

Biology-inspired Architecture for Situation Management

Kennie H. Jones, Kenneth N. Lodding, Stephan Olariu, Larry Wilson, and Chunsheng Xin

Abstract— Situation Management is a rapidly developing science combining new techniques for data collection with advanced methods of data fusion to facilitate the process leading to correct decisions prescribing action. Current research focuses on reducing increasing amounts of diverse data to knowledge used by decision makers and on reducing time between observations, decisions and actions.

No new technology is more promising for increasing the diversity and fidelity of observations than sensor networks. However, current research on sensor networks concentrates on a centralized network architecture. We believe this trend will not realize the full potential of situation management.

We propose a new architecture modeled after biological ecosystems where motes are autonomous and intelligent, yet cooperate with local neighborhoods. Providing a layered approach, they sense and act independently when possible, and cooperate with neighborhoods when necessary. The combination of their local actions results in global effects.

While situation management research is currently dominated by military applications, advances envisioned for industrial and business applications have similar requirements. NASA has requirements for intelligent and autonomous systems in future missions that can benefit from advances in situation management. We describe requirements for the Integrated Vehicle Health Management program where our biology-inspired architecture provides a layered approach and decisions can be made at the proper level to improve safety, reduce costs, and improve efficiency in making diagnostic and prognostic assessments of the structural integrity, aerodynamic characteristics, and operation of aircraft.

Index Terms—Autonomous, Intelligent, Distributed Systems, Sensor Networks, Situation Management.

Manuscript received September 1, 2006.

Kennie H. Jones is with the National Aeronautics and Space Administration, Hampton, VA 23681 USA (phone: 757-864-6720; e-mail: k.h.jones@nasa.gov).

Kenneth N. Lodding is with the National Aeronautics and Space Administration, Hampton, VA 23681 USA (e-mail: kennethnlodding@nasa.gov).

Stephan Olariu is with the Computer Science Department, Old Dominion University, Norfolk, VA 23529 USA (e-mail: olariu@cs.odu.edu).

Larry Wilson is with the Computer Science Department, Old Dominion University, Norfolk, VA 23529 USA (e-mail: wilson@cs.odu.edu).

Chunsheng Xin is with the Computer Science Department, Norfolk State University, Norfolk, VA 23504 USA (e-mail: cxin@nsu.edu).

I. INTRODUCTION

Situation management is a rapidly evolving science combining new techniques in *situation awareness* (i.e., collection of time/space/state data about a situation) with advanced techniques in data fusion to facilitate timely, correct decisions for action on the situation. Research in situation management is supported largely by military requirements. Military conflicts are increasing in mobility, velocity, complexity, and the dynamic nature of their situations. More effective methods are needed for situation monitoring, awareness, and control to provide command options, predict probable situation outcomes, and analyze potential threats and vulnerabilities.

However, such requirements are not limited to military applications. Indeed, any activity that requires correct, decisive action in a timely manner including homeland security, emergency/crisis management, manufacturing processes, financial management, medical, among many others, can benefit from the research in situation management [1].

The National Aeronautics and Space Administration (NASA) is researching how the techniques under development in situation management can contribute to future missions in Earth, planetary, and space science, mission control, and vehicle/equipment health management.

Current research in situation management is focused on developing techniques for fusion of large and diverse data with the goal of providing reduced and consolidated information manageable by humans. Current research in sensor networks, a promising tool for improving situation awareness, is concentrating on networking technologies with a centralized flavor. Our purpose here to identify the need for more research in architectures for the implementation of large scale, operational systems for situation management where many decisions are made autonomously without human intervention. Such systems are a requirement to meet the needs of future NASA missions.

In Section II, we define the goals of and requirements for situation management as related to this discussion. In Section III, we describe the current state of sensor network research and how the direction of that research limits their use for situation management. In Section IV, we describe the goals and requirements of a NASA project to develop technologies for aircraft health management, and their relationship to situation management. In Section V, we describe an

architecture for situation management modeled after biological systems that combines local autonomy with communal cooperation to effect timely action. In Section VI, we offer our conclusions.

