The *Smart Condo*: Visualizing Independent Living Environments in a Virtual World

Nicholas M. Boers, David Chodos, Jianzhao Huang, Pawel Gburzyński, Ioannis Nikolaidis and Eleni Stroulia
Computing Science Department
University of Alberta
Edmonton, Alberta T6G 2E8
Canada
Email: {boers, chodos, jianzhao, pawel, yannis, stroulia}@cs.ualberta.ca

Abstract—**Providing affordable, high-quality healthcare to the elderly while enabling them to live independently longer is of critical importance. In our *Smart Condo* project, we have deployed a wireless sensor network in an 850-square-foot condominium for assisted living. The sensor network records a variety of events and environmental parameters and feeds the related data into our web-based system. This system is responsible for inferring higher-order information about the activities of the condo's occupant and visualizing the collected information in both a 2D Geographic Information System (GIS) and a 3D virtual world. The architecture is flexible in terms of supported sensor types, analyses, and visualizations through which it communicates this information to its users, including the condo's occupant, their family, and their healthcare providers.**

I. INTRODUCTION

In the context of our service-systems activities (see http://ssrg.cs.ualberta.ca), we are working with colleagues in the Faculty of Rehabilitation Medicine, the Faculty of Arts, the Faculty of Education, and the Faculty of Pharmacy at the University of Alberta towards developing a "*Smart Condo*" that can support patients (and seniors) living in rehabilitation facilities and at home. Although these individuals are, in principle, able to live independently, they are still susceptible to harmful incidents related to physical infirmities or memory loss.

Technology can support independent living in a variety of ways. Medical devices are being redesigned for usability so that individuals can regularly monitor their own health at home. Computerized coaches help patients use these medical devices correctly and consistently. Special-purpose user interfaces allow patients to access information relevant to their condition in a timely manner. Communication devices facilitate contact with family, friends, and medical professionals. Special-purpose sensors, embedded in home-care devices, sense a variety of physiological parameters and allow patients to better control their conditions. Motion, acceleration, and pressure sensors, worn by the individual or embedded in surfaces, provide an unbroken information stream for mining and inferences about sleeping patterns, (ir)regular gait, and falls. This monitoring infrastructure, connected to a notification system, can trigger the attention of family and/or healthcare providers when necessary.

Clearly, the overall problem of at-home health monitoring and care delivery is socially important and technically challenging. Moreover, any proposed technology should be evaluated against the following requirements. First is the issue of privacy: patients, although they may appreciate the increased sense of safety that comes with the monitoring infrastructure, are hesitant to have their every move monitored. The challenge is to determine an acceptable trade-off between privacy intrusion and safety. Second is the issue of variability: patients come with different needs and, as their conditions progress, their needs change. This evolution of patient needs implies the need for an extensible assistive infrastructure that can evolve as necessary. Third is the issue of healthcare personnel training: the effectiveness of new technologies is limited by the knowledge of the people who are using them. There is a need for an education program to train health-sciences professionals in using at-home health monitoring and care technologies.

The main objective of our project is to create a comprehensive framework for affordable, high-quality, non-intrusive home assistance. We are developing the *Smart Condo* prototype in a physical space of roughly 850 square feet, representing a typical living area, and designing it according to Universal Design [1] principles. The *Smart Condo* has been instrumented by embedding numerous sensors connected by a Wireless Sensor Network (WSN). Information from the WSN is archived and processed by a web server that supports a range of REpresentational State Transfer (REST) [2] APIs through which the information is visualized in a variety of views. The most useful views are within a virtual world, currently Second Life (SL), which incorporates a three-dimensional model of the *Smart Condo* reflecting the layout and design of the real *Smart Condo* space. A re-engineered SL client accesses information regarding the patient's activities, as inferred by the sensor data stream, through the server's REST APIs and uses it to control an avatar that mirrors the real-world patient's activity.

The virtual world view of the patient's activity presents a sufficient degree of realism to be intuitive yet minimally intrusive. Specifically, although the system can approximate the patient's position and detect events, it does not record images, sound, or video. By strictly embedding sensors in the environment, the patient does not need to wear a badge,

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020

bracelet, or other position-tracking device. In this manner, our system strikes a balance between accurately representing the patient's condition and invading the patient's privacy.

Healthcare professionals can use this system to augment their patient care practices in several ways. First, they can use the system to monitor the person's activity in real-time, such as in the case of a head nurse responsible for the residents of a nursing-home wing, for example. At the same time, they would also be alerted in the case of a harmful event. Second, logs of the patient's activity, annotated with the sensor readings, can be replayed in an accelerated mode to allow quick viewing of large spans of time, thus providing an alternative mode of reviewing the patient's history. Third, interesting activity segments can be used for simulation-based training of health-sciences students; for example, students can be quizzed on a possible diagnosis based on the differences in a patient's activity patterns. Finally, health professionals can use the virtual world to observe while also interacting with the patient via text (keyboard input) or speech (headset), in order to instruct them on the use of their pills, for example.

