CELLULAR LOCALIZATION FOR AUTONOMOUS DRIVING

A Function Pull Approach to Safety-Critical Wireless Localization

Russ Whiton

©SHUTTERSTOCK.COM/CYBRAIN

Cellular localization has received marked attention in academia and in industry, as indicated by a rapidly expanding volume of literature specific to vehicular environments. Impressive results have been shown for a number of problems as measured by metrics such as accuracy and latency. This article identifies five important requirements for cellular localization for safety-critical systems with a particular focus on autonomous driving (AD) and puts them in the context of industrial and academic trends and standardization. We show why autonomous operation requires special consideration and suggest research directions toward novel and practically implementable solutions, drawing lessons from decades of work on satellite-based localization for aviation landing systems. In addition, we highlight the benefit of cellular localization technology for safety-critical autonomous systems, showing the utility of a satellite-navigation independent absolute localization sensor with error overbounding.

Digital Object Identifier 10.1109/MVT.2022.3208392
Date of current version: 20 October 2022

Introduction

The challenge of AD has resulted in a surge of industrial activity and academic publications in the past 10 years. Novel proposals have been made for the associated hardware and software components, system architectures and even behavioral and societal impacts. Using the levels defined in the Society of Automotive Engineers' (SAE) widely used J3016 taxonomy, companies, universities, and research institutes are now developing, testing, and validating levels 3 and 4 automated driving features on public roads.

Functional architectures for AD are typically split into blocks from sensing through to actuation, enabling the vehicle to navigate the environment both on a large scale (lane selection and route planning) as well as on a small scale (lane position and orientation) [1]. This is most frequently accomplished with a broad suite of sensing technologies and a world map. Localization entails determination of position, orientation, and velocity in an external reference frame, and provides context to AD beyond the horizon of the perception sensors. This context can be provided with a world map, which can be as simple as a road-level navigation map or as intricate as a high-definition (HD) map rich with detail about road classification, lane markings, signs, and other detectable objects.

Only global navigational satellite systems (GNSSs) provide absolute estimates of position at coordinates surveyed in the same global reference frame in which map data are typically stored. GNSS has well-known limitations, primarily a strong dependence on physical view of satellites, weak signals at the Earth's surface, and is increasingly easy to jam and spoof by unsophisticated malicious actors.

Cellular localization has been developed as an enabling technology for use cases ranging from factory automation to AD, recognized in industry consortia such as the 5G Automotive Alliance and Third-Generation Partnership Project (3GPP). The communications world sees the natural affinity for cellular systems with their widespread deployment, relatively high power, and large communication bandwidth and granular direction determination to address localization needs for increased robustness and for expansion of the operational design domain to locations where GNSS reliability is always low, such as urban canyons [2].

To take the next step toward deployment of cellular localization for safety-critical systems including AD, we believe it is necessary to take a function pull approach, in which the needs of the AD localization subsystem, and the role of cellular localization within that subsystem, are considered from an application and safety engineering standpoint. This is done in accordance with automotive industry-wide technical standards that allow safety principles to be shared across engineering teams with common vocabulary and practices. In doing so, we place new types of requirements on cellular localization that have not been addressed in prior literature, which has considered the technology push of cellular localization within the architecture of cellular networks in terms of signals, operating frequencies, and bandwidths [3], or the performance achievable in isolation for automotive use cases without regard to sensor fusion or safety engineering principles, quantified by metrics like accuracy and latency [4].

ONLY GLOBAL NAVIGATIONAL SATELLITE SYSTEMS PROVIDE ABSOLUTE ESTIMATES OF POSITION AT COORDINATES SURVEYED IN THE SAME GLOBAL REFERENCE FRAME IN WHICH MAP DATA ARE TYPICALLY STORED.

We highlight performance metrics that we see would be most valued for AD, primarily the ability to establish quantified confidence in a localization solution in a reference frame that can be aligned with maps. We also show why multipath propagation poses difficult challenges for classic overbounding approaches adapted from aviation [5], and more generally the extent to which error overbounding can be ported from aviation landing systems to automotive use cases [6].

New constellations of researchers across disciplines are necessary to achieve this vision. This encompasses the fields of classical navigation, different subdisciplines of wireless communications, geodesy, as well as safety engineering. To this end, we offer a list of research topics that need to be addressed to realize such a solution, drawing lessons from decades of work done on localization integrity in aviation.

