Evaluating local contributions to global performance in wireless sensor and actuator networks

Christopher J. Rozell and Don H. Johnson⋆

Department of Electrical and Computer Engineering
Rice University, Houston, TX 77025-1892
{crozell,dhj}@rice.edu

Abstract. Wireless sensor networks are often studied with the goal of removing information from the network as efficiently as possible. However, when the application also includes an actuator network, it is advantageous to determine actions in-network. In such settings, optimizing the sensor node behavior with respect to sensor information fidelity does not necessarily translate into optimum behavior in terms of action fidelity. Inspired by neural systems, we present a model of a sensor and actuator network based on the vector space tools of *frame theory* that applies to applications analogous to reflex behaviors in biological systems. Our analysis yields bounds on both absolute and average actuation error that point directly to strategies for limiting sensor communication based not only on local measurements but also on a measure of how important each sensor-actuator link is to the fidelity of the total actuation output.

1 Introduction

Recent interest in wireless sensor networks has fueled a tremendous increase in the study of signal and information processing in distributed settings. Energy conservation is very important for most interesting applications, which generally translates into minimizing the communication among sensors to preserve both individual node power and total network throughput. Consequently, recent sensor network research has primarily focused on adapting well-known signal processing algorithms to distributed settings where individual nodes perform local computations to minimize the information passed to distant nodes (e.g., [1–3]).

The goal of many proposed sensor network algorithms has been to get the information *out* of the network (via a special node connected directly to a more traditional data network) with a good trade-off between fidelity and energy expended. 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. 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

⋆ This work was supported by the Texas Instruments Leadership University Program.

and a network of actuator nodes capable of affecting the environment. While one possible strategy summarizes information for a system outside the network to determine actuator behaviors, greater efficiency should be achieved by determining actions through in-network processing. A more subtle issue is that processing and communication strategies optimizing sensor data fidelity may not yield the best results when actuation performance fidelity is the desired metric.

While WSANs are often discussed, quantitative analysis of their performance has not received much attention. Existing work can be found in areas such as software development models for WSANs [4] and heuristic algorithms for resource competition based on market models [5]. Other recent work [6] has used techniques from causal inference to evaluate specific actuation strategies. Most relevant is the recent work of Lemmon et al. [7] analyzing distributed control systems while considering the underlying communication network. A control system approach is certainly appropriate for some WSAN application models, but may use more communication resources (especially from actuators to sensors) and may require the sensors and actuators to operate in the same signal space.

Merging sensed information directly into actions without centralizing the information and decision making has rarely been considered in man-made systems. Fortunately, we have examples from biology that demonstrate the effectiveness of this strategy. 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. This shrewd strategy avoids creating a single point of vulnerability, so the system can function in the presence of isolated failures.

In neural systems, two types of behaviors exist, depending on whether there is "thinking" involved, which we call *conscious* and *reflex* behaviors. In conscious behavior, biological systems gather sensory information, make inferences from that information about the structure of their environment, and generate actions based on that inferred structure. In reflex behavior, a sensed stimulus directly generates an involuntary and stereotyped action in the peripheral nervous system before the brain is even aware of the stimulus [8]. An obvious example of a reflex behavior is the knee-jerk reaction achieved by a doctor's well-placed tap below the kneecap. A more subtle example is the eye position correction that allows our vision to stay focused on an object even when our head is moving.

WSAN applications have an analogous division, which we call *object-based* and *measurement-based* network tasks. For example, the canonical target tracking scenario is an object-based task because it involves using sensory measurements to infer information about objects in the environment. On the other hand, an application such as agricultural irrigation is a measurement-based task because sensor measurements directly contain all the necessary information — there is no underlying environmental object to try and infer. In this work we consider models of measurement-based WSAN applications. While measurement-based systems are simpler and possibly more limited than object-based systems, they provide an entry point for analyzing and designing WSAN algorithms.

WSANs are complex systems with many interacting layers of operation. There are significant communication and networking challenges in these systems that are the focus of current research efforts. While the biological reflex systems described earlier do not appear to adaptively change their communication strategy on short time scales, the nature of wireless networking may necessitate dynamic decisions to employ different communication strategies based on current network conditions. Networking strategies to limit communication in the system must weigh the cost of executing individual communication links against the detrimental effect of performing suboptimal information processing. The role of our present research is to analyze a distributed WSAN model for a broad class of applications. We want to determine the optimal information processing strategy and to quantify the effects of suboptimal strategies resulting from eliminating communication links. As a simple starting place for our analysis, we will use vector space methods to model sensors and actuators, leveraging the notion of *frame theory* to analyze systems of nodes with overlapping influence.

