



(19) **United States**

(12) **Patent Application Publication**
Chintalapudi et al.

(10) **Pub. No.: US 2011/0304503 A1**

(43) **Pub. Date: Dec. 15, 2011**

(54) **SIMULTANEOUS LOCALIZATION AND RF MODELING**

(52) **U.S. Cl. 342/357.29; 342/451**

(75) **Inventors: Krishna Kant Chintalapudi,**
Bangalore (IN); **Venkata N. Padmanabhan,**
Bangalore (IN); **Anand Padmanabha Iyer,**
Trivandrum (IN)

(57) **ABSTRACT**

The simultaneous localization and RF modeling technique pertains to a method of providing simultaneous localization and radio frequency (RF) modeling. In one embodiment, the technique operates in a space with wireless local area network coverage (or other RF transmitters). Users carrying Wi-Fi-enabled devices traverse this space while the mobile devices record the Received Signal Strength (RSS) measurements corresponding to access points (APs) in view at various unknown locations and report these RSS measurements, as well as any other available location fix to a localization server. A RF modeling algorithm runs on the server and is used to estimate the location of the APs using the recorded RSSI measurements and any other available location information. All of the observations are constrained by the physics of wireless propagation. The technique models these constraints and uses a genetic algorithm to solve them, thereby providing an absolute location of the mobile device.


(73) **Assignee: MICROSOFT CORPORATION,**
Redmond, WA (US)

(21) **Appl. No.: 12/815,973**

(22) **Filed: Jun. 15, 2010**

Publication Classification

(51) **Int. Cl.**
G01S 19/46 (2010.01)
G01S 3/02 (2006.01)