The manuscript is organized as follows: First, an introduction of AD localization architectures, safety-critical localization, and the state of cellular localization in literature and standardization is presented. The following three sections then identify novel requirements related to the gaps between the *pull* of AD functional needs and the *push* of cellular localization technology, formulating requirements and exploring the reasoning behind them. Finally, we conclude with an overview of the requirements and the functional benefits of their eventual realization.

State of the Art—Cellular Push and AD Pull

Localization for AD

Localization Function Goals

A localization function is responsible for providing the decision-making and actuation functions with the road geometry together with the position, orientation, and

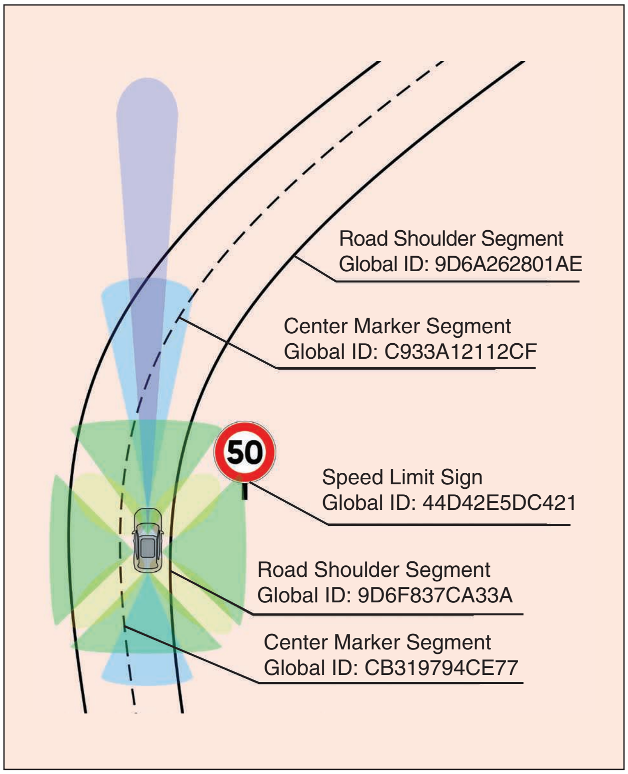

FIGURE 1 An HD map, which includes polynomial descriptions of dividing lines and road edges. Signage is also included as well as semantic information describing road type (rural road) and speed limit. The perception sensor suite provides a 360° understanding of the immediate vicinity and GNSS provides absolute location to narrow the search space and to help identify map irregularities.

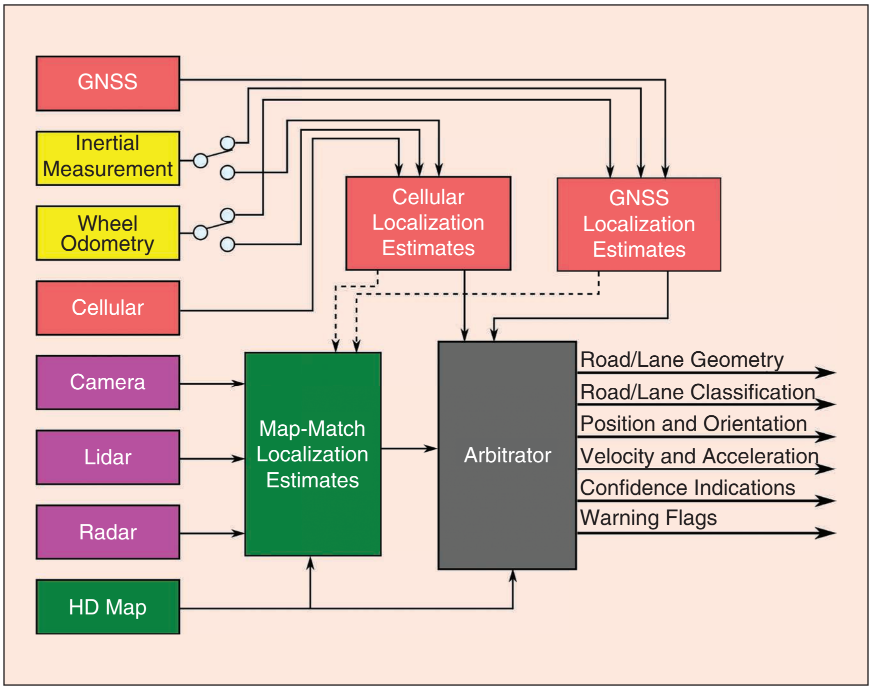

FIGURE 2 A possible implementation of an AD localization architecture including cellular localization. Alternative sensor fusion paradigms can be implemented at every stage, including feedback loops for deeper sensor coupling. Maximum separation of inputs is shown to prevent dependent failures.

movement of the vehicle inside that road geometry. This is closely related to the perception function that is responsible for identification and classification of static and dynamic objects immediately around the vehicle. The extent to which these two functionalities are separated in hardware and software is implementation dependent.

