



# Urban Localization *and* 3D Mapping using GNSS *Shadows*

Urban localization using GNSS is notoriously imprecise due to blockage of satellite signals by tall structures. However, this very phenomenon of blockage can be exploited to derive valuable (but very noisy) location information by “shadow matching” against 3D maps of the environment. This article describes a real-time, low-cost, cloud-based Bayesian localization and tracking solution that combines shadow information and raw GNSS position estimates with probabilistic motion models to substantially improve urban localization. The Bayesian framework also extends to enable synthesis of 3D maps from crowdsourced GNSS data.

**ANDREW IRISH, DANNY ILAND,  
AND UPAMANYU MADHOW**  
UNIVERSITY OF CALIFORNIA, SANTA BARBARA

Andrew Irish

It is by now well known that GNSS-based localization in built-up urban environments can be extremely inaccurate. This is a fundamental problem that hardware enhancements cannot solve.

A GNSS receiver estimates 3D location and timing from pseudoranges from four or more satellites, assuming that these pseudoranges correspond to direct line-of-sight (LOS) paths from each satellite. In urban canyons, however, the signal from a satellite to the receiver suffers from multipath propagation and shadowing.

Even if the LOS path is available, the pseudorange estimate may be corrupted due to the presence of alternative paths. Furthermore, the LOS is frequently blocked except for satellites at the highest elevations; so, the satellite is either unavailable, or the strongest path seen by the receiver is a reflection off a building, leading to a pseudorange often significantly larger than that for the LOS path. These errors in pseudoranges lead to large errors in localization (e.g., up to 50 meters in high-rise environments such as New York City).

Attempts to enhance GNSS localization using signals from other sources, such as cellular and WiFi, have had limited success: it is difficult to use geometric techniques to infer location in a complex propagation environment, and received signal strength measurements are subject to significant fluctuations due to multipath fading and shadowing.

## Shadow Matching

Over the past few years, several groups

of researchers have realized that we can turn shadowing in urban environments from a bug into a feature, by using information already available in the GNSS receiver regarding the signal-to-noise ratio (SNR) that corresponds to each satellite it sees. If the LOS path from the receiver to a satellite is blocked (i.e., the receiver is in the shadow of a structure), then this SNR is likely to be small. Conversely, if the LOS path is available, the SNR is likely to be high.

If we interpret these SNR measurements with the aid of a 3D map of the environment, then we can significantly reduce localization uncertainty by performing shadow matching, as shown in **Figure 1**. In its simplest form, the concept of shadow matching is summarized as follows:


### Shadow Matching with 3D Maps

1. Compute shadows of nearby obstacles using satellite ephemeris data and 3D maps
2. Classify each satellite as visible (LOS) or blocked (NLOS) based on measured SNR
3. Match the location of the GNSS device to areas inside/outside of the various shadows

Any GNSS-capable Android smartphone or tablet can provide, via the location application programming interface (API), its estimated position with uncertainty, as well as the satellite coordinates and SNRs. Thus, shadow matching can be performed entirely in software without requiring any changes in GNSS receiver hardware or firmware.

The article by Wang *et alia* listed in the Additional Resources section near the end of this article provides an introduction to the basic idea of shadow matching. Wang and his coauthors reported promising results in mitigating GNSS cross-street errors using relatively straightforward techniques.

While shadow matching clearly provides information that is useful for urban localization, naive application of this idea using deterministic algorithms has significant limitations. First, the concepts of “high SNR when not blocked” and “low SNR when blocked” are inherently probabilistic, because SNRs exhibit large fluctuations due to multipath fading. Second, the range of SNRs seen by a device depends on the hardware implementation and physical realization of the GNSS receiver. Third, the potentially large errors in the raw location estimates output by the GNSS receiver are not captured by the typical



