Extracting Co-locator context

Swaroop Kalasapur, Henry Song and Doreen Cheng

Abstract— **Having reliable context sources is very important for context-aware applications and the devices around a user can be a useful context for many applications. While the importance of 'devices around' as a context has been highlighted many times, to the best of our knowledge, there is no systematic mechanism reported to identify and capture this context. In this paper we report our mechanism for extracting patterns that we call co-locators from the surrounding Bluetooth devices. The co-locator context can then be used as an indication about the user's surroundings. The context can be directly associated with user behavior in the surroundings, which in turn can be used for recommendation purposes. The context can also be used as a supplemental location indicator in the absence of other location mechanisms such as GPS. We then show that our experimental results support the value of using extracted co-locator patterns as good supplemental location indicators for several applications.**

Index Terms—**Location Context, Co-locators, Semantic Location, Bluetooth Location.**

I. Introduction

THE present day personal devices have become powerful platforms capable of performing considerable amount of computation and communication. Context aware applications are fast becoming a reality on such devices. To enable such applications, availability of reliable context information is necessary.

One of the most precious contexts that have proven to be extremely valuable is the location of the user. While there are a number of existing mechanisms for location identification, each providing information at different levels of granularity, each of the existing mechanisms has its own practical impediments. For example, GPS based solutions do not work indoors and amidst high-rise buildings. WiFi based solutions require databases that are kept up-to-date. Since the aim of existing mechanisms is to provide location information to the masses, they do not differentiate their results for individual users. The richness in the association between user actions and their location is the Holy Grail which Location Based Services are going after.

In many instances, the user behavior is also dependant on their surroundings, even within the same location. For example, the user might be willing to take a phone call when they are in their office, but might not be as receptive to the phone call if they are in a meeting. By using just the location of the user, it is quite difficult to distinguish between the above scenarios. It is possible however, to distinguish the above based on the presence of other devices in the user's vicinity. For example, if there is good evidence that the user does not take any phone calls when surrounded by a particular set of devices, the decision can be automated on behalf of the user. Many of the users spend a significant amount of time in familiar surroundings, such as their work place or home. Under such familiar surroundings, the user's device will witness a similar set of devices repeatedly, over a period of time. With many of the personal devices being equipped with Bluetooth interface, it is common to see many Bluetooth devices in our everyday surroundings. These Bluetooth devices can be discovered by another Bluetooth capable device on demand. By extracting patterns within the discovered Bluetooth devices, it is possible to use such extracted patterns as a reliable context. The Familiar Stranger [11] project from Intel's lab at UC Berkeley has studied this phenomenon of people encountering same people repeatedly from a social sciences perspective. We attempt to take advantage of the presence of such familiar surroundings to model the familiar devices as a context.

In this paper, we focus on discovering such patterns among the devices around the user. We call this context as 'co-locators'. The co-locator context can be used as a supplemental location identifier, in the absence of a reliable positioning mechanism such as GPS. Also, the co-locator context can be directly consumed by context aware applications. If the users are willing to associate the discovered co-locator patterns with meaningful names, such as 'in a meeting', 'in the cafeteria', 'with design team', etc, such names can be directly consumed by a wide range of applications such as social networks that require presence information, auto-tagging applications for photographs, life streaming applications, etc.

We believe that our approach can be very valuable even in the presence of other location identification mechanisms. We have constructed a hierarchical location store that keeps track of the user locations based on identified granularity. At the highest level of the store, the zip-code based location as identified from the base station associated with the device is stored. At the next level, the street address retrieved from

Manuscript received March 27, 2009.

Swaroop Kalasapur is a Senior Engineer at Samsung Electronics R&D Center in San Jose, CA, USA (e-mail: s.kalasapur@sisa.samsung.com)

Henry Song is a Staff Engineer at Samsung Electronics R&D Center in San Jose, CA, USA (e-mail: hsong@sisa.samsung.com)

Doreen Cheng is a Principal Engineer at Samsung Electronics R&D Center in San Jose, CA, USA (e-mail: doreen.c@sisa.samsung.com)

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857

performing a reverse geo-code lookup is stored. In the lowest level of the location hierarchy, the co-locator patterns identified by the approach detailed in this paper are stored. When an application on the user device wants the location information, it can refer to the hierarchical store to retrieve the information. The users can optionally associate meaningful tags with the stored locations, at any level in the hierarchy. Such tags can be consumed directly by applications such as presence settings, photo-tags, etc.

