

Power-efficient Cyber-physical Late Arrival Detection with Smartphones

Danny Guymon, Brian Dougherty, Jules White
Dept. of Electrical and Computer Engineering
Virginia Tech
Blacksburg, VA, USA
Email: {dguymon, briandpd, julesw}@vt.edu

Mani Golparvar-Fard
Dept. of Civil and Environmental Engineering
Virginia Tech
Blacksburg, VA, USA
Email: golparvar@vt.edu

Abstract—Late arrival to jobs and meetings is a critical problem in a number of domains, such as building construction. Smartphones are a promising platform for detecting late arrivals and notifying affected parties due to the variety of sensors they possess, persistent Internet connection, and access to a user’s calendar data. This paper presents a power-efficient cyber-physical approach to indoor/outdoor late arrival detection, called LifeSaver. The paper provides the following three contributions to the study of late arrival detection and notification with smartphones: 1) it presents a cyber-physical approach for using accelerometer data to predict when a user is moving and make more power efficient localization decisions; 2) it shows how Bluetooth beacons can be used to provide a simplified meeting room localization mechanism for smartphones; and 3) it presents empirical results showing that the cyber-physical LifeSaver late arrival detection approach uses 43% less power compared to lateness detection approaches that rely on standard localization strategies.

Keywords—cyber-physical systems, smartphones, mobile computing, indoor/outdoor localization, 1

I. INTRODUCTION

Late arrival to a meeting or job is an important and prevalent problem in a large number of domains, such as shipping, disaster response, and construction. Unexpected late arrivals can increase shipping costs, delay dependent construction activities, cause customer dissatisfaction [1], [2], and increase construction costs [3], [4]. Because of the high cost of being late, an important problem is the automatic advance detection of a potential late arrival and notification of affected parties.

Late arrival detection is inherently a cyber-physical computing problem. First, an individual’s physical location must be determined and compared to the cyber-information dictating a user’s scheduled destinations throughout the day. If the person’s physical location does not match the cyber-destination information, then a cyber-physical computation must be used to determine the time it will take the person to get from his/her physical location to the destination listed in the cyber-information. Next, the amount of time the person needs to get to the destination must be used to determine if he/she will be late to the start of the job or task at the destination. If the person is unable to reach the destination on time, then a cyber-notification, such as a text message to

a person at the destination, providing his/her expected arrival time is needed so that alternate plans can be made.

In order to perform these cyber-physical steps, a system must be able to localize a user; have access to geographic information, such as maps; be able to read calendar information for the user’s appointments; and have one or more mechanisms for notifying affected parties. Smartphones, such as the iPhone and Android phone, are a promising platform for detecting late arrivals and notifying affected parties because they provide a package of sensors that travel with their user and measure physical phenomena, have access to a user’s cyber calendar and contact information, and have the ability to send cyber-notifications to contacts through text messaging, email, or other means. Moreover, smartphone platforms possess a variety of sensors that can be used for geolocation of a user in the physical world. For example, the iPhone’s sensors include an accelerometer, which monitors the phone’s physical acceleration in the x, y, and z-direction, GPS, which tracks the phone’s location on the earth, and a Bluetooth adapter, which can be used to detect other nearby Bluetooth devices or navigation beacons.

Open Problem: Power-efficient Indoor/Outdoor Smartphone-based Late Arrival Detection and Notification. Despite the enormous potential for using smartphones as a cyber-physical late arrival detection platform, little existing research has investigated approaches for detecting late arrivals on smartphones. A number of researchers have investigated approaches for user localization with smartphones, which is a major component of detecting late arrivals [5]–[11]. However, none of these approaches has addressed all of the cyber-physical challenges related to late arrival detection on smartphones.

There are three key cyber-physical issues that have not been addressed for late arrival detection on smartphones. The first issue is that localization sensors on a smartphone, such as GPS, consume significant power and can exhaust a smartphone’s battery in only a few hours. The second issue is that many indoor localization approaches rely on external hardware, pre-analysis of wireless signals in a building, or other environment conditioning steps to determine physical information. These setup steps may not be economical or possible to perform at all possible locations a user may

travel through. The third issue of late arrival detection is the requirement for both outdoor and indoor localization, which requires an approach that can seamlessly transition between the different localization approaches used in each environment.

Solution Approach \Rightarrow A Cyber-physical Approach for Detecting User Lateness. To fill this gap in research on smartphone-based cyber-physical late arrival detection, we have developed a smartphone-based late arrival detection and notification approach called LifeSaver. LifeSaver uses a combination of smartphone sensors to determine a user’s physical context, more intelligently power cycle localization sensors, and significantly reduce the power required to detect a late arrival. Moreover, LifeSaver uses a key insight from the domain of late arrival detection, that indoor localization only needs to determine if the user is in a meeting room, to simplify the indoor localization component of late arrival detection.

A key attribute of the LifeSaver approach is that it requires significantly less power to detect late arrivals than is required if existing localization approaches are relied upon. In Section V we present empirical data that we gathered from experiments that tested the operation of LifeSaver and compared its power consumption to late arrival detection with existing localization approaches. Our empirical results show that LifeSaver uses up to 42% less power than late arrival detection with existing localization approaches. In addition, our results show that LifeSaver is capable of seamless indoor/outdoor localization.

