

Prediction-based Monitoring in Sensor Networks: Taking Lessons from MPEG*

Samir Goel

Department of Computer Science
Rutgers, The State University Of New Jersey
Piscataway, NJ 08854
gsamir@cs.rutgers.edu

Tomasz Imielinski

Department of Computer Science
Rutgers, The State University Of New Jersey
Piscataway, NJ 08854
imielins@cs.rutgers.edu

ABSTRACT

In this paper we discuss the problem of monitoring data sensed in large sensor networks. A sensor typically runs on a battery having a limited lifetime. In order to increase the lifetime of a sensor it is important that the mechanisms used in monitoring them be energy-efficient. In this paper, we propose a new paradigm called Prediction-based monitoring for energy-efficient monitoring. We show that the paradigm can be visualized as a watching of a “sensor movie” and that concepts from MPEG may be applied to it. We have implemented the proposed algorithms in a testbed of Rene notes[Motes]. Experimental results show that the proposed solutions cut down the energy consumption by more than 5 times, considerably increasing sensor lifetimes, and thereby, the lifetime of the networks formed from these sensors.

1. INTRODUCTION

With the ever-increasing sophistication of hardware, it is becoming increasingly viable to deploy small units of autonomous nodes for performing various sensing tasks. Co-operative sensing, where a bunch of such nodes are thrown together to co-operate and perform a sensing task, has very useful applications in military and even common day-to-day situations. Some example applications include intrusion detection, gathering ground information like presence of snipers, tanks, etc. in the event of a war, or peak-hour traffic monitoring. Co-operative sensing raises several interesting research issues. A majority of these are rooted in the fact that sensors are typically limited by their energy reserves, communication bandwidth, and computational power. Of these, the energy constraint is the most crippling, because if the sensor runs out of battery, there can be neither communication nor computation. Another constraining characteristic that is specific to sensor networks is that the cost (human and others) of managing the sensor nodes far exceeds the cost of the nodes themselves. This is a result of three factors: firstly, the environment in which they are deployed could make human maintenance difficult (for e.g., consider the scenario where they are scattered over a high-altitude region). Secondly, it is expected that a reasonable-size sensor network will comprise of thousands of nodes. The sheer number of these nodes in a sensor network makes it difficult to manage. Lastly, increasing hardware integration and economy of scale are driving the prices of sensors down. All these factors make “use-and-throw” a very attractive option, where sensors that either have run out of battery or have failed are simply discarded and more sensors are thrown in to compensate for them.

Given these characteristics, the average lifetime of a sensor determines the cost of “running” a sensor network. Going by the current technology trend, although the computational power in a sensor is expected to follow Moore’s law, battery technology is only expected to improve by 2-3% per year[Mani, FPL]. Thus, the only possible way one can increase the lifetime of a sensor is by making use of mechanisms that are highly energy-efficient. Every operation at the sensor, be it transmitting data, receiving data, or performing a computation, consumes some energy. The goal of being energy-efficient, thus, translates into the problem of optimizing the number of these operations that need to be performed. At a sensor, the energy consumed in transmitting a packet is typically an order of magnitude higher than the energy consumed in receiving a packet. Also the energy consumed in receiving a packet is an order of magnitude higher than the energy consumed per instruction-execution [Wendi]. Given this relative cost, one may achieve a reduction in total energy consumption by cutting down the number of high-energy operations at the cost of an increase in the number of low-energy operations.

In this paper, our focus is on proposing mechanisms for performing monitoring in a sensor network in an energy-efficient way. Clearly the mechanisms associated with a traditional centralized database paradigm are unsuitable for our purposes. In such a system, a central server maintains a database of readings from all the sensors. Sensors update this server when their readings change. Monitoring operation is supported by the server, which maintains the current state of all the sensors involved in the operation. There are too many messages sent in such a system, making it very energy inefficient in many cases. We make two key observations to significantly improve the energy-efficiency of monitoring operation. Firstly, sensors in close proximity are likely to have correlated readings, and in a majority of the cases, one can predict the reading at a sensor given the knowledge of readings of sensors around it and their past history. An entity (for e.g., base-station, cluster-head) may exploit this observation and predict the set of readings that a sensor is going to see in the near future. These predictions are represented concisely as a “prediction-model” and sent to the sensor. The sensor now needs to transmit its sensed reading

* This research work was supported in part by DARPA under contract number N-666001-00-1-8953 and a grant from CISCO systems.

to the monitoring-entity only when it differs from the reading given by the prediction model by more than a certain pre-specified threshold. This mode of working gives us a new paradigm of operation in sensor networks. We call this the *PREdiction-based MONitoring (PREMON)* paradigm.

Our second key observation is that a snapshot of the sensor network may be visualized as an (optical) image – the readings of individual sensors correspond to intensity values of pixels in the image. Since a monitoring operation can be thought of as receiving a sequence of these snapshots on a continuous basis, one may visualize monitoring as watching a continuous sequence of corresponding images; in effect, watching a “video of sensed values”. Given this visualization, we explore if the concepts of MPEG [MPEG] (a standard for video compression) may be used for compressing this “video”. We show in section 3 that the MPEG encoder uses a paradigm that is an exact analogue of PREMON. This analogy with MPEG encoding gives us a convenient framework in which to visualize the problem. In addition, it also gives us a unique opportunity to adapt the well-established theory and algorithms of MPEG for use in sensor networks, and to examine the scope for cross-pollination between the two fields.

The PREMON paradigm prevents a sensor from unnecessarily transmitting all the readings that can be successfully predicted at the monitoring entity, thereby saving energy. This saving is obtained at the cost of extra computations at the monitoring-entity for generating prediction-models, and the extra cost of transmitting them. Given this tradeoff, clearly the effectiveness of the proposed paradigm is dependent on the accuracy with which prediction models can be generated and the percentage of readings that can be successfully predicted by them, without too much computational overhead. Based on the above observations, we propose algorithms in this paper and demonstrate their feasibility and effectiveness by implementing them on a test bed of real sensors. We show that it is feasible to generate prediction models for sensors and that significant savings can be achieved even with very simple methods of generating prediction models.

1.1 Background: Brief review of MPEG encoding process

In this section, we review MPEG-2 encoding briefly. Interested readers are referred to [MPEG] for details. MPEG-2 is a standard for audio and video compression. In this section, we limit ourselves to briefly describing only video compression. The video compression performed by MPEG falls in two classes:

1. *Spatial compression*: In this class of compression, a frame is encoded without reference to any other frame. This is based on the observation that in an image, there are many portions where the pixel value is the same. For e.g., “sky” in an image would be one such portion.
2. *Temporal Compression*: In this class of compression, a frame is encoded based on previous frames. We focus on this form of compression in the remainder of this section.

The idea of temporal compression is based on the observation that successive images in a video rarely have significant difference. MPEG exploits this observation by sending only the difference between the current image and the previous image instead of sending the entire current image. The decoder re-constructs the current image by adding the difference to the previous picture.




Fig. 1. Block Matching Algorithm in MPEG-2

The difference between successive images may become large if there is object movement in video. MPEG deals with this problem using *motion-compensation*. The encoder contains a motion estimator that measures the direction and distance of motion between successive images, and outputs these as *motion-vectors*. The motion-vectors determine the amount of horizontal and vertical shift required. These are sent to the decoder that uses them to shift data in the previous image to more closely resemble the current image. Since motion compensation can never be ideal, the encoder still needs to send an image-difference to take care of any shortcomings in motion-compensation. To compute this image-difference, it makes use of motion decoder logic. When outputting motion-vectors, the encoder also uses them locally in a same way that a decoder will, by feeding them into the motion-decoder logic to produce a predicted picture. This is

then subtracted from the actual current picture to produce a prediction error. Thus, the encoder encodes the current picture as a set of motion-vectors and a prediction error (fig. 2).

The decoder takes the previous picture, shifts it with the motion-vectors to recreate the predicted picture and then adds the prediction error to produce the actual picture. This process is shown in Fig. 2. Picture data sent as motion-vectors plus prediction error is called P-frame. The approach of sending a prediction error allows the motion estimation and compensation logic to be imperfect. Thus it allows for simple but inaccurate motion estimations in low-cost systems. If the prediction error is very large, the coder may simply transmit the current picture as it is. This constitutes an I-frame.

In order to generate motion-vectors in MPEG-2, the screen is divided into areas called macroblocks, which are 16x16 pixels square. Motion vectors are generated for each macroblock. One simple technique used for generating these is called *block-matching* (Fig. 1). In this technique, motion-vector for a macroblock is obtained by moving it around over the next consecutive picture looking for matching pixel values. When a match is found, the horizontal and vertical displacement needed to obtain it is used as a basis for computing motion-vectors for this macroblock. Thus the location of the boundaries of a macroblock is fixed and a motion-vector does not move the macroblock. Instead the vector tells the decoder where to find pixel data for the macroblock in another frame. The pseudo-code for the block-matching algorithm appears in figure 4 and 5.

1.2 Organization of the paper

This paper is organized as follows. In the next section we detail the model of sensor networks considered in this paper and state our assumptions about it. In section 3, we look at mechanisms based on PREMON paradigm for performing energy-efficient monitoring in sensor networks. We have implemented these mechanisms in a test bed of sensors. In section 4, we describe the experiments we performed and the results we obtained. In section 5, we propose extensions to the basic ideas in order to improve the accuracy of predictions. In section 7, we present the related work in the field and contrast our approach. In section 8, our plans for future work and present our conclusions in section 9.

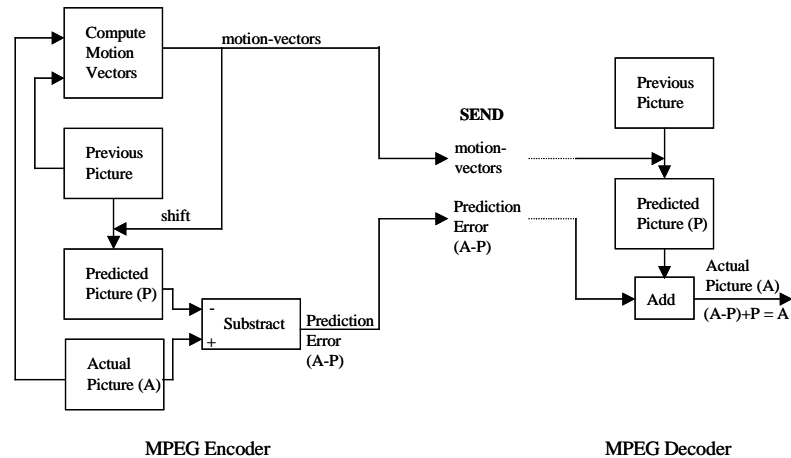


Fig. 2. MPEG encoding and decoding process

2. Model

There are many classes of sensor networks [lakme], each with its own characteristics. Mechanisms meant for one class may not be appropriate for other classes. Before we discuss the problem of optimizing these mechanisms, it is important that we define the class of sensor networks under consideration in this paper.