In this paper we discuss the architecture of our system and the experience with developing various components of the *Smart Condo* infrastructure. The remainder of this paper is organized as follows. Section II introduces relevant related work. Section III presents an overview of the architecture of our system from the sensor network to the 3D visualization. Section IV describes two development tools that we used. Section V discusses the actual setup of our *Smart Condo* today and focuses on the *Smart Condo* experience. Section VI proposes several potential directions for future work, and Section VII provides a summary and conclusions.

II. RELATED WORK

Sensed Environments: A number of WSN deployments have targeted environmental monitoring (see [3] for a survey), while few have been applied to healthcare settings, e.g., for wearable healthcare devices like heart [4] and other organ [5] monitoring. The CenceMe project [6] uses sensors and cellphones to publish sensed presence information (e.g., current activity and mobility) into popular social networking applications. In terms of sensing environments, [7] developed a system to recognize events in a living environment and provide aural feedback. The system required that the occupant wears an active radio-frequency identification (RFID) tag to help localization tasks and used acceleration sensors to detect doors opening and closing.

Our WSN-based environment sensing approach is similar, although instead of relying on RFID tags, we primarily use motion-detection sensors to estimate the patient's location. Note that the motion sensors constitute only a subset of all sensors deployed in our system. Its distinguishing innovative aspect is the visualization in a virtual world, which mirrors the real environment in which the sensor data are collected.

2D Visualization: Representative examples of 2D sensor-data visualization systems are Microsoft's SenseWeb and SensorMap [8]. SenseWeb is designed to provide shared sensing resources and querying and tasking mechanisms to other sensor applications, the most notable of them being SensorMap, which displays sensor network data on a Microsoft Virtual Earth map. In addition to presenting the most recent data acquired from the sensor network, the system makes it possible to retrieve historic data. The 3D capabilities of Virtual Earth can also be used to present the layout of sensors on 3D terrain. However, SensorMap and related tools are geared to outdoor 3D views with overlaid 2D information, hence a dense sensor deployment will result in cluttered representation. This means that the granularity of representation is unsuitable for data from multiple sensors collected within the limited space of a condo (or a similar living area). Secondly, sensor groups are predefined and contained, restricting the ability to compare sensor data across groups. Thirdly, the presentation of historic data is supported in a very limited fashion by the visual representation. The *Smart Condo* approach to integrating 2D and 3D views of sensed data using two related spatial metaphors, i.e., a 2D GIS as the view for large outdoor environments and a 3D virtual world for smaller indoor environments, is, to our knowledge, unique.

3D Visualization: To our knowledge, two projects are working at interfacing WSNs with 3D visualization. First, [9] describes a SL virtual hospital environment named the Razorback Hospital. It includes virtual artifacts and tools that are slightly ahead of current technology like smart pill bottles and shelves that know when they need to be restocked. The project also looks at supply-chain management, augmented by RFID-tagged inventory, forklifts, shelves, etc. A considerable amount of work has gone into creating virtual analogs of many of the real-world artifacts encountered in a healthcare setting. These range from equipment (IV poles and X-ray machines), to clothes (hospital gowns), and even internal organs. On the whole, the project makes a number of important contributions with respect to virtual healthcare with the overall objective of using the SL hospital for simulation-based training and process optimization purposes, but not actual patient monitoring and healthcare delivery. Second, the previously-mentioned CenceMe project has evolved to visualize its inferences within Second Life [10]. The project visualizes the real actions of sitting, standing, and running as yoga-floating, standing, and flying, respectively.

A few specific efforts overlap somewhat with the *Smart Condo* project. For example, the *Aware Home Research Institute (Georgia Tech)* [11] has pursued several projects related to home healthcare, such as Fetch [12], a mobile system that assists visually impaired people to locate misplaced objects, and Cook's Collage [13], which assists seniors in following recipes. Perhaps their closest project to the *Smart Condo*, however, is the Power Line Positioning (PLP) project [14]. They use a novel approach involving two special modules, one located at each end of the house, to generate tones that emanate through the building's preexisting wiring. Portable tags use the detected signal levels for localization; to locate a person, that person must carry one of these tags. In contrast, our use of passive infrared motion sensors attached to wireless

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020

nodes allows the occupant to roam freely without carrying a tag.

The *MavHome project* at the University of Texas at Arlington [15] takes a more active approach to helping the smart-home occupants. It uses sensors to monitor the state of the environment and analyzes the collected data to (a) identify lifestyle trends, through sequential pattern mining, (b) provide reminders to the home occupants, through prediction of future activities, and (c) detect anomalies in the current data, when the actual sensed events are considered unlikely according to the system's predictions. MavHome's power line control automates all lights and appliances, as well as HVAC, fans, and mini-blinds. Perception of light, humidity, temperature, smoke, gas, motion, and switch settings is performed through a sensor network developed in-house.