II. GOALS AND REQUIREMENTS FOR SITUATION MANAGEMENT

Jakobson et al. [1] described the critical aspects of situation management as “managing and controlling sources of information, processing real-time or near real-time streams of events, representing and integrating lower level events and higher level concepts, multi-source information fusion, information presentation that maximizes human comprehension, and reasoning about what is happening and what is important.” They defined a situation as a “large number of dynamic objects that change state in time and space and engage each other into complex spatio-temporal relationships”. The state of an object is expressed as a set of parametric values. The situation is a state of one or more objects at an assigned time. Therefore, two situations are distinguished by both the state of the objects and the time. Situations change as states and/or time changes [2]. This presents major problems for managing situations:

- *Information freshness*: If information is collected and time passes before a decision can be made using that information, the decision may be incorrect.
- *Information quality*: Information may be redundant, incomplete, or irrelevant.
- *Information overload or starvation*: Information may come in high volumes at times, little to none at other times.
- *Information fidelity*: Analysis of many situations requires massive inputs that can only be acquired in situ.

Jakobson et al. [1] used a chart shown in Figure 1 in their overview of situation management. They identified this as commonly used in the military and intelligent control. It does illustrate much of the process in situation management, depicting levels of complexity among tasks. *Sensors* collect information from *the world* and *effectors* act on it. A measure of increasing sensory perception is indicated by the vertical label *S*; increasing reasoning by the horizontal label, *R*; increasing control through the vertical label, *C*. The *levels* indicate tasks where *Level i+1* is more complex than *Level i* (e.g., requires more sensing and/or reasoning results in more control). The horizontal label, *R*, is related to time, as there is a lapse in time between sensing and control. It is this time that is required to apply fusion techniques to sensed data to determine action. One of the research objectives is to reduce the time required to apply reasoning (i.e., improve information freshness).

There are a number of concerns that are either not clear or not shown by this figure. Although higher level tasks appear to require more sensing and result in more control, all levels appear to require the same amount of reasoning. It seems that simpler tasks should require less reasoning in less time. It

appears that lower level sensing is passed on to higher level tasks but it is not clear how completion of lower level tasks affect higher level tasks. If *Level i* completes, and acts on the environment, before *Level i+1*, then its effect must somehow be factored into *Level i+1*'s process, setting up feedback loops. These concerns will be addressed in Section V.

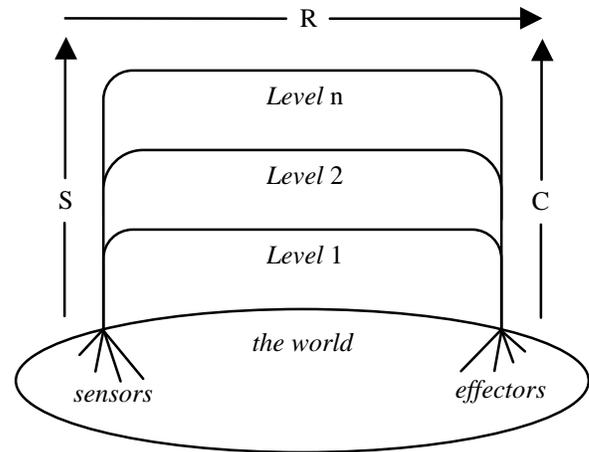


Fig. 1. The big picture of Situation Management.

The *R* scale currently depicts the major focus in situation management research. That is, given a set of sensor input, develop a reasoning algorithm that will determine the correct action in a required time. This in itself is a formidable task when the input set is large and diverse. It is understandable that these problems need to be solved first without architectural concerns. However, for a practical implementation of a large system, we envision an architecture where simpler tasks can process autonomously while feeding appropriate information to more complex tasks. This will be discussed in detail in Section V.

III. LIMITS OF CURRENT SENSOR NETWORKS RESEARCH

No new technology shows more promise for meeting these demands than sensor networks [3,4,5,6,7]. New techniques are rapidly reducing size and cost of devices that sense (or act), compute, and communicate under their own power. Low cost allows massive deployment, increasing the diversity and fidelity of observation. Small size enables unintrusive, inconspicuous deployment required for many situation management applications. However, the implementation of sensor networks simply as sources of input or as effectors will not solve, but actually contribute to, the problems listed in Section II. Additionally, Jakobson *et al.* [1], listed new problems posed by massive deployment of sensors:

- Configuration of the network changes frequently.
- Each node has limited computation, memory, power, and communication capabilities.
- There can be a high rate of node and link failure.
- Communication is limited by internode communication through multihop broadcasting.
- Functionality is enhanced by or requires an optimal spatial

distribution.