In [lakme], sensor networks have been classified according to their size and configuration. In this paper, we adopt as a model, the class of sensor network that is large, and has a non-deterministic¹ topology. Also, the sensors in this network may be multiple hops away from the nearest wired node. This, according to [lakme], is the most challenging class of sensor networks.

We assume that the sensors organize themselves in clusters, and that every cluster has a cluster-head². The exact mechanism for achieving this is orthogonal to our focus of supporting energy-efficient monitoring. A possible mechanism that may be used for this purpose has been proposed in [Wendi]. We assume the existence of some such mechanism.

¹ sensors are randomly thrown in the target field

² we use the words, cluster-head and base station, interchangeably in the paper to refer to the cluster-head

As is typically the case in these networks, we assume that sensors and base stations have limited energy and communicate sensed data over wireless link. Also, we assume that within a cluster, the base station and the sensors use TDMA at the link layer to communicate [Wendi]. The TDMA schedule is decided by the base station, which gets the first slot in the schedule. At the time the cluster is formed, the base station publishes the schedule, thereby communicating it to all the sensors in the cluster. They use it to appropriately switch between active and sleep mode.

We assume that the base stations in the sensor network communicate with the outside world via a few nodes on the wired network. We call these nodes the *access points* (wired base station in the figure) (fig. 3). We assume that every base station is statically pre-configured to communicate via one access point.

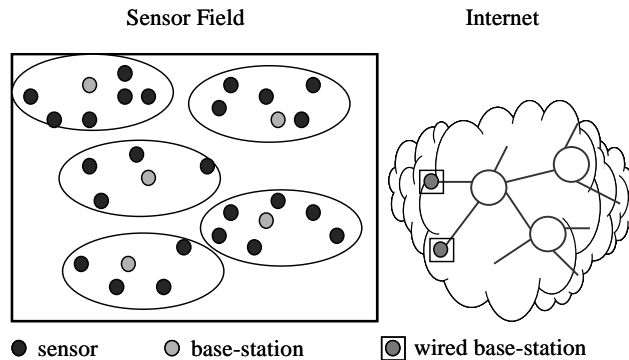


Fig. 3. Model of Sensor Network

3. Energy-Efficient monitoring

We start by describing a simple scheme for performing monitoring operations in sensor networks. We then describe how we may save energy by employing PREMON. Next, we present the details of the mechanisms based on the PREMON paradigm. Finally, we evaluate these mechanisms with simple experiments and show their effectiveness.

3.1 A naïve approach to monitoring

The simplest approach to supporting monitoring in sensor networks would be to have all the sensors report their sensed readings to the base station. The sensors update the base station whenever their reading changes. We refer to this mode of operation of sensors as “*update-mode*”. The base station collects these updates and sends them to an access point. Thus, at any point, a base station maintains the database of current readings of all the sensors in its cell. Also, the access points collectively maintain a database of current readings of all the sensors in the sensor field. Users interested in monitoring a query may register their interest with the appropriate access points[DataSpace]. We refer to this as the centralized approach.

3.2 Prediction Based Monitoring (PREMON)

“*It is not news if one can predict it*”.

It has been argued in [Pottie00] that a group of sensors (possibly of different types) is more effective in performing a sensing task than one powerful sensor. This argument favors a sensor network with a large number of sensors. In such a sensor network, a group of spatially proximate sensors are very likely to have correlated readings, where the correlation may be spatial, temporal, spatio-temporal, or a combination of these. We aim to exploit this correlation. We believe that in many instances one can predict a sensor’s reading based on the recent history of readings of the sensors around it and based on its own recent history. The sensor need not transmit its reading when it can be predicted by the monitoring entity. This gives us a new paradigm for operation in sensor networks. We call this the PREdiction-based MONitoring (PREMON) paradigm. The basic paradigm is as follows: an entity monitors the reading of sensors. Based on these, it generates *prediction-models*. A prediction-model is a concise representation of the readings that a sensor is expected to sense in the near future. These prediction-models are sent to the appropriate sensors. On receiving predictions sensors now send an update only when their readings differ from the one predicted by the prediction model (instead of sending an update every time their readings change).

3.2.1 Classes of prediction models

The prediction models may be classified into four categories:

- *Spatial*: These prediction models specify the reading at a sensor as a function of the readings at nearby sensors. Example: “*Reading at sensor X in time slot t is the same as reading at sensor Y during the same time slot*”.
- *Temporal*: These prediction models specify the reading at a sensor as a function of its readings in the past. Example: “*Reading at sensor X in time slot t is 2 greater than its reading in the previous time slot*”.

- *Spatio-temporal*: A combination of the above two. Example: “Reading at sensor X in time slot t is the same as the reading of sensor Y in the previous time slot”.
- *Absolute*: These prediction models directly specify the readings that a sensor is going to sense in the near future. The readings may be specified as a list of tuples $\langle \text{time}, \text{reading} \rangle$, or they may be specified by a function. Unlike spatial or temporal models however, the function neither depends on any past reading of the sensor, nor on the reading of any nearby sensor. Example: “Readings at sensor X in time slots t, t+1, and t+2 will be 32, 34, and 35 respectively”.

In some cases, one may need to represent spatial, temporal, or spatio-temporal prediction model as an absolute prediction model. This happens when the readings on which future readings depend, are not available at the sensor. For e.g., if a sensor only stores a limited amount of history, then readings needed to compute a temporal model may not be available. Similarly, if a sensor can hear only a few neighboring sensors, then the readings of sensors needed to compute a spatial model may not be available to it.

A prediction model is always qualified with a time interval in which it is valid. At the end of the time interval, the prediction model expires and the sensor reverts to the *update-mode* described earlier.

3.2.2 Key characteristics of PREMON

This paradigm trades increased computation (for computing prediction models) and slight increase in number of receptions (for receiving the prediction-models) for savings in number of typical transmissions. Note that the cost of transmission is orders of magnitude higher than the cost of computation and an order of magnitude higher than the cost of reception. If one is able to generate prediction-models without “too much overhead” and if they are correct a “majority” of times, we can achieve significant savings in energy at the sensors. Thus, performance of any system based on PREMON paradigm is a function of how computationally expensive is the operation of generating prediction-models, and how good is the “predictability” of the sensor readings. “Predictability” of sensor readings increases with the increase in degree of correlation among readings of sensors in close proximity, and with the increase in degree of correlation between recent history and the readings that a sensor is going to see in near future. At the same time, it also decreases with the decrease in the degree of correlation. For example, readings of sensors sensing wind direction, or readings of seismic sensors sensing earthquakes, have bad temporal correlation in general. Fortunately, a number of real world phenomenon like temperature, pressure, and motion, tend to have characteristics amenable to PREMON. We expect that generating a prediction model that makes precise predictions most of the times will be computationally expensive if not impossible. However, a more modest goal of generating a prediction model, which can predict within a small margin of error, should be achievable without too much computational overhead. Since “some” amount of error in reported data is tolerable in sensor networks (as is “small” delay), we expect the paradigm of prediction-based monitoring to work well in practice. The experimental results presented in section 4 bear this out.

3.2.3 Classes of Solutions

As we mentioned earlier, the correlation in readings of sensors in a sensor network may be spatial, temporal, spatio-temporal, or a combination of these. A wide range of techniques based on PREMON paradigm may be used to exploit each type of correlation. In this paper we focus primarily on exploiting *spatio-temporal* correlation. Our interest is in approaches that may be used over a wide range of sensor types. At the same time, we realize that the domain knowledge would play a significant role in giving final shape to the exact technique used, based on these approaches.

Spatio-temporal correlation is typically present in sensor networks sensing a moving phenomenon. For example consider a network of motion sensors. When an object passes through the field of these motion sensors, the readings of sensors along its path are very likely temporally correlated. In this paper, we propose techniques to exploit this correlation.

3.2.4 A New Framework for Visualizing the Problem

Let’s consider a specific example of a network of motion sensors (for e.g., magnetometer sensors) in a field. Assume that we are interested in monitoring the motion sensors lying in a small spatial region, $\langle \text{spatial-region} \rangle$. Assume for the sake of simplicity that all the sensors within the $\langle \text{spatial-region} \rangle$ fall within one cluster. Also assume that all the sensors operate in the *update-mode*.

When the monitoring operation begins, the sensors in the $\langle \text{spatial-region} \rangle$ transmit their current reading to the base station. Subsequently, they transmit updates to the base station as and when their readings change. Based on the data received by the base station, one can build an activity-map of the monitored region. Each pixel of this image represents a sensor in the monitored region, the intensity of the pixel indicates its reading, and its position in the image corresponds to the location of the sensors in the field. When visualized this way, the sensors may be seen as collectively sending a stream of activity-maps of the monitored region, one frame at a time. Since sensors send updates to the base station only when their reading changes, sensors may be seen as collectively sending the full frame initially, and then sending only the differences from the previous frame in subsequent frames. Figure 6 illustrates this view. Given this view of the monitoring operation, this whole process is analogous to how video is encoded in MPEG [MPEG]. The initial frame that contains data from all the sensors in $\langle \text{spatial-region} \rangle$ is analogous to an I-frame in MPEG. The subsequent frames (the difference-frames being sent by base station) are analogous to P-frames in MPEG (more precisely, P-frames with no motion-vectors). In fact, we can improve the analogy by performing prediction at the base stations of sensor networks.

```

/* Each frame can be thought of as composed of macro-blocks. Size of each macro-block is
 * blockWidth x blockWidth pixels. The function is passed two consecutive frames: frame1 and
 * frame2. For each macro-block in frame1, this function searches for the best-matching block
 * in frame2 (two macro-blocks are supposed to match if the sum of the difference in their
 * corresponding pixel values is less than a certain pre-specified threshold). Once a matching
 * block is found, the function records the displacement (in X and Y) as a motion-vector that
 * describes the movement of the macro-block from frame1 to frame2. The set of all motion
 * vectors considered are assumed to be stored in the array: PossibleMotionVectors
 */
function blockMatchingAlgo(frame1, frame2, blockWidth) returns Motion-Vectors

  foreach mblock in frame1 do
  {
    topLeftX1 = mblock.getTopLeftX();
    topLeftY1 = mblock.getTopLeftY();

    error = 0; minError = MAXINT;
    foreach mVector in PossibleMotionVectors do
    {
      <topLeftX2, topLeftY2> = applyMotionVector(topLeftX1, topLeftY1, mVector);
      if (isWithinFrame(topLeftX2, topLeftY2, frame1) {
        error = subtractMacroBlocks(frame1, topLeftX1, topLeftY1,
                                   frame2, topLeftX2, topLeftY2, blockWidth);

        if (error < minError) {
          bestMatchingVector = index;
          minError = error;
        }
      }
    }
    if (minError < ErrorThreshold)
      addNewMotionVector(motionVectors, topLeftX1, topLeftY1, mVector);
  }
  return motionVectors;

```

Fig. 4. Pseudo-code for block-matching algorithm in MPEG-2

```

/* This function finds the degree of match between the macro-block, starting at
 * <topLeftX1, topLeftY1>, of frame1, and macro-block, starting at <topLeftX2,
 * topLeftY2>, of frame2. Each macro-block is assumed to be of size -
 * blockWidth x blockWidth pixels. Function abs() gives the absolute value.
 */
function subtractMacroBlocks(frame1, topLeftX1, topLeftX2, frame2, topLeftX2,
                             topLeftY2, blockWidth) returns degree-of-match

  row1=topLeftX1; col1=topLeftY1;
  row2=topLeftX2; col2=topLeftY2;
  error = 0;
  while (row1 < row1+blockWidth) {
    while (col1 < col1+blockWidth) {
      error = error + abs(frame1[row1][col1] - frame2[row2][col2]);
      col1++; col2++;
    }
    row1++; row2++;
  }
  return error;

```

Fig. 5. Pseudo-code for block-matching algorithm in MPEG-2, cont'd...

