

ON EVALUATION NEEDS OF REAL-LIFE SENSOR NETWORK DEPLOYMENTS[†]

Alfred Park, Kalyan Perumalla*, Vladimir Protopopescu, Mallikarjun Shankar, Frank DeNap, Bryan Gorman
Computational Sciences and Engineering
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

*Corresponding Author: perumallaks@ornl.gov

KEYWORDS

Sensor Networks, Simulation, Plume Models

ABSTRACT

There is a rapidly growing need to evaluate sensor network functionality and performance in the context of the larger environment of infrastructure and applications in which the sensor network is organically embedded. This need, which is motivated by complex applications related to national security operations, leads to a paradigm fundamentally different from that of traditional data networks. In the sensor networks of interest to us, the network dynamics depend strongly on sensor activity, which in turn is triggered by events in the environment. Because the behavior of sensor networks is sensitive to these driving phenomena, the integrity of the sensed observations, measurements and resource usage by the network can widely vary. It is therefore imperative to accurately capture the environmental phenomena, and drive the simulation of the sensor network operation by accounting fully for the environment effects. In this paper, we illustrate the strong, intimate coupling between the sensor network operation and the driving phenomena in their applications with an example sensor network designed to detect and track gaseous plumes.

INTRODUCTION

Sensor networks are being designed for use in a variety of applications. Some of the applications of interest to us belong in the areas of border surveillance, chemical/biological agent monitoring, homeland security concerns in critical infrastructures and public/government assets. Design and development in these and other areas have been advancing the scalability and effectiveness of the sensor networks. As the complexity and scale of sensors networks increase, so does the need for their effective design and analysis. The complexity of overall sensor network operation is increased with the growing sophistication of data fusion techniques, tracking algorithms, sensor device capabilities, network device characteristics and amount of relevant detail in the phenomena being tracked.

The competing concerns of improved fidelity and performance abound in sensor network operations. While frequent sensor readings can help deliver better accuracy, they can strain the network and also compete with other sensors for network bandwidth. Increasing the wireless radio range of sensor nodes can help decrease the number of hops to a basestation and reduce latency, but it introduces greater interference and/or collisions in radio transmissions among neighbors. Ad-hoc packet routing, via a dynamically established mesh, can help alleviate set up efforts, but can also introduce complex delay and loss patterns into the sensor reading transmissions. Data fusion algorithms (Brooks et al. 1997) can help alleviate bandwidth requirements, but they are subject to various types of assumptions including those about both the sensed phenomena and the inter-sensor connectivity arrangements.

Critical sensor networks deployments often need an accurate a priori understanding of the expected performance of the deployment. Depending on the type of the network, an estimate of the “average case” may be sufficient. In other contexts, it may be necessary to plan for the unanticipated “worst case.” It becomes critical to take into account the spectrum of possible scenarios under which the sensor network deployment will have to operate. Limiting the testing to only a few scenarios could expose the application to unexpected operation, sometimes incurring a large amount of risk and very severe consequences. Our central contribution is demonstrating the specific need for higher-fidelity joint simulations of both the sensor network and the wider-area driving scenarios. We show in our simulation results that this exposes behaviors that traditional approaches (e.g., those that assume point processes for arrival rates, or those that assume non-global effects) may not capture.

Example: Sensor Networks For Plume Detection

Here, we focus on a type of sensor network that exhibits a critical type of operation and hence requires careful analysis of the range of relevant scenarios. Specifically, we consider sensor networks that are designed to sense air-borne chemical/biological agents and rapidly send the sensed observations to basestations/control centers.

[†] This manuscript has been authored by employees of UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. Accordingly, the United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

Clearly, reliable operation of such networks can be critical, especially if deployed for detecting lethal agents. Another important application is the use of such networks for fast and accurate inverse tracking of the location of bombs upon explosion.

EVALUATION PLATFORM

In order to satisfactorily evaluate the range of scenarios to which the plume detection sensor networks would be subjected, we adopt a high-fidelity simulation approach. A high resolution plume generation and dispersion model is used to simulate the phenomenon being sensed and tracked. The sensors are modeled in full detail at the level of analog-to-digital (A/D) conversion. The sensor network nodes connected to the sensors are simulated at the highest resolution possible, namely, by emulation of actual operating system and program execution. Wireless communication is simulated at bit level, and wireless signals are simulated using detailed interference, collision and loss models. Network operation is simulated with actual ad-hoc routing protocols executing on all sensor nodes.

