A Simulation Tool for ASCTA Microsensor Network Architecture [1,2]

Simon Woo, Esther Jennings, and Loren Clare

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109

{Sung.I.U, Esther.H.Jennings, Loren.P.Clare}@jpl.nasa.gov

Abstract—Advances in technology have made the large-scale deployment of low-cost networked sensors possible for situational awareness. We developed a Simulation Tool for the Advanced Sensors Collaborative Technology Alliance (ASCTA) Microsensor Network Architecture (STAMiNA) to evaluate the performance of networked sensor systems. This tool is built upon a commercial network simulator engine (QualNet), with extended capabilities to include both sensing and communication models in a discrete-event simulation environment. Using this tool, we can simulate target detection (sensing) and information dissemination (information fusion) via wireless communications under different parameters and scenarios, incorporating such metrics as target detection probabilities, false alarm rates, and communications load, and capturing effects of terrain and threat characteristics. An example of tool usage is presented illustrating the comparison of alternative microsensor network architectures such as localized-fusion, hierarchical-fusion, and distributed-fusion in the presence of false detection events. The trade-offs among these three different sensor architectures are examined under different fusion rules, sensing sliding window sizes, and false-event occurrence rates. Operating parameters that yield high detection and low false-alarm performance are examined.

TABLE OF CONTENTS

1. INTRODUCTION.....1
2. STAMINA SOFTWARE CAPABILITIES 2
3. SENSOR NETWORK ARCHITECTURE.....3
4. SIMULATION AND RESULTS 4
5. CONCLUSIONS 8
6. ACKNOWLEDGMENTS.....8
REFERENCES 8
BIOGRAPHIES..... 8

1. INTRODUCTION

Pervasive availability of sensors, especially microsensors and disposable sensors, promises significant advancements in situation awareness for the future military force. These benefits arise from the ability to *network* distributed sensors, thereby achieving synergy through fusion of information from different perspectives and operational improvements in performance by execution of corroborative actions. However, such systems exhibit high complexity, causing system performance prediction to be extremely challenging. There is a clear need to develop technologies for characterizing sensor networks to understand their potential use, based on key system parameters that capture the essential mission aspects. In this paper, we present a novel approach for deriving the sensing as well as communication performance of a distributed sensor system using a simulation environment. A microsensor network analysis tool has been developed at the Jet Propulsion Laboratory (JPL) as part of the Microsensors research program in the Advanced Sensors Collaborative Technology Alliance (ASCTA), a consortium of industry and university organizations working with the Army Research Laboratory. This tool is called Simulation Tool for ASCTA Microsensor Network Architecture (STAMiNA). It is built upon the commercial QualNet discrete-event simulation engine, which provides a highly capable simulation environment for wireless communications networks.

Using STAMiNA, users can define (1) the mission environment, including terrain features, (2) the sensed object set, including multiple threat objects, (3) the sensor placements, their modalities and their abilities to sense

[1] 1-4244-1488-1/08/$25.00 ©2008 IEEE
[2] IEEEAC Paper #1547, Version 4, Updated October 22, 2007

different object types, and false alarm rates, (4) (threat) object trajectories, (5) sensing as well as sensed data dissemination for information fusion, and (6) various network configurations and formations between sensors to examine the coupling of sensing and communication. With these features, the simulation tool can provide the overall system level performance of different sensor network architecture under different parametric conditions. The novelty of our extension is the capability to incorporate both sensing and communications in the same tool. Target sensing, fusion, and information dissemination via wireless communications are simulated in a common simulation environment, capturing the inter-related effects necessary to determine sensor network architecture performance.

The tool is used to support analyses of alternative sensor system architectures [1]. A comparison of localized, hierarchical, and distributed-fusion architectures is presented, in which sensing, false alarm, and communications performance measures are derived. We derive metrics including target detection and false alarm reports under different fusion rules, sensing sliding window sizes, and false-event occurrence rates.

This paper is organized as follows. A summary of the capabilities of the tool is given in Section 2. In Section 3, we present the three cluster-based sensor network architectures used to demonstrate the tool's capabilities. Section 4 presents the simulation performance comparison results. Concluding remarks are discussed in Section 5.

2. STAMINA SOFTWARE CAPABILITIES

STAMiNA models sensor node laydown, multiple sensor modalities, target behavior, cueing, multi-sensor fusion, and communications networking, and is capable of characterizing the performance of alternative sensor systems in terms of probability of detection and probability of false alarm metrics. These different simulation capabilities are greatly accelerated by Senor Media Access Control (Sensor MAC), which is the novel architecture built for sensing and information dissemination in STAMiNA. To better illustrate our tool, the following figure 1. is provided to compare the software structure of STAMiNA with the commonly known open systems interconnection (OSI) reference model (RM). Internal sensor network traffic can be generated and scheduled at the application layer. Sensor MAC is responsible for sensing and distributing detected data to other sensor nodes. Also, multi-modal and multi-target tracking is configured by Sensor MAC. False alarm, line-of-sight (LOS), and terrain effects are captured in the physical layer.