The *Sensorized Elderly Care Home* [16] is a system installed in a nursing home in Tokyo. This work is motivated by the desire to alleviate the routine workload of nursing personnel through automation. A sensor-based system is used for localizing patients in a nursing home, monitoring their status, and raising alarms as necessary so that nurses do not have to do routine rounds. The system assumes a relatively limited level of activity on the part of the patients. It relies on "Ultra Badge" transmitters, placed on wheelchairs, and receivers, placed in several locations in the nursing home, to monitor wheelchair movement. Furthermore, a set of transmitters and receivers have been placed on the ceiling to monitor the patient's head position on and around the bed. This latter functionality is not completely evaluated and in place.

The RAUM System (Universität Karlsruhe) [17] seeks to introduce ubiquitous computing into home and office environments by integrating small computers into everyday objects, such as doors, clothes, and furniture. Thus, these objects are turned into soft media devices, which can communicate with other such devices throughout the environment. Like the *Smart Condo* project, there is an interest in using sensors in a home environment to locate people (and other objects, as well). However, the application domain is different from that of the *Smart Condo*. Moreover, in contrast to the *Smart Condo*, the sensors in RAUM are in fact quite capable computing devices, e.g., being able to present information directly to the user.

The Ambient Assisted Living (AAL) Laboratory [18], developed by the Fraunhofer Institute for Experimental Software Engineering in Germany, is an apartment-like environment for developing, integrating, and analyzing ambient intelligence technologies. Currently, the AAL Lab supports the following specific scenarios: (i) monitored drinking via a computerized cup, (ii) monitored food quality via an RFID system built into the refrigerator, (iii) item location tracking, and (iv) fall detection. Although this system integrates a variety of interesting sensors, it does not, to the best of our knowledge, implement a 3D visualization.

Several other projects are at various stages of maturity, like the *Gator Tech Smart House* at the University of Florida [19], currently under construction, whose objective is to assist with the behavioral monitoring (and alteration) of elderly occupants or patients suffering from diabetes and obesity. Not motivated by healthcare needs, but still relevant, is [19] which describes a high-tech house with a variety of sensors and "smart" appliances assisting the occupants. Other projects, such as the WASP architecture [20], were not conceived strictly for smart home applications but focus on the software infrastructure for effectively integrating a population of wireless sensors to recognize events in a living environment and provide feedback. In contrast, our primary objective is monitoring and visualizing the living environment in a virtual world environment.

III. THE *Smart Condo* SYSTEM ARCHITECTURE

The *Smart Condo* system is based on a three-tier architecture (Fig. 1). The main resource is a sensor database (SensorDB), which constantly receives new (raw) data from the sensor network to which it is connected. The data is further processed by server-side scripts to infer higher-order information, such as the individual's location based on multiple motion sensors, and also stored in SensorDB. Around SensorDB we find a WSN-specific component (depicted by the WSN network, the sink, the operations support system, and the stream processing and location engine), and an application-specific component (depicted by the HTTP/PHP server and the path planning application). The HTTP/PHP server exposes a set of REST-based APIs to two clients: the SL view of the condo and its 2D floor plan associated with a location on the GIS component, called SensorGIS. Thus, sensor data are decoupled from their representation to the user, allowing for many different ways of visualizing the same sensor data. For instance, in the case of the SL visualization, a request is sent from within SL to the server for a set of readings over a certain period of time (e.g., the last ten seconds). The server then queries the database and returns the result to SL as XML data. The visualization engine created for this project displays the retrieved data as an avatar walking around the condo (the person moves from one location to the next at the time indicated by that sensor reading), sitting down (detected from pressure sensors on chairs and beds), opening and closing doors, etc. On the other hand, if the data is requested by a 2D web client, then the data will be shown as points on a map.

A. *Wireless Sensor Network (WSN)*

Each node in the *Smart Condo* WSN consists of several components built around a common transceiver module, the DM2200 [21]. This device contains an MSP430F148 microcontroller, a TR8100 radio transceiver, and associated circuitry and is powered by batteries. The microcontroller operates at a low clock speed ($\sim$4.5 MHz) and has limited memory (48 KB Flash and 2 KB RAM). The radio transceiver communicates at a low transmission rate (9600 bps). Each DM2200 is connected to one or more physical sensors, such as acceleration, tactile pressure, switch, or passive infrared motion sensors. One DM2200, attached to a more capable host (the sink), acts as the gateway between the WSN and the IP backbone network. Many sink nodes can be present but the small size of the condo environment is sufficiently covered by one sink node.

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020

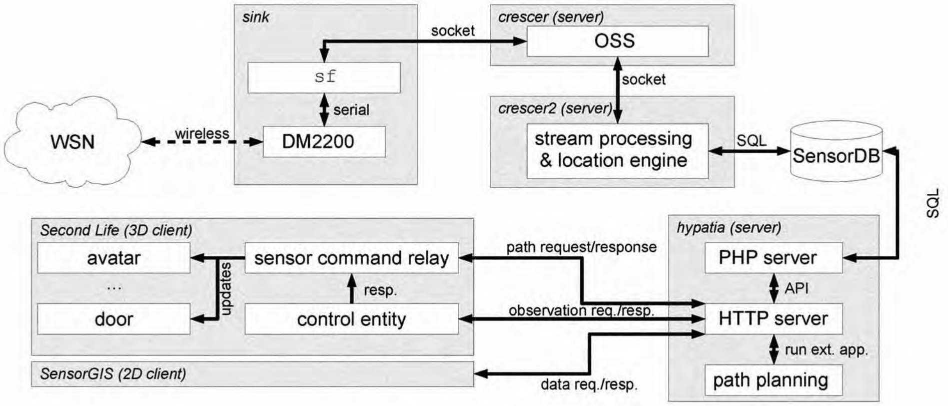

Fig. 1. The software architecture.