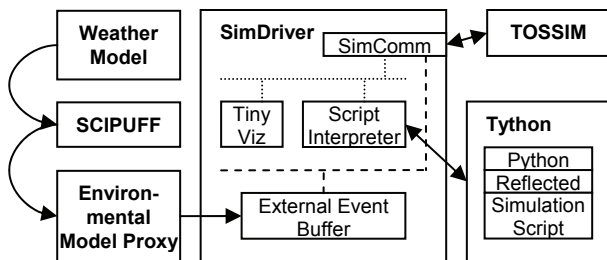


Figure 1: Experimental Infrastructure

The experimental platform illustrated in Figure 1 is composed of two high-fidelity simulators that a central SimDriver modules brings together: (a) a detailed plume generation and dispersion simulator, and (b) a detailed sensor network simulator. The output generated by the plume dispersion model is fed into the sensor network simulator to drive the sensors and sensor network operation. The high fidelity atmospheric dispersion simulation model generates concentration levels at sensor locations. These values are used to feed the sensor A/D interfaces at the sensor nodes positioned at the corresponding locations.

Second-order Closure Integrated Puff (SCIPUFF)

The high-fidelity plume dispersion simulator that we employ is the Second-order Closure Integrated Puff modeling tool (Sykes et al. 1998). SCIPUFF is capable of modeling three-dimensional, time-dependent concentration fields, in terms of “puffs” that are split and/or merged over time as necessary.

The SCIPUFF model has been widely used across many different domains and organizations. The U.S. Defense Threat Reduction Agency (DTRA) has incorporated the SCIPUFF model into the U.S. government-controlled Hazard Prediction and

Assessment Capability (HPAC) software. Aside from its use by the U.S. Army for military evaluations, SCIPUFF has also been recommended by the U.S. Environmental Protection Agency (EPA) as an alternative model for some regulatory applications.

TOSSIM and Tython

The TinyOS Simulator (TOSSIM) is a sensor network simulator based on the TinyOS operating system for wireless embedded sensor devices called motes (Hill et al. 2000; Levis et al. 2003). Motes typically have constrained communication bandwidth, limited computational power, and finite energy resources, but have relatively low power requirements. TOSSIM is unique compared to other sensor network simulators in that the simulation code is built directly from the same code that can be deployed to real motes. This allows application developers the expanded capability of testing unmodified application implementations in simulation. TOSSIM includes a high fidelity bit-level simulation of the radio stack which allows modeling of collisions at the bit-level and accurately capturing the hidden-terminal phenomena.

Tython (Demmer et al. 2005) provides control of TOSSIM simulations via a scripting interface. Tython is built on Jython, a Java implementation of the Python language, allowing users to exploit the full expressive power of Python to access the entire TinyOS Java tool chain. This allows powerful runtime manipulation of TOSSIM simulations such as dynamically adding external stimuli. We use this interface to incorporate SCIPUFF readings as input to simulated sensors attached to TOSSIM motes.

For the following experiments, we gathered concentration values and time information from SCIPUFF and directed it to a proxy in Tython, as shown in Figure 1. We inject this data with interrupts into TOSSIM buffered accordingly in Tython for processing as future discrete events.

Sensor Application and Network Routing

The TinyOS NesC application used in the study is one that activates a generic sensor and samples these values from the A/D Converter (ADC) periodically. Each sensor employs data fusion by sending average sample values over time to the base station node. We configured an ad-hoc multi-hop routing implementation using shortest-path-first and activate two-way link estimation to set up ad-hoc routes and route packets from source nodes to the base station node. Due to limited wireless bandwidth, sensors sample from the ADC every 250 milliseconds and store the sampled value into a circular buffer of length 8. Once the circular buffer is filled, a packet is emitted containing the average value of all readings in the buffer. This translates into one packet sent every two seconds, as to

not exceed the bandwidth available for both the application and the multi-hop ad-hoc routing protocol.

EXPERIMENTAL STUDY

The plume detection and tracking application is a fairly complex one, with many dimensions and parameters. Here we focus our study on a subset of the dimensions, to highlight the amount of variability of sensor network performance with varying set of scenarios and the need for careful and methodical evaluation of scenarios in a real-life deployment.