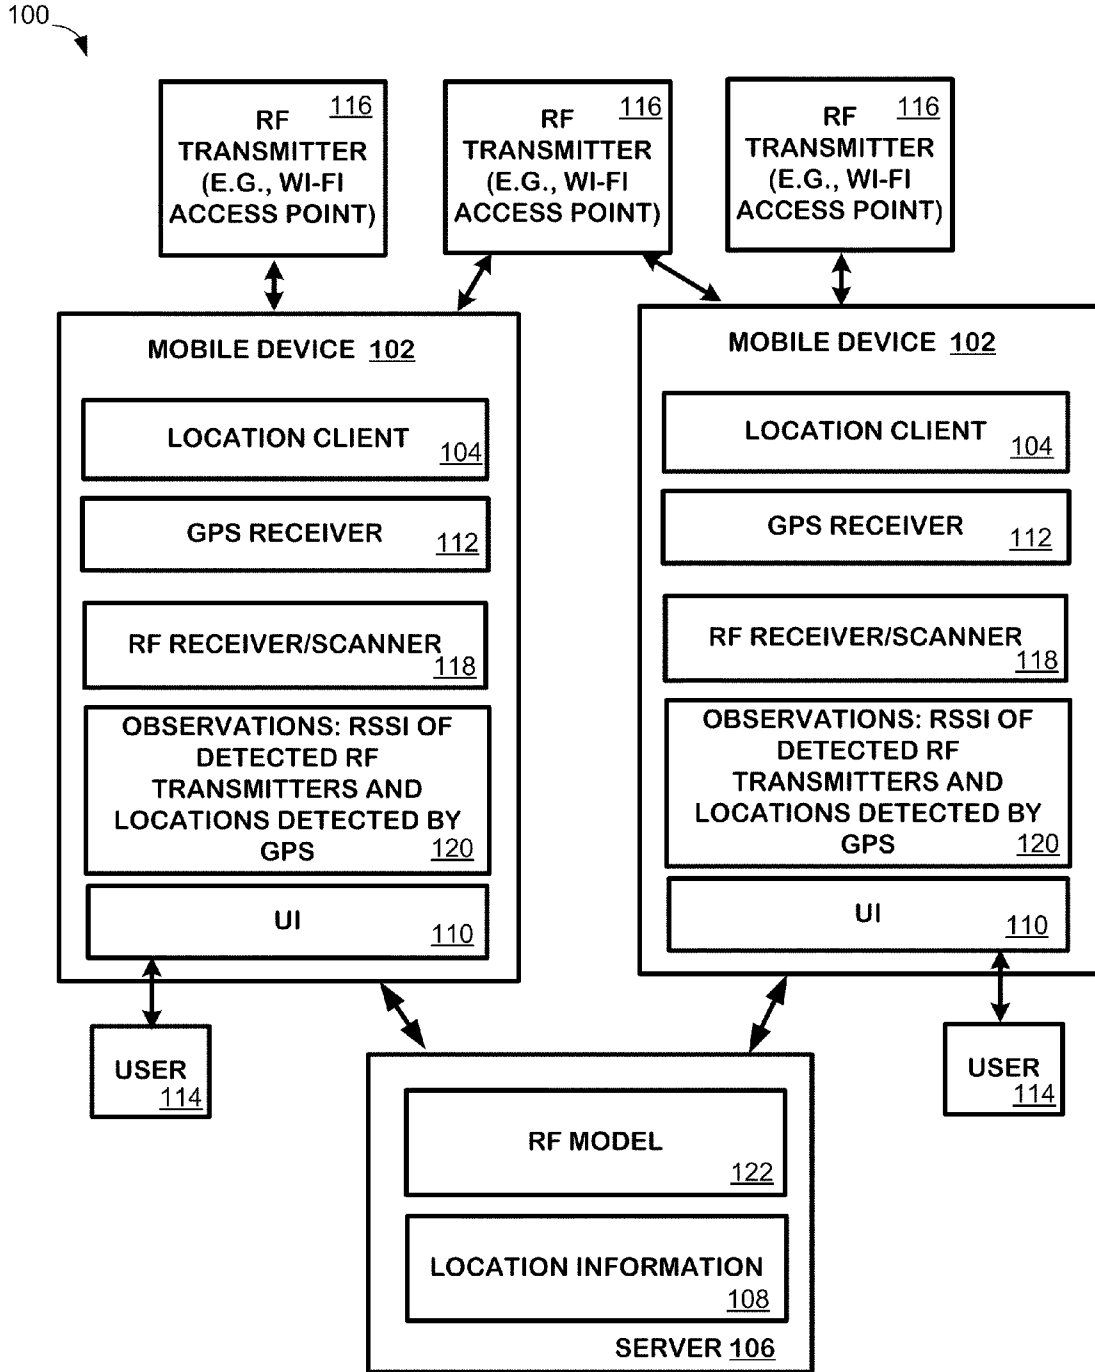


FIG. 1

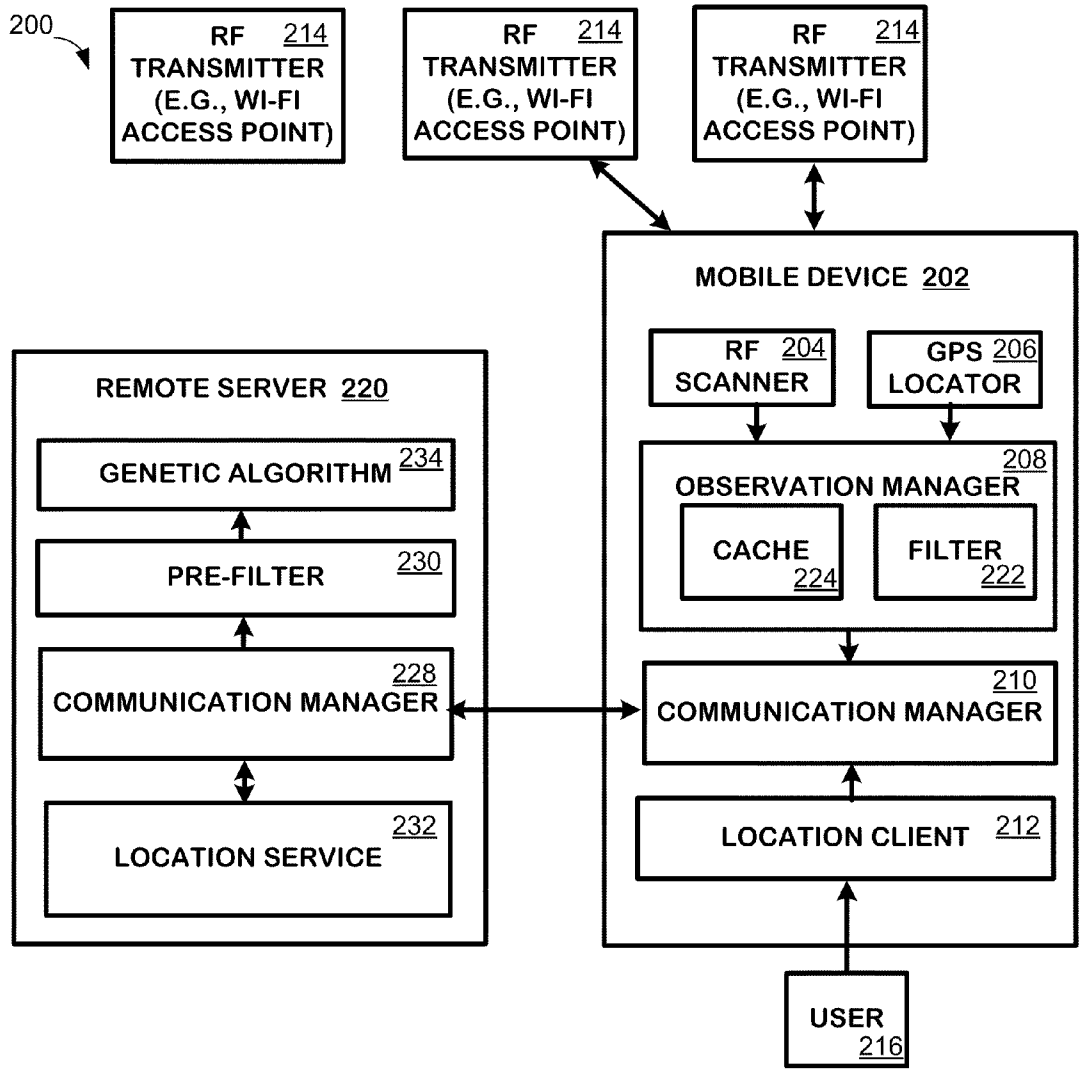


FIG. 2

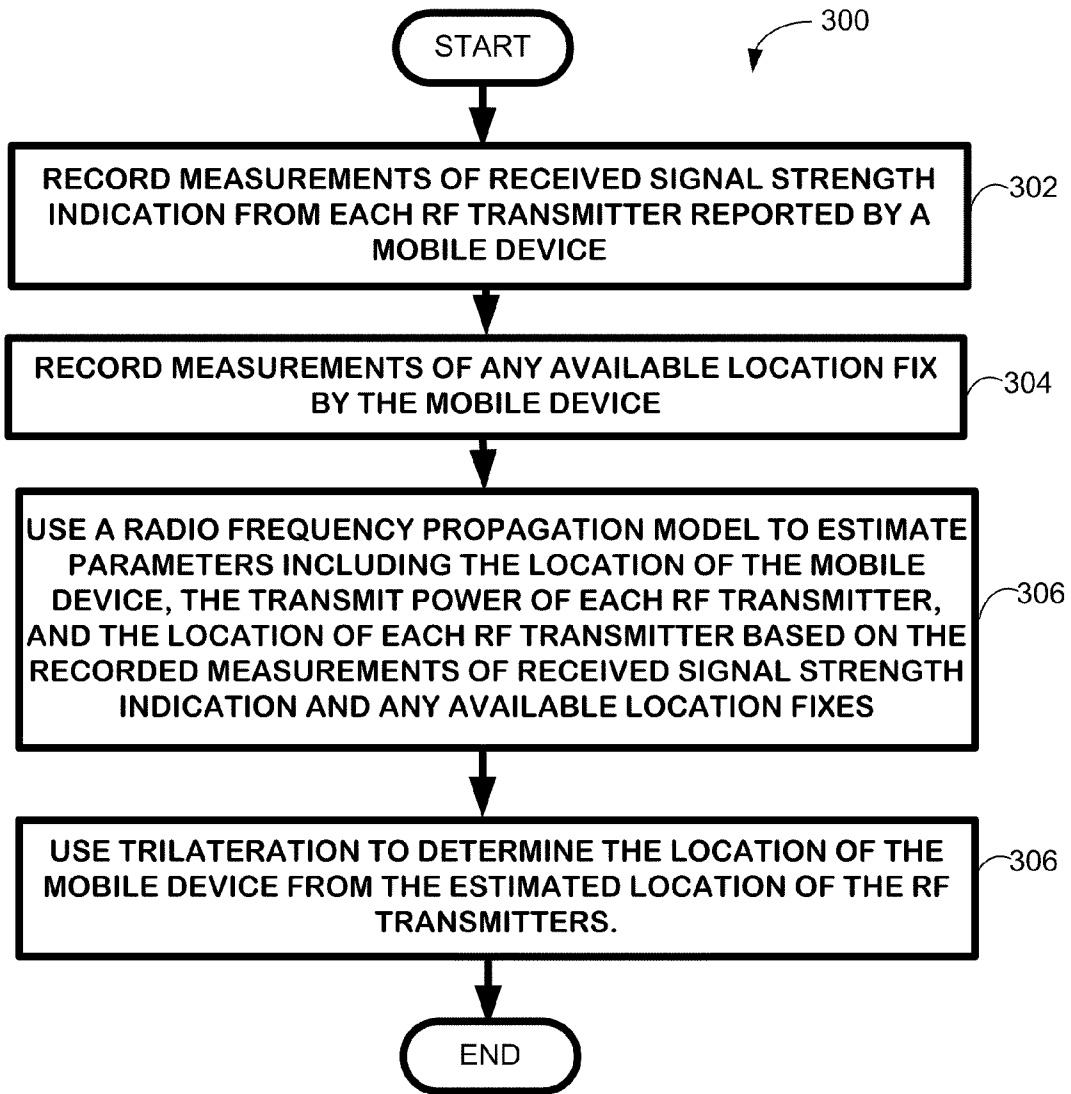


FIG. 3

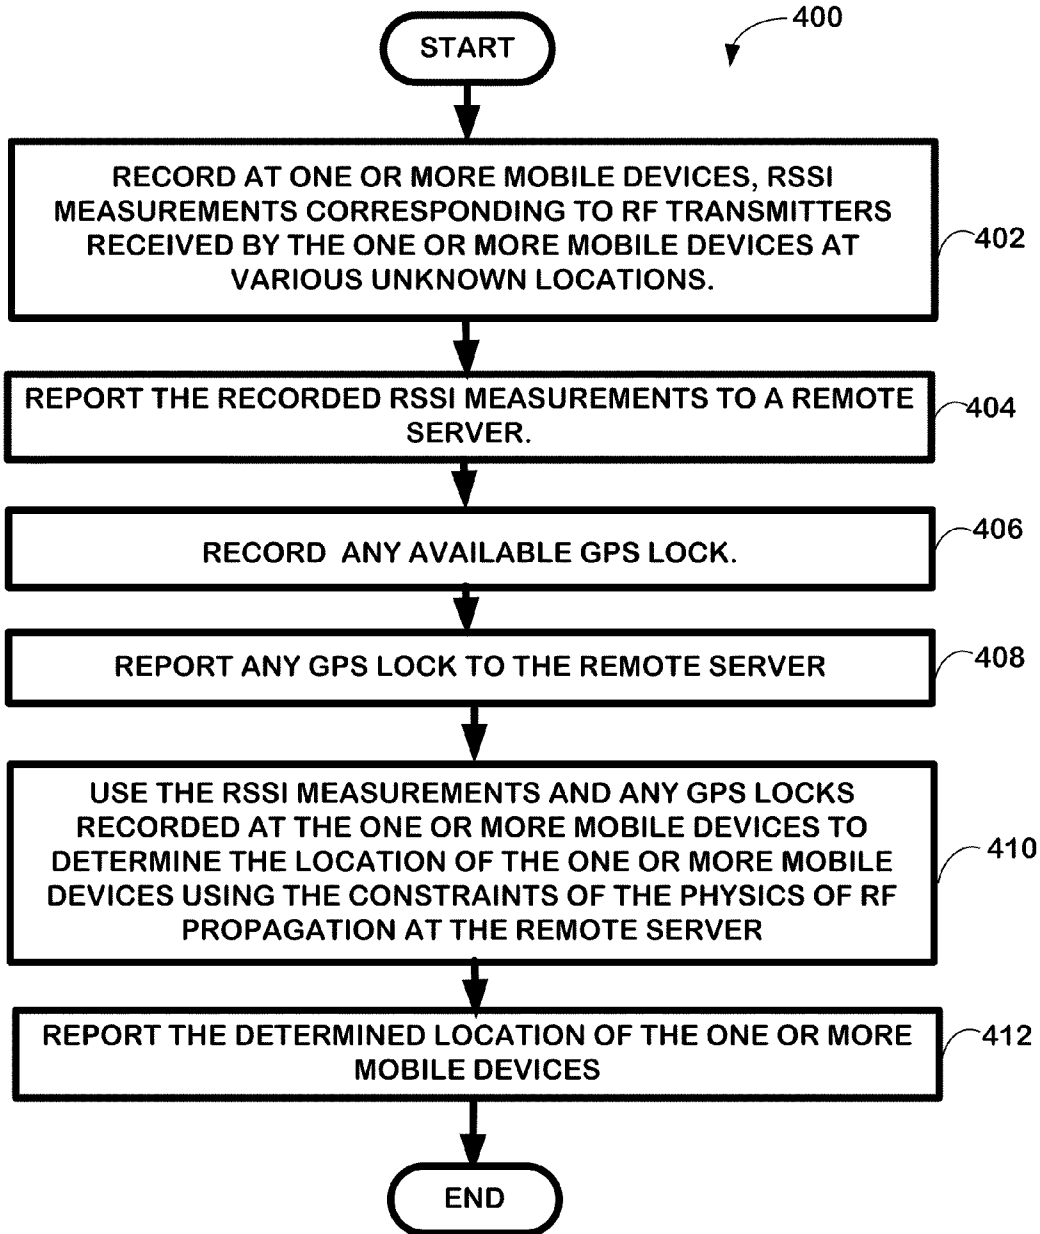


FIG. 4





FIG. 5

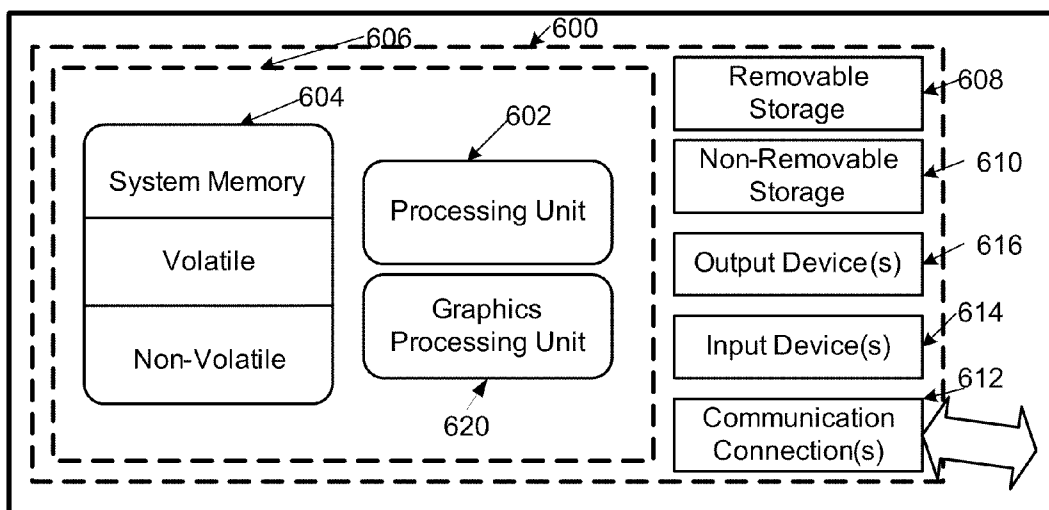


FIG. 6

SIMULTANEOUS LOCALIZATION AND RF MODELING

BACKGROUND

[0001] The need for location information to enable pervasive computing applications in indoor environments, coupled with the unavailability of GPS in such environments, has motivated a large body of research on indoor localization. In particular, there has been a focus on leveraging existing infrastructure, such as, for example, access points in a wireless local area networks, to enable indoor localization. The advantage of this approach is that the cost of deploying a specialized infrastructure for localization is avoided. While the near-ubiquity of wireless Local Area Networks (WLAN), also commonly named WiFi-networks, makes WLAN-based indoor localization attractive, there is a need for a significant degree of pre-deployment efforts such as, for example, building detailed Radio Frequency (RF) maps or RF propagation models based on surveys of the environment, with these approaches.

SUMMARY

[0002] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

[0003] The simultaneous localization and RF modeling technique described herein pertains to a method of providing simultaneous localization and radio frequency (RF) modeling.

[0004] In one embodiment of the simultaneous localization and RF modeling technique, the technique operates in an indoor space, such as an office building or a mall, with wide area WLAN coverage (or other RF transmitters) but where there is no knowledge of the physical layout, including the placement of wireless Access Point (APs) (or other RF transmitters). Users carrying Wi-Fi-enabled mobile computing devices, such as smartphones, traverse the space in normal course. The mobile devices record the Received Signal Strength (RSS) measurements corresponding to the APs in view at various (unknown) locations and report these to a localization server. Occasionally, a mobile device will also obtain and report a location fix, say by obtaining a Global Positioning System (GPS) lock at the entrance of the space or near a window. A simultaneous localization and RF modeling algorithm runs on the localization server and is used to estimate the location of the Wi-Fi APs using the recorded RSSI measurements and any other available location information. All of the observations reported to the server, even the many from unknown locations, are constrained by the physics of wireless RF propagation. The simultaneous localization and RF model models these constraints and then uses a genetic algorithm to solve them, thereby providing an absolute location of the mobile device (e.g., in terms of latitude and longitude).