The rest of the paper is organized as follows. In section II, we present some of the related work in the field of location identification, and also discuss the salient features our mechanism enables. In section III, we provide a discussion on the use of the co-locator context, and their applicability in context-aware applications. In section IV, we present our system architecture and algorithm for semantic location identification. Section V describes our experiments and results. We conclude this paper and discuss the future work in section VI.

II. RELATED WORK

There is a large body of work available for identifying the location of the device. The most popular mechanism employed in the devices today is the GPS sensor based mechanism. While the GPS sensors give accurate location information, their performance is not satisfactory when a device is indoors, or is in a highly urban environment with tall buildings. GPS devices also suffer from a high lag when turned on. Each GPS sensor will have to contact the satellites to obtain its position. This initial communication with the satellites can take quite a long time. Moreover, if the device is out of range of the satellites (ex: within a building), the sensors spend a significant amount of time and battery power trying to communicate with the satellites, before it reports the unsuccessful event.

Since many of the existing devices are equipped with WiFi interfaces for communication, researchers at University of Washington's Placelab have devised a way of finding out the position of a device based on the visible access points. By maintaining a database of Access points and the geographical locations where the access points are installed, [1] provides geographical address information to a device that reports the access points that are currently visible to it. While very useful, the mechanism proposed in [1] has a high cost of acquiring and maintaining the database of all access points. While extensive access point databases exist for popular urban areas, maintaining such databases can prove challenging in remote areas. A similar mechanism is used in [12] to position a wireless device, based on the observed signal strengths from various access points.

WiFi signals have also been used to pin-point the position of a wireless device within buildings [5]. Since the position of Access points within a building can be pre-determined, by constructing a reliable model of WiFi signals within the building, the position of a wireless device can be determined with a high degree of accuracy by comparing the signal characteristics observed by the wireless device with that stored in the model. Advanced processing techniques of triangulation, based on the strength of received signals are employed in such methods. Many techniques have also been proposed based on WiFi finger printing [8]. In finger printing technique, instead of relying on an exact map of a location, it creates a search index of radio fingerprints to latitude/longitude coordinates.

Triangulation techniques are also employed to locate a mobile device, based on the received signal strength from various base stations. The E-911 service uses a variant of such triangulation to locate a mobile device. But, the information about the location is currently not available on the device itself.

There are a number of research efforts that have used Bluetooth devices for the purpose of identifying the location of a user device. Most of these techniques are aimed at providing the position of a user device within a well known geographical area such as within a building. Bluetooth devices are strategically placed at various places in a building and their positions are recorded. When the position of a user device is to be determined, the user device performs a scan of visible Bluetooth devices. Based on the visible devices [2], or based on the observed signal characteristics from various devices [4], [9], [10]. Most of these techniques aim at providing highly accurate position information to the user devices. The high cost involved in configuring and maintaining an accurate positioning system is perhaps the greatest challenge in adapting such positioning mechanisms.

The closest work to our approach is reported in [2]. The approach taken in [2] is to place certain Bluetooth devices in the environment as location indicators. When a particular user is not able to gather the location through any other means, for example through cell-tower triangulation, or GPS, the Bluetooth device that was strategically placed is used as an indicator for location. But, as with most of the existing indoor location identification mechanisms, the major challenge is to place any location indicators in all possible places that a user can be at. The practical limitation of the above mechanism makes it only applicable in tech-savvy places such as university labs.

In our approach, we do not depend on the user's environment being strategically configured for location identification. We exploit the fact that most of the users have regular habits that they follow. Most of the users spend significant amounts of time under familiar surroundings such as their office, home, with friends, etc. By observing the Bluetooth devices around the user, and extracting reliable patterns around the user within such familiar places, it is possible to provide a low-cost solution for identifying user's location. Such extracted patterns can also be useful as an important context for context-aware applications.