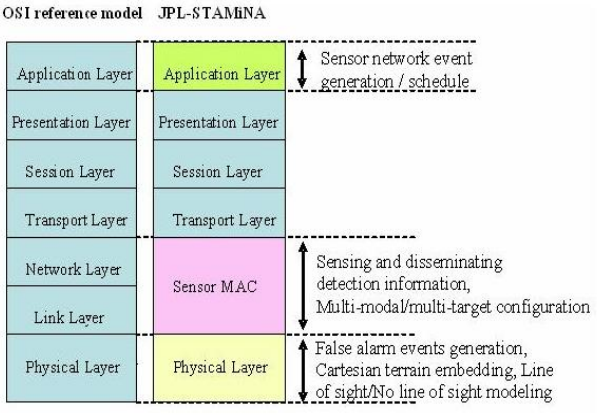

Figure 1. STAMiNA structural relationship to OSI-RM

In STAMiNA, sensor nodes are placed in either x-y (2D, flat terrain) or x-y-z (3D) positions and are generally assumed to be stationary. Stochastic spatial point processes with different correlation characteristics may be generated (e.g. "blue noise") for a more realistic node laydown in the 2-dimensional case. A general 3D terrain model that user can freely define may be used as an input to scenarios. Terrain blockage affects both communications signal and sensor signal propagation. Targets move according to a randomly or prescribed waypoint pattern and speeds may vary. Each target continuously (unintentionally) emits signals that are sensed when in range of the sensor nodes. The following figure 2. captures a simulation snapshot of a randomly generated terrain with several microsensor and clustehead sensors deployed over the region, and a tank (target) located near the center region.

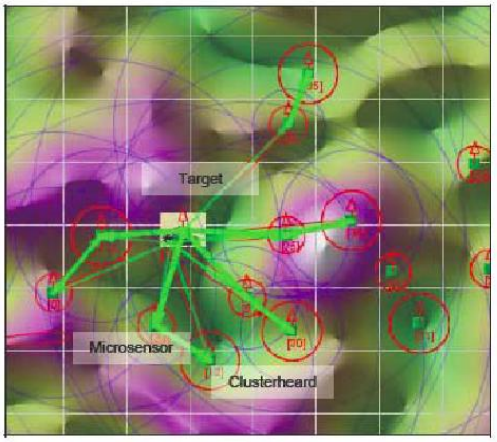

Figure 2. Simulation with randomly generated terrain

Each sensor node in STAMiNA contains at least one radio and at least one sensor. Each radio and each sensor is associated with a "channel". A further description of this software architecture was given in [2]. Different sensor channels correspond to different sensing modalities, such as acoustic, magnetic, etc., and each has its own defined attributes of range and false sensor detection event rates.

False alarm events are generated as an independent random process at each sensor based on uniform, exponential, and Gaussian distribution at the physical layer. STAMiNA can model both directional (bearing) and omni-directional sensing nodes. A general cueing mechanism is modeled that enables a target detection event to trigger communications. That is, as soon as a sensor detects a target, it forwards/broadcasts a target detection message to the destination(s) defined within the model.

Multi-hop communications enable sensor system information to be propagated among nodes and to end-users. Non-interfering inter-cluster and intra-cluster communications can be defined using orthogonal communication channels. Sensor nodes communicate using protocols based on models drawn from the rich QualNet wireless library (including MANET multi-hop networking choices). In addition, JPL has developed Disruption Tolerant Networking (DTN) communication models [3] that capture the intermittent connectivity that is known to arise in communications that are placed low to the ground such as microsensor nodes, and they can be used for sensor network evaluation.

Simulation outputs include real-time animation, statistics collection of metrics for target detection, false alarm occurrences, and communications/target detection forwarding activities that are logged in a statistics file for detailed post-analyses. These utilities greatly enhance users' abilities to process and analyze data to evaluate different sensor network architectures.

3. SENSOR NETWORK ARCHITECTURE

In this work, we specifically consider alternative cluster-based sensor network architectures, similar to those proposed in [4], as an example to demonstrate the tool's capabilities. The interactions between detection performance and communications load for given target trajectories are observed. Detection performance under different false event occurrence rates and varying sensing detection sliding window sizes and fusion thresholds are investigated. This provides quantitative comparisons among the proposed sensor network architectures.