**FIGURE 1** A simplified illustration of Shadow Matching.

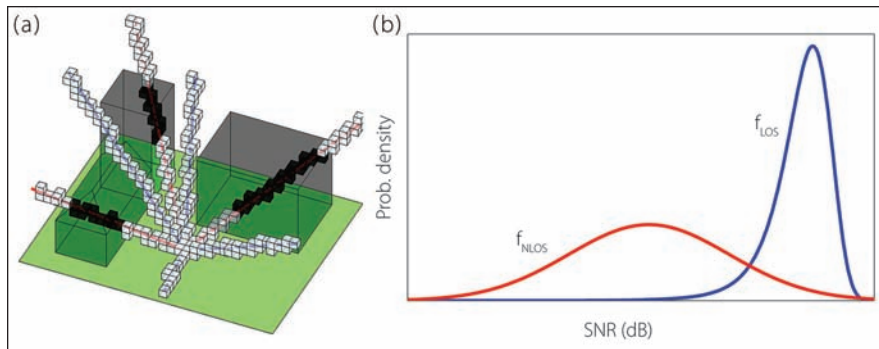


FIGURE 2 Modeling SNR measurements. (a) LOS/NLOS geometry (b) LOS/NLOS distributions

computations of uncertainty regions based on dilution of precision (DOP), because these implicitly assume that paths to all satellites are LoS. Hence, it is unclear how large an area to perform shadow matching over and how to fuse shadow matching information with GNSS location estimates. Finally, shadow matching is based on having access to accurate 3D maps, which are not readily available in many settings.

In the work described here, we address the preceding challenges through an adaptive Bayesian framework for inference, localization, and tracking with two complementary ingredients:

- **Real-time Localization:** Where accurate 3D maps are available, we develop nonlinear filters for localization based on the following ingredients: (a) probabilistic, rather than deterministic, modeling of shadow matching; (b) on-the-fly adaptation to device characteristics; (c) probabilistic fusion of information from shadow matching and GNSS location estimates; (d) mobility modeling using state space techniques. Because of significant modeling uncertainties and computational constraints, our solution needs to go well beyond classical particle filters.
- **3D SLAM:** When accurate 3D maps are not available, the same Bayesian framework, with additional computation, is used for generating and refining 3D maps using simultaneous localization and mapping (SLAM) based on crowdsourced GNSS data. Intuitively, we can do this by assigning likelihoods of blockage to

many crisscrossing receiver-satellite rays based on measured SNRs and then stitching these rays together, accounting for the uncertainty in the receivers' locations, into 3D maps. The framework can exploit "warm start" (e.g., based on public building data), if available.

Increased accuracy in urban localization clearly has many compelling applications, including car services, delivery services, navigation and guidance for vehicles (including for emerging automated and semi-automated operation), tourism, and hyperlocalized advertising.

We do not discuss any application in detail but instead focus on core technical aspects of our proposed solution in the remainder of this article. In particular, we describe the key concepts behind these algorithms in more detail, discuss cloud-based implementations with lightweight application layer modifications to mobile devices, and provide experimental results demonstrating the scalability and efficacy of these techniques.

Initial results based on these ideas, and the underlying technical details, were described in a series of publications last year. (See the articles by A. T. Irish *et alia* and J. T. Isaacs *et alia* in Additional Resources.) The experimental results reported here are based on significant enhancements to those techniques.

### Real-Time Localization and Tracking

In Bayesian estimation, we compute the conditional distribution of a quantity of interest given a set of measurements. This is termed a *posterior distribution*,

because we can only compute it after we see the measurements. In particular, we are interested in estimating the posterior distribution for the location of the GNSS receiver, based on measurements consisting of the GNSS position fixes and the satellite SNRs. In order to do this, we must model the conditional distribution of these measurements, conditioned on the actual location.

We consider a discrete time model, typically with samples spaced by one second. The true 3D location at time  $t$  is denoted by  $x_t$ , and the corresponding GNSS location fix is denoted by  $y_t$ . The satellite SNR measurements are denoted by  $z_{t,n}$ , where  $z_{t,n}$  is the SNR to satellite  $n$  at time  $t$ .