- Nodes are limited in situation recognition; the final operational situation must be compiled from fusion of input from a large numbers of nodes.

In previous work [8,9], we have identified several problems with the direction of current sensor network research that we summarize here. The sensor network community is largely concentrating on developing techniques for organizing a randomly deployed sensor network to establish a communication system for routing information to a centralized information *sink* (i.e., processor) in a severely energy constrained environment. This is favored approach because there is significant experience with hierarchical computer architectures. However, there are difficulties in building a real situation management system upon this model. Without a significant amount of redundancy, hierarchical system architectures are brittle and potentially possess multiple single points of failure. Sensed data is forced to travel thru multiple system levels and processes before it becomes available to influence, or effect change in the situation being monitored. This offers the potential for the decision being made at time T2 to be an inappropriate response for the information sensed at time T1.

These techniques are adaptations of the centralized networking technologies that are widely used in the Internet. The result of many designs is that data is sensed and passed via the multihop routing path to the sink. Furthermore real implementations are not reaching the massive numbers of nodes predicted [10,11,12]. Even simulations are in the order of hundreds of nodes. It is not evident that the techniques being developed will scale to the massive deployments that will be required of systems such as those needed by situation management. More importantly, we contend that the architecture where sensor nodes simply pass collected data to a central repository will not realize the full potential for sensor networks in such applications as situation management. Our alternative is proposed in Section V.

IV. GOALS AND REQUIREMENTS OF IVHM

NASA's Aviation Safety Program within its Aeronautics Research Mission Directorate is initiating a sub-program, Integrated Vehicle Health Management (IVHM) with the goals of improving safety, reducing costs, and improving performance in every aircraft class.

Concerning safety, an examination of recent aircraft incidents and accidents reveals the need for such an effort. The Commercial Aviation Safety Team (CAST) [13,14] identified system/component failure or malfunction as the third leading cause of crashes in examining commercial jet airplane accidents worldwide from 1987 through 2004. They found other leading causes to be fire, ice, fuel, wind shear, and lightning. These classes of accidents were identified as causing 24% of the total of all accidents. Loss of control caused 26%, but hardware and software failures are often contributing

factors in loss-of-control accidents.

Examining data from US-registered transport aircraft accidents from 1980 to 2001, the National Transportation Safety Board (NTSB) [15] found 52% of hardware-induced accidents were aircraft systems related, 36% were caused by propulsion system components, and 10% were caused by failures in the airframe. Of these, approximately one third each was due to problems with landing gear, turbine/turboprop engines, and flight control systems.

The Federal Aviation Administration (FAA) [16] examined 40,964 incidents involving US airplanes from 1998 through 2003. They found that about 67% of incidents were caused by a combination of system failure and malfunction, fire/smoke, and power loss.

CAST concluded that not all incidents and accidents were directly due to system failure but resulted from the failure of flight crews to:

- correctly interpret, process, and cross-check available, relevant data, assess failure modes and analyze effects;
- maintain aircraft system status awareness;
- understand the impact of inoperative or degraded systems.

Furthermore, sensors and other equipment failed to:

- accumulate and present adequate trend information;
- indicate impact and other damage;
- indicate and prevent icing;
- provide warning of unsafe flight critical systems.

In providing solutions for these problems, more information is not the only requirement. CAST was concerned that warnings and equipment failure announcements presented to the crew should not cause a nuisance that would contribute to crew complacency. They also identified that part of the solution is to provide real time information to ground crews.

In summary, the evidence reveals that a high percentage of aircraft accidents and incidents are caused directly by equipment failure or indirectly by the inability of crew to properly and timely manage situations of equipment failure. IVHM is tasked with developing technology that will mitigate both of these problems. We recognize that advances in both sensor network technology and situation management technologies are necessary to meet the goals of IVHM.

