Design of Trustworthy Fielded Sensor Networks

Gregory J. Pottie
Electrical Engineering Department
UCLA Box 951594
Los Angeles CA 90095-1594

ABSTRACT

Sensor networks are finding application as monitoring systems and as tools in the study of complex natural systems. In either situation, the primary goal is computation of some inference from the observations and available models. From this basic problem flows a broad set of practical and theoretical issues, among them assurance of data integrity, sufficiency of data to support the inferences made concerning models/hypotheses, deployment density, and what tools and hardware are required not just to take observations but enable a community of non-engineers to participate in and adapt a sequence of experiments as new observations are obtained. The resulting constraints for designing systems for such purposes are quite different from those commonly assumed in the infancy of wireless sensor network research, and even now in much ongoing systems research. We describe these constraints in light of experience in deploying sensor networks in support of scientific study at the Center for Embedded Networked Sensors (CENS).

Keywords: Sensor networks, modeling

I. INTRODUCTION

The promise of sensor networks is that they will provide an ability to enable remote interactions with the physical world at unprecedented physical scales, with a broad set of sensors, and at low cost[1]. A great deal remains to be done in order for this vision to be realized. Much of the sensor network community is working on problems with some variant of the following constraints: nodes are to be low cost, and thus large in number but with severe limitations on some combination of energy, sensor complexity, computational ability, memory capacity or communications rate. Dealing with severe constraints such as these is potentially tractable if the model of the physical phenomenon of interest is well-established, the purpose of the network is clear and the sensors themselves are reliable (trustworthy). This can be true for example in surveillance applications where accelerometers are the primary sensor; whereas the sensitivity of the individual sensors may be highly variable due to variations in coupling to the medium, with high density they may still reliably detect intruders. On the other hand, such variability would be extremely difficult to deal with in attempting to produce a scientific study of the vibrations due to pedestrians, absent some explicit calibration procedure and infrastructure. The pre-existence of a study or sequence of studies of exactly this sort is required for the surveillance application to be successful, or else the required node density cannot be computed. For reasons to be discussed, this investigatory phase is unlikely to be successful using highly constrained nodes. Moreover, there is implicitly some mobility mechanism to emplace the sensors, a means for locating them, and a model for sensor variability in the environment of interest. In some applications, these are immediately available while in others they are the dominant barriers to success.

Within the context of sensor networks as a means to draw useful inferences, there are many different optimizations that flow from (a) how broad a set of inferences is required (b) how well known is the physical model and (c) what kind of hardware is available and how well modeled it is. For example, in many military/security applications the sensing/system model can be very well known and reliable, but the inference problem can still be very difficult (e.g, establishment of human intent). This may require a long train of experimentation before the appropriate type of network, including sensors and processing, can be selected. In an environmental study by contrast, there is usually massive uncertainty in models at all levels,

Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security
and Homeland Defense VII, edited by Edward M. Carapezza,
Proc. of SPIE Vol. 6943, 69430U, (2008) · 0277-786X/08/$18 · doi: 10.1117/12.786264

driving system architectures towards intensive processing and standard communications networking so that effort can be devoted to ensuring measurements are trustworthy and also to the sequence of experiments required to produce inferences among interconnected sets of models. If the uncertainty is such that mobile elements are required (e.g., robots or graduate students), then their cost will often imply that higher end sensors, processors, and communications infrastructure may add only marginally to the total expense, in turn presenting new opportunities for exploration of heterogeneous systems.

Three examples of constraint sets and a sampling of the research problems associated with them are:

- Classic sensor networks: resource-starved nodes, with simple and potentially unreliable sensors, but which are well modeled; research problems include low-complexity security, efficient query/response, reliable multi-hop communications, consensus mechanisms for measurement reliability, cooperation for source/node localization, behavior in limits of single bit sensors and/or communication messages, robustness implications of flat networks, interplay of in-network processing with communications, dynamic re-tasking, low-footprint and robust databases
- Hybrid ensembles: resources and sensors flow from application demands, resulting in heterogeneous networks that may include mobile elements and communications infrastructure, but where there is high model or application requirements uncertainty. Problems include multi-scale sensing, data integrity, model creation and extension, fusion and inference from diverse sensing, adaptive sampling, security/privacy with unattended nodes, optimization with asymmetric resources/capabilities, serving multiple simultaneous users.
- Actuated systems: Mobile nodes bring questions of control theory to the fore. In general having the network support some control task such as factory production, industrial process control or environmental remediation is a natural extension that will nonetheless put greater pressure on the network to respond both quickly and reliably to user requests. In turn, more resources may go into infrastructure. Problems include new control abstractions to deal with the different time scales and reliability requirements for different functions of the network (e.g., tight for aircraft maneuvers, looser for the data the aircraft is gathering), reliable operation under intermittent communication, coordination of groups of sensors, actuators, processors and communication nodes including re-tasking

The greater the uncertainty concerning the phenomenon of interest, the greater the likelihood that no single experiment will suffice to answer the questions concerning how to model it. Similarly, it becomes very difficult to assure the reliability of the inferences drawn from a deployed sensor security system. Rather, a system must be devised so that a sequence of experiments can be carried out, with interactions between the domain experts and the system, resulting in re-tasking of elements, insertion of new equipment, or even complete redeployments. The objective of some of these iterations may be to characterize the equipment. For example, all biochemical sensors become fouled in the environment over time, and laboratory calibration will often be completely inadequate given the many confounding factors in natural environments. Electronic faults will occur, and one might not be able to afford the most expensive and reliable measurement equipment (and its associated procedures) at the desired density. This iterated procedure is described in section II. The data integrity problem consists of how to ensure reliable inferences can be made, even given that some elements are unreliable, and is the subject of section III. Given the state of biochemical sensors, it is the dominant design constraint in long-term sensor network deployments for ecological studies or providing alerts to the presence of bio-chemical species. We briefly also describe tools being developed in support of reliable experimentation. In section IV we address some issues specific to detection and localization problems, and in particular the value of cooperation among sensor nodes. Section V presents our conclusions.

II. THE DESIGN CYCLE

The main sensor network problem is the efficient computation of $P(X|Z)$, where X is the desired inference, Z is the set of observations, and the probability distribution P results from some model and set of observations. For example, Z may consist of data from multiple sensors yielding range information and the desired inference X may be the estimated position of one or more sources to within some desired accuracy. The computation of P requires knowledge of the background noise, signal propagation model and some characteristics of the sources, as well as some minimum number of independent sensors. To some extent, a larger number of sensors will assist in reducing the location uncertainty, but this will not overcome fundamental ignorance of the propagation conditions. In a more complicated version of the $P(X|Z)$ problem, the observations may be indirect in that they are themselves the product of a model, and the desired inference X may be a determination of which among several hypothesized physical models is most likely. It may be noted that sensor measurements are in fact always the result of some processing that employs a model of how the transducer operates. This model is usually limited to some particular range of physical conditions. But beyond this, the phenomenon of interest may not be directly observable with the available sensors, and so some phenomenon that is associated via some model is monitored instead. Thus for example in assessing whether to initiate the harvest, the color of a grape might be estimated using a digital camera and appropriate signal processing, given the impossibility of directly measuring sugar content at a distance. A set of experiments relating color and sugar content for the variety of interest would be required. One could then turn to the primary goal of examining alternative hypotheses on how to better grow grapes (or grow better grapes).

The combination of data and the models used in making inferences must meet reliability requirements. Noise, sensor faults, software and communication failures, miscalibration and drift will all contribute to unreliable or dropped data. With models of the fault mechanisms and the physical phenomena under study a network can be planned to detect and mitigate sources of unreliability and to make reliable inferences in the presence of bad data. Mitigation and detection means include sensors with independent fault mechanisms (e.g., sensors of different types), use of more expensive but trusted sensors/experimental procedures to audit lower cost but less reliable subsystems, and redundancy of observations with respect to the minimum required by the physical model. At one extreme, with a fully reliable model there is no need for observations, while at the other with a great deal of reliable data there is no need for an *a priori* model. In general there will be a design iteration as the level of trust in models of the physical phenomenon of interest and the subsystem used for studying it are improved.