2 Sensors and actuators

As an example reflex behavior that will shape our thinking about WSANs, we consider the crayfish visual system. The crayfish has a dorsal light reflex [9] where light movement in the visual field elicits predictable reflex movement in the eyestalk that attempts to keep a constant orientation of the visual field. The main visual representation (in neurons called "sustaining fibers") is comprised of sensory elements that sum light activity in overlapping spatial regions. All of the information available to the creature about the light stimulus is contained in this collection of sustaining fiber responses.

The crayfish eyestalk movement is controlled by a set of motorneurons, which send signals to several small muscles. Each muscle generates 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 in each motorneuron is determined directly from a processed combination of some sustaining fiber inputs. Though all of the motorneurons have to be coordinated to produce the desired total action, their distributed individual responses are generated directly from the distributed sustaining fiber 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 [10].

Our WSAN model will follow the principles seen in this example from the crayfish. Though the constraints facing biological systems are different from the constraints imposed by wireless networking, neural systems must also be very resource efficient and try to minimize communication (each neural signal generated means expending more metabolic energy). Biological systems must have solutions that do a good job (some would even argue optimal) at trading-off performance and efficiency, and we use them as a rough guide.

In our 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. We start with the simplest scenario where only this direct sensor-to-actuator communication is allowed. By eliminating inter-sensor and inter-actuator communication, we also eliminate the communication overhead necessary for such a scenario. It may be possible to improve system performance by allowing additional communication and cooperation, depending on the specific networking model and communication costs involved.

A major goal in any information processing strategy for WSANs is retaining good performance in the total actuation while reducing the communication burden from the sensors to the actuators. To analyze the performance of a WSAN under different design decisions, we use mathematical models based in the familiar tools and terminology of vector spaces.

2.1 Vector space models of sensors and actuators

Sensor network models often begin with a collection of sensors distributed over a 2-D spatial field limited to the spatial domain W (e.g., $W = [0,1]^2$). 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 $x(w)$, where $w \in W$ is a vector indicating location in the field.

Sensor measurement models often consist of averaging the stimulus field over non-overlapping spatial regions surrounding each sensor [11]. We generalize that notion by representing each sensor by a receptive field $s_k(w)$ over W that performs a weighted average over a spatial region. The sensor receptive fields are defined by the physics of the devices and could indicate sensors that are directional or have varying sensitivity over a region. Sensor measurements of the field are therefore given by

$$m_k = \int_W x(w) s_k(w) dw. \tag{1}$$

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. The measurement form given in (1) includes the special case of sensors averaging the field over disjoint local regions.

Recasting (1), the sensor measurements can be written as an inner product over the field W, $m_k = \langle x, s_k \rangle$. This vector space view of the sensor measurements indicates that with no further processing the measurements can represent any stimulus signal in the space $\mathcal{H}_x = \operatorname{span}(\{s_k\})$. 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 W. 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 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 W, denoted by a function $a_l(w)$. An actuator's influence field depends on the physics of the specific problem, and again may indicate actuators that are directional or have varying influence over a region.

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 W is $y = \sum_{l \in L} d_l a_l$, where, for simplicity (and to emphasize the vector space view), we drop the explicit notation of spatial location $w \in W$ from the actuator influence function $a_l(w)$ and the total actuation field $y(w)$. The collection of actuators can therefore cause any actuation field y in the space $\mathcal{H}_y = \mathrm{span}\left(\{a_l\}\right)$. 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 here 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 W. 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 functions spaces, and using general vector space definitions allows us to connect sensed inputs to actuation 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 on 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 sensor and actuator functions.

2.2 Frame theory

In section 2.1 we described the sensor measurement process as a projection of a stimulus field onto a collection of sensor representation functions. Similarly, we described actuators generating an effect as a weighted sum of individual actuator representation functions. In both the collections of sensors and actuators, the basic functions form a representation for a signal space ($\mathcal{H}_x$ and $\mathcal{H}_y$, respectively). The notion of representing a signal in terms of a collection of orthonormal basis (ONB) vectors is one of the most fundamental ideas in signal processing. Though the situation here is more complicated than an ONB, the collections of sensors and actuators are vectors that form a similar representation for their associated signal spaces. In this section, we will consider a general collection of vectors $\{\phi_j\}$

indexed over J. Fundamental results about this generic collection of vectors will be applied to the sensor and actuator representations in section 3.

An orthonormal basis has the property that any energy represented by the projection onto one vector will not be present in the projections onto any other vectors. As a consequence, reconstructing the signal from the projections is trivial; the projection coefficients simply weight the same vectors in the reconstruction. However, in general, collections of sensor receptive fields and actuator influence fields will not be orthogonal. In fact, in the most general case, these collections of functions may be linearly dependent and no longer form a basis.

