

AFRL-IF-RS-TM-2006-6
In-House Interim Technical Memorandum
October 2006



COGNITIVE NETWORK COMPUTER FOR ATMOSPHERIC MONITORING

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Rome Research Site Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-IF-RS-TM-2006-6 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

LOIS D. WALSH, Chief
Advanced Computing Technology Branch
Advanced Computing Division

/s/

JAMES A. COLLINS, Deputy Chief
Advanced Computing Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.</small>					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) OCT 2006		2. REPORT TYPE Interim		3. DATES COVERED (From - To) Jul 04 – Jul 06	
4. TITLE AND SUBTITLE COGNITIVE NETWORK COMPUTER FOR ATMOSPHERIC MONITORING				5a. CONTRACT NUMBER In-House	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 63789F	
6. AUTHOR(S) Douglas Holzhauer				5d. PROJECT NUMBER 459T	
				5e. TASK NUMBER CN	
				5f. WORK UNIT NUMBER CA	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFRL/IFTC 525 Brooks Rd Rome NY 13441-4505				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/IFTC 525 Brooks Rd Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TM-2006-6	
12. DISTRIBUTION AVAILABILITY STATEMENT <i>APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 06-739</i>					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The Air Force has an interest in advanced technology development aimed at monitoring environmental conditions. Weather is one specific environmental condition of significant interest to the military since it affects so many missions. Knowledge of atmospheric conditions in general, and knowledge about specific weather conditions in mission-critical regions, can be very beneficial to operational forces. However, specific atmospheric conditions are not always easy to obtain. Furthermore, timely processing and interpretation of weather data are critical issues. Current technologies can be used to provide information to users, but in order for it to be most useful, the information must be organized and managed in a timely fashion. Newer technologies related to cognitive science may process information into more significant and useful entities, toward the production of higher level concepts and “knowledge.” We describe a network of computationally rich sensor nodes grouped into small clusters and communicating with a console and then to the user (meteorologist).					
15. SUBJECT TERMS Artificial Intelligence, sensor networks, weather sensors					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 17	19a. NAME OF RESPONSIBLE PERSON Douglas Holzhauer
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)

TABLE OF CONTENTS

1	INTRODUCTION	1
2	OVERVIEW	2
2.1	MILITARY REQUIREMENTS	2
2.1.1	<i>Air Force Needs</i>	2
2.1.2	<i>Operational Scenario</i>	2
2.1.3	<i>Technical Problems</i>	3
2.2	RESEARCH AND DEVELOPMENT	3
2.2.1	<i>Technology Integration</i>	3
2.2.2	<i>Cognitive Computing</i>	3
2.2.3	<i>Global Patriot Demonstration</i>	3
3	APPLICATION TECHNOLOGIES.....	4
3.1	COMPUTATIONALLY RICH SENSOR NODE (FIGURE 1).....	4
3.1.1	<i>Weather Sensors</i>	4
3.1.2	<i>PASTA Nodes</i>	4
3.1.3	<i>Communication /OLSR</i>	4
3.2	BLACKBOARD ARCHITECTURE (FIGURES 4 AND 5)	4
3.2.1	<i>Control Shell</i>	5
3.2.2	<i>Repository/Objects</i>	5
3.2.3	<i>Knowledge Sources and Functions</i>	5
3.3	COGNITIVE COMPUTING CONCEPTS	7
3.3.1	<i>Information Processing</i>	7
4	FUTURE PLANS.....	8
5	REFERENCES	13

1 Introduction

Research and development in ad hoc sensor networks centers on the development and control of large wireless sensor networks. The individual sensor nodes have processors with limited memory, no long term storage and are designed to accomplish minimal analysis of the data. Typical of these sensor nodes are the “Motes” developed by Berkeley University under DARPA’s Smart Dust Program. The goal of this program was to develop small, disposable sensor nodes. Due to such research, sensor nodes are becoming increasingly smaller, but maintaining a limited analytical capability. This technology is appropriate for many applications. Indoor environmental monitoring and mechanical strain measurement are just two possibilities. However, there are other classes of problems where the phenomena being measured are complex (e.g., 3D, time varying), vary over large distances and produce large volumes of data that must be processed before the phenomena can be understood. The phenomena associated with the physical environment fit into this class.

