Confidence-based Data Management for Personal Area Sensor Networks *

Nesime Tatbul[§] Mark Buller[†] Reed Hoyt[†] Steve Mullen[†] Stan Zdonik[§]

[§] Brown University
{*tatbul, sbz*}*@cs.brown.edu*

[†] U.S. Army Research Institute of Environmental Medicine
{*mark.j.buller, reed.hoyt, stephen.mullen*}*@us.army.mil*

Abstract

The military is working on embedding sensors in a "smart uniform" that will monitor key biological parameters to determine the physiological status of a soldier. The soldier's status can only be determined accurately by combining the readings from several sensors using sophisticated physiological models. Unfortunately, the physical environment and the low-bandwidth, push-based personal-area network (PAN) introduce uncertainty in the inputs to the models. Thus the model must produce a confidence level as well as a physiological status value. This paper explores how confidence levels can be used to influence data management decisions. In particular, we look at power-efficient ways to keep the confidence above a given threshold. We also contrast push-based broadcast schedules with other schedules that are made possible by two-way communication.

1 Introduction

Data management has traditionally been reserved for large complex software environments in which huge amounts of data must be processed with limited resources. Modern database management systems (DBMS) that run on large back-office servers are the most well-known embodiment of this kind of technology. Researchers have recently realized that similar technologies are needed in smaller environments in which resource limitations are also an issue [9, 4]. In sensor-based applications, bandwidth and battery power are typically the scarce resources.

This paper looks at a real sensor-based application in which results are computed along with a confidence value. The data management game that we play here is to set transmission parameters (statically or dynamically) in order to achieve the highest confidence only when the application requires it. Our techniques use strategies that are informed by the confidence models to conserve bandwidth and power. We discuss these ideas and some possible approaches in terms of a military physiologic sensing application. Our main contribution is in the way that confidences can be used in this particular application.

The warfighter's workplace has unique occupational challenges: from mission demands, the environment, and combat injuries. Modern dismounted soldiers commonly engage in intense, mentally and physically demanding 3-10 day missions, often in rugged terrain or complex urban settings. Warriors carry heavy loads and are often food and sleep-restricted. Environmental conditions can vary widely in terms of ambient temperature, humidity, wind speed, barometric pressure, and the like. In non-war mode the military can suffer over 120 heat casualties a year [1]. Under or over hydration can decrement physical and cognitive performance, and increase the risk of heat injury, hyponatremia, or death [14, 15, 16]. Added to the harsh environment is the possibility of receiving a wound. Once a warfighter has become a casualty, it is critical that treatment is received quickly during the "golden hour", which is the short period of time when proper medical treatment can mean the difference between life and death. It has been suggested that 20% of these deaths could be prevented with rapid intervention [13]. Therefore, wearable physiological and medical status monitoring can play an important role in: sustaining physical and mental performance, reducing

* This work has been supported in part by the Military Operational Medicine Research Program, the Combat Casualty Care Research Program, and the Telemedicine and Advanced Technologies Research Center, US Army Medical Research and Material Command, Ft. Detrick, MD, 21704-5014; and by the NSF, under the grants IIS-0086057 and IIS-0325838. The opinions or assertions contained herein are private views of the authors and are not be construed as official or as reflecting the views of the US Army or Department of Defense. Citations of commercial organizations and trade names in this report do not constitute an official Department of the Army endorsement or approval of the products or services of these organizations.

Proceedings of the First Workshop on Data Management for Sensor Networks (DMSN 2004),
Toronto, Canada, August 30th, 2004.
`http://db.cs.pitt.edu/dmsn04/`

the likelihood of non-battle injuries such as heat stroke, and provide remote notification and medical status of a casualty.

In this paper, we first describe the Warfighter Physiologic Status Monitoring (WPSM) application in detail in Section 2, where we also present an example scenario and show how the sensor network behaves under this scenario. We discuss potential data management techniques that would improve the existing network in Section 3. We present some preliminary simulation results in support of these discussions. Section 4 summarizes related work in the area. We conclude the paper by discussing future directions in Section 5.

