Navigation in GNSS-Denied Environments: Signals of Opportunity and Beacons

Gregory L. Duckworth
QinetiQ North America
350 Second Avenue
Waltham, MA 02451
United States

gduckworth@foster-miller.com

Edward J. Baranoski
DARPA/STO
3701 N Fairfax Dr.
Arlington, VA 22203
United States

ed.baranoski@darpa.mil

ABSTRACT

Precision geolocation of individuals, vehicles, aircraft, ships, and weapons has revolutionized warfighting, commercial logistics, industrial farming, construction and surveying, and even personal auto travel through in-vehicle navigation systems. These applications have in common that they are predominantly carried out in open line-of-sight environments that can see satellite navigation assets. As we become more dependent on Global Navigation Satellite Systems (GNSS) in these applications, we have become more aware of their two primary shortcomings: i) lack of performance in urban, indoor, and underground environments, and ii) the ease with which they can be denied by intentional or unintentional interference. This paper provides a summary of the user needs for robust navigation and geolocation in GNSS-denied environments, some approaches and their advantages and disadvantages, and some recent results of indoor/urban geolocation research and development efforts. What we hope to show in this paper is that practical geolocation and navigation techniques for GNSS-denied environments and missions should consider serial and/or parallel hybrids of the fundamental approaches of i) direct measurements, ii) inertial or dead-reckoning, and iii) environmental correlation, to achieve real-time, robust, and lowest-cost results. We recommend a particular combination for the most general indoor/urban geolocation and navigation problem, and highlight what we believe are the most important remaining technical problems to be solved.

1.0 INTRODUCTION

Precision geolocation of individuals, vehicles, aircraft, ships, and weapons has revolutionized warfighting, commercial logistics, industrial farming, construction and surveying, and even personal auto travel through in-vehicle navigation systems. These applications have in common that they are predominantly carried out in open line-of-sight environments that can see satellite navigation assets. As we become more dependent on Global Navigation Satellite Systems (GNSS) in these applications, we have become more aware of their shortcomings. GNSS may be denied by attacks on the space or ground segments of the system, or by jamming the signals available at the user's receiver. This denial may also be the result of signal blockage or extreme signal multipathing by foliage, urban canyons, or use of the user receiver inside buildings. With the increasing urbanization of the world and the changing nature of global combat, the urban battlefield is becoming increasingly important, and we must maintain and improve our precision localization capabilities in this area.

Duckworth, G.L.; Baranoski, E.J. (2007) Navigation in GNSS-Denied Environments: Signals of Opportunity and Beacons. In *Military Capabilities Enabled by Advances in Navigation Sensors* (pp. 3-1 – 3-14). Meeting Proceedings RTO-MP-SET-104, Paper 3. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int.

For purposes of this paper, we distinguish navigation from geolocation in that navigation applications almost always involve movement of the user, and the need for the user to know where they are in real-time. Geolocation applications can include a stationary sensor that is emplaced, and needs to wake up and know and report where it is. The need for absolute (with respect to well-defined Earth/environmental coordinates) or relative positioning is a very mission-specific requirement, but with the increased digitization and geo-referencing of battlefield and domestic environments, it is almost always worthwhile to be able to tie back to standardized coordinates.

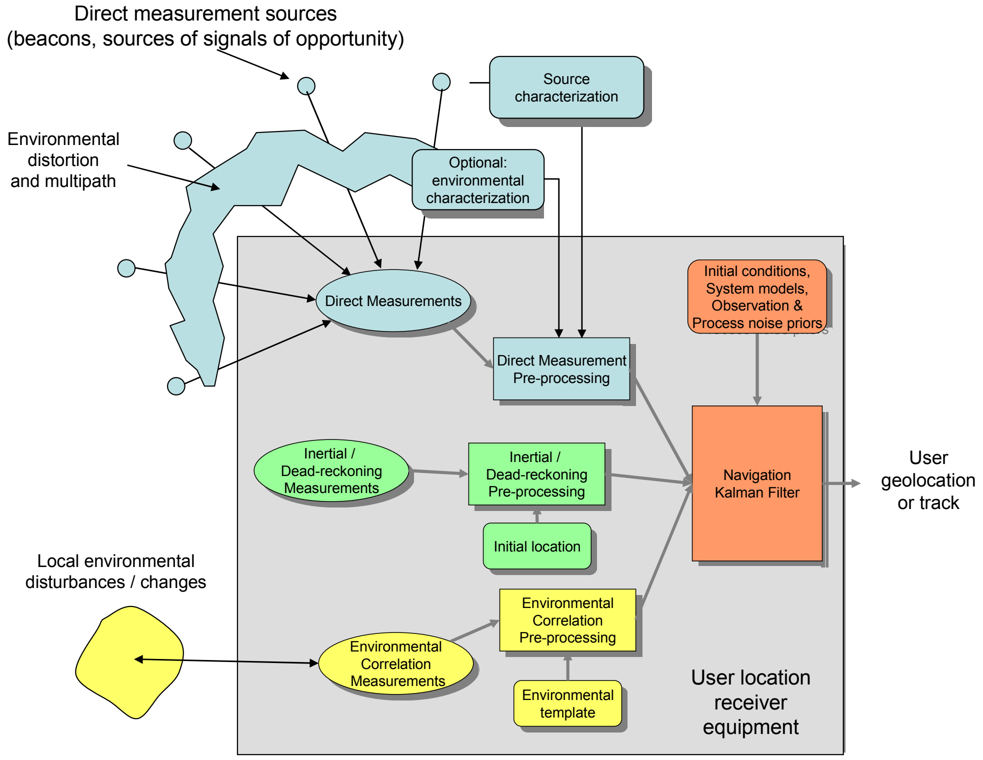

Figure 1. A generic Navigation/Geolocation system block diagram showing the three primary approaches to geolocation and navigation, and a top-level system component block diagram and nomenclature.

There are three fundamental approaches to geolocation and navigation, as shown in Figure 1:

1. Direct Measurements of amplitude, polarization, travel-time, phase, or angle-of-arrival of signals broadcast from known locations can be “triangulated” in the most general sense to get the location of the receiver. A major advantage of this approach is that if enough of these measurements are available and of good quality from known source locations, then the absolute geolocation problem is easily solved, instantly, and no prior knowledge of one’s location is required. Most importantly, there is no accumulation of location error over time from a known starting location, and no prior knowledge of the environment or pre-existing maps, environmental templates, or building layouts are required. GNSS systems use this approach, very successfully.