This paper provides the following contributions to the study of user late arrival detection with smartphones:

- We present a novel localization technique for mobile devices that is capable of determining device location seamlessly between indoor/outdoor environments without requiring additional preconfigured hardware.
- We provide a formal methodology for using smartphone sensors to acquire physical data, such as device location and nearby device IDs, and relating it to cyber model that contains schedule data, such as appointment time and location, to detect late arrivals.
- We present empirical results showing the impact of GPS and Bluetooth localization on the power consumption of mobile devices.
- We show that the Lifesaver technique can delay performing localization based on sensor data, thereby conserving device power by reducing GPS and Bluetooth activity.

The remainder of this paper is organized as follows: Section II describes a practical scenario clarifying the main objectives of this research effort, which we use as a motivating example throughout the paper; Section III discusses the challenges that we faced when addressing the two issues from above; Section IV covers our solution to the overall problem; Section V presents empirical results from the

benchmark experiments we ran; and Section VII presents concluding remarks and lessons learned.

II. MOTIVATING EXAMPLE

As a motivating example to illustrate the challenges of late arrival detection, we use a scenario from the construction industry. One type of construction late arrivals that we focus on in this paper is late contractor arrivals [12], which can cause significant delays and cost increases on a job site.




Figure 1. Example Schedule of Dependent Construction Tasks

A critical issue on construction sites is that many construction tasks directly depend upon the completion of prior tasks. For example, as seen in Figure II, initial framing of walls is dependent upon the completion of the foundation. Often, however, contractors, such as plumbers will arrive late, delaying work on a construction task and leaving workers assigned to certain dependent tasks idle.

If it was possible to predict when a contractor was going to be late, workers could be assigned to alternative tasks that are not directly dependent upon the late contractor. Studies have shown that in a sample of 215 completed commercial and industrial construction projects, 62% of them were late and 88% of these late projects had late contractors. A majority of the projects’ timelines to completion were extended by a month, but in some extreme cases a six-month extension was needed to finish the project [4]. A one to six month increase in construction time can add significant expense to a construction project [4].

III. CHALLENGES OF CYBER-PHYSICAL LATE ARRIVAL DETECTION

Although smartphones are a promising platform for user late arrival detection, Smartphone late arrival detection must address challenges related to the limited battery capacity of smartphones and issues associated with indoor localization. Because phones are usually only charged at the night when the user is asleep and many meetings occur indoors, power efficiency is critical for developing a smartphone late arrival detection system. Moore’s Law defines the rate of advancement for certain computing attributes, such as processor speed and storage density. However, battery power density has not increased nearly as fast as processor speed [13]. The remainder of this section explores three key challenges of

power efficient indoor/outdoor user late arrival detection on smartphones.

A. Challenge 1: Power Efficiency

GPS is the primary mechanism for precise user localization using smartphones. A key challenge of detecting late arrivals using smartphones is that smartphones have limited battery capacities that can quickly be exhausted using smartphone localization sensors, such as GPS. For example, the iPhone 4, the Droid X and the Nexus One have 1420 mWh, 1570 mWh and 1400 mWh battery capacity, respectively. A GPS sensor, alone, consumes roughly 140 mW of power when enabled [14], so neglecting all other power consumption on the phone, the battery is expected to last 10 - 11 hours for the iPhone, Droid X and Nexus One when GPS is turned on. In reality, an active GPS sensor also requires CPU and other power consumptive resources. When these other supporting resources are considered, battery life may be as short as 3-4 hours when GPS is enabled. Because users normally charge their phones at night while they sleep, running GPS all of the time is not a feasible approach for localization.

In the context of our construction scenario, described in Section II, the power consumption challenge becomes an issue when a late arrival detection application on the worker's smartphone uses up too much battery power throughout the day. The average construction worker's shift is eight hours, so in order to benefit the construction domain, an algorithm must not exhaust a smartphone's battery in less than this amount of time. Once the battery is depleted, late arrival detection of the smartphone's owner comes to a halt until the battery is recharged. Since most people recharge their phones at night while sleeping, they are less likely to recharge their phone during the day at work on a construction site. Section IV describes how we address this challenge by using a smartphone's accelerometer to determine user context an intelligently enable and disable GPS and other localization sensors to save power..

B. Challenge 2: Indoor Localization to a Specific Room Without External Hardware

The second challenge of smartphone late arrival detection is indoor localization without external hardware, such as Bluetooth beacons, WiFi hotspots and GPS repeaters. Figure III-B shows the inadequacy of outdoor GPS localization in an indoor environment due to physical barriers blocking signals. Indoor localization is the process of monitoring a person's location in a building. For late arrival detection, indoor localization must be accurate enough to determine if a user is in a specific meeting room.