Since we primarily are interested in just the visibility of a device, the mechanism used for identifying patterns is computationally inexpensive. The only added cost is in remembering the devices that are seen in user's surrounding

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857

environment. Since Bluetooth devices are becoming very common in most of the user devices, we believe that using the Bluetooth interface as a source of context is very attractive and useful.

III. The Context of Co-locators

It is very common for people to spend extended amounts of time in familiar places. In such places, they typically encounter a common set of devices, either carried by their colleagues, friends, family, or devices that are stationary such as computers, laptops, etc. In the absence of any mechanism to determine the location of the user, the surrounding devices can be used to deduce the user location. But since many of the devices are mobile, it is necessary to find a stable pattern that can reliably determine the location.

The co-locator based location identification can also be useful for devices that are not equipped with location sensing mechanisms, such as low-end mobile phones. Since the Bluetooth interface is very inexpensive, and is common in most of the mobile phones to enable hands-free communication, we can exploit their presence to build location identification mechanism using the co-locator patterns.

Apart from serving as supplemental location identification mechanism, the co-locator patterns can also be used directly as a context source. Within a single location, it can be expected that a user would behave differently when people around the user are different. For example, if the user is in a meeting with clients, an incoming phone call might disturb the user. But, if the user is at the work desk, an incoming phone call might be promptly answered. This distinction within the same location needs the knowledge about the user's surrounding environment and co-locator patterns can provide an important context that applications can make use of.

A. *Co-locator patterns based on Bluetooth devices*

Since we primarily employ the visible Bluetooth devices to derive co-locator patterns, we provide a brief description about the range of Bluetooth signal and its effect on the accuracy of location determination. Bluetooth specifications support multiple power classes. Based on the class of a particular device, the visibility ranges from about 1 meter to about 100 meters. Class 1 devices consume the maximum power, but also have a visible range of up to 100 meters. Class 3 devices are most power conservative, having visibility range up to 1 meter. Class 2 devices are in between both in terms of power consumption and in their range. Class 2 devices have a visible range up to 10 meters. Each Bluetooth device also implements one of the standardized profiles, making its features usable by other Bluetooth devices. Since we are interested only in the presence of a device, we do not go any further into Bluetooth profiles. Since a mobile device can see all visible devices, irrespective of their class, it will be necessary to determine the relative distance based on other means such as the class of the device, or the received signal strength. In our current experiments, we simply consider a visible device as being near, without delving further into the relative distance.

While the Bluetooth specifications do not require a device to transmit its MAC address to neighbors, most of the current devices do transmit this information. We employ the MAC address of a Bluetooth device as a unique identifier for the device. Since we are not targeting at sharing the collected data, apart from using them as patterns within the user devices, we believe that the problem of privacy will not be a challenge for its adoption.

To get reliable data on from which we can derive co-locator patterns, we have developed a mechanism that periodically scans for other Bluetooth devices in the vicinity. Although this consumes additional battery power, the utility of the co-locator patterns in context-aware applications justify the additional power consumed to gather the data. As reported in [2], keeping the Bluetooth interface turned on and performing periodic scans consume relatively less power, making it practical to employ the mechanism. In our experiments, we experienced that by keeping the Bluetooth interface turned *on*, and running an active scan every 15 minutes, we were required to charge the device every day.

B. *Hierarchical location identification*

We have designed a hierarchical location structure to attach semantic location identifiers. We employ all forms of location identification available through the device to form the hierarchical location structure. Since we primarily worked on mobile phones, we accommodated for cell tower information, GPS information and Bluetooth information for location determination. While WiFi can also be used for similar purposes in our architecture, we have not collected data with respect to WiFi devices and hence not a part of this paper. The hierarchical location structure is illustrated in Figure 1.

