Vision: Cloud and Crowd Assistance for GPS Urban Canyons

[1]Zhiyong Tan [2]David Chu [3]Lin Zhong
[1]NativeX, San Francisco, CA 94103 [2]Microsoft Research, Redmond, WA 98052
[3]Dept. of ECE, Rice University, Houston, TX 77005

ABSTRACT
Location is a core component of today's mobile devices and apps. While GPS is widely used, it continues to perform poorly when satellite visibility is obscured, e.g. in urban canyons. This paper explores the potential benefits of cloud- and crowd-assisted GPS for tackling urban canyons. Specifically, we present early results on utilizing high resolution environment scans such as Digital Surface Models (DSMs) to aid GPS calculations. Second, we suggest initial designs for leveraging location corrections from other mobile device users to enhance noisy GPS readings.

Categories and Subject Descriptors
C.2.1 [**Computer Systems Organization**]: Computer-Communication Networks – *wireless communication*

Keywords
GPS, mobile computing, localization, urban canyon

1. Introduction
GPS signals are not good at penetrating dense materials like walls and windows [1]. This is primarily because GPS signals, at 1575.42 MHz and 1227.60 MHz, have low skin depth when hitting building surfaces. As a result, signal from a GPS satellite can be completely blocked by buildings on the ground. Such circumstances are common in downtown areas where streets are surrounded by high rises as seen in Figure 1. This problem is known as the GPS *urban canyon*. When a mobile user in an urban canyon opens the map application, she is likely to see an estimated position with an uncertainty zone of a large radius as shown in Figure 2.

Because a GPS receiver needs signals from at least four satellites to determine its location accurately, when buildings block signals from too many GPS satellites, the receiver will suffer significant location inaccuracy. Moreover, even if line-of-sight with four satellites can be established, buildings and other tall structures limit the line-of-sight satellites to a set of closely-parallel lines, causing what is known as a high dilution-of-precision (DOP) and reducing accuracy. This is illustrated in Figure 3. As GPS receivers are now widely used by urban pedestrians and drivers, urban canyons have garnered increasing attention.

In urban canyons, the radius of uncertainty zone can be much larger than one or two blocks. That is, the real location of the receiver can be a few blocks away from the estimated location, the center point. Such a large error is problematic for both pedestrians and drivers seeking navigational assistance. For example, a pedestrian coming out of a building may have a problem determining

Figure 1: A typical urban canyon.

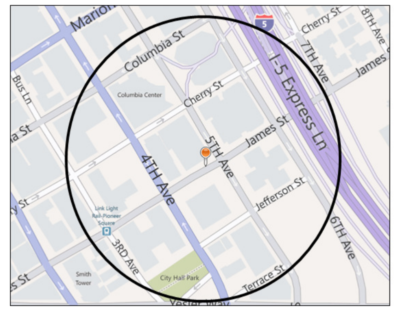

Figure 2: Large location uncertainty occurs in an urban canyon. Inaccuracy circle can be larger than shown here.

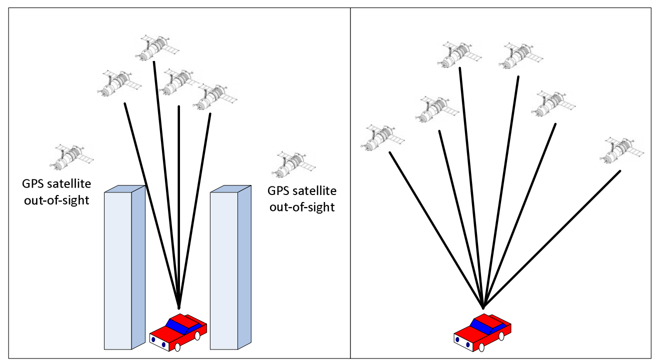

Figure 3: Illustration of Dilution-Of-Precision (DOP). The receiver on the left pane has poor DOP, because the satellite-receiver paths are close to parallel. The receiver on the right has good DOP.

which side of the building he came out of; the GPS navigation system in car may recalculate the route, mistakenly thinking the car is off the original route, and therefore confuse the driver following its guidance.

