Chapter 1

A SELF-CONFIGURING LOCATION DISCOVERY SYSTEM FOR SMART ENVIRONMENTS

Andreas Savvides
Electrical Engineering Department
Yale University
andreas.savvides@yale.edu

Mani Srivastava
Electrical Engineering Department
University of California, Los Angeles
mbs@ee.ucla.edu

Abstract This chapter describes our experiences with the design and use of a rapidly installable self-configuring beaconing system. Our system is comprised of a set of custom-designed wireless sensor nodes, the Medusa MK-2 nodes that can form a local coordinate system a few seconds after they are deployed. This results in a versatile, low-cost system based on battery operated devices eliminates the need for costly and time-consuming infrastructure installations. During the system bootstrapping phase, the Medusa Kindergarten - 2(MK-2) nodes localize themselves and then enter a service phase in which they act as "satellites that assist other objects in the room to determine their locations with a few centimeters of accuracy.

Keywords: Sensor networks, node localization, location aware systems

Introduction

As wireless embedded systems evolve, location awareness is becoming a fundamental requirement for small wireless devices to operate autonomously and it is the key enabler for many applications. This chapter reports our experiences in adding location awareness to a deeply instrumented system for observing the level of children interaction during early childhood education. This work was conducted as part of the Smart Kindergarten project [9, 8] at the University of California, Los Angeles. The Smart Kindergarten project aims to develop a deeply instrumented environment for studying child development in early childhood education. The overall goal of the project is to provide the necessary infrastructure support for investigating how learning takes place inside a classroom during a childs pre-school years. To do so the members of the Smart Kindergarten team have designed and implemented a wide variety of platforms and components that enable the close observation of student behavior in the form of speech, gestures and their interaction with toys and objects in a classroom setup. The Smart Kindergarten environment hosts a wide variety of devices that ubiquitously record the activities of children inside the classroom. A *Smart Table* [3] was developed to track multiple objects sitting on its surface. By identifying and localizing various objects on the table surface, the Smart Table measures the level of interaction of kids with a set of puzzle blocks sitting on the table. An elaborate speech processing system has been designed to record and process speech streams, and a rapidly deployable ad-hoc localization system has been designed to estimate and record high precision positions of students (within 10 centimeters) and objects within the classroom. This time-stamped information is propagated and stored in a backend database that can be used by the education researchers that need to process the data offline at a later stage.

This chapter focuses on the localization subsystem components. The uniqueness of this system lies in the new innovative design of an ecology of sensing devices, supported by a software infrastructure and a new set of location discovery algorithms specifically designed for this project. Instead of providing all the details of the project, this chapter highlights the main features of the smart beaconing system, and provides a set of references that describe each component in greater detail. The first half of the chapter describes the software and hardware components of the localization infrastructure, and the second half presents an overview of our results in the broader topic of ad-hoc node localization. The chapter concludes with summary of the lessons learned and a set of directions for future work.

1. Localization System Hardware Building Blocks

To closely observe student locations inside a classroom we have developed a location infrastructure shown in Figure 1.1. The infrastructure based on an ecology of wirelessly connected sensing devices that among