Future aircraft will be designed to sense, control, communicate, and navigate with increasing levels of autonomy. Automatic health monitoring combined with self-healing systems in aircraft will not only improve safety but also performance, reliability, and predictability of service while reducing costs. IVHM is tasked with developing many of the technologies required for such future aircraft. An essential component is improving diagnostic capabilities that form the basis for prognostics in airframe, propulsion, and other aircraft systems. IVMH is investigating fundamental failure physics and associated effects of damage and degradation caused by environmental hazards. The application

of situation management techniques using this knowledge steered by information collected real-time on operating characteristics will provide diagnostics that will determine the prognosis that is used for failure mitigation. From a systems perspective, IVHM also addresses challenges in communication and effective architectures to facilitate the integration of IVHM components with each other and with other vehicle systems.

IVHM represents a constrained microcosm of the larger problem domains more generally associated with the traditional study of situation management. The problems associated with providing a successful IVHM environment mirror the more general problems associated with situation management. At its most fundamental level of implementation, IVHM is the concept of instrumenting an aerospace vehicle with a web of large numbers of sensors to report on the health of its constituent parts. The simplest architecture for such a web would have the multitude of sensors sending their data to a central computer for processing. The single central processor must perform all situation management functions: data reduction, conversion, fusion, and presentation to the pilot or flight engineer. This approach has a number of shortcomings:

- Single point of failure is at the central computer.
- Potentially large amounts of uninteresting data (e.g. current jet exhaust temperature) will unnecessarily burden the central computer.
- Data that is most meaningful when viewed in conjunction with other data, rather than in isolation must be segregated and fused with other information for presentation. This consumes processor cycles and memory that could be used for other tasks.
- Data that is not currently needed, or not important to the current mission, but needs to be made available to a different user at a later time will still flow to the central computer, further absorbing available processor capabilities: cycles and memory.

In Section V, we provide an alternative model, which is inspired by biological ecosystems.

V. AN ARCHITECTURE FOR SITUATION MANAGEMENT

As we will demonstrate, we consider sensor nodes as organisms interacting with their environment. They function autonomously, yet cooperate with local neighboring sensor nodes. This architecture provides many advantages over the conventional, centralized approach:

- Multiple, independent processors reduce the chance for catastrophic failure through single point of failure.
- Multiple sensors may be logically joined to provide virtual sensors, which offer the potential for inferring data that would be otherwise unavailable.
- Local sensors can hold data for alternate users, such as maintenance workers, without burdening the central processor, or flight crew with unnecessary, yet

important detail.

- Local sensor groups can rapidly make decisions in highly dynamic situations, such as flight anomalies, to respond to anomalies with suggested recovery or mitigation actions.

We recognize that IVHM is more tightly constrained than larger situation management environments, such as in battle management. However, the benefits of the ecological approach, which are clear for IVHM, seem to us to be appropriate for large-scale deployments of sensors and processors in an isomorphic manner. Distributing decision making to the proper level in the network will be the real challenge. To meet this challenge, we believe conventional techniques of computer science are not sufficient. We look to the success of living systems to provide architectural models. In the next sections, we describe two such models and how they apply to our architecture.

A. Biological Model: Animal Response to Injury

We believe that a fundamental part of the solution to building a successful situation management system is in the architectural design of the base sensor network. In our view, the sensor network, rather than simply being the foundation upon which the situation management pyramid, with all of its' information flowing upward for processing, is in reality the initial layer for filtering, processing, and possibly responding to locally detected events. Our sensor network model is biologically inspired, having more in common with a multicellular organism than a traditional hierarchical computer network. Raw data can be shared among the sensor cells and locally processed and fused, constrained to the onboard communication, processing and memory capabilities of the individual sensor nodes. As appropriate, cells may simply pass data up to the next level of organization, or may respond immediately to local events. A good analogy to the architecture is the multi-level mechanism by which the human body protects itself: a cut in the skin occurs, while nerves carry the information upward, local chemistry immediately responds to the situation: blood coagulates and white cells gather at the trauma site. At some intermediate level, the body generates signals to the finger-hand-arm combination to pull away from the offending object to prevent further damage. At the topmost level, awareness of the situation invokes further review by the individual who then decides whether the cut can be safely ignored, requires simple treatment (e.g., application of a bandage), or demands professional medical attention. Finally, at the highest level, the environment is assessed for action that may prevent injuries in the future. Key to this model is that local decisions at each level promote an appropriate response for the global benefit, with a minimal time delay. The measure of the success of our architecture is costs of time between observation and action, power consumption, and communication and in the mitigation of single point of failure compared with a more centralized approach. In the next section, we describe another biological model for autonomous