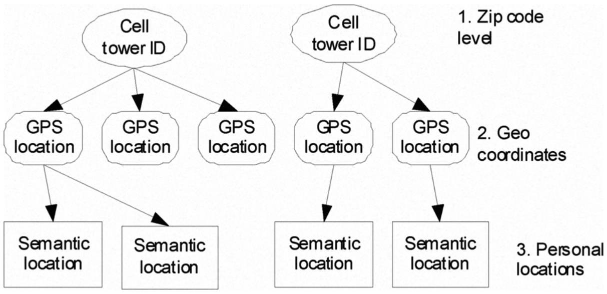

Fig. 1. Hierarchical location organization.

The cell tower ID is the easiest and cheapest location identifier. Since a mobile phone is already associated with the base station, the ID of the base station is readily available on the mobile phone. There are existing databases that provide geographical information based on a cell tower ID. But, the granularity of the resolved geo-coordinate is accurate up to the Zip-code level. This information forms the highest identifier within the hierarchical location scheme.

The GPS sensor provides accurate geo-coordinates for a device. The result from the GPS sensor is in <*Latitude, Longitude*> format. There are existing web services that perform reverse geo-coding, taking the geo coordinates and returning a postal address corresponding to the coordinates. The result of such geo coding populates the second level of the location hierarchy. The accurate information provided by GPS

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857

comes at a cost. GPS is a power-hungry sensor. If the GPS sensor is left always powered on, it rapidly consumes power, which is a precious resource on mobile devices. Due to this reason, most of the mobile devices turn the sensor off when not in use. Therefore, the GPS sensor will have to find its position by contacting the satellites when turned on. This operation of finding satellites and acquiring location information can take significant time when the device is in the communication range of the satellites. Another problem with GPS is that it does not work indoors and in urban areas with tall buildings, making it difficult for users to get the location information. This is one area where we believe that our work can supplement the location information.

At the lowest level of the location hierarchy are the semantic locations formed by extracting patterns from visible Bluetooth devices. In the presence of multiple, clearly distinguishable patterns, a single physical location can consist of a number of semantic locations.

The identification of the top levels within the location hierarchy is straight forward. In the next section, we describe the mechanism we have employed to form the patterns from visible Bluetooth devices.

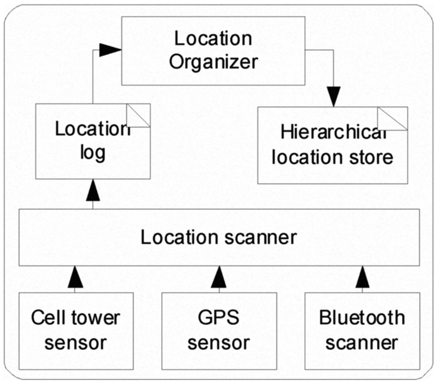

Fig. 2. Overall system architecture.

IV. SYSTEM ARCHITECTURE AND METHEDOLOGY

To construct the hierarchical location information, we developed a set of modules that collect information and process it on the mobile device. The modules were designed to periodically collect (a) cell tower to which the mobile device is associated (b) GPS information and (c) visible Bluetooth device information. The collected information is then processed into the hierarchical structure as shown in Figure 1. The periodic scanning for location information is scheduled by a module that stores the collected information into a log. The log is periodically processed on the mobile device to extract location information and populate the hierarchical location store. The overall architecture is shown in Figure 2.

The periodic scanning produces location logs that contain the recorded information along with time stamps. The Bluetooth devices addresses are stored separately with an incremental ID associated with each address. Therefore, within the log, the IDs associated with the devices are stored, and not the actual MAC addresses. The collected data is XML encoded for better readability. The collected log is periodically analyzed to populate the hierarchical location store. The top 2 levels of the hierarchical location structure are populated through lookups to web services. We have employed Zonetag [13] from Yahoo for looking up the location of cell towers and Google's reverse geo-coding web service for GPS locations. The result of the cell tower based lookup is in zip-code level granularity, while the results from GPS lookups are at a postal-address level granularity. Once the top 2 levels of the hierarchy are populated, the log data corresponding to the Bluetooth devices are analyzed to extract patterns.