[0005] It should be noted that the technique is not limited to only indoor environments and can be used anywhere there are RF transmitters. Additionally, any RF transmitter device can be used instead of only Wi-Fi access points, such as, for example, global system for mobile communications (GSM) transmitters. Also, although no knowledge of the physical

layout of the space being mapped is necessary, if layout information is available it can be used to improve simultaneous localization and RF modeling using the technique.

DESCRIPTION OF THE DRAWINGS

[0006] The specific features, aspects, and advantages of the disclosure will become better understood with regard to the following description, appended claims, and accompanying drawings where:

[0007] FIG. 1 is an exemplary architecture for employing one exemplary embodiment of the simultaneous localization and RF modeling technique described herein.

[0008] FIG. 2 is another exemplary architecture for employing another exemplary embodiment of the simultaneous localization and RF modeling technique.

[0009] FIG. 3 depicts a flow diagram of an exemplary process for employing one embodiment of the simultaneous localization and RF modeling technique.

[0010] FIG. 4 depicts a flow diagram of another exemplary process for employing another embodiment of the simultaneous localization and RF modeling technique.

[0011] FIG. 5 depicts a diagram showing how knowing enough distances between access points and mobile devices allows a unique determination of their relative locations.

[0012] FIG. 6 is a schematic of an exemplary computing environment which can be used to practice the simultaneous localization and RF modeling technique.

DETAILED DESCRIPTION

[0013] In the following description of the simultaneous localization and RF modeling technique, reference is made to the accompanying drawings, which form a part thereof, and which show by way of illustration examples by which the simultaneous localization and RF modeling technique described herein may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the claimed subject matter.

[0014] 1.0 Simultaneous Localization and RF Modeling Technique

[0015] The following sections provide an overview of the simultaneous localization and RF modeling technique, as well as exemplary architectures and processes for employing the technique. Exemplary computations for creating and using the RF model employed in one embodiment of the technique for localization are also described.