Simple driver support features such as lane departure warnings require estimation of lateral (cross-road) parameters like lane position. Higher-level AD requires longitudinal estimates (along road) and global estimates for lane selection and route-level navigation. Even cooperative maneuvering based on information sharing among vehicles is best understood in the context of both lane geometry and static infrastructure; traversing intersections or performing merging is done based on the rules and geometry of the road.

The understanding of the road comes from a geospatial database of map data. In the years since the Defense Advanced Research Projects Agency AD Grand Challenges, so-called *HD* maps have grown to include more types of information, with competing offerings in development from a number of map suppliers and functional architectures for crowd sourcing [7]. Figure 1 shows an example of an HD map segment.

Sensors, including lidar, radar, and cameras, match observations with map data in a two-step process of feature extraction and data association [8]. Absolute localization serves to narrow the search space for the initial search, bound probabilities of association, and detect anomalies in the map database. An AD localization architecture, with a proposed integration of cellular localization, is shown in Figure 2. Absolute localization is supplemented through proprioception, and the fusion of these sensors produces global position and orientation estimates and seeds a map-matching algorithm. An arbitrator block looks for consistency among the sets of estimates and forwards the expected road geometry and confidence estimations for decision-making and actuation, serving in part as a safety mechanism.

There are at least two important lessons to draw from this architecture. First, absolute positioning is solely responsible for long-term unbiased estimates of position, even if observations are noisy or low rate [9]. This is complementary to high-rate sensors such as inertial measurement units (IMUs), which can provide ≥100-Hz observations, or odometry. This means that the traditional metric of latency is not highly consequential for an absolute positioning sensor in a fused system.

Second, the centrality of GNSS shows why it can be a bottleneck for localization functionality. In anticipation of this need, investment in GNSS for mass-market applications has improved performance tremendously in recent years [10]. Among the chief improvements enabling mass-market use are new GNSS constellations coming online, updates to the broadcast satellite signals themselves, and production of inexpensive dual-frequency, multiconstellation receivers, and state–space representation (SSR) correction services including regional atmospheric models. However, even with these improvements, fundamental limitations of satellite navigation justify investment in alternative technologies.

Safety-Critical Localization

For safety-critical systems, error distribution tails are of paramount concern because infrequent extreme errors, if multiplied by the number of hours of operation, can result in significant risk exposure in aggregate. Failures are inevitable and they need to be addressed systematically over the complete lifetime of the vehicle.

To provide an example outside the localization domain, traction battery overheating in an electric vehicle can lead to chemical fires. There are many plausible scenarios that could lead to overheating after the vehicle leaves the factory. The cooling system might incur a leak stemming from mechanical damage or a broken coolant pump. An incorrect control unit software flash or software bug might impact cooling system control or cause an incorrect assessment of battery aging effects. None of these failures are effectively modeled as an extrapolation of nominal operation because, as with most practical engineering problems, tail errors are decidedly non-Gaussian even if such approximations may be of utility for normal operation. Detection, mitigation, and recovery mechanisms are necessary to prevent rare failures leading to injury or death. A battery module temperature sensor (redundant sensors in practice) can help flag a module before it overheats. Disconnecting the overheating module from the rest of the battery pack allows for fail-safe operation. The goal in the end is to achieve integrity, real-time assessment of when a system is safe for use and base control decisions on the most trustworthy of redundant systems.

Accomplishing integrity at an AD level requires structured ways of thinking about safety, as prescribed through multiple industry standards including ISO-26262 and ISO/PAS-21448 [12]. Before looking at the sensor level, hazard assessment and risk analysis (HARA) entails comprehensively listing potentially dangerous outcomes at a vehicle level, then assessing the risk of each in terms of severity, likelihood, and controllability. Functional safety goals are developed from the HARA results, and subsequently safety requirements are generated which can be met with hardware and software elements.

ABSOLUTE LOCALIZATION SERVES TO NARROW THE SEARCH SPACE FOR THE INITIAL SEARCH, BOUND PROBABILITIES OF ASSOCIATION, AND DETECT ANOMALIES IN THE MAP DATABASE.