In general, the physical environment size relative to a WSN node's transmission range determines the number of required sinks.

Sinks decouple the main software-architecture components from the lower layer protocols used by the WSN (which, in the case of the DM2200, are proprietary). Each sink acts as a gateway of the *unaltered* sensor data to multiple data-processing applications. The routing of the WSN nodes is configured such that the sink(s) is(are) the destination of all WSN data. When a node transmits its observations to a sink, it does so directly in a single hop. Our approach contrasts with more common multi-hop topologies where intermediate nodes forward packets towards the sink (multi-hop communication). To support multi-hop communication, a subset of forwarding nodes must periodically wake up and enable their receivers regardless of whether they have data to send or forward. Maintaining an active radio receiver consumes substantial current (∼5.5 mA) compared with sleeping (∼68 μA). In a confined environment like the *Smart Condo* where a single hop suffices, we save a significant amount of energy by restricting the topology. Each node runs the PicOS operating system [22]. To aid in the software development, we rely extensively on an emulation environment named VUE^2 [23]. We describe VUE^2 in more detail in Section IV-B.

B. Sensor Data Interface

Whenever the sink receives a packet from the WSN, it encapsulates it in the TinyOS [24] serial message format and transmits it over the serial connection to the PC. On the PC, an existing TinyOS program named `sf` forwards between the serial port and a network socket, essentially giving one or more clients access to the serial port. This program forwards all data arriving over the serial port to each connected client and vice versa.

The purpose of this part of the architecture is to offload *all* data processing on more capable servers. The sink node, even though more powerful than a WSN node, is still assumed to have limited resources of its own, except for abundant energy, as it is not battery powered.

Another component, called the Operations Support System (OSS) serves as a link between the raw sensor data and any application logic that needs to process sensor data in a particular fashion. The OSS allows the same WSN to be shared among multiple applications, where each application would have its own servers perform processing on the raw data. As a consequence of the serial to socket forwarding, it receives all packets received at the sink. The OSS (written in Perl) parses the messages and generates acknowledgments for nodes in the WSN. In our application, the OSS is connected to a single other component, namely the stream processing and location engine. Specifically, the OSS provides it with data that include: (i) time-stamp, (ii) network ID, (iii) node ID, (iv) sensor ID, (v) sensor type, and (vi) event type.

C. Sensor Data Stream Processing

The stream processing and location engine stores raw observations in a database (SensorDB) and performs data stream mining to draw conclusions about the environment's state. The engine adds its conclusions to the database; these entries are subsequently retrieved by the 3D visualization engine. The most complicated conclusion drawn from the data stream mining is the location of the *Smart Condo* occupant. Note that, at this point, we assume that there is a single occupant; RFIDs could be used to overcome that limitation, as in [7]. When a person moves within view of a passive-infrared motion sensor, the node sends out a binary value to indicate that it has registered a movement. Viewable regions of the motion sensors may overlap (i.e., several of them may produce output at the same time). As readings arrive from the OSS, the stream-processing engine stores them in a queue: each new reading corresponds to a 3-tuple <time-stamp, event type, node ID>. Multiple readings with the same time-stamp but different node IDs are mapped into a single queued query with multiple node IDs. Within the stream processor, an age threshold determines when to remove old items from the queue. Its value depends on the maximum expected delay for receiving a reading from the WSN. This way we can be confident that all the readings with the a given time-stamp have arrived before we make a decision.

When the occupant is in a region where sensors overlap, the readings of the various motion sensors covering the region may

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020

have slightly different time-stamps; this is why the stream-processing engine is configured to look at a window of time-stamps. After waiting sufficiently long (typically ~5 seconds) for all observations with a given time-stamp to arrive, the process of inferring the location of the occupant begins. Note that to support localization, we populate (in advance) an auxiliary database table with an entry for each region (single and overlapped). We have developed a special-purpose intersection analysis tool (Section IV-A) to support the construction of this table. Each region has an entry in the table containing the node IDs that can sense the region and the corresponding region's centroid. We note that this is one of many possible approaches to location inference. Since it would be easy to deliver the same raw data to multiple sensor stream processors, one could imagine a more refined solution where different location algorithms run in parallel.