2. Inertial and Dead-reckoning using accelerometers and gyros, or velocity measurements and compasses, respectively, provide another method of determining ones location in a previously unmapped environment. The advantage of these methods is that the user can be virtually independent of the environment since there are no direct-measurement signals which must travel from known locations to the user, and potentially be corrupted or blocked by the environment between. The disadvantage of these methods over direct-measurements is that they require an independently-determined initial location, and their error grows with elapsed time from the beginning of the mission at that known location. In practice, it has been found that only extremely high value missions, such as submarine navigation, and other strategic platforms and weapons, can afford the cost, size, weight, and power required for the quality inertial measurement units (IMUs) need to achieve the desired mission accuracy and longevity performance. Dead-reckoning techniques sacrifice some environmental independence to reduce inertial navigation's double integration of (mostly environmentally independent) instrument errors to a single integration of heading and velocity measurements that may be influenced by the local environment—such as magnetic anomalies, motion of the volume or surface from which the velocity is being measured, or reflectance anomalies of these scatterers.

3. Environmental Correlation Navigation exploits specific pre-mapped characteristics of the mission environment. Users entering the environment determine where they are by recognizing these pre-mapped characteristics or features, and determine where they are with respect to them. With respect to other methods, this has the disadvantage of requiring pre-mapping of the environment for the cue that is used. Pre-mapping to get the navigation template is usually laborious, and often impossible in denied areas. It can sometimes be done using standoff remote sensing techniques to get the environmental template, but many environmental signatures are only map-able, or provide the fidelity required for the desired performance, at a mapping-stage standoff similar to the final navigation user. Thus, environmental correlation navigation methods can be valuable for navigation in specialized applications, but are not as general as the two other approaches. These methods often require user movement within the environment to obtain enough data to do template matching, but not always.

Almost any successful geolocation/navigation technique can be found to use one or more of these three approaches. Even so-called "surveying" approaches that establish and then propagate baselines via sequences of precision angle and/or distance measurements are hybrids of these approaches. For example, the height and location of Mount Everest by Lambton, Everest, and Waugh was originally determined by establishing a precision baseline on the shore at sea level, referenced in global coordinates by its proximity to the sea (environmental correlation) and celestial navigation (a direct measurement method). Then a laborious sequence of direct measurements of angle and distance walking away from the coast was used to determine the peak's height and location with respect to this starting baseline. This is essentially dead-reckoning using first integrals of angle and distance steps. In this example, outstanding results were achieved through a very long and carefully executed serial-hybrid of all three approaches.

Due to the demands of different environments and missions, no universally applicable system is likely to emerge. However, the authors feel that any practical approach is likely to use optimum combinations of these techniques to compensate for the weaknesses in low-cost / operationally-deployable implementations of each individual approach, and the very high cost of a robust solution using only one approach. What we hope to show in this paper is that practical geolocation and navigation techniques for GNSS-denied environments and missions should consider specific user requirement-dependent serial and/or parallel hybrids of these fundamental approaches to achieve real-time, accurate, robust, and lowest-cost results.

2. USERS: ENVIRONMENTS, MISSIONS, AND REQUIREMENTS

The performance needs of urban warfighters (human or robotic), and emergency personnel are remarkably similar. They need to know where they and their comrades are on a scale that is tactically significant—to the room and floor level in buildings, and on what side of the street when outdoors. Their requirements are similar, though much more difficult to achieve, when underground in tunnels, sewers, and underground facilities. From these needs one can see that achieving reliable real-time and indefinite mission duration 3-D location accuracy at this level should be the goal for tactically useful indoor, underground, and urban geolocation and navigation systems.

All three geolocation/navigation approaches outlined in the introduction, and serial and parallel permutations and combinations of them, are applicable in some contexts. The most general solutions that fit the needs of warfighters, emergency personnel, and sensor and weapons platforms meet the following requirements:

- Support indefinite mission duration.
- Support real-time 3-D location performance appropriate for tactical mission requirements.
- Support localization inside residential and most commercial buildings.
- System must be logically extensible to urban canyons and all current stationary and mobile GNSS users under GNSS-denied conditions.
- Support localization from power-off condition and require no separate initialization of user equipment with starting location, etc.
- Continuous location tracking during the mission to maintain performance must not be required. Performance must be able to be re-established after losing track during the conduct of the mission.
- System must operate in arbitrary, previously unknown and non-pre-mapped or characterized environments.
- The system should be low cost—available to individual dismounts and emergency personnel when needed for their mission.
- Support individual isolated user terminals. Multiple users in communications range of one another must not be required for the system to achieve its performance.
- System must not require user motion in the environment to work.

These requirements clearly argue strongly for a “direct-measurement” approach as the core of the solution, since it is the only technique that can meet some of these requirements.

It is also clear that systems which fall short of this mark are also valuable until we can get to this technical ideal. There are existing solutions that use direct RF range measurements from Signals of Opportunity (SoOP) (e.g. Rosum [1]) but do not achieve the fine location accuracy desired by some. Their accuracy is typically 30-40 meters in urban environments, which is useful for some users and applications. There are existing inertial systems (e.g. Applanix [2]) that can provide single-meter levels of accuracy over mission durations of an hour, but are quite expensive (more than $100k/user). There are also very inexpensive inertial systems (e.g. TRX [3]) that when supplied with accurate indoor plans, and perhaps augmented with pre-installed RF tags or location-tagged images of features, can provide reasonably reliable localization over long indoor missions for applications in which the building is pre-mapped or prepared.

Each of these commercial or near-commercial systems has users that can usefully employ them at their performance and cost points. However, as motivated above, we believe that affordable, robust, meter-level non-GNSS, indefinite mission duration systems that will work in non-pre-mapped interior and urban environments or buildings can be developed. The rest of the paper will indicate how we think this can be done.