A collection of M vectors $\{\phi_j\}$ forms a *frame* [12] for $\mathcal{H}$ if there exist constants $0 < A \leq B < \infty$ so that Parseval's relation is bounded for any $x \in \mathcal{H}$,

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

In general, there will be more vectors than are necessary to represent $\mathcal{H}$ ($M > N$, where $N = \dim(\mathcal{H})$), meaning that the frame is redundant. When the frame vectors are normalized $\left\| \phi_j \right\|^2 = 1$ (which we assume here), the frame bounds measure the minimum and maximum redundancy of the system and satisfy $A \leq \frac{M}{N} \leq B$. Frames were originally introduced in 1952 in the context of nonharmonic Fourier series [13] and later played a key role in wavelet theory [14]. They have recently been used in many other areas, including filterbanks [15], image processing [16], communications [17], coding [18] and machine learning [19].

The frame condition given above guarantees that the analysis coefficients obtained from projecting a signal onto the frame vectors contains all of the information necessary to synthesize (or reconstruct) the signal. Mathematically, the analysis coefficients are generated through the frame analysis operator $\Phi : \mathcal{H} \to l^2$, which is given by $(\Phi x)_j = c_j = \langle \phi_j, x \rangle$. In vector notation, the collection of all analysis coefficients is given by $c = \Phi x$. For finite dimensional frames (as in practical systems), the operator Φ is a matrix multiplication.

The adjoint of the frame analysis operator is the frame synthesis operator, $\Phi' : l^2 \to \mathcal{H}$, given by $\Phi' c = \sum_{j \in J} c_j \phi_j$. Because of the dependency present between frame vectors, the same set of vectors cannot generally be used for both analysis and synthesis. Even though Φ' and Φ are inverse operations in an ONB, in general Φ will not have a unique inverse. Therefore, the usual reconstruction will not work, $x \neq \Phi' \Phi x = \sum_{j \in J} \langle x, \phi_j \rangle \phi_j$. Instead, the pseudoinverse operator $\Phi^* = \left(\Phi' \Phi \right)^{-1} \Phi'$ is used for reconstruction, $x = \Phi^* \Phi x = \left(\Phi' \Phi \right)^{-1} \sum_{j \in J} \langle x, \phi_j \rangle \phi_j$. Equivalently, we can view the reconstruction as using a different set of vectors $\{\widetilde{\phi}_j\}$ called the dual set, $x = \sum_{j \in J} \langle x, \phi_j \rangle \widetilde{\phi}_j$. While there are an infinite number of sets of dual vectors that will work, the canonical dual set is given by $\widetilde{\phi}_j = \left(\Phi' \Phi \right)^{-1} \phi_j$. These dual vectors are also a frame for $\mathcal{H}$, with lower and upper frame bounds $\left(\frac{1}{B}, \frac{1}{A} \right)$, respectively. Importantly, the frame and dual set are interchangeable in the reconstruction equation,

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

The frame bounds are related directly to the eigenstructure induced by the frame vectors: $A = \lambda_{\min}$ and $B = ||\Phi'\Phi|| = \lambda_{\max}$, where $\{\lambda_i\}$ are the eigenvalues of $(\Phi'\Phi)$. When a collection of vectors has frame bounds that are equal, $A = B = \frac{M}{N}$, it is called a *tight frame*. When a frame is tight, the dual vectors are simply rescaled versions of the frame vectors, $\widetilde{\phi}_j = \frac{1}{A}\phi_j$. A collection of vectors is an orthonormal basis if and only if it is a tight frame with $A = B = 1$.

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. When we apply frame theoretic models to the analysis of sensor and actuator networks, we want to know the impact of reducing communication costs by using approximate coefficients in the reconstruction.

Stated generally, we need to calculate a bound on the maximum error when a perturbation p_j is added to each frame coefficient c_j in the reconstruction, $\hat{x} = \sum_{j\in J}(c_j + p_j)\widetilde{\phi}_j$. Perturbations may include removing the coefficient from the reconstruction, $p_j = -(c_j)$. The error resulting from these perturbations is

$$||x - \hat{x}||^2 = \left|\left|\sum_{j\in J} p_j \widetilde{\phi}_j\right|\right|^2. \tag{2}$$

We recall that the dual set $\{\widetilde{\phi}_j\}$ is also a frame for $\mathcal{H}_x$, and we denote the analysis operator for the dual frame to be $\widetilde{\Phi}$. Note that the error signal recast in matrix notation is $(x - \hat{x}) = \widetilde{\Phi}'p$, where p is the perturbation vector $p = \left[p_1 p_2 \dots p_{|J|}\right]'$. Linear algebra can yield a bound on the error,

$$\left|\left|\widetilde{\Phi}'p\right|\right|^2 = \left|\langle p, \widetilde{\Phi}\widetilde{\Phi}'p\rangle\right| \leq \left|\left|\widetilde{\Phi}\widetilde{\Phi}'\right|\right| \cdot ||p||^2.$$