In addition, urban canyons can lead to high energy usage by GPS receivers. GPS receivers are known to be power-hungry, especially when they are searching for satellite signals. Without knowing it is inside an urban canyon, a GPS receiver in a mobile device may keep on searching for the fourth satellite needed for an accurate location, draining battery unnecessarily.

In this work, we propose using both the cloud and crowds to assist localization in urban canyons. First, we propose augmenting the capabilities of the GPS receiver with additional data from the Digital Surface Model (DSM), a high resolution 3D scan of the urban environment. By using a DSM, the existence of line-of-sight for every satellite to a receiver at any point in the DSM can be determined. Constructing a DSM of a region with urban canyons requires Light Detection and Ranging (LIDAR) from aerial surveys, which can be expensive especially if high accuracy (up to 50 mm vertical accuracy) is demanded. Recently, lower-cost methods such as stereographic photogrammetry have become much more commonplace due to DSM usage in online mapping applications.

MCS'14, June 16, 2014, Bretton Woods, New Hampshire, USA.

http://dx.doi.org/10.1145/2609908.2609950

We use the cloud's computation resources to calculate possible signal availability given a DSM and positions of the satellites. Distinct from previous work [2], [3], [4], [5], [6], [7] we also consider signal reflections, since line-of-sight signal propagation alone is insufficient to characterize GPS communication. We then propose several ways in which seeing or not seeing particular satellites can be useful for users.

We have performed preliminary tests of our DSM-assisted GPS in several areas of downtown Houston. Experimental verifications with line-of-sight approaches used in prior work were done, verifying their low accuracy and generalizability. In addition, we also attempted to consider geometrical factors beyond line-of-sight and specular reflection, by searching for heuristics involving the DSM which could lead to GPS satellite visibility prediction in urban canyons, and which can be applied with good accuracy across a wide range of possible situations.

The generalizability of all approaches were extensively tested by generating predictions of satellite visibility across multiple times spaced apart in the satellite orbit period and for multiple locations, then experimental verification of all predictions by comparison with satellite visibility data collected at the simulated times and locations.

Second, we propose the use of crowdsourced GPS readings to help in urban canyon situations. The main idea is that user-corrected GPS positions can be shared amongst users. To this end, we have built a prototype mobile app that lets users share GPS corrections. The main technical challenge lies in determining when one users' GPS correction should be applied to another users'. Here, we present an early design, sky map matching, which attempts to further the goal of crowdsourcing GPS reading corrections.

2. Related Work

Many approaches have been attempted to address the urban canyon problem. GPS devices with inertial measurement units (GPS-IMU) have been an active field of research for years and are widely used (see, e.g. [3]). An IMU allows navigation in regions of poor GPS signal receptivity like urban canyons. A major disadvantage of IMU is that they suffer from accumulated error. When unable to access GPS signals for a prolonged period, a receiver will accumulate considerable error.

In driving navigation units, a road map is used to improve location estimation by assuming that the receiver is located on a road. If the location estimation states that the user is off the road by less than an arbitrary distance (on the order of 50 ft), the location estimation is adjusted to the nearest location on a road. However, this method cannot be used in pedestrian GPS receivers (such as those in smartphones) and will not work when the error is larger than a street block, as the unit may erroneously point to a wrong road [9]. Many of these existing solutions are often used simultaneously in a GPS unit. Our proposed solution serves to further reduce the inaccuracy of user positioning in urban canyons. It can be used in combination of existing solutions.

By use of a precise clock, the number of satellites required can be reduced from 4 to 3. This classic solution has not been widely implemented due to the high cost, large size, and high power consumption of atomic clocks. The recent commercial development of chip-scale atomic clocks [10] has led to greater interest in their use in GPS receivers. The US military is currently investigating the use of chip-scale atomic clocks for rapid recovery of location estimates after signal loss due to the receiver moving under blockages and then back into (sufficiently) open sky, or from signal jamming.

Several groups have attempted to use assumptions and approximations for drivers to continue obtaining location estimates while driving through urban canyons. Chang et al. [11] combined a pseudorange predictor, receiver clock bias predictor, an altitude-hold algorithm, and assumed constant vehicle speed. Their technique allowed accurate horizontal tracking through a straight tunnel. Cui and Ge [12] adopted an approach similar to map-matching, by using the receiver position on a map prior to signal loss and assuming that the receiver continues to travel on the same road. They used a mathematical model to represent the road curvature, and extended Kalman filter for location estimates. Their simulation provides location estimates accurate to 3 meters using a road consisting of a convex loop.