action and communal cooperation among component subsystems

B. Biological Model: An Ecosystem

A situation management system can be thought of as a *community of organisms* interacting with their *environment* forming an *ecosystem*. Each sensor node, sinks, and other subsystems represent an organism in the community. At birth (i.e., at deployment time) these organisms are endowed with genetic material, containing, among others, an initial state and rules by which they interact with each other and the environment. The state and the rules may change as these organisms interact with the environment, reflecting their dynamic adaptation to conditions in their neighborhood. Additionally, the organisms may remember and record their interaction with the environment by storing information in their limited on-board memory. Memory and its use to change state or rules are considered learning. Changing state conditions or rules based on learning demonstrates some level of cognition.

A major feature of this architecture is to use learning and cognition of each organism, their autonomy, and their cooperation to benefit the functionality of the system for situation management. In particular, we are interested in using these attributes to enable local decisions based on local information that effect global results. Limiting decisions to localities is important for reasons of scalability and autonomy. Local decisions allow distributed control. In turn, distributed control through local decisions provides a natural redundancy affording fault tolerance – as some organisms fail, others will continue to make decisions.

C. Biological Model Applied to Situation Management

The scope of problems that an IVHM sensor system must detect and respond to is extremely wide and diverse. At one end of the event spectrum are problems requiring an immediate response to prevent catastrophic failure, while at the other end are events allowing a delayed response (e.g., preventative maintenance requests). Between the extremes lie a large class of problems, some diagnostic and some prognostic, to which an integrated health management system is required to identify, process, and react appropriately.

For those problems where the sensor-effector network must respond to an event as quickly as possible, event detection, identification, resolution and response activity should be as close as possible to the event trigger, where distance is measured in time delay. The rationale for this is diagrammed in the Figure 2. Here two responses are made to the event. The first response is made by the sensing node autonomously, assuming all fusion and reasoning is done by that node. The second response is made using the conventional method of sensor networks: multihop messaging to a central sink where fusion and reasoning is performed, decision is made, and commands sent back to an effector.

A time delay can lead to two distinct problem situations. In

the first instance, the response to a critical situation arrives too late to correct the problem, due to the long processing time, and although a proper response was generated, a catastrophic failure still results. The second instance occurs when the response generated for the original event is no longer appropriate because the situation has altered from the sensed state. The diagram illustrates the simple truth that the greater the separation between the source event and response event, the greater the delay in resolving the problem. It is the desire to resolve this time delay problem, along with the goal of avoiding an implementation flawed by a single point of failure architectural style, which argues for pushing intelligence out from a central computer to *smart* sensor-effector nodes.