Note that because the singular values of $\widetilde{\Phi}$ are the square roots of the eigenvalues of both $\left(\widetilde{\Phi}\widetilde{\Phi}'\right)$ and $\left(\widetilde{\Phi}'\widetilde{\Phi}\right)$, it follows that $\left|\left|\widetilde{\Phi}\widetilde{\Phi}'\right|\right| = \left|\left|\widetilde{\Phi}'\widetilde{\Phi}\right|\right|$. Because the dual set is a frame for $\mathcal{H}$ with upper frame bound $\left(\frac{1}{A}\right)$ and because of the relationship between the eigenvalues of $\left(\widetilde{\Phi}'\widetilde{\Phi}\right)$ and the frame bounds, we can finally write a useful bound (alluded to in [20]) on the reconstruction error

$$||x - \hat{x}||^2 \leq \frac{||p||^2}{A}. \tag{3}$$

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 equation (3) is consistent with probabilistic robustness results when stochastic noise is added to frame coefficients [18].

3 Connecting sensors to actuators

Following our example of reflex behavior, actuators must generate activity using received sensors measurements without communicating with other actuators. The overlapping actuator influence fields 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 must coordinate behavior (without communication) to account for the the action field components covered by the other sensors.

3.1 Generating optimal actuation

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 $\{\tilde{a}_l\}$. Note that the dual sets $\{\tilde{s}_k\}$ and $\{\tilde{a}_l\}$ aren't realized directly in physical systems. For example, the sensor receptive field dual functions $\{\tilde{s}_k\}$ may have spatial characteristics that would be impossible to build into any type of real-world sensor.

To generate coordinated behavior in the actuator network, we must necessarily start with the ideal solution for generating actions. Each WSAN has an application specific goal that defines its existence. For example, a system might use sensed rainfall to order the diversion of floodwater or the delivery of irrigation to meet specified conditions. Though the actions necessary to achieve the goal depend on the specific observed stimulus, the goal itself is stimulus independent. To quantify this application goal, we assume that for any measured stimulus field x there is a mapping $T : \mathcal{H}_x \to \mathcal{H}_y$ that defines the ideal action field response, $y = Tx$. The mapping T would be determined as a design specification for the WSAN in advance. While it may be possible to reconfigure a WSAN to perform a different application (with a different goal) on long time scales, we assume that the goal (as quantified by T) stays fixed.

An ideal actuator network would have each node determine action coefficients $\{d_l\}$ to generate the optimal response $Tx = \sum_{l \in L} d_l a_l$. Drawing on the frame theory results from section 2.2, the coefficients weighting the action influence field vectors are given by the inner products between the action *dual* vectors and the action signal that we are trying to generate,

$$d_l = \langle \tilde{a}_l, Tx \rangle. \tag{4}$$

To determine the optimal action coefficients, consider first the reconstruction equation for the stimulus field based on the sensor measurements,

$$x = \sum_{k \in K} m_k \tilde{s}_k. \tag{5}$$

Substituting equation (5) into equation (4), the optimal action coefficients are

$$d_l = \langle \tilde{a}_l, T \sum_{k \in K} m_k \tilde{s}_k \rangle = \sum_{k \in K} m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle. \tag{6}$$

The conversion from sensor measurements $m = \left[m_1, m_2, \ldots, m_{|K|}\right]^{'}$ to actuator intensity coefficients $d = \left[d_1, d_2, \ldots, d_{|L|}\right]^{'}$ in matrix form is $d = Vm$, where

$$V = \begin{bmatrix} \tilde{a}_1^{'} T\tilde{s}_1 & \tilde{a}_1^{'} T\tilde{s}_2 \cdots & \tilde{a}_1^{'} T\tilde{s}_{|K|} \\ \tilde{a}_2^{'} T\tilde{s}_1 & \ddots & \vdots \\ \vdots & & \\ \tilde{a}_{|L|}^{'} T\tilde{s}_1 & \cdots & \tilde{a}_{|L|}^{'} T\tilde{s}_{|K|} \end{bmatrix}.$$

The expression in equation (6) (or equivalently the entries of V) illuminate the form of the actuator intensity coefficients necessary to generate the optimal total action Tx. Unfortunately, each 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 an 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 links would be necessary. While a portion of this burden could be reduced through broadcast communication, some sensor-to-actuator links may involve several communications in a multi-hop routing scheme. Any realistic networking scheme will have to eliminate some of these communication links based on their communication cost and their contribution to the total actuation performance. Intuitively, some sensor measurements will be more important than others in determining an actuators behavior. 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. Using the frame theory results presented in section 2.2 along with the vector space model of sensor and actuator networks, we have tools for analyzing the effects of eliminating communication links on the total actuation performance.