3. A SOLUTION METHODOLOGY EXAMPLE

The DARPA Robust Surface Navigation (RSN) [4] and Sub Surface Navigation (SSN) [5] were initiated and run by the authors to first extend the information advantage of the GPS-equipped US warfighters and platforms to GPS-denied interior and underground environments, respectively, and then to provide a GPS-independent capability to navigate and geolocate over broader urban environments and ultimately globally. Each program required that any approach include a direct-measurement core technology—so that the user could wake up and navigate anywhere independent of the environment or any prior knowledge of it, and do so for arbitrarily long periods of time without inertial or dead-reckoning drift. This direct measurement technology could use purposely-deployed beacon emitters and/or exploit signals-of-opportunity (SoOP) in the environment, such as commercial broadcast TV, radio, and telecommunications signals from ground or space-based assets. The advantages of exploiting existing natural or cultural SoOP are obvious—such as less expense, less deployment logistics, and no overt emitters. Non-direct approaches were also encouraged at a non-developmental level if they were valuable for augmenting the direct measurements in order to increase the accuracy and robustness, or decrease the cost of the system. A generic block diagram of almost any geolocation/navigation system capable of exploiting each of these approaches was shown previously in Figure 1, and defines the nomenclature used in the remainder of this paper.

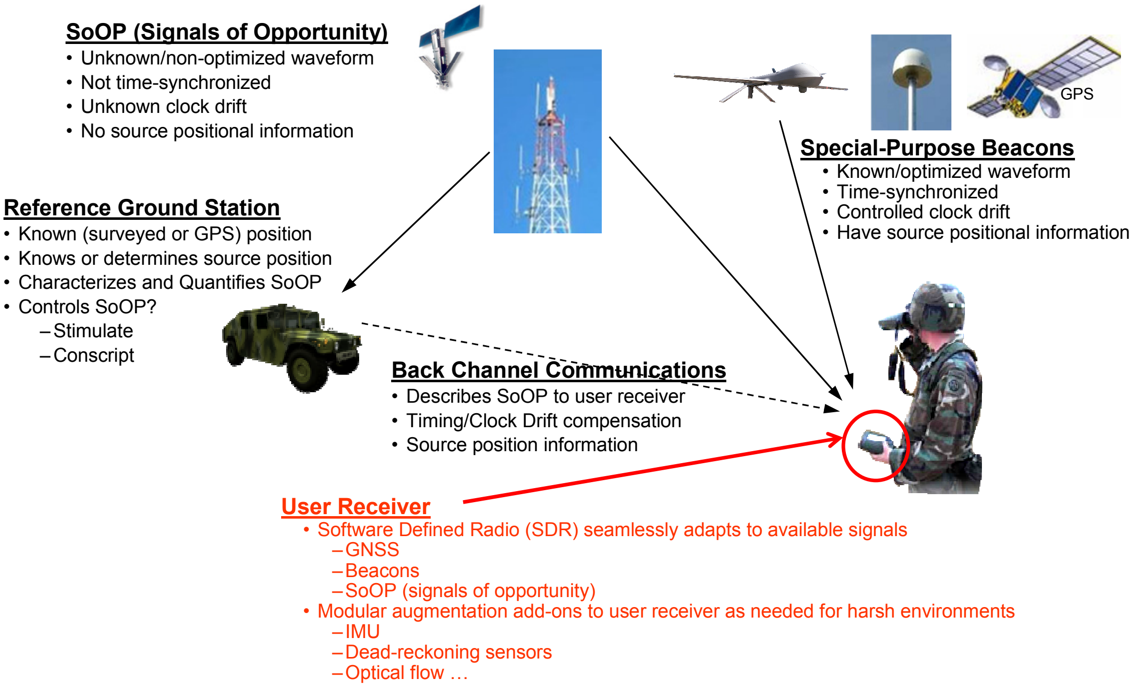

Figure 2. Navigation/Geolocation system architecture that can potentially meet the demanding requirements outlined in the text.

A system architecture as shown in Figure 2 could satisfy these requirements. Since it is direct-measurement based, it does not have a mission time limitation due to inertial navigation system (INS) or dead-reckoning drift. Figure 2 also shows that there are two fundamental ways to get the needed direct measurements on single building through global scales—through purpose-deployed beacons (such as GNSS satellites and terrestrial analogs, such as LORAN stations), and through exploiting “signals of opportunity” (SoOP) in the environment, such as commercial TV and radio transmissions and telecommunications, both terrestrial and space-based.

The two goals of the RSN program:

- provide geolocation capability where it does not now exist in interior/urban environments, and
- provide a robust backup to GPS globally

are initially being served by development of beacon-based systems and algorithms that will combat the penetration and multipath problems of building interiors. When it is clear that a working solution using beacons is possible, it will be determined if and when SoOP can be used as well, and what performance and area coverage advantages they can provide. It will also be determined what the requirements are on the SoOP calibration reference ground station network. These stations serve a purpose similar to GPS differential correction ground stations, but are more important since they must completely characterize the SoOP sources, not just provide minor corrections to enhance GPS performance degraded by ionospheric distortion. An initial assessment of the advantages and disadvantages of beacons and SoOP, from a system perspective, is shown in Figure 3.