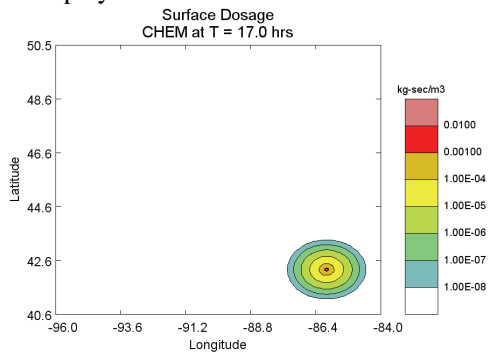


Figure 2: Plume Propagation with No Winds

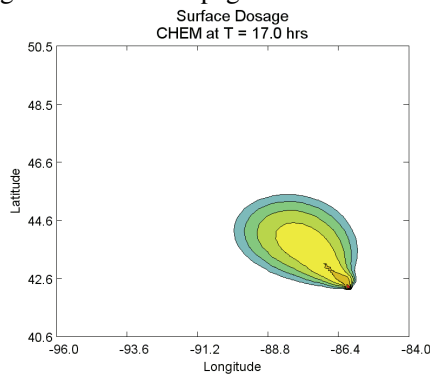


Figure 3: Plume Propagation with (Moderate) Winds

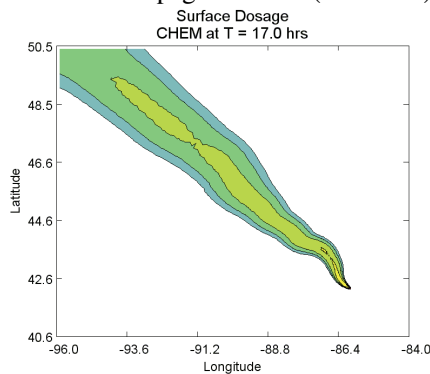


Figure 4: Plume Propagation with Turbulent Weather

Scenario Parameters

Weather and Dispersion Effects

In our experiments, the weather model in SCIPUFF was varied with three representative scenarios: *no winds*, *some winds*, and *turbulent* weather conditions, whose snapshots are illustrated in Figure 2, Figure 3 and Figure 4 respectively. The weather model scenarios respectively exhibit no precipitation and no

winds, no precipitation with variable light to moderate winds up to 8 mph, and heavy rain with variable winds up to 40 mph respectively. Temperatures for all cases fell between 70 to 77 °F. The *no winds* case is used to represent a base model as shown in Figure 2 due to near uniform spread of the gas in all directions indicating a simple dispersion model which is sufficient for comparing to more complex scenarios.

The SCIPUFF model generates concentration values for a single source release with a duration of 300 seconds in the middle of the area modeled, with a release time at 300 seconds. The initial 300 seconds allow the network routing protocol to build routes to the base station.

Sensor Node Layout




Figure 5: A 5×5 Sensor Network Layout

Two flat geographic areas were modeled: a *dense* area measuring 98.42ft² and a *sparse* area measuring 196.85ft². For each of these areas, the sensors were arranged as a 5×5 grid of sensor nodes at ground level as shown in Figure 5. Nodes are numbered contiguously, with 0 at zero on left-top corner, 4 on right-top corner, 20 in left-bottom and 24 at right-bottom corner (node 13 is the closest to the source).

Wireless Communication Effects

The radio propagation model was varied for each case using either a *binary disc model* or an *empirical lossy model*. In the *binary disc model*, any pair of nodes within 50 feet of each other would be guaranteed to receive packets from each other (assuming no collisions during the packet transmission). In the *empirical lossy model*, two-way loss rates are assigned from observed values as assigned by TOSSIM's LossyBuilder program.

Simulation Results

Sensor Reading Variation due to Weather

In the sparsely populated sensor node grid configuration (*sparse* 5×5), a maximum of one other node can hear packet transmissions from the sender in each direction. This configuration allows analysis of weather effects as perceived by the base station with minimal impacts from the wireless network due to packet loss from collisions.

One of the clear observations from the experiments is the significant effect of weather on the phenomenon

sensed at the base station. Figure 6 shows mean concentration readings for the gas of interest for each of the three weather conditions, corresponding to the values sent by the mote closest to the plume source. The sensor readings as received by the base station are plotted over time.