TABLE I
BLUETOOTH DEVICES WITHIN A TIME WINDOW

Time stamp	*Bluetooth devices*
11/12 9:38	102 107 109 110 112
11/12 9:43	102 104 107 109 110 112 114 116 117
11/12 9:48	102 107 109 110 112 116 117
11/12 9:53	102 107 109 110 112 116 117
11/12 9:58	102 107 109 110 112 117
11/13 9:38	102 107 109 110 112 116
11/13 9:43	102 107 109 110 112
11/13 9:48	102 107 109 110 112 117
11/13 9:53	102 106 107 109 110 112 114 115 116 117
11/13 9:58	102 107 109 110 112 116 117

A. Extracting Co-locator patterns from Bluetooth device log

The basic idea of using visible Bluetooth devices as a location indicator depends on the fact that there is a possibility of seeing the same devices at the same location. The co-existence of devices is therefore the pattern we are interested in. We have devised a simple, frequency based co-occurrence pattern extraction mechanism to identify core groups among visible Bluetooth devices. The co-occurrence pattern is captured by aggregating all the occurrences over a time window. The length of the window can be adjusted based on the requirements. For example, to extract hourly patterns within a work-week, all data corresponding to the week days over a week are combined together and sliced into hourly intervals. An example of the logs extracted (only the Bluetooth portion) within the same 30 minute interval over two days is shown in Table I. Recall that the logs have already been sorted, based on the higher levels within the location hierarchy. Therefore, the Bluetooth device data we will be examining will correspond to those seen within a location indicated by GPS.

The logs are then transformed into a co-occurrence matrix within the considered time window. The co-occurrence matrix simply captures the count of devices appearing together within the window of time. Based on the number of samples within a given time window, if a device has been seen for a large number of times, the corresponding device can be employed as a good indicator for the location. Similarly, if a set of devices are seen together with high co-occurrence counts, the set can be used as a pattern to indicate the location. An example of the co-occurrence matrix is shown in Table II.

Once the co-occurrence matrix has been created, we can extract any existing patterns that can be employed as

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857

indicators to the location. We employ an approach based on heuristics for this purpose. If a particular device has been present for a large portion of the time window, it can be used as a good indicator. If any devices appear together for an extended period of time, we can employ the group of devices as a pattern for identifying the location. To capture the above factors, we have defined two threshold values. The first, T_c is defined as a percentage of the total number of log records being inspected within the considered time window. This limits our examination to only those devices that have been seen more than T_c times within the considered window. The second threshold T_i is a defined on only those elements that cross T_c, to determine the co-occurrence among elements. We follow the procedure listed below to extract the pattern that indicates a location.

a) Eliminate all elements that were seen less than T_c times (entries in the principle diagonal of the co-occurrence matrix).

b) For the remaining elements within the matrix, group together elements that cross the co-occurrence threshold of T_i.

c) Form the list of all element groups extracted in step b.

d) If no elements can be extracted after step c, this indicates that the data is unstable, and no clear identifier can be extracted for the location.

The above procedure first eliminates those elements that do not contribute to a stable pattern. Once we have identified those elements that appear within the considered time window, we then need to examine sets of such elements that appear together, as achieved by step b.

TABLE II

EXAMPLE CO-OCCURRENCE MATRIX

device	102	104	106	107	109	110	112	114	115	116	117
102	**10**	1	1	10	10	10	10	2	1	6	7
104	-	**1**	0	1	1	1	1	1	0	1	1
106	-	-	**1**	1	1	1	1	1	1	1	1
107	-	-	-	**10**	10	10	10	2	1	6	7
109	-	-	-	-	**10**	10	10	2	1	6	7
110	-	-	-	-	-	**10**	10	2	1	6	7
112	-	-	-	-	-	-	**10**	2	1	6	7
114	-	-	-	-	-	-	-	**2**	1	2	2
115	-	-	-	-	-	-	-	-	**1**	1	1
116	-	-	-	-	-	-	-	-	-	**6**	6
117	-	-	-	-	-	-	-	-	-	-	**7**

From the shown example in Table II, the extracted pattern following the above procedure, with the thresholds T_c and T_i both set to *80%* will be a string “102 107 109 110 112 117”. This set of Bluetooth devices can now be used as an indicator for the user’s location.

