UCLA
Technical Reports

Title
Augmenting Film and Video Footage with Sensor Data

Permalink

Authors
Norman Makoto Su
Heemin Park
Eric Bostrom
et al.

Publication Date
2003

Augmenting Film and Video Footage with Sensor Data

Norman Makoto Su[*], Heemin Park[†], Eric Bostrom[*], Jeff Burke[‡], Mani B. Srivastava[†], Deborah Estrin[*]

[*]Department of Computer Science,
[†]Department of Electrical Engineering,
[‡]School of Theater, Film and Television
University of California, Los Angeles
{normsu,ebostrom,destrin}@cs.ucla.edu, {hmpark,mbs}@ee.ucla.edu, jburke@hypermedia.ucla.edu

Abstract— With the advent of tiny networked devices, Mark Weiser's vision of a world embedded with invisible computers is coming to age. Due to their small size and relative ease of deployment, sensor networks have been utilized by zoologists, seismologists and military personnel. In this paper, we investigate the novel application of sensor networks to the film industry. In particular, we are interested in *augmenting* film and video footage with sensor data. Unobtrusive sensors are deployed on a film set or in a television studio and on performers. During a filming of a scene, sensor data such as light intensity, color temperature and location are collected and synchronized with each film or video frame. Later, editors, graphics artists and programmers can view this data in synchronization with film and video playback. For example, such data can help define a new level of seamless integration between computer graphics and real world photography. A real-time version of our system would allow sensor data trigger camera movement and cue special effects. In this paper, we discuss the design and implementation of the first part of our embedded film set environment, the augmented recording system. Augmented recording is a foundational component for the UCLA Hypermedia Studio's research into the use of sensor networks in film and video production. In addition, we have evaluated our system in a television studio.

I. Introduction

It remains to be seen what the real-world wireless sensor network applications (WSNA) will be. Much of the ongoing motivation for WSNs are for monitoring natural phenomenon such as seismic propagation in buildings [1], animal habitats [2] or red-tide concentrations [3]. In such applications, WSNs are an ideal fit because they are minimally intrusive to their surrounding environment, as well as being disposable and easily deployed. As of late, however, wireless sensor networks have also been viewed as an enabler of ubiquitous environments. One can imagine sensors being deployed in groceries, schools or shopping malls. The requirements for a ubiquitous environment parallel those of the aforementioned monitoring applications. Indeed, these sensors will, as Mark Weiser prophetically stated, "weave themselves into the fabric of everyday life until they are indistinguishable from it" [4]. Sensor networks promise to provide a way to actuate on and record the dynamic events that are part of our everyday lives.

We believe WSNs will also prove immensely useful for *filmmaking* and *media production*. In traditional production, these networks will enable cost savings through efficient and detailed data gathering as well as offer increased expressive capabilities for the creative team through decision support and control of existing automated equipment. Our research seeks to explore the potential application of WSN deployments on film and television sets to support the daily work of production. Using WSNs, we can gain the following:

Seamless Integration: Any devices that are deployed on the set or on the performers ought to provide as little interference as possible to actors, directors and technicians.

Mobility: In addition, filming is often done in various harsh environments where deploying large equipment can be difficult. Thus, robust, lightweight and easily configured equipment is ideal.

Increased Expressiveness: Special effects and complex shots are heavily scripted and precisely choreographed to achieve the desired results. In the future, data from a WSN could be used to adjust these processes on-the-fly according to creative rules defined by the director or cinematographer.

In this paper, we describe the first step in realizing a ubiquitous film set where sensors are deployed on a set to monitor and actuate based on lighting conditions, acoustic events, director cues and performer gestures. We have a built a system infrastructure to *augment* film footage with sensor data. Augmented footage is a sensor data record synchronized with the film camera's 'privileged' visual record[1]. At this early stage, the system supports display of sensor data corresponding to each frame when video proxies from a shot are reviewed.

While recording a scene, sensors will collect vital data that is *synchronized* with each film or video frame. In this post-facto synchronization scheme, one can later view the recorded film and observe its corresponding sensor data. Such synchronized data can be utilized, for example, by digital artists to add CG to scenes according to the environment. In the future (as yet unimplemented), a real-time version of this augmented recording system could actuate camera movements and special effect cues. From hereon, "augmented recording system" shall be abbreviated as ARS.

This paper is organized as follows: Section II describes related work, section III describes the requirements for ARS, and sections IV–VIII describes our system architecture and

---

[1]The film camera's viewpoint is 'privileged' because for film production, all that matters at the end of a shoot is what has been captured by that camera.

modules. Section IX has our TV studio evaluations which is followed by Section X, our conclusion and future works.

II. RELATED WORK

Previous work has been done at UCLA's Hypermedia Studio in theatre production. In a production of Eugene Ionesco's *Macbett* [5], a complete system for controlling theatrical lighting and sound based on performer position and movement was developed. For example, tracking devices were embedded into the witch characters' staves. At the performer's own initiative, she or he could conjure up lightning strikes or whirling winds by throwing the staff in the air or swirling it, respectively.

At MIT's Media Lab, Joesph Paradiso's Responsive Environments Group's work on interactive environments that, for example, create music based on user gestures, foot positions and jumping is an example of an environment embedded with sensors [6]. The focus here is more on building sensors which can dynamically change environments, rather than on recording changes in environments. However, dynamically changing camera movement, etc. is a future goal of our embedded film set.

