energy generated is then funneled into the storage modules for use during times of low power generation, enabling the network to run without human intervention after the initial construction of the network.

B. *Active Christmas Tree*

The active Christmas tree is our example application using an Active Tape, as shown by the picture[1] in Figure 3. This ubiquitous computing application makes the tree sense its surroundings and respond to the environment. This application presents an excellent test bed for expressing the design challenges that active tapes solve. The tree's hardware consists of 14 Mica motes with accelerometer and sounder enabled sensor boards, connected by a power bus to form the active tape. We also attached a storage module to allow for external power interruption where one module provides about 2.5 minutes of power. All nodes use wireless communication.

[1]An animated GIF version can be found at:
http://www.cs.rutgers.edu/dataman/ActiveTape/Tree/

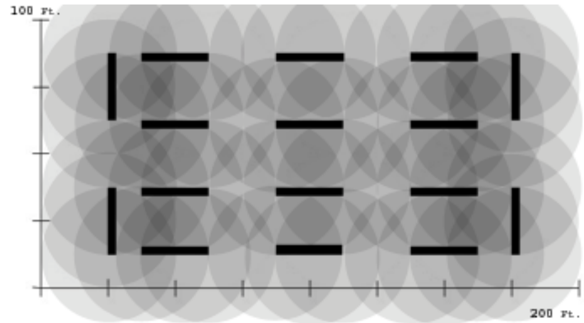

Fig. 4. Arrangement of active tapes for Localization. Each line shows the location of a tape with 3 landmarks. Circles show the expected region where 95% of the packets from a given landmark would be received.

The overall design of the tree can be abstracted into a simple four state machine, where the states are: sleep, happy, unhappy, and panic. During the sleeping state the LEDs attached to the motes blink at a slow rate, as the tree does not detect anyone within range to view the light display. When a radio signal is received from a mote carrying human, the tree responds by increasing the rate of the blinking. This accelerated blinking rate will continue for one minute after the tree stops receiving the human's radio signal and then will time out into the sleep state. The unhappy state is triggered by the tree registering an accelerometer event, signifying an attempt at moving or modifying the tree. During this state, the tree will only blink its yellow LED at the same rate as the happy state in order to warn whoever is moving the tree that it is being disturbed. This state can lead to either the sleep state by a one minute time out or to the happy state by the receipt of the human's radio signal. The final state is the panicked state, which is triggered by the receipt of a security packet emitted from another mote hidden somewhere near the exits of the room. When this state is triggered, the tree's LEDs will blink in unison and the sounder will chirp. The state will persist until the tree receives the human's wireless packet or the motes' battery backups are drained.

A challenge still presented by the active Christmas tree that future active-tapes could solve is the creation of a unified state machine. When completed, the tree should be able to enter a global state using a weighted voting algorithm. This global state needs to be sensitive enough to detect attempts at tampering with the tree while still maintaining sufficient wireless bandwidth to allow the tree to enter the happy and panic states without interruption. The possible solution that active tapes present here is the addition of a data bus to the tape, which allows intra-tree communication over the motes' I^{2}C , leaving the radio free for external communication.

C. *Localization Tapes*

Our hypothetical localization application deploys a set of fixed landmarks, each emitting beacons used by mobile, wireless nodes to localize. Our current approach for localization uses a combination of location fingerprinting[3] and signal strength to distance mapping. The localization application has a high duty cycle, as each landmark sends a maximum radio strength packet once per second. It consumes a moderate amount of power, and the landmarks are in fixed positions, thus making it an ideal candidate for active tapes.

Figure 4 above demonstrates our planned layout of active tapes to cover one floor of our building using stationary motes as landmarks. To obtain a sufficient sample or fingerprint, any given point should be covered by 3-4 landmarks. A study on packet reception vs. distance shows that for Mica motes to achieve a 95% packet reception efficiency, the motes should have a radio range of 25 feet [18]. According to this result, the arrangement of the tapes should be as follows: 3 motes attached to each tape of 25 feet. Given these dimensions, 16 active tapes could be constructed to have most of the coverage area covered by 3 or more landmark motes. From Figure 4 it can be seen that there are certain points near the edges of the grid that have 6 different landmarks covering that location, whereas points near the outer edges of the grid are covered only by 3 landmarks. The coverage of the grid can be explained by how dark a particular region is: lighter shades indicate less coverage while darker shades indicate more coverage. The arrangements of the tapes encapsulate the coverage area optimally with the given dimensions in mind.