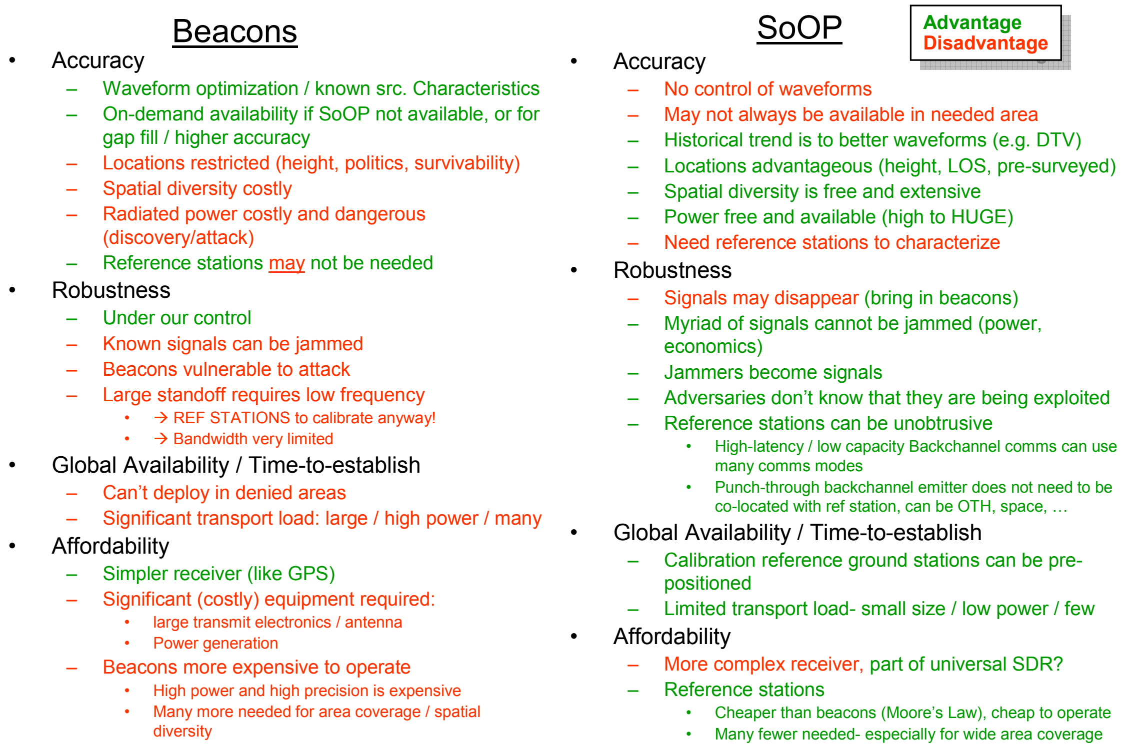

Advantage
Disadvantage

Beacons

- Accuracy
  - Waveform optimization / known src. Characteristics
  - On-demand availability if SoOP not available, or for gap fill / higher accuracy
  - Locations restricted (height, politics, survivability)
  - Spatial diversity costly
  - Radiated power costly and dangerous (discovery/attack)
  - Reference stations may not be needed
- Robustness
  - Under our control
  - Known signals can be jammed
  - Beacons vulnerable to attack
  - Large standoff requires low frequency
    - → REF STATIONS to calibrate anyway!
    - → Bandwidth very limited
- Global Availability / Time-to-establish
  - Can’t deploy in denied areas
  - Significant transport load: large / high power / many
- Affordability
  - Simpler receiver (like GPS)
  - Significant (costly) equipment required:
    - large transmit electronics / antenna
    - Power generation
  - Beacons more expensive to operate
    - High power and high precision is expensive
    - Many more needed for area coverage / spatial diversity

SoOP

- Accuracy
  - No control of waveforms
  - May not always be available in needed area
  - Historical trend is to better waveforms (e.g. DTV)
  - Locations advantageous (height, LOS, pre-surveyed)
  - Spatial diversity is free and extensive
  - Power free and available (high to HUGE)
  - Need reference stations to characterize
- Robustness
  - Signals may disappear (bring in beacons)
  - Myriad of signals cannot be jammed (power, economics)
  - Jammers become signals
  - Adversaries don’t know that they are being exploited
  - Reference stations can be unobtrusive
    - High-latency / low capacity Backchannel comms can use many comms modes
    - Punch-through backchannel emitter does not need to be co-located with ref station, can be OTH, space, …
- Global Availability / Time-to-establish
  - Calibration reference ground stations can be pre-positioned
  - Limited transport load- small size / low power / few
- Affordability
  - More complex receiver, part of universal SDR?
  - Reference stations
    - Cheaper than beacons (Moore’s Law), cheap to operate
    - Many fewer needed- especially for wide area coverage

Figure 3. Relative advantages and disadvantages of purpose-deployed beacon-based, and Signal of Opportunity-based geolocation and navigation systems.

The authors believe that it is likely that a hybrid approach using both direct-measurement and inertial/dead-reckoning methods will be needed to meet all the system requirements with robustness and minimum cost. Figure 4 shows that the shortcomings in each of these methods can be compensated by the other to a great degree, and medium-performance / medium-cost implementation of each will lead to a robust high-performance result. By robust, we mean that it will meet performance goals most of the time in most environments. Figure 4 motivates the use of direct measurements for a drift-free "wake up and navigate" solution based on robust inversion ("triangulation") of beacon- or SoOP- based measurements.

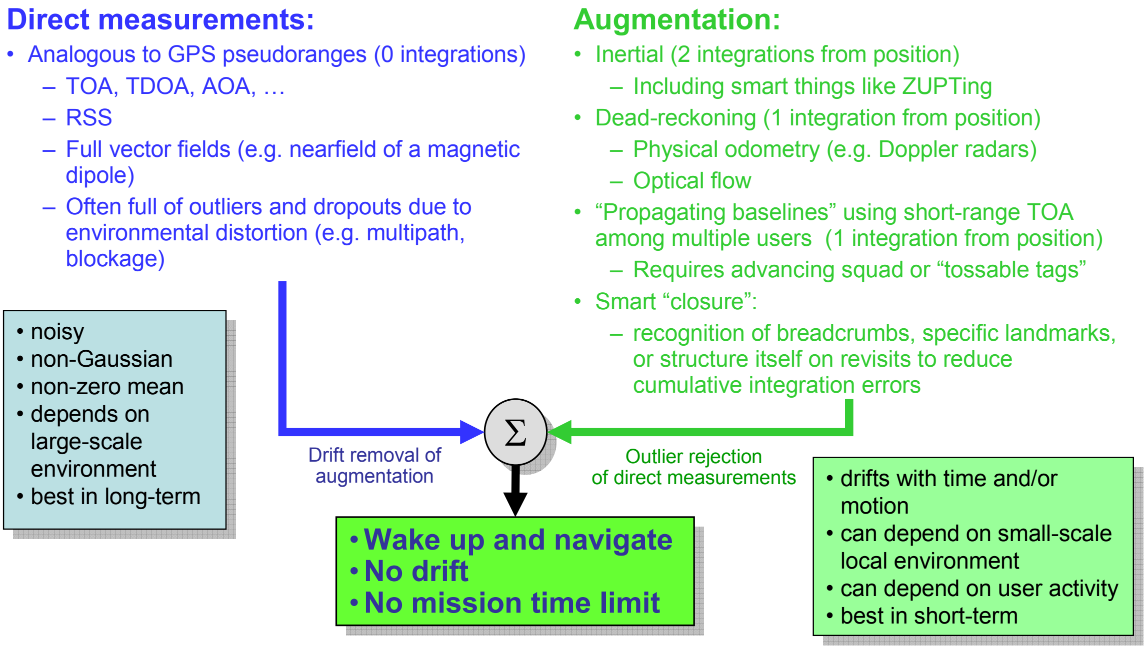

Figure 4. The complementary aspects of direct measurement and inertial/dead-reckoning solutions will lead to a robust solution meeting the desired requirements.