**Modeling GNSS Location Fixes.** The measurement model for the GNSS location fix is, in principle, straightforward:

$$y_t = x_t + e_t,$$

where the covariance of the error  $e_t$  can be estimated using standard DOP-style computations. However, these standard computations assume that the paths seen by the GNSS receiver are LOS, whereas many of the dominant paths in urban environments are actually strong reflections, with the LOS path blocked. Thus, in order to obtain satisfactory localization performance, we need to modify the model for  $\{e_t\}$ , resulting in a non-Gaussian statistical model.

### Modeling Satellite SNR Measurements.

We represent the 3D map  $m$  using an occupancy grid: space is divided into binary-valued "voxels" or "cells," with  $m_i = 1$  if the  $i$ th cell is occupied, and  $m_i = 0$  if the  $i$ th cell is unoccupied. If the ray from a hypothesized location to a satellite crosses only unoccupied cells, it is classified as LOS. If a ray passes through at least one occupied cell, it is classified as NLOS (or shadowed).

Figure 2 illustrates this approach. On the left, we illustrate ray tracing: blue/red lines represent LOS/NLOS signal paths to satellites, light/dark grid cells approximate empty/occupied space. On the right, we show typical SNR distributions that have been found to yield good performance: Rician for LOS, and log-

normal (with smaller mean and higher spread) for NLOS.

Note that we have made a drastic simplification in our model for NLOS rays, in that the distribution does not depend on the number of occupied voxels the ray crosses. Although it might be possible to improve performance using a more detailed model, our simplification leads to a significant reduction in computational complexity, while yielding excellent localization performance. Additional improvements can be obtained by introducing dependence of the LOS/NLOS SNR distributions on satellite elevation.

**Tracking framework.** We use a standard linear state space model for pedestrian motion, with the state consisting of position and velocity. However, because the measurement model is nonlinear and non-Gaussian, a standard Kalman filter cannot be used. Simple extensions such as the extended Kalman filter also do not work, as the posterior distribution of the state is often multimodal.

We therefore use a particle filter. Roughly speaking, it operates as follows: At each time  $t$ , the posterior distribution of the state is approximated by a discrete probability mass function putting weights at a set of hypothesized state values, or particles. These particles are propagated probabilistically to obtain a new set of particles and weights at time  $t + 1$ , based on the dynamics of the motion model, and the new set of measurements.

The particle filter has by now become a standard tool, but we have needed to make a number of modifications in order to account for modeling uncertainties.

Figure 3 shows example pedestrian results for Bush Street in downtown San Francisco. Figure 3(a) shows the mean trails and 68 percent confidence circles for the GPS reported fix (red) and the particle-filtered estimate (blue), with the ground truth path in yellow. Note that GPS makes cross-street errors which are corrected by our algorithm.

