Power Scheduling for Wireless Sensor and Actuator Networks

Christopher J. Rozell
Rice University
Houston, Texas 77251-1892
crozell@rice.edu

Don H. Johnson
Rice University
Houston, Texas 77251-1892
dhj@rice.edu

ABSTRACT

We previously presented a model for some wireless sensor and actuator network (WSAN) applications based on the vector space tools of frame theory. In this WSAN model there is a weight associated to each sensor-actuator link denoting the importance of that communication link to the actuation fidelity. These weights were shown to be useful in pruning away communication links to reduce the number of active channels. Inspired by recent work in power scheduling for decentralized estimation, we investigate the optimal allocation of system resources for achieving a desired actuation fidelity. In this scheme, each sensor acquires a noisy observation and sends a message to a subset of actuators using an MQAM transmission strategy. The message sent on each sensor-actuator communication link is quantized with a variable number of bits, with the number of bits optimized to minimize the total network power consumption subject to a constraint on the actuation distortion. We show analytically and verify through simulation that performing this optimal power scheduling can yield significant power savings over communication strategies that use a fixed number of bits on each communication link.

1. INTRODUCTION

Recent interest in wireless sensor networks (WSN) has led to increased research in many areas central to distributed data processing, including novel information processing, communications, and networking strategies. Energy conservation to increase the functional lifespan of the WSN is very important because the battery-powered sensors may be difficult (or impossible) to access on a large scale for maintenance. This energy conservation principle generally translates into minimizing the communication among sensors to preserve both individual node power and total network throughput. Basic research in communications and networking can improve the efficiency in transporting data from one location to another in the network. However, these aspects are limited by the underlying information processing strategies that determine the type and amount of data that needs to be transported within the network to achieve the application goal. Consequently, much of the recent sensor network research has focused on adapting well-known signal processing algorithms to distributed settings where individual sensor nodes perform local computations to minimize the information that needs to be passed to more distant nodes (e.g., [2, 11, 20, 23]).

Many WSN algorithms start with the assumption that there is information contained in the sensor measurements that must be communicated to a destination (called a sink, or a fusion center) that is accessible by the WSN operator. In these applications, the sink represents the notion that the data collected in the WSN is only valuable to the extent that it can be removed from the network. However, in many applications the implicit assumption is that the information coming out of the network will be used to monitor the environment and *take action* when necessary. One example of this is in agricultural irrigation. Soil moisture measurements are taken precisely so irrigation levels can be adjusted to improve growing conditions while conserving as much water as possible.

A significant and natural extension to the sensor network paradigm is a wireless sensor and actuator network (WSAN). A WSAN consists of a network of sensor nodes that can measure stimuli in the environment and a network of actuator nodes capable of affecting their local environment. With direct communication from sensors to actuators, WSANs can achieve the application goals without aggregating the information in a single location or removing it from the network. In this way, WSANs can approach the robust and efficient ideal embodied in the mantra: "the network is the processor". Using only local communications may become even more important in applications where the sensor nodes are deployed under a layer of soil, thereby deteriorating the wireless channel and significantly reducing the realistic communication distances [1]. While WSAN applications open the door for in-network processing, optimal WSAN information processing strategies may be very different from the strategies developed for WSNs. Sensor processing and communication strategies that blindly optimize sensor data fidelity may not yield the best results when actuation is involved. Information strategies in the WSAN must be designed with the final actuation performance fidelity in mind.

The task of merging sensed information directly into actions efficiently but without centralizing the information and decision making is difficult primarily because of the redundancy among the sensors and actuators. The coordination necessary to resolve this overlap must be built into the behavior of each individual node if no centralized controller will be used. Fortunately, distributed sensing and actuation is modeled to us in biology where neural systems perform a chain of tasks very similar to the needs of WSANs: sensing, analysis, and response. Furthermore, evidence indicates that neural systems represent and process information in a distributed way (using groups of neurons) rather than centralizing the information and decision making in one single location. We have recently presented a WSAN model based on neural reflex behaviors and leveraging the vector space tools of frame theory to analyze the effects of sensors and actuators with overlapping regions of influence [24]. In this model (summarized in section 2), we showed that such a WSAN could exactly implement linear actuation strategies through direct communication from sensors to actuators and without centralizing the information. We also demonstrated that each sensor-actuator link has an associated importance value that can be used to efficiently determine which communication links are least important to the final actuation fidelity and can be eliminated to conserve power.

In the present work we consider the optimal power scheduling problem for this WSAN model in an inhomogeneous sensing environment. To optimally allocate the resources in the WSAN scenario, it is critical to consider not only the importance of a single sensor measurement value to the final actuation fidelity, but also the reliability of that measurement and the energy required to communicate it to its destination. These factors represent competing interests and must all be considered jointly when determining the optimal power allocation (i.e., fidelity) for each communication link in the WSAN application. Inspired by recent work for power scheduling in decentralized estimation [25], we propose and solve the optimal power scheduling problem for this WSAN model. The solution is based on an MQAM communications scheme with a fixed bit-error rate for the wireless environment and a uniform quantization scheme for sensor measurements with varying local SNRs. The resulting optimization problem is relaxed to a convex program that can be solved explicitly and evaluated numerically. We demonstrate its effectiveness in significantly reducing the total power expended over uniform resource allocation schemes.

2. BACKGROUND

Neural reflex behaviors provide a starting place to think about merging sensed information directly into action with overlapping sensor and actuator elements. As an example, we consider the crayfish dorsal light reflex [19] where light movement in the visual field elicits predictable eyestalk movement. The main visual representation is comprised of a collection of sensory neurons which each represent the sum of light activity in overlapping spatial regions. The crayfish eyestalk movement is controlled by muscles that individually generate movement in one specific direction. As with the sensory units, the muscle movement directions also overlap in the movement space (i.e., muscle movements are not "orthogonal"). Most importantly, the activity of each muscle is determined directly from a processed combination of sensory neuron activity with no centralized controller. Though all of the muscles have to be coordinated to produce the desired total action, their distributed individual responses are generated directly from the distributed sensory representation and without a centralized decision-making structure. Previous research has shown that even in this critical behavior, the contributions of each sensory unit to the total action are simple and essentially linear [14].