Thus, a base station may perform motion estimation based on last couple of frames, and use it to predict a few frames in the future. The base station would then send the prediction-models (motion-vectors) to the sensors. Sensors now need to transmit only when their reading differs from the reading predicted by the prediction-models. This is very similar to what happens in MPEG where apart from the motion-vectors, the difference between the predicted and actual frame is also sent. If a significant fraction of the predictions are correct, then substantial energy savings may be achieved at the sensor.




Fig. 6. Illustration of MPEG analogy

Although we have considered the visualization and MPEG analogy in the context of a network of motion sensors, the proposed framework is very general and it can be applied to a network of any type (or types) of sensors. This is because the algorithms used in MPEG encoding do not make any assumptions about the underlying data, i.e., in some sense they are “blind” to the specifics, and can be applied to any form of image. It is important to note that the readings of sensors need not necessarily be single values like temperature. Each reading may, in general, be a set of values (for e.g., imagine a network of tiny cameras, where each reading of the camera is an image) and still the same algorithms may be applied. However there are differences between sensor monitoring and MPEG many of which in fact work in the overall favor of PREMON. In the next section we elaborate on these differences.

3.2.5 Differences between MPEG and PREMON

Monitoring operation in sensor networks does not have some of the constraints that are present in video encoding. We summarize the fundamental differences below along with their effect on the algorithms:

- The most prominent difference is that live video encoding has hard real-time constraint whereas monitoring operation in sensor networks has soft real-time constraint. The hard real-time constraint causes the MPEG encoder to generate motion-vectors based on a small number of previous image frames. Also, it uses very simple algorithms for the purpose. Thus, one may view motion-vectors in MPEG as based on patterns occurring on a very small time-scale. In sensor networks, since the real-time constraint is relaxed, the algorithms may look for patterns not only on small time-scale, but also on medium and large time-scales. Furthermore, the algorithms in sensor networks may afford to spend more time in generating prediction-models.
- In sensor networks, energy of the sensors is the most precious resource. This is not so in video encoding. As a result, the PREMON paradigm attempts to save energy by generating prediction-models at base station. While in MPEG, motion-vectors are a bandwidth-efficient way of encoding the current picture based on the previous picture. Although MPEG tries to be as bandwidth efficient as possible, this requirement is only soft. Whereas, in sensor networks, extreme energy-efficiency is a requirement. Thus, there is a need for improving the accuracy of predictions and the frequency with which such predictions are made.
- The resolution of the image that MPEG works on is very high compared to that of the image that PREMON works on. Given the small radio-range of a sensor (approximately 100 feet in outdoor environment), and that all the sensors in a cluster need to be in the radio-range of the base station, we expect that the number of sensors (sensing a moving phenomenon) will be very small compared to the number of pixels in an image frame of a video. This small size significantly reduces the computation involved in performing motion-estimation.
- The frame rate in MPEG is an order of magnitude higher compared to the expected frame rate in PREMON. This allows techniques in PREMON to do motion estimation/compensation based on a set of previous frames.

- Since sensors are rarely hand-placed in the target environment, the arrangement of sensors in the field is going to be random. Thus the image visualization of the sensor network will have non-uniform placement of pixels in the image. On the other hand, the pixels in a video have a very uniform pattern. In section 3.2.6, we discuss techniques to handle this.

In the next section, we describe the architecture of the system.

3.2.6 Architecture

We adapt the architecture of MPEG encoder for our purposes. In this section, we describe in detail the processing at the base stations and at the sensors.

Processing at the Base Station

The MPEG-like encoding at the base station involves the following processing steps:

- Collect updates from the sensors. The updates received in the same TDMA cycle form a frame (of the sensor stream). They are fed into a prediction model generation unit.
- Prediction model generation unit looks at a set of successive frames and computes prediction models. A prediction model predicts the direction of motion for a “macro-block” in the frame for next few time units. This prediction model is sent to the sensors.
- In the next TDMA cycle, the base station assumes that the prediction model predicted readings of all sensors correctly and forms the next frame. The sensors whose readings don’t match with the prediction model will send an update thereby overriding their assumed readings in the frame.
- If sending the prediction model resulted in fewer updates, then the base station encodes the frame as a set of prediction models, followed by the updates necessary to override wrong predictions (prediction errors). It sends this to an access point.

Going by the MPEG analogy, if a base station is low on energy, it may send out “low-quality images” in order to cut down the amount of data being sent. Here, by “low-quality” we mean less accurate information of the sensor readings. For e.g., the base station may divide its cluster in spatial blocks and may only send average reading of each block to the access point.

Processing at the Sensor

A sensor is in update mode by default – it sends an update to the base station whenever its reading changes. When it receives a prediction-model from the base station, it alters its behavior slightly. For the duration for which the prediction-model is valid, it compares its sensed reading with the reading predicted by the prediction-model. It sends the former to the base station only when it differs from the predicted reading by more than some pre-configured margin of error. When the prediction-model expires, the sensor reverts back to the default update-mode.

Computing a Prediction Model

In this section we discuss prediction-model generation mechanism in more detail. Although there may be many different ways of generating prediction models, our focus is on using mechanisms that are generally applicable to a wide range of sensor networks.

We use the block-matching algorithm in MPEG to compute prediction models. Note that the position of sensors in a cell is not uniform, and hence, in the corresponding visualized image the pixels will not be uniformly positioned. The block-matching algorithm cannot be applied directly to such an image. In order to correct this problem, one may use the sensor readings to interpolate (or extrapolate) the readings at regular grid points in the image using one of the several standard techniques [Interpolation]. As a first cut, we employ a very simple interpolation mechanism. This mechanism assigns the reading of the closest sensor to a grid point. If more than one sensor happen to be closest to a grid point, the mechanism assigns the average of their readings to the grid point. Since the sensors may potentially be unevenly distributed in the region, it is possible that a grid point in the visualized image does not have any sensor “nearby”. We deal with this problem by labeling such a grid point as “transparent”. The significance of such a labeling will become clear shortly.

After interpolation, the visualized image will have pixels at regular intervals. Some of these pixels may be “transparent”. To deal with the “transparent” pixels we modify the standard block-matching algorithm. As in standard block-matching algorithm, the image is divided into macro-blocks. For each macro-block in a frame, the algorithm looks for a match by moving it over the next consecutive frame. If a match is found, it computes the motion-vectors. During this process, we give special treatment to “transparent” pixels. We assume that a “transparent” pixel matches any other pixel. Also, we try to compute the motion vectors for a macro-block only when the percentage of transparent pixels in it is less than a certain threshold. Fig. 8 gives the pseudo-code for the modified block-matching algorithm. When determining the degree of match between two macro-blocks, the original MPEG algorithm (fig 5) would simply subtract the intensity values of the corresponding pixels and sum these up. However, in a sensor network, each reading of a sensor may not be a single value but a set of values. In such a case, a simple subtraction will not work. Instead, we compute the degree of match between the two readings. The way this is computed is dependent on the type of sensor. Figure 9 gives the pseudo-code for the modified algorithm.

In order to generate a prediction model, the base station works with the four most recent frames. It first applies the above algorithm to frames 1 & 2. Let's assume that this gives us the set, M , of motion-vectors. Each motion-vector in M indicates "motion" (the case of no motion can be seen as a special case of "motion"). For each motion vector in M , the base station checks frames 2 & 3, and 3 & 4, to see if the motion continues and if it can be described by the same motion-vector. If so, we say that the motion-vector "holds". As an example, consider the four consecutive frames shown in figure 7. The motion vector $\langle dx, dy \rangle$ "holds" for these four frames. If a motion-vector does not hold, it is discarded. If a motion-vector holds, the base station assumes that the motion described by it will continue for the next few frames in the future, and generates an absolute prediction model based on it. This model contains the readings of sensors in the macro-block of frame 1. Depending on the type of sensors, this information may be encoded in a more efficient form. For example, in our test bed, the magnetometer sensors [Motes] output a binary value – LOW/HIGH. LOW indicates that no (metallic) object is present in the "vicinity". HIGH indicates that an object is in the vicinity of the sensor. A macro-block of readings of magnetometer sensors consists of 0s and 1s. In this case, instead of sending the entire macro-block, one can find the largest rectangular block of 1s and send the co-ordinates of this rectangle (8 bytes of information), thereby substantially reducing the data that needs to be sent. The prediction model would then be valid for only the sensors within the rectangle and not for the entire macro-block.

This model is sent to the geographic area containing the sensors that are going to see the motion in the next four time frames. Note that at any time instant only a portion of the target area is going to see the motion. In order to indicate this to the sensors, we mention the target of the model as of type GEO_TEMPORAL. This allows the sensors in the target region to compute their predicted value correctly.