The Air Force has interest in sensor networks for environmental monitoring, as knowledge of atmospheric conditions would increase air drop precision and all-weather landing safety. Atmospheric phenomena are characteristically complex with slowly changing spatial gradients. Because of the complexity, mathematical models supported by small numbers of observations can not quantify environmental variations with 100% accuracy. In many cases there is a need for such extreme accuracy. Large ad hoc sensor networks can give operational troops the real time high accuracy environmental data needed. Current research, by both DARPA and AFRL’s Sensors Directorate, is leading to the development of sensors that are suitable for environmental monitoring. By using these sensors as the constitutive components an ad hoc sensor network could be formed. Inter-nodal communication can be done using the IEEE 802.11 commercial standard, which meets the one to two kilometers range required by the scale of weather phenomena. The physical size of these sensors is quite large (compared to the tiny Motes) as the measurements have to be taken at a distance from the ground. Research into making the sensors sufficiently rugged, to enable them to be air dropped into forward air drop zones, and to be power aware, to conserve their battery prime power, is continuing at DARPA and AFRL.

A network of this scale exposes two problems: 1) The essence of ad hoc sensors is that they are randomly distributed, air dropped is implicit, which means we can now obtain data from regions that were previously closed to us. However this also means the sensor must autonomously determine the quality of the data that it is sending back to the Command and Control (C2) system. This is dependent upon the sensor node having specific location and environment knowledge and some level of internal cognitive ability, and 2) Even though ad hoc sensor networks promise the capability of furnishing an unprecedented amount of data, the amount is so large that it would surpass the capability of even a midsized system’s wireless communication channel and the ability of (current) C2 systems to process the data into the decision quality information in a timely manner. Our solution to these information management problems is to convert the sensor network into an autonomic wireless network computer and to accomplish much of the information processing within the ad hoc network.

This research project examines the development of a medium sized (hundreds of nodes) autonomic wireless network computer and concentrates on the measurement and processing of atmospheric observations. Individual network nodes comprise several heterogeneous sensors.

Emphasis is placed on the process of network computing so scalability issues will be addressed but not experimentally demonstrated. Sensor nodes use commercial Wireless Fidelity (WiFi) for communication. This project leverages research of AFRL's Sensors Directorate and the DARPA's Power Aware Computing and Communication project for the sensors and computers that form the sensor node hardware. Placing this research into an operational context reveals needs particular to the operational military environment. This includes examining the "information last mile" to assure that ad hoc sensor networks can become an integral part of the command and control system. Research addresses issues that enable the development of autonomous sensor nodes and network computing.

2 Overview

2.1 Military Requirements

In making the case for the development of an intelligent sensor network that would provide the Air Force with weather observations, we have reviewed literature and interviewed combat meteorologists to validate the need for such a system. Indeed most sensor network research has emphasized the detection of personnel or vehicles by radar, acoustics or IR sensors. In this section we make the case for a weather sensor network.

2.1.1 Air Force Needs

The Air Force mission is greatly affected by the weather conditions where it operates. Almost all types of missions are affected by the weather but the poor weather conditions affect some more than others. Some of the missions that are more weather dependent include precision air drops of material or combatants, unattended air vehicles, low level air operations by helicopters or fixed wing aircraft in mountainous terrain, helicopter landing and takeoff in high altitudes and temperatures, propagation of chemical weapons and spills and large fire fighting activities. In each of these cases the weather conditions inordinately affect the success of the mission associated with it.