The WSAN model presented in [24] follows the principles seen in this example from the crayfish. In this model, a collection of sensors measuring overlapping spatial regions gather information about a stimulus field. A collection of actuators have individual environmental effects that overlap and must be coordinated. Each actuator determines its individual contribution to a behavioral goal through a combination of the sensor measurements. The actuators do not communicate directly with one another but their behavioral decisions take into account what the other actuators in the network will be doing. To analyze the performance of a WSAN under different design decisions, we use mathematical models based on the familiar tools and terminology of vector spaces.

2.1 Vector space sensor and actuator models

Sensor network models often begin with a collection of sensors distributed over a 2-D spatial field limited to the spatial domain Ω (e.g., $\Omega = [0,1]^2$). Sensors are indexed by $k \in K$, and are located either irregularly or on a regular grid. The spatial region being sensed contains a stimulus field, denoted by $x(\omega)$, where $\omega \in \Omega$ indicates location in the field. Sensor measurement models often consist of averaging the stimulus field over non-overlapping spatial region surrounding each sensor [21]. We generalize that notion by representing each sensor by a receptive field $s_k(\omega)$ over Ω that performs a weighted average over a spatial region. The sensor receptive field would be defined by the physics of the device and could indicate sensors that are directional or have varying sensitivity over a region. The receptive field $s_k(\omega)$ may also incorporate the physics of the environment surrounding the sensor. For example, a sensor taking a measurement by averaging readings over a time window may have a receptive field corresponding to the spatial averaging induced by the propagation properties of the surrounding medium, corresponding to the Green's function [4] for the physical process being measured. In the agricultural irrigation scenario, a soil moisture sensor measurement over a time window represents a spatial average of the true moisture that is defined by the diffusion behavior [18] of the soil.

Ideal sensor measurements of the field are therefore given by

$$m_k = \int_{\Omega} x(\omega) s_k(\omega) d\omega = \langle x, s_k \rangle. \tag{1}$$

We will not assume any particular arrangement or shape of the sensor fields; in general we expect sensors to be irregularly spaced and have highly overlapping receptive fields. This vector space view of the sensor measurements indicates that the measurements can represent any stimulus signal in the space $\mathcal{H}_x = \text{span}(\{s_k\})$. The space $\mathcal{H}_x$ represents a restricted class of fields that is consistent with the resolution of the sensors. For example, $\mathcal{H}_x$ may be a space of spatially bandlimited functions over Ω. The actual stimulus field in the environment may not be in $\mathcal{H}_x$, but the sensors have a limited resolution (depending on design and placement of the sensors) that precludes them from sensing an unrestricted class of signals. Therefore, we assume that $x \in \mathcal{H}_x$, though in reality x only represents the component of the true environmental field within the sensing resolution of the network.

Just as individual sensors have local but overlapping regions of sensitivity, actuator networks are composed of individual actuators that each can affect the environment through (possibly overlapping) local regions of influence. Actuators are indexed by $l \in L$, and again are located either irregularly or on a regular grid. Whereas each sensor is represented by a receptive field, each actuator is represented by a influence field over Ω, denoted by a function $a_l(\omega)$. As with the sensor receptive fields, an actuator's influence field depends on the physics of the specific device and the surrounding medium. In the agricultural irrigation example, the actuator influence field may represent the water delivery pattern of the sprinkler elements. Additionally, the influence field function $a_l(\omega)$ may also incorporate the physics of the environment, such as the water absorption or runoff rates of the soil.

Each actuator responds with an intensity that indicates how strongly it acts on the environment. We will model an actuator's intensity d_l as weighting its influence function. The resulting total actuation field over Ω is

$$y = \sum_{l \in L} d_l a_l.$$

The collection of actuators can therefore cause any actuation field y in the space $\mathcal{H}_y = \text{span}(\{a_l\})$. The space $\mathcal{H}_y$ represents a restricted class of fields that is consistent with the resolution and placement of the actuators (e.g., a class of spatially bandlimited signals, etc.).

It is critical to note that the collection of sensors $\{s_k\}$ and actuators $\{a_l\}$ do not share many characteristics; they can have different numbers of elements at different locations over Ω. Most importantly, individual sensor and actuator functions can have different shapes and even involve different modalities (e.g., temperature sensors and water delivery actuators). Consequently, $\mathcal{H}_x$ and $\mathcal{H}_y$ can be *very* different spaces of functions, and defining them in terms of general vector spaces allows us to make connections between the sensed inputs and the resulting actuated outputs.

In order to design effective communication strategies between sensors and actuators, we need methods to analyze the relationship between individual node activity (m_k and d_l) and the resulting impact to signals in $\mathcal{H}_x$ and $\mathcal{H}_y$. The analysis is complicated because of the overlap between both individual sensor receptive fields and actuator influence fields; in short, the representational elements are not orthogonal. We appeal to the tools of *frame theory* to analyze systems of linearly dependent elements that will apply to both the collections of sensor and actuator functions.

2.2 Frame theory

The collections of sensor receptive fields and actuator influence functions form a representation for a signal space in the environment ($\mathcal{H}_x$ and $\mathcal{H}_y$, respectively). Though the overlap among the sensors and actuators makes the situation more complicated than with a familiar orthonormal basis (ONB), we can still use familiar vector space methods to learn how to deal with these collections of elements. In this section, we will consider a general collection of vectors $\{\phi_j\}$ indexed over J. Fundamental results about this generic collection of overlapping vectors can then be applied to the sensor and actuator representations.