In the present scope of study, two types of sensor nodes are considered: clusterhead sensor nodes and microsensor nodes. Clusterhead sensor nodes have more processing power and larger sensing range than microsensor nodes. Although microsensor nodes have relatively reduced capabilities, it is assumed that their unit cost is much less than than that of a clusterhead sensor node. Hence, a larger number of microsensor nodes can be deployed. Microsensors and clusterhead sensor nodes not only can detect targets, but they can also communicate with each other. In each cluster, there are multiple microsensor nodes and a single clusterhead. The clusters are not necessarily disjoint.

Based on these assumptions, three different cluster-based architectures are considered as shown in Fig. 3, where a large circle represents a clusterhead, and a small circle represents a microsensor. The dotted circles indicate the communication ranges of clusterhead sensor nodes.

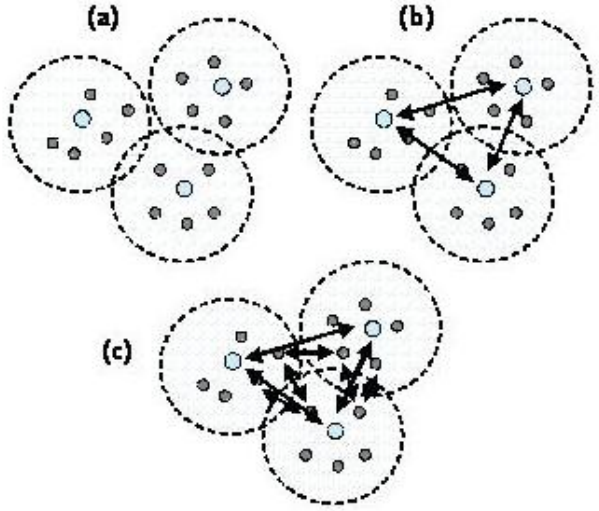

Figure 3. Cluster based sensor network employing (a) localized-fusion (b) hierarchical-fusion and (c) distributed-fusion.

Figure 3(a) depicts a localized-fusion architecture, where the microsensor nodes only communicate with other nodes within its cluster; there are no inter-cluster communications and the clusters are assumed to be disjoint. This architecture is used as a baseline for comparison. Figure 3(b) shows a hierarchical-fusion architecture, where microsensors communicate with their respective clusterheads and the clusterheads handles inter-cluster communication. The arrow lines depict the inter-cluster communications. Figure 3(c) shows a distributed-fusion architecture, where each microsensor and clusterhead can communicate to sensor nodes in other clusters either directly or multi-hopped. Hence, this is a truly distributed sensor network architecture, where detected information from one sensor node is propagated to all other sensor nodes.

The distributed fusion-architecture is more robust and resilient, since detected information takes multiple paths to get to a clusterhead, providing spatial diversity. However, this architecture is the most prone to false detection events and incurs the greatest communications load. Hierarchical-fusion sits in between the localized-fusion and distributed-fusion schemes and can filter target detection decisions at the clusterheads. Hence, hierarchical fusion may stop certain false alarm reports from propagating through the entire network.

In [5], three different information aggregation methods were described: majority voting (MV), distance weighted voting (DWV) and confidence weighted voting (CWV). In [6] and [7], spatial, temporal, and confidence dimensions were proposed to decrease false alarm rates in data aggregation. Our data aggregation/fusion model implements the temporal dimension by using an adjustable time window; the spatial dimension is implemented via clustering and intra-/inter-cluster interactions. Hence, a target detection decision is made at the clusterheads by gathering temporal detection information from microsensors. The signals emitted by a target are received at microsensors according to a predefined sampling rate. Target detection rules at each clusterhead are as follows: A clusterhead decides that a target is present if it receives consecutive detection reports from a single sensor, or simultaneous detection reports from multiple sensors that sum to greater than a *threshold* (θ), where θ is a predefined parameter at the clusterheads. The larger θ value will reduce the number of detection report events. The intention is to filter out and minimize the false detection events (system level false alarm rate).