2 The WPSM Application

2.1 The Sensor System

The Medical Research and Material Command (MRMC) under its Warfighter Physiologic Status Monitoring - Initial Capability (WPSM-IC) program is developing what is essentially a wellness monitor for each soldier. This system is comprised of a medical hub which hosts a personal area network of physiologic and medical sensors and a number of algorithms. The algorithms estimate the state of the warfighter in the following areas: *Thermal*, *Hydration*, *Cognitive*, *Life Signs*, and *Wound Detection*. Each area has four potential states that are coded by color. *Green* represents normal-no action is required; *Yellow* means requires attention; *Red* calls for immediate action; and *Blue* indicates a system fault. For each area's state, the hub also estimates a confidence level. Confidence refers to the accuracy level of the state estimated by a model.

The states for each medical and physiologic area are based upon input to the state algorithms from a number of sensors distributed around a warfighter's body, uniform and equipment, as well as outputs from other algorithms resident in the medical hub. Figure 1 shows a schematic of the current WPSM-IC sensor system and the physical placement of sensor equipment on a warfighter. The ingestible thermometer pill is network-enabled, and measures the temperature of the stomach and intestines, which is usually a good indication of body core temperature. The fluid intake monitor measures the amount of fluid consumed through a bladder-style canteen. The life sign detection sensor (LSDS) is an integrated system with multiple parameters and algorithms including heart rate, respiration rate, body orientation, actigraphy [1], and skin temperature. The LSDS also has an integrated ballistic impact detection device which provides an alert when on-body acoustic signals are detected that indicate the probability that a ballistic projectile has impacted the warfighter. The sleep performance watch treats sleep as a consumable quantity, measures it, and uses an algorithm to equate this to apparent cognitive readiness. The soldier also carries a GPS and other technologies which report his geographic location.

[1] Actigraphy is a measure of activity patterns [7].

Figure 1: WPSM-IC Sensor System

The sensors are connected to the medical hub by a proprietary wireless RF network [12]. The network was developed with a number of key requirements unique to a military operational environment. The network needed to be very low power, to allow miniature physiologic sensors to run for weeks without the need of battery recharge or replacement, and also have an ability to reject cross talk and interference from similar networks borne by other soldiers when congregated in close proximity to each other. In addition, the network had to provide a low profile signature to avoid detection.

The current network uses a detuned (low detectability) 40MHz radio frequency (RF) carrier. Digital data are transmitted from sensors to the medical hub utilizing a pseudo random push transmission scheme. Sensors are factory set with an identification number (ID) and random number table seed. Sensors are supplied operating in a deep sleep mode and are activated through an infrared (IR) port, by a medical hub. Activation associates a particular sensor with a particular hub. The sensor in a series of initial transmissions sends its transmission schedule (based upon its ID and random number seed) and clock information to the hub. Knowing this information, the hub is able to keep itself in a sleep mode, powering up fully only when it knows to expect a transmission from an associated sensor. This reduces power consumption in the hub ($\sim$0.1% duty cycle) and also guards, to some degree, against cross talk from other sensors. This "push-only" scheme has the benefits of allowing sensors to only carry transmission circuitry which is activated on a known schedule, rather than both a transmitter and receiver. In a "polled" scheme, a sensor would need to constantly power the receiver circuitry to listen for data polls, and hence consume more power. Sensors in the current network sample every 15 seconds and transmit data at 2400 baud on average every 15 seconds. The transmission interval can vary from 3 seconds to 27 seconds according to the pseudo random schedule with each transmission time interval having an equal probability of occurrence. Each sensor message is 240 bits long.

Model	Skin Temp.	Heart Rate	Actigraphy	Geo-Location	Resp. Rate	Pill	# Sensors
TSkin	$\surd$						1
Threshold	$\surd$	$\surd$					2
Model1			$\surd$				1
Model2			$\surd$	$\surd$			2
Model3		$\surd$	$\surd$	$\surd$	$\surd$		4
TCore						$\surd$	1

Table 1: Models for estimating thermal state

2.2 Example Scenario: Estimating the Thermal State

In this paper, we focus more closely on the warfighter thermal state, and the sensors and models which allow thermal state and its confidence to be determined. In what follows, we describe an example scenario for estimating the thermal state of a soldier.