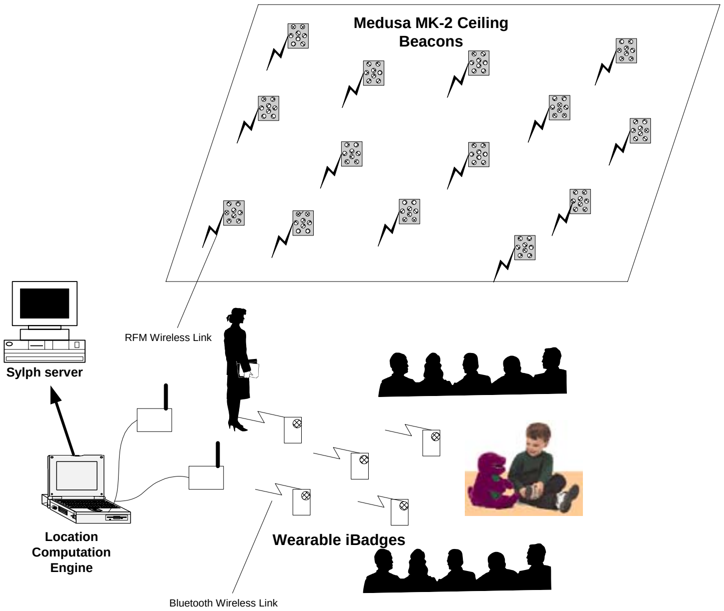

Figure 1.1. The Smart Kindergarten Localization Infrastructure

other tasks are responsible for extracting detailed location information from the classroom environment. The students are tracked with the help of a custom designed wearable device, the *iBadge* [2], that is able to obtain its location with the help of a set of smart beacons, the Medusa MK-2 beacon nodes attached on the classroom ceiling. Other objects in the room are tracked with an object tag which is implemented with a Mica wireless sensor node designed at UC Berkeley. All the devices are battery operated and the system is designed to be rapidly deployable and self-configuring. Operation in a typical classroom setting will proceed as follows. First the ceiling beacons are evenly placed on the classroom ceiling. During an initial *bootstrapping phase*, the ceiling beacons form a local coordinate system by measuring the horizontal distances to each other using their onboard ultrasonic distance measurement system. This process takes a few seconds to complete and the locations of the ceiling beacons are stored on a workstation that serves as the *location computation engine*. Once the bootstrapping phase is completed, the ceiling beacons enter a *service mode*. When in service mode, the beacons synchronize among themselves to broadcast a combination of radio and ultrasound reference signals into the classroom space at a frequency of approximately 12 reference signals per second. The iBadges and ob-

Figure 1.2. The Smart Kindergarten Localization Platforms

ject tags, use these signals to compute their distances to the beacons (the Medusa MK-2 nodes) and transmit their time-stamped distance measurements to the location computation engine that computes node locations and stores them in the Sylph server [10] for future processing.

While it is possible to pursue a distributed implementation where beacons could transmit their signaling in a more ad-hoc manner and devices in the room could estimate their locations individually, we decided on a centralized implementation for performance purposes. By having beacons transmit their reference signals in a synchronous manner, the localization process can achieve a higher rate, that is, devices inside the room can receive location updates at a faster rate. For the same reason, the devices inside the classroom propagate their measurements to the location computation engine, to achieve higher rate processing of the measurements. In addition to performance issues, this setup was also selected to favor scalability. By having a limited set of ceiling beacons transmitting the reference information, and allowing wearable devices and object tags to make passive measurements, one can accommodate a very large number of devices in the same room without compromising the location update rate of the system. Figure 1.2 depicts the platforms

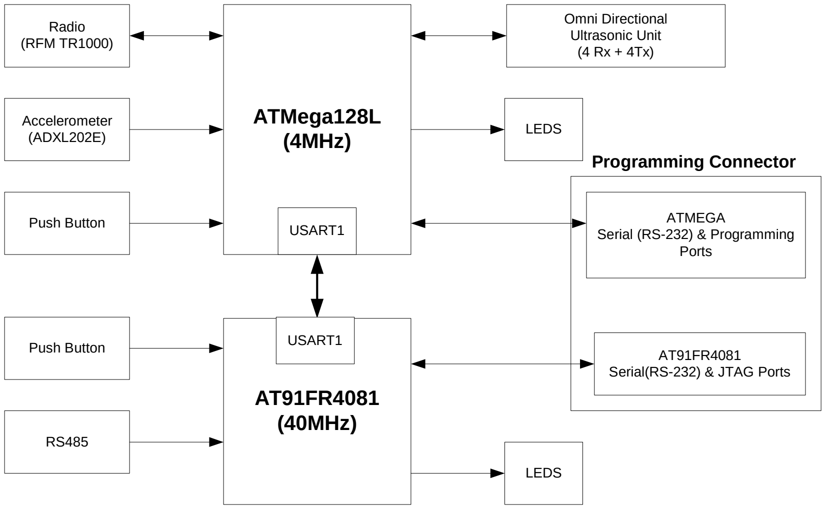

Figure 1.3. Medusa MK-2 architecture

designed for this system. A more detailed description of the platforms is provided in the subsections that follow.

Ceiling Beacons: *Medusa MK-2*

The Medusa MK-2 node [6] is a low cost, low power wireless sensor node, specifically designed to act as a ceiling beacon node. The MK-2 node has two on board microcontrollers, an ATMega128L and an AT91FR4081, both from ATMEL. The ATMega128L is an 8-bit microcontroller running at 4MHz and it is responsible for running frequent tasks that do not require any heavy weight computation. Such tasks include driving the on-board RFM TR1000 radio, handling a lightweight MAC layer and driving the ultrasonic ranging system. The different software components are implemented as tasks inside the PALOS embedded operating system that will be described in section 2.0.