The big problem in urban/indoor environments is that these measurements are corrupted compared to the unobstructed line-of-site paths provided to outdoor users of GNSS systems. In these tougher environments, the RF paths are non-line of sight, and fade due to blockage, and are often highly multipathed, so that the errors in time delay "pseudoranges" are non-zero mean, non-Gaussian, and often contain outliers due to missed direct path arrivals and detection of multipath only. The beacon and waveform designers will achieve the best compromise among building penetration, waveform bandwidth, and beacon/antenna complexity. They will make high-resolution direct-path (the most easily modelled and inverted for location) non-multipath-corrupted arrivals as available as possible. The big problem is identification and removal of any corrupted direct-path measurement data. This can be done in two ways:

- collect a large highly-redundant set of measurements and sort through it to find the limited observations that are (self)-consistent with a single estimated location, and
- use inertial/dead-reckoning (DR) augmentation to remove measurements that are inconsistent with its (low drift) short-time extrapolation from the previous trusted location estimate to the current estimate.

A particular advantage of augmentation with reasonable performance inertial/DR measurements is that the highly redundant set of direct measurements does not need to be taken at the same time—the inertial/DR data and navigation Kalman filter (see Figure 1) can filter outliers and integrate direct measurements over the short-term stability interval of the inertial/DR sensors. In addition, outlier rejection can use a non-causal smoothing filter formulation that can take advantage of "future data" and use hind-sight to eliminate grossly inconsistent measurements based on more recent direct and inertial/DR data. The current location estimate can be based on up-to-date inertial/DR navigation integrating forward over a short interval from the last direct measurement data "fix" that is old enough to have been cleaned up by this smoothing filter.

No matter how the direct measurements are collected, it is clear that propagation path diversity is a key to multipath mitigation, second in importance only to high-bandwidth waveforms for high-resolution time-of-arrival/angle-of-arrival measurements. The need is to get enough good measurements to determine location, and to identify and eliminate bad "outlier" measurements. The physics we can exploit is this: urban/interior true direct paths are reasonably stable in arrival time and angle—they may fluctuate in amplitude and be completely blocked, but their arrival time is physically predictable from short-term beacon/user motion estimation using inexpensive low-SWAP (size, weight, and power) inertial/DR augmentation. Over time, the unstable/scintillating multipath signals are readily identifiable by their inconsistency with the short-term predictions, as are direct arrivals that are corrupted by close-in multipath distortion of their arrival times or angles, and both may be rejected from the localization inversion in the direct-measurement pre-processing or navigation Kalman filter. Somewhat counter-intuitively, this is what can make a system with moving users and/or moving beacons actually simpler than the completely stationary case. In essence, for moving beacons or users, one may get the same spatial propagation path diversity with many fewer beacons with the simple addition of a low-to-medium performance inertial/DR augmentation subsystem, and collecting and integrating the diverse paths over time. In a short-range/multipath-dominated environment, removal of large errors is much more important to the solution quality than SNR or averaging.

With the above design considerations:

- waveform center frequency and bandwidth tuned for the desired compromise between structural penetration (e.g. SNR) and range/multipath resolution,
- high spatial diversity through multiple beacons/SoOP sources and source/user motion to get enough uncorrupted direct-path arrivals to provide a well-determined (low-DOP) estimate for the user location, and to allow direct paths to be separated from multipath through data/estimate consistency measures, and
- inertial or dead-reckoning augmentation to allow outlier rejection and temporal smoothing of noisy direct measurement estimates,

the authors believe that the very general indoor/urban 3-D location problem meeting the requirements above can be solved.

Before we go on to discuss some current results of activities related to this approach, we comment briefly on what is not in the design prescription above: the exploitation of multipath arrivals, not just their mitigation. The holy grail of this approach would be a system that could also estimate the building structure and provide it to the user location equipment. The user location algorithms in these could then use the full richness of the multipath arrival structure in a "matched field" sense as used in sonar processing, to get better performance, not worse, in the face of multipath. There are programs, such as DARPA's VisiBuilding effort, that are seeking to use proximal radars to probe and estimate building interior structures. When these come to fruition, they may be coupled into interior localization solutions. However, in the near term the work being done on blind exploitation of multipath by the Charles Stark Draper laboratory [6] is the most practical effort we have seen to exploit multipath in real time. This approach recognizes that reasonably accurate empirical models can be developed to estimate the relationship between multipath arrival times, the last multipath reflection point, and arrival user location. These models are estimated using auxiliary "nuisance" parameters in the navigation Kalman filter that estimates and tracks location, and that are fed associated direct arrivals, multipaths, and (optional, but useful) inertial or dead-reckoning sensor measurements. Once these parameters are estimated to a sufficiently low variance, they are often valid over reasonable spatial extents of motion by the user (dependent upon the structural characteristics of the environment, and the range to the beacons/SoOP sources—farther is better for the approximations), and direct paths may disappear, but the user location estimation continues as long as the associated multipath arrivals continue to come in. The limits of applicability of this technique are still under study. It is not yet certain that for most environments of interest whether the additional information in the multipaths is sufficient to estimate the time-varying nuisance parameters to the degree needed before their multipath arrivals disappear, or the empirical approximation breaks down, due to user motion-induced geometry changes.

In summary, the Draper multipath exploitation approach, and the RSN program efforts by Argon ST and Boeing-led teams, all sponsored by DARPA, are currently seeking to apply these principles using novel waveform and beacon designs, and multi-layered outlier-rejection algorithms and navigation filters. We have high hopes that these efforts will point the way to balanced system designs getting the best of direct measurement and inertial/DR approaches to define a robust and affordable capability for 3-D in-building and urban localization. We also hope that this may be seamlessly extensible to SoOP-based direct measurements allowing the technique to be applicable at regional and global scales.