The patterns thus extracted for all considered time windows are now stored within the hierarchical location store. An example snapshot of the storage is shown in Figure 3.

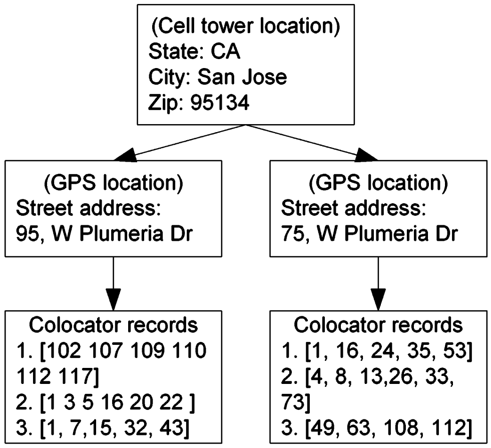

Fig. 3. Snapshot of the Hierarchical location store.

The value used as the threshold can have a significant impact on the extracted patterns. For example, if the user has very regular patterns, such as for a professional who spends most of the work day at a desk, higher threshold value can yield good co-locator patterns. But, if the environment around the user is highly dynamic, setting a high threshold might lead to very few co-locator patterns. Based on the values used for T_c and T_i, the patterns extracted can be different. For example, if we use a larger value for both thresholds, we can only extract patterns that are very significant. If the logs correspond to a highly dynamic environment, such as a shopping mall, higher values for the thresholds will result in very few or no patterns being recognized.

B. Looking up the Co-locator traces

When the device needs to know its location, or a context-aware application needs to consume the co-locator context, it can perform lookups on the hierarchical location store. Since the location store can be quite large, we can utilize the available cell tower information to arrive at the hierarchy we are interested in. Since each mobile phone always has the knowledge of the cell tower it is associated with, we can use this information without additional scanning cost. Within the hierarchy for a particular zip-code (arrived at through the cell tower information), we can look at all the co-locator records to extract the closest match. The procedure is similar to the one used for building the co-locator record. The logs collected over the past n time units, for example the last 60 minutes can be used to construct the co-occurrence matrix for the current time, as described in section *IV A*. Based on the constructed co-occurrence matrix, we can employ a similar procedure to extract the co-locator pattern for the current log over past *n* time units. This newly extracted pattern is then matched against those stored, to identify the closest match. The closest matching record from the hierarchical store is computed by a simple string similarity measure, with respect to a threshold. If the newly constructed co-locator record has a match exceeding a threshold (60% for example), we can conclude that the current co-locator context is similar to that stored.

V. EXPERIMENTAL RESULTS AND EVALUTATION

To evaluate our co-locator based context identification

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857

mechanism, we conducted a user study to collect data. For the user study, employed Samsung Blackjack II mobile phone equipped with GPS and Bluetooth. We had 6 participants in the study. We developed a software module that collected Cell tower data, GPS information and visible Bluetooth device information, along with other things. To obtain the relevant data, our software is equipped to communicate with the corresponding hardware resources [14]–[16] and query for the required information. We conducted the user study for a period of 6 months, during which the participants went about their daily activities. The experiment was designed to be completely un-obtrusive, and hence the participants were not required to do anything but use their phones as regular users. The software was designed to run in the background collecting relevant data and send it back to us. We have evaluated the co-locator context formation mechanism on the collected data set.

The collected data is then prepared for analysis by grouping the data based on date and time. To analyze the data, we considered data collected over one month at a time. We first extracted the Zip-code level location and GPS based location information from the logs and arranged the logs into the hierarchical structure shown in Figure 1. At this level, only the top 2 levels of the hierarchy are populated. Next, within each location, we grouped all data corresponding to weekdays together into one set and those corresponding to weekends into a separate set. Within the groups, we re-ordered the data based on the time slices, so that we can consider all the data within the same time interval, between 8:00 Am – 11:59 Am on weekdays for example. Once the time-folded data is prepared, we performed the pattern formation procedure detailed in section *IV A* to extract the co-locator records. We employed the threshold values of T_i and T_c both set to 80% for our evaluation.