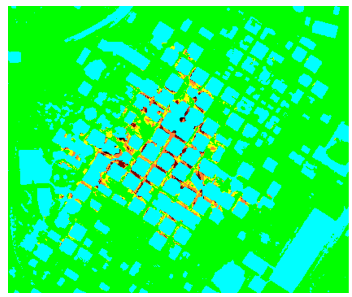

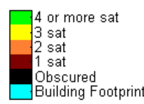

Figure 4 (Left): Simulation of S2 results for downtown Houston.

Figure 5 (Right): Locations of receiver samples in downtown Houston.

3. Cloud Assistance with DSMs

Next, we explain in principle how the availability of a GPS satellite can be obtained via simulation off-line in the cloud.

3.1 GPS Satellite Signal Predictions

We denote the set of satellites whose signals are observed by a GPS receiver at a given time and location as M, and identify interesting sets of satellites, S1, S2, S3 and S4 that can aid in more accurate and energy efficient GPS utilization. Note that all sets are location- and phase-dependent which we omit for purposes of exposition.

3.1.1 LOS Satellites without DSM (S1)

First, with the satellite ephemeris information, one can obtain the set of satellites, S1, whose signals will be available at a given time at any location, without considering terrestrial structures. As the ephemeris inaccuracy of a GPS satellite increases with its age, and the ephemeris message requires at least 30 seconds to transmit, the time-to-first-fix (TTFF) of a receiver is shorter if the receiver has been recently used.

Let the set S1 consist of satellites with a minimum of 10 degree elevation above the horizon, consistent with an elevation mask of 10 degrees set in most GPS receivers. S1 can be obtained by using line of sight ray tracing comparing the position of each satellite relative to the map center. Exact receiver location is not required. Thus, we obtain $M \subseteq S1$.

S1 can also be derived from the almanac component of the GPS message. However, transmission of the almanac takes 12.5 minutes and is thus relevant only for decades-old receivers with slow GPS signal acquisition. Modern receivers, with their faster acquisition time, usually search for all GPS satellites.

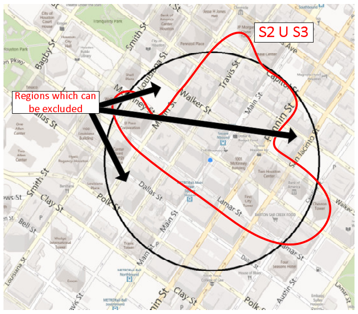

Figure 6: A proposed approach to reducing the location estimate accuracy radius. Current GPS apps have large error estimates in urban canyons. By obtaining S3, portions of the accuracy circle can be excluded.

3.1.2 LOS Satellites with DSM (S2)

Second, if the terrestrial structure information (DSM) is also given, one can obtain the set of satellite, S2, whose signals will be available line of sight at a given time to any location. This is the first logical step in considering how geometry affects satellite visibility, as explored in prior work.

To derive S2, first mark the receiver and S1 positions on the DSM. Next, draw a straight line between the receiver and each satellite. If there is line-of-sight between a receiver and satellite, the line does not intersect any buildings in the DSM. S2 consists of the satellites which have line-of-sight to the receiver. Therefore, $S2 \subseteq S1$, and furthermore $S2 \subseteq M \subseteq S1$ due to non-line of sight satellites as discussed in the next section.

The long distance from satellite to receiver (20200 km at zenith) ensure that errors in satellite elevation angle from broadcast ephemeris inaccuracy (approximately 1 m) are trivial. The accuracy of determining S2 is strongly dependent on DSM accuracy. The DSM used in this study had 5 ft. horizontal resolution and 1 ft. vertical resolution [13].

In order to efficiently calculate S2, we used Bresenham's line algorithm to determine the DSM pixels intersected between the receiver's location (assumed to be 1 m above ground) and the satellite. This computation is well-suited for cloud offload, especially as the spatial and temporal resolution of satellite visibility is increased.