3. EXAMPLE RESULTS

At this time, the RSN program is in a Phase 1 competition, so the detailed approaches and interim results are still team proprietary and cannot be reported here. However, there have been some very promising results from other government funded and privately funded efforts that use some of the principles advocated above, and show the promise of integrated high spatial diversity solutions. The work by Draper Laboratory discussed earlier in the paper is a promising algorithmic approach that may be able to make most effective use of both direct and multipath arrivals, particularly when combined with a moderate performance inertial/DR system to average over direct measurement fluctuations and outages, and particularly to aid in the removal of large direct measurement outliers.

3.1 Combating Multipath

Effectively obtaining the quantity and quality of direct, and discrete identifiable multipath measurements needed to achieve tactical-level interior localization is still an unsolved problem. Unfortunately, as discussed earlier in this paper, the statistics of these arrivals are highly non-Gaussian and non-zero-mean in most environments using practical waveforms with moderate-standoff beacons with good building penetration capability in the upper VHF to UHF bands. This is because the multipath inter-arrival times are less than the resolving power (~inverse bandwidth) of the waveforms, and any detected arrival is biased by the overlapping multipath impulse responses. Higher-frequency ultra-wideband (UWB) waveforms with better resolving power are possible, but have both building penetration and potential interference and regulatory issues. Thus, in most cases using moderate bandwidth waveforms we cannot "average our way out" of the problem in a standard navigation Kalman filter combining direct measurements and smoothing inertial/DR. Instead, the algorithms must recognize the few uncorrupted arrivals (usually direct path, but potentially later-arriving multipaths using approaches like Draper's) and separate and exploit these "needles in a haystack". These unbiased measurements, and these measurements alone, must be integrated into the smoothed solution by the inertially/DR-driven navigation Kalman filter.

So what does it take to get enough good non-biased measurements to combat the multipath problem? The competitive DARPA RSN program is developing unique multichannel waveform and beacon structures, both moving and stationary, to get the needed spatial diversity and unbiased glimpses. In addition, a DHS-sponsored effort by the Worcester Polytechnic Institute (WPI) has shown recent work that provides a constructive proof through demonstration of a measurement geometry and algorithm suite that can suppress multipath enough to allow sub-meter location accuracies with 5 meter inverse-bandwidth waveforms. In reference [7] they show that a 60 MHz bandwidth “multi-carrier-wide-band” (MC-WB) waveform at 440 MHz center frequency with c/60 MHz = 5 m conventional resolving power can be used in moderately complex interior environments to achieve 0.4 to 1 meter rms location accuracies. They do this with a combination of “super-resolution” state space-based TDOA estimates [8, 9] and currently proprietary inversion algorithms to go from a large number (13-16) of well-distributed beacon equivalents. For the WPI system, the beacons are replaced by stand-off receivers, and the user device is a transmitter, providing the equivalent measurement geometry to what we have discussed thus far, but with the transmitters and receivers reversed. This system approach has advantages for the non-LPI needs of emergency first-responders, but it is generally unacceptable for warfighters to emit and give away their positions to a potentially unfriendly intercept system. Example results are shown in Figures 5 and 6, and are reported in detail in [7].

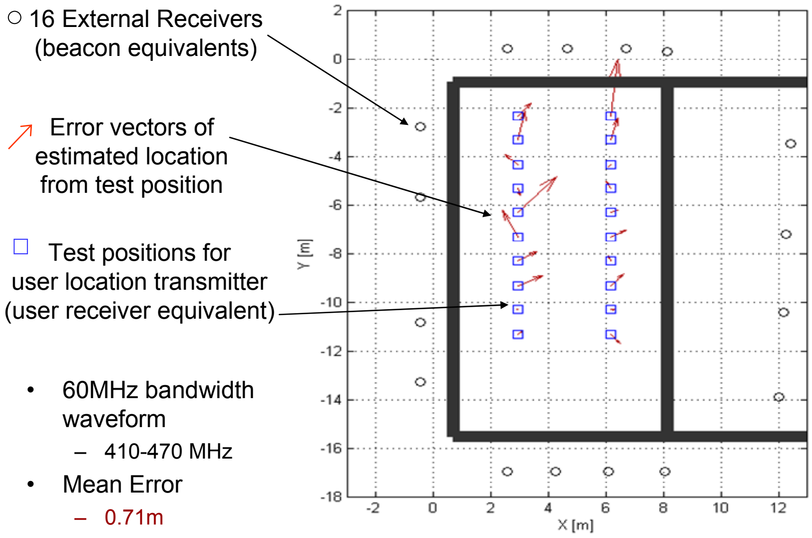

Figure 5. WPI interior measurement results. Highly spatially diverse 60 MHz (5 m equivalent pulse length) measurements provide sub-meter performance in an interior environment with multiple metallic obstructions and reflectors. (used by permission)

Note that in Figure 5 that the performance of the system generally improves as the user location inside the building moves further from multipath-causing walls, allowing more of the beacon (receiver, in WPI implementation) antennas to have multipath-free direct arrivals, and thus be available for error-reducing averaging after the multipath-induced outliers are rejected. A greater number of available beacons also increases the probability that a good beacon geometry is available that reduces the “DOP” (dilution of precision) due to poor “triangulation” geometries—essentially poorly-conditioned inverse estimation problems due to near-singular systems of equations to solve.

In Figure 6 the WPI data also show the potential for a non-linear decrease in performance with diminishing waveform bandwidth in a multipath-controlled, rather than SNR-controlled, environment.