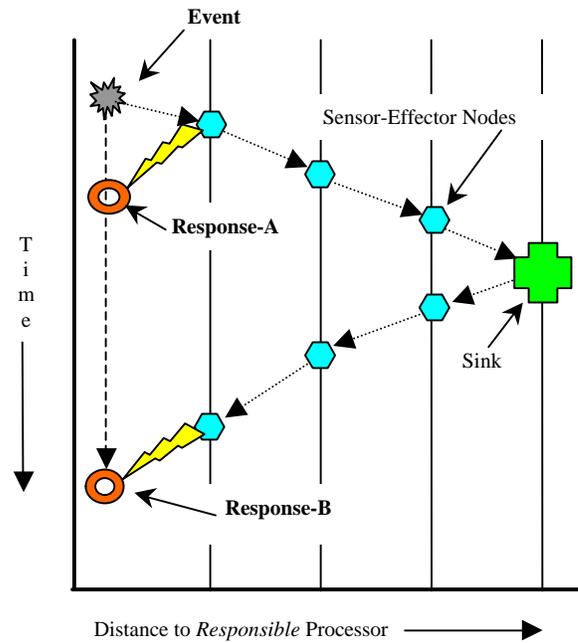


Fig. 2. Local control vs. central control

The broad continuum of potential prognostic and diagnostic event states to be detected and responded to for an aerospace vehicle lends itself to a structured response solution in which the appropriate handler attempts to resolve the problem, while still allowing for refined inputs to higher level resources. In this scenario, the event is rapidly filtered at the sensor node to determine whether to immediately react and simultaneously pass the event information onward, or to simply pass the event information onward, with no local reaction. The model is loosely based on that described for the human body's response to damage. Applying this concept to Figure 1 is shown in Figure 3.

The challenge of this architecture is to design subsystems that can function at appropriate levels. The objective of situation management is often described as transforming large amounts of diverse data into information to be used by humans to make better decisions. An astute reader may observe Figure 3 and deduce that such an architecture cannot function if humans are to be directly involved at each level. If lower levels are to act quickly, they must also act autonomously.

These subsystems must be designed to function from sensing to fusion and reasoning to deciding and acting without human intervention. Such an architecture is imperative for the complexity of IVHM and we contend it is for other situation management applications. Pilots cannot be directly involved in every decision that will be made in a successful IVHM implementation. In the next two sections, we describe specialized organizations of sensor networks that we consider imperative in the implementation of this architecture.

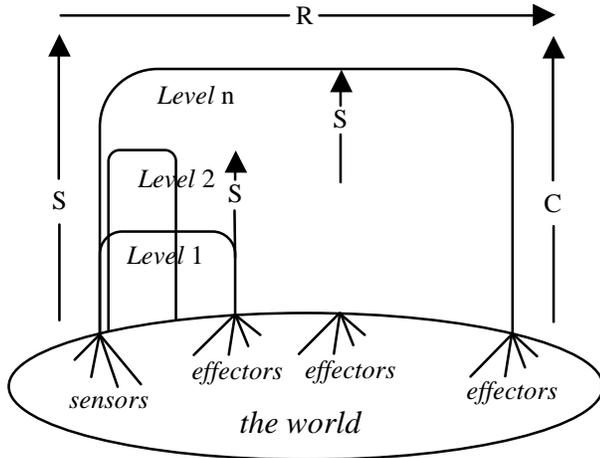


Fig. 3. Revised big picture of Situation Management

D. Virtual Sensors

Physical sensors measure a property, condition, or state of a process or object. In contrast, *virtual sensors* infer a value based upon a combination of a software algorithm and input from “other” physical sensors. “Other” here refers to sensors that do not directly measure the attribute of interest, but instead measure other attributes, which are used to calculate an estimated value for the property of interest.

Intelligent Virtual Sensors (IVS) offers the possibility of developing unique sensing devices which combine information from multiple sources to infer environmental conditions where there is no sensing device available to directly measure the phenomenon of interest, or where environmental conditions exceed the capabilities of available measuring devices, or where the placement of a measurement device is precluded by physical space limitations, or where the introduction of a suitable device might compromise structural integrity of the local environment.

The intelligent virtual sensor is unique in that it offers the possibility of extending the reporting capability of a traditional sensor into a prognosticating ability. By combining and analyzing a number of different measurement variables available to it and its neighbors and applying AI techniques to analyze the information, it is possible that an IVS could infer some future state for the monitored environment. This would be extremely valuable in such applications as predicting aircraft structural or engine failure.

E. Communal Cooperation Within Sensor Networks

This section summarizes our continuing work in which we look at sensor networks in a novel way, motivated by our belief that in order to scale to massive deployment, sensor networks can benefit from lessons learned from the way biological ecosystems are organized. Indeed, in the presence of a massive deployment, sensor networks must behave as a community of organisms, where individual sensor nodes, or *motes*, operate asynchronously and autonomously in parallel. We focus on fundamental characteristics of future sensor networks that are not demonstrated by current implementations, yet are imperative for optimal use in situation management. More specifically, we demonstrate that in such a model, fully distributed data aggregation and integration can be performed in a scalable fashion in massively deployed sensor networks, where individual motes operate based on local information, making local decisions that are aggregated across the network to achieve globally-meaningful effects.