The augmented footage concept itself is similar in spirit to the Smart Kindergarten [7] and Classroom 2000 [8] projects in that it collects a variety of time-stamped data from human inhabited environments. However, ARS requires a much more stringent synchronization requirement.

There already exists a large body of research based on image processing. Indeed, one might wonder whether cameras may make the use of sensors a moot point. With image processing one can infer location, lighting, etc. Sensors, however, can provide a granularity of accuracy hard to achieve purely with image processing techniques. Moreover, cameras require line of sight with the scene being filmed. Should we wish to monitor multiple areas of interest in a film set, numerous cameras are required. On the other hand, sensors have a small form-factor, making them easy (and cheap) to deploy. Finally, there are many phenomenon that can not be easily inferred from image processing. Factors such as temperature or wind speed, which may prove to be interesting to installation art are easily measured with sensor networks. We do, however, believe vision technologies will provide a valuable role, *with* WSNs, in the future embedded film or TV set. For example, image processing can provide head pose/gaze direction measurements [9], motion recognition [10], and tracking [11].

III. SYSTEM REQUIREMENTS

To give a clearer idea of the requirements to consider when building an ARS, let us give an example scenario for video production:

1) Sensors are placed in a film set, as well as on performers.
2) A video camera starts recording a scene.
3) While recording:
   a) A *timecode generator* signal encodes frame number data onto an audio track of the video.

   b) Sensors collect frame stamped data in synch with the timecode generator.
4) The crew can view the footage with sensor readings mapped to frames. In other words, while playing back footage, the sensor data collected at each frame can be viewed.

Firstly, it is apparent that *time synchronization* with both the camera and sensors is necessary. Secondly, we require sensor data for each video frame. This requires a very high sampling rate demanded by TV standards. Thirdly, we wish to provide an easily extendible system encompassing a wide range of sensor devices. The final requirement is an interface to the system that enables users to control and manage the augmented recording system.

Preliminary mining of sensor data in augmented footage would support two access methods: 1) random access of film or video frames to provide additional scene information during film footage playback that is not apparent from footage alone and 2) sensor data-based queries of augmented footage about high-level events and lower-level environmental conditions.

Currently, only item 1 is implemented via the GUI client.

IV. SYSTEM ARCHITECTURE

In this ARS, we separate software into four distinct modules:

Sensor Node Neighborhood: Each sensor node neighborhood is an autonomous cluster with one basestation and several neighboring sensor.

Serialport Server: The serialport server directly issues commands to the various basestation sensors, interpolates missing data and will time synchronize sensors.

Sylph Server Middleware: Sylph [12] is a middleware solution for easily allowing higher level queries from Jini clients to manage sensors. This server translates messages between sensors and clients.

Jini Client: This is a Java client adhering to Jini standards. It provides an interface for the user to control ARS, store information on a database and to later inspect frame synchronized sensor data.

Figure 1 shows an overview of the key components and their connections. The lightning bolts indicate an Internet connection. The clouds represent a wireless neighborhood. Here, we describe the sequence of events leading to successful storage of frame-synched data.

1) A Jini client connects to a well-known lookup service, locates the Sylph server managing ARS and issues a specially formatted query string.
2) The Sylph server will translate the query, issue (if any) intermediary replies back to the Jini client and proceed to forward the query to the serial port server.
3) The serial port server will multiplex the message to the appropriate base station sensors.
4) The base stations will inform its surrounding sensors to begin transferring data synched with a global timecode generator.
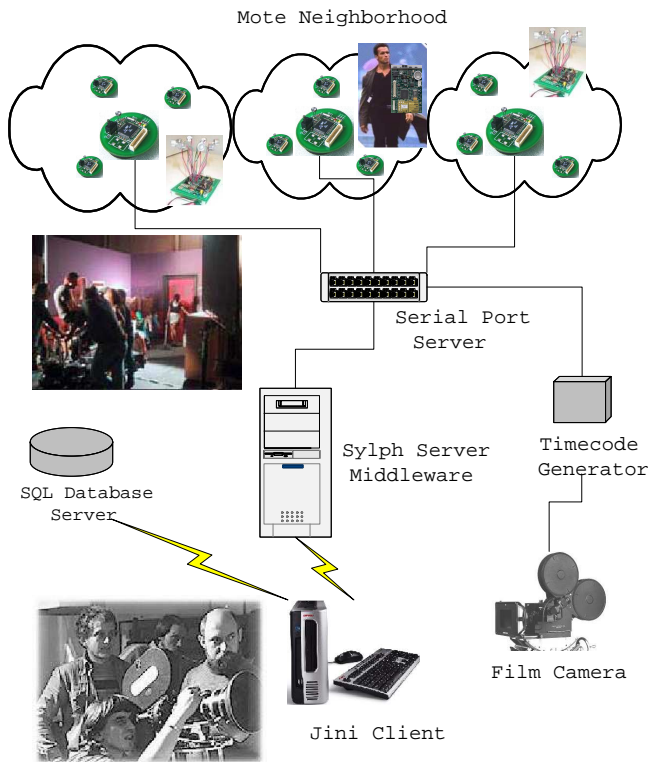
Fig. 1. System Architecture

Fig. 2. Palos Task Flow

5) As data streams back, the base stations will collect data and send it to the serial port server.
6) The serial port server will perform error checking and interpolation on the data, format it and send it off to the Sylph server.
7) The Sylph server will announce to the Jini client that data is available.
8) The Jini client will store the information onto a (possibly remote) database for later perusal.