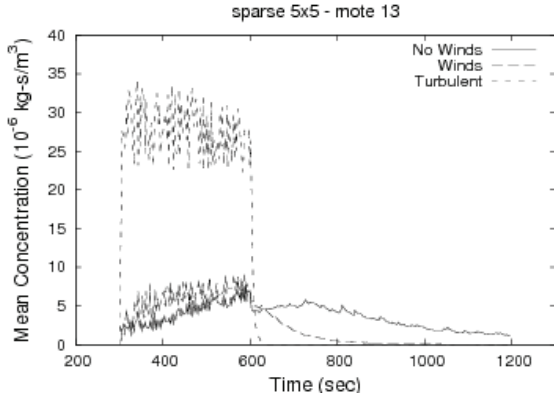


Figure 6: Sensor Readings at Release Source

As expected, the *no winds* case shows a gradual rise in observed readings due to plume release and then a steady decline due to plume dispersion. In the *winds* case, the concentration of the gas rises faster with greater variability in the readings. However, after the gas release stops at 600 seconds, the readings decline towards zero faster than the *no winds* case due to dispersion of the gas aided by winds. The mean concentration rises sharply with high fluctuation for the duration of the release in the *turbulent* case due to high winds and heavy downpour dispersing the gas quickly. Upon end of the release, mean concentration values for the gas decline rapidly, again due to the high winds, reaching undetectable levels far quicker than the *no winds* or the *winds* cases.




Figure 7: Sensor Readings 65.61ft from Source in Opposite Wind Direction

Figure 7 shows sensor readings for the mote situated furthest east of the source, opposite of the general wind direction. Mean concentration values are lower for the *winds* case and significantly lower for the *turbulent* case than observed at the source release and at mote locations which fall in the general wind direction.

Network Performance Variation due to Node Density

In the densely packed sensor grid (*dense 5x5*), sensor nodes can be in range of up to two other sensor nodes in

its broadcast radius. Although this can reduce the number of hops a packet must traverse to the base station node, the probability of bit collisions increases. This scenario provides some insight into network effects on readings received by the base station as compared to the sparsely connected sensor network.

The sensor readings of the dense network show similar patterns to the sparse sensor network grid configuration although with higher concentration values due to reduced inter-sensor distance to the source than the *sparse* case. However, observations indicate small anomalies in the readings received by the base station.

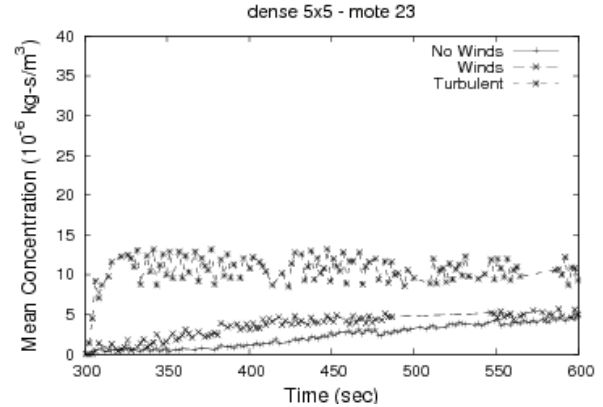


Figure 8: Sensor Readings showing the Effect of Packet Loss

Figure 8 shows the source release time in more detail, illustrating large periods of time where no data for mote 23 was received by the base station. In the *winds* case, no sensor readings were received for approximately 60 seconds, a significant portion of the release time. If any fluctuations in sensor reading had occurred during this time frame, they would not have been perceived by the base station. These gaps in sensor readings are due to packet losses from bit collisions that were not present in the sparse scenarios.

Wireless Communication and Network Effects

Binary Disc Radio Propagation Model: Under the best-case network operation where packets can be received by other nodes if within 50 feet assuming no bit-collisions occur during packet transmission, weather effects on sensor readings can be analyzed with minimal impact from network effects by using the *binary disc model* as illustrated in the previous scenarios.

Empirical Lossy Radio Propagation Model: The following scenarios evaluate the impact of a lossy radio model exhibiting bit-error rates with some probability mirroring real-world observations. This is achieved by sampling by distance from a Gaussian packet loss probability distribution matching observed data (Levis et al. 2003).

Figure 9 and Figure 10 show dramatically different results than what was observed in the *binary disc model*. Although the rough outline of trends can be determined from the figures, many critical points are missing from each sensor node location. For instance, at the source release mote 13, there is no indication of a rise in concentration, or a rapid decrease in

concentration for the *turbulent* case as was observed in Figure 6. In the *no winds* and *winds* cases for mote 13, no data is received until 100 seconds after the source release begins.