3.1.3 NLOS Satellites with DSM (S3 & S4)

Third, if one further considers reflection of GPS signals by terrestrial structures (e.g. buildings, elevated highways, etc.), one can obtain the set of satellites, S3, whose signals travel along a non-line of sight path at a given time to a particular location. S3 is of interest because previous work by Taylor et al. [13], as well as our own investigation, reveals that S2 tends to present an overly pessimistic picture of GPS satellite visibility, with fewer satellites than what is actually detectable at ground truth.

Observe that $S3 \subseteq S1$ as well as $(S2 \cup S3) \subseteq S1$. In the general case, reflection is challenging because the degrees of specular and diffuse reflection are dependent on complex factors such as building material and density. As a convenient abstraction, we use a single parameter θ to denote how aggressively reflection is considered about the specular reflection path. A larger θ indicates more reflective paths, reflective paths of more reflection, and more diffusive reflection are considered. Therefore, we have $S3(\theta_1) \subseteq S3(\theta_2)$ if $\theta_1 \leq \theta_2$. Ideally, we should have $S2 \cup S3 = M$. In practice, we need to find θ_0 such that $S2 \cup S3(\theta_0) \approx M$. It is important to note that θ_0 is determined not only by the simulation details but also by the sensitivity and computational power of the GPS receiver. For our current study, we calculated θ_0 empirically based on an iterative search to minimize false positives and false negatives.

To aid the search, we hypothesized that detection probability of satellites without line-of-sight was proportional to the total area of possible reflection surfaces or to signal reflection angle. The area of possible reflection surfaces was calculated by ray tracing the DSM points with line-of-sight to both the satellite and the receiver. Reflection angle was determined using ray tracing techniques.

Finally, while it can be hard to derive $S3(\theta)$, $S3(\infty)$ can be derived much more readily. With the terrestrial structure information and satellite ephemeris information, one can easily derive S4, the subset of S1 whose signals will be definitely blocked by the terrestrial structures. Therefore we have $S3(\infty)=S1/(S2 \cup S4)$. We have $S2 \subseteq M \subseteq S1/S4 \subseteq S1$.

3.2 Uses for GPS Satellite Information

In this section, we propose multiple ways to utilize GPS satellite information for various purposes when satellite signal prediction is available at various levels.

3.2.1 S1 is Available

Current GPS receivers attempt to acquire all satellites, generating the pseudorandom number (PRN) codes for all satellites. A smartphone can download the almanac from the server, reducing the number of satellites to be acquired from 31 to 8-12. As GPS is a power-hungry application, searching for fewer satellites will improve energy efficiency. Searching can terminate early once M=S1 have been found, thereby saving energy. This strategy is particularly effective if $|M| > 4$ such that all satellites in S1 are contributing line-of-sight position information.

3.2.2 S2 is Available

Next, S2 can be determined from S1 and ray tracing with the DSM. Figure 4 shows an example of S2 for downtown Houston at a given time, which was determined using the technique described in Section 3.1.2. An immediate application of determining S2 will be to inform receivers which satellite signals can be trusted more than others. Specifically, S2 satellites should be favored in computing location because they are known to be line-of-sight. If the receiver's precise location is known immediately prior to entering an urban canyon, determination of receiver location within the urban canyon may be improved by application of a weighting function with weight biased toward S2. In our investigation, we confirmed that satellite visibility is dependent on the extent of deviation of line-of-sight from the satellite.

In a scenario where the user does not need to enter the dead zone, but is simply driving through en route to a destination, the receiver may direct him to avoid the regions of urban canyon where $|S2| < 4$, preventing him from losing his way in the urban canyon.

3.2.3 S4 is available

Given a list of possible urban canyon dead zones, it can be useful to narrow down the list of zones in which the receiver may be. S4 can be used to reduce the number of dead zones in which a receiver is located. If a receiver acquires any satellite in S4, this indicates that it is not in that region at all. In this case, it is the absence rather than the presence of a satellite that provides location information. In essence, we can rule out locations where $M \subseteq S1/S4$ does not hold. This information can be used to improve navigation, e.g. routing despite the lack of sufficient satellites in a canyon.