Many previous indoor localization approaches have relied upon preconfigured external hardware, such as custom-made Bluetooth beacons or wireless routers and GPS repeaters, in order to track a user's location indoors. Each of these



Figure 2. GPS Signals Blocked in a Construction Environment

approaches requires preconfiguration of the indoor environment or additional hardware. For example, in WiFi localization approaches, the cell signals through the building must be mapped and the area of overlap between access points measured. In Bluetooth localization approaches, Bluetooth beacons are laid about a floor plan requiring custom-made containers to focus the signal for increased detection range. In accelerometer-based indoor localization, floor plans are downloaded to the mobile device and the user specifies a starting location. Extra hardware increases the cost and preconfiguration of the environment and may not be practical due to the lack of robustness and portability to other environments.

Referring back to our motivating example in Section II, construction workers already have a great deal of equipment to use and keep track of, so the introduction of fragile and expensive headsets or beacons for late arrival detection and localization would not be ideal in a construction site environment. Also, construction sites are dynamically changing and so approaches that rely on measurements of external physical phenomena, such as WiFi signals, would have to be continually recalibrated as the building is built and the signal propagation properties on the site changed. Section IV-E describes how we address this challenge by using the smartphone's Bluetooth adapter for indoor localization without preconfigured external hardware.

C. Challenge 3: Seamless Indoor/Outdoor Localization Transitions

The third challenge of user late arrival detection is the difficulty in accomplishing seamless indoor/outdoor localization transitions. Seamless transitions are needed because they ensure continual location tracking regardless of the

changes in the surrounding environment. While there have been many research efforts focused on indoor and outdoor localization, there are very few that seek to accomplish seamless handoff between the two localization methods. One solution, called Streamspin, uses a combination of GPS and WiFi hotspots to achieve seamless indoor/outdoor localization handoff on a mobile PDA device [15]. Even though Streamspin accomplished seamless indoor/outdoor localization transitions, the algorithm depended upon GPS, required preconfigured WiFi maps and consumed 35% of the device’s battery per hour, which would exhaust the battery too quickly for many late arrival detection domains. We can see that this challenge is not mutually exclusive from the other two, thus making seamless indoor/outdoor localization that much more difficult to address. One may be able to accomplish seamless localization handoff, but the accomplishment comes at the expense of power consumption and simplicity in design.

Construction workers may walk in and out of the building during construction. In order to be able to track the late arrivals of workers or of specific job-related tasks, it is necessary to accomplish seamless indoor/outdoor localization handoff in a user late arrival detection algorithm. For example, a plumber is going to work on several construction projects in the same day. On some sites, he may be doing plumbing work inside a partially finished structure where the GPS fix on his location may be periodically lost and found. At other parts of the day, he may be installing pipe as part of the foundation for a building and has a continual GPS signal. A late arrival detection algorithm with seamless localization would be able to keep track of where the plumber was at any point of the day in order to determine if he was running late to a job based upon his work schedule. Section IV-D describes how we address this challenge by developing a GPS-based algorithm for detecting indoor/outdoor localization handoff.

IV. SOLUTION

LifeSaver is a cyber-physical approach to late arrival detection with smartphones. LifeSaver uses a decision tree to drive a state-machine using physical events as inputs for cyber decisions on the smartphone. LifeSaver addresses the three challenges described in Section III: (1) low power consumption with an accelerometer-based algorithm to determine when to start localization, (2) relative indoor localization without extra preconfigured hardware beyond what’s available on the smartphone, and (3) seamless indoor/outdoor localization handoff with a GPS-based algorithm. These challenges are overcome by LifeSaver through the use of two critical power efficiency and indoor localization techniques.

Context-based Power Cycling. As described in Section III-A, power efficiency is a major challenge in smartphone-based late arrival detection. In order to conserve

battery power, LifeSaver uses a context-based approach to power GPS and Bluetooth sensors on and off.

Indoor Localization Approach. Indoor meetings require a smartphone to localize a user to a specific room to determine if the user has entered the meeting room. However, when a user is inside the building containing a meeting room, it is typically not necessary to know exactly which room they are in. Once a user is in the building where a meeting is taking place, knowing whether the user is in the designated meeting room or not is sufficient. In most cases, the travel time within a building is insignificant.

Because it is not necessary to know exactly where a user is within the building of a meeting location, LifeSaver uses a simple Bluetooth rendez-vous beacon to determine when a user has reached a meeting room. When the first participant in a meeting reaches the designated meeting room, they activate a Bluetooth beacon that broadcasts a unique ID assigned to each meeting. When the devices of other participants in the meeting receive the Bluetooth broadcast, they automatically assume that they have reached the meeting room. This localization approach is similar to prior Bluetooth-based indoor localization approaches, but does not require the installation of external hardware. The approach relies on the short range of Bluetooth communications to determine if a user has reached a meeting room.