Figure 3(b) shows the SNR likelihood surface at the beginning of the trail: GPS

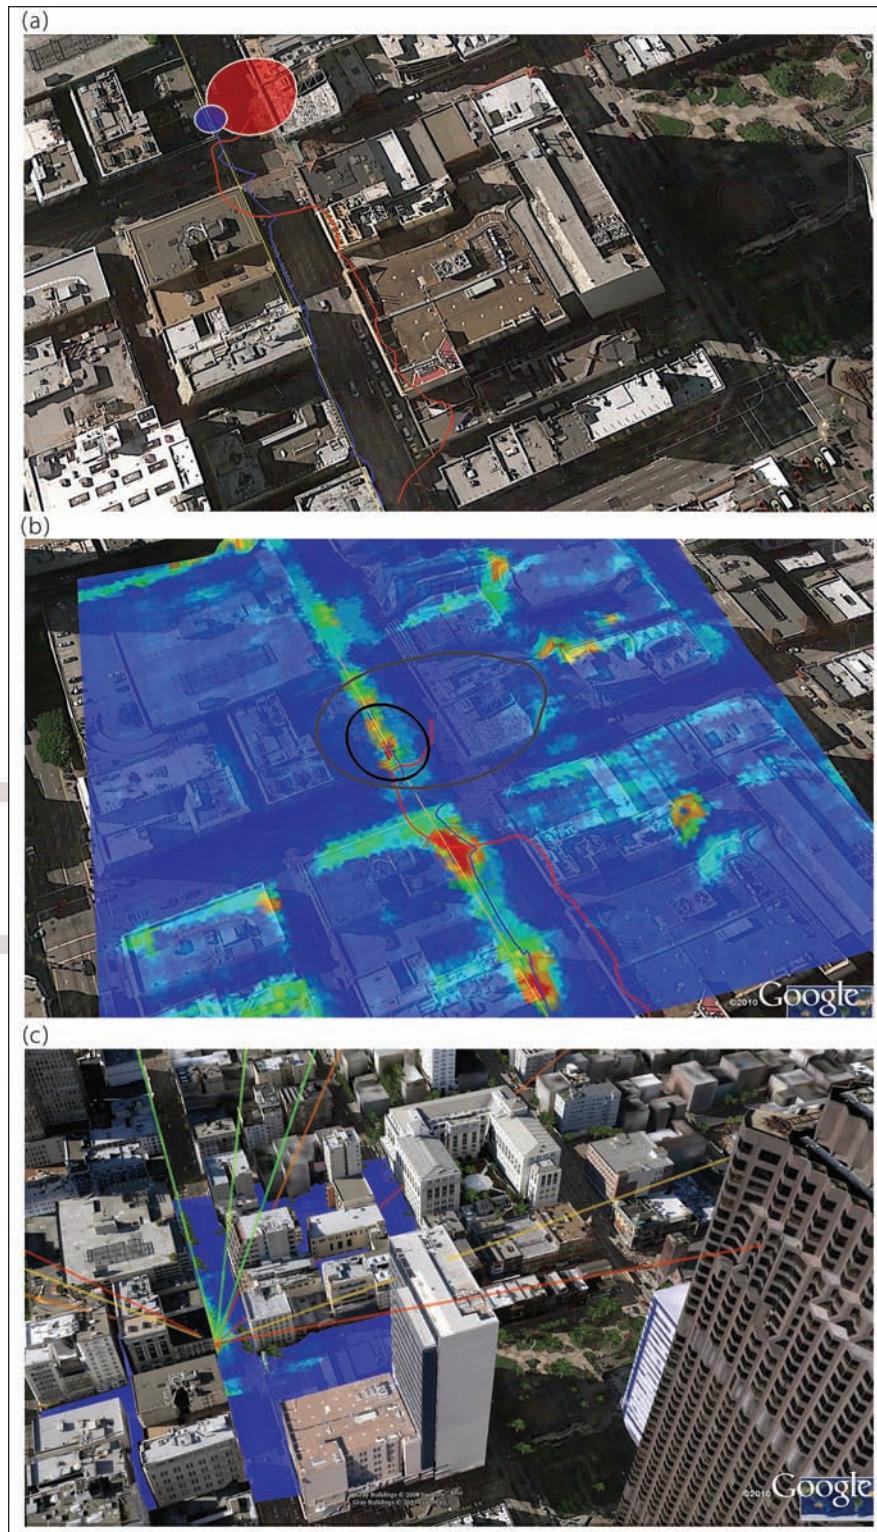


FIGURE 3 Example pedestrian results in San Francisco (Bush Street). (a) Mean trails and 68% confidence circles (b) SNR likelihood surface and error ellipses (c) Rays and composite (GPS x SNR) likelihood surface

starts out with a cross-street error, which our algorithm is able to correct because the SNR likelihood has a strong peak on

the correct side of the street. The large greyish-green ellipse is the  $3\sigma$  uncertainty estimated by GPS around its (incor-




FIGURE 4 Example vehicular results in San Francisco




FIGURE 5 Graphical model for Shadow SLAM. (a) Measurement scenario (b) Factor graph

rect) estimate, while the smaller black ellipse is the  $3\sigma$  uncertainty estimated around the location estimate provided by the particle filter.

Finally, Figure 3(c) shows the composite SNR/GPS likelihood surface, which exhibits a peak at the correct location. We also show the satellite rays colored likely LOS (green) to likely NLOS (red) according to our SNR model; from the image, we can see that these do correspond to unblocked and blocked rays, respectively.