In our localized-fusion algorithm, we adapted the MV algorithm in [5] because it only supports intra-cluster communication. Let s_i (t, t+Δ, r) be positive target detection at sensor i for the time period from t to t+Δ, at a sampling rate of r, where Δ is defined as a window size. If a target is detected from t to t+Δ at sensor i, then s_i (t, t+Δ, r) = 1. Otherwise, s_i (t, t+Δ, r) = 0. Let n be the number of nodes in a cluster. At time t, we define the mth clusterhead's localized-fusion, hierarchical fusion, and distributed fusion target detection reports as $D^m{}_{LF}(t)$, $D^m{}_{HF}(t)$, and $D^m{}_{DF}(t)$ as follows:

$$D^m{}_{LF}(t) = \begin{cases} 1, \sum_{i=1}^{n} s_i(t, t+\Delta, r) \geq \theta \\ 0, \text{otherwise} \end{cases},$$

$$D^m{}_{HF}(t) = \begin{cases} 1, \sum_{i=1}^{n} s_i(t, t+\Delta, r) + \alpha \cdot \sum_{k=1}^{K-1} D^k_{LF}(t) \geq \theta \\ 0, \text{otherwise} \end{cases},$$

and,

$$D^m{}_{DF}(t) = \begin{cases} 1, \sum_{i=1}^{N} s_i(t, t+\Delta, r) \geq \theta \\ 0, \text{otherwise} \end{cases},$$

where $D^k_{LF}(t)$ is the localized target detection made by kth neighboring cluster (value of either 0 or 1) at time t, α is a weighting factor parameter that allows us to incorporate distance weighting or confidence weighting with respect to each cluster, K is the total number of clusters, and N is the total number of sensors deployed in the simulation.

Detection fusion occurrences $D^m{}_{LF}(t)$, $D^m{}_{HF}(t)$, and $D^m{}_{DF}(t)$ and aggregates of these at each clusterhead are collected for execution runs of simulation scenarios. An architecture with more frequent target detection reports implies greater robustness and accuracy about situational awareness, since it generally is able to filter out false sensing events while providing substantial target detection information.

The sensing sliding window parameter Δ has an important affect on sensor network system performance. A larger Δ will cause more detection reports, allowing more false reports to go through. However, choosing too small a Δ may cause legitimate detection events to be rejected. These trade-offs are shown in the next section.

4. SIMULATION AND RESULTS

The performance of three different architecture is evaluated by comparing (1) the number of target detection occurred, $D(t)$, against different thresholds, θ, (2) sensor detection sliding window sizes, Δ, and (3) false event occurrence rates, P_{fa}. Three different sensor network architectures are simulated in a randomly generated 1500 meter x 1500 meter terrain. The terrain contains arbitrarily generated humps (hills) as blockage to signals, where some of sensors may not detect the target even if the target is within the detection range. Each architecture contains two clusters, where each cluster is composed of a single clusterhead and seven microsensors. For the example scenarios considered here, two clusters are adequate to capture the detection and communication performance of the different sensor network architectures. Since a larger network can be decomposed into smaller clusters, we can easily extend and generalize the result obtained from this work to larger networks. We model the target to follow a near straight line path with different starting and ending points and may pass through a cluster or traverse between clusters (on cluster borders). For the sake of brevity, we only provide the simulation result from a single target trajectory in the present study. Microsensors have a sensing range of 200 meters and clusterhead sensors have a sensing range of 400 meters and are assumed to be omnidirectional. We chose 802.11b DCF as the MAC communication protocol among other available communication protocols in our tool for both inter- and intra-cluster network communications. The detection sampling rate, r, is set to once per second. The sensing sliding detection window is varied parametrically from one to five seconds. We compare the target detection and information dissemination performance with the same sensor deployments.

The following Figure 4 shows that communication improves *target awareness time*, defined as an indicator for the instant whether the sensor node network either detects a target or is notified about the detection of a target for a given instant.