2.1.2 Operational Scenario

In the current context, the weather prediction that is used to make mission decisions is accomplished by the Air Force Weather Agency using sophisticated models such as MM5, Mesoscale Model 5.¹ Using the new Weather Research & Forecasting model, which incorporates observations into the MM5 model, can increase the prediction accuracy. Current observations that are fed into MM5/WRF are acquired by satellites or by land based radar. Naturally, the closer the observations that are used as inputs to MM5/WRF the more accurate the prediction will be. Even when used in combination, these sensors can exhibit serious limitations when the area of interest is located in a remote isolated region. Cloud cover can mask the lower elevation weather parameters, and the curvature of the earth quickly negates ground radars from observing lower portions of the troposphere. In the case of mountainous terrain, large geographical changes over small distances can prevent even the best models from accurately determining the weather.² Solutions to the last scenario has been to install the automated weather observation system (AWOS) in mountain passes.³ Of course this system has constant power supplied and is constantly reporting weather conditions. Such a system can not exist in a tactical combat location. However that functionality is what is needed.

2.1.3 Technical Problems

The development of a tactical weather network has many problems associated with it. The current tactical weather station, the AN-TQM-53 is by its very nature not up to the task because of its packaging, cost and lack of air drop capability. But our literature research has shown that both the Navy and the Army have been doing research and development leading to the realization of an air dropped weather station that is battery powered. There are several problems with this apparent solution.

Normally when a weather station is positioned, the meteorologist uses their understanding of geography and meteorology to optimize the location for weather measurements. This is not the case when they are air dropped. Even a suitable or marginal location for the weather station can not be assured. Secondly the battery will not last forever as was the case with AWOS. Finally, but related to the battery limitation, the more sensors that are used, the more data will be transmitted. As a sensor network grows in size it very quickly overextends its ability to transmit all the data it is sensing.

2.2 Research and Development

The CNAS solution to the technical problems is to add to the Navy and Army solution, not come up with something entirely new. To complement their development work, AFRL's research is embedding power aware computers into the tactical weather station to perform onboard processing to accomplish network and power control, quality assessment of the observation data, and cognitive information processing of the observation data into regionally significant information.

2.2.1 Technology Integration

To accomplish the research there is a necessity to approach from a systems viewpoint and assure that the integration of the various components is correctly done. Hardware integration is needed for the diverse components that make up the sensor pod; sensor, communication antenna, embedded computer, battery, and packaging. Software integration is needed for the sensor data acquisition, ad hoc networking, and cognitive information processing. This includes the operating system and the higher order languages needed to support the information processing. Of course there needs to be compatibility between the software and the hardware.

2.2.2 Cognitive Computing

By making use of the computing resources available on the nodes, low level data will be integrated and fused into a smaller amount of information, thus greatly reducing the amount of communications. This will be accomplished by having several expert systems analyze at different levels of abstraction and integration. Intra-nodal processing will include the integration and fusing of complex data from the node's heterogeneous sensors. Inter-node processing will integrate and fuse nodal information from geographically separated sensors. Finally by integrating High Performance Computing and sensor networks, phenomena based models will be used as complex events that trigger the reporting of sensor observations.

2.2.3 Global Patriot Demonstration

In order to assure the operation of the system we participated at two National Guard Patriot exercises. This gave a goal for each instantiation of the system, allowed for debugging in an operational environment and gave feedback from uniformed troops.

3 Application Technologies

This project involves the development of computationally rich sensor nodes, blackboard architecture and advanced computing concepts integrated in a network fashion to illustrate an elementary cognitive computing application. In order to accomplish our research goals, we have chosen a selection of off the shelf, state of the art, and research results for the individual parts of the new system. By leveraging existing components, we are able to concentrate on the information oriented research.

3.1 Computationally Rich Sensor Node (Figure 1)

3.1.1 Weather Sensors

We selected a off the shelf weather station for use in the CNAS node. The Climatronic's TACMET weather station (Figure 2) is extensively used by the Army in tactical situations. This unit sends its observation every second through a serial port.