Vehicular applications require more sophisticated motion models, and knowledge of the road network can also be exploited. However, the essential ingredients, in terms of the measurement model and the particle filter, are the same. Figure 4 shows example vehicular results with the 68 percent probability tunnels for both the GPS reported estimate (red) and our filtered estimate (blue), along with a road centerline clamping of the filtered estimate (green).

### SLAM for CrowdSourced 3D Maps

For 3D mapping, we are interested in using the noisy GNSS position measurements  $y$  and the SNR measurements  $z$ , to estimate the 3D map  $m$ . In particular, we wish to estimate the probability of each map cell being occupied or not, given all the measurements. That is, we wish to compute the *marginal posterior distributions*,  $p(m_i | y, z)$  for each cell  $i$ . However, because the GPS fixes are noisy, we do not know the paths  $x$  followed by the devices. Thus, in order to estimate the map, we must also estimate quantities of the form  $p(x_i^j | y, z)$ , where  $x_i^j$  is the position of a particular device  $j$  at time  $t$ . In the robotics community, this is referred to as the SLAM problem.

Our approach for Bayesian estimation of the map is to use a *factor graph* representation of the measurements and the quantities to be inferred, which allows us to employ *loopy belief propagation* to approximately compute the map marginal posteriors. While a detailed exposition of factor graphs is beyond the scope of this article, we provide intuition into their applicability in our context.

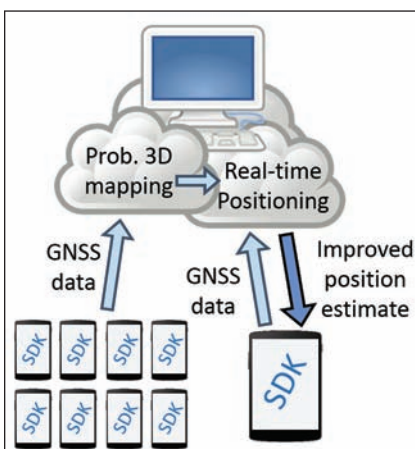
Figure 5 illustrates the construction of the factor graph based on measurements corresponding to a single GNSS receiver over two consecutive time periods. The variables that we wish to perform inference on, denoted by circles, are the binary-valued map cell occupancies  $\{m_i\}$  (which is our primary interest in shadow-based SLAM) and the true locations  $\{x_i\}$  as a function of time  $t$  (which are, for the mapping task, nuisance variables to be averaged out).

The factors represent the statistical information about the variables provided by the measurements. In Figure 5, the factor  $g_t$  represents  $p(y_t | x_t)$ , the likelihood function corresponding to the GPS fix at time  $t$ , and the factor  $f_{t,n}$  represents  $p(z_{t,n} | m, x_t)$ , the likelihood function corresponding to the SNR to the  $n$ th satellite at time  $t$ .

Belief propagation can now be implemented by the standard sum-product



**FIGURE 6** Shadow SLAM in downtown Santa Barbara. (a) Mapped area and GNSS traces (b) Horizontal slices of estimated 3D map



**FIGURE 7** Cloud-based Architecture

algorithm described in the article by F. Kschischang *et alia*, passing messages back and forth between nodes. A message is simply a function whose domain is the set of values taken by the variable. Thus, for a binary-valued map cell  $m_p$ , a message is a binary vector.

For a position variable  $x_p$ , we simply quantize its possible set of values to a discrete set (e.g., with  $K$  possible values), and a message is then a  $K$ -dimensional vector corresponding to the values of the function evaluated at these  $K$  points. A message along an edge from a variable node to a factor node is simply a product of messages coming along all of the other edges.

A message from a factor node  $f$  to a variable node  $v$  is more complicated: it involves products of messages coming

into that factor node, but also involves summing over (marginalizing out) all variables other than  $v$ . We refer the reader to the article by A. T. Irish *et alia* (2014a) for details, but it is worth mentioning that our simplified SNR model (in which the NLOS model for SNR is independent of the number of occupied map cells intersected by the ray) is key to tractable computation of messages from SNR factor nodes to position and map variable nodes.