[0016] 1.1 Overview of the Technique

[0017] There has been a focus on leveraging existing infrastructure (e.g., wireless access points) to enable indoor localization, the advantage being that the cost of deploying a specialized infrastructure for localization is avoided. Existing solutions, however, require extensive pre-deployment effort, for instance, to build detailed RF maps or RF propagation models based on surveys of the environment. The simultaneous localization and RF modeling technique described herein provides a novel localization system that leverages existing infrastructure but does not require any explicit pre-deployment effort. Although much of this specification describes simultaneous localization and RF modeling by employing a wireless local area network (WLAN), commonly called a Wi-Fi network, and associated wireless access points in an indoor environment, it is to be understood that the

technique's embodiments can be employed with any system of RF transmitters in both indoor and outdoor environments.

[0018] In one embodiment, in an indoor environment, the simultaneous localization and RF modeling technique relies on three basic assumptions: (i) that there are enough wireless APs to provide excellent coverage throughout the indoor environment, (ii) that users carry mobile devices, such as smartphones and netbooks, equipped with Wi-Fi receivers, and (iii) that occasionally a mobile device obtains an absolute location fix, say by obtaining a GPS lock at the edges of the indoor environment, such as at the entrance or near a window. In one embodiment of the simultaneous localization and RF modeling technique, users simply sit, stand, or move around in the indoor environment in normal course. While they do so, each user's mobile device records the received signal strength (RSS) from the wireless APs visible to it at various (unknown) locations, and reports this information, along with the occasional location fix when available, to a central localization server. The server uses this data to simultaneously learn the characteristics of the RF propagation environment to build an RF model and to localize the users' devices. Localization is performed in terms of absolute coordinates: latitude and longitude.

[0019] A key advantage of embodiments of the simultaneous localization and RF modeling technique is that they do not require any prior knowledge of the RF environment, including the location and transmit power of the APs, information that is often not readily available in settings such as malls and multi-tenant office buildings, where APs have been deployed by many different entities reducing the calibration effort needed for indoor localization. Another advantage is that embodiments of the simultaneous localization and RF modeling technique described herein do not require any explicit user participation to aid the localization process. In particular, users are not required to indicate their current locations, even when building the RF model used in localization. Finally, the simultaneous localization and RF modeling technique only needs measurements of the APs by the mobile node and does not require any distance measurements between mobile nodes. So even a single mobile node that traverses the space of interest over time can generate sufficient data for the simultaneous localization and RF modeling.

[0020] By not requiring any pre-deployment effort or explicit user participation, the simultaneous localization and RF modeling technique has the potential of enabling practical and viable localization. For instance, a simultaneous localization and RF modeling server in a cloud could automatically construct an RF model for, and thereby enable localization in, an indoor space in any part of the world, based just on the measurements reported by the simultaneous localization and RF modeling technique's clients in the space of interest. Such automated operation brings forth a number of novel challenges, including filtering the very large number of measurement reports to identify a small and useful subset, and efficiently computing the RF model despite the very large space of possibilities.

[0021] 1.2 Exemplary Architecture.

[0022] As shown in FIG. 1, one embodiment of the simultaneous localization and RF modeling technique employs a client-server architecture **100**. Mobile computing devices (laptops, cell phones, netbooks, and so on) **102** act as clients and connect to a simultaneous localization and RF modeling server **106**. Both the mobile computing devices **102** and the server **106** can be implemented in a general purpose comput-

ing environment which will be discussed in greater detail with respect to FIG. 6. The server **106** is responsible for providing location information **108** to the one or more mobile devices **102**. In one embodiment of the technique, the system can support multiple mobile devices **102** to communicate with the server **106** and obtain location information. One embodiment of the technique provides a User Interface (UI) **110** on the mobile device **102** to let the user **114** send location queries. The UI **110** displays the current location of the user **114** on a floor map in response to the query. In one working embodiment of the technique, all communication between the mobile device **102** and the server **104** has been implemented over TCP sockets. This implementation uses a Wi-Fi interface to communicate, however, a cellular interface could be used as well.

[0023] In one embodiment, a location client **104** is a piece of software that is installed on the mobile devices **102**. The location client **104** has two important tasks. First, it provides location information to mobile applications residing on the mobile device **102**, and second, it assists the server **106** in the construction of an RSS map of the local environment.

[0024] In one embodiment of the technique, the location client **104** first checks to see it can obtain its location using a GPS receiver **112**. If no GPS signal is available, it scans its environment for wireless APs **116** in its view using a RF scanner **118**. The scanner **118** scans for a few seconds (3 seconds in one implementation) collecting beacons from each AP **116** it sees. It then transmits its observations **120**, in one embodiment a list of the mean and standard deviation of the RSS seen from various APs **116** over the collected data to the server **106**. Using the mean reduces errors due to multi-path while the standard deviation is used to determine the quality of the data. The server **106** then uses a computed RF model **122** of the RSS environment to determine the location of the mobile device in real time and responds to the location client **104**.

[0025] For new indoor spaces, the RF model **122** must first be generated before location queries can be answered. For this, the client or clients **104** scan for wireless APs **116** in their range and transmit the observed mean and standard deviation of RSS to the server **106** in the same manner as a location query. While in one implementation the mobile devices **102** periodically perform a scan and push the information to the server **106**, this can be easily implemented as a pull, where the server **106** requests a scan. Pull is desirable since mobile devices then need scan only when necessary and this can lead to significant energy savings.

[0026] The server **106** has two important functions. First, it responds to location queries from the location clients **104** in real time, and second, it uses a simultaneous localization and RF modeling algorithm to construct and maintain the RF model **122** for the environment in question.

[0027] The data from a large number of mobile devices can be very large and in many cases is not useful for creating the RF model **122**. Thus, in one embodiment, the simultaneous localization and RF modeling technique performs pre-filtering and selects only a useful subset of the available data to construct the model. Additionally, the server **106**, may potentially reside in a cloud and leverage its computing resource to construct the RF model **122**.

[0028] FIG. 2 provides another exemplary architecture **200** for employing the simultaneous localization and RF modeling technique. A mobile device **202**, as depicted in FIG. 2 comprises five components namely, a RF scanner **204**, a GPS

locator **206**, an observation manager **208**, a communication manager **210** and a location client **212**. The RF scanner **204** periodically scans for wireless access points **214**. In one working embodiment, the RF scanner **204** scans for about 3 seconds collecting roughly 30 beacons from each AP **214**. Collecting multiple beacons allows for computing average RSS values and reduces the errors due to multi-path effects. Computation of the standard deviation can be used to determine the quality of the data. For example, a mobile user **216** in motion will see higher standard deviations in the RSS compared to a stationary observer. The GPS locator **206** is triggered periodically along with the RF scanner **204**. It uses the device's internal GPS information to obtain the latitude and longitude of the mobile device **202**. The GPS locator **206** is used to transmit known locations opportunistically whenever the user **216** is located near balconies, entrances or doors and obtains a GPS lock.

[0029] The communication manager **210** handles all communication between the client and the server. In one working embodiment, the technique uses a Wi-Fi interface to communicate.

[0030] The observation manager **208** collects raw data from the RF scanner **204** and filters it in a filter **222** before transmitting it to the server **220**. The observation manager **208** computes the average and standard deviation of the RSS over all the readings for each AP. In one working embodiment, average RSS readings less than -85 dB are deemed unreliable and dropped. Low RSS readings can lead to large location accuracies due to decreased spatial resolution arising out the exponential nature of signal decay. The observation manager **208** also suppresses transmission of RSS readings with a high standard deviation. Transmitting data from the same location over and over again when the mobile user is stationary can be wasteful. To avoid this, in one embodiment of the technique, the observation manager **208** maintains a cache **224** of recently transmitted values and suppresses potentially duplicate information by checking against the cache.

[0031] The location client **212** provides an interface to send queries to the server **220** issued from the user **216**. Upon receiving a query request from the user **216**, the location client **212** first attempts to use the GPS locator **206** to obtain its location. Upon failing, it uses the RF Scanner **204** to obtain the list of APs **214** seen and the corresponding RSS values. These values are sent in the query to the server **220** using the communication manager **210** on the mobile device **202**.