In general, we observed that with more data available, the co-locator data grew to be more accurate, especially for individuals who adhere to a fairly regular schedule. But, we also observed that for individuals spending time in more dynamic environments, such as students, even with a large amount of data available, the co-locator context does not grow in accuracy. Therefore, we decided to slice the data purely based on time, rather than the amount of available data. Thus, we employed the time folding and slicing mechanism explained earlier.

From the collected data, we found that the mobile device did not have the GPS information for more than 50% of the time. This could be due to the fact that the corresponding user was inside a building where GPS is unusable. From the experimental evaluation, we were able to provide the location information for about 60% of the time instances where GPS is unavailable. This alone improves the availability of location information about 30% of the times.

Based on the collected data, we also observed that individual users have different patterns when it comes to locations. In many of the locations, we were able to form multiple co-locator patterns that could be used as supplemental location information. The findings from our experiments are listed in Table III. In the table, *A* represents the total number of locations as determined through GPS co-ordinates. *B* is the number of co-locator patterns identified within all the locations. *C* is the number of locations which contain more than 2 co-locator patterns. *D* is the percentage of instances when GPS was un-available and Co-locator pattern can be used as a supplemental location identifier.

TABLE III

EXPERIMENATL RESULT SUMMARY

User	*A*	*B*	*C*	*D*
1	40	17	3	64%
2	93	49	6	53%
3	110	53	5	40%
4	60	27	7	81%
5	53	15	3	28%
6	64	24	2	51%

The numbers in column C indicate that within the same geographical location, there were multiple patterns identified through co-location. Such patterns represent a richer context as compared to just the location co-ordinates. If we can identify multiple patterns within more locations, the co-locator pattern can be used not just as a supplemental location identifier, but also as a context indicating user's situation.

To use the co-locator patterns directly as context source, a context-aware application can query the hierarchical store to retrieve the pattern corresponding to those devices visible in the recent past, for example the last 30 minutes. Alternatively, a separate user interface can be provided so that users will associate meaningful names to the identified co-locator patterns. Such names can be directly consumed by various applications.

While the co-locator patterns can directly be used as supplemental source of location information, they are not always accurate. Consider the scenario where an individual goes to the cafeteria along with the other members of the same team. In such a case, since the same set of devices is seen at the cafeteria, as at the individual's office, co-locator pattern might directly conclude that the individual is still at office. But, this is issue can easily be handled, if the cafeteria is located at a different geographical location (as identified through cell tower or GPS information), or there are additional levels within the hierarchy, for example based on WiFi access points. The resolution of co-locator context thus depends on the hierarchical structure and the components that build this structure.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a mechanism to derive the co-locator context from visible Bluetooth devices around a user device. The resulting co-locator pattern can be a supplemental source for location identification, and also can be a valuable context source for various applications.

Along with identifying locations, the co-locator mechanism provides a personal store for managing the familiar locations that a user frequently visits. Just by associating the co-locator patterns with meaningful tags, it becomes possible for a variety of applications like life streaming, social networking

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857

updates, photo tagging, etc to consume such tags that carry personal meaning to the user.

Current mechanism for collecting visible Bluetooth device data is purely periodic. This can result in additional power consumption, which is not desirable for most users. To avoid this, we would like to experiment with adaptive mechanisms for data collection. One possible way for adapting the data collection is to integrate a graceful back-off mechanism. When the environment does not change rapidly, we can reduce the frequency of scanning. The frequency can be brought back to normal when we detect major changes in the environment, or a change in location is detected through other means such as GPS.

Along with the adaptive scanning of co-locator devices, it would also be valuable to have a good mechanism for determining the threshold values. Currently, we rely solely on heuristics to set the threshold values. But, as mentioned in section *IV A*, the threshold values play an important role in extracting useful co-locator patterns. We would like to investigate this issue further and aim to devise a mechanism that can adjust the thresholds based on individual users.