D. SensorDB

SensorDB is the repository of raw and processed sensor information. It contains five basic tables: (a) NODES to record the sensor node IDs deployed in the network along with their type (e.g., an observer versus a sink); (b) NODE_TYPES to annotate each node type with a description; (c) SENSORS to annotate each sensor type with a name, description, and range of admissible values; (d) LOCATIONS to identify the physical location of each node's sensors; and (e) READINGS to store all raw readings received from the network. Two additional tables, INTERSECT and INFERENCES, support the 3D visualization. INTERSECT results from our intersection analysis tool for overlapping passive-infrared sensor regions. The localization procedure uses data from this table. The sensor data stream processing component populates the INFERENCES table.

E. 2D Visualization: SensorGIS

The *Smart Condo* is equipped with a GIS component which visualizes sensor network data in the context of a map. Sensor nodes are shown as markers on the map. Using OpenLayers [25], SensorGIS can flexibly switch among different types of maps such as Google Maps,[1] Yahoo Maps,[2] Microsoft Virtual Earth,[3] and NASA Global Mosaic,[4] as well as proprietary maps, if necessary. For instance, to provide a better view of the interior of the building where a sensor network is deployed, the SensorGIS user can superimpose an image of the floor plan of the building on the map, showing the location of each sensor node installed in the building.

Besides showing the locations of the sensor nodes, SensorGIS also visualizes sensor readings over time. As the user navigates the map, he/she may select markers and see the latest readings of the sensors attached to the corresponding node in a panel attached to SensorGIS. Statistics (maximum, minimum, average, standard deviation, and variance) of a sensor within a user-defined period can be retrieved and displayed in a sortable table, and historical readings can be plotted as a graph. Furthermore, users can select a group of sensors, either by drawing a polygon on the SensorGIS map or by clicking on individual markers, and see their most recent readings in a view that enables comparisons among these readings. One can also easily obtain the statistics of the sensor readings over a specified period of time.

All queries issued by SensorGIS to the server use AJAX (Asynchronous JavaScript and XML), which enables a high degree of interaction and responsiveness of the interface.

F. 3D Visualization: SL

The visualization component, which shows the subject's position using a virtual (SL) representation of the condo, is itself composed of several parts: a conversion utility that can turn a 2D blueprint into a 3D home, an automated character module that moves the occupant's avatar through the *Smart Condo* area, a path planning algorithm to guide the occupant's avatar around obstacles as it moves between locations, and a control system to coordinate these components.

The first step in creating the SL visualization is to create a representation of the physical space of the *Smart Condo* in the virtual world. This is achieved by creating a simplified version of the blueprint, indicating the location of walls and furniture, and then using that information to generate virtual walls and furniture in the appropriate locations. We have developed a simple web-based tool for the user to input the wall and furniture locations, which then stores these locations in a database. These locations can either be loaded from a vector-based image file or, if necessary, entered manually via a web form. Once this input process is complete, a program written within SL reads the information from the database and creates walls and furniture in the specified locations. Finally, color and textures may be added, through the SL tools, to improve the appearance of the virtual condo. At the same time, the wall and furniture locations are used to create a grid-based obstacle map, which guides the path planning algorithm. Fig. 2 (top) shows a top-down view of the constructed SL condo.

Once the virtual space has been created, the occupant's avatar can be placed within that space and given instructions on where it should move. These instructions are accessed by the control system through the server's API in the form of geographic (latitude/longitude) co-ordinates and associated time-stamps. These data are then converted in both space and time: the spatial co-ordinates are translated to the virtual-world co-ordinates and their associated time-stamps are translated to match the replay speed chosen by the user. Thus, the control system creates a list of local destination co-ordinates along with the appropriate delay between each destination. Before these co-ordinates are used to move the occupant's avatar, however, they are converted to a set of intermediate points using the path planning algorithm. This component, a C implementation of a potential-fields algorithm [26], ensures that the character does not walk through any obstacles on its way to a destination.

---

[1] See http://maps.google.com/.

[2] See http://maps.yahoo.com/.

[3] See http://www.microsoft.com/VirtualEarth/.

[4] See http://worldwind.arc.nasa.gov/.

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020

Fig. 2. The *Smart Condo* rendered in SL (top) and a snapshot of the occupant's avatar in the SL condo (bottom).

The automated character module that controls the occupant's avatar was created using LibSL, an open-source project that provides access to SL functionality via a C#-based API. Via LibSL, the module is able to control an SL avatar and cause it to follow instructions received from the control system rather than from a human user. The module is set up to receive a set of co-ordinates from the path planning algorithm, and to follow that path to reach each destination. The module also instructs the character to sit upon arrival at a chair, for example, and to open a door upon receiving a corresponding sensor message. Fig. 2 (bottom) is a screenshot of the occupant avatar in the SL condo.

G. REST Style Integration

As we mentioned at the beginning of Section III, the server-side intermediary program uses a REST-style API [2] to allow the SL client (as well as other clients) access to the sensor data. REST is based on a common understanding between the server program and its clients about the format and structure of data being exchanged. Simply put, the components agree on a set of HTTP GET-based parameters to be sent by the client, and the structure of an XML document to be sent in response by the server. In this case, the SL client issues a request to the intermediary program for sensor data when the user asks to view the subject's location over a certain period of time. The intermediary program receives this request and sends an appropriate database query to the sensor database. The results are returned to the program, processed and converted into an XML document, and then returned to the SL engine (via an additional processing script that simplifies the data further for use in the SL programming environment). These data include the location, time, and sensor type for each sensor reading within the requested time window.