[0032] The server **220**, as depicted in FIG. 2 comprises four components namely, a communication manager **228**, a pre-filter **230**, a location service **232**, and a genetic algorithm (GA) module **234**. The location service **232** is responsible for handling location queries from various users. Upon receiving a query (comprising RSS values seen from various APs **214**) it forwards the request to a GA module **234**. The GA module **234** then uses the current estimate of the AP parameters to convert the RSS values into distances. Location is then estimated in the location service module **232** using trilateration and sent back to location client **212** using the communication manager **228**. Trilateration is a method of determining the relative positions of objects using the geometry of triangles in a similar fashion as triangulation. Unlike triangulation, which uses angle measurements (together with at least one known distance) to calculate the subject's location, trilateration uses the known locations of two or more reference points, and the measured distance between the subject and each reference point.

[0033] In one embodiment, the pre-filter **230** is responsible for analyzing the received RSS data received periodically from various users and presenting only a representative subset to the GA module **234**. This is important to make the technique practical, since solving the equations with thousands of locations may require impractically long amounts of time. In one embodiment, the pre-filter **230** uses specialized algorithms described later in Sections 2.1 and 2.2 to reduce the volumes of data to a manageable set of equations.

[0034] The GA module **234** continuously runs a GA (described later in more detail) to solve the equations, improving the estimates of the AP parameters in each generation. The GA is inherently amenable to massively parallel computation, since each of the operations such as crossover and mutation, for example, can be parallelized.

[0035] 1.3 Exemplary Processes Employed by the Simultaneous Localization and RF modeling Technique.

[0036] The following paragraphs provide descriptions of exemplary processes for employing the simultaneous localization and RF modeling technique. It should be understood that in some cases the order of actions can be interchanged, and in some cases some of the actions may even be omitted.

[0037] FIG. 3 depicts an exemplary computer-implemented process **300** for performing localization of one or more mobile devices. As shown in block **302**, measurements of received signal strength indication (RSSI) are recorded from each Radio Frequency (RF) transmitter reported by a mobile device. Additionally, as shown in block **304**, measurements of any available location fix by the mobile device are also recorded. As shown in block **306**, a RF propagation model is used to estimate parameters that include the location of the mobile device, the transmit power of each RF transmitter, and the location of each RF transmitter based on the recorded measurements of received signal strength indication and any available location fixes. Trilateration is then used, as shown in block **308**, to determine the location of the mobile device from the estimated location of the RF transmitter and the estimated parameters.

[0038] FIG. 4 depicts another exemplary computer-implemented process **400** for determining the location of one or more mobile devices in a space. As shown in block **402**, RSSI measurements corresponding to RF transmitters received by the one or more mobile devices at various unknown locations are recorded at one or more mobile devices. The recorded RSSI measurements are sent to a remote server, as shown in block **404**. Any available GPS lock is also recorded at the one or more mobile devices, as shown in block **406**. The recorded GPS locks are also reported to the remote server, as shown in block **408**. As shown in block **410**, the RSSI measurements and any GPS locks recorded at the one or more mobile devices are used to determine the location of the one or more mobile devices using the constraints of the physics of RF propagation at the remote server. The determined location of the one or more mobile devices is reported back to one or more of the mobile devices, as shown in block **412**.

[0039] 1.4 Creation and Use of a RF Model.

[0040] An overview of the technique and related art, as well as exemplary architectures and processes employing the technique, having been discussed, the following sections provide exemplary computations used in one embodiment of the technique to create and use a RF model (e.g., as shown in FIG. 1, block **122**) to determine mobile device location. In general, the model is created using a set of LDPL equations and then these equations are used to find the location of one or more

mobile devices. In one embodiment a genetic algorithm is used to solve the set of equations.

[0041] 1.4.1 Localization using the Simultaneous Localization and RF Modeling Technique

[0042] To provide a physical intuition to the working of the simultaneous localization and RF modeling technique described herein, example scenarios depicted in FIG. 5 are referred to. In scenario 1-A, two mobile nodes (M1 502 and M2 504) have measured their distances (d_{11} , d_{12} , d_{21} , d_{22}) 506, 508, 510 and 512, respectively, from two APs (AP1 516 and AP2 514) with the hope of determining their locations relative to the two APs. This is, however, not possible, since the same set of distances allows for several different possibilities for relative locations (an alternative scenario 518 is depicted in Scenario I-B). For a set of three APs 520, 522, 524 and three mobile user locations 526, 528, 530 (as depicted in Scenario-II), it can be shown that knowing all nine pairs of distances between APs and mobile users allows for only one possible realization for their relative positions. Such a structure is deemed localizable, (or globally rigid) i.e., the entire set of locations can be translated, rotated and reflected (flipped) but not distorted in any manner if all distances are to be preserved. Localizability is a well studied area and conditions for localizability have been well studied. In general given “enough” distance constraints between APs and mobile devices, it is possible to establish all their locations in a relative sense. Knowing the absolute locations of any three non-collinear mobile devices then allows determination of the absolute locations of the rest.

[0043] In practice, however, the distances between mobile devices and APs can only be inferred from RSS values.

$$p_{ij} = P_i - 10\gamma_i \log d_{ij} + R \quad (1)$$

[0044] In Eqn 1, the j^{th} mobile user located at a distance d_{ij} (e.g., measured in meters) from the i^{th} AP sees a signal strength of p_{ij} (measured in dBm). P_i is the RSS from the i^{th} AP at a distance of one meter (referred to as transmit power henceforth). The path loss exponent γ_i captures the rate of fall of RSS in the vicinity of the i^{th} AP. The higher the value of γ_i , the steeper is the fall of RSS with distance. The need for having a different γ_i for each AP arises from the fact that rate at which RSS falls with distance depends on the local environment. RSS from an AP that is located in an area surrounded by walls, people and other obstacles might decay at a much faster rate compared to the same from other APs in the indoor environment that enjoy relatively freer signal propagation. R in Eqn 1 is a random variable that hopes to capture the variations in the RSS due to multi-path effects, asymmetries in the physical environment (e.g., obstructions) and other imperfections in the model itself.

[0045] Based on a log distance path loss (LDPL) model d_{ij} can be computed as,

$$d_{ij} = 10^{\left(\frac{P_i - p_{ij}}{10\gamma_i}\right)} \quad (2)$$

[0046] Eqn 2 assumes the a priori knowledge of P_i and γ_i . The simultaneous localization and RF modeling technique, takes a novel approach to estimating these. Given a set of RSS observations between APs and mobile users (p_{ij}), the technique treats P_i and γ_i as unknowns in addition to the unknown

locations of APs and mobile users. It then solves the set of simultaneous equations formed by the LDPL model for each RSS observation.

[0047] Assume that there are m APs and n unknown locations on a floor. For simplicity assume that all the m APs are visible from each of the n locations (this assumption will be relaxed later in this section). The total number of RSS observations and hence the total number of LDPL equations will be mn . Assuming 2D locations, each of the n locations has two unknowns namely the x and y coordinates. Each of the m APs has four unknowns namely P_i , γ_i and its 2D location. The total number of unknowns is thus $4m+2n$.

[0048] While mn grows in a quadratic fashion, $4m+2n$ grows linearly. This suggests that given enough locations (such that $mn > 4m+2n$), there will be eventually enough constraints in the system of equations to make the system uniquely solvable. Closer examination however, reveals that the system of LDPL equations is scale, translation, rotation and reflection invariant. In other words, a solution to LDPL equations will yield locations that are a scaled, translated, rotated and/or reflected version of the true locations. Knowing three true, non-collinear locations (either AP or mobile users) then, all the other true locations can be determined.

[0049] 1.4.2 The Nature of LDPL Equations

[0050] LDPL equations (Eqn 1) are a system of simultaneous non-linear equations. This section explains the nature of solutions to LDPL equations.