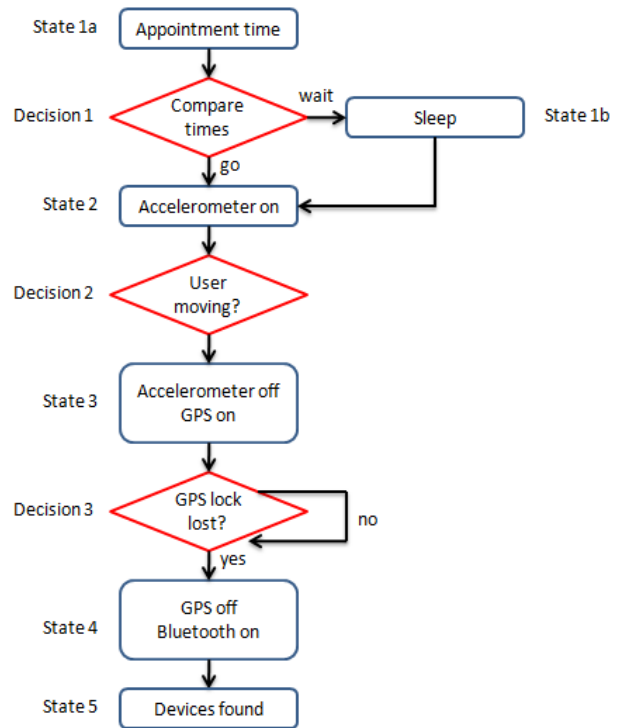


Figure 3. LifeSaver’s State-machine and Decision Tree

A. Formal Model of LifeSaver’s Cyber-physical Late Arrival Detection Algorithm

Late arrival detection requires the processing of physical stimulus data from a smartphone’s sensors, estimation of a user’s location/context, and determination of whether a user will reach a meeting location by its start time. LifeSaver’s late arrival detection algorithm is based on a 6-tuple state representation that encodes physical stimulus data and location/context estimation:

$$Ls_k = \langle t_k, D(t_k), C(t_i), \phi(t_i), \rho(t_i), \sigma(t_i) \rangle \quad (1)$$

where:

- $t_k \subset T$ is the time of the future scheduled meeting M_k , from the user’s calendar. The set of all future meeting times is T .
- $D(t_k)$ is the latitude and longitude of the meeting, M_k , at time t_k .
- $C(t_i)$ is the last known latitude and longitude of the device at time t_i , where $t_i < t_k$.
- $\phi(t_i)$ is a flag indicating whether or not the device is believed to be indoors or outdoors at time t_i . If the device is indoors, $\phi(t_i) = 1$.
- $\rho(t_i)$ is a flag indicating whether or not the device is believed to be currently moving. If the device is believed to be moving, $\rho(t_i) = 1$.
- $\sigma(t_i)$ is a flag indicating whether or not the device has received a Bluetooth beacon signal from the device of a user who has arrived at the meeting location. If the signal has been received, $\sigma(t_i) = 1$.

Late arrival detection is modeled as a function, $Late(Ls)$, where the determination of whether or not a user is late is based on the device’s current location, the estimated time to reach a future meeting location, and whether or not a Bluetooth beacon signal has been received for the given meeting:

$$L(M_k) = \begin{cases} t_i + \Delta(D(t_k), C(t_i)) \leq t_k + \zeta \wedge \alpha = 0 & 0 \text{ (a),} \\ t_i + \Delta(D(t_k), C(t_i)) + \epsilon \leq t_k + \zeta & 0 \text{ (b),} \\ \sigma(t_i) = 1 & 0 \text{ (c)} \\ t_i + \Delta(D(t_k), C(t_i)) + \zeta > t_k & 1 \text{ (d),} \\ otherwise & 1 \text{ (e)} \end{cases} \quad (2)$$

where:

- M_k is a future meeting.
- t_i is the current time on the user’s device.
- $\Delta(D(t_k), C(t_i))$ that estimates the travel time from the user’s current last known location to the meeting location.
- ζ is a lateness threshold factor that determines how many minutes past a meeting’s start time the user can arrive without being considered late. This factor can

be varied based on meeting importance, local customs, and other factors.

- α is a flag that indicates whether or not arriving at the meeting destination is sufficient to have arrived at the meeting or if the attendance requires reaching a specific room.
- ϵ is a time allowance to walk from the entrance of a building to a meeting room. We use a fixed value for ϵ . In practice, ϵ will vary based on the building, however we assume that in most cases ϵ is negligible and will not make a substantial impact. If needed, this value could be dynamically assigned based on the building and room that the meeting is scheduled in.

The LifeSaver late arrival detection algorithm considers a user to be on-time in three primary cases. In case (a), the meeting is not taking place indoors or within a specific room ($\alpha = 0$) and the current time plus the travel time from the user’s current location to the meeting location is less than the meeting start time of the meeting (t_k). Case (b) focuses on situations in which the user has not reached the destination location yet and the meeting is occurring in a specific room. In this situation, if the user can reach the destination location and walk to the room before the meeting start time, the user is considered on-time. The final on-time case, case (c), marks the user as being on-time because a Bluetooth beacon signal for the meeting has been received from another attendant’s smartphone. In all other cases, the user is expected to arrive late to the meeting.

B. LifeSaver’s State-based Sensor Activation Model

The first high-level component of LifeSaver is the state-machine, which is represented by the blue rectangles in Figure IV. The red diamonds are the decisions made by the decision tree which we will discuss later. The state-machine has five states: (1) appointment setup/application sleep, (2) accelerometer enabled, (3) GPS enabled, (4) Bluetooth enabled and (5) destination reached.