3.1.2 PASTA Nodes

The Power Aware Sensing Tracking and Analysis (PASTA) computer was developed by DARPA in the Power Aware Computing and Communication program. (Figure 3) This computer was designed from the ground up to be an embedded sensor computer and perform on board data processing. Originally envisioned as a low level signal processor, the PASTA computer did have the ability to run higher level languages because of its Linux operating system. Conversations with the PASTA developers, Integrated Sciences Institute, indicated that indeed the possibility of running LISP, the Artificial Intelligence language, was possible. The PASTA computer also has a very good suite of IO ports and contained both serial and USB.

3.1.3 Communication /OLSR

We chose the use of WiFi or IEEE 802.11 as the wireless connection for the nodes. The WiFi antennas we used were connected to the PASTA computer through the USB interface. Built on this were the TCP/IP protocol and the Optimized Link State Routing (OLSR) program. OLSR allowed for messages to be hopped between nodes that were not within range of one another, when other nodes were in between.

3.2 Blackboard Architecture (Figures 4 and 5)

The blackboard framework used in the CNAS nodes is based on GBBopen, an open source generic blackboard builder.⁴ GBBopen was customized for use in this unique environment and is itself a LISP program. The blackboard framework is resident on each of the CNAS nodes to accomplish the temporal and spatial data abstraction that transforms the data into decision quality information within the network. The blackboard framework comprises a blackboard repository, knowledge sources, functions, object classes and triggering events. The blackboard framework is designed to 1) automatically assemble the network from the CNAS nodes and program the network to uniformly activate the wireless communication and sleep to conserve the battery power and 2) accomplish the in-network processing of the sensor data. Each part of the CNAS' Blackboard System is reviewed and explained in the following sections.

3.2.1 Control Shell

The agenda shell controls which knowledge source is placed in the activation. To accomplish this, each knowledge source is triggered by one or more events. All the knowledge sources, triggering events are shown in table 1. If two are activated at the same time, then the rating number (1-100) determines the queue arrangement. CNAS Blackboard system also uses GBB predefined event, start-control-shell.

The start-control-shell command is the event that triggers the determine-network-structure and determine-clusters knowledge sources.

3.2.2 Repository/Objects

The blackboard repository is a major feature of this instantiation. The blackboard repository defines the way that the data is stored and retrieved for use. The blackboard repository structure, classes and events are defined in the file CNAS.lisp.

3.2.3 Knowledge Sources and Functions

There are eight Knowledge sources and three functions used in the CNAS Blackboard. Functions are continually running and Knowledge sources event driven. Figure 6 has a listing of the Knowledge sources and their activation events.

sensor-reading

The objective of this function is to read the data stream coming from both the Tacmet's sensors through the serial port and reformats it into an event that encapsulates the sensor data. This even is written on the blackboard repository but is overwritten each time a new event is created, so it is not persistent. Both sensor and cluster head nodes use this function, since both can have a Tacmet weather station. The Tacmet weather station sends data very second. Air Force Weather Agency has a minimum reporting time of five minutes. The sensor-reading function does a five minute time average for each of the measured weather parameters. This results in the writing of a tacmet-five-minute-event.

time-event-generator

The objective of this function is to read the internal clock and creates time events on the blackboard repository. Time events include the wake-up-event, sleep-event, the send-tacmet-observation-event and the send-cluster-observation-event. During the daytime (between 6:00am and 6:00pm), the wake-up-event occurs five minutes before each hour and the sleep-event occurs at the hour. This event is used by the communication-wakeup and the communication-sleep knowledge sources. Four minutes into the wake-up-event either the send-tacmet-observation-event or the send-cluster-observation-event are generated based on whether a sensor or cluster head node is using this function.

socket-listener

The socket-listener periodically examines the WiFi socket to see if any messages have come in for the node. Two types of messages are used, those for network control and those for transferring node and cluster observations. Each of these is cast into the appropriate event on the blackboard repository.