[0051] 1.4.2.1 Solving for the Parameters of a Solitary AP

[0052] It is well known that distances from at least three known non-collinear locations are necessary to uniquely determine an unknown location (using trilateration). A corresponding issue in the simultaneous localization and RF modeling technique is, what is the minimum number of known locations at which RSS measurements must be taken in order to uniquely determine the four AP parameters namely P_i , γ_i and the 2D location?

[0053] If P and γ are known, then a RSS measurement p can be converted into the distance d between the x and the AP using the equation Eqn 2. A minimum of three such RSS measurements at known locations are then required to uniquely establish the location of the AP. In the absence of the knowledge of P and γ , then, two additional measurements (constraints) will be required to uniquely determine their values. In other words five RSS measurements are required to uniquely determine an AP. It also follows that that given RSS measurements from only four known locations, there will be two possible solutions for (P , γ and location) that satisfy the set of four RSS measurements.

[0054] 1.4.2.2 The Notion of Co-Circular Dependency

[0055] It is well known that three or more collinear locations cannot be used in trilateration to determine an unknown location. A similar yet slightly different situation arises in solving LDPL equations when all the locations where RSS observations were taken are co-circular with respect to the AP i.e., lie on a circle centered around the AP's location. In this special case, while the location of the AP can be ascertained as the centre of the circle passing through these locations, it is impossible to uniquely determine both P and γ . In the simultaneous localization and RF modeling technique thus, one embodiment avoids observations that have almost the same RSS values from the same AP.

[0056] 1.4.3 Localizability

[0057] While a necessary criterion for the existence of a unique solution to the system of simultaneous equations is

that the number of equations should be greater than or equal to the number of variables, this is by no means a sufficient condition that guarantees a unique solution. For example, if there are 6 APs (AP₁ through AP₆) and 20 mobile user locations from which RSS observations were made, an edge between an AP and a mobile user location is drawn if and only if the AP can be seen from that location. The number of LDPL equations in this system is 61, and the number of variables is 58. However, from the very structure it is clear that the system is not localizable since the group of APs, AP1 through AP3 is free to rotate about to the AP4 through AP6.

[0058] The localizability question in the simultaneous localization and RF modeling technique is as follows: given a set of RSS measurements at some unknown and known locations, is it possible to determine a unique set of coordinates for all the unknown locations by solving the corresponding the simultaneous localization and RF modeling technique's equations? The question is extremely relevant because, in practice not all APs may be visible from all locations in the indoor environment. Unfortunately, determining the necessary and sufficient conditions under which a set of LDPL equations has a unique solution is still an open problem.

[0059] However, for most practical scenarios, it is possible to determine whether or not a system of LDPL equations can be uniquely solved by making sure that following three conditions are satisfied.

[0060] C1: Each unknown location must see at least 3 APs.

[0061] C2: Each AP must be seen from at least 5 locations (known or unknown)

[0062] C3: The Jacobian of the system of LDPL equations must have a full rank (equal to the number of variables) for a random choice of the AP parameters and unknown locations.

[0063] Conditions C1 and C2 follow from the discussion in Section 1.5.2.1 and 1.5.2.2. Condition C3 essentially linearizes the system of LDPL equations into the form $Jy=k$, where J is the Jacobian, y is a vector containing all unknown parameters and k a constant vector. The lack of full rank then exposes any insufficient coupling in the equations.

[0064] 1.5 Computational Challenges

[0065] While there are several different approaches to solving a set of over-determined equations, in one implementation, the technique attempts to find a solution that minimizes the least mean absolute error,

$$J_{Ez} = \frac{1}{N} \sum_{ij} |P_{ij} - P_i^0 + 10 * \gamma_i \log t_{ij}| \quad (3)$$

In Eqn 3, N is the total number of equations.

[0066] J_{Ez} is a non-linear objective function and does not appear to allow for an analytical closed form solution. Optimization schemes such as the Newton Raphson Method or Gradient Descent (GD) are iterative schemes that start from an initial guess and find the closest local minimum. It was found that such schemes fail to find a solution to J_{Ez} since the number of local minima in J_{Ez} is immense. The other alternatives are search techniques such as simulated annealing or genetic algorithms. While genetic algorithms can search the solution space efficiently, they can miss local minima that might provide a reasonably good solution. Consequently, to obtain the benefits of both these approaches, in one imple-

mentation of the technique, a hybrid approach is used that used gradient descent to refine the solutions generated by a GA.

[0067] Solving LDPL equations using the GA can take a few minutes to several hours depending on the size of the problem. However, these equations need only be solved once (or periodically once every a few days to refresh the model) to determine the AP locations, their transmit powers and the path loss exponents. Once having estimated the model, new location queries can be answered through standard techniques such as trilateration (by converting RSS measurements to distances in Eqn 2) in real-time. It should be noted that although various embodiments of the simultaneous localization and RF modeling technique described herein employ a LDPL model, other derivatives of the LDPL model could be employed with the technique with other parameters that might also be estimated.

[0068] 1.5.1 The Genetic Algorithm (GA)

[0069] The Genetic Algorithm (GA) employed in one embodiment of the technique starts by picking an initial set of solutions (initial generation) randomly and refining them using gradient descent. A solution consists of a vector of values of all the unknowns to be solved in the LDPL equations. The fitness of each solution is then evaluated by computing

$$\frac{1}{J_{Ez}}$$

or the solution. Thereafter, consecutive generations of solutions are generated in the following manner:

[0070] 1. 10% of the solutions with the highest fitness are retained.

[0071] 2. 10% of the solutions are randomly generated.

[0072] 3. 60% of the solutions are generated by picking two solutions S_1^{old} , S_2^{old} from the previous generation (parent solutions) and mixing them using a random convex linear combination. In other words,

$$S^{new} = a \cdot S_1^{old} + (1-a) \cdot S_2^{old} \quad (4)$$

Here, a is a random vector with each element independently randomly drawn from (0,1), $\mathbf{1}$ is vector with all its elements equal to 1 and \cdot represents a vector dot product. The newly generated solution is then refined using a Gradient Descent (GD) technique.

[0073] 4. The remaining 20% solutions are generated by randomly picking a solution from the previous generation and perturbing it (e.g., using perturbation based mutation) by adding (or subtracting) random values (drawn from an exponential distribution to allow occasional large perturbations) to all the locations, P_i and γ_i . The solutions are then refined using the GD.

[0074] As generations evolve, solutions with higher fitness are discovered. The process terminates the GA when no improvement is found for ten consecutive generations.

[0075] 1.5.2 Reducing the Search Space

[0076] When the solution space is extremely large, a randomly picked solution is likely to be far from the optimal solution. Consequently it may take a large amount of time before the GA discovers the optimal solution. Narrowing the search space can dramatically reduce running times. The most obvious way to narrow search space is to limit the search space of the variables. For example, knowing the dimensions

of the floor one can limit the search of the locations to within the floor perimeter. For AP transmission powers the technique chooses a generous search space namely $(-50,0)$ dBm. For γ_i the technique chooses the search space $(1.5,6.0)$. It is believed that these ranges will be accommodating for most practical indoor deployments.

[0077] Another way to narrow the search space is to leverage constraints inherent to the problem. Given m APs and n locations, as previously discussed there are a total of $4m+2n$ variables to be picked randomly. However, having picked all the $4m$ AP parameters the $2n$ unknown locations can be determined through trilateration (using Eqn 2 to convert RSS to distances). Thus, the GA needs to randomly pick only $4m$ unknowns rather than $4m+2n$. Eliminating each variable to be picked randomly exponentially reduces the search space.

[0078] The search space can be further reduced by using the already determined (or known) locations. For example, suppose that a particular AP can be seen from three known locations. Then after randomly picking P and γ , its location can be uniquely determined. In general, the underlying constraints in the LDPL equations can be listed as follows:

[0079] R1: If an AP can be seen from five or more fixed (or determined) locations, then all four of its parameters can be uniquely solved.

[0080] R2: If an AP can be seen from four fixed locations, there exist only two possible solutions for the four parameters of the AP.

[0081] R3: If an AP (say i^{th} AP) can be seen from three fixed locations. Then after picking γ_i randomly from $(1.5,6.0)$ there exist only two possible solutions that satisfy the observations for the AP location and its transmission power.