In the following sections we will go into greater detail with the software modules.

V. SENSOR NODE NEIGHBORHOOD

The sensors are Crossbow mica motes [13] with the standard sensor boards equipped with a thermistor, light sensor, microphone and accelerometer. In addition, we have developed a simple MAC (medium access control) protocol for high speed transfers. We chose motes because they are lightweight, easy to prototype on, have wireless radios and there exist many sensor boards for it. In addition, it models the future vision of what many researchers believe such nodes will be—power constrained and limited in processing power. Another alternative include HP's iPAQs. However, they lack the portability and range of sensor boards that motes have.

A. Software Platform

We are running an operating system developed by the NESL group, called PALOS (**P**seudo-Re**AL**time **O**perating **S**ystem)
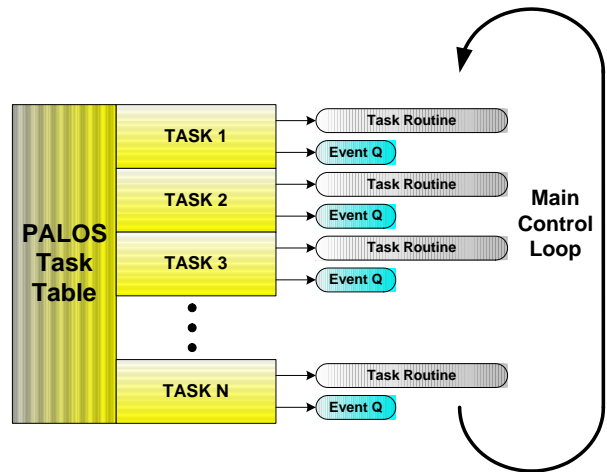
[14]. Currently, the bulk of today's mote programming is done with UC Berkeley's excellent TinyOS. However, for quick prototyping, we found that PALOS provides a distinct advantage. When programming with PALOS, one defines well-defined *tasks*. Each task has its own *event queue*. During the program execution, a global scheduler, by priority, runs through each event queue, dequeuing and passing the events to their respective tasks. By event, we mean a data structure which stores an identifier as well as some arbitrary data. Figure 2 shows the scheduler's cycle. Each task can also *send* messages (i.e. events) to other tasks by pushing an event into the appropriate task's queue. Programming with PALOS requires minimal use of macros and reduces the likelihood of deadlock or race conditions by enforcing a priority based schedule. However, the tradeoff is a degradation in performance. Finally, code from TinyOS is easily ported to PALOS—several libraries are based on TinyOS's.

In ARS, for example, we define the following tasks for a basestation mote:

UART Task: The UART Task waits for a fixed size data packet from the serial server and decides on an appropriate action. For example, when it receives a STOP_SENDING packet, it will tell the BMAC Task to send a packet wirelessly telling the neighboring motes to stop sending.

BMAC Task: The BMAC Task receives messages from the UART Task to send packets over the wireless link. It also sends messages to the Data Collection Task when it receives a wireless packet. So, for example, upon receiving a sensor reading from a neighboring mote, the BMAC Task will notify the Data Collection Task of this incoming packet.

Data Collection Task: The Data Collection Task will receive messages from the BMAC Task and forward the packet onward to the Serial Port Server. For example, it will receive an event with the data packet from the BMAC Task and then forward the data packet to the Serial Port Server.

The following code snippet shows us registering the UART

and BMAC task:

```
1  UART0_Event uart0EventQ[UART_Q_SIZE];
2  BMAC_Event bmacEventQ[BMAC_Q_SIZE];
3
4  void globalInitTask() {
5      BMAC_id = PalosTask_register(StopWatch_init,
6                                   BMAC_task,
7                                   BMAC_XCNTR,
8                                   BMAC_RCNTR,
9                                   sizeof(StopWatch_Event),
10                                  BMAC_Q_SIZE,
11                                  (void *)bmacEventQ);
12     UART0_id = PalosTask_enroll(UART0_task,
13                                  UART0_XCNTR,
14                                  UART0_RCNTR,
15                                  sizeof(UART0_Event),
16                                  UART_Q_SIZE,
17                                  (void *)uart0EventQ);
18 }
```

Listing 1.   Registering Tasks in PALOS

Lines 1,2 are the queues for their respective tasks. As data streams in from the UART or BMAC, the appropriate task routine pushes an event onto these event queue. The `globalInitTask()` function registers the two tasks (lines 5–17), informing PALOS to pass the popped event to either `BMAC_task` or `UART0_task`.

When the main control loop, for example, reaches the UART0's event queue, it will pop the event in a FIFO fashion and pass it to a handler, `UART0_task`:

```
1  CHAR UART0_task(void *ev) {
2      UART0_Event *event = (UART0_Event *)ev;
3      UCHAR *recvData = event->data;
4      switch(recvData[0]) {
5          case OPCODE_NUM_NEIGHBORS:
6          ...
7      return PALOS_TASK_DONE;
8  }
```

Listing 2.   An Event Handler in PALOS

The `ev` (line 1) parameter is the event popped off the queue. As evidenced by the code, a UART0 event has a `recvData` member. We then do something with that data, and finally return `PALOS_TASK_DONE` (line 7). `PALOS_TASK_DONE` tells the scheduler to pop the event off the queue permanently. Other variants include `PALOS_TASK_NEXT`, pop the event off the queue and immediately handle the next event from the *same* queue (i.e. for this example, we would immediately process the next UART event waiting in the queue) and `PALOS_TASK_KEEP`, repush the event onto the queue and process the next event queue normally.