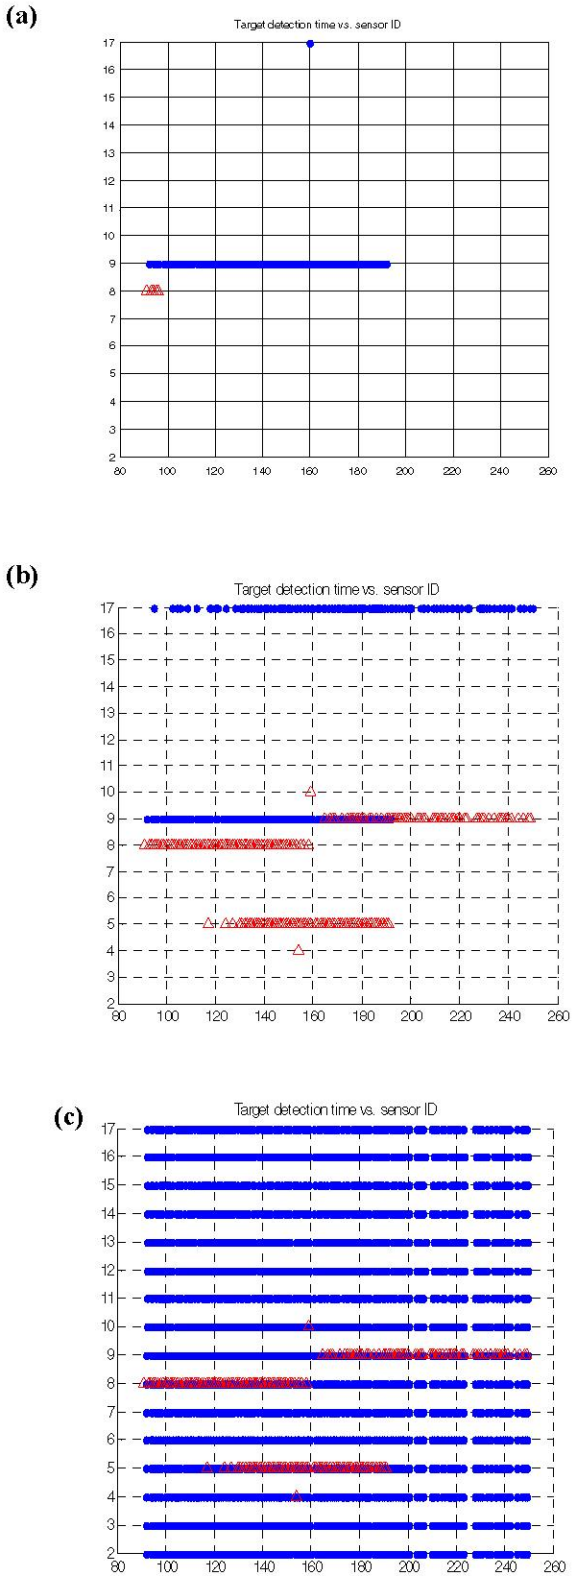

Figure 4. Target awareness time using (a) localized-fusion, (b) hierarchical-fusion, and (c) distributed-fusion.

In this simulation, sensor nodes numbered 9 and 17 are clusterheads. Sensor node IDs from 2 to 9 form one cluster and 10 through 17 form the other cluster. The x-axis is the simulation time (in sec), and the y-axis is the sensor ID.

The triangle marks show the target detection time by each sensor. The circle marks show the time instances when a sensor is informed of target detection by other sensors. Figure 4 provides the *target awareness time* with respect to the three different fusion architectures. Figure 4(a) shows target detection using localized-fusion. We clearly observe the detection separation among clusters; the target is only detected by nodes in cluster 1. Figure 4(b) shows target detection and forwarding of detection information using hierarchical-fusion. The target is detected by nodes in cluster 1 and communicated to the clusterhead (node 17) of cluster 2. Figure 4(c) shows target detection using distributed-fusion. Target detection at one sensor is immediately propagated through all reachable sensor nodes, whereas in hierarchical-fusion, clusterhead serves as a filter for decision making and forwarding target detection to other clusters.

It is clear from Figure 4 that the distributed fusion case results in the best situational awareness, although it stimulates the greatest communications load on the network and correspondingly consumes the most energy. The communication performance is captured in the following Figure 5. In Figure 5, the total number of transmissions incurred at each sensor node to relay target detection information are shown.

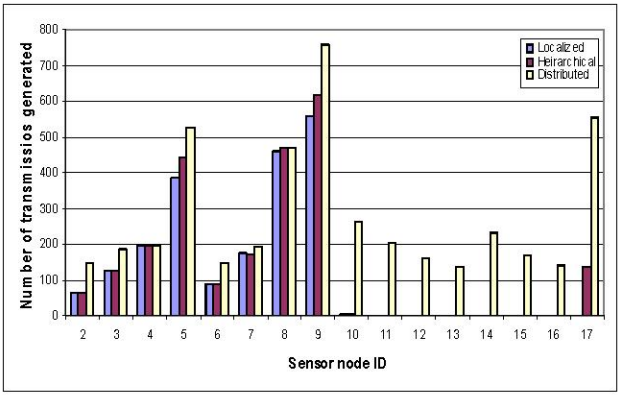

Figure 5. Communication loads incurred at each sensor upon target detection for localized-fusion (LF), hierarchical-fusion (HF), and distributed-fusion (DF).

The x-axis is the sensor ID and the y-axis is the number of transmissions by each sensor node. For sensor node IDs 9 through 17, no relay communications occurred in the localized-fusion architecture, since there is no inter-cluster

communication. On the other hand, relay traffic occurred at clusterhead (sensor ID 17) in the hierarchical-fusion architecture for inter-cluster communication. Overall, the distributed-fusion architecture yields approximately twice the communication load as the hierarchical-fusion architecture to fully propagate the target detection information. The hierarchical-fusion scheme requires only 12 percent more communication overhead than the localized-fusion case to achieve inter-cluster communication. This increased communication load leads to better detection capability by the sensor network.