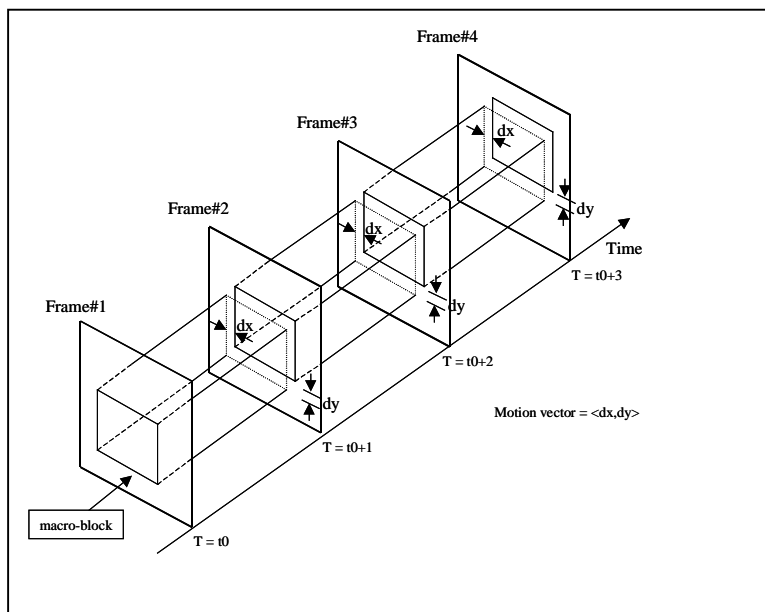


Fig .7. Movement of a macro-block in four consecutive frames

A special case occurs when a motion-vector indicates "no motion" (i.e., the horizontal and vertical displacement are both zero) and it holds for four most recent frames. This indicates that the readings of sensors in a macro-block are not changing over time. In such a case, instead of sending the entire macro-block, the base station simply sets a flag to indicate that the values are not expected to change for the next few frames. Figures 10-11 give the pseudo-code for generating prediction models.

Communicating a Prediction Model

Any prediction model can be communicated to sensors by specifying four components in a message:

- *Type*: Prediction models may be one of the four types:
 - Absolute
 - Spatial
 - Temporal
 - Spatio-temporal
- *Model*: The representation of model changes with the type. It may be represented as a list of tuples $\langle \text{time}, \text{reading} \rangle$, or as a function (for e.g., polynomial).
- *Destination*: The destination of the prediction model decides the sensors for which it is targeted. The destination may be specified as a broadcast address (the targets are all the sensors in the cell), sensor id (the target is a specific sensor), or a spatial polygon (the targets are all the sensors that are physically present inside the polygon).

- TTL: This field gives the time duration for which the prediction model is valid.

3.2.7 Effect of losses

Packet losses are quite common on wireless links. They are especially common in sensor networks because of primitive channel encoding used. We briefly look at their effect on the sensors operating in PREMON paradigm.

When a prediction is lost on the wireless link, the sensor continues to operate in default mode. The only effect of this loss is that the sensor will end up transmitting its reading even though it was expected/predicted at the base station. A more interesting case occurs when updates are lost while a prediction is active at a sensor. Suppose that a prediction is active at a sensor and the sensor's reading doesn't match with the predicted value. This would cause the sensor to send an update to the base station. If this update is lost on the wireless link, the base station would wrongly assume that the sensor's reading conformed with the prediction --- a false positive. We believe that this problem is no worse than the one caused by loss of updates (from sensors operating in default mode) on the wireless link, which causes the base station to not know what is going on in the sensor field. Since it is quite common to have group of sensors in close proximity in sensor networks, these problems caused by loss of updates will be controlled by the inherent correlation in their readings. As long as at least one of the updates of sensors from a group monitoring the same phenomenon reaches the base station, we will be able to effectively monitor the sensor field despite channel errors.

4. Experiments: Setup and Results

In this section we present the results from experiments that we have conducted to evaluate the PREMON paradigm. We first describe the experimental setup.

4.1 Experimental Setup

Our experimental setup consists of a set of light sensors. Each light sensor is fitted on a Rene Mote [Motes], a tiny computation device (about the size of a quarter) with communication capability. A mote has an 8-bit micro-processor (ATMEL AVR 90LS8535 running at 4 MHz) and a small flash memory (8KB program memory and 512 bytes data memory). The motes run a tiny microthreaded operating system called TinyOS [TinyOS]. A set of tools [Tools] allow the motes to be extensively and easily programmed using a combination of C language and a component description language. One of the main design criteria for motes was low power consumption. The mote runs in one of the three modes: active, idle, and inactive. In active mode, with processor, sensor, and radio running, it consumes 19.5 mA of current and can run for around 30 hrs on 2 AA batteries. In inactive mode, the mote consumes only 10 μ A of current and can run for a

```

/* This function is a slight variation of the block-matching algorithm of MPEG (fig 4).
 * The difference is that a macro-block in the image visualization of sensor data can
 * have transparent pixels. If their number exceeds a threshold
 * (TransparentPixelThreshold) then the macro-block is not processed
 */
function newBlockMatchingAlgo(frame1, frame2, blockWidth) returns Motion-Vectors

  foreach mblock in frame1 do
  {
    if (getNumTransparentPixels(mblock) > TransparentPixelThreshold) continue;

    topLeftX1 = mblock.getTopLeftX();
    topLeftY1 = mblock.getTopLeftY();

    error = 0; minError = MAXINT;
    foreach mVector in PossibleMotionVectors do
    {
      <topLeftX2, topLeftY2> = applyMotionVector(topLeftX1, topLeftY1, mVector);
      if (isWithinFrame(topLeftX2, topleftY2, frame1) {
        error = newSubtractMacroBlocks(frame1, topLeftX1, topLeftY1,
                                       frame2, topLeftX2, topLeftY2, blockWidth);

        if (error < minError) {
          bestMatchingVector = index;
          minError = error;
        }
      }
    }

    if (minError < ErrorThreshold)
      addNewMotionVector(motionVectors, topLeftX1, topLeftY1, mVector);
  }
return motionVectors;

```

Fig. 8. Pseudo-code for the variation of block-matching algorithm of MPEG-2.

The difference from the original algorithm (Fig. 4) appears in bold

year [TinyOS]. It consumes 1 μJ (on an average) for transmitting one bit and 0.5 μJ to receive 1 bit. Furthermore it consumes 0.8 μJ of energy for executing 208 machine cycles (roughly 100 instructions) [TinyOS]. These motes have a radio (RFM TR1000) on-board that works at 916 MHz. The radio can support a maximum bit rate of 19.2Kbps (on/off keying). The light sensors are not soldered on the mote. Instead, they sit on a tiny sensor board that can plug onto the Rene Mote via a small onboard connector (fig 12). The light sensor measures light intensity and gives a reading between 0 and 7, where 0 represents lowest level and 7 represents the highest level.

```

/* This function finds the degree of match between the macro-block, starting at
 * <topLeftX1, topLeftY1>, of frame1, and macro-block, starting at <topLeftX2,
 * topLeftY2>, of frame2. Each macro-block is assumed to be of size -
 * blockWidth x blockWidth pixels. Function degreeOfMatch() gives a quantitative
 * measure of the degree to which two pixel values match each other.
 */
function newSubtractMacroBlocks(frame1, topLeftX1, topLeftX2, frame2, topLeftX2,
                               topLeftY2, blockWidth) returns degree-of-match

    row1=topLeftX1; col1=topLeftY1;
    row2=topLeftX2; col2=topLeftY2;
    error = 0;

    while (row1 < row1+blockWidth) {
        while (col1 < col1+blockWidth) {
            if (frame1[row1][col1].type() == TRANSPARENT //
                frame2[row2][col2].type() == TRANSPARENT)
                continue;
            else
                error = error + degreeOfMatch(frame1[row1][col1] - frame2[row2][col2]);
                col1++; col2++;
            }
            row1++; row2++;
        }
    }
    return error;

```

Fig. 9. Pseudo-code for the variation of block-matching algorithm of MPEG-2.

The difference from the original algorithm (Fig. 5) appears in bold

```

/* This function takes a set of four consecutive frames and computes motion-vectors
 * that holds for all these frames. It assumes that these motion-vectors will hold
 * for next 4 frames (in future), assuming constant speed of movement. Based on this
 * assumption, it translates the motion-vectors into prediction-models.
 */
function generatePredictionModels(frames[], blockWidth) returns Prediction-Models

    validMotionVectors = newBlockMatchingAlgo(frames[0], frames[1], blockWidth);

    foreach vmVector in validMotionVectors do {
        topLeftX1 = vmVector.topLeftX;
        topLeftY1 = vmVector.topLeftY;
        mVector = vmVector.mVector;

        <tempTopLeftX, tempTopLeftY> = applyMotionVector(topLeftX1, topLeftY2, mVector);
        <topLeftX2, topLeftY2> = applyMotionVector(tempTopLeftX, tempTopLeftY, mVector);
        error = newSubtractMacroBlocks(frame[0], topLeftX1, topLeftY1,
                                       frame[2], topLeftX2, topLeftY2);

        if (error < ErrorThreshold) {
            <topLeftX3, topLeftY3> = applyMotionVector(topLeftX2, topLeftY2, mVector);
            error = newSubtractMacroBlocks(frame[0], topLeftX1, topLeftY1,
                                           frame[3], topLeftX3, topLeftY3);

            if (error < ErrorThreshold) {
                newModel = translateVectorToModel(frame[0], vmVector, blockWidth);
                addNewModel(newModel, predictionModels);
            }
        }
    }
    return predictionModels;

```

Fig. 10. Prediction-model generation algorithm

```

function translateVectorToModel(frame, vmVector, blockWidth) returns Prediction-model
mVector = vmVector.motionVectorIndex;
topLeftX = vmVector.topLeftX;
topLeftY = vmVector.topLeftY;
newModel = NULL;

if (mVector.displacementX == 0 && mVector.displacementY == 0) {
    /* Send a temporal model */
    newModel.type = TEMPORAL;
    newModel.model = {FUNCTION, f(x, t) = f(x, t-1)};
    newModel.destination = {GEO, topLeftX, topLeftY,
                            topLeftX+blockWidth, topLeftY+blockWidth};
    newModel.TTL = 4; /* valid for next 4 time units */
    return newModel;
}

/* Send an absolute model */
macro_block = extractMacroBlock(frame, topLeftX, topLeftY, blockWidth);
newModel.type = ABSOLUTE;
newModel.model = {ABSOLUTE, size=blockWidth*blockWidth, macro_block};
newModel.TTL = 4; /* valid for next 4 time units */

destPolygon = computeDestPolygon(topLeftX, topLeftY, mVector, newModel.TTL);
newModel.destination = {GEO_TEMPORAL, destPolygon};

return newModel;

```

Fig. 11 . Pseudo-code for translating motion-vectors into prediction models




Fig.12 Rene Mote and sensor board




Fig. 13 A motion-sensor assembly

4.1.1 Creating a motion-sensor

We create a motion-sensor for our experiments by using a mote fitted with a light sensor (we call this *sensor-mote*) and a focused light source. The sensor-mote is placed in a small box in such a manner that light from the pin-hole falls directly on the light sensor. The light source is placed at a distance from the box, with its light falling on the pin-hole (Fig. 13).