The MAC protocol is built on top of a physical layer which was ported from Wei Ye's communication stack [15]. While Berkeley's stack allows a longer transmission range, Ye has shown that his stack achieves high reliability (near 100%) and effective collision avoidance mechanisms with greater than 4 simultaneous transmitters. Our BMAC protocol (**B**asic **MAC**) is a very simple, barebones contention based protocol relying on carrier sense and ACKs—it does not use CTS and RTS. Our motivation behind this was to develop a reliable, quick protocol for sending short packets.


Fig. 3.   Packet Loss Test Results

B. Clustering

We chose to implement several "clusters" of mote neighborhoods. The reasoning behind not simply having one central basestation mote is simple—limited bandwidth. Let us assume we need approximately 20 motes deployed in the film set and packets of 30 bytes are transferred every frame. Motes have an approximate maximum transfer rate of 10kbits per sec. This would require transferring 30 bytes per 1.7 ms which is far greater than the 10kbits per sec. In ARS, we define a cluster as a basestation and its neighboring motes. Our system currently only utilizes a one-hop topology.

In general, the maximum number of neighbors allowable for each basestation depends on the environment, as well as the bandwidth. We conducted empirical tests in our lab to determine what sort of range is applicable. With a packet size of 30 bytes or, equivalently, 13 frames of data, we placed motes around a basestation (at approximately 4 inches away) and had each neighbor mote send packets at a rate of approximately $\frac{1}{29.97} \times 13 \approx 433$ms. The test stopped once 2000 packets were successfully from any neighboring mote. Results are shown in Figure 3. In ARS, we current would like to limit our data loss rate to 10%. Given the results, setting the neighborhood to 5 or 6 motes seems reasonable.

C. Basestation

The basestation motes multiplex on messages from the serial port server. Via a broadcast, it tells the neighboring motes to start/stop sending data. The basestation can also assign *roles* to neighboring motes (e.g. "record light" or "record sound"). The basestation will listen for data from the neighboring motes and forward them to the serial port server.

1) POT Calibration Phase: Before the basestation begins requesting data from its neighbors, it performs a potentiometer (radio power) calibration phase. The POT calibration phase serves three purposes: discovering neighboring motes' ids, assigning each neighbor mote a basestation and reducing the total area of the cluster. The later allows us to minimize interference between multiple clusters.

Note that a higher potentiometer setting implies a shorter range. For sake of clarity, we shall only refer to power settings from now on. The algorithm's input is X, the number of

```
1  /* Need succ/trials to go on valid list */
2  succ <- 3
3  trials <- 4
4  neigh_found <- 0
5  init_power <- 0
6  /* constant factor to increment power */
7  power_inc <- 5
8  neigh <- Make-Zero-Hash(X)
9  valid <- Make-False-Hash(X)
10
11 BaseStation-Algo(X)
12     Adjust-Power(0,X,init_power)
13
14 Adjust-Power(t,X,power)
15     Broadcast-Change-Power(power,valid)
16     Start-Timer()
17     while not Timer-Expired()
18         do msg <- BMAC-Read()
19             if neigh[msg.src] = 0
20                 then neigh_found <- neigh_found + 1
21             neigh[msg.src] <- neigh[msg.src] + 1
22             if neigh_found = X
23                 then return
24             else if t = trials
25                 for k in neigh.keys
26                     do if neigh[k] >= succ
27                         then valid[k] = true
28                 neigh <- Make-Zero-Hash(X)
29                 Adjust-Pot(0,X,power+power_inc)
30             else Adjust-Pot(t+1,X,power)
```

Listing 3.   Basestation POT Calibration Pseudocode

```
1  Neighbor-Algo(my_id)
2    loop
3       do msg <- BMAC-Read()
4          for i in msg.valid
5             do if i = my_id
6                then break
7          Set-Power(msg.power)
8          Unicast-Power-Set-ACK(msg.src)
```

Listing 4.   Neighbor POT Calibration Pseudocode

neighbors, we want a basestation to find. The basestation broadcasts a message (at full power) requesting all neighboring motes to set their power to a minimal value and to reply back. The basestation adds those neighbors whose reply successfully reached it to its `valid` list. Then, the basestation rebroadcasts the same message, except this time, the power value is increased by some constant factor and the broadcast contains the `valid` list. Those neighbors who are already on the `valid` list ignore the message. The other neighbors will increment their power and reply back. This process is repeated until we find X neighbors. All neighboring motes will remember the basestation which calibrated it and only reply to that basestation. For full details, the reader can see the pseudocode in Listings 3 and 4. In the current implementation, we require that for a neighbor to be pushed onto the `valid` list, it must successfully transmit to the basestation `succ` out of `trials` times. Our scheme is similar to Kumar's COMPOW work [16].


```
OPCODE|FrameRef|Role|Address|FrameRef+0 Data|FrameRef+1 Data|…
```

Fig. 4.   Neighbor→Basestation Packet Format

Fig. 5.   Filtering Mechanism

D. Neighbors

Upon receiving a START_SENDING message from the basestation, the neighboring motes will buffer readings for N frames. All data is time stamped with the mote's *own* internal film or video frame counter which is from 2 to MAX_INT. This counter is based on the mote's internal clock. We start the mote from frame 2 due to our latency measurements (described later). Of course, clock drift and skewing will occur; a solution is described later. After $N = 13$ frames are read, the readings are sent out in one packet to the basestation via unicast. A frame number F is attached indicating the frame number of the first data item such that such a packet will have F to $F+N$ frames of data. The packet format is shown in Figure 4.