In order to characterize detection performance, we computed the number of target detection reports for each sensor network architecture, while varying the sensing detection sliding window size. This allows capturing the effect of Δ on target detection performance for a fixed threshold. Based on target detection reports at each sensor node, we computed the number of target detections reported at each clusterhead. We used *threshold* = 1, 2, and 3 respectively. We varied the sensing detection sliding window size as 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, and 5 sec, and recomputed the corresponding D_{LF}, D_{HF}, and D_{DF} of each sensor network algorithm. We set $\alpha = 1$ for simplicity.

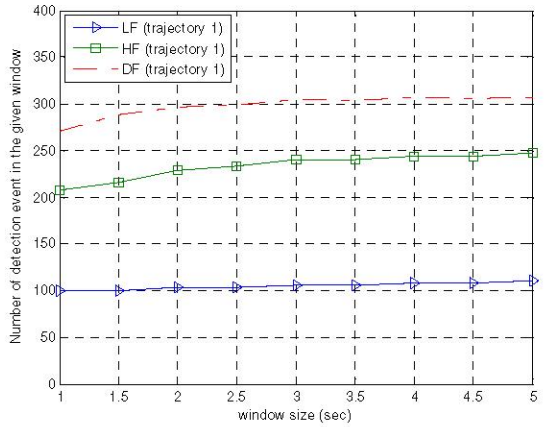

Figure 6. Target detection reports at clusterheads with *threshold* = 1, Δ = 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, and 5 sec.

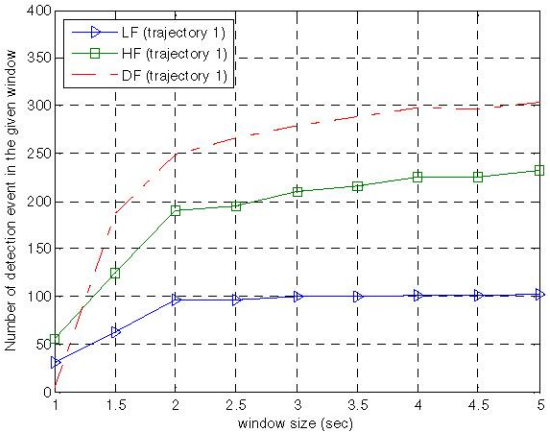

Figure 7. Target detection reports at clusterheads with *threshold* = 2, Δ = 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, and 5 sec.

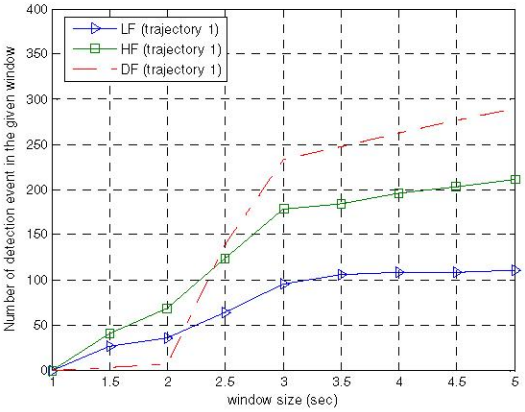

Figure 8. Target detection reports at clusterheads with *threshold* = 3, Δ = 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, and 5 sec.

The total simulation time, T, was set to be 250 sec. Figures 6, 7, and 8 show aggregated $\sum_{t=1}^{T}\sum_{m=1}^{2} D_{LF}{}^{m}(t)$, $\sum_{t=1}^{T}\sum_{m=1}^{2} D_{HF}{}^{m}(t)$, and $\sum_{t=1}^{T}\sum_{m=1}^{2} D_{DF}{}^{m}(t)$ with respect to different threshold values and different window sizes. The x-axis is the sensing detection sliding window size and the y-axis is the number of detection reports at the clusterheads (sum of reports at each clusterhead). We observe that larger sensing detection sliding window size allows more detection events to occur, regardless of threshold and sensor network architectures. The distributed architecture yields the highest detection occurrences as the sensing detection sliding window size increases. However, the cost of doubled communication load only resulted in 20 to 35 percent increase in detection reports, comparing distributed-fusion with hierarchical-fusion. When $\Delta \leq 3$ and *threshold* ≥ 2, the detection performance of distributed-fusion is worse than any other scheme. This is because data

is sampled at every second and it is meaningful to have minimum window size = 2 to receive two reports at the clusterhead when *threshold* = 2. The same was found to be true when *threshold* = 3. From the simulation, three seconds window size is abundant to receive reports from other sensors.