Figure 9: Sensor Readings at Release Source
sparse 5x5 - mote 13

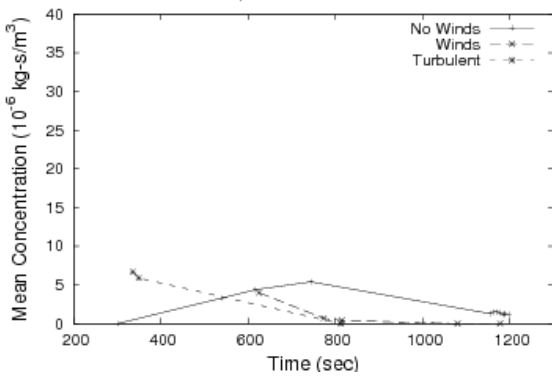


Figure 10: Sensor Readings 65.61ft from Source in Opposite Wind Direction

Figure 10 shows no discernable trends for the *winds* and *turbulent* cases as the data received by the base station is too sparse except for the *no winds* case. The low number of received packets is due to the distance mote 23 is away from the base station node requiring a minimum of five hops in the sparse grid. Each hop incurs a probability of transmission failure either due to collisions or corruption at the bit-level. The relationship between distance and number of packets received is observed in Figure 9 and Figure 10 with the base station receiving more packets from closer nodes.

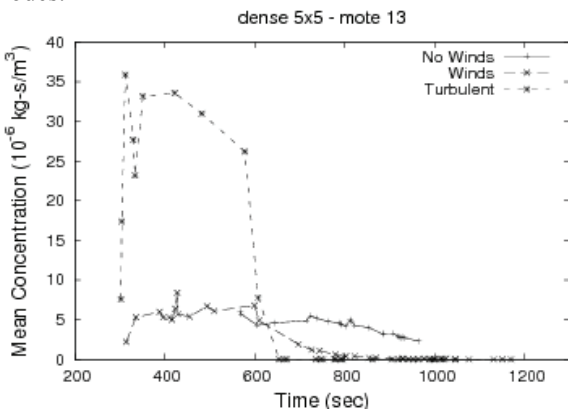


Figure 11: Sensor Readings at Release Source

With the decrease in distances between nodes in the dense scenarios, the bit-error rate decreases resulting in higher reception rates across all nodes as illustrated in Figure 11. Although this improves the rate at which packets are received by the base station from all sensor nodes, there are deficiencies still present such as in the *no winds* case in Figure 11 where the first data point is received near the end release time and thus no noticeable pattern of concentration increase is observed.

A high fidelity sensor network simulator with realistic radio propagation models can reveal phenomena such as low packet reception as seen in the empirical lossy cases. These effects combined with different weather models from the environmental simulator produces highly variable results as perceived by the base station.

Aggregate Network Performance

Although some of the network effects were observed when visualizing concentration values for various cases, analysis of the interplay and effects of network layout, sensor node proximity, and constrained bandwidth is of interest to evaluate the cumulative impact of the simulation results.

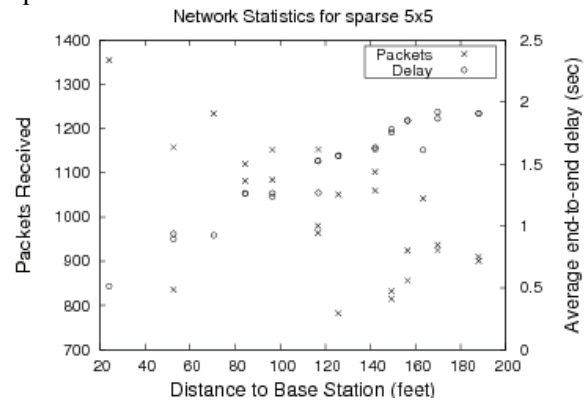


Figure 12: Binary Disc Model for Sparse 5x5 Grid

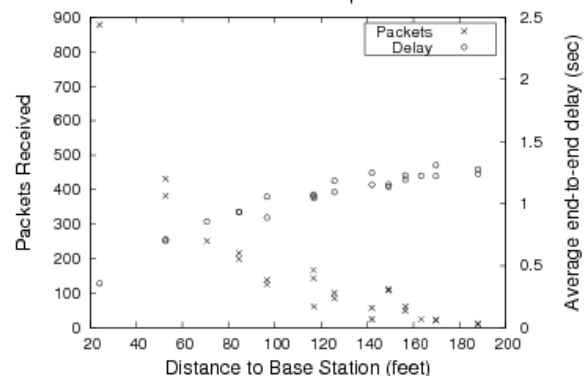


Figure 13: Empirical Lossy Model for Sparse 5x5 Grid

For the binary disc radio propagation models, Figure 12 shows an expected inverse relationship in packets received as the absolute distance between the base station and the sensor location increases. Moreover, the observed average end-to-end delay increases as distance increases between the base station and mote. Figure 13 illustrates a rapid decline in packet

reception rates as the distance between sensor location and base station increases for the empirical lossy radio propagation models. This is due to the added bit-error rates which can cause packets to be corrupted and dropped in addition to normal packet loss due to collisions.