The best and most confident method to assess thermal state is direct measurement of core body temperature by using the network-enabled ingestible pill. When core body temperature is greater than $39.5°$C, there is a high probability that the warfighter is in thermal strain. However, this method is impractical for continual use. Thus, these devices are reserved for use during high thermal stress missions, while encapsulation in nuclear, biological, and chemical protective suits, and/or if use is indicated by other algorithms or medics.

When a core temperature pill is not being used, WPSM-IC plans to use variants of two basic types of models to provide an estimate of thermal state. The simplest model is the Threshold Model [2] that takes inputs from two sensors measuring skin temperature and heart rate. Under very low and high skin temperatures, the confidence in states produced by this model is higher than otherwise. For mid-values of temperature, knowing heart rate values improves confidence. The second model is a first principles model similar to the USARIEM Scenario Model [6], that takes metabolic rate, environmental conditions, clothing configurations and biometric data as inputs to estimate core body temperature. Metabolic rate and the environmental conditions are key drivers of this model. From the current system, metabolic rate can be derived independently from heart rate, respiration rate, actigraphy, and geo-location readings in multiple ways with different confidence levels. Based on these, Table 1 summarizes six alternative models to estimate thermal state together with the sensors they are using. TSkin Model is a simplified version of the Threshold Model, using only the skin temperature sensor. The Threshold Model additionally uses the heart rate sensor. Models 1-3 represent variants of the first principles model where metabolic rate is derived using different sets of sensors: Model1 uses just actigraphy; Model2 uses both actigraphy and geo-location; Model3 uses actigraphy, geo-location, heart rate and respiration rate. Finally, TCore Model uses the core temperature pill. Each alternative model has complex algorithms that map sensor values to physiologic states with certain confidence levels. The details of these algorithms are outside the scope of this paper.

Our thermal state estimation problem consists of three major dimensions that determine the confidence levels:

1. **Model:** The first factor is the model, and hence the set of sensors, that participate in the state computation. Input from a greater number of sensors usually increases the confidence in the state. This is not true when the core temperature pill is used. However, the core temperature pill is unique in that it is a consumable sensor, with a costly logistics and resupply train.

2. **Latency:** The second factor is the latency of sensor messages. As readings get older, their relevance and usefulness to the models and state algorithms decay. Thus, a latency decay function or "shelf-life" is defined for each sensor. This function maps latency values measured in seconds to decay coefficients. For our example scenario, all sensors are simply assumed to have the following exponential decay function[2]:

$$2^{-(\lceil latency/15 \rceil - 1)}, \text{ where } latency > 0$$

   For example, a heart rate reading of age 20 seconds has a decay coefficient of 0.5, i.e., a state computation that uses this heart rate value would have its confidence level degraded by 0.5. When multiple sensors are involved in a model computation, we simply use their average latency to compute the decay coefficient. If sensors had different latency decay functions, then we would take an average of their individual decay coefficients.

3. **State:** Finally, the third determinant of confidence is the output state. For our thermal state estimation problem, the Green state can be determined with higher certainty than the Yellow and Red states.

Next, we present confidence assignments on two of the dimensions, Model and State. The latency dimension is based on the decay function provided above.

As mentioned earlier, physiologic models are also affected by the physical environment. In Table 2, we illustrate a detailed work environment scenario. The first two columns of this table show nine different environment-activity combinations. Work environment conditions are

[2] In general, it is more realistic to choose different decay functions for different sensors. For example, heart rate readings would certainly age faster than ambient temperature readings.