In general, collections of sensor receptive fields and actuator influence fields are not orthogonal and may be linearly dependent. When a representation system has vectors that are linearly dependent, the collection of vectors technically no longer forms a basis. A collection of M vectors $\{\phi_j\}$ is said to form a *frame* [3, 10, 12] for the space $\mathcal{H}$ if there exist constants $0 < A \leq B < \infty$ such that Parseval's relation is bounded for any $x \in \mathcal{H}$,

$$A\,||x||^2 \leq \sum_{j \in J} |\langle \phi_j, x\rangle|^2 \leq B\,||x||^2.$$

Because of the dependency present between frame vectors, the same set of vectors cannot generally be used for both analysis and synthesis,

$$x \neq \Phi'\Phi x = \sum_{j \in J} \langle x, \phi_j \rangle \phi_j.$$

Instead, a set of dual vectors are used for the reconstruction,

$$x = \sum_{j \in J} \langle x, \phi_j \rangle \widetilde{\phi}_j.$$

While there are an infinite number of sets of dual vectors that will work for reconstruction, the canonical dual set is calculated to reduce the error due to corrupted coefficients as much as possible (i.e., the dual vectors represent the pseudoinverse operation) [17]. The set of dual vectors $\{\widetilde{\phi}_j\}$ is also a frame for $\mathcal{H}$, with lower and upper frame bounds $\left(\frac{1}{B}, \frac{1}{A}\right)$, respectively. Importantly, the roles of the frame and dual set in the reconstruction equation are interchangeable,

$$x = \sum_{j \in J} \langle \phi_j, x \rangle \widetilde{\phi}_j = \sum_{j \in J} \langle \widetilde{\phi}_j, x \rangle \phi_j.$$

In an ONB, perturbing a measurement coefficient (including removing it entirely) has a proportional impact on the reconstruction — the energy in the reconstruction error is the same as the energy in the perturbation. The redundancy present in a frame can provide a measure of robustness to perturbations that is not present in orthonormal systems, but it also makes the effect of such perturbations harder to analyze. In [24], we show that when a perturbation p_j is added to each frame coefficient c_j in the reconstruction,

$$\hat{x} = \sum_{j \in J} (c_j + p_j)\,\widetilde{\phi}_j,$$

the total error is bounded by the inverse of the lower frame bound

$$||x - \hat{x}||^2 \leq \frac{\sum_{j \in J} p_j^2}{A}. \tag{2}$$

In words, the perturbation energy is reduced in the reconstruction by at least the minimum redundancy in the set of frame analysis vectors $\{\phi_j\}$. The upper bound in (2) is consistent with probabilistic robustness results when stochastic noise is added to frame coefficients [15].

2.3 WSAN actuation strategies

A specific WSAN application can be defined by its goal. Given an environmental signal, the goal defines the optimal actuated response. For example, in the agricultural irrigation scenario the goal would be to keep the measured soil moisture values close to a pre-determined set-point. In a insect repellent application, the goal may be to deliver insecticides based on the measurements of temperature and humidity sensors. For any measured stimulus field x, we assume that there is a mapping $T : \mathcal{H}_x \rightarrow \mathcal{H}_y$ that defines the ideal actuation field response, $y = Tx$. The mapping T would be determined as a design specification for the WSAN in advance, and we assume that it remains fixed over a duration between WSAN calibrations. Such a control law based strictly on the current measurements obviously doesn't encompass every interesting WSAN application (e.g., target tracking and pursuit). However, such a control strategy does cover a number of significant settings, and our analysis will be limited to applications that fall into this class.

The mapping T may reflect either an open- or closed-loop control scheme, depending on the nature of the application and whether the actuator activity is reflected directly in the sensors. In a closed loop scenario, the linear mapping T is analogous to a proportional controller [13]. As one specific example, assume that the user has specified a desired environmental field x_0 as a set point for the system. For example, a farmer may specify the desired moisture levels of a field containing several different types of crops. In this case, we may want the current environmental field with the addition of the actuated field to be equal to the set point, $x_0 = x + y$. One simple actuation strategy would to apply the identity mapping ($T = I$) to the difference of the desired and the actual environmental field, $y = x_0 - x$.

Following our example of reflex behavior, actuators must generate their own activity using measurements received from the sensors and without communicating with the other actuators. The overlapping influence fields of the actuators prevent a purely greedy approach where each actuator generates the locally optimal activity. Nearby actuators could be nearly identical and wildly overcompensate their actions in a greedy approach. Sensors and actuators must together have coordinated behavior that accounts for the components of the action field that the other nodes must be covering.

To formalize this notion of coordination, we draw on our discussion of frame theoretic models for sensors and actuators in section 2.2. We assume that the collection of sensors represented by $\{s_k\}$ form a frame for $\mathcal{H}_x$ with frame bounds (A_s, B_s) and with dual functions given by $\{\tilde{s}_k\}$. Similarly, we assume that the collection of actuators represented by $\{a_l\}$ form a frame for $\mathcal{H}_y$ with frame bounds (A_a, B_a) and with dual functions given by $\{\tilde{a}_l\}$.

Given a specified actuation function T and a current environmental field x, the WSAN would like to generate actuation coefficients $\{d_l\}$ such that:

$$y = Tx = \sum_{l \in L} d_l a_l.$$

Drawing on frame theory, we show in [24] that the optimal actuation coefficients can be generated from a weighted sum of the sensor measurements

$$d_l = \sum_{k \in K} w_{k,l} m_k, \tag{3}$$

where $w_{k,l} = \langle \tilde{a}_l, T\tilde{s}_k \rangle$. Unfortunately, each ideal coefficient d_l is a sum including sensor measurements s_k over all $k \in K$; each individual actuator would require knowledge of *every* sensor measurement in order to generate its own optimal actuation intensity.