State 1a: Appointment setup. State 1a involves the start up of the application and the user scheduling an appointment in the phone by use of a time widget, such as the prototype time widget used for our experiments and shown in Figure IV-C. Once the user selects the time of the appointment, the meeting data is stored for later use by the decision tree in the cyber model. Extraction of geolocation data from natural language was not a component of the research presented in this paper. We assume that appointments with GPS coordinates are available from a service, such as Google Calendar. In our experiments, we used a destination specification widget to capture location data for appointments.

State 1b: Sleep. Transitioning to State 1b is determined by the decision tree in Decision 1 as seen in Figure IV. If the decision tree determines that the application must sleep, then the main thread goes to sleep until the current time is

within γ minutes of the appointment time selected in State 1a. For example, if the user scheduled an appointment for 11:30 AM, $\gamma = 5min$, and the current time is 11:05 AM, then the application would sleep for twenty minutes and then wake back up at 11:25 AM. In our experiments, a buffer time of five minutes was shown to effectively detect late arrivals in most cases without excessively consuming battery power.

State 2: Accelerometer on. An important energy saving step of LifeSaver is that it attempts to only localize the user if movement is detected. The LifeSaver formal model captures perceived motion in $\rho(t_i)$. A smartphone's accelerometer is used to calculate $\rho(t_i)$.

In State 2, the application is woken up and the smartphone's accelerometer is enabled to record physical acceleration experienced by the phone caused by the user's movement. For example, if the user is carrying the phone as he/she walks around, the accelerometer will record the changes in acceleration in the x, y, and z-direction. This physical data is used by the decision tree to enable the localization sensors. LifeSaver uses a threshold function based on the total movement recorded over a configurable duration, β . If the recorded motion exceeds a threshold, the smartphone is considered to be in motion and $\rho(t_i) = 1$. Otherwise, $\rho(t_i) = 0$.

State 3: Accelerometer off/GPS on. At regular intervals, LifeSaver must attempt to determine the device's current location and update $C(t_i)$. LifeSaver uses a smartphone's GPS sensor to derive a location for $C(t_i)$. GPS is also used to make a determination of whether or not the device is indoors and assign $\phi(t_i)$.

In State 3, the accelerometer is disabled and acceleration data is no longer recorded since the decision tree has decided that the user is walking around. GPS is enabled on the smartphone to establish an initial outdoor location fix. Once a location fix is established, location updates are sent to the smartphone periodically as the user walks to update $C(t_i)$. When a new location update is sent to the phone, a background count down timer is reset and initiated by the decision tree. If this timer reaches zero, then the decision tree determines that the user is no longer outside and has entered a building ($\phi(t_i) = 1$).

State 4: GPS off/Bluetooth on. For indoor meeting locations, a Bluetooth rendez-vous scheme is used to detect when a user has reached the meeting location. The first meeting participant to enter the meeting location enables a Bluetooth beacon on their device that broadcasts the unique ID of the meeting. When other devices discover this Bluetooth beacon, they update $\sigma(t_i)$ to 1. The function $\sigma(t_i)$ serves as an indicator that the user has been localized to the meeting room.

State 4 begins the Bluetooth discovery process to update $\sigma(t_i)$. This state is reached when the phone no longer receives location updates causing the count down timer from State 3 to reach zero. The smartphone's GPS is disabled and

the Bluetooth adapter is turned on. The adapter is put into discovery mode to detect the meeting's Bluetooth beacon.

State 5: Devices found. In State 5, if the correct Bluetooth broadcast is discovered, then the user is at the destination.

C. Address Challenge 1: Conserving Power with Motion-based Sensor Power Cycling



Figure 4. Main Screen of the LifeSaver Implementation

Rather than immediately turning on the GPS to attempt to acquire a location fix, we took a more power conservative approach. When the application is sleeping, power consumption is at a minimum. Once the application wakes up, the accelerometer is activated to determine if the user is moving before the GPS is turned on to begin outdoor localization. The accelerometer provides data on the changes in the phone's acceleration, and we equated roughly twenty seconds of continual accelerometer updates as evidence of the user walking around, assuming that he/she is heading towards the destination by exiting the building. Once the accelerometer detects these physical changes, the data is stored in the cyber application so that the decision tree can turn off the accelerometer and then enable GPS to establish an initial outdoor location fix.

We used the accelerometer in this manner to shorten the time we needed the GPS to establish an initial location fix. By default, the GPS on the smartphone does not have a timeout, meaning, if the phone's GPS is on and the phone is sitting on a desk somewhere inside a building, the phone will continually attempt to connect to a satellite to establish a location fix. Since this operation is costly in terms of power

consumption, delaying activation of the physical sensors, such as GPS and accelerometers, until the current time is within five minutes of the appointment time stored in the cyber model, can substantially reduce power consumption.

The user confirms the appointment by pressing the confirm button which then transitions the state machine to State 2 to put the application to sleep.