The co-locator based mechanism can co-exist with other location identification mechanisms. Since it fits nicely into the location hierarchy, it provides an additional level of granularity that can prove to be useful for a number of applications.

In the future, we plan to include the co-locator pattern based mechanism in a few applications. Currently we have examined the capabilities of the co-locator based mechanism offline. We would like to include the mechanism and a suite of applications that can consume the data on top, to get feedback from users.

There are a few other interesting aspects we can extract from the collected Bluetooth data. Many applications can benefit from knowing the state of the user, for example if the user is moving fast, or is stationary. This can easily be extracted from inspecting the recent log of location data. Continuous change in cell towers, and/or GPS data can indicate that the user is driving along. Once such information is identified, it is also possible to identify co-passengers based on the visible Bluetooth devices within the same time window, as detailed in this paper. Many peer-to-peer applications can greatly benefit from such information. It is also interesting to extract the co-locator context based upon other available information. For example, based on the calendar entry that suggests a meeting, it is possible to identify co-locators around the time of the meeting and associate such extracted patterns with the meeting event. Many personalization applications can consume such information to assist users.

References

[1] Lamarca, et.al., "Place lab: Device positioning using radio beacons in the wild," in Proceedings of the Third International Conference on Pervasive Computing, May 2005.

[2] N. Eagle, Pentland, A., "Reality mining: sensing complex social systems," Personal Ubiquitous Computing. 10, 4 (Mar. 2006), 255-268.

[3] J. Letchner, D. Fox, and A. LaMarca, "Large-Scale Localization from Wireless Signal Strength", In Proceedings of the National Conference on Artificial Intelligence (AAAI-05).

[4] P. Bahl and V. N. Padmanabhan, "RADAR: An In-Building RF-Based User Location and Tracking System", in *Proceedings of IEEE INFOCOM* 2000, Vol. 2: 775-784

[5] M. Ciurana, F. Barceló, and S. Cugno, "Indoor tracking in WLAN location with TOA measurements," In Proceedings of the 4th ACM international Workshop on Mobility Management and Wireless Access (Terromolinos, Spain, October 02 - 02, 2006). MobiWac '06. ACM, New York, NY, 121-125.

[6] G.V. Zàruba, M. Huber, F.A. Kamangar, and I. Chlamtac, "Indoor location tracking using RSSI readings from a single Wi-Fi access point," Wirel. Netw. 13, 2 (Apr. 2007), 221-235.

[7] A. Harder, L. Song, and Y. Wang, "Towards an Indoor Location System Using RF Signal Strength in IEEE 802.11 Networks," In Proceedings of the international Conference on information Technology: Coding and Computing (Itcc'05) - Volume II - Volume 02 (April 04 - 06, 2005). ITCC. IEEE Computer Society, Washington, DC, 228-233.

[8] K. Kaemarungsi, P. Krishnamurthy, "Modeling of indoor positioning systems based on location fingerprinting," INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies, Vol 2. pp. 1012 – 1022.

[9] S. Feldmann, K. Kyamakya, A. Zapater, Z. Lue, "An Indoor Bluetooth-Based Positioning System: Concept, Implementation and Experimental Evaluation," International Conference on Wireless Networks 2003: pp. 109-113

[10] M. Rodriguez, J.P. Pece, C.J.Escudero, "In-building location using Bluetooth," In Proceedings of the International Workshop on Wireless Ad Hoc Networks, 2005.

[11] Familiar Strangers: http://www.paulos.net/research/intel/familiarstranger/index.htm

[12] Loki: www.loki.com

[13] Yahoo Zone tag: http://developer.yahoo.com/yrb/zonetag/locatecell.html

[14] RIL Reference: http://msdn.microsoft.com/en-us/library/aa920441.aspx

[15] GPS intermediate driver: http://msdn.microsoft.com/en-us/library/bb202086.aspx

[16] Bluetooh on Windows Mobile: http://msdn.microsoft.com/en-us/library/aa916877.aspx

Digital Object Identifier: 10.4108/ICST.MOBIQUITOUS2009.6857
http://dx.doi.org/10.4108/ICST.MOBIQUITOUS2009.6857