A scenario where every sensor in the network communicates its measurement to every actuator would present an unreasonable communication burden on the network — approximately $|K| \cdot |L|$ communication paths would be necessary. Because of sensor noise and the need to limit the communication costs, we will only be able to generate approximate coefficients $\{\hat{d}_l\}$ at each actuator, which will generate the approximate actuation field $\hat{y}$. In [24] we showed that the MSE in the actuated field is upper bounded by a term proportional to the total MSE in the actuation coefficients,

$$||\hat{y} - y||^2 \leq B_a \sum_{l \in L} \left(d_l - \hat{d}_l \right)^2.$$

However, it is intuitive to think that relative to an individual actuator's behavior, some sensor measurements will be more important than others. For example, a moisture sensor spatially located a long distance away from the influence field of a specific irrigation actuator will likely have very little relevance on that actuator's optimal behavior coefficient. We also showed in [24] that the weights $\{w_{k,l}\}$ serve as an importance value indicating how critical a given communications link is to the actuation fidelity. In numerical examples, these relative weights served as a good indicator to which wireless links can be eliminated to reduce communications costs.

3. OPTIMAL POWER SCHEDULING

We will consider a set of k sensors, each recording the ideal measurement corrupted by additive noise

$$m_k^s = m_k + n_k^s = \langle x, s_k \rangle + n_k^s.$$

We assume that the sensor noise is zero mean with variance $\mathcal{E}\left[(n_k^s)^2\right] = \sigma_k^2$, but the distribution is otherwise unknown. Note that in general, the sensors can be inhomogeneous with different local signal to noise ratios. We further assume that the measurements lie in the range $m_k^s \in [-A, A]$. Note also that depending on the form of the actuation law, the sensor may record the difference from a nominal value (or a set point) instead of an absolute measurement. We will not notate this explicitly because a mean shift is irrelevant to the subsequent noise analysis.

The ideal actuation coefficient for the l^{th} actuator is a weighted combination of the ideal sensor measurements (shown in equation (3)), where the weights $\{w_{k,l}\}$ are determined by the actuation strategy as described in section 2.

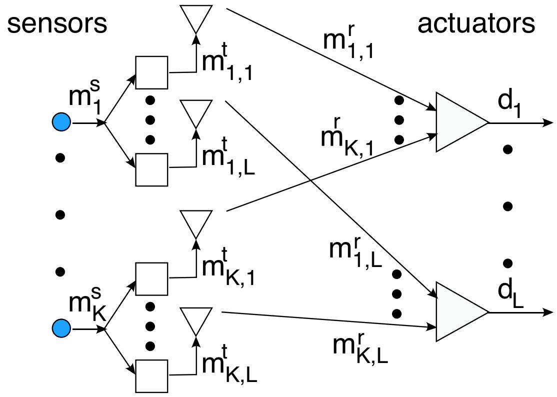

Figure 1: A schematic of the communications pathways in the WSAN model. The k^{th} sensor makes a noisy reading m_k^s. The quantized measurement $m_{k,l}^t$ with $B_{k,l}$ bits of precision is sent to the actuator l^{th} actuator. After incurring bit errors in the wireless channel, the measurements $m_{k,l}^r$ are received by the actuator and used to form the approximate actuation coefficient $\hat{d}_l$.

However, each sensor must quantize its measurement and send it with some finite precision along a noisy communications channel to each actuator. An inherent rate-distortion tradeoff exists for a single actuator coefficient: higher fidelity in the estimated actuator coefficient requires higher precision in the received measurements, but sending more bits expends more energy. The optimal power scheduling problem results from realizing that not all sensor coefficients are equally reliable or important to a given actuator, and the variable distances among the nodes creates inhomogeneity in the cost of each communication link. The fidelity of each communicated message can be chosen through using a variable number of bits (including zero) on each communication path to optimally use the system resources. To set up the optimal power scheduling problem, we must have both a model for the communications scheme that calculates the relative transmission power needed to communicate each bit and a model of the distortion that calculates the relative error reduction for each bit of precision.

3.1 Communication energy model

We first consider the transmission energy required to communicate the measurement m_k^s from sensor k to actuator l. Following the previous work of Cui, et. al [7,8] and Xiao, et. al [25], we assume that sensors communicate to an actuator over a wireless link using an MQAM strategy and following a common multiple access scheme (e.g., TDMA). If sensor k sends $B_{k,l}$ bits of precision about its measurement to actuator l, the MQAM scheme will have a constellation with $2^{B_{k,l}}$ symbols. We also assume an AWGN wireless channel, and that the transmission scheme is designed with a fixed bit error probability of p.

We expect that nodes using an MQAM scheme will incur a transmission energy proportional to the number of symbols being transmitted. In previous work [7,8,25], an upper bound on the transmission energy is estimated to be

$$P_{k,l} \leq \kappa d_{k,l}^{\alpha} \left(2^{B_{k,l}} - 1\right) \ln\left(\frac{2}{p}\right),$$

where $d_{k,l}$ is the distance between the sensor and actuator, α is the fading coefficient of the transmission medium, and κ is a constant depending on several aspects of the communication environment (e.g., channel SNR, antenna gain, coding gain, etc.). This expression applies for coded or uncoded MQAM, though κ will change when error correcting codes are used. We take κ to be a constant for all communication links in the system and we estimate transmission power for a measurement by the upper bound

$$P_{k,l} \propto d_{k,l}^{\alpha} \left(2^{B_{k,l}} - 1\right).$$

If the communications scenario is inhomogeneous, additional link-dependent constants can be introduced without significant difficulty. While the constants being left out of this expression are critical for determining the exact amount of transmission energy required, they are not necessary for comparing the relative effectiveness of two competing resource allocation strategies.