To further explore the detection performance with different *thresholds*, Figure 9 is provided. The x-axis is the threshold and the y-axis is the number of detection events occurred normalized by the number of detection events occurred with the same window size and *threshold* = 1.

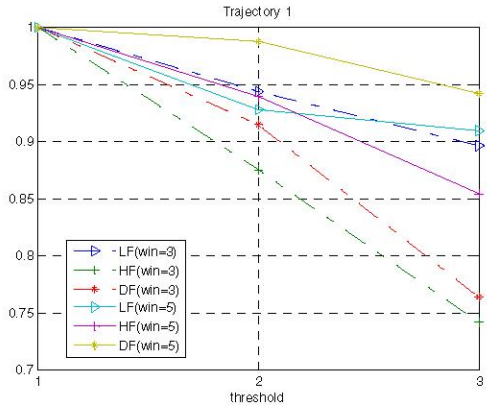

Figure 9. The Normalized detection performance with different threshold values and different window sizes.

The reason for normalizing is that we can clearly observe the relative target detection occurrence reductions due to increased threshold values with fixed sensing detection sliding window sizes. As the threshold value is increased from 1 to 3, the number of detection reports decreases and the percentages drops in all three sensor system architectures. The large threshold value suppresses more detection events from being reported. The reductions were more significant for window size = 3 than window size = 5. By increasing the threshold, about 5 to 50 percent reductions have been observed depending on the architecture and the choice of window size, threshold value, and target trajectories. In general, we can clearly examine that detection performance increases as window sizes increases and threshold decreases.

In addition, we generated independent random false sensing events at each sensor with a probability, P_{fa}, of 0.005 to observe the detection performance under varying false sensing event rates. The combinations of two different *thresholds* and three sensing detection sliding window sizes are used to observe the performance impact caused by increased false detection event rates. Figure 10 shows the total number of false target detection reports received at clusterheads with respect to different sensing detection sliding window sizes and false detection event rates.

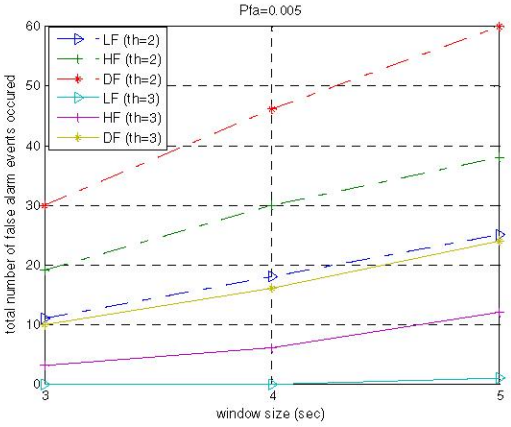

Figure 10. The total number of false target detection reported at clusterhead of each architecture.

As we increase the detection sliding window size, more false target detections are reported in clusterheads, due to the large Δ. Also, higher false detection event generation rates at each sensor increases the total number of false reports for all three architectures. We can observe that the number of false alarms occurred at clusterheads with distributed-fusion architecture is almost 1.5 times more than the hierarchical-fusion architecture and twice more than the localized-fusion architecture when the *threshold* is small (2) and the window Δ is large (5). Therefore, we verified from Figure 10 that the distributed-fusion architecture is more vulnerable to false detection events than the hierarchical-fusion or localized-fusion architectures.

Figure 11 captures the false alarm report suppression with increasing the decision threshold. In Figure 11, the x-axis is the threshold and the y-axis the total number of false target detections reported by the clusterheads. The false detection event rate is fixed at P_{fa} = 0.005.

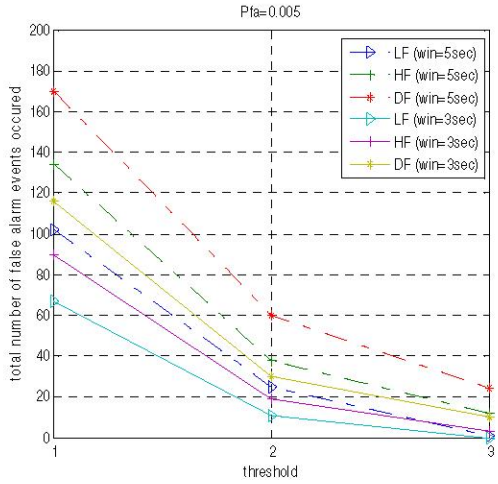

Figure 11. Different false event detection rates vs. the number of false alarms reported by clusterheads.