[0082] R4: If an AP can be seen from two fixed locations, then, having picked P_i and γ_i randomly, there will be only two possible locations for its location.

[0083] R5: If an AP can be seen from one location only, then, after picking P_i and γ_i randomly, the AP can only lie on a circle of radius given by Eqn 2 centered about the known location.

[0084] R6: If the parameters for three (or more) APs have been fixed, then all unknown locations that see all these APs can be exactly determined using trilateration.

[0085] Constraints R0-R6 indicate that after selecting only a few of the variables randomly, the rest can be deterministically computed. The smaller the number of randomly chosen variables, the smaller is the search space.

[0086] Based on these underlying constraints in EPL equations, a Random Solution Generation Algorithm (RSGA) was developed. RSGA attempts to minimize the number of variables that need to be randomly picked among the entire set of variables in a given set of equations in a greedy fashion. The essential idea behind RSGA is to start by determining AP parameters with as many known locations as possible. Upon determining AP parameters for three or more APs, locations that can see these APs can be determined using trilateration. These determined locations in turn can be used to determine the parameters of some other APs. The procedure continues until all APs and locations have been determined. The pseudo code for RSGA is provided below:

-continued

```

change = false
for i = 1 to mdo
  if i ∉ C then
    set O' = pij | j ∈ Ldone
    if |O'| ≥ 1 then
      {ci, Pi, γi} = APRandomInit(l, O', Ldone)
      Cdone = Cdone ∪ {i}
      change = true
      if l < base then
        return true
      end if
    end if
  end if
end for
if change = false then
  if l > 0 then
    change = RSGA(Ldone, Cdone, O, l - 1, base)
    if l < base then
      Return change
    end if
  else
    Return change
  end if
end if
for j = 1 to ndo
  if j ∉ L then
    set O' = pij | i ∈ Cdone
    if |O'| ≥ 3 then
      xj = Trilaterate(O', C)
      Ldone = Ldone ∪ {j}
      change = true
    end if
  end if
end for
until change = true
Return change

```

[0087] RSGA is a recursive algorithm that takes five inputs. L_{done} is the set indices of all locations where RSS observations were taken that have been determined so far. C_{done} is the set of indices of APs with their AP parameters determined. O is the set of RSS observations $p_1 \dots p_{n,1} \dots p_{m,1}$. l is the recursion level which searches for the constraint R_l . $base$ takes a value of 5 if there are 5 or more known locations, otherwise it takes the value of the number of known locations. The entire procedure begins by initializing L_{done} with indices of all the known locations, C_{done} as an empty set and $l=base$. The function $APRandomInit(l, O', L_{done})$ finds a set of random AP parameters given l determined (or known) locations based on constraint rule R_l . For each value of l in $APRandomInit(l, O', L_{done})$ a different strategy is used to determine the random AP parameters based on constraints R2-R6. For example, in case of R2 or R3 ($l=4,5$), the value of the AP parameters is determined through an exhaustive search over several combinations of P and γ in combination with trilateration. In case of R4 ($l=3$), γ is chosen randomly and P is searched exhaustively to determine local minima in the mean absolute error. At $l=0$, all AP parameters are generated randomly.

[0088] 2.0 Challenges in Real Environments

[0089] In this section the novel practical challenges addressed by one embodiment of the technique are addressed.

[0090] 2.1 Selecting the Right Set of Access Points

[0091] In one working embodiment of the simultaneous localization and RF modeling technique, the sheer number of wireless APs that could be seen on a given floor in test deployments came as a rather unexpected surprise. For example, in one deployment one could see a total of about 160 APs across a single office floor. A large fraction of these APs actually belonged to neighboring office buildings. Another interesting

observation was that often each AP was configured with multiple Service Set Identifiers (SSIDs) and appeared as different APs. Running the technique on all observed APs would constitute a computational hardship at the cost of incremental gains. Consequently, an automated AP filtering mechanism is employed by one embodiment of the technique that analyzes the data and selects the most suitable APs. Note that the technique has no information as to if these APs belonged to the indoor environment in question or not.

[0092] Several naive approaches to AP selection can be designed based on desirable properties such as coverage, low standard deviation in RSS, and high average signal strength. For one embodiment of the technique however, the goal was to minimize the number of selected APs while preserving performance. To this end, an access point selection algorithm, herein called APSelect, was developed. The essential idea behind APSelect is to select each AP to provide information that other selected AP do not. APSelect however, uses a more approximate and simpler approach. It starts by computing a similarity metric for each pair of APs based on observed RSS. APs with the most similar data are then clustered together. A representative AP is then elected from each cluster.

[0093] To compute the similarity metric, all the RSS observations p_{ij} from the i^{th} AP at the j^{th} location are first normalized to lie within the range (0,1) by dividing them by 100 (since observed RSS typically lie in the range 0 to -100 dBm) and then their mean is subtracted from each reading to give a normalized RSS observation $p_{ij}^{\text{normalized}}$. When RSS readings are not available at certain locations, these gaps are filled with a reading of -100 dBm assuming that the RSS is below the receive threshold of the receiver. The similarity metric is then computed as,

$$\lambda_{ij} = 1 - \frac{1}{n} \sum_k |\hat{p}_{ik} - \hat{p}_{jk}| \quad (5)$$

λ_{ij} essentially measures how similar RSS readings were across all locations (known and unknown) for the two APs.

[0094] For clustering similar APs, hierarchical clustering is used. Initially each AP represents a single cluster and at each following step two of the most similar clusters are merged. The similarity between two AP clusters is computed as the average similarity between all inter-cluster pairs of APs. For selecting a representative AP within a cluster all APs are first ranked in the order of the number of known locations that can be seen by the AP. An AP that can see a larger number of known locations is given a higher priority. Among APs with same priority, an AP which has the highest average similarity to rest of the APs in its cluster is selected as the cluster head. In one embodiment the minimum number of clusters is selected while ensuring that no two clusters have a similarity greater than 90%.

[0095] 2.2 Selecting a Subset of Locations

[0096] For creating the RF model, ideally RSS observations come from several different locations across the indoor environment. For example, data from a single mobile user who actively ventures to different places across the indoor space to different locations while his mobile device transmits the RSS information to a server. Alternatively, data can come from a large number of mobile device users spread out across

various locations of the indoor space. Given the large amount of data, then, the technique culls out the “useful” subset of data.

[0097] To tackle this problem a location selection algorithm, herein called LocSelect, was developed. LocSelect works in a similar manner as APSelect, except that each new location is selected such that the RSS information it provides has minimum overlap with other selected locations.

[0098] 2.3 Receiver Gain Differences

[0099] While ideally all mobile devices should measure the same RSS, they often do not. The average RSS received by different mobile devices at the same location, during the same time interval and oriented along the same direction at line-of-sight from an AP can have large differences. Such differences can arise from the difference in receiver gains of these mobile devices and calibration errors. Since several different mobile devices can participate in the simultaneous localization and RF modeling technique, such differences can potentially increase localization errors. It has been suggested to maintain a database of pre-measured differences in receiver gains among various well known mobile devices. However, it has been found that there are gain differences of 2-7 dB even among two devices of the same model.

[0100] To account for these gain differences an additional unknown parameter—the receiver gain G , for each user, is introduced in one embodiment of the technique. In other words the LDPL model becomes,

$$p_{ij}^k = P_r - G^k + 10 * \gamma_i \log d_{ij}^k + R \quad (6)$$

Here, the superscript k indicates data specific to the k^{th} mobile device and G^k is the receiver gain of the mobile device. G^k can then be estimated using the genetic algorithm along with all the other parameters.