The nodes do a very simple filtering scheme (Figure 5). The rule is—if the last frame of the previous packet is equivalent to all the data of the current packet to be sent, discard and don't send the packet. Equivalence may not be exact; for example, light readings off the ADC often fluctuate ± 2 even though the light environment is unchanging.

Additionally, if a packet is waiting to be sent (for example, if carrier sense saw that the channel was busy or it is waiting for an ACK), we drop all further readings until the successful (or we give up after retrying some $k = 3$ number of times) transmission of the packet. This is handled by the network layer by BMAC.

VI. SERIAL PORT SERVER

The serial port server waits for connections from the Sylph Middleware Server. It then translates and forwards these messages to each basestation mote. On initial startup, it sends a NUM_NEIGHBORS command to each of its basestation motes to find out how many neighbors surround each basestation and their respective IDs. Besides multiplexing messages from the Sylph Middleware Server, the serial port server has three important software modules, the Timecode Decoder Module, Interpolation Module, and Time Synchronization Module.

A. Timecode Decoder Module

The SMPTE, or Society of Motion Pictures and Television Engineers, timecode standard is an international industry standard timecode signal used in films to achieve time sychronization with frames. A good introduction can be found at [17].
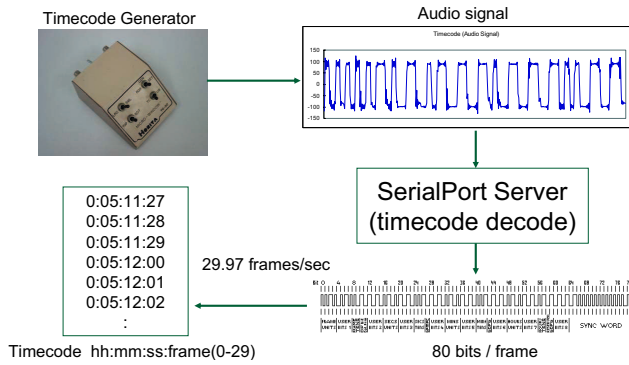
Fig. 6. SMPTE Decoding Process

```
OPCODE|Hour:Min:Sec:Frame|MoteID, Data|MoteID, Data|...
```

Fig. 7. Serial Port Server→Sylph Sensor Module Packet Format

In the evaluation, we use a Mini SMPTE Timecode Generator which generates LTC (longitudinal time code) audio signals (Figure 6) that are typically recorded on a special audio track on a tape, called the "strip". For television, such a signal has a sampling rate of 29.97 fps. Film has a rate of 24 fps. The audio signal encodes hours, seconds, minutes and frames that have passed since the timecode generator was started. Frames cycle from 0 to 29, while hours, seconds and minutes continue to iterate. How the device generates 29.97 fps is out of the scope of this paper; instead, refer to [17].

The timecode decoder module discretizes the SMPTE time-code and generates an event for each decoded period. Such information is used by the Timecode Synchronization Module and Interpolation Module.

B. Interpolation Module

When a Jini Client request information from ARS, it expects to get back information for each and every frame. The interpolation module is responsible for combining all the data received at the multiple basestations and sending it to the Sylph Middleware Server. Since the network connection between the serial server and sensor module is assumed to be much more reliable (i.e. Internet or LAN), we pack all sensor readings into one packet. One packet contains one frame of data for all motes. Figure 7 shows the exact packet format. As mentioned in Section V, motes will not send redundant data. The interpolation module will interpolate missing data. Consider that we may receive packets from a mote for frames (these frame numbers are the *mote's* internal frame count) 2 to 15 and then 19 to 30. In such a case, the interpolation module will fill in frames 16 to 18 with the same data as in frame 15.

The interpolation module has an internal buffer with two threshold values, MAX_BUFFER_SIZE and MAX_QUEUE_SIZE. The two sizes allow us to continue receiving packets during the interpolation process. The internal buffer is filled with data from each of the basestations. This buffer queue is of size MAX_QUEUE_SIZE and once the


Fig. 8. Interpolation Process

queue reaches MAX_BUFFER_SIZE, the interpolation module will construct packets from the last interpolated frame (i.e., the last frame it sent up to the Sylph Middleware Server) to the current frame. By comparing the mote's internal frame number with the frame number calculated from the reference frame counter, the module can combine the data packets into the format described above and send it. Figure 8 illustrates the interpolation module process.

C. Time Synchronization Module

Our time synchronization scheme is based on two factors—that there exists an upper bound on time drift for the motes and that broadcast messages are received at approximately the same time at all receivers. Other time sychronization schemes exist for WSNs [18]. Based on tests conducted at our lab [19], the time drift between two mica motes does not exceed 7 μs per second. Since the duration of one frame is $\frac{1}{29.97}$ seconds, we have:

$$T_{sync} = \frac{\frac{1}{29.97}}{7 \times 10^{-6}} = 4766.7 frames$$

Hence, if we sync the timers in the sensor nodes every 4766.7 frames, then the time drift shall not exceed one frame. Taking a conservative route, we set T_{sync} to be 4500 frames and thus we resync approximately every 2.5 minutes.