A complete integrity solution in the localization domain entails generating real-time estimates of localization parameter error overbound and assigning quantified risk of violation of the overbounds that can help meet functional safety goals. In the parlance of the aviation world in which these concepts have been pioneered for localization estimate overbounding, the error overbound is referred to as a protection level and the quantified risk of violation is the integrity risk. Overbounding failures are called integrity failures, but errors are of course not observable in real time.

While it does not necessarily need to be formulated this way, the protection level is typically compared to an application-specific error alert limit, a threshold that can lead to identifiable dangerous operation such as incorrect determination of road, lane, or compass direction. This narrows the development and validation scope by excluding too-large or too-small failures. An overbound of 10 cm when the true error is 11 cm is not likely to be dangerous, nor is an overbound of 110 m when the true error is 111 m because the data will not be used downstream at this magnitude of uncertainty. An alert limit defines whether the system is available and suitable for use by the application, just as a temperature threshold may lead to the exclusion of a battery module and force the combined system to rely on the remaining modules.

Figure 3 illustrates the idea of localization integrity for lateral position estimates by showing two localization systems operating simultaneously for a single time epoch, with an integrity failure at this epoch for one of the systems. System 1 (blue) is unavailable; the lateral position overbound is too large to perform lane determination. This is safe operation, but unavailable for lane determination. System 2 (green) is available; the overbound indicates it is suitable for lane determination, but from the right part of the figure we can see that judgment was an overestimation of confidence (misleading information); the true position is outside the protection level, in this case incorrect lane determination. GNSS error overbounds have proven to be frequently unavailable and even unsafe in urban areas [11], which has prompted significant interest in cellular localization.

Technology Push of Cellular Localization, in Three Categories

Wireless localization, and cellular localization specifically, has a decades-long history. Readers are referred to [13] for an overview of fundamental wireless positioning techniques as well as a survey-of-surveys on wireless localization from 1977–2017. A number of new technologies spanning in maturity from recent commercial deployment to algorithms showing promise in simulations have been developed in recent years, which we generalize into three broad categories.