		TSkin			Threshold			Model1			Model2			Model3			TCore
Env.	Work	G	Y	R	G	Y	R	G	Y	R	G	Y	R	G	Y	R	G/Y/R
cool	low	80	76	72	90	85.5	81	95	90.25	85.5	95	90.25	85.5	95	90.25	85.5	100
warm	low	80	76	72	90	85.5	81	95	90.25	85.5	95	90.25	85.5	95	90.25	85.5	100
hot	low	60	57	54	80	76	72	95	90.25	85.5	95	90.25	85.5	95	90.25	85.5	100
cool	med	70	66.5	63	90	85.5	81	95	90.25	85.5	95	90.25	85.5	95	90.25	85.5	100
warm	med	50	47.5	45	70	66.5	63	80	76	72	90	85.5	81	95	90.25	85.5	100
hot	med	40	38	36	60	57	54	70	66.5	63	80	76	72	90	85.5	81	100
cool	high	40	38	36	60	57	54	90	85.5	81	95	90.25	85.5	95	90.25	85.5	100
warm	high	20	19	18	40	38	36	60	57	54	70	66.5	63	80	76	72	100
hot	high	5	4.75	4.5	20	19	18	50	47.5	45	60	57	54	75	71.25	67.5	100

Table 2: Work Environment models to estimate thermal state and their confidence levels

measured independently from the soldier (e.g. through a weather station) and they are external to the soldier's personal area sensor network. However, they directly affect the confidence achieved by the models. For each environment-activity combination, confidence levels for six alternative models are shown. Note that these values are representative values. Each model can estimate the Green (G) state with the highest confidence. If a Yellow (Y) is computed, this confidence degrades by 0.95; if a Red (R) is computed, it degrades by 0.90. TCore Model is an exception as its confidence for all states is perfect due to its being a direct measure of thermal state. Note that as the environment moves from cool to hot, and as activity moves from low to high, more types of sensors may be needed to maintain a high confidence about the soldier's thermal state.

The application requires different confidence levels depending on soldier's state. Table 3 shows the required thresholds for our example. If a Green state is reported, its confidence has to be at least 50. If a Yellow state is observed, a confidence value of at least 70 is required. Finally, if soldier's state is reported to be Red, a confidence value of at least 80 has to be provided. In other words, the application requires higher confidence for more important events. The goal is to operate the sensor network in such a way that it delivers state estimations with sufficient confidence levels.

State	Confidence Threshold
Green	≥ 50
Yellow	≥ 70
Red	≥ 80

Table 3: Required confidence thresholds for each state

2.3 The Push-Only Transmission Scheme

We simulated the existing push-only sensor network on CSIM [11]. We ran the alternative models of the example scenario through the simulator, using one model and one environment-work pair at a time. We assumed that the soldier is in the Green state. We make the following important observation: Models requiring more sensors do not always achieve better confidence levels. Models periodically compute states based on what has most recently been received from the participating sensors. When more sensors are present in the network, the frequency of packet collisions and message drops increases. When the most recent measurement from a sensor is missing, state computation at the hub has to rely on a stale earlier reading from that sensor. As mentioned before, stale data degrades the confidence level associated with each instance of model output. Figure 2 shows how each model behaves under three of the environment-activity conditions. Model3, using four sensors to estimate metabolic rate, achieves better confidence than other models (except TCore) in the (warm, high) and the (hot, high) cases which represent relatively high intensity conditions. To generalize this notion, delivering high-confidence for different detection goals (i.e., thermal stress, wound detection, etc) demand different models.

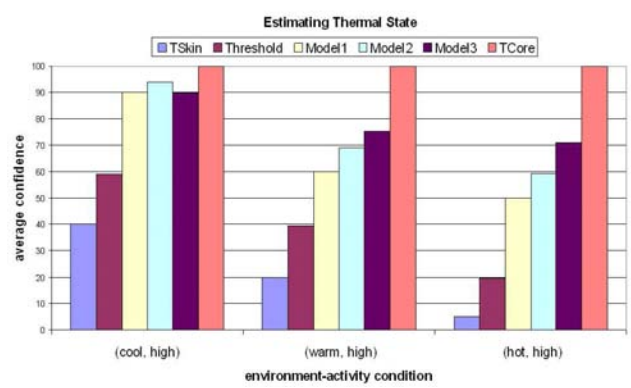

Figure 2: Simulation of the push-only scheme

3 Confidence-based Data Management

In the WPSM context, data management largely concerns the scheduling of data transmissions. Frequent transmission can in principle improve latency, but over zealous transmission can waste power and increase the odds of a collision (i.e., lost data). In what follows, we discuss techniques for optimizing this tradeoff. We use confidence modeling as the primary way to inform these decisions.