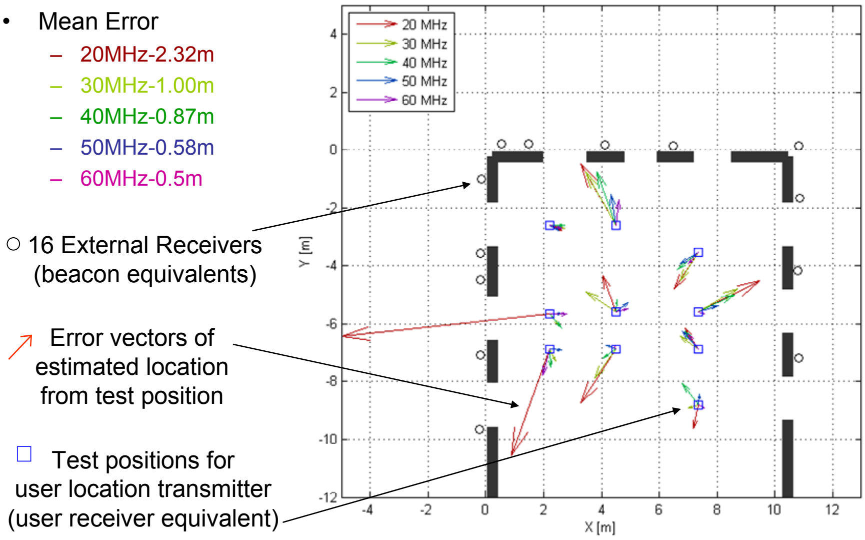

Figure 6. WPI interior measurement results. Results at different waveform bandwidths show that probability of outliers causing large error in location goes up with lower bandwidth, especially at the lowest bandwidth tested. (used by permission)

When not controlled by multipath, the rms Cramer-Rao lower bound estimation performance should scale proportional to bandwidth, since the TDOA standard deviation accuracy scales with this. For multipath-dominated environments, this bound is not reached until the arrival of interest is free and clear of overlapping multipath arrivals. This is still true even with more advanced super-resolution / multi-arrival hypothesis algorithms that explicitly try to compensate for the multipath (e.g. RAKE-like methods). Thus, one gets a thresholding phenomenon in which performance quickly degrades when significant non-resolvable (Δt < 1/bandwidth) multipath comes into play. This is particularly true when there are few, or non-moving, beacons and the spatial degrees of freedom available to separate clean arrivals from multipath distorted arrivals are not available. The excellent performance of the WPI measurements is thus attributable to:

- the relatively wide bandwidth of their waveforms (60 MHz),
- their "super-resolution" algorithms for further resolving beyond 1/bandwith,
- the large number (O(15)) of beacons that provide more opportunities for clean arrivals, and
- their algorithm for location estimation from overdetermined measurements which uses additional outlier rejection to preferentially select clean arrivals.

It is the authors' experience that "L1 norm" [10, 11] estimation algorithms using over-determined datasets provide an excellent method for automatic outlier rejection without *ad hoc* selection of outlier rejection thresholds, which generally require more solid knowledge of the statistical distribution of errors than is available. These L1 solutions automatically find the median or "robust" solution [12] and reject statistical outliers well by selecting only the N data points maximally consistent with each N-dimensional location hypothesis, fitting them exactly, and letting the errors on all others float. This provides a convenient indicator of which data are considered potential outliers. These may then be culled further to remove those on the tails of the distribution, and all remaining points re-inverted using a least-squares (L2) algorithm. This is the optimal procedure when the culled data are approximately Gaussian. Finally, while some readers may find it worrisome that so many beacons may be required to achieve good performance close to that governed by the resolving power of the waveforms (α 1/(bandwidth•√snr)) rather than multipath interference, we are encouraged that a constructive proof has been given showing that algorithms can be developed that exploit available spatial degrees of freedom well enough to achieve these results using waveforms whose bandwidth under-resolves the multipath interference inter-arrival time.

3.2 Exploitation of Signals of Opportunity (SoOP)

The WPI results discussed in the previous section are particularly encouraging if one's ultimate goal is the exploitation of Signals of Opportunity at the building or urban canyon/regional level. This is because in most cases, while SoOP waveforms will have much lower bandwidth than desired to resolve all multipath, they have high signal to noise ratios, and there are many SoOP in the environment with diverse characteristics and high spatial degrees of freedom are often available. A quick informal assessment of the readily available bands in the Washington DC area indicates that 75 strong signals are available in just the standard radio and TV broadcast bands. These do not even include satellite broadcast and communications signals (e.g. TV, radio, and internet distribution, Globalstar, Iridium, …), cellular telephony, and the increasing RF metropolitan area network signals. It should be noted that elevated satellite signals have special importance, providing much needed high-elevation measurements needed for the vertical component of low-DOP 3-D location estimation without specially placed beacons.

To illustrate the potential of SoOP, we cite the work by Rosum Corporation [1] that has developed a fielded capability to exploit TV transmissions for urban/in-building localization. For their markets, they have not yet tried to combat multipath to the room and hallway level, but they have solved the difficult system-level problems of characterizing uncontrolled SoOP to make them exploitable for localization and timing, and have dealt with the acceptance issue of user-receiver cost through advanced ASICs to implement their user equipment. The Rosum system concept and architecture is shown in Figure 7.

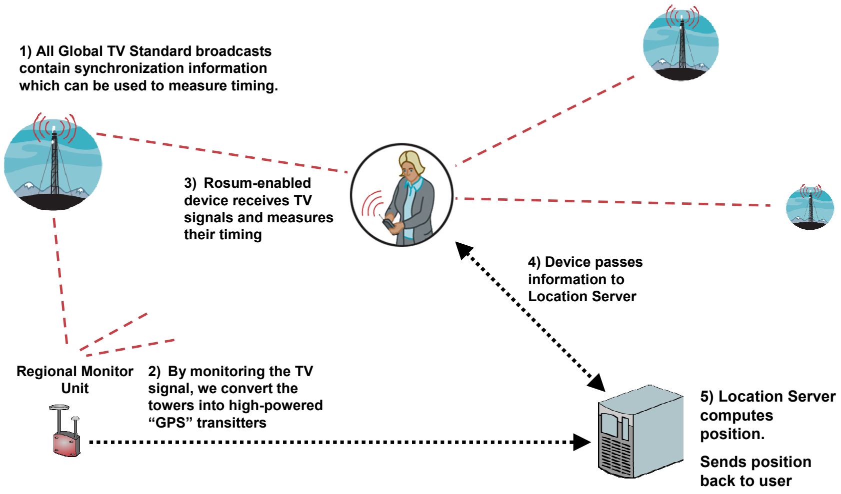

Figure 7 . ROSUM System concept and architecture. (used by permission)

By using the powerful low-frequency (50-750 MHz) TV signals, they have very high SNR to penetrate deep into buildings. This, and the fact that there are easily exploitable synchronization signals embedded in the TV waveform structure, provides a compelling reason to use the signals. However, their low bandwidth (~5 MHz) does mean that multipath is likely the dominant source of error in most scenarios. However, even with these ~60 m pulse length signals, outdoor performance is 16 m (50% CEP), or about ¼ of a conventional range cell. As seen in Table 1, performance in various urban indoor and outdoor environments is inversely proportional to urban complexity, and this is typically achieved with a small number of channels (3-5) being used at any one time. Since minimal outlier filtering can be done without fully exploiting a much larger number of SoOP channels, or exploiting user receiver motion, this leaves much room for potential improvements to reach the tougher tactical goals of military and first-responders.