node-observation-writer Knowledge Source

This knowledge source is activated by the tacmet-five-minute-event. It takes the five minute time averaged sensor data and creates a persistent object on the blackboard repository. This results in the accumulation of sensor data objects that can then be used for processing.

determine-clusters Knowledge Source

The objective of this knowledge source is to evaluate the location of each node, use the user defined cluster extent to determine which cluster the node is in and then have each node determine if it is a sensor node or a cluster head node. Every node has knowledge of all the other nodes in the cluster so once the cluster head node is elected the sensor nodes will send their observations to it every hour. Both the sensor and cluster head nodes use this knowledge source. This knowledge source creates or updates node objects. This knowledge source is triggered by the start-control-shell event.

communication-wakeup and communication-sleep Knowledge Sources

These two complementary knowledge sources either “wakes” or “sleeps” the WiFi antenna. They are activated by wake-up-event and sleep-event. The WiFi antenna is turned on for five minutes of every hour during the daytime, and not at all during the night.

node-observations-transmitter Knowledge Source

The objective of this knowledge source is to read the node-observation and reformat its data and send a data stream to the cluster head node. The sensor nodes use this knowledge source. This knowledge source does not create or update any blackboard objects. This knowledge source is triggered by the send-node-observation-event.

cluster-observations-transmitter Knowledge Source

The objective of this knowledge source is to read the cluster-observation and reformat its data and send a data stream to the regional node. The cluster head node uses this knowledge source. This knowledge source does not create or update any objects. This knowledge source is triggered by the send-cluster-observation-event.

cluster-abstraction Knowledge Source

This knowledge source is used on the cluster head nodes to transform the node level observations into environmental information about the region that the cluster of nodes has been emplaced in. This knowledge source creates or updates cluster observation objects. This knowledge source is triggered by the wake-up-event, as are the three communications knowledge sources. The communication knowledge sources move data from the spatially distributed CNAS sensor nodes to the cluster head node. This results in the node-observation objects being instantiated or updated on the cluster head’s blackboard. After it is activated by the wake-up-event, the cluster-abstraction knowledge source waits for four minutes for all the node-observation objects to be updated and then reads all the node-observation objects that have a time that is greater than the wake-up-event time. It then calculates the average, minimum and maximum value of each of the sensor parameter (observations) and instantiates or updates the cluster-observation objects on the blackboard. The sensor parameters include temperature, pressure, u-wind-vector, and v-wind-vector. The humidity observations are converted to dew point, again with the average, minimum and the maximum values.

node-observation-receiver Knowledge Source

The objective of this knowledge source is to obtain data from the WiFi communication and reformat it into a node-observation object on the blackboard. The cluster head node uses this knowledge source. This knowledge source creates or updates node-observation objects. This knowledge source is triggered by the node-observation-message-event.

node-announcement-receiver Knowledge Source

The objective of this knowledge source is to obtain data from the WiFi communication and reformat it into a node-observation object on the blackboard. The sensor and cluster head nodes use this knowledge source. This knowledge source creates or updates node objects. This knowledge source is triggered by the node-announcement-message-event.

3.3 Cognitive Computing Concepts

The term *cognition* is derived from the Latin word *cognoscere*, which fundamentally means *to know*, *knowledge* and *knowing*. Other loose definitions of cognition can include related concepts such as awareness, reasoning, thinking and understanding. But these concepts are at least once removed from the root concept, which is “to know.” Everyone agrees that cognition is a very complex process, and research to model the cognitive process using computers (sometimes termed “cognitive computing”) can involve many disciplines and encompass a wide variety of approaches, algorithms and architectures. For instance, much work in the field of artificial intelligence has involved advanced computer science concepts targeted at reasoning methods and knowledge representation. The neural network area has drawn heavily on mathematical methods to develop learning algorithms and generalization techniques. Cognitive science draws from an even more diverse mix of disciplines, including so-called “soft” sciences such as psychology and sociology. Much debate exists as to which approach is appropriate and whether or not a resulting system is cognitive.