3.2 Limiting communication costs

Each entry of the matrix V indicates a communication link from a sensor to an actuator. Before blindly reducing communications, a networking scheme must know the importance of each possible communication. In a sensor network, performance is often judged by assessing the fidelity of the information removed from the network at representing the original sensor measurements (or the underlying stimulus field). However, the only performance metric of any consequence in a WSAN is the fidelity of the resulting total action.

To quantify the importance of individual communications, we must determine how the total actuation performance is affected when a communication is not executed. We quantify this notion of importance through the results described in equation (3). Consider the case where for actuator l, a subset of sensor nodes $E_l \subset K$ do not transmit their measurement coefficient to this actuator. Instead of optimal actuator intensity coefficients (see equation (6)), actuators form approximate intensity coefficients using the received sensor measurements

$$\hat{d}_l = \sum_{k \in (K \setminus E_l)} m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle. \tag{7}$$

The approximate actuator intensities generate a total action field approximating the desired optimal action Tx,

$$\hat{y} = \sum_{l \in L} \hat{d}_l a_l.$$

Generating a total action field with the approximate coefficients $\{\hat{d}_l\}$ is equivalent to performing a frame reconstruction with perturbed coefficients, as described in section 2.2. Subtly, the actuator frame vectors are performing synthesis, meaning that dual vectors (with lower frame bound $\frac{1}{B_a}$) are now the analysis set. Therefore, equation (3) relates the fidelity of the approximate actuator intensity coefficients to the fidelity of the resulting total action field,

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

Using equations (7) and (6), we can write the total action field error in terms of individual sensor coefficients *not* communicated to actuator nodes

$$||Tx - \hat{y}||^2 \leq B_a \sum_{l \in L} \left| \sum_{k \in E_l} m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle \right|^2 \tag{8}$$

$$\leq B_a \sum_{l \in L} \sum_{k \in E_l} \left| m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle \right|^2. \tag{9}$$

As we see in equation (9), the networking strategy for sensor node k can use the value of $|m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle|^2$ to quantify the maximum contribution it would make to the total action error by *not* communicating its measurement to actuator l. The bound in equation (9) can be used to set a threshold γ guaranteeing an absolute upper limit on the actuation error.

Importantly, the form of the error bound in equation (9) isolates each communication link as an independent term so that no communication overhead is required to determine the absolute worst actuation error that can be incurred by eliminating a communication link[1]. In applications where a WSAN must respond quickly to critical but rare events (e.g., a fire suppression system), an

[1] We are assuming that the setup phase of the WSAN has given nodes information about the relative locations of their neighboring nodes that can be used to calculate the necessary inner product.

absolute bound on the actuation error computed locally is probably appropriate. To ensure that the actuation error is within an absolute tolerance, the active communication links between sensors and actuators will necessarily change depending on the input signal. While this dynamic decision making doesn't impose a large computational burden on the sensor nodes, the underlying communications network must be able to handle large fluctuations in demand for resources.

Because the sensor and actuator fields overlap and form a frame (instead of an orthonormal basis), the contributions from two different sensor measurements to an actuator coefficient could, in effect, "cancel" each other. Because the error bound provided in equation (9) is expressly written in terms of local sensor node measurements, this bound favors a conservative interpretation rather than accounting for these interactions. Given a specific communication and networking scenario, it may or may not be advantageous to allow sensors to explicitly communicate to calculate a tighter error estimate (based on the original error expression in equation 2) and coordinate their communication accordingly. While the frame theoretic analysis paradigm introduced here would allow such an analysis, it would necessarily be specific to the application details (particularly the communication and networking scenario).

In many settings, designing around an absolute error constraint results in a system that is too conservative in its average behavior. To analyze the average actuation error one must assume a stochastic model for the measurements, such as assuming that the sensor measurements have zero mean ($\mathcal{E}\left[m\right] = 0$) and covariance matrix Γ_m. The covariance matrix Γ_m will be determined by a combination of the the sensor receptive field properties and the distribution assumed on x within the signal space $\mathcal{H}_x$. Only the first two moments of the distribution on m are relevant, so we need not assume Gaussian distributions.

Average WSAN performance is much easier to calculate if we recast equation (8) using matrix notation. We first need to write approximate actuator coefficients in equation (7) in terms of a perturbation of V, which captures the ideal transformation from sensor measurements to actuator coefficients. Let the approximate actuator coefficient be given by $\hat{d} = \left(V + \widetilde{V}\right) m$, where the matrix $\widetilde{V}$ is defined to remove inactive communication links:

$$\left(\widetilde{V}\right)_{k,l} = \begin{cases} -\left(\tilde{a}_1' T \tilde{s}_1\right) & \text{if } k \in E_l \\ 0 & \text{if } k \in (K \setminus E_l). \end{cases}$$

Incorporating this definition into equation (8) and taking the expectation of both sides lets us bound the average error

$$\mathcal{E}\left[||Tx - \hat{y}||^2\right] \leq B_a \mathrm{Tr}\left[\widetilde{V} \Gamma_m \widetilde{V}'\right], \tag{10}$$

where $\mathrm{Tr}\left[\cdot\right]$ is the trace operator.