The decision to use REST was particularly important for our project because of the restrictions imposed by the SL programming environment. This environment limits communication with external web-based applications to a URL-based request and the receipt of the response as a string of text which can then be parsed by the limited text-processing functions. Fortunately, these limited capabilities are sufficient for a REST-based system, and thus the SL client is able to communicate effectively with the intermediary program.

IV. DEVELOPMENT EXPERIENCE

A. Intersection Analysis Tool

To monitor the subject, multiple types of sensors can be deployed in the *Smart Condo*. By fusing the information provided by sensors automatically, the monitoring system is able to visualize the patient's current condition. One of the key elements in this information fusion process is localization.

Though each sensor detects environmental changes independently and has its own detection zone, different zones might overlap. It is likely that a single movement of the patient triggers multiple readings from different sensors. In such a case, those readings will indicate that the patient is somewhere within the intersection of the detection zones and therefore his/her location can be computed. However, while deploying the sensors, it might not be obvious whether and where a detection zone will intersect with other zones. Thus to help us determine the locations of those intersections, an analysis tool was developed.

The analysis tool imports the floor plan of the *Smart Condo* and emulates the network layout based on user inputs. The coverage of each sensor and detection zone intersections are visualized on the floor plan. The user can adjust the sensor layout easily and re-compute the visual coverage when that layout changes. After the user comes up with a satisfactory layout, he/she can save it for further adjustment and/or task the tool to export the respective table entries into the database for localization lookup. This way, the analysis tool can save the network developer much effort in identifying the intersections.

B. Emulation

To test the stability and reliability of the *Smart Condo* software, we first set up an emulation environment. By using SMURPH [27], its wireless extensions [28], and its implementation of the PicOS API [23], we can develop the application source code once and then compile it for both the actual hardware nodes and the simulation environment. Therefore, we can switch back and forth between the real and virtual environments with little effort. Furthermore, we built our Operations Support System so that it can interface with both the real and simulated systems. In this manner, we can also thoroughly test our OSS, including the stream processing and location engine and other related software, without actually deploying any hardware. Needless to say, authoritative testing by emulation is invaluable, especially as a prerequisite to a large scale *production* WSN deployment, with sensor nodes embedded inside the furniture, appliances, etc.

V. THE *Smart Condo* EXPERIENCE

Our *Smart Condo* is located in a single large room (850 ft^2) in a building on the University of Alberta campus. As part of an undergraduate course, teams of Industrial Design and Occupational Therapy (OT) students converted this area into a six-room condominium. Inside each room, they created prototypes for appliances, furniture, and other fixtures. Inside this space, we have deployed our sensor network, which currently consists of nineteen nodes. In terms of motion sensors, we have deployed six passive-infrared motion sensors for spot detection (Panasonic AMN43121) and seven passive-infrared motion sensors for wide-area detection (Panasonic AMN44121). Spaced throughout the condominium, the thirteen motion sensors give us adequate coverage of the unit to successfully locate an occupant.

Two of the chairs within the condo have pressure sensors (FlexiForce A201, 1 pound), which allow us to detect when someone sits on either chair. We attached reed switches to the front door and the door of the microwave, to determine whether they are open. We attached an accelerometer to the front door to detect knocking. We are currently placing electric current sensors within the unit to detect whether devices (e.g., a waffle maker) are turned on.

After placing sensors within the unit, we worked with our colleagues in occupational therapy (OT) to evaluate our work. A test subject followed a number of scripts within the unit while the OTs evaluated the virtual representation. The *Smart Condo* contains a number of video cameras (which were installed in the room for a different project) which we used to verify inferences about the location of the occupant. Our OT colleagues matched the script with their observations of both the video feed and the virtual representation. The scripts included such actions as (a) moving from room to room, (b) sitting on chairs, and (c) opening and closing doors. The initial reaction of our colleagues was positive: they were very satisfied with the fidelity of the virtual representation with respect to the real world, and felt that this would be a useful tool for monitoring patients undergoing rehabilitation.

VI. FUTURE WORK

A. SensorGIS

A marker group manager is essential to improve the 2D (SensorGIS) visualization of large numbers of sensors, at different levels of geospatial detail, so that markers can be aggregated when the user zooms out. This technique is especially useful in representing the relationship among markers and, at the same time, prevents filling the map with a large number of markers.

In SensorGIS, a marker group would correspond to a sensor group. While the Google Maps marker group manager mainly deals with appearance based on zoom level, the marker group manager in SensorGIS must take care of both appearance and all the existing query or management operations. For instance when a marker denoting a sensor group is selected, group query functions should apply on this marker.

A user interface for managing marker groups should also be provided so that users can change, delete, or add members in a group. This interface must communicate with the server and manage the database tables containing the group member list.