Our research approach centers on a few key points regarding the cognitive process. The first key point comes from the root of the word cognition, which basically means “to know.” Our focus here is on the ability to transform data, which is low level, into information, and then into knowledge, which is higher level. Another key point focuses on the human ability to abstract, which means to capture the essential qualities or aspects of something (i.e., based on signals present). In this context, abstraction involves a hierarchy ranging from physical and formal to mental and conceptual (most abstract). Human languages are good examples of representations used to describe all levels of abstraction. These two aspects of cognition, the ability “to know” and to abstract, drive our research effort. We also realize that the cognitive process involves a many-to-one reduction of signal activity, as exemplified by the neuron. Both the abstraction process and the data–information–knowledge transition involve this many-to-one reduction of signals, and our network architecture naturally capitalizes on this.

3.3.1 Information Processing

CNAS is a hierarchical system consisting of sensor nodes, cluster head nodes and regional node(s). Each node is designed to communicate with other nodes, if possible, and take part in cluster management and network operations. At the lowest level of the hierarchy, sensor nodes

gather atmospheric data and produce weather observations. A cluster is a grouping of sensor nodes located in a geographic region of interest. Each sensor node within a cluster provides weather observations and some pre-processed information to the cluster head node. The cluster head collects information from all sensor nodes reporting in its group and reduces the multiple reports into an averaged observation. At present, the cluster head is determined simply as the operational node closest to the center of the geographic region defined by the cluster. The regional node performs functions similar to a cluster head node (if multiple clusters are present) except at a higher level. The regional node also serves as the primary interface between the console or user and the rest of the network, handling system messages and commands. As for the kinds of functions and processing techniques proposed toward the goal of cognitive-like processing, this includes averaging of parameters, expected ranges, thresholds, filtering, generalizations and the establishment of relationships or links among sensed parameters and other constructs formed at higher levels within the system. An overarching goal of this processing will be to capture, represent and process signals in time and across various levels of abstraction, thereby helping to recognize their underlying order and organization.

4 Future Plans

The last year of the project will continue development of both the hardware and the software. The hardware advance will incorporate a new sensor board that will be internal to the CNAS pod. This new sensor will include a global positioning system (GPS) in the sensor suite, in addition to the temperature, humidity, and pressure environmental sensors. The GPS antenna will enable each of the sensor nodes to determine its location and from that determine which cluster it is in and whether it should assume the cluster head role.

The 2006 system included the use of the cognitive blackboard framework with minimal onboard processing. This next stage of development builds on that capability and will address two problems associated with the sensor network, assessing data quality and converting the large amount of data into actionable information. Assessing the quality of data will be accomplished by using the multiple sensor nodes to collectively assess each other's data to eliminate wrong or unusable data. The transformation of the data into information will be accomplished by coding cognitive algorithms, such as pattern matching, into the information processing.

Next year's technology demonstration will be at the joint US/Australia exercise. We will be incorporating the hardware upgrade but not the software advances into this demonstration.