A system designer could use equation (10) to characterize (on average) how important a communication link between a specific sensor and actuator pair is to generating the total actuation field. Using this information, a WSAN design could choose *a priori* which communication links between sensors and actuators will be active in the network. Such a scheme has the disadvantage that it may not react well to events that are large deviations from the usual behavior. The advantages to this type of non-adaptive communication scheme in a WSAN are that the communication resources are used more efficiently most of the time, the network can count on a limited communication burden for any stimulus field, and the real cost of executing individual communication links (through a possibly multi-hop network) can be easily integrated into generating an optimal strategy. Also, it is worth noting that the bound in equation (10) is tighter than the bound in equation (9) (because it is based directly on equation (8)), reflecting the fact that all of the communication links can be considered jointly when designing the system for average error performance.

4 An example WSAN system

As an illustrative example, consider a WSAN operating a fire suppression system in an office building with four research labs. Each lab contains expensive equipment, so there is a strong desire to localize the fire suppression to minimize water damage to adjacent labs. The building space is covered with a network of 21 temperature sensors (modeled with radially symmetric, exponentially-decaying receptive fields) and 13 actuators (modeled with an oriented and exponentially decaying influence field), all illustrated in Fig. 1. This WSAN has 273 possible communication links from the sensor nodes to actuator nodes. In this example we assume an equal communication cost for each link (i.e., we would like to use as few links as possible regardless of which links are in use).

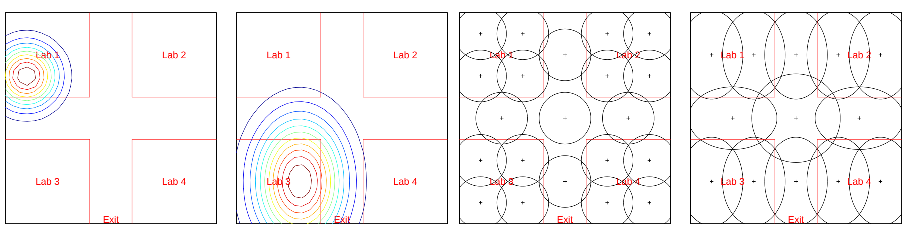

Fig. 1. Contour plot of example sensor *(Far left)* and actuator *(Middle left)* nodes. Layout and shape of the sensor *(Middle right)* and actuator *(Far right)* nodes.

We specified a function T mapping the temperature inputs to an imaginary desired fire suppression output. To illustrate that this mapping may be spatially varying, we note that fire activity in all labs will induce fire suppression activity along a path to the main exit. We used two sample temperature fields indicating a fire in different labs areas (shown in Fig. 2, along with optimal responses). As discussed in section 3, the quantity $|m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle|$ determines the importance of each communication link (sorted and plotted in Fig. 3 for these test signals). In these signals, a threshold of $\gamma = .2$ allows approximately 15 of the 273 possible communication links to be active, and $\gamma = .05$ allows approximately 40 active communication links. The resulting active communication links are shown in Fig. 3. Close examination of the connection diagrams shows that some communication choices are non-obvious; the most important sensor to a particular actuator is not always the one with heavily overlapping influence functions.

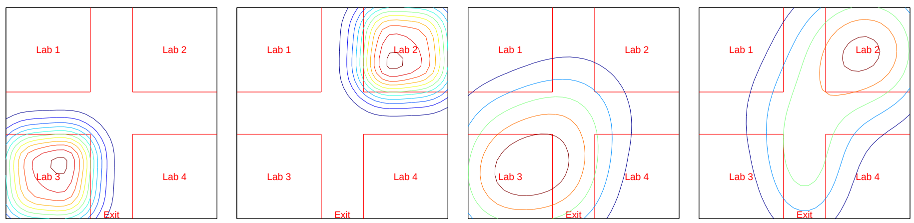

Fig. 2. Contour plots of sample temperature fields for test signal 1 indicating a fire in lab 3 *(Far left)* and test signal 2 indicating a fire in lab 2 *(Middle left)*. Contour plots of optimal actuation responses to the two test scenarios *(Middle right and far right, respectively)*. Different spatial response characteristics keep the main exits clear.