B. 3D Visualization

Much of the future work in this area centers on the representation of different types of sensor readings within a virtual environment. For example, when a pressure sensor registers a reading corresponding to someone in the real world sitting down, this causes the avatar in the virtual environment to sit down, as well. However, one can easily imagine sensor readings that correspond to more complex real-world actions or activities, and the virtual environment will need to convey these actions realistically and effectively. One example is a heat sensor located over a stove element. Depending on the type of stove, this could be represented by glowing burner coils for an electric stove or a circular flame for a gas stove.

Another important area that we plan to explore is two-way communication between the virtual environment and the real world. At present, the virtual environment is used to receive and display information. However, the patient could also be a participant in the virtual environment, and could be provided with tools to interact with medical staff or family members who are logged into the system. SL has a number of built-in communication tools, including the ability to chat using either a keyboard or a microphone and headset.

C. Adding Interaction

We aim to expand our system beyond passive monitoring by adding the possibility of interaction with the patient. In [29], the authors suggest using a TV (a familiar interface) as a means for such interaction. Potential applications could include medicine reminders, personal assistant functionality, and home automation. A familiar remote control could allow the patient to interact with the system. More technologically savvy patients could also take advantage of the SL 3D virtualization for interaction.

VII. SUMMARY AND CONCLUSIONS

In this paper, we discussed our initial work in the *Smart Condo* project. In collaboration with colleagues from Occupational Therapy and Industrial Design, we have developed a model condo, designed according to Universal Design principles, within which we have embedded a wireless network with a variety of sensors. Information from the sensor network is archived in a server which supports a range of REST APIs. Using these APIs, the information is visualized in a 2D GIS and a 3D virtual world. Although there has been substantial research on using sensor networks to monitoring environments for assisted living purposes, we believe that our approach to

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020

this problem is innovative in several respects. First, we are experimenting with a variety of sensor types, and we are looking into integrating commodity home-care device sensors in our networks. Second, we have developed an integrated software architecture with a component for collecting and archiving sensor-network data which is unaffected by changes in the sensor network topology, a stream-data mining component for synthesizing raw sensor network data into higher-order information, and a set of APIs through which the information can be provided to different clients for different types of visualization. Finally, we are using a virtual world, i.e., SL, as a highly realistic visualization of the condo and the activities of its occupant. This visualization is intuitive and easy to use for all healthcare professionals, who do not need to interpret graphs to infer information about the patient. At the same time, the system is minimally intrusive, since the patient's appearance is not monitored or recorded.

This work is in its initial stages, but based on our preliminary evaluation, we have established the effectiveness of the approach and the overall correctness of the virtual-world visualization. This work was originally done in the context of an undergraduate course with Occupational Therapy and Industrial Design students, and we are working towards further evolving the course to take into account the potential of sensor networks. In this manner, we are confident that such deployments will be effectively used by health professionals.

ACKNOWLEDGMENT

The authors are grateful to their colleagues, Drs. Lili Liu, Robert Lederer, and Sharla King, whose tireless work, insightful comments, and humor were essential in making this happen. Thanks are also due to the funding agencies, NSERC, Alberta Advanced Education and Technology, iCORE, and IBM, that sponsored this research.

REFERENCES

[1] The Center for Universal Design, *The Principles of Universal Design, Version 2.0*. Raleigh, NC, USA: North Carolina State University, 1997.

[2] R. T. Fielding and R. N. Taylor, "Principled design of the modern web architecture," *ACM Trans. Interet Technol.*, vol. 2, no. 2, pp. 115–150, 2002.

[3] J. K. Hart and K. Martinez, "Environmental sensor networks: A revolution in the earth system science?" in *Earth-Science Reviews*, Oct. 2006, vol. 78, no. 3-4, pp. 177–191. [Online]. Available: http://www.sciencedirect.com/science/article/B6V62-4KBVWX7-2/2/9131e5d84373382ba80cb07e1eaece48

[4] V. Shnayder, B. Chen, K. Lorincz, T. Jones, and M. Welsh, "Sensor networks for medical care," in *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems*, vol. 2, no. 4, 2005, pp. 314–314.

[5] R. Bader, M. Pinto, F. Spenrath, P. Wollmann, and F. Kargl, "BigNurse: A wireless ad hoc network for patient monitoring," in *Pervasive Health Conference and Workshops, 2006*, Nov.–Dec. 2006, pp. 1–4.

[6] E. Miluzzo, N. D. Lane, S. B. Eisenman, and A. T. Campbell, *Smart Sensing and Context*, 2007, ch. CenceMe – Injecting Sensing Presence into Social Networking Applications, pp. 1–28.

[7] H. Kanai, T. Nakada, Y. Hanbat, and S. Kunifuji, "A support system for context awareness in a group home using sound cues," in *PervasiveHealth 2008: Second International Conference on Pervasive Computing Technologies for Healthcare*, Jan.–Feb. 2008, pp. 264–267.

[8] "SenseWeb Project." [Online]. Available: http://research.microsoft.com/nec/SenseWeb/

[9] C. W. Thompson and F. Hagstrom, "Modeling healthcare logistics in a virtual world," *IEEE Internet Computing*, vol. 12, no. 5, pp. 100–104, 2008.