Figure 1. Computationally rich sensor node with portable weather station



Figure 2. TACMET II weather station




Figure 3. PASTA computer

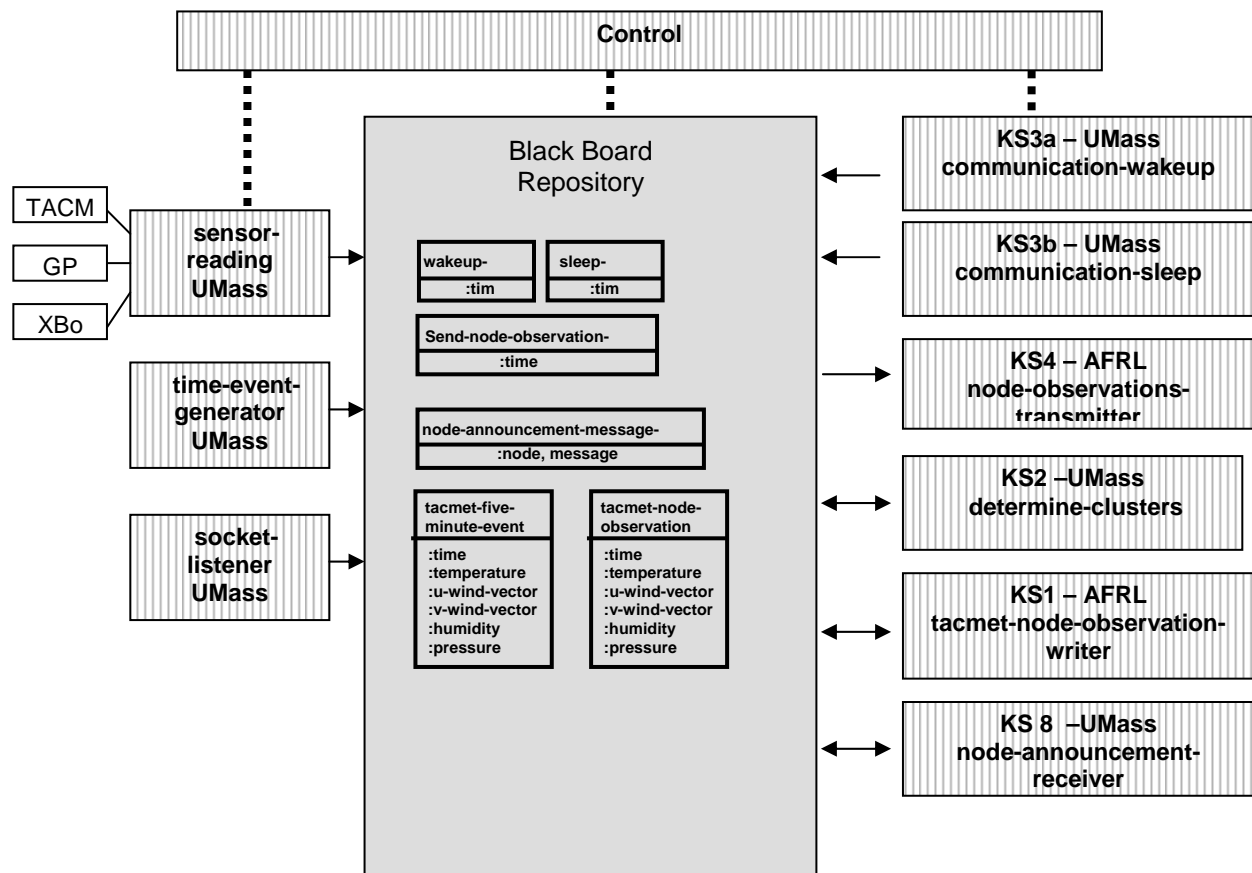


Figure 4. Sensor Node Blackboard

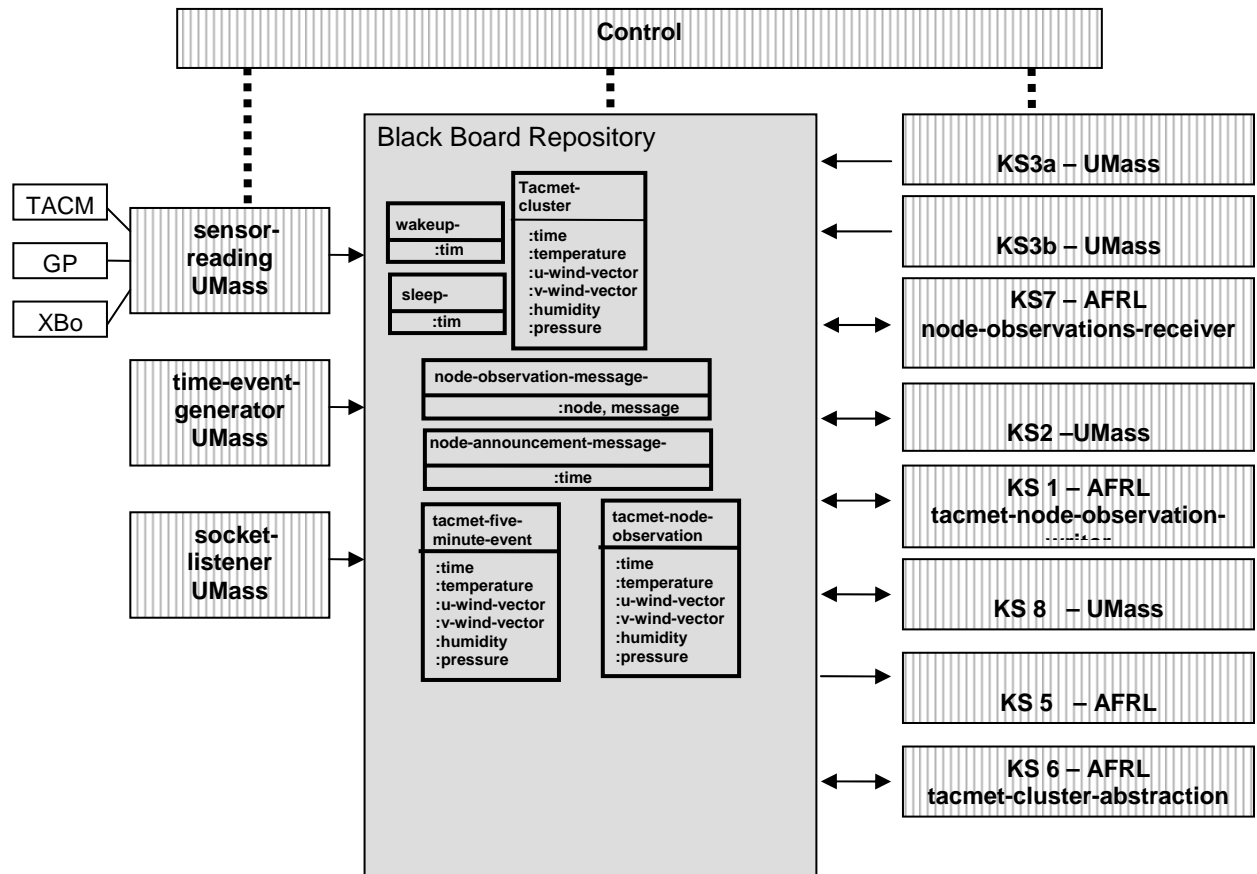


Figure 5. Cluster Head Node Blackboard

	Knowledge Source	trigger event
1	tacmet-node-observation-writer	tacmet-five-minute-
2	determine-clusters	start-control-
3a	communication-wakeup	wakeup-event
3b	communication-sleep	sleep-event
4	node-observations-transmitter	send-node-observation-
5	cluster-observations-transmitter	send-cluster-observation-
6	tacmet-cluster-abstraction	send-cluster-observation-
7	node-observation-receiver	node-observation-message-
8	node-announcement-receiver	node-announcement-message-

Figure 6. Knowledge Sources and Triggers

5. References

¹ Mesoscale Model 5, <http://www.mmm.ucar.edu/mm5/>

² Clark, D.A., Matthews, M.P., An Integrated Sensor Testbed for Support of Theater Operations in Areas of Complex Terrain, Battlespace Atmospheric and Cloud Impacts on Military Operations, 25-27 April 2000, Fort Collins, CO.

³ Aircraft Owners and Pilots Association, <http://www.aopa.org/asf/wx/mountain.html#high>

⁴ Corkill, D., GBBopen Reference Manual, Version 0.9, November 6, 2005