Active tapes provide another improvement for localization compared to the standard model with the idea of a power bus. With the power bus in place, the nodes can be on for longer periods of time than they could if batteries were used. With the current configuration of power consumption, wireless nodes would last for approximately 6 days after which battery replacements would be required. Localization qualifies as an application with a relatively high duty cycle, and the up front cost of wired networks is absorbed by this lengthy duty cycle. As one can see from Figure 5, localization becomes relatively more expensive as time goes on. The relative low expense that the power bus provides can play a crucial role to localization if long-term samples are needed.

IV. Application Cost Studies

In order to give the models from Section II real world grounding, we examine three different applications using our model. The first two are the active Christmas tree and localization application as described in Sections III-B and III-C. The third is a habitat monitoring project from the University of California at Berkeley's Great Duck Island project [12]. The applications constitute a broad spectrum in that they have a range of duty cycles, power consumptions, and network lifetimes.

Each application possesses a different challenge to the construction of its network in terms of environment, power consumption, and duty cycle. Great Duck Island was in operation, outdoors, for approximately nine months, which reflects its low duty cycle and low power consumption. The localization application represents a system with a relatively high duty cycle, resulting in medium power consumption. The active Christmas tree serves as an example of a very high duty cycle application with a high power consumption.

For each of these applications, we determined the power consumption of each sensor node based on the same per-unit power draws described in Section II-A.3 and enumerated in Table II. We also used additional information provided in the documentation of the Great Duck Island project's modifications to the standard mote [12]. We then combined the per-unit current draws with the duty cycles for each unit described in Table III to arrive at an overall application current draw, I_{total}. Having computed the expected current draw, we were able to compute the total cost for a given network lifetime for both the wired and wireless approaches. Across all of our models the per-node battery costs (C_{batt}) are \$3, the labor cost (C_{labor}) is \$10/hr, down time cost ($C_{downtime}$) is \$0, wires (C_{wire}) are \$0.05/ft ,wired electricity ($C_{electricity}$) is \$0.05/KWh, the battery capacity is 2268 mA hours and the voltage over wired bus is 3.3 volts.

Each case study uses the different components of the mote for its specific purpose. The mote board is comprised of different components that draw a certain amount of current for each operation undertaken. However, according to which application is running on the mote, the components may or not be used completely. As in the case of Great Duck Island, which is an outdoor application, the power consumption is minimal because it has a low duty cycle. From Table III for the Great Duck Island application, the consumption is very low because the duty cycle for the processor, radio, and weather board are extremely low. On the other hand, Table III shows the active Christmas tree is a high power consumption application because it extensively uses the processor, LEDs, and radio. The localization application consumes a moderate mount of power by only using the processor and the radio. Each application thus has its unique current draw characteristic, which is a key factor in distinguishing the relationship between wired and wireless costs.

Figure 5 shows the relative deployment and maintenance cost of the three applications normalized to the wired case for a range of network lifetimes. Recall that because our model does not contain a fixed term relative to the number of nodes, the relative difference as a function of the number of nodes is constant.

Great Duck Island as described in [12] shows that the traditional sensor network model is quite cost effective for this application. However, our models show that if the experiment extended over years the costs involved to replace the networks batteries every year would quickly eclipse the cost of wiring the network.

For the localization application, the network lifetime is the determining factor in the cost of running the network. However, its lifetime requirement is not as clear as for the Great Duck Island application. Figure 5 shows that for a lifetime of 1 day, the standard wireless model is the clear winner. The

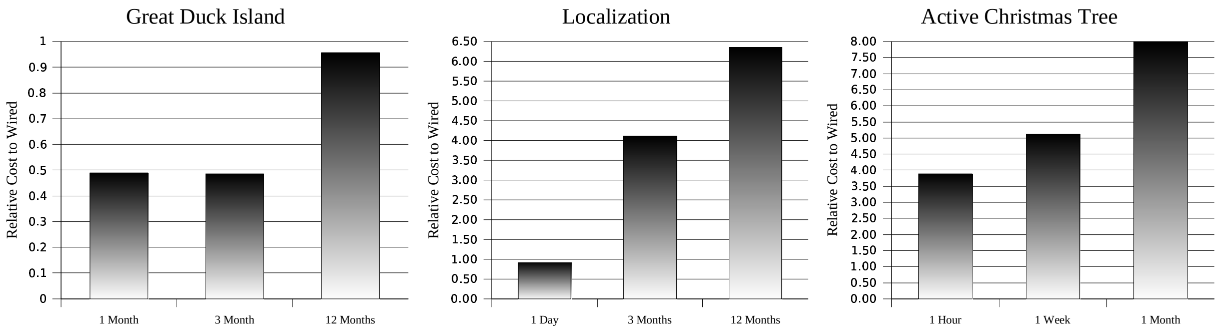

Fig. 5. Comparing relative costs for three different applications of sensor networks for varying durations.