D. Addressing Challenge 2: User Arrival Detection with Bluetooth

second challenge described in Section III-B. The second challenge involves the difficulty in providing accurate indoor localization down to a specific room without using pre-configured, external hardware. We used the Bluetooth adapter already built into the smartphone to achieve relative indoor localization of the user. By discovering recognized Bluetooth devices in the nearby area based upon device name and MAC address, we're alleviated from the need of precise localization, which is hard to do without costly additional hardware. When the adapter sees other Bluetooth devices, the decision tree determines if they are friends/co-workers

E. Addressing Challenge 3: Detecting Outdoor-to-Indoor Transition

Previously, in Section III-C we discussed the challenge of seamless indoor/outdoor localization. The decision tree transitions the state-machine to State 3 by enabling the smartphone's GPS after roughly twenty seconds of movement captured by the accelerometer. GPS serves two purposes once it is activated. First, GPS can effectively localize targets when outdoors. Again, we assume that the user has to walk outside in order to reach the destination, so once he/she steps outside, GPS can establish a position lock and provide position updates periodically.

The second function of the GPS is to provide the decision tree with the point in time when the user leaves the outdoors and enters the destination building. We accomplished this by using a timeout counter, which is reset every time the smartphone receives a position update from the satellite. The two scenarios of when the phone would cease to receive these notifications is either when the phone is no longer outside or if the user stops walking.

In both of these events, the timeout counter would reach zero, thus signaling the decision tree to disable GPS and transition to State 4 by enabling the Bluetooth adapter, as seen in Figure IV. This scheme provides a seamless transition from outdoor localization to indoor localization without any manual input from the user and without the user even knowing.

V. RESULTS

A. Experimental Platform

We tested the functionality of LifeSaver by running real-world tests with the application. We were interested in two

things: (1) how well LifeSaver provided a solution to the three challenges described in Section III, and (2) how much power could be saved in comparison to two other localization solutions.

Before conducting our experiments, we made three assumptions about our testing schedule. The first assumption was that the user starts inside a building once the application is started. Even though the application will work if the user starts outside, we kept this constant through the tests.

The next assumption we made was that the user would walk around outside at some point. LifeSaver's current algorithm needs an initial GPS location lock since acquiring the first location update is an important step for the decision tree to drive the state-machine, seen in Figure IV. Our last assumption was that the user's destination is another building so that we could test the seamless indoor/outdoor localization.

With these assumptions in place, we began real-world tests with the application installed on a Droid smartphone with Android version 2.2.2. The testing schedule was set up to parallel the state transitions described in Section IV.

The user begins by sitting in an office and selects a time for the next appointment from the main activity screen. Figure IV-C is a screenshot of the main activity screen. Once the user selects the appointment time, a secondary activity is initialized and awaits for user confirmation of the scheduled appointment. Note that all the state information is provided on the screen keeping the user updated of any state transitions if he/she so chooses to watch the application.

The application sleeps until the current time is within five minutes of the appointment time, then the accelerometer is turned on. The accelerometer-based algorithm for turning on localization sensors was successful in activating the phone's GPS after the user walked from indoors to outdoors. This algorithm is also a contributor LiveSaver's relatively low power consumption.

We tested LifeSaver's seamless localization hand-off feature by monitoring how the GPS-based algorithm handled the lost location fix when the user moves from outdoors to indoors. When the user was walking around outside, the phone received periodic location updates from a satellite. When the user moved from outside to inside, GPS was disabled and Bluetooth was then immediately enabled as desired.

We tested LifeSaver's relative indoor localization feature by placing other Bluetooth devices in an indoor environment and had the user walk around with the phone in Bluetooth discovery mode as described in State 4 in Section IV. The application was able to detect all the physical Bluetooth devices and then confirm the user's arrival at the destination by comparing the device names and MAC addresses to those already recognized by the phone.

Once we confirmed that the application worked properly, we prepared quantitative experiments to compare the power

consumption of our application to that of single-sensor localization approaches the use a single sensor or adapter for the same duration of time it took our state-machine to reach the last state.

The two sensors in consideration were also used in our application: GPS and Bluetooth. The GPS application continuously displayed position updates to the screen every second and the Bluetooth application kept the smartphone in discovery mode and displayed all discovered devices to the screen. For both of these applications, we ran them for the same amount of time it took for LifeSaver to go from start to finish. We describe the results from these experiments below.

B. Experiment 1: Pure GPS Late Arrival Detection vs. LifeSaver

The first experiment was designed to compare the battery usage on the smartphone of our application versus that of one with just the GPS activated. We wrote code that enables the smartphone’s GPS and continuously displays the GPS coordinates of the user to the screen. To monitor battery consumption, we used a battery monitoring widget available on the Android market that shows the battery consumption of each running process as a percentage. Typical processes that contribute to battery usage of the smartphone include the backlight for the display, call standby and WiFi.

Before running our application and the GPS application, we rebooted the phone to clear all battery usage history to establish a common starting point. We ran our application for six minutes, which was the typical amount of time needed to go from start to finish. We reset the system to clear the battery history and then ran the GPS application for the same amount of time.

Hypothesis: Motion-based Power Cycling is More Power-efficient than Continually Using GPS. Our hypothesis was that running GPS outside would have a higher battery usage percentage than LifeSaver and running GPS without acquiring an initial position lock by staying indoors would yield an even higher percentage than both outdoor GPS and LifeSaver. We assumed that the constant overhead in attempting to establish a lock would consume more power than walking around outside with valid position updates. When GPS cannot establish an initial position lock, it continues to do so without any timeout.

Experiment 1 Results: Pure GPS Late Arrival Detection vs. LifeSaver Power Consumption.