The preceding procedure has been successfully used to construct 3D maps of the campus of the University of California, Santa Barbara (UCSB), as well as of downtown Santa Barbara. Here we will present sample results for the latter based on about 25 hours of input data from four Android devices. **Figure 6(a)** shows an aerial view, with GNSS traces in red and mapped region outlined in yellow. **Figure 6(b)** shows the cell occupancy estimates for 2D slices at multiple heights.

We employed a “warm start” with prior information based on 2D OpenStreetMap data employed to initialize the map for the first couple of slices. Note, however, that, given enough data, 3D maps can be generated without any prior information. (This was the case for the 3D maps generated for UCSB campus; see A. T. Irish *et alia* (2014a) for details).

## Architecture and Implementation

Shadow-based localization and 3D SLAM can be easily deployed on existing mobile devices via a software development kit (SDK), which performs a very simple function: send GNSS data to “the cloud” at regular intervals and receive improved position information (along with an estimated uncertainty) back. The GNSS data includes the estimated latitude and longitude coordinates, along with the azimuth, elevation, and SNR of

each satellite in view. As shown in **Figure 7**, all of the heavy lifting for both real-time localization and 3D mapping is done by means of Internet cloud computing.

The preceding architecture naturally scales to include additional sources of information when available — for example, inertial navigation, WiFi, and cellular. Such information can be sent from mobile device to cloud, and can be incorporated into the Bayesian computations done in the cloud. Such extensions are also under development.

## Conclusions

The accurate urban localization provided by enhancing GPS with shadow matching has the potential to create significant commercial benefits, including for car/taxi services, delivery services, location-based advertising, e911, tourism, and automated and semi-automated navigation.

Several academic research groups worldwide have demonstrated the promise of shadow matching over the past few years. However, a key innovation in the approach described here is that, in contrast to the ad hoc algorithms used previously, it is based on a Bayesian inference framework that systematically accounts for significant modeling uncertainties and nonlinearities, and seamlessly extends to incorporate additional sources of information as they become available.

The approach proposed here has been demonstrated to provide significant improvements in localization over raw GNSS estimates in challenging high-rise environments such as downtown San Francisco, with latencies of the order of 100 milliseconds. With current (and ever decreasing) cloud computing costs and the inherent parallelizability of the core computation bottlenecks (such as ray tracing and particle sampling techniques), providing such a service is estimated to cost less than one U.S. penny per hour per device.

On the mapping side, shadow-based SLAM provides a method for creating 3D maps from scratch and an effective means of refining them from a warm start. The former is probably most important for military applications.

For most urban environments worldwide, 2D maps together with public municipal data on building heights can be used to build a coarse 3D map, which can subsequently be refined by shadowbased SLAM. Even in areas where accurate 3D building models are available, shadow-based SLAM provides the ability to continuously monitor and capture changes in the environment (e.g., new construction, demolition). It also captures features such as trees and lampposts. As vehicles migrate towards automated operation, a fine-grained mapping of features in the environment become increasingly important for timely decision making in navigation and path planning.

While the cloud-centric architecture described previously in the “Architecture and Implementation” section enables immediate deployment of shadow-matching technology on mobile devices via a lightweight SDK, it is also of great interest to explore means of efficiently migrating some of the algorithmic functionalities to a mobile device. Accomplishing such a migration would reduce reliance on network connectivity by exploiting the computational resources available on today’s smart phones. Self-contained operation with limited or no network connectivity is also an important feature for incorpo-

rating these advances in urban localization into vehicular navigation systems.