[10] M. Musolesi, E. Miluzzo, N. D. Lane, S. B. Eisenman, T. Choudhury, and A. T. Campbell, "The Second Life of a sensor: Integrating real-world experience in virtual worlds using mobile phones," in *Proceedings of the 5th Workshop on Embedded Networked Sensors (EmNets'08)*, Charlottesville, Virginia, Jun. 2-3 2008.

[11] C. D. Kidd, R. Orr, G. D. Abowd, C. G. Atkeson, I. A. Essa, B. MacIntyre, E. D. Mynatt, T. Starner, and W. Newstetter, "The Aware Home: A living laboratory for ubiquitous computing research," in *CoBuild '99: Proceedings of the Second International Workshop on Cooperative Buildings, Integrating Information, Organization, and Architecture*. London, UK: Springer-Verlag, 1999, pp. 191–198.

[12] J. A. Kientz, S. N. Patel, A. Z. Tyebkhan, B. Gane, J. Wiley, and G. D. Abowd, "Where's my stuff?: design and evaluation of a mobile system for locating lost items for the visually impaired," in *Assets '06: Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility*. New York, NY, USA: ACM, 2006, pp. 103–110.

[13] Q. Tran, G. Calcaterra, and E. Mynatt, "Cook's collage: Deja vu display for a home kitchen," in *Proceedings of HOIT: Home-Oriented Informatics and Telematics*, 2005, pp. 15–32.

[14] S. N. Patel, K. N. Truong, and G. D. Abowd, "Powerline positioning: A practical sub-room-level indoor location system for domestic use," in *Proceedings of UbiComp 2006: Ubiquitous Computing*, 2006, pp. 441–458.

[15] D. J. Cook, "Health monitoring and assistance to support aging in place," in *Journal of UCS, Volume 12, Issue 1*, 2006, pp. 15–29.

[16] T. Hori, Y. Nishida, H. Aizawa, S. Murakami, and H. Mizoguchi, "Sensor network for supporting elderly care home sensors," in *Proceedings of IEEE, Volume 2*, 2004, pp. 575–578.

[17] F. Hupfeld and M. Beigl, "Spatially aware local communication in the RAUM system," in *Interactive Distributed Multimedia Systems and Telecommunication Services*, 2000, pp. 293–319.

[18] M. Becker, E. Werkman, M. Anastasopoulos, and T. Kleinberger, "Approaching ambient intelligent home care systems," in *Pervasive Health Conference and Workshops, 2006*, Nov.–Dec. 2006, pp. 1–10.

[19] S. Helal, W. Mann, H. El-Zabadani, J. King, Y. Kaddoura, and E. Jansen, "The Gator Tech Smart House: A programmable pervasive space," *Computer*, vol. 38, no. 3, pp. 50–60, 2005.

[20] L. Atallah, B. Lo, G.-Z. Yang, and F. Siegemund, "Wirelessly accessible sensor populations (WASP) for elderly care monitoring," in *PervasiveHealth 2008: Second International Conference on Pervasive Computing Technologies for Healthcare*, Jan.–Feb. 2008, pp. 2–7.

[21] "RF Monolithics: DM2200." [Online]. Available: http://www.rfm.com/products/pdf/dm2200-916vm.pdf

[22] E. Akhmetshina, P. Gburzyński, and F. Vizeacoumar, "PicOS: A tiny operating system for extremely small embedded platforms," in *Proc. of ESA'03*, Las Vegas, Jun. 2003, pp. 116–122.

[23] N. M. Boers, P. Gburzyński, I. Nikolaidis, and W. Olesinski, "Supporting wireless application development via virtual execution," in *Proc. of Workshop on Wireless and Unstructured Networking: Taming the Stray (WAHOC'08)*, Wisla, Poland, Oct. 2008.

[24] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler, *Ambient Intelligence*. Springer Berlin Heidelberg, 2005, ch. TinyOS: An Operating System for Sensor Networks, pp. 115–148.

[25] "OpenLayers." [Online]. Available: http://www.openlayers.org

[26] M. Greytak, "Numerical potential field path planning tutorial." [Online]. Available: http://web.mit.edu/mgreytak/OldFiles/Public/NPFwebsite/

[27] W. Dobosiewicz and P. Gburzyński, "Protocol design in SMURPH," in *State-of-the-art in Performance Modeling and Simulation*, J. Walrand and K. Bagchi, Eds. Gordon and Breach, 1997, pp. 255–274.

[28] P. Gburzyński and I. Nikolaidis, "Wireless network simulation extensions in SMURPH/SIDE," in *Proc. of the 2006 Winter Simulation Conference (WSC'06)*, Monetery, California, Dec. 2006.

[29] M. Ghorbel, F. Arab, and M. Mokhtari, "Assistive housing: Case study in a residence for elderly people," in *PervasiveHealth 2008: Second International Conference on Pervasive Computing Technologies for Healthcare*, Jan.–Feb. 2008, pp. 140–143.

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.6020
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6020