Model	Average Confidence	% Drop
Model 1	64.92	0
Model 2	72.73	1.98
Model 3	77.75	5.79
All Models	79.22	5.71

Table 4: Model redundancy simulation results

3.1 Exploiting Redundancy

Physiologic states can be estimated with higher certainty by allowing redundancy at several levels.

Model Redundancy. All alternative models to estimate a particular state can run concurrently. As we have demonstrated, various factors like sensor values and latency decay may cause one model to achieve higher confidence than another. By running the models simultaneously, one can obtain multiple state estimations at different levels of certainty and the one with the highest confidence can be picked. Table 4 shows preliminary results from a model redundancy simulation for a changing work environment scenario. We again assume that the soldier is in the Green state. Models 1-3 are run both separately and all together. The work environment is initially set to (cool, high) and then gradually changed to (warm, high) and (hot, high). When all models are redundantly run together, the average confidence is the highest. Models 1 and 2 have fewer drops due to fewer sensors sharing the channels. *Model 3* and *All Models* use four sensors and they both experience higher percent message drop due to collisions. Note that *All Models* loses around the same percent of messages as *Model 3* alone, but achieves higher average confidence.

Data Redundancy. Sensor readings can be transmitted multiple times. A sensor message not only contains the most recent reading, but also the previous reading as well. This type of redundancy is useful when the model to be computed not only requires the most recent sensor value, but a valid sensor reading every certain time period. This increases the probability that a reading will get through.

Obviously, allowing redundancy has drawbacks in terms of resource consumption. Running all models at the same time increases network traffic and message loss. Similarly, repeating readings in multiple messages increases message lengths, thus consuming bandwidth and expending additional battery power. Therefore, the degree of redundancy has to be adjusted based on a tradeoff between desired level of confidence, variability of the conditions affecting confidence, and resource consumption.

3.2 Adjusting Sampling Rates

In the current deployable network, sensors come with factory-set transmission schemes. Thus, their sampling and transmission periods are not adjustable. However, we believe that confidence levels and network lifetime could be considerably improved by dynamically adjusting these sensor parameters to match the requirements of the physiological models. Thus, we foresee a need to incorporate two-way communication into future sensor designs. Of course, we must be able to show that the extra power needed to run the receiver is worth it.

In general, sensors reporting with high frequency feed low-latency values into the models, but messages are more likely to get dropped due to collisions. In the extreme case, high data rates can translate into high latency as well as extensive energy consumption. On the other hand, low-frequency transmissions seldom get dropped and use power economically, but they may not refresh the models as often as needed. Each sensor's sampling rate should be adjusted between these two extremes based on model requirements.

One thing to consider is the sharing between running models. There are five different areas where state estimation is needed. Each area may also run multiple models concurrently. Each model requires readings from a certain subset of the sensors. Sensors could be ranked based on how many models they are feeding. Also, importance of a state could be considered. For example, Wound Detection may be more important than Cognitive State. Sensors involved in Wound Detection should have higher rank. Sensors of high rank should have shorter sampling and reporting periods.

A second consideration is the latency decay functions of the sensors. A cumulative latency decay function could be defined based on functions of all sensors involved in a model computation. This function would indicate how often that model has to be refreshed to preserve its confidence level. As mentioned before, some sensors can have stricter latency requirements than others. For example, heart rate readings age faster than temperature readings. This implies that the heart rate sensor must update more often, illustrating the notion that refresh periods are application dependent.

3.3 Bi-directional Data Communication

The sensor network used in the described application is designed to be push-only, where data flows in a single direction, from sensors to the hub. Sensors do not have any receivers, but only transmitters. The rationale behind this kind of a setup is threefold. First, it uses less power since no sensor wastes battery by listening to the network. Second, message loss is small since collisions are expected to occur less frequently. Last but not least, push-only sensors are much cheaper to build. However, this design limits many potential optimizations that could be performed at the receiver hub.