Many sources of trust can be invoked in sensor networks. Often, the physical processes that are the objects of scientific observation will be either stochastic or so poorly understood that they must be modeled as such. Clearly, the more predictive the model the more the process itself can be used to test the system that is observing it. For example, spatio-temporal smoothness assumptions allow for over-sampling and thus permit checking measurements against each other for inconsistency. If however little is known about the process then some combination of a greater degree of reliability in the observing system and deployment density is required. Thus, knowledge of the process assists both with planning what density of nodes is required assuming all are reliable and also what increase in density is needed given a model of their unreliability. With a poor model of the reliability of the nodes, an even larger penalty must be paid in deployment density. This is usually judged unacceptable, and consequently controlled experiments with trusted nodes are typically conducted to characterize processes, while others are conducted with controlled conditions to develop a model of the reliability of the nodes. Clearly trusted instruments (in general, external audits) are required for both processes. In this way, uncertainty margins are reduced. Bootstrapping without some trusted experiments is futile: without external validation, it is impossible to know whether observations support the desired inference.

Consider further the role of uncertainty in the sensing/control system design cycle, depicted in Figure 1. Based on some set of hypotheses or control requirements arising from a physical model or set of models, a class of desired outcomes is formulated. An experimental system is then designed to produce these outcomes (e.g., select the most likely hypothesis). How tightly the system can be designed depends both on the breadth of the requirement set and the uncertainties in the model of the physical processed to be studied/controlled, as well as the uncertainty in the behavior of the system itself. Ambiguity in requirements and models in general results in redundancy in design. Thus in practice multiple design iterations are performed to progressively reduce uncertainty in the system and physical models. This allows later designs to be deployed on larger scales at considerably reduced cost. This process implicitly includes humans in the loop, supplying intuition based on past experience, and judgment on how design goals should be modified based upon evolving experience. Consequently, human interfaces should be integral to the design, rather than afterthoughts. Tools are required at all steps to help interpret data and assist the decision process on design modifications and formulation of the next experiment.

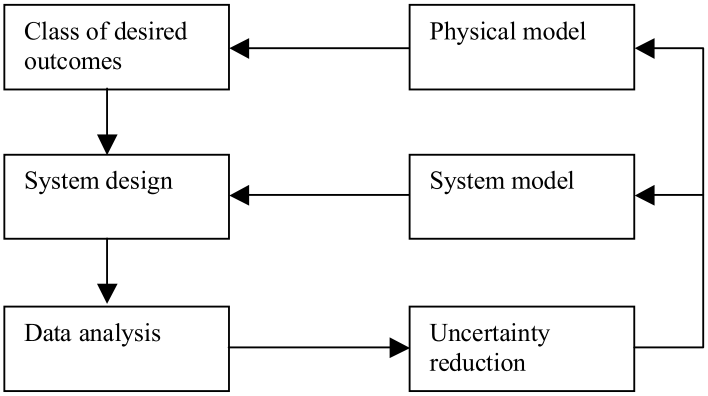

Figure 1: The Design and Deployment Cycle

Consider in particular a situation in which there is some trusted procedure including a combination of models, physical apparatus and operational protocols for observing a physical system at a particular spatio-temporal scale. To be able to draw reliable inferences from observations at broader scales at reasonable cost, new models and system components must build outwards from the trusted system. This trusted system may be used for example in intensive early validation tests of the new components or periodically to audit the system performance. This is the basic engine of the expansion of scientific knowledge and the production of complex systems. The overriding practical consideration is of course how to make the validation less painful, given that the creation of functioning complex systems tends to be one part design and one thousand parts debugging, the latter being impossible if there are too many unmodeled confounding factors. Thus, we re-use that which has already been debugged (hardware, software and models), innovating the least required for each step in the design iteration.