This approach mitigates many of the problems identified for massively deployed networks. Furthermore, this approach, with the benefit of rapidly increasing computational, power, and communication capabilities of motes, will mitigate many of the problems identified for managing situations by allowing much of the decision making task to occur closer to the source of information, improving the freshness of information and decisions.

It is not our claim that the examples presented here directly apply to any problem of situation management, but rather that the application of sensor networks to situation management may benefit from our alternate approach. Particularly in adversarial scenarios, centralized control of *in situ* networks, with its complex infrastructure and single points of failure, may be quickly defeated. Our decentralized approach is much more robust and fault tolerant and would, therefore, be much more difficult to disable.

Our basic premise takes advantage of the true nature of the broadcast environment in a sensor network: each sensor node can only directly communicate with neighbor nodes within its limited transmission distance. Any information it obtains from other nodes must come indirectly through multi-hop routing. Our algorithms then limit information exchange among neighborhoods of nodes that can communicate. This exchange occurs as if they are the only nodes in the network. The interested reader must obtain details of the algorithms from references; only summaries are provided here.

Suppose sensors nodes are deployed with randomly distributed value. The objective is for the network is for all nodes change this value to the average of all nodes [17]. This task would be accomplished simply by conventional means. All values could be sent via multi-hop routes to a sink that would compute the average and distribute this value to all nodes. This centralized approach has distinct disadvantages that we have documented. Our approach has each node communicate only with its neighborhood, calculate an average,

and set only the members of the neighborhood to the neighborhood average. Through a series of iterations, each node in the network eventually is set to the global average at initial distribution.

By a similar method, we solved the problem of majority rule [18]. If at distributions, each node is set to a binary value, the objective is for each node to reset itself to the initial majority value. Again, this is easily solved by a centralized approach, but this has been documented to be a difficult problem if only local neighborhoods are examined [19,20,21].

We have also compared our similar approach to time synchronization to conventional, more centralized approaches [22] and shown that it can be accomplished using only local information and local actions.

An essential requirement for scalability, we have proven that our approach does not require more time and does not require more energy per node for a solution as network sizes grow. This is not true for a centralized approach.

In all of these examples, global tasks are accomplished by interactions limited to local neighborhoods. Such must be the design if subsystems will function autonomously. This is the direction in which we believe sensor network research must proceed to facilitate the kind of systems required for IVHM.

VI. CONCLUSION

Situation management has application in so many areas and will provide the vehicle to realize much of the promised automation that has thus far not been possible. However, to realize full potential, techniques must not be limited to fusion of data into information to be used directly by humans. Fusion techniques must provide reasoning sufficient for decisions to be made and actions to be taken autonomously at a subsystem level.

NASA's IVHM program is an example of the need for this autonomy and cooperation among subsystems. Unmanned space vehicles have a stronger need: it is impractical for future space missions of such vehicles to be controlled by humans. They must be able to operate autonomously, adapting to possibly unexpected environmental challenges. The situation management community is already providing part of the solution. But the current work must be embedded in an architecture that can fulfill these needs.