In order to be able to detect rapid events, the sensor-mote samples the light sensor 8 times per second. However, in order to make efficient use of the channel bandwidth, the sensor-mote transmits the light sensor readings once per second. Each transmission contains the eight most recent readings of the light sensor.

When no activity occurs, the light sensor measures a constant reading of 7. When an object passes between the light source and the box containing sensor-mote, thereby obstructing the light path, the readings of the light sensor changes. Fig 14 shows the typical fluctuation in light sensor readings as an object crosses the light path. These fluctuations are reported the next time the mote transmits. The next set of readings transmitted by the mote will report these changes in the light readings.

4.1.2 Setup

Our experimental setup consists of four motion sensors placed along a busy corridor (in our case we had placed them on the third floor of our department's building) (fig. 20r). Another mote acts as the base station. All motion sensors report to the base station. Following the PREMON paradigm, another mote, serving as base station, analyzes the readings of the motion sensors and generates predictions about the future readings of the motion sensors. The motion sensors store the predictions received and transmit their readings only when they don't match with the predictions.




Fig. 14 Typical light sensor readings when an object passes by

For the base-station-mote, all four sensors reside in the same cluster. It divides the cluster into four macro-blocks such that there is one sensor per macro-block. It makes predictions based on the last two frames received.

The base station makes two types of predictions. We briefly discuss each type of prediction and describe how the domain knowledge has shaped the exact technique implemented in the mote to generate them.

1. *Constant values predictions:* These predict that the reading of a sensor is not going to change. Since sensor readings change only when an object is right in front of them, in our experiments they see periods of activity (a person passing by) followed by relatively long periods of inactivity. When there is no activity, the sensor-motes read a constant value. When the last two frames indicate no change in the reading of a macro-block, the base-station-mote generates this prediction. It always associates a TTL of infinity with this type of prediction.

The packet carrying constant prediction contains two fields: type of prediction and the stable value of the sensor. It is 5 bytes in length (including 3 bytes of header).

2. *Movement prediction:* These predict that a sensor is about to see a pattern in its reading. As we stated earlier, the sensor-motes transmit the sampled readings of the light sensor at a maximum rate of once every second. Thus, the frame rate at the base-station-mote is 1 frame/sec. Assume that a person starts from one end of the corridor and walks to the other end, without stopping, at almost constant speed. Also assume that there is no other activity in the corridor. The person may take longer than one second to cross the light path of next motion-sensor from the instant the previous sensor detects him if he is walking leisurely. Thus, the base-station-mote will see a frame with one macro-block reporting activity, followed by a sequence of frames (*no-activity frames*) that do not show any activity, followed by a frame (*activity-frame*) that shows activity indicating that the person finally passed by the next motion sensor. The base-station-mote timestamps each frame. When applying the block-matching algorithm, it examines only the set of most recently received activity frames, discarding all the intervening no-activity frames. The time difference between the most recent activity frame and least recent activity frame in the examined set is not allowed to exceed a pre-defined constant value, `ALLOWED_STALENESS`. This constant determines the maximum number of frames considered by the block-matching algorithm. In the block-matching algorithm, the base-station-mote considers only three motion-vectors $\langle dx=0, dy=0 \rangle$, $\langle dx=1, dy=0 \rangle$, and $\langle dx=-1, dy=0 \rangle$, corresponding to *no-motion*, *motion from left-to-right* and *motion from right-to-left*, respectively. For each macro-block in the first frame, the base-station-mote examines all the remaining frames in the set to look for a matching macro-block. The algorithm generates a prediction after finding just one match. The prediction is generated only for the next sensor along the direction of motion.

In order to determine if two macro-blocks match (fig 9), it is important to specify how one can determine the degree to which readings of two sensor-motes match. Since each transmission of sensor-mote consists of 8 readings, the simplistic approach based on subtracting the two readings will not work. In order to determine the degree of match, we observe that as a person crosses the light path, the readings of the light sensor are going to drop from its stable value, fluctuate a bit, and finally, return back to the original stable value (see fig. 14). We refer to the string of values between the two stable readings as *pattern-string* (for e.g., the reading showed in fig 14 contains the pattern string: 3,0,0,0,3). We compute the degree of match between two sensor-mote readings by computing the absolute value of the difference between the lengths of the pattern-string contained in them.

In the prediction, the base-station-mote specifies the interval in which the next sensor, along the direction of motion, is most likely to see the pattern, and the expected length of the pattern-string. The time interval is determined based on the assumption that the person moves at a constant speed. The expected length of the pattern-string is set to the average of the length of the pattern-string seen at the two motion-sensors.

A movement-prediction is said to hold at a sensor if it sees a pattern-string during the specified time interval $\pm \Delta$, and the length of the pattern-string is within ∂ of the specified length. Thus, if the specified time interval is $[t_1, t_2]$, then we allow the pattern to occur in the interval $[t_1 - \Delta, t_2 + \Delta]$. Also, if the expected length of the pattern-string is $(l_1 + l_2)/2$, then all pattern-strings with length in the interval $[(l_1 + l_2)/2 - \partial, (l_1 + l_2)/2 + \partial]$ match with the prediction.

When a sensor detects a pattern as predicted, it starts expecting constant reading from the light sensor, as if it has just received a constant prediction. Surely, if an object has passed by, the light sensor will read the value 7 indicating absence of any perturbations.

A packet carrying movement prediction contains 5 fields: the type of prediction, the lower and higher time interval, the expected length of the pattern, and the expected stable reading of the sensor once the event has occurred. The size of the prediction packet is 8 bytes (including the 3 bytes of header).

4.2 Result and Analysis

As our initial set of experiments, we allowed the setup described above (fig. 15) to run for half an hour. We ran two such experiments. Table 1 contains the results averaged over both the experiments (note that sensors were assigned addresses: 4, 8, 12, and 14, from left-to-right (fig. 15)). Since these experiments are preliminary ones, it is important that the results presented in this section be treated qualitatively rather than quantitatively.

These experiments give us enough trace data to determine the energy consumption at each sensor-mote and the base-station-mote, in the following three cases:

1. *Case#1: All Sensor-motes operate in default mode:* In this case, the sensor-motes transmits its readings periodically (once every second).
2. *Case#2: Base-station-mote generates only constant-value predictions:* In this case, the sensor transmits its reading whenever it changes.
3. *Case#3: Base-station-mote generates both constant-value and movement predictions:* All sensor-motes follow the techniques presented in this paper.

The energy consumed in the mote for each of the basic operations is listed in table 2³. In all the calculations below, we assume that there are no packet losses.

Table 1. Table summarizing the results of experiments

Sensor Addr.	Number of Predictions received		Number of transmissions saved due to:	
	Constant Predictions	Movement Predictions	Constant Predictions	Movement Predictions
4	81	60	972	595
8	77	66	924	645
12	86	54	1020	525
14	72	58	864	720

4.2.1 Case#1: Default mode

In this mode, the energy consumption at each sensor-mote is easy to compute. Since the experiment runs for 30 mins, each sensor will send 1800 (30x60) updates (at the rate of once/sec) during the experiment. Therefore, the energy consumed at each sensor-mote during the entire run of the experiment is:

$$E_s = 1800 \text{ (updates)} \times T_u = 158.4 \text{ mJ}$$

The base-station-mote in this mode receives 1800 updates from each sensor during the entire run of the experiment. Therefore, the energy consumed at the base-station-mote is:

$$E_b = 4 \times 1800 \times R_u = 316.8 \text{ mJ}$$

The total energy consumed in the system is given by:

$$TE_1 = 158.4 \times 4 + 316.8 = 950.4 \text{ mJ}$$

4.2.2 Case#2: Constant-value predictions only

Let's first consider the sensor with address 4.

$$\text{Total Energy consumed } (TE_4) = \text{Energy consumed in transmitting updates} + \text{Energy consumed in receiving predictions}$$

The base-station makes 60 movement predictions. In the absence of these, there will roughly be 60 instances where an object will pass by this sensor (assuming, for simplicity, that predictions were correct). In each of these instants, the sensor will generate 3 updates for each

³ Table 2 also assigns a symbol for each of the basic operation. We use these symbols in the rest of the paper.

movement predicted by the base-station-mote -- one update is generated when it detects a change in reading. The other two are generated because base-station-mote does not make a constant value prediction until it has seen two consecutive sets of constant readings. Thus, the sensor will see 141 (81+60) constant value predictions and it will save 415 (595 – 60x3) transmissions of updates.

Thus,

$$TE_{S4} = [1800 - (972 + 415)] \times Tu + 141 \times Rcp = 39.16 \text{ mJ}$$

Following similar approach, the total energy consumed at sensors with addresses 8, 12, and 14 is:

$$TE_{S8} = 411 \times Tu + 143 \times Rcp = 39.02 \text{ mJ}$$

$$TE_{S12} = 435 \times Tu + 140 \times Rcp = 41.08 \text{ mJ}$$

$$TE_{S14} = 396 \times Tu + 130 \times Rcp = 37.40 \text{ mJ}$$

The total energy consumed at base-station-mote in receiving updates and sending predictions is:

$$TE_p = (413+411+435+396) \times Ru + (141+143+140+130) \times Tcp = 94.98 \text{ mJ}$$

We assume that the cost of generating each constant-value prediction is negligible.