The SYNC command is sent from the basestation motes. This command instructs the neighbors to stop sending data, reset their internal frame counters to 2 and then immediately
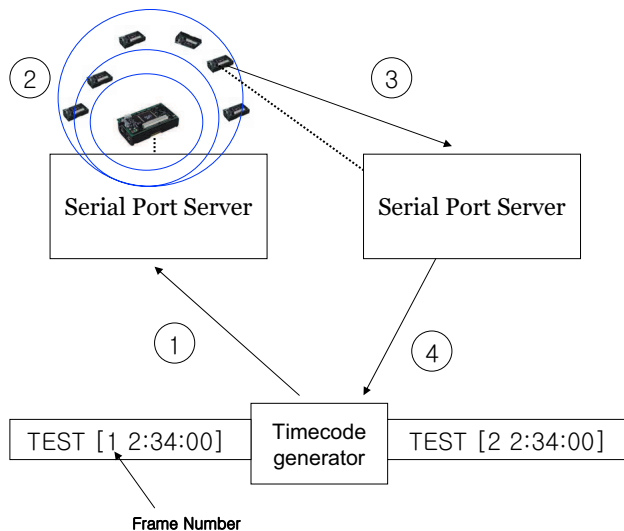
Fig. 9. Latency Test

Fig. 10. Jini Client GUI Interface

start sending. Of course such a command is broadcast to the neighbors, and as BMAC has no way of knowing of a successful transmission of a broadcast, ARS will instruct the basestations to continue sending SYNC until they receive data marked with frame #2 from all its neighbors.

There are two frame counters to consider. The timecode generator is the reference frame counter. Upon sending a `SYNC` command, the Serial Port Server records the time from the timecode generator. Any packets received afterwards are in reference to that recorded frame counter.

However, how can we ensure that the motes will start sending data at the exact same time a `SYNC` command is executed? For this, we constructed the experimental setup detailed in Figure 9. We hooked up a basestation mote and sensor node to serial ports 1 and 2, respectively. At a certain instance of time, the current timecode address is transferred to the basestation mote through serial port 1. Then the basestation mote will broadcast this packet to the sensor nodes. Once the sensor node receives this packet, the contents of the packet are printed out through serial port 2. When receiving packets from serial port 2, we compare the contents of the packet and the current timecode address. We found that the latency is at most 3 frames after 20 trials. This time corresponds about 100 ms. This latency includes the transmission time over the wireless channel, all the processing time in the motes, the serial port server processing and any delay in serial cables. To utilize this latency information, we set the initial frame value in sensor node using this latency value. Even such a small skew of 2 or 3 frames (0.1 seconds) can have a detrimental effect on some applications of ARS—lip synchronization, for example, with audio becomes noticeable with only a few frames of skew (100-150 ms) [20].

VII. Sylph Server Middleware

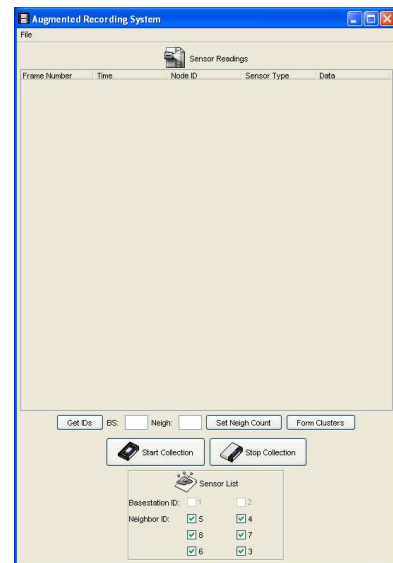The Sylph system handles registration, multiplexing and higher-level query translation for sensor networks. Sylph was chosen because it is particularly well-suited for supporting multiple types of sensor nodes and allows one to write minimal new Jini code for each new type of sensor. It has been successfully used in sensor network projects at UCLA such as Smart Kindergarten. On the lower-layer, Sylph is based on using SET and GET operators on attributes. Once one defines an appropriate *Sensor Module* for their sensors and its attributes, it is easy for a Jini client to change or get these attribute values. For example, in our Augmented Recording Sensor Module, we define attributes such as light, period and command. To begin sending data, the Jini client queries a service discovery domain for sensors and then issues queries to each of them of the form "READ LIGHT EVERY 30 SECONDS". Such a query gets translated to SET and GET requests—"SET PERIOD 30 SECONDS" and "SET COMMAND=STARTSENDING". Thus, Sylph provides a common interface for Jini clients to interact with sensor nodes.

VIII. Jini Client

The client will receive data for each frame and for each mote and will store this information at a database. We use Microsoft Access to store our information and access/store information by using a RMI-JDBC bridge driver.

In designing a GUI to couple with the Jini Client, our first phase will be to design a GUI that will allow us to "scrub" through film or video frames. In other words, one can fast-forward, or rewind through frames and, at the same time, sensor data corresponding to the displayed frames will be shown. The client allows one to create clusters of nodes, find their neighbors and to start or stop collection of data. Figure 10 shows the client. The table displays sensor data when an audio signal and database are specified. The bottom checkboxes allow the user to filter out certain sensor node readings.