The receiver hub is the only point in the network that has a complete view of all the sensors and all the physiological models with their confidence requirements. As such, it can make the best judgement about how to deliver high confidence states in an efficient way. However, in a push-only scheme, it has no control over sensor transmissions. The hub must be able to "pull" from the sensors as needed.

With a two-way communication model, we can accumulate minimal sensor readings in order to populate the lower-confidence models. Typically, the amount of data and the latency requirements are lower for low confidence results. In this situation, if we get an alert for a thermal stress event with a low confidence, we can then contact the sensors to collect more data in order to feed the higher confidence models. Thus, we only spend bandwidth and power when it is needed. In other words, in the normal operating case, it is best to run lean at the expense of confidence. When an important but low confidence event is observed, we expend more resources to confirm or deny it. We now illustrate this point on our work environment example presented in Section 2.2. As shown in Table 3, our application has different confidence requirements depending on the soldier's state. These requirements can be met in multiple ways using alternative models. For example, if the soldier is in the Green state and under the (hot, high) condition, Models 1-3 and TCore Model can deliver enough confidence ($\geq$ 50). Among these models, Model1 is the most desirable one. First of all, Model1 uses only one sensor. Thus, network bandwidth does not have to be shared with other sensors. The network lifetime with one sensor would be much longer as the energy consumption at the hub is proportional to the the number of sensors it is communicating with. Finally, the actigraphy sensor used by Model1 is a much cheaper alternative than using the core temperature pill. If we apply this heuristic of "using as few sensors as possible" to all condition and state combinations in Table 2, we end up with model preferences shown in Table 5.

To show the performance benefit of this heuristic, we considered a scenario where the soldier is in a (hot, medium) environment and is initially in the Green state. Then his state gradually changes to Yellow and Red. Model3 delivers enough confidence for all of these states. Therefore, we ran one simulation where only Model3 is used. In a second simulation, we started out with Model1 and changed to Model2 only when soldier entered Yellow state, when Model1 can not deliver enough confidence. Similarly, when the soldier's state changed to Red, we switched from Model2 to Model3 so that confidence is above the required threshold. This second run simulates the behavior of a hub pulling from sensors as necessary. Initially, it only pulls from the actigraphy sensor; then the geo-Location sensor is added; and finally, heart rates and respiration rates are pulled. We further assumed that the main determinant of network lifetime is the battery at the hub which is about 1800 mAHrs. Additionally, we assume that each sensor that is turned on has a current draw of 50mA; i.e, if this sensor is left on for an hour, it will consume 50mAHrs of the total 1800mAHrs battery. Then we compared these two simulations in terms of network lifetime. The first simulation runs out of hub battery in 9 hours whereas the second one can survive more than 14 hours. This simple scenario clearly demonstrates how a pull-based model could conserve energy based on model and situation-specific confidence requirements.

In a way, bi-directional communication enables switching between alternative estimation models dynamically. As such, it is a much more efficient alternative to the redundancy approach proposed in Section 3.1.

Env.	Work	G	Y	R
cool	low	TSkin/Model1	TSkin/Model1	Model1
warm	low	TSkin/Model1	TSkin/Model1	Model1
hot	low	TSkin/Model1	Model1	Model1
cool	med	TSkin/Model1	Model1	Model1
warm	med	TSkin/Model1	Model1	Model2
hot	**med**	**Model1**	**Model2**	**Model3**
cool	high	Model1	Model1	Model1
warm	high	Model1	Model3	TCore
hot	high	Model1	Model3	TCore

Table 5: Model preferences based on the number of sensors

Two-way communication is also more flexible than the sampling rate adjustment approach discussed in Section 3.2. Sensor transmission rates can effectively be adjusted by changing the pull frequency at the hub.

4 Related Work

There is a growing body of research on sensor network data management. TinyDB [18] and Cougar [4] are two example query processing systems for multi-hop sensor networks. These systems emphasize in-network processing of declarative queries to reduce data communications and battery usage. TinyDB especially focuses on acquisitional aspects of query processing like where, when and how often data should be collected from the sensors [9]. Sensor sampling rates are adjusted based on event and lifetime specifications of queries. Cougar uses sensor update and query occurrence probabilities for view selection and location on top of a carefully constructed aggregation tree [4]. Scheduling techniques to overcome collisions in the sensor network are also explored in this project. These systems are designed to serve monitoring applications that span a larger or difficult to reach geographical area than the personal area case, where multi-hop sensor communication is a necessity (e.g., habitat monitoring).