The AT91FR4081 processor features an ARM7TDMI core supporting the 32-bit ARM and the 16-bit THUMB instruction sets. This processor runs at 40MHZ and has 128Kbytes of on-chip RAM and 1Mbyte of FLASH. This processor is dedicated to less frequent but more computation intensive tasks such as the least squares estimation algorithms used to estimate node locations during the bootstrapping phase. The location computation in a distributed fashion inside the sensor node or it can take place on the sensor node using the algorithm described in [6] or at the central location computation engine. Figure 1.3 shows a block diagram that describes the MK-2 architecture.

In addition to the ultrasonic ranging board, the Medusa MK-2 node includes a 2-axis MEMS accelerometer to detect node movement, two push-buttons that can serve as user interfaces, as well as GPS interface to receive position and timing information from an additional outdoor node placed outside the classroom. The two processors communicate with each other through a UART and each processor has an additional UART attached to the programming connector pins that allows the processors to communicate with external devices.

The Ultrasonic Ranging Subystem. The ultrasonic ranging system is a separate accessory board attached to the Medusa MK-2 node. This board carries four receiver-transmitter pairs of 40KHz, 120-degree beam angle ultrasonic transducers. Three transducer pairs are aligned in a circular pattern on the board perimeter, to provide a 360-degree angle. These transducers are slightly bent so that the ultrasonic signal transmissions coverage starts from the plane facing the node enclosure base and expands towards the plane perpendicular to the node base. The fourth receiver transmitter pair is perpendicular to the board to cover the region directly below the node. This transducer alignment produces a hemispherical coverage pattern that allows the ceiling beacons to measure distances to each other on the horizontal ceiling plane. At the same time, the ultrasonic transmissions are also transmitted in the room space, allowing the indoor localization of objects in 3D space. Each ultrasonic transmitter is driven by a separate general purpose I/O line from the ATMega128L microcontroller. The receiver circuit uses a two-stage op-amp amplifier followed by a threshold comparator and outputs a digital signal upon the detection of an ultrasonic signal. Each receiver output is wired to an external interrupt line that interrupts the ATMega128L microcontroller each time an ultrasonic transmission is detected.

Distance is measured by recording the time difference between the arrival of a radio pulse and an ultrasound pulse. The effective measurement range can range from a few centimeters to approximately 20 meters. In the MK-2 implementation, the maximum measurement range is about 4 meters. We found this range to be suitable for localizing objects in a room for several reasons. First, the transmission range is long enough to reach sensor nodes lying on the room floor. Since the transducers beam pattern has a lobe shape, the maximum horizontal range is slightly less (around 2.5 - 3 meters). This allowed the experimentation with the position calibration of smart beacons by using multihop measurements.

One of the main design parameters of the smart beacon system is multihop operation. This decision is based on four main reasons. First, the ability of the system to operate using multihop measurements, can help reduce localization latency. When low power ultrasonic transmissions are used, the beacons have to wait for a smaller interval for a any reverberations to die out before transmitting a new signal. Second, lower power ultrasonic transmissions; prolong the battery lifetime of each beacon. Third, multihop operation makes the system more scalable. This is desirable when deploying the ceiling beacons in rooms or over corridors in there is no full measurement connectivity between beacons. Multihop operation can localize beacons in the presence of obstacles by using indirect line-of-sight measurements.

Wearable units and object tags: *iBadge* and *Mica Ranger*

A wearable wireless badge, the iBadge [2], is equipped with a speech-processing unit, two radios (a Bluetooth and a low power RFM radio from RF Monolithics), 3D axis magnetic and acceleration sensors, temperature and humidity sensors and an ultrasonic ranging using to assist with localization of the iBadge. All the sensor measurements are transmitted through a Bluetooth link to a backend server, where all the data is stored in the Sylph middleware infrastructure. Sylph supports a set of abstractions that allow the easy connection of the sensors to the infrastructure and provides a suitable API that allows researchers to perform data mining on the collected data.

The object tag is implemented using a Mica sensor node from Crossbow [11]. It carries a custom design ultrasonic ranging board that is compatible with the ultrasonic ranging system on the Medusa MK-2 and iBadge nodes. The main difference of the object tag to the iBadge is that it can also act as a beacon by featuring both an ultrasonic transmission and reception. This was specifically implemented to facilitate localization in the presence of obstructions inside a room. The placement of object tags at different places inside the classroom, we can increase the probability that other objects and badges in the room to can be localized. This is because the object tags can transmit beacon signals from other locations in the room, to reach devices that have limited line-of-sight to the ceiling beacons.