The total energy consumed in the system is given by:

$$TE_2 = 39.16 + 39.02 + 41.08 + 37.40 + 94.98 = 251.64 \text{ mJ}$$

4.2.3 Case#3: Constant-value and movement predictions

We assume that the energy consumed in computing each movement prediction at the base-station-mote is given by Cm .

$$TE_{S4} = [1800 - (972 + 595)] \times Tu + 81 \times Rcp + 60 \times Rmp = 24.04 \text{ mJ}$$

$$TE_{S8} = [1800 - (924 + 645)] \times Tu + 77 \times Rcp + 66 \times Rmp = 23.98 \text{ mJ}$$

$$TE_{S12} = [1800 - (1020 + 525)] \times Tu + 86 \times Rcp + 54 \times Rmp = 25.88 \text{ mJ}$$

$$TE_{S14} = [1800 - (864 + 720)] \times Tu + 72 \times Rp + 58 \times Rmp = 22.30 \text{ mJ}$$

The total energy consumed at base-station-mote is the sum of the energy consumed in receiving updates, generating predictions, and sending predictions. For the sake of simplicity, we assume for all updates, except ones that result in constant-value predictions, the amount of processing performed by the block-matching code is the same. This is clearly the worst-case assumption. The total energy consumed at the base station is given by:

$$TE_p = (233+231+255+216) \times Ru + (81 + 77 + 86 + 72) \times Tcp + (60 + 66 + 54 + 58) \times Tmp + (233+231+255+216 - 81 - 77 - 86 - 72) \times Cm \\ = 935 \times Ru + 316 \times Tcp + 238 \times Tmp + 619 \times Cm = 69012 + 619 \times Cm$$

We can determine the value of Cm to a good approximation as follows: The tools provided with the motes allow one to generate a map of the mote's memory. Using this memory map, one can determine the size of any function that is part of the program. We determined that the block-matching function (for computing movement prediction) occupied approximately 1,200 bytes of program memory. In the worst case, every byte of will be a separate machine instruction. Thus, the function has 1,200 machine instructions. Since we consider only two motion-vectors ($\langle dx=1, dy=0 \rangle$ and $\langle dx=-1, dy=0 \rangle$) (ignoring the cost of constant-value prediction), at worse, all the instructions in the block-matching code will be executed twice for every update received. Thus, the worst-case cost of Cm is the energy consumed in executing 2×1200 instructions. This is given by:

$$Cm = 2400 \times 0.8/100 \mu\text{J}/\text{update-processing} = 19.2 \mu\text{J}/\text{update-processing} \quad \dots (1)$$

With this worst-case computation, we can compute the energy consumed at the base-station-mote in the worst case. This is given by:

$$TE_p = 69012 + 619 \times 19.2 = 80.9 \text{ mJ}$$

This represents an upper bound on the energy consumed at the base-station-mote. The total energy consumed in the system is given by:

$$TE_3 = 24.04 + 23.98 + 25.88 + 22.30 + 80.9 = 177.1 \text{ mJ}$$

Operation (Symbol)	Packet Type	Packet Size (bytes)	Per-bit cost ($\mu\text{J}/\text{bit}$)	Cost (μJ)
Transmit (Tu)	Update	11	1	88
Receive (Ru)	Update	11	0.5	44
Transmit (Tcp)	Constant value prediction	5	1	40
Receive (Rcp)	Constant value prediction	5	0.5	20
Transmit (Tmp)	Movement Prediction	8	1	64
Receive (Rmp)	Movement Prediction	8	0.5	32

Table 2. The cost of basic operations in the mote

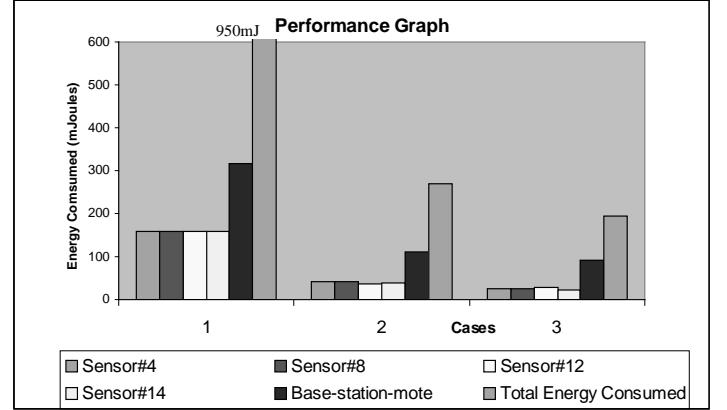


Fig. 15. Performance Graph

Thus, the total energy consumption in this case is close to **5.3 times** lesser than one in case#1, and **30%** lesser than one in case#2. We expect case#3 to perform significantly better than the other cases when more sensors are involved.

4.3 When is movement prediction more energy efficient ?

In this section, we establish conditions under which a scheme generating motion-prediction and constant-value prediction (case#3) will perform better than the one generating only constant-value prediction (case#2).

Let us assume that the average success rate for movement predictions is s (expressed as percent of the total number of movement predictions), and the cost of computations in order to generate a movement prediction is C_m .

Assume that a motion-sensor is in steady state (i.e., the light sensor reading is constant over time). When an object passes by, the motion-sensor will see a pattern and report this to the base-station. Assuming that the object took less than one second to pass by, the reading of this motion-sensor in the next frame would be its constant stable value. It will report this constant value for two consecutive frames before the base station generates a constant-value prediction. Thus, whenever an object passes by a motion-sensor, it results in three updates being sent to the base station. If the base-station is able to correctly predict the movement of an object, it can save these three transmissions. But in order to predict correctly, it needs to process the readings of two motion-sensors that this object triggered before passing the current motion-sensor. Thus, the total energy consumed in this case is:

$$\begin{aligned}
TE_3 &= s \times \text{energy consumed when the prediction succeeds} + (1-s) \times \text{energy consumed when the prediction fails} \\
&= s \times (\text{energy consumed in sending, receiving, and processing updates sent by the two motion sensors} + \\
&\quad \text{energy consumed in sending constant value predictions to the two sensors} + \\
&\quad \text{energy consumed sending and receiving a movement prediction}) + \\
&\quad (1-s) \times (\text{energy consumed in sending, receiving, and processing updates sent by the two motion sensors} + \\
&\quad \text{energy consumed in sending constant value predictions to the two sensors} + \\
&\quad \text{energy consumed in sending and receiving a movement prediction} + \\
&\quad \text{energy consumed in sending and receiving update due to failed prediction} + \\
&\quad \text{energy consumed in sending constant value predictions to the sensor}) \\
&= s \times (2 \times (132x3 + C_m + (40+20)) + (64+32)) + (1-s) \times (2 \times (132x3 + C_m + (40+20)) + (64+32) + 132x3 + (40+20)) \dots \quad (2)
\end{aligned}$$

In case#2, the total energy consumed would be:

$$\begin{aligned}
TE_2 &= \text{energy consumed in sending and receiving updates from the three motion-sensors} + \\
&\quad \text{energy consumed in sending constant value predictions to the three sensors} \\
&= 3 \times (132x3 + (40+20)) \dots \quad (3)
\end{aligned}$$

In order for case#3 to perform better than case#2, the following condition must be satisfied:

$$TE_3 \leq TE_2$$

From equations (2) and (3), we get,

$$s \times (2 \times (132x^3 + Cm + (40+20)) + (64+32)) + (1-s) \times (2 \times (132x^3 + Cm + (40+20)) + (64+32) + 132x^3 + (40+20)) \leq 3 \times (132x^3 + (40+20))$$

Simplifying, we get,

$$s \geq (2 \times Cm + 96)/456 \quad \dots (4)$$

Substituting the value of Cm from equation (1), we get,

$$s \geq 0.294$$

Thus, even if the base station predicts successfully 30% of the time, the scheme with movement prediction (case#3) will perform better than the scheme with just constant-value prediction (case#2).

Solving equation (4) for Cm, we get,

$$Cm \leq (456 \times s - 96)/2 \quad \dots (5)$$

From table 1, we observe that the average success rate is approximately 75%. Substituting this value of s in equation (5), we get,

$$Cm \leq 123 \mu J$$

Thus, as long as we consume less than 123 μJ of energy in computing a processing an update, the scheme with movement prediction (case#3) will perform better than the scheme with just constant-value prediction (case#2). Given that 100 instructions consume 0.8 μJ of energy in the mote, 94.5 μJ is enough to perform 15,375 instructions. This, we believe, is more than enough to generate movement prediction to a good accuracy.

5. Improving quality of predictions by extracting patterns at medium and large time-scales

In the previous section, we looked at a possible way of extracting pattern at smaller time-scale and generating prediction models based on that. We believe that one may improve the accuracy of the predictions by extracting patterns that occur frequently at medium and larger time-scales and using them to improve predictions on smaller time-scales. As an example, consider a sensor network consisting of acoustic sensors thrown over a battlefield. As a tank approaches an acoustic sensor, the sensor will record higher amplitude in the sensed reading. The amplitude of the sensed reading will reach a peak value and will start decreasing as the tank passes by. If one were to look only at pattern on a smaller time-scale, by seeing a steady increase in amplitude of the reading, one would predict that the amplitude of the sensed reading would continue to rise. Similarly, by noticing a steady decline in amplitude, one would predict that the amplitude of the sensed reading will continue to decrease. However, by processing the data over a longer time period, one would be able to extract the signature pattern (steady rise followed by steady decline) of a tank. If this information is given to the base stations, they will be able to identify the pattern just by looking at the steady increase in amplitude. They would then predict not only the steady rise in amplitude but also the steady decline following that. Assuming that the signature pattern is unique enough, this would substantially reduce the number of prediction misses.

However, extracting patterns on larger time-scales can be a computationally intensive task. We believe that some powerful node in the wired network that does not have any energy constraints will do this processing. The architecture for the system is illustrated in figure 16. However, extracting patterns on larger time-scales can be a computationally intensive task. We believe that some powerful node in the wired network that does not have any energy constraints will do this processing. The architecture for the system is illustrated in figure 16.

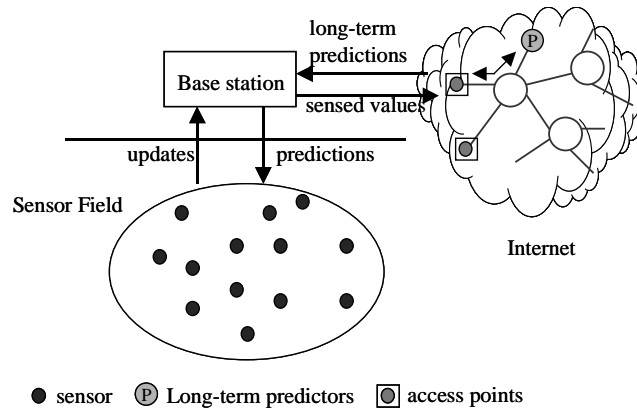


Fig. 16 Architecture of the system with wired nodes (longer-term predictors) extracting patterns at longer time-scales

We propose two approaches for extracting patterns at the medium and large scales. One makes use of data mining techniques and the other makes use of compression theory. In this section, we explain each of these approaches.