	% Component Active		
Component	Tree	GDI	Local
Processor (full op)	100%	3%	100%
Processor sleep	0%	97%	0%
LEDs (each)	33%	0%	0%
Radio Rx	100%	0.011%	0%
Radio Tx	0%	0.011%	100%
Radio sleep	0%	99.99%	0%
Sensor (full op)	100%	-	0%
Sensor sleep	0%	-	100%
Flash write	0%	0%	0%
Flash read	0%	0	0%
Flash sleep	100%	100	100%
Weather Board (full op)	0%	3%	0%
Weather Board sleep	0	97%	0%
I_{total}(mA-hr)	38.8	.2854	18.01
n	10	150	48
d (ft)	1	31 (avg)	30
α_T	1 week	9 months	3 months
β (months)	0.08	10.88	0.17
$\alpha_{replace}$ (minutes)	6	30	6
α_{setup} (minutes)	12	60	30

TABLE III

The component contributions to the overall current draw of a node as well as their corresponding values, and sample parameter values for each application.

actual cross-over point is at 10 days. For longer deployments ranging from months to years, however, the cost of changing batteries clearly outstrips the cost of running wires.

The active Christmas tree possesses a very high duty cycle, to detect fine-grained movement, and power consumption, to blink lights. Our models show that the cost of running even for only an hour, a standard wireless network is many times greater than when using an active-tape. Thus this application shows there are situations where the standard sensor network model is not appropriate.

V. Related Work

A seminal paper on the subject of sensing for ubiquitous computing made no sharp distinction between wired and wireless [17]. Rather, the important distinction from traditional computing was that the nodes sensed and interacted with the physical world. The work thus presented two challenges: handling mobility in a transparent way as well as coping with massive changes in bandwidth when moving between wired and wireless networks.

More recent work on wired sensor networks consists of electronic textiles [14]. These works explored how networks can be woven into the fabric of clothing. These works differ from ours because the natural interconnect structure of clothing lead them to explore mesh interconnects rather than simple buses. However, until clothing is constructed of conductive materials, it will be much cheaper to retrofit existing clothing with a simple bus rather than a mesh.

Another work in the domain of wearable computing makes a good case for replacing batteries with renewable energy sources [15]. Our active tape architecture is ideal for similar situations, where distributing power across multiple nodes is a concern.

Other works examining wired sensors were in ubiquitous computing environments. One was a sensing table and lazy-susan in [13]. Here, Ethernet was used inside the table as the wired network to send data between the lazy-susan and various RFID tag readers, and the outside world. A second work examined using geographic information for routing video data over an IEEE 1394b (FireWire) based sensor network [5]. A similar argument for the utility of using geographic routing approach in wired sensor networks was made in [11]. None of these works, however, quantitatively compared wireless to wired networks.

A more recent study specifically mentions the low-spatial density requirement for wireless networks [9]. That work used a wide-area, outdoor seismic monitoring as an example application. As in [10], the works outlined some conditions where wireless sensors make sense, such as low-power (and thus low-duty cycle), sparse networks far from wired power. However, an open question is if the sensor-networks are defined by such requirements, or if sensing the physical world physical world is the primary requirement and wireless vs. wired applies only in the appropriate circumstances.

VI. Future Work

Our initial exploration into the design space of active tapes focused on power sharing. Although power sharing alone presents a convincing argument for active tapes, we see 3 future directions for active tapes: active power management, data buses, and mixed organizations.

a) Active Power Management.: Our applications simply added storage and generation units to the bus to get the proper level of power to all the units. More sophisticated designs could manage available power across the entire tape, activating and de-activating sensors in response to available energy. In addition, the power generation and storage units could communicate their states along the tape. Such tapes, as they adjusting to environmental conditions, would in effect act as a mini power grid.

b) Data Buses: In addition to high bandwidth sensor network applications, such as [5], there are also low bandwidth applications that can benefit from data buses. Data buses greatly ease higher-level problems such as time-synchronization, flooding, and serialization. A data bus can be especially attractive if the cost of a power bus infrastructure has already been determined to be more cost effective over its wireless counterpart. In that case, a data bus can be added with little effort. . For example, we built an active tape with a data bus using the motes' I^2C interface using a cable with only 2 additional wires.

c) Mixed organizations: In regions of local density or proximity to external power, tapes can be used as a building block. Economies of scale can make it attractive to concentrate energy storage in larger batteries and distribute the power over a set of local tapes.