**Additional Resources**

- [1] Ben-Moshe, B., and E. Elkin, H. Levi, and A. Weissman, “Improving Accuracy of GNSS Devices in Urban Canyons,” in *Proceedings of the 23rd Canadian Conference on Computational Geometry*, 2011
- [2] Groves, P., “Shadow Matching: A New GNSS Positioning Technique for Urban Canyons,” *Journal of Navigation*, 64(3):417–430, 2011
- [3] Irish, A. T. (2014a), J. T. Isaacs, F. Quitin, J. P. Hespanha, and U. Madhow, “Belief Propagation Based Localization and Mapping Using Sparsely Sampled GNSS SNR Measurements,” in *Proceedings of IEEE International Conference on Robotics and Automation*, 2014
- [4] Irish, A. T. (2014b), J. T. Isaacs, F. Quitin, J. P. Hespanha, and U. Madhow, “Probabilistic 3D mapping based on GNSS SNR measurements,” in *Proceedings of IEEE International Conference on Acoustics and Signal Processing*, 2014
- [5] Isaacs, J. T., and A. T. Irish, F. Quitin, U. Madhow, and J. P. Hespanha, “Bayesian Localization and Mapping Using GNSS SNR Measurements,” in *Proceedings of IEEE/ION Position Location and Navigation Symposium*, 2014
- [6] Kim, K., and J. Summet, T. Starner, D. Ashbrook, M. Kapade, and I. Essa, “Localization and 3D Reconstruction of Urban Scenes using GPS,” in *Proceedings of IEEE International Symposium on Wearable Computers*, pp. 11–14, 2008
- [7] Kschischang, F., and B. Frey, and H.-A. Loeliger, “Factor Graphs and the Sum-Product Algorithm,” *IEEE Transactions on Information Theory*, 47(2):498–519, 2001
- [8] Park, M., and H. Kim, S. Lee, and K. Bae, “Performance Evaluation of Android Location Service at the Urban Canyon,” in *Proceedings of IEEE International Conference on Advanced Communication Technology*, pp. 662–665, 2014
- [9] Thrun, S., “Simultaneous Localization and Mapping,” in *Robotics and Cognitive Approaches to Spatial Mapping*, pp. 13–41, Springer, 2008
- [10] Wang, L., and P. D. Groves, and M. K. Ziebart, “Urban Positioning on a Smartphone: Real-Time Shadow Matching Using GNSS and 3D City Models,” in *Proceedings of ION GNSS+*, 2013

**Authors**



**Andrew Irish** is a Ph.D. candidate in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara (UCSB). His research interests include mapping and localization, machine learning, signal processing, and wireless communication. He is a

founder of ShadowMaps, a UCSB spinoff commercializing localization and 3D mapping based on GNSS shadow matching, where he is taking the lead in implementing and scaling GNSS SNR-based localization and mapping techniques. In addition to his time at ShadowMaps, he has acquired more than two and half years of industry research and development experience, which includes time spent at Qualcomm, InterDigital, and the National Research Council of Canada. He graduated with a bachelors of engineering in electrical engineering from McGill University, in Montreal, Canada.



**Danny Iland** is a Ph.D. candidate in the Department of Computer Science at the University of California, Santa Barbara (UCSB). His research focuses on cellular and wireless networking,

off-grid rural networks, and network measurements and analysis. He is a founder of ShadowMaps, where he is developing mobile applications and scalable cloud services for localization and mapping based on GNSS shadow matching. He has previously worked on localization improvement of aerial imagery at Goodrich ISR Systems, as a software developer at IBM Software Group, and as a security engineer at iSec Partners. He graduated with a bachelors of engineering in computer science from the Rochester Institute of Technology.



**Upamanyu Madhow** is a professor of electrical and computer engineering at UCSB. His research interests span wireless communication, signal processing, bio-inspired design, and

machine learning. He actively engages in technology transfer of his academic research. He is a founder of ShadowMaps, where he is working on algorithms, intellectual property, and commercialization strategy for localization and mapping based on GNSS shadow matching. He has co-founded two prior companies, Bymobile (now part of Citrix) and Wirama (now part of Checkpoint Systems) in mobile data and RFID localization, respectively. He is a Fellow of the IEEE, recipient of the National Science Foundation Career award, and co-recipient of the IEEE Marconi prize paper award in wireless communications. He is the author of two textbooks in communication systems, published by Cambridge University Press. He holds 14 U.S. patents. 