3.2 Communication distortion model

3.2.1 Quantization noise

Computation and communication systems must represent measurements in terms of a finite precision bitstream. Ideally, a real-valued measurement in the range $m_k^s \in [-A, A]$ is represented by an infinitely long bitstream $\{b_j\}, b_j \in \{0,1\}$,

$$m_k^s = 2A\left(\sum_{j=1}^{\infty} b_j 2^{-j}\right) - A.$$

The k^{th} sensor quantizes its measurement for transmission to the l^{th} actuator by truncating this series at $B_{k,l}$ bits. The resulting quantization error thus creates an additive noise term

$$m_{k,l}^t = m_k^s + n_{k,l}^q.$$

Assuming m_k^s is uniformly distributed over $[-A, A]$ and uniform quantization is employed, it is straightforward to calculate the first two moments of the quantization noise

$$\mathcal{E}\left[n_{k,l}^q\right] = 0, \qquad \mathcal{E}\left[\left(n_{k,l}^q\right)^2\right] = \frac{\left(\Delta_{k,l}\right)^2}{12} = \frac{A^2}{3} 2^{-2B_{k,l}}.$$

The uniform distribution of the measurements and the specific form of the quantizer are not critical for the setup of the optimization problem established in section 3.3. If a given application has a more appropriate model, the variance of that quantization noise can be used instead. Note also that although the same measurement is being communicated from sensor k to every actuator, the quantization noise term is different for each of those transmissions. In fact, some of those transmission link may use zero bits (i.e., no transmission actually occurs) when the transmission would be particularly costly relative to the benefit of transmitting the measurement.

3.2.2 Transmission noise

Ideally the quantized measurement would reach the destination exactly as it was transmitted. However, communication with finite power on a wireless channel will inevitably suffer errors. When sensor k transmits its quantized measurement $m_{k,l}^t$ to actuator l, the received measurement is denoted $m_{k,l}^r$. We can similarly denote the inaccuracy caused by transmission errors as an additive noise term

$$m_{k,l}^r = m_{k,l}^t + n_{k,l}^t.$$

To send its quantized measurement to actuator l, sensor k transmits the bit sequence $\{b_j\}_{j=1}^{B_{k,l}}$ along a noisy channel. The received bit sequence $\{b_j^r\}_{j=1}^{B_{k,l}}, b_j^r \in \{0,1\}$ is a Bernoulli random variable with the conditional distribution based on the bit error rate

$$P\left[b_j^r = b_j | b_j\right] = 1 - p$$
$$P\left[b_j^r = (b_j \oplus 1) | b_j\right] = p,$$

where $\oplus$ denotes modulo 2 addition.

A uniform distribution of m_k^s over the dynamic range translates to an assumption that b_j is a Bernoulli random variable with probability 1/2. Consequently, the expectation of the transmission noise also zero, $\mathcal{E}\left[n_{k,l}^t\right] = 0$. The variance of the transmission noise can be calculated as:

$$\mathcal{E}\left[\left(n_{k,l}^t\right)^2\right] = \frac{4pA^2}{3}\left(1 - 2^{-2B_{k,l}}\right)$$

3.3 Optimal bit allocation

As mentioned earlier, the power scheduling problem is fundamentally a rate-distortion tradeoff. In one form of that problem, we desire to meet a specified distortion criteria while expending the minimal amount of transmission energy across the WSAN. The vehicle for making this tradeoff is to assign the bit allocations on each sensor-actuator link $\{B_{k,l}\}$ to make optimal use of the power resources. As in [25], we will also minimize the ℓ^2 norm of the total power as a compromise measure between minimizing the total power (the ℓ^1 norm) and the maximum power (the ℓ^∞ norm).

We consider for now a single actuator coefficient d_l, which is approximated at the actuator using the received measurements

$$\hat{d}_l = \sum_{k\in K} w_{k,l} m_{k,l}^r = \sum_{k\in K} w_{k,l}\left(m_k + n_{k,l}\right),$$

where $n_{k,l} = n_k^s + n_{k,l}^q + n_{k,l}^t$ represents the total distortion from the ideal sensor measurement. As stated earlier, the results from [24] indicate that reducing the error in the actuator coefficients causes a proportional decrease in the upper bound of the error on the actuation field. Therefore, we will consider only the fidelity of a single actuation coefficient here, which can be used to bound the fidelity of the total actuation field. If desired, the resource allocation problem outlined here can be applied jointly to all of the actuator coefficients.

Making the typical assumption that the noise sources due to the sensor, quantization and bit errors are independent, we calculate the variance of the total noise source as the sum of the component variances

$$\mathcal{E}\left[(n_{k,l})^2\right] = \sigma_{k,l}^2$$
$$= \sigma_k^2 + \frac{A^2}{3}2^{-2B_{k,l}} + \frac{4pA^2}{3}\left(1 - 2^{-2B_{k,l}}\right). \quad (4)$$

Accounting for the multiplication of the weights necessary to calculate the actuation coefficient, the variance of $\hat{d}_l$ is given by

$$\mathcal{E}\left[\left(d_l - \hat{d}_l\right)^2\right] = \sum_{k\in K} w_{k,l}^2 \sigma_{k,l}^2$$

This distortion model clearly has some limitations. It is likely that the noise terms will not be completely independent, especially at very low bit rates. We will see in a simulation described later that the independence assumption does not have a large effect on the results for distortion ranges of interest. Also, the sensor noise is included in the distortion calculation even when no bits are used on the communication link and the measurement is not communicated at all. We will also see that this inaccuracy in the model does not have a large effect as long as the sensor noise is small compared to the dynamic range of the sensor so it is overwhelmed by the quantization error in the low bit-rate regime.