REFERENCES

- [1] G. Jakobson, L. Lewis, C. Matheus, M. Kokar, and J. Buford, "Overview of Situation Management at SIMA 2005", retrieved May 9, 2006 from <http://vistology.com/papers/SIMA2005.pdf>
- [2] G. Jakobson, J. Buford, and L. Lewis, "Towards an Architecture for Reasoning about Complex Event-Based Dynamic Situations", International Workshop on Distributed Event-Based Systems (DEBS '04) at 26th International Conference on Software Engineering (ICSE 2004), Edinburgh, Scotland, UK, May 24-25, 2004.
- [3] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Next century challenges: Mobile support for Smart Dust", Proc. ACM MOBICOM, pp. 271-278, Seattle, WA, August 1999.
- [4] B. Warneke, M. Last, B. Leibowitz, and K. Pister, "SmartDust: communicating with a cubic-millimeter computer", IEEE Computer, 34(1), pp. 44-55, 2001.
- [5] V. V. Zhirnov and D. J. C. Herr, "New frontiers: self-assembly and nano-electronics", IEEE Computer, 34(1), pp. 34-43, 2001.
- [6] I. F. Akyildiz, W. Su, Y. Sankarasubramanian, and E. Cayirci, "Wireless sensor networks: A survey", Computer Networks, 38(4), pp. 393-422, 2002.
- [7] D. Culler, D. Estrin and M. Srivastava, "Overview of sensor networks", IEEE Computer, 37(8), pp. 41-49, 2004.
- [8] Jones, K., Lodding, K., Olariu, S., Wilson, L., and Xin, C., "Biology Inspired Distributed Consensus in Massively-Deployed Sensor Networks", Proceedings of 4th International Conference on AD-HOC Networks & Wireless (ADHOC NOW 2005), Cancun, Mexico, October 6-8, 2005.
- [9] Jones, K., Lodding, K., Olariu, S., Wilson, L., and Xin, C., "Biology Inspired Approach for Communal Behavior in Sensor Networks", presented at Minitrack on Wireless Sensor Networks and Applications, Proceedings of 39th Hawaii International Conference on System Sciences (HICSS-39), Kauai, Hawaii, January 4-7, 2006.
- [10] A. Mainwaring, J. Polastre, R. Szewczyk, and D. Culler, "Wireless Sensor Networks for Habitat Monitoring", Intel Research, IRB-TR-02-006, 2002 ACM International Workshop on Wireless Sensor Networks and Applications, June 10, 2002, Retrieved April 5, 2004, from <http://www.greatduckisland.net>
- [11] D. Lammers, "Embedded projects take a share of Intel's research dollars", EE Times, August 28, 2001. Retrieved April 5, 2004, from <http://today.cs.berkeley.edu/800demo/eetimes.html>
- [12] UCB/MLB 29 Palms UAV-Dropped Sensor Network Demo, 2001, University of California, Berkeley. Retrieved April 5, 2004, from <http://robotics.eecs.berkeley.edu/~pister/29Palms0103>
- [13] Commercial Aviation Safety Team (CAST) Safety Plan, October 7, 2004.
- [14] Commercial Aviation Safety Team (CAST) Update, February 2006.
- [15] NTSB accident database, Retrieved May 15, 2006 from <http://www.ntsb.gov/ntsb/query.asp>
- [16] FAA AIDS database, Retrieved May 15, 2006 from http://www.nasdac.faa.gov/portal/page?_pageid=33.36440&_dad=portal&_schema=PORTAL
- [17] Jones, K., Lodding, K., Olariu, S., Wilson, L., and Xin, C., "Energy Usage in Biomimetic Models for Massively-Deployed Sensor Networks", presented at 1st International Workshop on Mobile Ad-hoc and Ubiquitous Sensor Networks (MAUSN 2005), Proceedings of 3rd International Symposium on Parallel and Distributed Processing and Applications (ISPA 2005), Nanjing, China, November 2-5, 2005.
- [18] Jones, K., Lodding, K., Olariu, S., Wilson, L., and Xin, C., "Biomimetic Models for an Ecological Approach to Massively-Deployed Sensor Networks", Proceedings of ISCA 18th International Conference on Parallel and Distributed Computing Systems (PDCS 2005), Las Vegas, Nevada, September 12-14, 2005.
- [19] M. Mitchel, J. Crutchfield. and R. Das, Computer science application: Evolving cellular automata to perform computations, in T. Bäck, D., Fogel and Z. Michalewics, (Eds.), Handbook of Evolutionary Computation, Oxford University Press, 1997.
- [20] M. Mitchell, An Introduction to Genetic Algorithms, The MIT Press, 1999.
- [21] J. Epstein, Learning to be thoughtless: Social norms and individual computation, Center on Social and Economic Dynamics Working Paper No. 6, revised January 2000.
- [22] Jones, K., Lodding, K., Olariu, S., Wilson, L., and Xin, C., "An Ecosystem Model for Massively-Deployed Sensor Networks", Proceedings of 2nd International Conference on Intelligent Sensors, Sensor Networks and Information Processing, Melbourne (ISSNIP 2005), Australia, December 5-8, 2005.