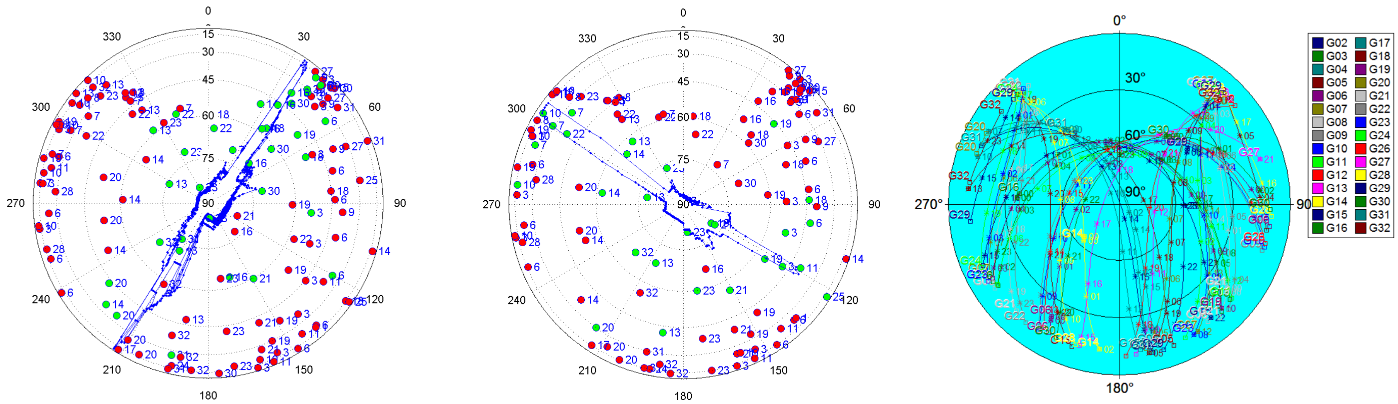

Figure 7: (Left and Center) Sky plots of data from two downtown Houston locations (imagine looking skyward). Green spots represent satellites visible in measurement. Red spots represent satellites nonvisible in measurement. Blue line indicates DSM (building) envelope. (Right) Simulated sky plot of the same day.

3.2.4 $S3(\theta_0)$ is Available

In current smartphone GPS applications, when the number of acquired satellites is < 4, the user is presented with a 1 km accuracy radius centered on the last known location. Knowledge of which regions have S3 < 4 can reduce the accuracy radius to the extent of the urban canyon where S3 < 4, aiding the user by narrowing the range of possible locations he is in. See Figure 6 for an illustration. Similarly, if we have $S2 \cup S3(\theta_0) \approx M$, we can further improve the energy efficiency and location accuracy of the GPS. In this case, we can terminate satellite search early.

3.3 Early Results

Here we report early results for validating our solutions for deriving S1, S2 and S3.

3.3.1 Measurement Setup

The measurement setup consists of a Garmin GN3SV3 receiver connected to a GPS ANT-555 antenna. A software receiver was chosen for full control of the GPS signal processing, unbiased by correction algorithms which may be present in commercial units, such as map-matching or use of WiFi or GSM.

Data collection was done at 15 points in downtown Houston, shown on Figure 5. The points were chosen based on the heights of surrounding buildings. Sampling times were during the day between 8 AM and 8 PM, with each sample taking at least a minute to capture an entire ephemeris record. A total of 273 measurements were taken.

3.3.2 S1 and S2

S1 can be easily computed. Surprisingly, we have not encountered any commercial implementations of S1, despite the energy savings which are possible by searching only for a subset of satellites instead of the entire constellation.

All experimental data indicate that S2 can be determined fully. Figure 4 illustrates one such example using the downtown Houston DSM. Note that in much of the active center of the city, fewer than four satellites are in S2. The small number of satellites comprising S2 indicate that effort must be undertaken to determine S3 for useful applications (such as power-saving applications) to be created.

3.3.3 S3

One primary challenge is to determine the reflections which are sufficiently strong for detection. Figure 7 shows sky plots of measurements at two downtown Houston locations out of the 16 locations we measured. Examining the sky plot, we make several observations.