Fig. 11. Deployment in a TV Studio (Circles are neighboring motes, squares are basestation motes)

Fig. 12. Example of gradual light changes in the studio.

IX. EVALUATION

We deployed 10 sensors (2 basestations, each with 4 neighbors) in a television studio. Figure 11 shows the overall layout of the studio. The test area was 1232x115 centimeters. The studio has lighting controls accessed from a control room. The control room allows the operator to finely adjust the light intensity of certain overhead lamps. It also allows one to quickly flash lights. We conducted three tests: 1) gradually changing the light intensity, 2) flashing the light quickly and 3) having a person cast shadows over sensors. The metric here is the latency involved between an event captured on video and on the ARS database. In the tests, the clustering algorithm took an average of 2.59 minutes to complete. In three tests, we recorded for approximately 2.5, 2.5 and 10.0 minutes. The total amount of data transferred from neighbors to basestation is 1114884.0 bytes (recall that the packet size is 31 bytes):

$$(2.5 \times 60) \div \frac{1}{29.97} \times 31 = 139360.5 \; bytes$$
$$139360.6 \times 8 = 1114884.0 \; bytes$$

For 10.0 minutes, this is 4459536 bytes. Database sizes were about 390 MB (using MS Access) for all tests.

We only tested with light readings. However, our system can easily be extended to work with other sensors such as temperature, acoustics, etc. that can be attached to our nodes.

A. Gradual Light Intensity Changes

Figure 12 shows footage of a gradual change of lighting. The maximum light intensity value is 255. In these three frames, the values were (from left to right) 235, 200 and 150.


Fig. 13. Graph of light changes from Figure 12.

Figure 13 shows the corresponding values from the leftmost frame to the rightmost frame in the aforementioned figure. From 7 to 8 seconds, there is a large drop. This is due to the SYNC command issued from the basestation. In our tests we found that the broadcast mechanism was far from reliable. Any broadcast to request a SYNC *must* be received by all the neighboring motes, before this ARS can continue. Otherwise, if a SYNC is not acknowledged (i.e., all packets have reset their frame counter to 0), the serial server will continue to broadcast. We found that continual transmissions can sometimes last almost 10 seconds (in Figure 13's case, 2 retries of SYNC were required). Perhaps polling, or unicasting, each neighbor to SYNC would be better, since ACKs are utilized.

B. Sudden Light Intensity Changes

We found it difficult to gauge the latency purely by sight (i.e., watching the video and recorded sensor data stream side-by-side). So, with this test, we stopped the video player at the frame the instant the flash occurs. Of course, such a flash is not truly instantaneous, but we would expect the data to show a sharp increase in light intensity within a short period. Figure 14 shows the first frame we chose. The graph at Figure 15 starts off at that frame. As shown, the light intensity does not start rising until about frame 28. This indicates a latency of nearly one second. This is a rather large latency (when working within the frames metric) and should be fixed. However, in another test we left running for 5 minutes, the latency never actually exceeds 2 seconds and the average latency is 1.3 seconds. So while the SYNC is working to prevent the time drift, there is is still a significant latency warranting investigation.

C. Casting Shadows

This test was done to gauge the responsiveness of ARS. It simulates what would actually be done on a set where human actions in the frame cause changes in the sensor data using light. By flashing lights such that come facing forwards into the performer, the shadow cast will cover the nodes. Since we know beforehand the locations of the motes, we can infer what areas have performers in it. Figure 16 shows one covering a mote explicitly. Again, as with the test beforehand, latency never exceeded 2 seconds.

X. CONCLUSION AND FUTURE WORK

We have described our implementation of the Augmented Recording System and a brief evaluation in a television studio.

Fig. 14. Beginning and ending frames for a light flash.

Fig. 15. Light intensity recorded, where frame 1 equals Figure 14.

Fig. 16. A "performer" casting shadows over the motes.

Our system is fully functional and demonstrates a maximum latency time of 2 seconds. ARS integrates the Sylph Middleware Server to allow easy interoperability with standard Jini clients. We have used tools such as PALOS to easily design mote programs. With augmented footage users can view film or video frames mapped to sensor data readings. This is the first step to realizing a film set embedded with wireless sensor networks at UCLA's Hypermedia Studio.

Our future work is as follows:

Better Synchronization: The current maximum latency of 2 seconds is too long. We are seeking to reduce the latency to less than 1 frame. Java may be a bottleneck and we are investigating alternative languages to implement the serial server in. A smarter scheme of interpolation would also lessen the processing time that may be causing latency.

Calibration: During our experiments it quickly became evident that despite equivalent phenomenon, two different sensor nodes would output readings which were different from each other. The degree of difference varied widely and some sort of normalizing of readings is needed. Work done by Bychkovskiy, et al. [21], begins to address this issue.

More Sensors: Light intensity is certainly not the only reading we are interested in. Sensors that detect color temperature, location, audio, etc. will be especially applicable to film or television.

Semantic Indexing of Video Streams: We would like to allow the user in, some sort of high level language or macro, to express a certain *interest* in browsing parts of video streams where some property holds. For this, we are considering WSN runtime environments such as SensorWare [22] which provide TCL bindings.