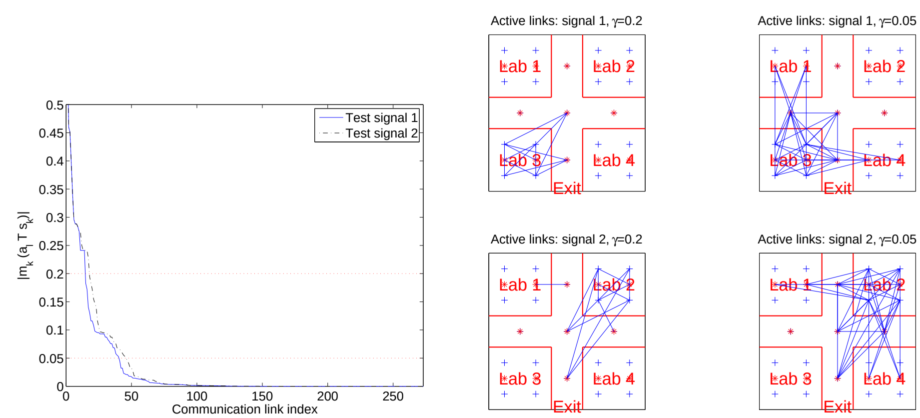

Fig. 3. *Left:* The importance measurements of each communication link ($|m_k \langle \tilde{a}_l, T\tilde{s}_k \rangle|$) are sorted and plotted for the two test signals. *Right:* Connection diagrams for the two test signals under the two thresholds in the example system. Sensor nodes are marked with a blue (+) and actuator nodes are marked with a red (*). Active connections from a sensor to an actuator are denoted by a blue line.

The actuation response is generated for both test signals using threshold values of $\gamma = .2$ and $\gamma = .05$, and the resulting total actuation fields are plotted in

Fig. 4. The reduced communication scheme based on the thresholds resulted in the number of active communication channels and associated percentage errors given in Table 1. The principles discussed in section 3 allow the WSAN to generate excellent approximations to the optimal actuation field by using local rules to activate only a fraction of the communication links. Interestingly, if we activate the same number of links using the more intuitive measure $|m_k \langle a_l, Ts_k \rangle|$, the resulting actuation error increases by roughly an order of magnitude.

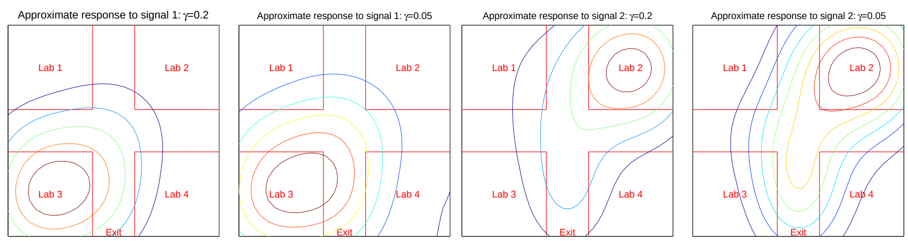

Fig. 4. Contour plots of actuation responses when using only a subset of possible communication links (determined by thresholding each link's importance to the total actuation). Approximate responses to test signal 1 are shown when using 14 and 40 communication links (*Far left and middle left*). Approximate responses to test signal 2 are shown when using 17 and 45 communication links (*Middle right and far right*).

Table 1. Results from the example WSAN fire suppression system

	$\gamma = .2$	$\gamma = .05$
Active links	14	40
Relative error	2.22%	0.04%

Test signal 1

	$\gamma = .2$	$\gamma = .05$
Active links	17	45
Relative error	2.46%	0.15%

Test signal 2

5 Conclusions and future work

WSANs are often discussed as a logical extension to sensor networks, but there is little research investigating sensor and actuator systems working in concert together. While algorithms that reduce communications and ensure data fidelity for sensor measurements are important for many applications, they are not the ultimate arbiter for obtaining good actuation performance. The total system must be designed and managed with the final actuation goal in mind. Our frame-theoretic WSAN model illustrates one strategy for taking such a holistic information management view with actuation fidelity as the relevant metric.

The analytic tools we present characterize the effect of eliminating an individual communication link between a sensor and an actuator, both in terms of absolute (for specific sensor measurements) and average actuation error. Choosing a networking strategy for eliminating communication links is both difficult and non-intuitive. While intuition would indicate that the relationship between the activation fields of a sensor and an actuator are the relevant quantity characterizing the importance of the communication between those two nodes, our work shows that it is the relationship between the mathematical *duals* of the activation fields that captures this inherent importance. It is through these dual functions that the relationship of the whole sensor network to the whole actuator network can be accounted for in local communications between pairs of nodes. Characterizing the importance of individual communication links to the overall goal points directly to how a networking strategy could weigh the costs and benefits of each communication link to achieve the desired balance between performance and energy efficiency. The value of our analysis is highlighted in an example WSAN system where link activations based on the sensor and actuator duals performed an order of magnitude better than activations based on the simple overlap of the sensor and actuator receptive field functions.

Today we are only seeing the beginning of work in information management in WSANs. In this work, we have given explicit upper bounds on actuation error that can be determined locally with no cooperation between the sensors. We have also indicated how this analysis framework could be used in a specific application and networking scenario to investigate the benefits of allowing local sensor coordinate their communications to an actuator. Finally, we have also derived analogous average error bounds that could be used to design static networking strategies for applications where that approach is more appropriate.