5.1 Data Mining approach

Consider the sensors in a cell. Assume that every time the reading sensed at a sensor changes, it transmits its new reading to the base station of the cell. The base station, in turn, relays these to some long-term predictor via an access point. Based on the time series data collected at a long-term predictor, one can form a database, where each record of the database contains the readings of all the sensors in the cell at a certain instant of time. One can apply one of the existing data mining algorithms [FD92, FD93] for extracting functional dependencies satisfied by the database. This extraction process is constrained by the fact that the left-hand side of a functional dependency can only comprise of readings of sensors that are in close proximity to the sensor whose reading appears on the right-hand side. Mathematically, let $S(x,y)$ represents a sensor located at location $\langle x,y \rangle$ and let $S(x,y,t)$ represents its reading at time t . Then data mining techniques may be applied to extract functional dependencies of the form:

$$\{S(xk,yk,tj), S(xk+1, yk+1, tj), \dots, S(xw, yw, tj)\} \Rightarrow S(xm, ym, tj) \quad \dots (1)$$

where, the sensor $S(xm, ym)$ is in radio range of the sensors $S(xk,yk), S(xk+1, yk+1), \dots, S(xw, yw)$.

We say that the functional dependency in equation (1) holds iff:

$$S(xk,yk,t) = S(xk, yk, t'), S(xk+1, yk+1, tj) = S(xk+1, yk+1, t'), \dots, \text{ and } S(xw, yw, t) = S(xw,yw,t') \Rightarrow \\ |S(xm, ym, t) - S(xm, ym, t')| < e,$$

for any $t, t' \in \{t1, \dots, tn\}$, where e denotes the acceptable margin of error.

For every functional dependency extracted, one may apply regression techniques (linear or non-linear) to obtain the function P such that: $|S(xm, ym, t) - P(S(xk,yk,t), S(xk+1, yk+1, t), \dots, S(xw, yw, t))| < e, t= t1, \dots, tn \quad \dots (2)$

The computed function forms the spatial prediction model for the sensor $S(xm, ym)$. It may be relayed to the sensor $S(xm,ym)$ via the base station. On receiving the prediction model, the behavior of the sensor will change. It will now start listening for transmission from sensors $S(xk, yk), S(xk+1, yk+1), \dots, S(xw, yw)$ and compute the function $P()$. It will not transmit its reading to the base station if actual sensed reading differs from value returned by function $P()$ by less than the acceptable margin of error.

In equation (2), the readings of sensors $S(xk, yk), S(xk+1, yk+1), \dots, S(xw, yw)$ are being used as arguments in the spatial prediction model for sensor $S(xm, ym)$. It is possible that the reading of one of these sensors $S(xi, yi), i \in \{k, \dots, w\}$ also follows some model M and can be predicted. However, if one were to predict its reading, the sensor $S(xi, yi)$ will not transmit its reading at time instants when it conforms with model M . At these time instants, sensor $S(xm, ym)$ will not have access to the reading of $S(xi, yi)$ and therefore, will not be able to compute its prediction function $P()$. One way around would be to have the base station send its estimate of the reading of $S(xi, yi)$ based on the model M . However, these estimates may differ from actual reading by as much as e , the acceptable margin of error. If the function $P()$ is "tolerant" to estimated arguments (i.e., if it can predict within the acceptable margin of error even when the value of one or many of its arguments are only estimates), this approach may work correctly. However, we choose not to take such an approach since this would restrict function $P()$ only to certain function families.

It is easy to verify that the above functional dependencies are governed by the same set of axioms as regular functional dependencies in relational databases.

Let,

$U = \{S(xi, yi, t): \text{set of all readings of all the sensors in the cell at time instant } t\}$, and

$DET = \{V1, \dots, Vn\}$, where $Vi = S(xi, yi, t)$ and $DET \subseteq U$.

Now let us assume that by data mining the database of the sensor field we have determined that the sensor field satisfies the set F of functional dependencies (similar to equation 1). Then, one may view DET a **key** of the sensor field S *if and only if* the functional dependency, $DET \Rightarrow U$, is a consequence of F , and no subset of DET has this property.

Note that existence of a functional dependency f does not require that we actually know the function that will allow us to compute the readings of the sensor $S(x, y, t)$ from the readings of other sensors which occur in the body f ; the functional dependency only states that such a function exists.

The interpretation of DET is as follows: If we always know precisely the readings of sensors in DET then we can determine the readings of all other sensors. Thus DET is the minimal set of sensors that are necessary in order to predict the readings of all the sensors of the sensor field. There may be many such minimal sets, just as there can be many keys in relational databases. One criterion of choosing a specific set could be its cardinality (smallest). We may also decide to group sensors in the sensor field into classes (categories) and have each category have its own "local key".

One may relax the constraint that the arguments to the spatial prediction model for a sensor $S(x_m, y_m)$ needs to be readings of those sensors that $S(x_m, y_m)$ can listen to. We may do this by requiring the base station to broadcast the readings of sensors in DET to all the sensors in the cell.

Temporal and spatio-temporal patterns may be extracted in a similar way using data mining techniques. For extracting spatio-temporal patterns, the long-term predictor builds a database of sensor readings, where each tuple has the following form:

$\langle S(x_1, y_1, t_1), S(x_1, y_1, t_1-1), \dots, S(x_1, y_1, t_1-n), S(x_2, y_2, t_1), S(x_2, y_2, t_1-1), \dots, S(x_2, y_2, t_1-n), \dots, S(x_w, y_w, t_1), S(x_w, y_w, t_1-1), \dots, S(x_w, y_w, t_1-n) \rangle$

Thus, each tuple stores the last n readings of all the sensors in the cell. Using the techniques described above one may generate spatio-temporal prediction models.

5.2 Compression theory approach

The visualization of this time-series data will be a sequence of image frames. Consider dividing each frame into macro-blocks. Since our interest is in extracting temporal patterns, we focus on each macro-block one at a time. Consider one particular macro-block in a frame. We look at the same macro-block in every frame. We examine how the readings of sensors in this macro-block are changing from one frame to another. Based on the readings of sensors in a macro-block, we build a dictionary of pattern of sensor readings in a macro-block, and assign each unique macro-block a unique codeword. We start with the macro-block in the first frame in the sequence. We add it to the dictionary and assign it a unique codeword. For each subsequent frame in the sequence, we see if the macro-block under consideration is already present in the dictionary. A macro-block is said to match one in the dictionary if the sum of the absolute value of difference in their corresponding sensor readings is less than a certain threshold. If a macro-block doesn't match any entry in the dictionary, we add it to the dictionary and assign it a unique codeword. As we are building this dictionary, we also generate a sequence of corresponding code words. This sequence has a one-to-one correspondence with the sequence of frames – the i th codeword in the sequence corresponds to the macro-block under consideration in the i th frame. By the time we get to the last frame in the frame sequence, we would have a corresponding sequence of code words. If we treat these code words as alphabets of some language L , the entire sequence of code words will form a “sentence” of L . The problem of extracting temporal pattern then boils down to extracting a frequently occurring pattern of alphabets in this sentence. This may be accomplished by using one of the standard dictionary-based compression algorithms [LZ, LZW]. The pseudo-code for this approach appears in figure 17-19

```

/* This function takes as input a sequence of frames, a dictionary mapping unique macro-blocks
 * to unique codewords, and a macro_block number. We assume that all the macro-blocks of a frame
 * are numbered according to their position in the frame. This function looks at the same macro-
 * block (identified by macroBlockNum) in each frame in the sequence. Based on this sequence, it
 * generates a corresponding sequence of codewords. For each macro-block, the dictionary is
 * consulted. If an entry exist for the macro-block, the corresponding codeword is put in the
 * codeword-sequence. If no entry exist, a new entry for the macro-block is inserted in the
 * dictionary and a unique codeword assigned to it.
 */
function generateCodewords(Vector frameSeq, int macroBlockNum, Dictionary MBDict)
    returns code-word sequence
    foreach currentFrame in frameSeq do {
        /* extract the macro-block # macroBlockNum from the current frame */
        currentMacroBlock = getMacroBlock(currentFrame, macro_blockNum);

        /* search for the macro-block in the dictionary */
        index = search(MBDict, currentMacroBlock);

        if (index == FAILURE) {
            /* LOOKUP FAILURE:
             * insert a new entry for this macro-block. The function returns
             * the codeword assigned to this */
            codeWord = addNew(MBDict, currentMacroBlock);
        }
        else {
            /* LOOKUP SUCCESS:
             * get the corresponding codeword. 'index' contains the index of
             * the matching macro-block in the dictionary */
            codeword = getCodeWord(MBDict, index);
        }

        /* add the codeword corresponding to the macro-block to the
         * sequence of codewords being generated */
        addNew(codeWordSeq, codeword);
    }
    return codeWordSeq;

```

Fig. 17 Pseudo-code for generating a sequence of code words from a sequence of macro-blocks

```

/* this function searches the dictionary for an entry whose macro-block
 * matches with the macro_block
 */
function search(Dictionary MBDict, Image macro_block) returns result

    foreach dictEntry in MBDict do {
        if (match(dictEntry.MacroBlock, macro_block)) {
            /* found a match. Return the index of the entry */
            return dictEntry.index;
        }
    }

    /* no match for macro_block found in the dictionary */
    return FAILURE;

```

Fig. 18 Pseudo-code for searching the dictionary for an entry with matching macro-block

```

/* This function takes two macro-blocks and checks if they match. Two
 * macro-blocks are said to match if the sum of the absolute difference in
 * their corresponding pixel values is less than MatchErrorThreshold.
 * The macro-blocks are assumed to be of size blockWidth x blockWidth
 */
function match(Image macroBlock1, Image macroBlock2) returns result
    int row1, col1, row2, col2;
    int error;

    row1=0; col1=0; row2=0; col2=0;
    error = 0;

    while (row1 < blockWidth) {
        while (col1 < blockWidth) {
            error = error + abs(macroBlock1[row1][col1] - macroBlock2[row2][col2]);
            col1++; col2++;
        }
        row1++; row2++;
    }
    if (error < MatchingErrorThreshold)
        return SUCCESS;
    else
        return FAILURE;

```

Fig. 19 Pseudo-code for determining if two macro-blocks match

6. Quality of Monitoring

The constraints on the quality of the answer to a monitoring query required by a user govern the amount of energy savings that may be achieved. The quality of answer is controlled by two parameters:

1. *Error tolerance*: Controls the amount by which the received answer may deviate from its actual reading
2. *Delay tolerance*: Controls the interval from the time a sensor's reading changes to the time it is received at the user.

The more the error tolerance, the less precise the prediction-models need to be, and hence, more is the potential for saving energy. In the extreme case when the error tolerance is infinite, the sensors do not have to transmit anything. Also, the more the delay tolerance, the more time is available to generate prediction-models. This increases the potential for generating close-to-perfect prediction models at the base station, which, in turn, increases the potential for saving energy.

7. Related Work

Researchers in sensor networks have looked at various ways of saving energy of sensors. In particular, [Spin, Leach] have proposed exploiting redundancy in sensed readings of sensors that are in close proximity. However, the approaches proposed in [Spin] rely on sensors to figure out among themselves the redundant part in their sensed information and suppress it from transmission.