III. PRODUCING TRUSTWORTHY INFERENCES

In this section, we describe several research efforts aimed at ensuring experiments produce trustworthy results, even given unreliable components. A highly motivating experience was an experiment in Bangladesh in which the objective was to study the mechanism whereby arsenic is mobilized to contaminate rice paddies[2]. A sensor network was deployed, and it was not realized until considerably later when the data were analyzed that a large fraction of the sensors at one point malfunctioned or were miscalibrated. (Here the equipment performing the calibration is the trusted system, and the domain scientists involved had definite hypotheses that were to be tested). While useful results were obtained, the quality of the scientific data could have been considerably improved if there were automated detection mechanisms or even a means for researchers in the field to have the necessary tools to analyze results in progress, and thus be able to take remedial action. Given the large logistical cost in this and many other deployments, research into such questions subsequently became a high priority at CENS. Here we describe some research into data integrity and tools to assist scientists and other users in the field.

Detection of faults and design for resilience to faults while gathering the required data have a common theoretical basis. The major components are: (a) a model of the phenomenon of interest (b) a model of the detection system operation under normal conditions (c) a list of the known faults, with corresponding models for behavior under each of them and (d) a trusted mechanism for detection of known and unknown faults and/or restoration of normal conditions. There will always be some uncertainty regarding the parameters of the various models, including prior probabilities, moments of the distributions, etc. The problem is to design a detection system that is robust to the uncertainties while being parsimonious in resource usage. These are conflicting goals, resolvable to some extent through the use of mechanisms that reduce the uncertainties, and thus the redundancy required. Put another way, for the data to be trusted, a high degree of trust is required in some composition of the process model, fault model, and auditing system, with in some sense a surplus of trust in one domain fungible to make up for uncertainties in the others.

When there are no faults, one can use correlated measurements to assist in calibration[3]. Explicit investigation of fault and calibration models can improve the reliability and generality of these techniques. A simple linear model[4] can characterize common faults and calibration errors such as offsets, gain, stuck-at faults and degradation in sensor sensitivity. Remarkably, if sensors are calibrated up to an unknown gain and offset, then if a phenomenon is even slightly oversampled, the solution of a system of linear equations suffices to provide calibration[5]. That is, neither a very dense network nor extensive ground truth or modeling is required. The principle of operation is to perform an orthogonal projection from the space of measurements to a subspace in which the correctly calibrated sensor readings must lie.

When sensors are redundantly deployed with respect to the variations of the phenomenon of interest, they may to some extent verify each other's operation. This can be approached as a data fusion problem[6-10]. It is also possible to correct faulty data using models[11-14]. In the following, we concentrate on consensus problems[15] in which sensors are judged based upon the extent they agree with a model derived from the majority of the sensors. A Bayesian framework is particularly useful in this situation since it is possible to make use of all information available, make approximations when it is not, and refine the models as more information is gathered.

Without ground truth, certain assumptions must be made about the model of the phenomenon in order to proceed. Implicitly, this requires some physical model of the phenomenon based upon past trusted observations. For example, if the phenomenon is modeled to be smoothly varying over the set of K sensors, then variations in both data trends (first derivative) and offsets may be used to distinguish faulty sensors. In our first approach to this problem[16] we find the $K/2$ sensors that agree most strongly, and use this set to create the model of "normal" operation. Since the process evolves over time, the agreeing set is updated, with smoothing to limit jumps in membership. The maximum *a posteriori* (MAP) criterion[17] is used to determine membership, with prior probabilities of a sensor being faulty updated in each step. Some details to consider are the noise model, accounting for differences in measurement times of the different sensors, sampling window duration for smoothing, and the thresholds to set for determining unacceptable data offsets. These all require some knowledge of the physical process of interest. Having then determined which sensors belong to the agreeing subset, the next step is to estimate which of the remaining sensors are faulty. A Bayesian approach may again be used, comparing sensor behavior against a model (linear in our case) derived from the sensors in the agreeing subset, by computing a likelihood moving average. In the end, a Neyman-Pearson test[17] can be performed under the assumption that this likelihood is Gaussian, so that either fault detection or false alarm rate is met. In some experiments with correlated temperature data in a cold air drainage experiment[18] in which one of four sensors was faulty, use of data trends alone resulted in the faulty sensor being correctly detected in 82% of the windows, while with offsets taken into account the detection probability rose to 99%, at the expense of an increased false alarm rate.