For individual actuator coefficients, the optimal resource allocation problem is expressed as an optimization of the bitrates $\{B_{k,l}\}$ over the field $\mathcal{X} = \mathbb{Z}_+$ subject to a distortion constraint of D_0,

$$\min_{B_{k,l}\in\mathcal{X}} \quad \sum_{k\in K} P_{k,l}^2$$

$$\text{s.t.} \quad \mathcal{E}\left[\left(d_l - \hat{d}_l\right)^2\right] \leq D_0.$$

Writing this optimization problem in terms of the energy and variance terms calculated from the communication and distortion models described above, the optimal resources allocation problem is fully stated as

$$\min_{B_{k,l}\in\mathcal{X}} \sum_{k\in K} d_{k,l}^{2\alpha}\left(2^{B_{k,l}} - 1\right)^2$$

$$\text{s.t.} \quad \sum_{k\in K} w_{k,l}^2\left(\sigma_k^2 + \frac{A^2}{3}\left(4p + (1-4p)\,2^{-2B_{k,l}}\right)\right) \leq D_0$$

This optimization is an integer program, whose solutions are are generally NP-hard. As in [25], we relax the space of the optimization problem to be $\mathcal{X} = \mathbb{R}_+$, where we obtain a convex program for which we can write an explicit solution. This relaxed problem will no longer provide an optimal solution, but by converting the optimal real-valued answers into integer values we will see that even introducing this suboptimality still allows for significant relative gains over uniform bit allocation methods. Using the ceiling function for this conversion ensures that the targeted distortion criteria D_0 is still met. However, using a less stringent rounding function may produce a better rate-distortion tradeoff curve even though the individual distortion criteria are not guaranteed to be met. In particular, using a rounding instead of a ceiling function generates many more links using zero bits, thereby possibly saving the overhead energy needed to power up the communications circuits on the sensor.

We use Lagrangian methods to solve this optimization for $B_{k,l} \in \mathbb{R}$. Making the substitution $y_{k,l} = 2^{B_{k,l}}$ for nota-

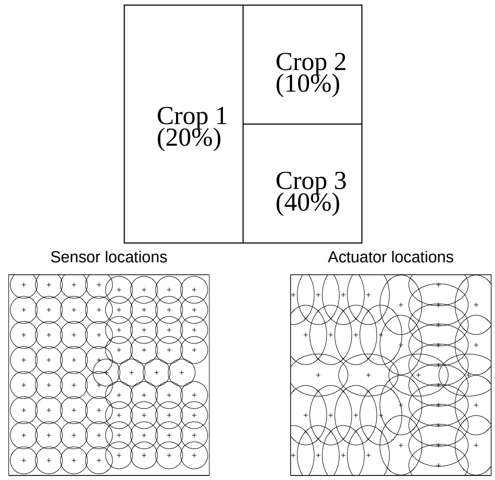

Figure 2: An example WSAN for a stylized agricultural irrigation application. A field 100m on each edge has three crops with varying desired soil moisture percentages. A collection of moisture sensors (bottom left) and irrigation actuators (bottom right) are deployed, with a single contour denoting where the influence field has mostly vanished.

tional convenience, the Lagrangian function is

$$L = \sum_{k \in K} d_{k,l}^{2\alpha} \left(y_{k,l} - 1\right)^2 + \left(\lambda \sum_{k \in K} w_{k,l}^2 \left(\sigma_k^2 + \frac{A^2}{3}\left(4p + (1-4p)\, y_{k,l}^{-2}\right)\right) - D_0\right),$$

where λ is the Lagrange multiplier constant. Differentiating L with respect to $y_{k,l}$ and setting it equal to zero we see that the optimal values of $y_{k,l}$ are those that meet the constraints and are a solution of the quartic equation

$$y_{k,l}^3 \left(y_{k,l} - 1\right) - \lambda \left(\frac{w_{k,l}^2 A^2 (1-4p)}{3 d_{k,l}^{2\alpha}}\right) = 0.$$

While analytic solutions to the quartic equation are cumbersome, numerical methods can easily solve this equation for specific numeric values. There are four roots to this quartic equation, only one of which is real and positive so as to be a viable solution to our optimization problem. Moreover, we find that the real and positive solutions for $y_{k,l}$ are strictly increasing in λ, which suggests an optimization strategy similar to the well-known water-filling solution problem for multi-channel communications [5]. We can increase λ to increase the number of bits allocated to each channel until the distortion criteria D_0 is met with equality. We see that communication links with measurements more critical to the actuation coefficient (measured by $w_{k,l}^2$) will be assigned bits more quickly, and communication links expending more transmission energy (measured by $d_{k,l}^{2\alpha}$) will be penalized and assigned bits more slowly.

4. NUMERICAL RESULTS

To demonstrate the utility of the power scheduling problem described in section 3.3, we examine a stylized WSAN setup