As seen in Figure V-B, keeping GPS active during the six-minute test consumed more power than LifeSaver. Running the GPS without an initial position lock constituted 68% of the smartphone’s battery use and GPS with position lock was slightly lower at 62%. These figures are more than double the 30% battery usage level of LifeSaver.




Figure 5. Battery Usage of LifeSaver vs. Pure GPS Late Arrival Detection

C. Experiment 2: Pure Bluetooth Late Arrival Detection vs. LifeSaver

This experiment was designed to determine whether keeping the Bluetooth adapter on the smartphone in discovery mode consumes more power than LifeSaver. When the adapter is put into discovery mode, it sends out discovery requests for roughly twelve seconds to find nearby Bluetooth devices that the requesting smartphone can pair and connect with. The twelve second duration for discovery is the default amount, but our test application kept the phone in discovery mode for a full six minutes.

Hypothesis: Bluetooth Consumes More Power than Power Cycling GPS and Bluetooth Based on Movement. Our hypothesis was that Bluetooth would consume more battery than LifeSaver. We expected to see Bluetooth’s power consumption even higher than that of the GPS solution because of the comparatively more overhead involved in Bluetooth discovery.

Experiment 2 results: Pure Bluetooth Late Arrival Detection vs. LifeSaver Power Consumption.

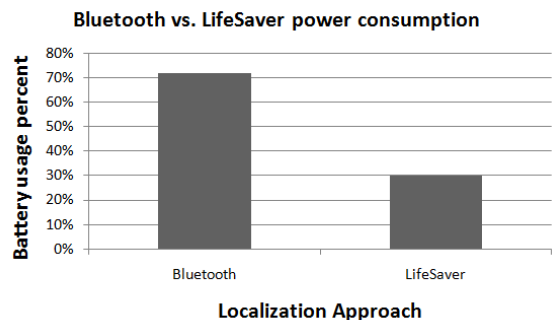


Figure 6. Battery use of LifeSaver vs. Pure Bluetooth Late Arrival Detection.

As seen in Figure V-C, running the smartphone’s Bluetooth adapter in discovery mode for six minutes consumed even more battery usage than both the GPS application we wrote and LifeSaver.

D. Analysis of Results

Considering all the power consumption experiment results, LifeSaver's power consumption is sufficiently less than the other localization approaches. Our solution contributed to 30% of the phone's battery use during the six-minute duration of keeping the application running. We saw dramatic increases in battery use by the GPS application when the user was walking around outside with a position lock and even more when the user was walking around inside without any initial position lock.

Even though GPS is a useful hardware tool for localization, using GPS for long durations of time on a mobile device consumes a considerable amount of power and shortens the life of the phone. With LifeSaver, we trim down the impact of GPS by implementing the accelerometer-based algorithm for enabling GPS discussed in Section IV.

The Bluetooth test application required the most battery power to discover other Bluetooth devices in the area. Walking around with the phone continually discovering devices does consume a great amount of power, but LifeSaver uses this feature sparingly to determine user arrival to the destination.

This data is important because it sets the bar for optimizing our application even further in the future. Some optimizations we are interested in implementing involve making our approach more robust to accommodate even more late arrival scenarios, which we will discuss later. We designed some baseline benchmark experiments to rate the performance of our application with the metric of power consumption and our LifeSaver yielded favorable results.

VI. RELATED WORK

Accurate methods of indoor and outdoor localization comprise a significant portion of the issues in detecting late arrivals. This section provides a taxonomy of related work. Related works are divided into four categories: indoor WiFi localization, indoor accelerometer with floor plan localization, indoor Bluetooth localization and outdoor GPS localization.

Indoor WiFi Localization. Indoor WiFi localization is the use of WiFi hotspots and WLAN mapping to achieve indoor location tracking down to a relatively small distance, such as between offices in a building. One of the solutions that uses indoor WiFi localization is Streamspin, which is a mobile localization application used on PDAs [15]. The application compares the signal strengths of the wireless routers in the building to a precompiled database of signal strength-distance pairs stored for each wireless point. As the user walks around, the application on the PDA takes readings of the emitting WiFi signal and does a comparison with the data in the database. LifeSaver, however, does not require additional preconfigured hardware to accomplish indoor localization. Our solution uses the Bluetooth adapter to

discover other nearby recognized smartphones to determine if the user is in the correct location.

Lifetag is another WiFi-based localization solution that is very similar to Streamspin [8]. Instead of a PDA, the user carries a WiFi sensing device that only collects access point IDs and signal strength at the moment of collection and then compares these pairs to those in a precompiled database to map where the user is at in an urban environment or academic building, for example. LifeSaver does not need a map of the indoor environment to accomplish indoor localization, but rather uses the Bluetooth information of nearby smartphones to establish localization.

Indoor Bluetooth Localization. Similar to LifeSaver's use of Bluetooth as a means of location tracking, Cheung et al. accomplish Bluetooth location tracking using beacons and a smartphone. [7]. Custom-made short-range Bluetooth beacons were constructed and all set to discovery mode so that the smartphone navigating through the maze of beacons could keep track of its current location.