Given these encouraging results, many refinements and improvements are possible. A number of parameters must be set in the above procedure; these rely upon human judgments on what actually is the underlying physical model. Better knowledge of that model (e.g., from ground truth resulting from more reliable instruments, accumulating knowledge over multiple experiments, etc.) can yield improved performance. Additionally, the binary division of sensors into faulty or non-faulty is an artifact of having weak assumptions concerning the physical model. What would be more useful is an estimate of the likelihood of a sensor being faulty, given the data and the model, with such estimates being updated as confidence in a model improves.

Faults can also occur at a network level. The Sympathy tool[19] operates by comparing the quantity of data collected at a sink against schedules or expected levels, and then based on past behavior determines likely causes, such as crashes, reboots, lack of connectivity, poor paths through the network, etc., so that faults may be localized. This tool has been used with success in a wide set of sensor deployments, but it has limitations such as the need for expert knowledge of the faults so that a static decision tree can be composed, and the resulting false classifications that occur when conditions are different from what was assumed in building the model. In the Confidence tool the goal was to be more robust to modeling errors and also to deal with data faults. Since ground truth is generally not available at the desired level of detail, an unsupervised learning technique is required. Here, *k*-means clustering is used to learn the feature space in which faults appear as anomalies. The method results in a much lower rate of misdiagnosis of faults, speeding up identification of network faults a factor of three compared to Sympathy.

Given the capability of detecting faults and calibration errors, the obvious next step is to remediate the situation through some combination of adjustment of the data analysis, repair/recalibration of sensors, and redeployment of the network. Additionally, having gathered some data, it may be apparent that the sensor node density is insufficient to answer the question of interest, or new questions occur to the end-users. In all of these cases, it is essential therefore that the network be easily re-tasked (both physically and with new software), and that human interactions are simplified since decisions of these types are well beyond what machines can reliably deal with in complicated settings (e.g., in the natural world). A number of rapidly deployable and re-taskable networks exist[2], enabling the iterated experimentation described in section II. This ability to adapt, using human experience and intuition together with re-taskable hardware is an important means of dealing with uncertainty in models of the equipment and the phenomena of interest alike.

To take advantage of these opportunities for redeployment, there has been considerable research into optimal placement of sensors. Classically, given a model and an objective function one performs an optimization that determines the sensor layout. When the model is uncertain an incremental deployment becomes necessary[20]. In general, what is needed is a technique that is robust to modeling errors. For example, we may begin with the hypothesis that the true model lies within some set, and then incrementally deploy sensors to sharpen the differences between these models, allowing an early decision on which is most likely (or indeed, if there will sufficient confidence in any). This problem aligns quite closely with the practice of experiment design in which there are usually a number of good hypotheses provided by the expert end-users, and is far more tractable than the problem of both designing the experiment and determining the model with unsupervised learning and few priors.

IV. COOPERATION VERSUS DEPLOYMENT DENSITY

Now suppose physical models are well-known and the sensors are reliable. Some natural questions which then arise in planning a deployment are what density of nodes is required for reliable inferences, and how many nodes should cooperate in producing such inferences (e.g., for the detection and localization of a source). A variation on these questions is whether it is better to deploy a few high-performance nodes (e.g., with adaptive beamforming capabilities) or a large number of small nodes (e.g., with only proximity detection). The answer to the latter question depends in large measure on the resource costs assigned to the different classes of nodes, but it is quite possible to compute the relative numbers and the extent to which cooperation among the nodes in making a given inference is helpful.

Consider a scenario whereby nodes are deployed to detect the presence of a source with known characteristics, with the only impairments being simple distance loss (e.g., second or fourth power loss with distance) and additive noise. In such situations, given high precision the optimal fusion rule is maximal ratio combining[21]. Essentially, the final decision is a weighted and coherent sum of the readings from the participating sensors, with the weighting proportional to the SNRs of the individual readings. The result is a composite SNR that is strictly better than that of the best (closest) sensor. In practice of course when there is the possibility of multiple confounding sources and other impairments we should not include within the cooperating set of sensors those with low SNR. Inclusion also implies a resource cost, and so it is in general better to stop when the fused SNR is above some desired threshold, or in general, when the inference meets the quality of service requirements. Thus, the question is the balance between the marginal utility of including the readings of additional sensors in the fused decision, against the cost of doing so.