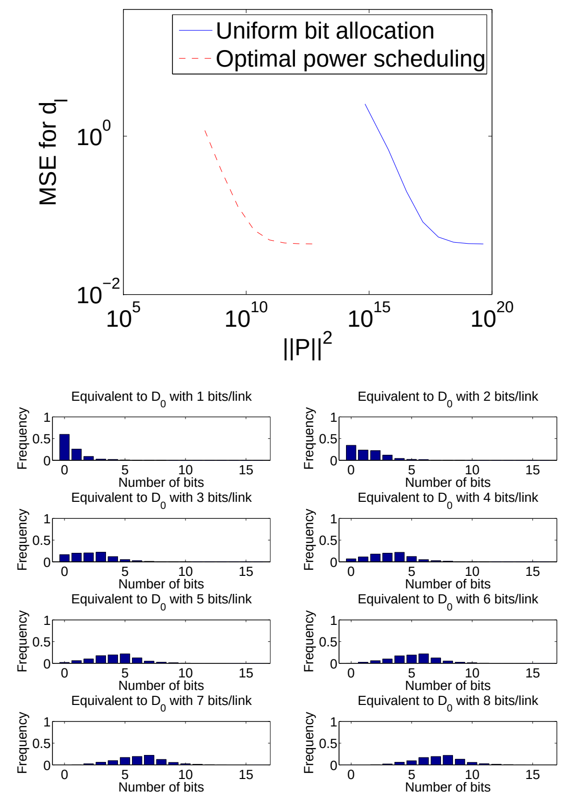

Figure 3: The rate-distortion curve comparing MSE estimation error for a single actuation coefficient d_l and the total network power needed to achieve that distortion are shown in the top plot. The bit allocations resulting from the optimal power scheduling algorithm achieve the same distortion using significantly less total power than a uniform bit allocation strategy. The bottom plot indicates the distribution of bits on each communication link. Each pane corresponds to the distortion equivalent of a uniform bit allocation scheme. The bars indicate the fraction of links using that number of bits.

for agricultural irrigation. Figure 2 denotes the location of different crops with varying optimal soil moisture needs, as well as the location of 68 moisture sensors and 38 irrigation actuators within a 100m square field. The sensor receptive fields are modeled by a circular Gaussian function, corresponding to the Green's function for homogeneous diffusion medium with infinite boundaries [4]. The actuator influence fields are modeled by an elliptic Gaussian function, representing a shaped water delivery pattern. Denoting the ideal moisture levels depicted in figure 2 as x_0, the ideal sensor measurements are relative to the nominal measurement, $m_k = \langle x_0, s_k \rangle - \langle x, s_k \rangle$. One choice of actuation function to try and achieve the set point is identity $T = I$, making the weights $w_{k,l} = \langle \tilde{a}_l, \tilde{s}_k \rangle$. We use simulation parameters of $\alpha = 3.5$ and $p = 10^{-3}$. To focus on the noise aspect related to bit allocation, we also assume a common sensor SNR of $\sigma_k^2 = 10^{-4}$ for all k.

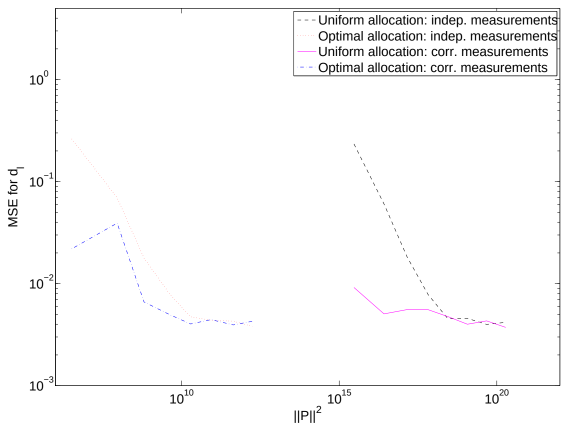

Figure 4: The rate-distortion curves calculated through simulated measurements. Measurements were simulated directly through independent random generation. Correlated measurements were also simulated by generating random environmental fields from a piecewise linear class and taking the resulting sensor measurements.

Focusing on a specific actuator coefficient d_l, we first examine the theoretical improvement of the power scheduling algorithm 3.3 over a uniform bit allocation scheme. Figure 3 shows the rate distortion curve produced for the expected MSE $\mathcal{E}\left[\left(d_l - \hat{d}_l\right)^2\right]$ when each link from a sensor to the l^{th} actuator used the same number of bits. For comparison, the optimal power scheduling algorithm of 3.3 was run to find optimal bit allocations and the resulting calculated rate-distortion curve is also shown in figure 3. We note that the bit allocation algorithm is able to reduce the network power consumption by several orders of magnitude for the same distortion over a uniform bit allocation strategy. The rate-distortion curves shown are averages over all actuator coefficients in the example setup. Figure 3 also indicates the distribution of bit assignments in the optimal power allocation scheme. Each pane corresponds to a distortion equivalent to a uniform bit allocation scheme, with the bars denoting the fraction of sensor-actuator communication links using that number of bits. The distribution of bit allocation levels across a range of values both above and below the corresponding uniform allocation indicates that power savings are being achieved not only through a reduction of bits but also through allocating those bits to the communication links that are the best combination of important and reliable. Also note that several links are allocated zero bits, meaning that no communication occurs over this link.

To verify our noise distortion model and the calculated rate-distortion curves shown above, we simulated the calculation of actuation coefficients from random measurements in the WSAN example application. Figure 4 shows the rate distortion curves calculated when independent random measurements were generated and used to calculate actuation coefficients (after adding sensor noise, quantizing, and simulating transmission errors). We see that the rate-distortion curves are essentially the same as those theoretically calculated and shown in figure 3. To explore the effect of sensor correlation on our distortion estimate, we performed the same experiment when the measurements are correlated. In this experiment we generating random environmental fields from a class of piecewise linear functions. The resulting correlated measurements were used to calculate the actuator coefficients in the same manner and the rate-distortion results are also shown. We see that these correlations do not play a significant effect on the distortions calculated from the theoretical model.

5. CONCLUSIONS AND FUTURE WORK

WSANs present an exciting opportunity to realize truly distributed in-network processing. The coordination among overlapping sensor and actuator nodes without a centralized controller presents, however, a daunting challenge. We have shown here that part of this coordination can be achieved through the optimal power scheduling algorithm problem. The solution of this convex optimization yields a resource allocation strategy that greatly reduces the network's power demands compared to a uniform resource allocation. While this optimization involves centralized information and processing, the nodes operate independently once they are given their optimal bit allocations for each communications link. The optimization strategy can therefore be done off-line and updated with a simple transmission to each sensor whenever the nodes are recalibrated. In this way, the necessary coordination is built into the pre-calculation and no centralized controller is needed to implement the control strategy.