2. Software Infrastructure Support

The localization infrastructure is supported by three main software components, the Palos embedded operating system running on all the wireless platforms (MK-2, iBadge and Mica Ranger), the measurement protocol stack that drives the ultrasonic measurement system and the location computation engine that fuses distance measurements to estimate node locations.

Palos

Palos (for Power Aware Lightweight OS) is a lightweight pseudo real-time non preemptive OS developed by Sung Park at the Networked and Embedded Systems Lab at UCLA to support the Smart Kindergarten platforms. Palos implements a lightweight pseudo real-time multitasking scheme where different functionalities are specified as a set of well defined tasks. The operating system consists of three main entities, the OS Core the hardware abstraction layer(HAL), a manager and a set of tasks. The HAL implements the drivers to microcontroller specific peripherals (e.g UART and SPI drivers) and platform specific components such the RFM radio drivers and the drivers for the ultrasonic ranging subsystems. Palos tasks are similar to threads on legacy operating systems. Tasks are registered with the OS core during system initialization and are either scheduled to execute periodically according to a set of user defined intervals. Each task can communicate with other tasks by depositing an event in each others event queue. Finally, the OS core runs the main control loop that determines the order and frequency of task execution and maintains the task event queues. When a task has turn to execute, any events previously deposited in the task queue are sequentially passed to the task. The operating system does not support task preemption but provides a set of mechanisms for prioritizing, stopping and resuming tasks. The stripped Palos core of the operating system occupies 956 bytes of microcontroller FLASH and 548 bytes of RAM.

Distance Measurement Stack

To handle the communication and measurement process among the ceiling beacons during the bootstrapping phase, and the distance measurement between beacons and the mobile entities during the service phase, we have developed a lightweight layered distance measurement stack, shown in figure 1.4. The bottom layer is an integrated layer that acts as a driver for the radio and the ultrasonic ranging hardware. The BMAC layer implements a simple CSMA medium access control protocol that provides reliable communication on the RF communication channel. The ranging coordination layer is implemented on top of BMAC and is responsible for the coordination of ultrasonic transmissions between nodes. This layer is responsible for ensuring that ultrasonic transmis-

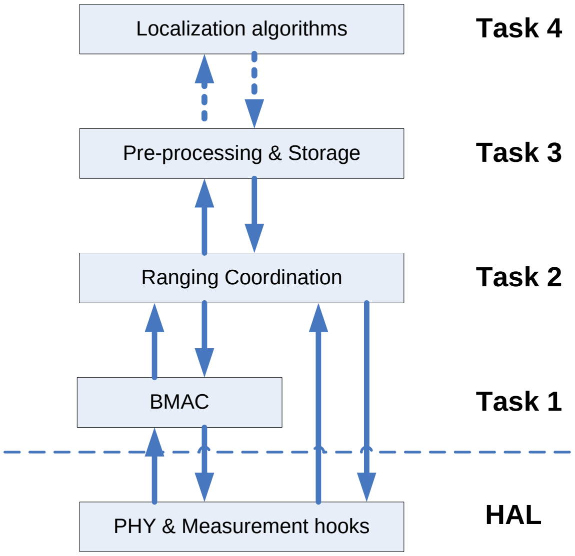

Figure 1.4. Measurement Protocol Stack

sions from different nodes do not interfere with each other. This is done by applying an arbitration protocol that determines the ordering that nodes transmit their reference signals. Finally, an additional layer above the ranging coordination layer performs an initial filtering on the measurement data and classifies the data into a set of appropriate data structures in which they are forwarded to the location computation engine, or in the case of a distributed implementation they are passed to the localization algorithm running in a separate task.

Location Computation Engine

3. Location Discovery Algorithms

The two main location discovery algorithms devised for the Smart Kindergarten environment are *iterative multilateration* [4] and *collaborative multilateration* [6]. Both algorithms are primarily used for localizing the ceiling beacons during the bootstrapping phase. In iterative multilateration, a node localizes itself if it can receive reference signals from three other nodes that already know their locations. This allows the node with unknown location to estimate its position and then provide reference signaling for other nodes with unknown locations to estimate their locations. Given sufficient densities and an initial number of beacon nodes, iterative multilateration will localize all the nodes in the network.