- Detected satellites (green dots) can be outside the DSM envelope. Some can be very far outside (e.g. #6 in the Left Skyplot).
- Undetected satellites (red dots) can be within the DSM envelope.
- The simulated skyplot (Fig 5, Right) is not correlated with actual visibility (Fig 5, Left & Center) in urban canyons.

We are in the process of developing an appropriate criteria to classify visible from invisible satellites. We are incorporating building geometry by considering the relationships between (1) satellite visibility and the surface area of possible reflections, and (2) satellite visibility and the angle of reflection, and (3) satellite visibility and satellite azimuth and elevation. With simple threshold-based classification criteria, experiments have yielded roughly equal numbers of true positives (satellite detected both in simulation and measurement) and false positives (satellite detected in simulation but undetected in measurement). With tuning, the false positives for a particular sampling location can be minimized to approximately ¾ of the true positives. Generalizing such tuning across locations is an active area of investigation.

4. Crowd Assistance for Urban Canyons

Cloud connectivity also implies that a receiver can tap into data collected from other receivers. The possibility of crowdsourced GPS measurements opens up a large design space. We are investigating one specific promising formulation: (1) can users provide corrections to GPS-based location measurements, and (2) can these corrections be automatically shared and applied amongst users for overall improved localization? For example, consider a pedestrian in a canyon who is prompted to correct her receiver's position by means of placing a pushpin corresponding to her true location on a map interface. The user submits this correction together with a *fingerprint* of her receiver's data to the cloud. In its most basic form, a fingerprint can simply be a receiver's observed set of visible satellites. When a subsequent user finds himself in an urban canyon, his device can query the fingerprint database for corrections submitted by other users. If there is a match between fingerprints, the user automatically benefits from the correct position.

To examine whether users can successfully correct GPS-based location measurements, we have developed a mobile app for pedestrians that shows a list of places nearby the current estimated location. Places include shops, restaurants and public buildings and landmarks. The user can look around to spot a place which matches one listed, and select a relative orientation, e.g. *On My Right*.

Figure 8: Screenshots of the User-in-the-Loop mobile app. The left pane shows the receiver inside a building, which is an inaccurate position estimate. The middle pane shows a list of nearby locations and their relative orientations to the user from which the user can make a selection. The right pane shows the corrected position, which is more accurate than the initial location.

Screenshots of the app are shown in Figure 8. Our hope is that if such an interface is easy to use, users already engaged with mobile maps (e.g. for pedestrian navigation) will find submitting corrections to be low overhead.

4.1 Sky View Matching

Provided a database of location fingerprints and crowd-sourced corrections, the question is then how can fingerprint queries be matched efficiently? We introduce *sky view matching* as a means to achieve this.

The premise of sky view matching is that a user u_1 that observes satellites $\vec{s}_1$ can refine her position based on satellite observations $\vec{s}_2$ from user u_2. The simplest case is when u_1 and u_2 submit GPS queries at the nearly the same time. Assuming that coarse-grained localization is available (e.g. from cell base station ID), if $\vec{s}_1 = \vec{s}_2$, then u_1 can simply utilize the position correction submitted by u_2.

Progressing in sophistication, when u_1 and u_2 submit GPS queries that are not at the same time, the change in satellite position must be reconciled. We propose a function $patch_t(s) \rightarrow p$ that converts a set of satellite *s* at time *t* to a set of sky patches *p*. A sky patch represents a small area of visible sky. When applied to a $\vec{s}_2$ at time t_2, the result $\vec{p}_2$ then identifies a time-independent set of sky patches visible to a u_2 at their correct position. Subsequently, when a future user u_1 applies the patch function at time t_1, the match of two users' sky patches allows u_1 to use the corrected position supplied by u_2.

The basic $patch_t()$ can convert a satellite to a sky patch simply by reporting patches as $< elevation, azimuth, tolerance >$ tuples. The match function used is $match_\tau(\vec{p}_1, \vec{p}_2)$ and includes a tolerance degree τ. Two satellites whose elevations (and azimuths) are within τ degrees are considered a match. A smaller τ places a tighter bound on satellite positioning similarity while bigger τ allows greater usage of crowd-sourced data at the expense of noisier matches. Future investigation remains with regards to the appropriate crowd density, temporal resolution and spatial resolution required for successful application.