[0101] In one implementation of the technique a novel approach is taken to estimate G^k . The basic idea is that if two users happen to transmit RSS observations from the same (or nearby) locations, then, the difference in gains between the two users can be estimated by subtracting their observations. In practice, however, the technique does not know the locations from which the RSS data was obtained (unless a corresponding GPS reading was obtained). Hence, the technique must rely on the RSS observations to ascertain that the observations from the two users were indeed taken from the same or close locations. To this end, the technique uses the fact that proximity between two locations can be inferred from the likelihood of the RSS readings seen at these locations. Suppose the RSS readings seen by two users are $\langle p_1^1, p_2^1, \dots, p_m^1 \rangle$ and $\langle p_1^2, p_2^2, \dots, p_m^2 \rangle$ respectively, then the effect of the receiver gains can then be eliminated by computing the vectors $\langle p_2^1 - p_1^1, \dots, p_m^1 - p_1^1 \rangle$ and $\langle p_2^2 - p_1^2, \dots, p_m^2 - p_1^2 \rangle$. The technique uses the distance between these vectors to infer the likelihood that these readings came from locations that were close by. By creating pairs of readings from the two users' RSS data, that are most likely to have come from the same location, the technique estimates the difference in gains from between the users' receivers. These estimates can then be used as constraints in the genetic algorithm.

[0102] When a new mobile user enters the indoor space, the trilateration based localization (after using AP parameters to convert RSS to distances from the APs) can lead to errors if the receiver gains are not taken into account.

[0103] 3.0 The Computing Environment

[0104] The simultaneous localization and RF modeling technique is designed to operate in a computing environment.

The following description is intended to provide a brief, general description of a suitable computing environment in which the simultaneous localization and RF modeling technique can be implemented. The technique is operational with numerous general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable include, but are not limited to, personal computers, server computers, hand-held or laptop devices (for example, media players, notebook computers, cellular phones, personal data assistants, voice recorders), multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0105] FIG. 6 illustrates an example of a suitable computing system environment. The computing system environment is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the present technique. Neither should the computing environment be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment. With reference to FIG. 6, an exemplary system for implementing the simultaneous localization and RF modeling technique includes a computing device, such as computing device 600. In its most basic configuration, computing device 600 typically includes at least one processing unit 602 and memory 604. Depending on the exact configuration and type of computing device, memory 604 may be volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. This most basic configuration is illustrated in FIG. 6 by dashed line 606. Additionally, device 600 may also have additional features/functionality. For example, device 600 may also include additional storage (removable and/or non-removable) including, but not limited to, magnetic or optical disks or tape. Such additional storage is illustrated in FIG. 6 by removable storage 608 and non-removable storage 610. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Memory 604, removable storage 608 and non-removable storage 610 are all examples of computer storage media. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by device 600. Computer readable media include both transitory, propagating signals and computer (readable) storage media. Any such computer storage media may be part of device 600.

[0106] Device 600 also can contain communications connection(s) 612 that allow the device to communicate with other devices and networks. Communications connection(s) 612 is an example of communication media. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more

of its characteristics set or changed in such a manner as to encode information in the signal, thereby changing the configuration or state of the receiving device of the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. The term computer readable media as used herein includes both storage media and communication media.

[0107] Device 600 may have various input device(s) 614 such as a display, keyboard, mouse, pen, camera, touch input device, and so on. Output device(s) 616 devices such as a display, speakers, a printer, and so on may also be included. All of these devices are well known in the art and need not be discussed at length here.

[0108] The simultaneous localization and RF modeling technique may be described in the general context of computer-executable instructions, such as program modules, being executed by a computing device. Generally, program modules include routines, programs, objects, components, data structures, and so on, that perform particular tasks or implement particular abstract data types. The simultaneous localization and RF modeling technique may be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

[0109] It should also be noted that any or all of the aforementioned alternate embodiments described herein may be used in any combination desired to form additional hybrid embodiments. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. The specific features and acts described above are disclosed as example forms of implementing the claims.

1. A computer-implemented process for performing localization of one or more mobile devices, comprising:

using a computing device for:

recording measurements of received signal strength indication (RSSI) from each Radio Frequency (RF) transmitter reported by a mobile device;

recording measurements of any available location fix by the mobile device;

using a RF propagation model, prepared without prior knowledge of the physical space in which the mobile device is located or the location and power of the RF transmitters, to estimate parameters comprising: the location of the mobile device, the transmit power of each RF transmitter, and the location of each RF transmitter based on the recorded measurements of received signal strength indication and any available location fixes; and

using trilateration to determine the location of the mobile device from the estimated location of the RF transmitters and the estimated parameters.

2. The computer-implemented process of claim 1, wherein determining the location of the mobile device further comprises:

translating the RSSI measurements reported by the mobile device into distances from the corresponding RF transmitters whose locations are estimated with the RF propagation model, and

using the distances from at least three RF transmitters to compute the location of the mobile device.

3. The computer-implemented process of claim 1 further comprising creating the RF propagation model using a log distance path loss model.

4. The computer-implemented process of claim 1, further comprising using a combination of a gradient descent approach and a genetic algorithm to estimate the parameters using the RF propagation model.

5. The computer-implemented process of claim 1 wherein the location of any available location fix is obtained by a global positioning system device.

6. A computer-implemented process for performing localization, comprising:

- using a computing device for,
 - recording at one or more mobile devices, RSSI measurements corresponding to RF transmitters received by the one or more mobile devices at various unknown locations;
 - reporting the recorded RSSI measurements to a remote server;
 - recording any available GPS lock at the one or more mobile devices;
 - reporting any GPS lock to the remote server;
 - using the RSSI measurements and any GPS locks recorded at the one or more mobile devices to determine the location of the one or more mobile devices, without prior knowledge of the physical space in which the mobile devices are located or the location and power of the RF transmitters, using the constraints of the physics of RF propagation at the remote server; and
 - reporting the determined location of the one or more mobile devices.

7. The computer-implemented process of claim 6 wherein the RF transmitters are located in an indoor environment.

8. The computer-implemented process of claim 6 wherein the RF transmitters are located in an outdoor environment.

9. The computer-implemented process of claim 6 wherein only one mobile device traversing the environment over time is used to obtain the location of the mobile device.

10. The computer-implemented process of claim 6, wherein the constraints are modeled using a log distance path loss model.

11. The computer-implemented process of claim 10 wherein the constraints modeled using the log distance path loss model are solved using a genetic algorithm.

12. The computer-implemented process of claim 9, wherein the constraints further comprise:

- location of each RF transmitter;
- transmit power of each RF transmitter; and
- a decay constant of the log distance path loss model.

13. A system for determining locations of a Wi-Fi capable device in an indoor environment, comprising:

- a general purpose computing device;
- a computer program comprising program modules executable by the general purpose computing device, wherein the computing device is directed by the program modules of the computer program to,
 - periodically scan for Wi-Fi access points using a Wi-Fi enabled mobile device;
 - transmit a list of discovered Wi-Fi access-points and associated received signal strength indication (RSSI) measured by the mobile device to a remotely located server;
 - transmit any GPS locations discovered by the mobile device to the remotely located server;
 - use a RF model of the indoor environment, created without prior knowledge of the space the mobile device is located in or the location and power of the Wi-Fi access points, using the list of discovered Wi-Fi access points and associated received RSSI and the discovered GPS locations to provide the location of the mobile device to user of the mobile device.

14. The system of claim 13 wherein the location of the mobile device is provided to the user without action by the user.

15. The system of claim 13 wherein the location of the mobile device is determined without any prior knowledge of the location of the mobile device, the location of the Wi-Fi access points or knowledge of the RF power of the Wi-Fi access points.

16. The system of claim 13 wherein the RSSI values are used to determine relative distances between Wi-Fi access points and wherein the relative locations between the Wi-Fi access points are used to determine the location of the mobile devices.

17. The system of claim 13 wherein determined locations of the Wi-Fi access points are used to find unknown locations of the Wi-Fi access points.

18. The system of claim 17 wherein only a subset of the determined locations are selected for use to find the unknown locations, wherein the subset of the selected locations are selected by removing redundant determined locations.

19. The system of claim 13 further comprising a module for using a subset of the list of discovered access points that do not provide redundant information when using the created RF model to provide the location of the mobile device to a user.

20. The system of claim 13 wherein gains of the mobile devices are estimated by using RSS readings of pairs of mobile devices that are most likely at the same location.

* * * * *