Enhanced Activity Detection Via Sensor Fusion: In detecting certain activities on the film set, one type of sensor may not be enough. Our current system simply collects data from various sensors. Future versions will utilize fuse these sensor readings to give a more meaningful interpretation of events occurring

in the environment—for example, vision sensors (cameras) for localization could complement other sensors in a film set..

Currently, the lag between an event and storing it on the database is too large for real-time actuation. However, should ARS be able to fulfill the real-time requirements, equipment positions or settings could be automatically calibrated based on events on the film set, or even scripted events entered in some sort of macro language.

Finally, script supervision, or continuity, ensures that different shots and scenes be logically sound and consistent within a spatiotemporal context. For example, for practical reasons, certain scenes that are at the end of a film may be filmed first. Perhaps ARS would be able to allow one to easily merge different augmented footage frame (or different footage for the same time frame in several shots of different angles) to allow checks for continuity.

ACKNOWLEDGEMENT

REFERENCES

[1] S. Baher, P. M. Davis, and G. Fuis, "Separation of Site Effects and Structural Focusing in the Santa Monica Damage Zone from the Northridge Earthquake," in *Bulletin of the Seismological Society of America*, vol. 92, no. 8, Dec. 2002.

[2] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein, "Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet," in *The Tenth International Conference on Architectural Support for Programming Languages and Operating Systems*, Oct. 2002.

[3] A. Requicha, D. Caron, C. Zhou, W. Zhang, and X. Liu. Detection and Identification of Marine Microorganisms. [Online]. Available: http://cens.ucla.edu/Project-Descriptions/Microorganisms_Identification/index.html

[4] M. Weiser, "The Computer for the 21st Century," *Scientific American*, pp. 94–104, Sept. 1991.

[5] J. Burke, "Dynamic Performance Spaces for Theater Production," *Theatre Design & Technology (U.S. Institute of Theatre Technology)*, vol. 38, no. 1, 2002.

[6] J. Paradiso, C. Adler, K. yuh Hsiao, and M. Renolds, "The Magic Carpet: Physical Sensing for Immersive Environments," in *Proceedings of the CHI '97 Conference on Human Factors in Computing Systems, Extended Abstracts*, 1997, pp. 277–278.

[7] M. Srivastava, R. Muntz, and M. Potkonjak, "Smart Kindergarten: Sensor-based Wireless Networks for Smart Developmental Problem-solving Environments," in *Proceedings of the ACM SIGMOBILE 7th Annual International Conference on Mobile Computing and Networking*, July 2001.

[8] G. D. Abowd, "Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment," in *IBM Systems Journal*, vol. 38, no. 4, 1999.

[9] Y. Matsumoto and A. Zelinsky, "Algorithm and Real-time Implementation of Head Pose and Gaze Direction Measurement System Using Stereo Vision," in *Proceedings of the IEEE 4th International Conference on Face and Gesture Recognition (FG2000)*, 2000, pp. 499–505.

[10] C. Rao, A. Yilmaz, and M. Shah, "View-Invariant Representation and Recognition of Actions," in *International Journal of Computer Vision*, vol. 50, no. 2, 2002, pp. 203–226.

[11] S. L. Dockstader and A. M. Tekalp, "On the Tracking of Articulated and Occluded Video Object Motion," in *Real-Time Imaging (Special Issue)*, vol. 7, no. 5, pp. 415–432.

[12] A. Chen, R. R. Muntz, S. Yuen, I. Locher, S. I. Park, and M. B. Srivastava, "A Support Infrastructure for the Smart Kindergarten," *Proceedings the 6th International Symposium on Wearable Computers*, 2002.

[13] TinyOS: Hardware Designs. [Online]. Available: http://webs.cs.berkeley.edu/tos/hardware/hardware.html

[14] PALOS Sourceforge Page. [Online]. Available: http://palos.sourceforge.net

[15] W. Ye, J. Heidemann, and D. Estrin, "A Flexible and Reliable Radio Communication Stack on Motes," USC Information Sciences Institute," Technical Report ISI-TR-565, Sept. 2002.

[16] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power Control in Ad-Hoc Networks: Theory, Architecture, Algorithm and Implementation of the COMPOW Protocol," in *Proceedings of the European Wireless Conference - Next Generation Wireless Networks: Technologies, Protocols, Services and Applications*, Feb. 2002, pp. 156–162.

[17] SMPTE EBU timecode by Phil Rees. [Online]. Available: http://palos.sourceforge.net

[18] J. Elson, L. Girod, and D. Estrin, "Fine-Grained Network Time Synchronization using Reference Broadcasts," in *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI 2002)*, Dec. 2002.

[19] S. Ganeriwal, private communication, 2002.

[20] C. Sreenan, "To Use or Avoid Global Clocks?" in *The IEEE Workshop on Multimedia Synchronization (Sync '95)*, May 1995.

[21] V. Bychkovskiy, S. Megerian, D. Estrin, and M. Potkonjak, "Colibration: A Collaborative Approach to In-Place Sensor Calibration," in *2nd International Workshop on Information Processing in Sensor Networks (IPSN'03)*, 2003, pp. 301–316.

[22] A. Boulis, C.-C. Han, and M. B. Srivastava, "Design and implementation of a framework for efficient and programmable sensor networks," in *The First International Conference on Mobile Systems, Applications, and Services (MobiSys 2003)*, May 2003.