Table 1 . Rosum system performance as a function of environment. Accuracy is 50% CEP. (used by permission)

		Accuracy in Meters		
Environment	Metro Area	Overall	Outdoor	Indoor
Rural-Urban	Nashua, NH	***23***	*16*	*29*
Rural-Urban	Needham, MA	***27***	*27*	*28*
Suburban-Urban	Santa Clara, CA	***28***	*28*	*29*
Urban	Washington, DC	***37***	*35*	*37*
Suburban-Urban	Edison, NJ	***38***	*34*	*41*
Dense Urban	Cambridge, MA	***47***	*40*	*52*
Extreme Urban	Manhattan, NY	***63***	*42*	*75*

Also encouraging is the fact that Rosum has found that a very sparse set of “Regional Monitoring Units” (the calibration reference stations of Figure 1) is sufficient for characterizing the SoOP in an area, currently on the order of one per metro area covered. Clearly as performance expectations are raised, the required density will go up. As part of their system offering, Rosum also has special purpose beacons that may be installed as needed to improve DOP or signal availability. These use the same signal structure as TV, and are tuned to use unallocated channels in the area to be covered. They can be a useful adjunct because, for example, TV towers will appear around the horizon for most users. To get vertical resolution for true 3-D localization in a building, a beacon may need to be placed on top of it or a nearby structure to obtain a good measurement geometry.

4. CONCLUSION

In the previous sections we have discussed three examples of advanced algorithms and system implementations that exploit spatial diversity of beacons and signals of opportunity. These show that the penetration and multipath problems of indoor/urban localization can be dealt with through the design principles we outlined in the introduction of the paper. We feel that additional emerging multipath suppression direct measurement system concepts and algorithms coming from DARPA and other Navy, Army, Air Force, and DHS programs will continue to improve indoor/urban geolocation systems meeting the requirements outline above, especially when optimally combined using robust (outlier-rejecting) navigation filter implementations and/or L1-based location “inversion” estimation techniques.

5. ACKNOWLEDGEMENTS

The authors would like to thank Professors R. James Duckworth and David Cyganski of WPI, and Mr. Jon Metzler and Mr. Todd Young of Rosum Corporation for permission to use the figures from their recent work. We would also like to thank Dr. Don Gustafson and Dr. John Elwell of Draper Labs for their discussions and insights. We would also like to thank the US Army CERDEC and the Air Force Research Laboratory for their technical, administrative, and requirement support for the DARPA RSN and SSN programs.

6. REFERENCES

[1] Arik Hesseldahl, “Rosum: Taking GPS Indoors”, Business Week Online, 26 April 2007, http://www.businessweek.com/technology/content/apr2007/tc20070426_529453.htm

[2] Eric Lithopoulos and Louis Lalumiere, “The GPS-denied Navigation and Mapping System”, Proceedings of WPI Precision Indoor Personnel Location and Tracking for Emergency Responders, Second Annual Technology Workshop, August 6-7, 2007 Worcester Polytechnic Institute, Worcester, MA. http://www.ece.wpi.edu/Research/PPL/Workshops/2007/agenda.html

[3] Carole Teolis, “TRX Sentinel Tracking System”, Proceedings of WPI Precision Indoor Personnel Location and Tracking for Emergency Responders, Second Annual Technology Workshop, August 6-7, 2007 Worcester Polytechnic Institute, Worcester, MA. http://www.ece.wpi.edu/Research/PPL/Workshops/2007/agenda.html

[4] Proposer Information Pamphlet, DARPA BAA05-32 “Robust Surface Navigation”, http://www.fbo.gov/servlet/Documents/R/399682/583229

[5] Proposer Information Pamphlet, DARPA BAA05-33 “Sub-Surface Navigation”, http://www.fbo.gov/servlet/Documents/R/399670/583221

[6] D. E. Gustafson, J.M. Elwell, and J. A. Soltz, “Innovative Indoor Geolocation Using RF Multipath Diversity”, IEEE PLANS Conference, San Diego, April 25-27, 2006.

[7] D. Cyganski, J. Duckworth, S. Makarov, W. Michalson, J. Orr, V. Amendolare, J. Coyne, H. Daempfling, S. Kulkarni, H. Parikh, B. Woodacre, “WPI Precision Personnel Locator System – Indoor Location Demonstrations and RF Design Improvements”, Institute of Navigation, 63rd Annual Meeting, Cambridge, MA, April 2007.

[8] D. Cyganski, J. A. Orr, R. Angilly and B. Woodacre, “Performance Limitations of a Precision Indoor Positioning System Using a Multi-Carrier Approach”, Proc. Institute of Navigation, National Technical Meeting 2005, January 24-26, San Diego, CA

[9] B. D. Rao, K. S. Arun, “Model Based Processing of Signals: A State Space Approach,” Proc. IEEE, vol. 80, no. 2, pp. 283-309, Feb. 1992.

[10] G.L. Duckworth, "A Robust Algorithm for Tracking of Drifting Acoustic Arrays in the Arctic," Conference Record, 21st Asilomar Conference on Signals, Systems and Computers, November 2-4, 1987, Pacific Grove, CA, Maple Press.

[11] G.L. Duckworth, D.C. Gilbert, J.E. Barger, "Acoustic counter-sniper system", Command, Control, Communications, and Intelligence Systems for Law Enforcement, SPIE Proceedings Vol. 2938, SPIE International Symposium on Enabling Technologies for Law Enforcement and Security (invited paper), Boston, MA, 19-21 November 1996

[12] Jon F. Claerbout, Fundamentals of Geophysical Data Processing: With Applications to Petroleum Prospecting, Blackwell Science, Inc., 1985.