Finally, it may be possible to further refine location by identifying similarities between the auxiliary data collected by u_1 and u_2. In addition to DoP mentioned above, a typical GPS receiver also reports auxiliary data consisting of per-satellite SNR. This auxiliary data is often discarded by the time it reaches the application layer. One proposed approach of incorporating auxiliary data is in the form of a more detailed fingerprint. Taking this auxiliary data in conjunction with the sky patches, we hope to understand whether two users whose error data is similar (e.g. SNR vectors are similar) determines whether two users are in a similar location.

5. Conclusion

Cloud- and crowd-assisted approaches to the GPS urban canyon problem appear promising. DSMs provide a useful starting point in assessing satellite visibility, and should be considered in conjunction with additional non-line-of-sight signal propagation models. Crowdsourcing a sufficient volume of corrected GPS traces is an interesting next step.

6. Acknowledgments

The authors thank Erjola Buzi and Zhe Ren for their assistance in early development. The authors also thank Paramvir Bahl for helpful discussions.

7. References

[1] R. Klukas, O. Julien, L. Dong, E. Cannon and G. Lachapelle, "Effects of building materials on UHF ranging signals," *GPS Solutions,* vol. 8, no. 1, pp. 1-8, 2004.

[2] E. Costa, "Simulation of the Effects of Different Urban Envi-ronments on GPS Performance Using Digital Elevation Mod-els and Building Databases," *IEEE Transactions on Intelligent Transportation Systems,* vol. 12, no. 3, pp. 819-829, September 2011.

[3] A. Steed, "Supporting Mobile Applications with Real-Time Visualisation of GPS Availability," in *Proceedings of Mobile HCI 2004*, LNCS 3160Springer LNCS 3160, 2004.

[4] J. Li, C. H. Jarvis and C. Brunsdon, "The use of immersive real-time 3D computer graphics for visualisation of dilution of precision in virtual environments," *International Journal of Geographical Information Science,* vol. 24, no. 4, pp. 591-605, 2010.

[5] B. Ben-Moshe, E. Elkin, H. Levi and A. Weissman, "Im-proving Accuracy of GNSS Devices in Urban Canyons," in *Proceedings of CCCG'11*.

[6] J.-Y. Han and P.-H. Li, "Utilizing 3-D topographical infor-mation for the quality assessment of a satellite surveying," *Ap-plied Geomatics,* vol. 2, no. 1, pp. 21-32, 2010.

[7] J. Li, G. Taylor, D. Kidner and M. Ware, "Prediction and visualization of GPS multipath signals in urban areas using LiDAR Digital Surface Models and building footprints," *In-ternational Journal of Geographical Information Science,* vol. 22, no. 11-12, pp. 1197-1218, 2008.

[8] J. Paek, J. Kim and R. Govindan, "Energy-efficient rate-adaptive GPS-based positioning for smartphones," in *Proceed-ings of the 8th international conference on Mobile systems*, and services, New York, NY, USAapplications, and services, New York, NY, USA, 2010.

[9] R. Toledo-Moreo, D. Betaille and F. Peyret, "Lane-Level Integrity Provision for Navigation and Map Matching With GNSS, Dead Reckoning, and Enhanced Maps," *IEEE Trans-actions on Intelligent Transportation Systems,* vol. 11, no. 1, pp. 100-112, March 2010.

[10] W. D. Jones, Chip-Scale Atomic Clock, IEEE Spectrum, 2011.

[11] T.-H. Chang, L.-S. Wang and F.-R. Chang, "A Solution to the Ill-Conditioned GPS Positioning Problem in an Urban Environment," *IEEE Transactions on Intelligent Transportation Systems,* vol. 10, no. 1, pp. 135-145, March 2009.

[12] Y. J. Cui and S. S. Ge, "Autonomous vehicle positioning with GPS in urban canyon environments," in *IEEE* ICRA vol. 2, vol.2, 2001.

[13] G. Taylor, J. Li, D. Kidner, C. Brunsdon and M. Ware, "Modelling and prediction of GPS availability with digital pho-togrammetry and LiDAR," *International Journal of Geo-graphical Information Science,* vol. 21, no. 1, pp. 1-20, January 2007.