These types of questions have been investigated for the scenario of localization of sources[22]. The utility of a fused decision was taken to be the inverse of the location ambiguity. With range-only information, the utility of having additional sensors rises quickly until the point that the ambiguities in the solution disappear (e.g., 3 sensors in a plane, or four in 3-space). Afterwards, the marginal utility of including additional sensors quickly declines due to two factors: a) the diminished marginal contribution of the nth sensors as compared to the first few, as by then many sensors have contributed to the utility and b) the decline of SNR with distance, so that the best sensors (i.e., the close ones) dominate the utility. Thus, even though utility might not be bounded, its increase in the number of sensors is at best logarithmic, resulting in extremely high marginal cost versus the utility obtained. Thus, relatively small numbers of sensors beyond the minimum needed for computing location (e.g., on the order of 5-7) suffice for giving most of the utility. Put another way, very large increases in the number of cooperating sensors are required for overcoming an insufficient density of deployment, and thus a density that is high enough for relatively few needing to cooperate on a given decision is a good design choice. The transition from low to high utility occurs quite sharply, so that many reasonable sensor selection algorithms yield similar results. These conclusions are not much altered with more challenging propagation conditions such as multipath fading. In this case, the number of sensors required is increased in the usual fashion for providing diversity against fading, and thereafter the utility saturates.

Thus, when high-precision measurements are available, the cooperating group should be relatively small. This does not mean that the total number of sensors in the network itself must be small; to cover a large region, many will be required. High-quality decisions on specific events within the network could however be made through the cooperation of relatively few sensors in the vicinity of each event, and we argue that the network should be designed with this objective in mind. A higher density than the minimum required for reliable decisions results in needless expense, and a lesser density results in exponentially increased communications. The sharpness of the transition from insufficient quality of inference to saturation of utility indicates the value of the knowing the models: one can then design deployment densities for the sweet spot. With uncertainty, or temporal-spatial variability, a safety margin in deployment density will be required, escalating costs.

A related question we are now investigating is whether there is a similar saturation phenomenon at work in the number of bits of precision in the measurements of the individual sensors. While classically this question is posed in terms of the communication or signal processing resources available to nodes, a more practical motivation is that sensors have will often have limited accuracy, with large expense in calibrating to many bits of precision. Additionally, the relevant variations in the phenomenon of interest may be in the last few bits of the measurement dynamic range. The solution of problems of this type requires a joint optimization of the quantization levels and the weights assigned to these levels. In the simplest case of 1-bit decisions (e.g., for proximity detection), the weights are trivially 1 or 0, but the threshold in SNR for inclusion of a sensor in the fusion decision requires an optimization. As in the high precision case, there will be further tradeoffs involving the deployment density that will depend on the propagation model and the required quality of the fused decisions. Matters of course become much more deeply complicated with uncertainty in the models and with unreliable nodes.

V. CONCLUSION

The certainty available concerning the phenomena to be sensed and the network performing the sensing has a central role in determining the experimental approach. Networks of relatively simple nodes can be successful if everything is well-characterized (i.e., apparatus, phenomenon of interest, and goals). However, matters are seldom so clear in deployments outside laboratories, and even more so in scientific experiments, where the experiment is undertaken precisely because the phenomenon is only partly known, and the sensors are often subject to many environmental degradations. This motivates research into how to assure that the data can reliably support the inferences to be drawn, and leads to the notion of sensor networks and associated tools that can support iterated experimentation with humans in the loop to guide the sequence as data is collected and analyzed. Indeed, in most applications it is advantageous to plan for the inevitable interaction between the designed system and the human operators with their analytical abilities and evolving goals. It is not so much a matter of designing the best system for the problem at hand as designing it to be able to evolve with increased experience and changed priorities. A broad range of very interesting research questions arise, greatly expanding the universe of sensor network problems.