Indoor Accelerometer with Floor Plan Localization. Some researchers have even explored the possibility of using an accelerometer as a means of indoor localization. Using an accelerometer in this fashion does require a floor plan of the area being navigated to cross reference in order to have some level of accuracy [16]. SmartSLAM is an indoor localization solution that uses the smartphone's accelerometer and floor plan in combination with WiFi mapping. The approach uses high-level statistics to process the accelerometer data to match the floor plan and synchronize with the WiFi mapping monitored by the phone. LifeSaver does not utilize the accelerometer or maps for indoor localization. Again, LifeSaver only depends upon the discovery of other smartphones to determine if the user is in the right location. LifeSaver, though, does use the accelerometer to determine when to enable the localization sensors on the phone.

Outdoor GPS Localization There has been extensive research on the power consumption and accuracy of GPS showing the mobile GPS is both expensive in terms of power consumption and inaccurate in urban areas. Also, GPS is virtually useless in indoor settings because of the signal attenuation caused by physical barriers. These are the reasons for research efforts attempting to find other means of indoor and outdoor localization in combination with GPS that we have previously discussed [8], [15]. LifeSaver does use GPS for outdoor localization, but the solution implements a power-saving accelerometer-based algorithm to decrease the time that GPS is enabled for outdoor localization.

VII. CONCLUDING REMARKS & LESSONS LEARNED

Late arrivals at meetings and jobs can have a number of substantial negative impacts. For example, on a construction site, construction activities are typically dependent upon one another and the late arrival of a contractor can leave workers idle and lead to cost increases. Some of this negative impact

can be mitigated, however, if late arrivals can be predicted in advanced and affected parties notified.

Smartphones are a promising platform for detecting late arrivals since they are nearly always carried by their users, have access to a user's calendar data, and have a variety of sensors, such as GPS, that can be used to aid in localization. A key challenge, however, is that localization approaches require significant power and can quickly exhaust a smartphone's battery. This paper presented an approach, called LifeSaver, that intelligently power cycles localization sensors based on motion data and utilizes Bluetooth beacons as a lightweight indoor localization mechanism for detecting arrivals in meeting rooms. Our empirical results show that this combination of approaches consumes roughly 43% less power than late arrival detection with existing approaches.

An implementation of LifeSaver is available in open-source form from <http://code.google.com/clearmobile>.

REFERENCES

- [1] J. Ehrlich and W. Lilegdon, "Making better manufacturing decisions with aim," in *Proceedings of the 29th conference on Winter simulation*. IEEE Computer Society, 1997, pp. 552–558.
- [2] J. Wagner, "Maximizing Profits in a Tight Market Through Optimization."
- [3] K. Nassar, "Evaluating Contractor Performance: Application to the Dubai Construction Industry."
- [4] M. Bassett and R. Engineers, "Client time expectations and construction industry performance."
- [5] N. Ohmori, M. Nakazato, and N. Harata, "Gps mobile phone-based activity diary survey," in *Proceedings of the Eastern Asia Society for Transportation Studies*, vol. 5, 2005, pp. 1104–1115.
- [6] Y. Lui and S. Tian, "Research of indoor GPS signals acquisition algorithm," in *4th International Conference on Wireless Communications, Networking and Mobile Computing*, 2008, pp. 1–4.
- [7] K. Cheung, S. Intille, and K. Larson, "An inexpensive bluetooth-based indoor positioning hack," *Proc. UbiComp06 Extended Abstracts*, 2006.
- [8] J. Rekimoto, T. Miyaki, and T. Ishizawa, "Lifetag: Wifi-based continuous location logging for life pattern analysis," in *Location-and context-awareness: third international symposium, LoCA 2007, Oberpfaffenhofen, Germany, September 20-21, 2007: proceedings*, vol. 4718. Springer-Verlag New York Inc, 2007, p. 35.
- [9] G. Dedes and A. Dempster, "Indoor gps positioning-challenges and opportunities," in *Vehicular Technology Conference, 2005. VTC-2005-Fall. 2005 IEEE 62nd*, vol. 1. IEEE, 2005, pp. 412–415.
- [10] F. van Diggelen, "Indoor gps theory & implementation," in *Position Location and Navigation Symposium, 2002 IEEE*. IEEE, 2002, pp. 240–247.
- [11] Y. Nakazato, M. Kanbara, and N. Yokoya, "Localization system for large indoor environments using invisible markers," in *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*. ACM, 2008, pp. 295–296.
- [12] K. Van Oorschot, J. Bertrand, and C. Rutte, "Field studies into the dynamics of product development tasks," pp. 720–739, 2005.
- [13] O. S. Industries, "Optical sensors in smart mobile devices," 2010.
- [14] A. Carroll and G. Heiser, "An analysis of power consumption in a smartphone," in *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*. USENIX Association, 2010, pp. 21–21.
- [15] R. Hansen, R. Wind, C. Jensen, and B. Thomsen, "Seamless indoor/outdoor positioning handover for location-based services in streamspin," in *Mobile Data Management: Systems, Services and Middleware, 2009. MDM'09. Tenth International Conference on*. IEEE, pp. 267–272.
- [16] H. Shin, J. Chon, and H. Cha, "SmartSLAM: Constructing an indoor floor plan using smartphone," 2010.