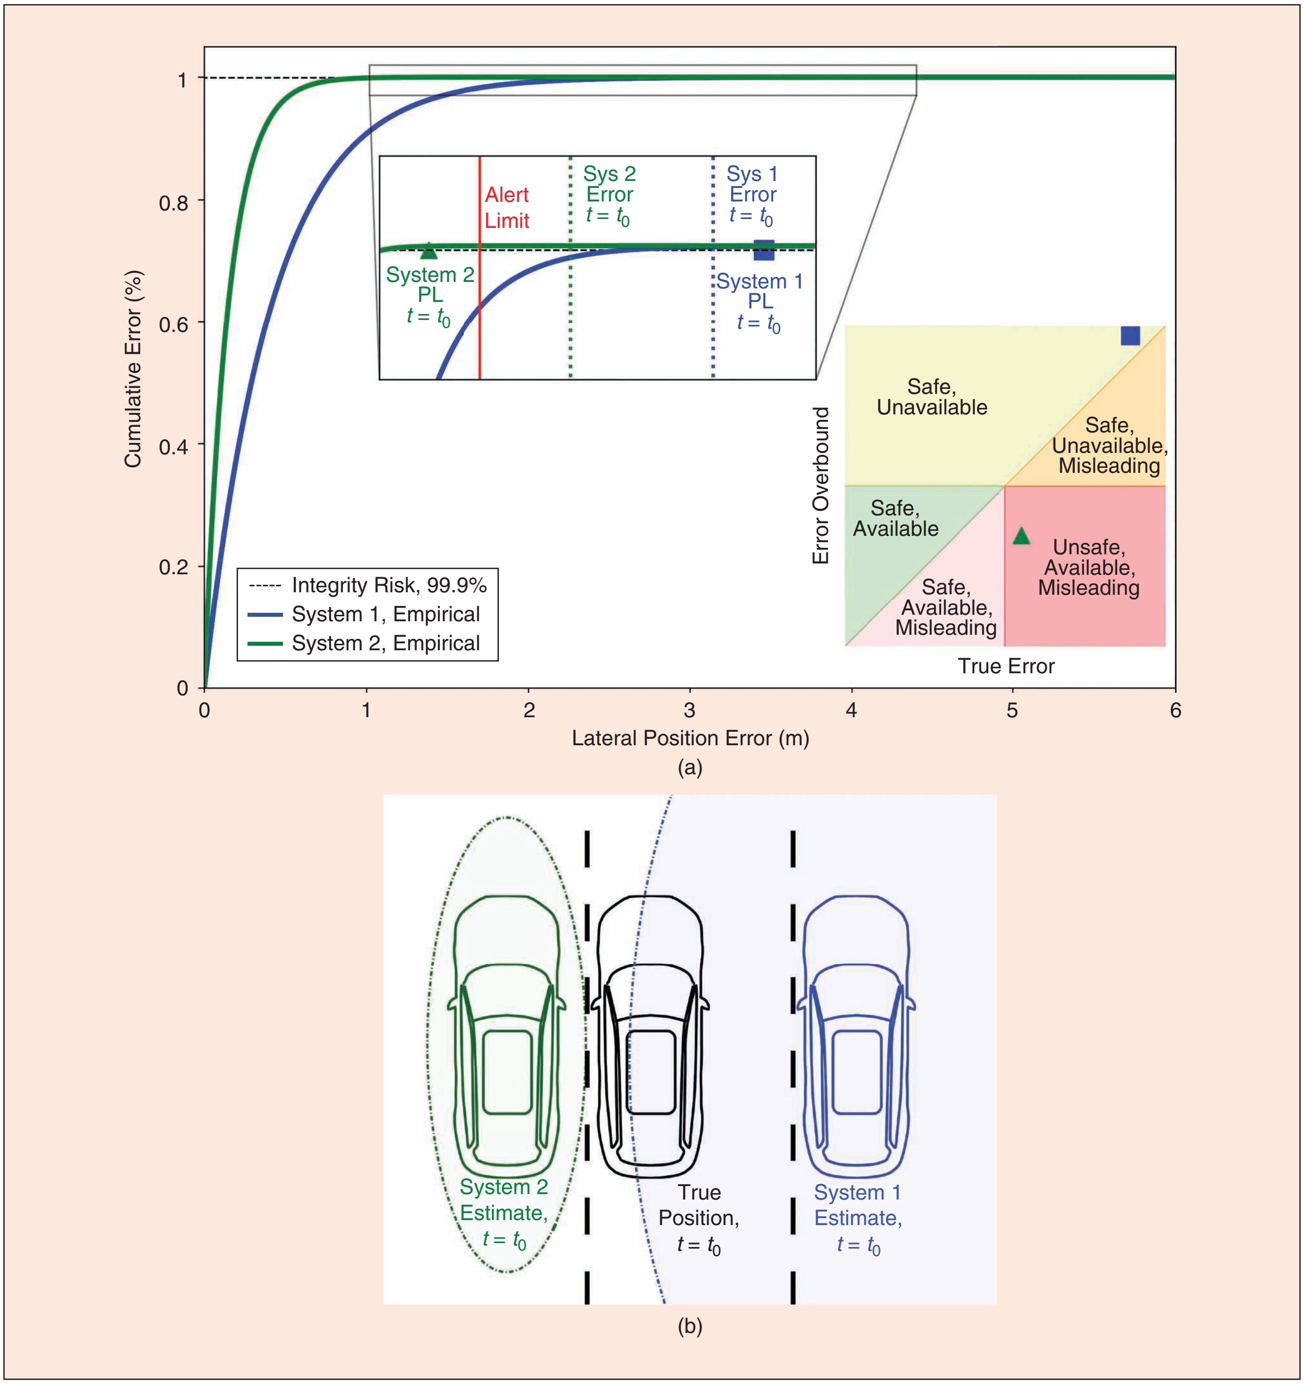

FIGURE 3 A single localization epoch including instantaneous error and protection levels for two systems at time t_0. (a) A backdrop of aggregate system performance shows that System 2 has better *average accuracy* than System 1. The inset graph shows a given navigation epoch t_0, where System 2 has smaller position error, but incorrectly underestimates the error overbound, leading to an integrity failure. Error distributions for localization can also include orientation or velocity estimation errors. Readers are referred to [6] and [11] for longer discussions of the Stanford Diagram. (b) Illustration of scenario. Note that the actual lane positions are incidental; the erroneous error bound determination below the Alert Limit would be considered an Integrity Failure regardless of the on-road position.

Machine Learning

Interest in statistical methods for wireless localization has increased since the year 2010 [14]. Supervised methods are most common. Location in space is used as the label and a representation of the wireless channel, a channel fingerprint, as the data. Feature engineering holds promise in finding channel representations that are well-suited for scalability and transferability. Additionally, cellular signals are just one possible input; other types of signals such as Wi-Fi can be integrated into the same estimation framework for improved performance.