As we can observe, increasing the threshold eliminates the false reports more than half, particularly, for high false detection event occurrence rate coupled with a large sensing detection sliding window size. Additional investigations (not shown above) using STAMiNA have revealed that for varying P_{fa}, the false alarm report rate scales in both abscissa and ordinate dimension with respect to the threshold parameter.

5. CONCLUSIONS

In this paper, we presented the sensor network simulation tool STAMiNA, which is suitable for evaluating and characterizing different sensor network architectures. We provided examples determining performance for alternative sensor network architectures. The system level detection and false alarm performance were analyzed and trade-offs characterized as a function of decision threshold and sensing window parameter sizing. The simulation results illustrated STAMiNA's capacity to provide accurate system level performance for the complex inter-relationships that occur among sensing and communications networking, and facilitate architectural level decisions for networked sensor system design. We numerically show that the hierarchical-fusion architecture yields the good detection performance as well as robustness against false-alarm. Hence, we can conclude that the hierarchical-fusion is the reasonable sensor network architecture incorporating the advantages of localized and distributed-fusion architecture.

Further research is needed to route target detection information more efficiently in a harsh communication environment. An especially interesting future research pursuit is to explore Disruption Tolerant Networking (DTN), where a store-and-forward architecture is suitable for such operating conditions.

6. ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology and was sponsored by the Army Research Laboratory Advanced Sensors Collaborative Technology Alliance and the National Aeronautics and Space Administration.

REFERENCES

[1] Andree Filipov, Nassy Srour, and Mark Falco, "Distributed and Disposable Sensors at ARL and the ASCTA," in *Proc. IEEE Aerospace Conf.,* March 2004, Big Sky, MT.

[2] L.P. Clare, E.H. Jennings, and J.L. Gao, "Performance Evaluation Modeling of Networked Sensors", in *Proc. IEEE Aerospace,* Vol. 3, March 2003.

[3] C. Krupiarz, E. Jennings, J. Pang, J. Schoolcraft, J. Segui and L. Torgerson, Spacecraft Data and Relay Management Using Delay Tolerant Networking, in *Proc. AIAA SpaceOps,* June 2006.

[4] X. Wang, H. Qi and S. Beck, "Distributed Multi-target Detection in Sensor Networks", Chapter 14 in *Distributed Sensor Networks* (S. Iyenegar, R. r. Brooks Eds), Chapman & Hall, CRC Press, 2005.

[5] T. Clouqueur, K. K. Saluja, P. Ramanathan, "Fault Tolerance in Collaborative Sensor Networks for Target Detection", in *Proc. IEEE Transactions on Computers,* Vol. 53, Issue 3, March 2004, pp. 320—333.

[6] M. Ding, D. Chen, A. Thaeler, and X. Cheng, "Fault-tolerant Target Detection in Sensor Networks", in *Proc. IEEE WCNC,* 2003.

[7] T. Sung, L.-H. Chen, C.-C. Han, M. Gerla, "Reliable Sensor Networks for Planet Exploration", in *Proc. IEEE Networking, Sensing and Control,* 2005.

BIOGRAPHIES

Simon Woo *is a Member of Technical Staff in the communications networks group at Jet Propulsion Laboratory (JPL), conducting space communications and networking research. Prior to joining JPL in 2005, he interned at Entropic Communications in San Diego to evaluate protocol performance. Also, he worked in the wireless computing and communications group and pre-silicon validation group at Intel Corp. He earned his BSEE degree from Univ. of Washington, Seattle in 2003 and MSEE degree from Univ. of California, San Diego in 2005, specializing in communication theory and systems. His primary research interests are the design, simulation, and analysis of communication protocols.*

Esther Jennings *is a research staff with the Communication Systems and Research Section of the Jet Propulsion Laboratory. She received a Ph.D. in Computer Science from Luleå University of Technology, Sweden in 1997. From 1997-1999, she was a postdoctoral fellow at the Industrial*

Engineering Department at Technion, Israel Institute of Technology. From 1999-2001, she has been an assistant professor at the Computer Science Department of California State Polytechnic University, Pomona. Her research interests are in distributed graph algorithms, reliable multicast protocols, energy-efficient algorithms for wireless network and algorithm simulations.

Loren Clare *is the supervisor for the Communications Networks Group at the Jet Propulsion Laboratory. He obtained the Ph.D. in System Science from the University of California, Los Angeles in 1983. His research interests include wireless communications protocols, self-organizing systems, network systems design, modeling and analysis, and distributed control systems. Prior to joining JPL in May 2000, he was a senior research scientist at the Rockwell Science Center, where he acquired extensive experience in distributed sensor networks, satellite networking, and communications protocols for realtime networks supporting industrial automation.*