An interesting phenomenon that only occurs in the binary disc radio propagation models is the set of observations with 1.5 second average end-to-end delays (Figure 12). This pattern is observed due to 100% accurate packet reception for nodes at the edge of radio range assuming no collisions, while in the empirical lossy models, many of these packets would be lost due to bit errors. Using models that compromise accuracy such as the *binary disc model* can display anomalies such as receiving packets which would, in reality, have very small probability of successful reception.

From these observations, sensor network application implementations could incorporate a sophisticated packet acknowledgement and re-transmission scheme to improve network performance and ultimately providing more accurate sensor readings.

SUMMARY AND FUTURE WORK

The experimental results provide clear evidence that environmental models that ignore broader-scope effects such as weather and orientation may not suffice to provide realistic plume propagation data. Detailed weather models, in conjunction with high fidelity atmospheric modeling, drive accurate simulation results for evaluating complex interactions between the environment and the sensor network.

Careful consideration of node placement in sensor network deployment is also required, as shown by simulation. Intelligent organization of sensor nodes providing adequate coverage while maintaining balanced node density in an area of interest can minimize side effects such as loss of information due to packet collisions. Moreover, these results emphasize the importance of utilizing a high fidelity sensor network simulator with accurate radio propagation models. Less realistic models can compromise the overall accuracy of the simulation by including unrealistic packet transmissions and receptions which, in turn, could affect other packets through collisions or dropped packets due to bandwidth limitations.

In general, although existing simulations of distinct but interdependent systems exist, it is our contention that a scenario-informed approach to simulation at high-fidelity will expose unexpected behaviors that simpler assumptions will not capture. Our high-fidelity simulation-based evaluation approach can be generalized into an integrated environment of a rich variety of models, simulations and live equipment in a framework to configure, simulate/emulate, test and evaluate sensor-based designs, as shown in Figure 14. A unified, flexible Testing and Evaluation (T&E) framework for chemical/biological (CB) sensors is envisioned which can be built on top of existing

standards to the extent possible. Within this T&E framework, the Department of Defense High Level Architecture (HLA) can be leveraged as a component for runtime integration to facilitate incorporation of existing simulators.




Figure 14: Generalized chemical/biological sensor network testbed framework

The interplay of CB dispersion, weather effects and sensor communications can be evaluated together ultimately in the context of actual operations. To this end, it is important to integrate them with operations simulators such as semi-automated forces (SAF) simulators, facilitating end-to-end, mission-specific evaluations. A similar need also exists in T&E of live devices with respect to their operational response and effectiveness in the larger context.

Design of such a comprehensive framework enables a wide range of T&E possibilities. Scenarios of operation can be effectively simulated and measures of performance can be obtained for individual subsystems or for the entire operation as a whole. The framework can help uncover possible realms of operation that are otherwise hard to capture with small, isolated testing.

Acknowledgements

The authors would like to thank Ronald Lee for information on SCIPUFF operation and Auroop Ganguly for background on related work in plume modeling.

REFERENCES

- Brooks, R. R., et al. (1997). *Multisensor Fusion: Fundamentals and Applications with Software*, Prentice Hall.
- Demmer, M., et al. (2005). *Tython: A Dynamic Simulation Environment for Sensor Networks*, UC Berkeley: 6.
- Hill, J., et al. (2000). *System Architecture Directions for Networked Sensors*. Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems, ACM Press.
- Levis, P., et al. (2003). *TOSSIM: Accurate and Scalable Simulation of Entire TinyOS Applications*. Proceedings of the 1st International Conference on Embedded Networked Sensor Systems, ACM Press.
- Sykes, R. I., et al. (1998). *PC-SCIPUFF Version 1.2PD Technical Documentation*, Titan Corporation, Titan Research & Technology Division, ARAP Group: 181.