Iterative multilateration suffers from two main problems. First, it can get stuck in regions of the network that do not have sufficient beacon densities. If the process reaches a point where none of the nodes with unknown locations can obtain reference information from at least three beacon nodes, localization can get stuck. To alleviate this problem, we developed collaborative multilateration. Collaborative multilateration enables multiple nodes to share their beacon information and jointly estimate their positions using the beacon node locations and a set of inter-node distance measurements.

Collaborative multilateration uses a three-phase process. During the first phase, the nodes compute a set of initial estimates by forming a set of bounding boxes around the nodes. The nodes then organize themselves into over-contrained groups in which their positions are further refined using least squares. The refinement phase can use one of two possible computational models, centralized and distributed. The centralized computation model requires beacon positions and distance measurement information for the entire network. The distributed computation model is an approximation of the centralized computation model in which each node is responsible for computing its own location by exchanging information with its one-hop neighbors. The key attribute that makes distributed collaborative multilateration possible is the *in-sequence* execution within an over-constrained set of nodes. In distributed collaborative multilateration, each node executes a multilateration using the highly uncertain position estimates of its neighbors, and the corresponding distance to each of its neighbors. The consistent sequence of execution of multilaterations among each nodes results in the formation of a global gradient that allows each node to compute its own position estimate locally by following a gradient with respect to the global constraints.

Location Computation Engine

The location computation engine is the software the fuses distance measurements to derive node locations. This software is responsible for filtering the measurements and selecting the set of beacon nodes that will give the best position estimate based on a set of geometric criteria. Once the positions are computed they get stored in the central system database for further processing.

4. Conclusions and Lessons Learned

The design of the Smart Kindergarten localization infrastructure was a challenging and rewarding experience. Although the initial system achieves some level of functionality, this localization system is still under constant refinement. The design of customized wireless sensor node provided valuable insight in the design of low power platforms using off-the-shelf components. This process also revealed several sources of overhead in the low volume production of experimental systems in an academic environment. The purchasing of components from multiple vendors and unexpected lead-times in part delivery introduced unforeseen delays in the development our platforms. The experience derived from this project is driving the development of a new generation of localization platforms and algorithms. For more details and the latest developments in this platforms we refer the reader to our websites [1] and [8].

References

[1] ENALAB Website, http://www.eng.yale.edu/enalab

[2] S. Park, I. Locher, A. Savvides, M. B. Srivastava, A. Chen, R. Muntz, S. Yuen, *Design of a Wearable Sensor Badge for Smart Kindergarten*, Proceedings of the International Symposium on Wearable Computing, October 2002

[3] P. Steurer and M. B. Srivastava, *System Design of Smart Table*, Proceedings of IEEE International Conference on Pervasive Computing and Communications, PerCom 2003, March 2003.

[4] A. Savvides, C-C Han, and M. Srivastava. *Dynamic Fine-Grained Localization In Ad-Hoc Networks Of Sensors*. Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking, July 2001. pp. 166 - 179

[5] A. Savvides and M. B. Srivastava, *A Distributed Computation Platform for Wireless Embedded Sensing*, Proceedings of International Conference on Computer Design (ICCD), Freiburg, Germany, September 2002

[6] A. Savvides, H. Park and M. B. Srivastava, *The n-hop Multilateration Primitive for Node Localization Problems*, Proceedings of Mobile Networks and Applications, 8, 443-451, 2003

[7] A. Savvides, W. Garber, R. L. Moses and M. B. Srivastava, *An Analysis of Error Inducing Parameters in Multihop Sensor Node Localization*, to appear in IEEE Transactions on Mobile Computing, 2004.

[8] Smart Kindergarten Website http://nesl.ee.ucla.edu/projects/smartkg

[9] M. Srivastava, R. Muntz and M. Potkonjak, *Smart Kindergarten: Sensor-based Wireless Networks for Smart Developmental Problem-solving Environments*, Proceedings of the ACM SIGMOBILE 7th

Annual International Conference on Mobile Computing and Networking, Rome, Italy, July 2001

[10] A. Chen, R.R. Muntz, S. Yuen, I. Locher, S. Park. M.B. Srivastava, *A Support Infrastructure for the Smart Kindergarten*, IEEE Pervasive Computing, vol. 1, no. 2, April-June 2002, pp. 49-57

[11] Crossbow Website, http://www.xbow.com