By allowing multiple sensors to share the same battery, tapes reduce the discretization effect of batteries. The tape can thus smooth the "staircases" in Figure 1. Depending on the capacities of the batteries used, such an effect could be substantial. For example, tapes are the most appropriate organization in wearable computing environments, where environmental factors favor tapes. The cost of running wires inside clothing is fairly low, the lifetime of the network is long (months to years), the duty cycles for monitoring can be high (e.g., for breathing and heart rate), and finally, the density is fairly high (only a few feet separate units).

VII. Conclusions

In this paper we proposed a novel bus-based approach to the deployment of sensor networks. A key contribution of our "active tapes" is a quantitative model that allows network designers to make better decisions when implementing their networks. We thus introduced a cost model for comparing wired and wireless systems and analyzed the deployment and maintenance costs for each. In addition to exploring the abstract design space, we applied our models to three applications. We found that the wide variety of densities, power-consumptions, and lifetimes of these applications called for examining the standard sensor network model. We also provided an initial implementation of active tapes as a proof-of-concept in the form of a power bus and used it in a real application, an active Christmas tree. Being able to plug in various power generation and storage units into the bus also allows for the self sufficiency of a sensor network without the overhead of providing each node with its own power source.

We saw that the initial, potentially high, setup costs for wired networks could be offset and amortized over long application lifetimes. These are the same long duty cycles that force wireless solutions to limit the number of nodes in a system or limit "wake time" for all nodes. We have shown the intuition that replacing batteries is costly, and have quantified how costly in contrast with active tape implementations.

Acknowledgments

We would like to the thank Kevin Wine for his invaluable advice. Without it, building our prototype would not have been possible.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A Survey on Sensor Networks. *IEEE Communications Magazine*, 40(8):102– 114, Aug. 2002.

[2] J. Aslam, Q. Li, and D. Rus. Three Power-aware Routing Algorithms for Sensor Networks. *Wireless Communications and Mobile Computing*, 3(2):187–208, Mar. 2003.

[3] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *INFOCOM*, March 2000.

[4] S. Chandra. Wireless Network Interface Energy Consumption Implications of Popular Streaming Formats. In *Multimedia Computing and Networking (MMCN'02)*, San Jose, CA, Jan. 2002.

[5] V. Chandramohan and K. Christensen. A First Look at Wired Sensor Networks for Video Surveillance Systems. In *Proceedings of the 27th Annual IEEE Conference on Local Computer Networks (LCN02)*, Tampa, FL, Nov. 2002.

[6] J.-C. Chen, K. M. Sivalingam, and P. Agrawal. Performance Comparison of Battery Power Consumption in Wireless Multiple Access Protocols. *Wireless Networks*, 5(6):445–460, 1999.

[7] Crossbow Technology Inc. Mote Battery Life Calculator Spreadsheet. Available at http://www.xbow.com/Support/tinyos.htm.

[8] N. Economides. The Economics of Networks. *International Journal of Industrial Organization*, 14(2), Mar. 1996.

[9] J. Elson and D. Estrin. *Wireless Sensor Networks: A Bridge to the Physical World*, chapter in Wireless Sensor Networks. Kluwer, 2004. in press.

[10] D. Estrin, D. Culler, K. Pister, and G. Sukhatme. Connecting the Physical World with Pervasive Networks. *IEEE Pervasive Computing*, 1:59–69, Jan. 2002.

[11] J. Heidemann. Using Geospatial Information in Sensor Networks. In *Proceedings of the CSTB workshop on Intersection of Geospatial Information and Information Technology*, Oct. 2001.

[12] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson. Wireless Sensor Networks for Habitat Monitoring. In *ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02)*, Atlanta, GA, Sept. 2002.

[13] O. Omojola, E. R. Post, M. D. Hancher, Y. Maguire, R. Pappu, B. Schoner, P. R. Russo, R. Fletcher, and N. Gershenfeld. An Installation of Interactive Furniture. *IBM Systems Journal*, 39(3/4):861–879, 2000.

[14] E. R. Post, M. Orth, P. R. Russo, and N. Gershenfeld. E-broidery,Design and Fabrication of Textile-based Computing. *IBM Systems Journal*, 39(3/4):840–860, 2000.

[15] T. Starner. Human-Powered Wearable Computing. *IBM Systems Journal*, 35(3/4):618–629, 1996.

[16] H. Varian. Differential Pricing and Efficiency. *First Monday*, 1(2), Aug. 1996.

[17] M. Weiser. The Computer for the 21st Century. *Scientific American*, pages 94–104, Sept. 1991.

[18] A. Woo, T. Tong, and D. Culler. Taming the Underlying Issues for Reliable Multhop Routing in Sensor Networks. In *ACM SenSys*, Los Angeles, CA, Nov. 2003.