We see three problems in this approach. Firstly, it will take N transmissions for a group of N sensors to figure out all the redundant information, which is costly and may outweigh the advantage of removing redundancy. Secondly, the approach works well when the sensed readings are simple (like temperature, pressure, etc.). In cases when the sensed readings are complex (like image from a camera,

assuming a network of cameras[MIT]), it is not clear how one may determine the redundant part given the limited computational power at the sensors. Thirdly, at the very best, this approach only minimizes the redundancy in information being sent. It may work well in sensor networks like network of temperature sensors where the sensed reading is not expected to vary much in a small spatial region. However, the assumption of redundancy is not true in general. Consider a network of motion sensors. Assume that the reading at a sensor (for e.g., magnetometer sensor) is a function of its distance from a moving object. In such a case, the readings of sensors in close proximity will not be that same but will be correlated. We distinguish between redundancy and correlation. We say that readings at two sensors have *redundancy* if there is a part that is common to both. For example, if two temperature sensors in close proximity sense 65F as the temperature, we say that there is redundancy in their information. Although the word “correlation” has many connotations, for the purpose of this paper, we say that the readings at two sensors are *correlated* if given the reading of one sensor one may derive the reading at another sensor with reasonable accuracy. In other words, whenever the reading at one sensor is a function of the readings at the other sensor we say that two sensors have correlated readings. One may think of redundancy as a special case of correlation where one may derive reading at a sensor by employing an identity function (i.e., $f(x) = x$). Returning back to our example of sensor network of motion sensors, if the reading at a motion sensor is a deterministic function of its distance from the closest moving object, then readings at two sensors A & B are correlated (not redundant). The reading at A may be derived given the reading at B, the location of two sensors, and the location of the object. Lastly, this approach only talk about eliminating redundancy in readings sensed in the same time frame. In our proposed scheme, we specify mechanisms for identifying correlation in readings across time frames and try to eliminate the need for transmitting these if they can be derived from older readings. We expect this to translate into significant savings in energy.

In [Leach], authors deal with the problem of how to self-organize sensors that are thrown together, so that a meaningful sensing task can be accomplished. The paper prescribes compressing information at the nodes acting as cluster-head, before sending to a central station.

However, the paper does not specify any mechanism for performing this compression operation. Also, this approach does not prevent sensors from sending redundant information.

In [Estrin] the emphasis is on being able to “route” data sensed at sensors to where the interest is. It does not deal with cutting down the information being relayed.

8. Future work

We plan to explore the following future directions:

1. *Learning in PREMON:*

We plan to explore how one can incorporate learning patterns of movement based on history, and use them to guide the prediction-generation logic.

2. *Duty Cycles:*

A possible way of reducing the energy usage in a sensor network is by reducing the duty cycle of sensors that are not seeing frequent activities. Thus, we would like to partition a sensor network into regions based on the level of activity they see. For sensors in these different types of areas, the duty cycle requirement is clearly different. We plan to look at protocols for dynamically partitioning a sensor field into these different areas and adjust the duty cycles of sensors appropriately to achieve near-optimal energy usage.

3. *Extracting patterns at larger time-scales:*

We intend to evaluate the effectiveness of the mechanisms proposed in section 5 for extracting patterns at larger time-scales and using them to make more accurate predictions.

4. *Predicting aggregates:*

In this paper, we have looked at predicting values of sensors. However, in many instances one is interested in the aggregate value over a small spatial region (e.g., what is the average temperature in DataMan Lab). The challenge is to apply the PREMON paradigm to predict the values of aggregates.

5. *Simulations:*

We are in the process of performing simulations to analyze the performance of the proposed algorithms over large sensor networks.

9. Conclusions

In this paper, we have presented energy-efficient mechanisms for performing monitoring operations in sensor networks. For energy-efficient monitoring, we have proposed a new paradigm called Prediction-based monitoring (PREMON). It exploits the spatial, temporal, and spatio-temporal correlation likely to be present in readings of spatially proximate sensors. We have shown that this paradigm is analogous in many ways to the one followed by MPEG encoder. This close analogy allows us to adapt the theory and algorithms of MPEG for use in sensor networks. We have proposed an algorithm for generating prediction models based on the block-matching algorithm in MPEG. We have implemented the proposed algorithm on a testbed of sensor-motes. The results of our experiments have shown that utilizing the PREMON paradigm helps in significantly cutting down the number of transmissions from the sensors, thereby achieving substantial energy savings.

Acknowledgements

We would like to acknowledge the help of Vihari Komaragiri, Sudeept Bhatnagar, and Anirudha Bohra, for giving insightful comments to the earlier drafts of this paper.

References:

- [DataSpace] T. Imielinski and S. Goel, "DataSpace -- querying and monitoring deeply networked collections in physical space", IEEE Personal Communication Magazine, Special issue on "Networking the physical world", October 2000.
- [Cougar] Ph. Bonnet, J. Gehrke, and P. Seshadri, "Querying the physical world", IEEE Personal Communication Magazine, Special issue on "Networking the physical world", October 2000.
- [Leach] W. Heinzelman, A. Chandrasekaran, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor network", Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS '00), January 2000.
- [Spin] J. Kulik, W. Rabiner, and H. Balakrishnan, "Adaptive protocols for information dissemination in wireless sensor networks", Proceedings of 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99), Seattle, WA, 1999
- [Lakme] L. Subramaniam and R.H. Katz, "An architecture for building self-configurable systems", IEEE/ACM workshop on Mobile Ad Hoc Networking & Computing (MobiHOC '00), Boston, August 2000.
- [Pottie00] G.J. Pottie and W.J. Kaiser, "Wireless Integrated Network Sensors", Communications of ACM, Vol. 43, No. 5, May 2000.
- [Pottie99] K. Sohrabi, J. Gao, V. Ailawadhi, and G.J. Pottie, "A self-organizing wireless sensor network", Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing, September 1999.
- [Estrin] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A scalable & robust communication paradigm for sensor networks", Proceedings of 6th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '00), Boston, August 2000.
- [Wolfson] O. Wolfson, S. Jajodia, and Y. Huang, "An Adaptive Data Replication Algorithm", ACM Transactions on Database systems, Vol. 22, No. 2, June 1997.
- [MPEG] J. Watkinson, "MPEG-2", Local Press, 1999.
- [MIT] "Forest of Sensors Project", AI Lab, MIT, <http://www.ai.mit.edu/projects/vsam/>
- [FPL] J. Rabaey, "Silicon Platforms for the Next Generation Wireless Systems – What Role does Reconfigurable Hardware Play?", Proceedings FPL 2000, Austria, August 2000.
- [TinyOS] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, K. Pister, "System Architecture Directions for Networked Sensors", Submitted for publication.
(<http://tinys.millennium.berkeley.edu/papers/tos.pdf>)
- [Tools] Tools for programming Rene Motes.
<http://tinys.millennium.berkeley.edu/release/tos-latest.tar.gz>
- [Motes] TinyOS: An operating system for Networked Sensors, (<http://tinys.millennium.berkeley.edu/>)
- [Mani] Mani Srivastava, "Energy Efficient Wireless Systems", in DIMACS Summer School on Foundations of Wireless Networks and Applications, August 7-18, 2000.
- [Datamining] R. Agrawal, T. Imielinski, A. Swami, "Mining Association Rules between Sets of Items in Large Databases", Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data.
- [LZW] Khalid Sayood, Introduction to Data Compression, Second Edition, Morgan Kaufmann Publishers.
- [Interpolation] R.D. Harding, D.A. Quinney, Simple Introduction to Numerical Analysis : Interpolation and Approximation, Adam Hilger.
- [FD93] Iztok Savnik and Peter A. Flach. Bottom-up induction of functional dependencies from relations. In Gregory Piatetsky-Shapiro, editor, *Proceedings of the AAAI Knowledge Discovery from Databases Workshop (KDD)*, 1993.
- [FD92] Mannila, H., Raiha, K.J. et al.: *Discovering Functional and Inclusion Dependencies in Relational Databases*. International Journal of Intelligent Systems, Vol 7, 1992.

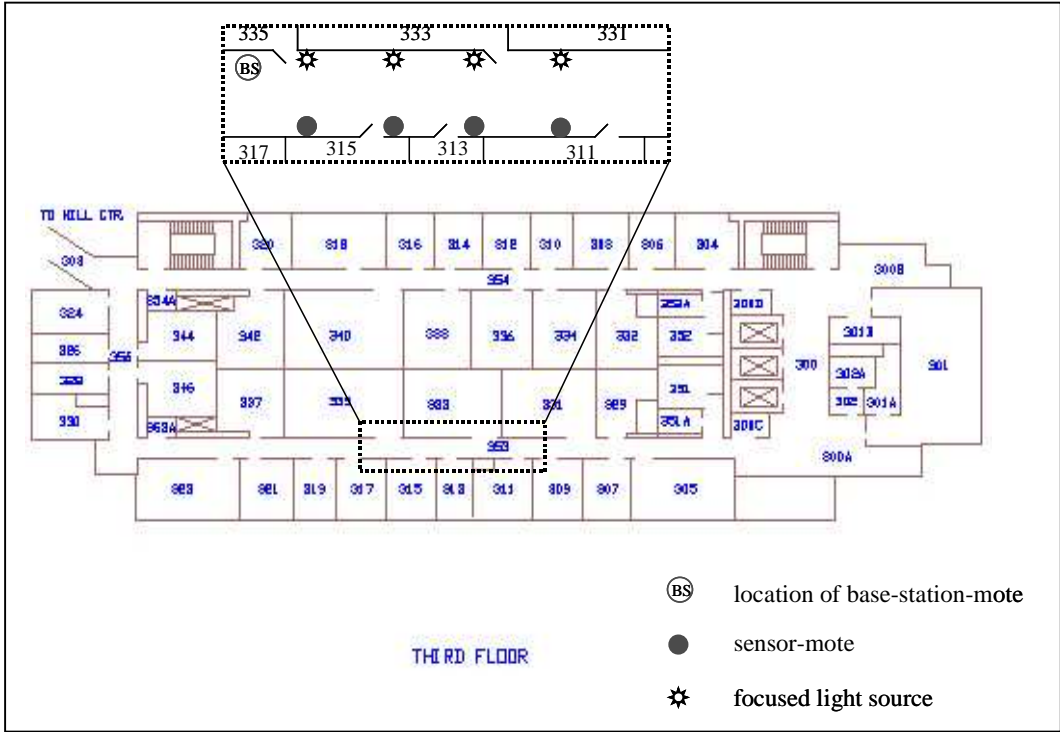


Fig. 20 Floor plan indicating placement of motion-sensors