We are currently working on many extensions to this work. We have considered the case where perfect (analog) coefficients are sent on active communication links. While real systems would have to use quantized coefficients, we believe that typical quantization schemes would have only a second order effect relative to other actions taken to limit communication (such as eliminating communication links). However, it is more interesting to consider a variable rate communication scheme where some links could send coefficients with variable fidelity. Such variable rate schemes could be particularly interesting as we consider incorporating information about the variable networking costs of different communication links. We are working to more tightly integrate the costs and benefits of individual communication links to find optimal strategies for determining which links to activate dynamically and with minimal overhead.

Finally, our system model considers a single actuation response to a set of sensor measurements. This is something of an open-loop system because the sensors don't necessarily receive any direct feedback from the actuators. This generality is appealing in many senses; our model allows sensors and actuators to live in separate signal spaces and it may be possible that actuation is not directly observable by the sensors. However, in many practical applications, future sensor measurements will be affected by actuator behavior even when they operate in different signal spaces (e.g., fire suppression actions will reduce the temperature measured by sensors). We are working on methods for extending this work to consider the dynamic properties of such an implicit feedback system.

References

1. Nowak, R.: Distributed EM algorithms for density estimation and clustering in sensor networks. IEEE Transactions on Signal Proc. **51** (2003) 2245–2253
2. Rabbat, M., Nowak, R.: Distributed optimization in sensor networks. In: Intl. Symposium on Information Proc. in Sensor Networks (IPSN), Berkeley, CA (2004)
3. Blatt, D., Hero, A.: Distributed maximum likelihood estimation in sensor networks. In: Intl. Conf. on Acoustics, Speech, and Signal Proc., Montreal, Canada (2004)
4. Liu, J., Chu, M., Liu, J., Reich, J., Zhao, F.: State-centric programming for sensor and actuator network systems. IEEE Pervasive Computing Magazine **2**(4) (2003)
5. Gerkey, B., Mataric, M.: A market-based formulation of sensor-actuator network coordination. In: AAAI Spring Symp. on Intel. Embed. and Dist. Sys. (2002)
6. Coates, M.: Evaluating causal relationships in wireless sensor/actuator networks. In: Intl. Conf. on Acoustics, Speech, and Signal Proc., Philadelphia, PA (2005)
7. Lemmon, M.D., Ling, Q., Sun, Y.: Overload management in sensor-actuator networks used for spatially-distributed control systems. In: Proceedings of the ACM Sensys Conference. (2003)
8. Kandel, E., Schwartz, J., Jessell, T.: Principles of Neural Science. Third edn. Appleton & Lange, Norwalk, CT (1991)
9. Neil, D.: Compensatory eye movements. In Sandeman, D., Atwood, H., eds.: The Biology of Crustacea, Neural Integration and Behavior. Academic Press, New York (1982) 133–163
10. Glantz, R., Nudelman, H., Waldrop, B.: Linear integration of convergent visual inputs in an oculomotor reflex pathway. J. of Neurophys. **52**(6) (1984) 1213–1225
11. Nowak, R., Mitra, U., Willett, R.: Estimating inhomogeneous fields using wireless sensor networks. IEEE J. on Selected Areas in Comm. **22**(6) (2004) 999–1006
12. Christensen, O.: An Introduction to Frames and Riesz Bases. Birkhauser, Boston, MA (2002)
13. Duffin, R., Schaeffer, A.: A class of nonharmonic Fourier series. Transactions of the American Mathematical Society **72**(2) (1952) 341–366
14. Daubechies, I.: Ten Lectures on Wavelets. Society for Industrial and Applied Mathematics, Philadelphia, PA (1992)
15. Bolcskei, H., Hlawatsch, F., Feichtinger, H.: Frame-theoretic analysis of oversampled filter banks. IEEE Transactions on Signal Proc. **46**(12) (1998) 3256–3268
16. Candès, E., Donoho, D.: New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities. Communications on Pure and Applied Mathematics **57**(2) (2004) 219–266
17. Strohmer, T., Heath Jr., R.: Grassmannian frames with applications to coding and communcations. Applied and Comp. Harmonic Analysis **14**(3) (2003) 257–275
18. Goyal, V., Kovačević, J., Kelner, J.: Quantized frame expansions with erasures. Applied and Computational Harmonic Analysis **10** (2001) 203–233
19. Gao, J., Harris, C., Gunn, S.: On a class of support vector kernels based on frames in function hilbert spaces. Neural Computation **13** (2001) 1975–1994
20. Balan, R., Casazza, P., Heil, C., Landau, Z.: Density, overcompleteness, and localization of frames, I. Theory. Preprint (2005)