There are many open areas for future work in WSAN information management. Our treatment of the optimal power scheduling problem assumed a simple network topology where every node can communicate wirelessly to every actuator. Our optimal power allocation scheme results in many of these communication paths not being used (i.e., they are allocated zero bits), meaning that in operation the sensor and actuator fields would not be fully connected. However, the optimization based on the assumption that each sensor node *could have* directly communicated with each actuator is clearly unrealistic when the spatial extent of the WSAN is large and multiple hops would be needed to make a connection between two nodes. Future work will be needed to incorporate these more complicated network topologies.

We envision several direct extensions to optimal power scheduling problem addressed in this paper. Recent research shows that altering the bit-interval durations to create lower bit error rates for the most significant bits can result in reduced error of the received measurement [16]. Also, in our previous work [24], we showed that sensors could also adaptively make decisions about communicating on each link based on their current measurement values. It may be beneficial to explore a hybrid strategy between these two approaches where the resource allocation is a combination of off-line computations and adaptive decision making. WSAN information processing strategies may also benefit from recent work and analysis in cooperative MIMO communications [6, 22] and distributed estimation using analog communication [9]. Finally, WSAN research will critically depend on methods for determining actuation laws based on physical environmental models and optimal control theory.

6. ACKNOWLEDGMENTS

This research was partially supported by the Texas Instruments DSP Leadership University program. The authors would like to thank Ilan Goodman for several helpful discussions during the course of this work.

7. REFERENCES

[1] I. Akyildiz and E. Stuntebeck. Wireless underground sensors networks: Research challenges. *Ad Hoc Networks*, 4:669–686, 2006.

[2] D. Blatt and A. Hero. Distributed maximum likelihood estimation in sensor networks. In *International Conference on Acoustics, Speech, and Signal Processing*, Montreal, Canada, May 2004.

[3] O. Christensen. *An Introduction to Frames and Riesz Bases*. Birkhauser, Boston, MA, 2002.

[4] R. Courant and D. Hilbert. *Methods of Mathematical Physics*, volume I. John Wiley & Sons, New York, 1937.

[5] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., New York, NY, 1991.

[6] S. Cui, A. Goldsmith, and A. Bahai. Energy-efficiency of MIMO and cooperative MIMO techniques in sensor networks. *IEEE Journal on Selected Areas in Communications*, 22(6):1089–1098, August 2004.

[7] S. Cui, A. Goldsmith, and A. Bahai. Joint modulation and multiple access optimization under energy constraints. In *IEEE Global Telecommunications Conference (GLOBECOM)*, volume 1, pages 151–155, 2004.

[8] S. Cui, A. Goldsmith, and A. Bahai. Energy-constrained modulation optimization. *IEEE Transactions on Wireless Communications*, 4(5):2349–2360, September 2005.

[9] S. Cui, J. Xiao, A. Goldsmith, Z. Luo, and V. Poor. Energy-efficient joint estimation in sensor networks: Analog vs. digital. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, Philadelhia, PA, March 2005.

[10] I. Daubechies. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1992.

[11] V. Delouille, R. Neelamani, and R. Baraniuk. The embedded triangles algorithm for distributed estimation in sensor networks. In *IEEE Workshop on Statistical Signal Processing*, St. Louis, MO, 2003.

[12] R. Duffin and A. Schaeffer. A class of nonharmonic Fourier series. *Transactions of the American Mathematical Society*, 72(2):341–366, March 1952.

[13] G. Franklin, J. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Addison-Wesley Publishing Company, Reading, MA, 1986.

[14] R. Glantz, H. Nudelman, and B. Waldrop. Linear integration of convergent visual inputs in an oculomotor reflex pathway. *Journal of Neurophysiology*, 52(6):1213–1225, December 1984.

[15] V. Goyal, J. Kovačević, and J. Kelner. Quantized frame expansions with erasures. *Applied and Computational Harmonic Analysis*, 10:203–233, 2001.

[16] D. Johnson and H. Rodriguez-Diaz. Optimizing physical layer data transmission for minimal signal distortion. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Hong Kong, April 2003.

[17] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, New York, NY, second edition, 1999.

[18] T. Marshall, J. Holmes, and C. Rose. *Soil Physics*. Cambridge University Press, New York, 1979.

[19] D. Neil. Compensatory eye movements. In D. Sandeman and H. Atwood, editors, *The Biology of Crustacea, Neural Integration and Behavior*, pages 133–163. Academic Press, New York, 1982.

[20] R. Nowak. Distributed EM algorithms for density estimation and clustering in sensor networks. *IEEE Transactions on Signal Processing*, 51:2245–2253, August 2003.

[21] R. Nowak, U. Mitra, and R. Willett. Estimating inhomogeneous fields using wireless sensor networks. *IEEE Journal on Selected Areas in Communications*, 22(6):999–1006, August 2004.

[22] H. Ochiai, P. Mitran, H. Poor, and V. Tarokh. Collaborative beamforming for distributed wireless ad hoc sensor networks. *IEEE Transactions on Signal Processing*, 53(11):4110–4124, November 2005.

[23] M. Rabbat and R. Nowak. Distributed optimization in sensor networks. In *International Symposium on Information Processing in Sensor Networks (IPSN)*, Berkeley, CA, April 2004.

[24] C. Rozell and D. Johnson. Evaluating local contributions to global performance in wireless sensor and actuator networks. *Lecture Notes in Computer Science*, 4026:1–16, 2006. *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS)*, San Francisco, CA, June 2006.

[25] J. Xiao, S. Cui, Z. Luo, and A. Goldsmith. Power scheduling of universal decentralized estimation in sensor networks. *IEEE Transactions on Signal Processing*, 54(2):413–422, February 2006.