More relevant to our problem are quality-driven approaches. As an example, TiNA exploits temporal coherency tolerance specifications of users in in-network processing to trade off between result quality and energy conservation [17]. Sensor readings are reported only if they differ from an earlier value by a certain amount. Another example is the QUASAR project [8], which also exploits applications' tolerance to imprecision to minimize resource consumption. As a more closely related work to ours, a model-driven approach for data acquisition in sensor networks has been recently developed by Deshpande et al [5]. A probabilistic model of the sensor network data is created based on a history of readings from sensors and correlations between them. Queries can be approximately answered based on this model. If confidence requirements can not be met by the model alone, then the sensors in the network need to be queried. The model is also refined as more readings are received. In the application that we consider, multiple complex models exist to estimate physiological states of a soldier. Each model uses a different set of sensors. These models and their confidence levels are well-defined. Rather than building and refining the models, we concentrate on efficient data acquisition from sensors to estimate states with acceptable confidence using alternative models.

There is some related work on data management for personal area sensor networks as well. For example, a recent work proposes a query processing system for healthcare bio-sensor networks [3]. Patient heart rates are monitored using electrocardiogram (ECG) and accelerometer sensors. Multiple ECG sensors have to be worn for a complete measurement of the electrical activity of the patient's body. Furthermore, if the patient moves, ECG signals may be corrupted. Therefore, readings from an accelerometer sensor have to be correlated with ECG readings for a more reliable result. This application has similar sensor network uncertainty concerns as ours. However, the focus of this work is on query processing at the base station. We believe our confidence-based approach could be used at the data acquisition phases of this system to improve query results. In the same domain, CodeBlue is a wireless communications infrastructure for medical care applications [10]. It is based on publish/subscribe data delivery where sensors worn by patients publish streams of vital signs and geographic locations to which PDAs or PCs accessed by medical personnel can subscribe. Secure and ad hoc communication, prioritization of critical data, and effective allocation of emergency personnel in case of mass casualty events are major emphases of this project.

Finally, wireless sensor networks are also a subject of recent research in the networking community. Of particular relevance to our work are MAC (Medium Access Control) protocols that determine when and how the network nodes coordinate to share a broadcast channel [19]. Collision avoidance is a major concern in these protocols. S-MAC is one such protocol where sensor nodes periodically sleep to reduce energy consumption by avoiding idle listening [20]. While such protocols make the underlying network more reliable in power-efficient ways, they are unaware of the application-specific requirements, like confidence levels in our WPSM-IC application.

5 Future Directions

In this paper, we presented a challenging sensor network application which can highly benefit from various data management strategies as evidenced by our initial simulation results. We are currently working on making these strategies operational on the real network. In the future, we are planning to extend this work in several directions. A potential research direction involves treating sensor readings as continuous waveforms with integrity constraints. If sensor values could be noisy or erroneous, earlier values could verify or deny confidence of the latest value. We could also decide when to pull from sensors based on what values have recently been received. If recent heart rate readings suggest that the heart rate could not have gone beyond normal threshold since the last reading, then we do not need to receive a new heart rate report.

WPSM-IC is currently concerned with dismounted warriors and the management of their personal area networks. The goal at this point is to create a summary of the soldier's physiological state at the hub. In the future, this state information would be disseminated to other battlefield units. This might include mobile medics who are deployed in the theater of operation or to advanced field hospitals that are prepared to deal with both prevention of potential casualties as well as management of known casualties of various kinds. The information that is uploaded beyond the individual soldier would be used for some form of remote triage.

The remote triage problem, of course, comes with its own technical challenges. Similar reports from more than one co-located soldier might be an indication of a particular kind of attack. Physiological status reports from many soldiers can be used to prioritize treatment. In these cases, the medic might find that the reported confidence level is not high-enough to warrant the deployment of an ambulance. Instead, the medic may contact the soldier's hub to amplify the confidence to some given level. This might require a great expenditure of resource, but in an emergency, the investment is likely worth it.