ACKNOWLEDGEMENT

Much of the material in sections I and II is a result of discussions at the IPAM 2007 Workshop on Sensor Networks. The formulation of the primary sensor network problem as being the determination of $P(X|Z)$ is due to Hugh Durrant-Whyte. The research summarized in section III is the product of the research of many colleagues in CENS, while section IV summarizes work conducted by Yu-Ching Tong. Funding was provided principally by NSF through STC funding of CENS and NETS-NOSS-0520006

REFERENCES

[1] Pottie, G.J. and Kaiser, W.J., [Principles of Embedded Networked Systems], Cambridge University Press, (2005).

[2] Ramanathan, N., Balzano, L., Estrin, D., Hansen, M., Harmon, T., Jay, J., Kaiser, W., and Sukhatme, G.,"Designing wireless sensor networks as a shared resource for sustainable development," First Int'l Conf. on Information and Communication Technologies and Development, (2006)

[3] Bychkovskiy, V., Megerian, S., Estrin, D., and Potkonjak, M., "Colibration: A Collaborative Approach to In-Place Sensor Calibration," Proc. IPSN'03, 301-316 (2003).

[4] Balzano, L., and Srivastava, M.B., "Fault in Sensor Networks," TR-UCLA-NESL-200606-01 (2006.)

[5] Balzano, L., and Nowak, R., "Calibration by Signal Subpsace Assumption," Proc. IPSN'07, (2007).

[6] Rogova, G., "Reliability in information fusion: literature survey," 7th International Conference on Information Fusion, (2004).

[7] Marzullo, K., "Tolerating Failures of Continuous-Valued Sensors," ACM Transactions on Computer Systems (TOCS), 8(4), (1990).

[8] Prasad, L., Iyengar, S.S., Kashyap, R.L., and Madan, R.N., "Functional characterization of fault tolerant integration in distributed sensor networks," IEEE Trans. Syst., Man, Cybern., 21(5), 1082-1087, (1991).

[9] Prasad, L., Iyengar, S.S., Rao, R.L. and Kashyap, R.L., "Fault-tolerant sensor integration using multiresolution decomposition," Physical Review E, 49(4), 3452-3461, (1994).

[10] Clouqueur, T., Saluja, K.K. and Ramanathan, P., "Fault tolerance in collaborative sensor networks for target detection," IEEE Trans. Comput., 53(3), 320-333, (2004).

[11] Elnahrawy, E. and Nath, B., "Cleaning and Querying Noisy Sensors," Proc. WSNA'03, (2003).

[12] Elnahrawy, E. and Nath, B., " Context-aware sensors," Proc. 1st European Workshop on Wireless Sensor Networks (2004).

[13] Jeffrey, S.R., Alonso, G., Franklin, M.J., Hong, W. and Widom, J., "Declarative support for sensor data cleaning," 4th Intl. Conf. on Pervasive Computing, (2006).

[14] Mukhodpadhyay, S., Panigrahi, D. and Dey, S., "Model based error correction for wireless sensor networks," Proc. IEEE SECON, 575-584, (2004).

[15] Fischer, M.J., "The consensus problem in unreliable distributed systems (a brief survey)," in Fundamentals of Computation Theory, 127-140 (1983).

[16] Ni, K. and Pottie, G., "Bayesian selection of non-faulty sensors," IEEE ISIT, (2007).

[17] Van Trees, H.L [Detection, Estimation and Modulation Theory, Part I], Wiley, (1968).

[18] Balzano, L. and Graham, E., "Cold air drainage data," Data from sensors 90, 200, 202, 203 at James Reserve, CA, mar 9-13, online at http://sensorbase.org, (2006)

[19] Ramanathan, N., Chang, K., Kapur, R., Girod, L., Kohler, E., and Estrin, D., "Sympathy for the Sensor Network Debugger," ACM SenSys'05, (2005)

[20] Rahimi, M.H., Hansen, M., Kaiser, W., Sukhatme, G.S. and Estrin, D., "Adaptive sampling for environmental field estimation using robotic sensors," IEEE/RSJ International Conf. on Intelligent Robots and Systems, 747-753, (2005).

[21] Proakis. J.G., [Digital Communications, 4th Ed.], McGraw-Hill, (2001).

[22] Tong, Y-C. and Pottie, G.J., "The marginal utility of cooperation in sensor networks," Information Theory and Applications Workshop, (2008)