References

[1] Army Medical Surveillance Activity. `http://amsa.army.mil/`.

[2] L. Berglund, M. Yokota, and M. Kolka. Non-Invasive Physiological Hyperthermia Warning System. Technical report, USARIEM, 2004 (in process).

[3] C.-M. Chen, H. Agrawal, M. Cochinwala, and D. Rosenbluth. Stream Query Processing for Healthcare Bio-sensor Applications. In *IEEE ICDE Conference*, April 2004.

[4] A. Demers, J. Gehrke, R. Rajaraman, N. Trigoni, and Y. Yao. The Cougar Project: A Work-In-Progress Report. *Sigmod Record*, 32(4), December 2003.

[5] A. Deshpande, C. Guestrin, S. Madden, J. Hellerstein, and W. Hong. Model-Driven Data Acquisition in Sensor Networks. In *VLDB Conference*, Toronto, Canada, September 2004.

[6] A. P. Gagge, A. P. Fobelets, and L. G. Berglund. A standard predictive index of human response to the thermal environment. *ASHRAE Transactions*, 92(2B):709–731, 1986.

[7] R. Hoyt, M. Buller, J. DeLaney, D. Stultz, K. Warren, M. Hamlet, D. Schantz, W. Matthew, W. Tharion, P. Smith, and B. Smith. Warfighter Physiologic Status Monitoring (WPSM): Energy Balance and Thermal Status During a 10-Day Cold Weather US Marine Corps Infantry Officer Course Field Exercise. Technical Report T-02/02, DTIC Number A396133, USARIEM, October 2001.

[8] I. Lazaridis, Q. Han, X. Yu, S. Mehrotra, N. Venkatasubramanian, D. V. Kalashnikov, and W. Yang. QUASAR: Quality-Aware Sensing Architecture. *Sigmod Record*, 33(1), March 2004.

[9] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. The Design of an Acquisitional Query Processor for Sensor Networks. In *ACM SIGMOD Conference*, San Diego, CA, June 2003.

[10] D. Malan, T. Fulford-Jones, M. Welsh, and S. Moulton. CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care. In *International Workshop on Wearable and Implantable Body Sensor Networks*, April 2004.

[11] Mesquite Software, Inc. CSIM18 Simulation Engine. `http://www.mesquite.com/`.

[12] Mini Mitter, Inc. Physiological and Behavioral Monitoring for Humans and Animals. `http://www.minimitter.com/`.

[13] R. F. Bellamy. The Causes of Death in Conventional Land Warfare: Implication for Combat Casualty Care Research. *Military Medicine*, 149:55–62, 1984.

[14] S. J. Montain and E. F. Coyle. Fluid ingestion during exercise increases skin blood flow independent of increases in blood volume. *Journal of Applied Physiology*, 73(3):903–910, 1992.

[15] S. J. Montain and E. F. Coyle. Influence of graded dehydration on hyperthermia and cardiovascular drift during exercise. *Journal of Applied Physiology*, 73(4):1340–1350, 1992.

[16] S. J. Montain and M. N. Sawka and W. A. Latzka and C. R. Valeri. Thermal and cardiovascular strain from hypohydration: Influence of exercise intensity. *International Journal of Sports Medicine*, 19(2):87–91, 1998.

[17] M. A. Sharaf, J. Beaver, A. Labrinidis, and P. K. Chrysanthis. TiNA: A Scheme for Temporal Coherency-Aware in-Network Aggregation. In *3rd ACM MobiDE Workshop*, September 2003.

[18] TinyDB: A Declarative Database for Sensor Networks. `http://telegraph.cs.berkeley.edu/tinydb/`.

[19] W. Ye and J. Heidemann. Medium Access Control in Wireless Sensor Networks. In C. S. Raghavendra, K. M. Sivalingam, and T. Znati, editors, *Wireless Sensor Networks*. Kluwer Academic Publishers, 2004.

[20] W. Ye, J. Heidemann, and D. Estrin. An Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *21st International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOMM 2002)*, New York, NY, September 2002.