Triangulation and Multilateration

New measurements, including 3GPP Release 16's multicell round-trip time, downlink angle of departure, and uplink angle of arrival, were introduced to serve a multitude of envisioned future use cases such as factory automation and vehicle-to-everything communication [3]. Enhanced bandwidth (up to 400 MHz with 120-kHz carrier spacing) results in time resolution capabilities superior to those of previous generations for positioning reference signals.

Opportunistic position estimation through observation of cellular pseudoranges (GNSS-like) has been demonstrated in academic literature for cellular transmitters, and error overbounding has even been applied (see, e.g., [5] and references therein). However, direct application of these methods has several limitations, which are discussed further in the requirements sections. Treating angles and delays of arrival as representative of actual transmitter–receiver geometry is sensitive to the soft-onset problem [14], when the direct path is much weaker than multipath components or not estimable. In the worst case, this can lead to hazardous operation when the receiver mistakenly considers multipath to be the direct path. Classic error overbounding approaches based on residuals testing such as [5] struggle in urban areas because of the strong correlation of errors [11]. Other mitigation strategies include statistical tests for non-line-of-sight identification [14].

Simultaneous Location and Mapping Based on Multipath Component Reflections

Multipath components can be decomposed in both space and time, with larger antenna aperture and larger signal bandwidth, respectively. Work in propagation modeling led to the insight that such reflections can be used as physical references in a simultaneous location and mapping (SLAM) problem formulation as virtual transmitters. This method has attracted significant attention in the academic literature (see, e.g., [15]). It is therefore possible to use multipath information advantageously rather than as a problem to be mitigated. This holds great promise for solving positive ranging biasing, but introduces new challenges, as we discuss in our requirements.

Requirement Area 1: Reference Frames

Geospatial data (the data stored in a vehicle's map database) is expressed with reference systems, reference frames, and geodetic datums. Common reference systems for map data include the international terrestrial reference system (ITRS) and the world geodetic system (WGS). Representative and well-defined datums are critical for precision applications (see, e.g., Section II.F in [10]).

Different reference frames for the same vehicle-transmitter (and virtual transmitter) geometry are illustrated in Figure 4, which also includes map data in the form of road geometry and signage. Estimates of Figure 2 are expressed in geodetic coordinates, Earth-centered, Earth-fixed Cartesian coordinates in this case. Map-matching entails feature extraction in the vehicle-centered coordinate system (sometimes expressed as *body frame* [9]), and data association with the global identifier, such as in Figure 1. Wireless communication systems are powerful in their ability to easily multiplex globally unique identifiers making the data association problem simple; unique transmitter IDs can be embedded with observations, realized in GNSS as the pairing of the navigation message with pseudorandom noise sequences.

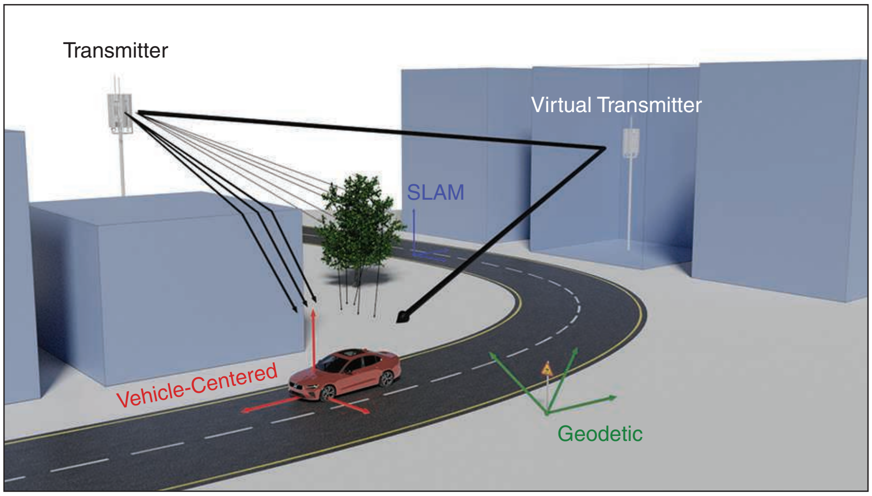

FIGURE 4 A vehicle, transmitter, virtual transmitter, and other objects in a map database with different wave propagation mechanisms including reflection, diffraction, and scattering. Expressed equivalently in three reference frames: Geodetic, Vehicle-Centered, and SLAM. Note that items in the map database such as signs and road marker indications are surveyed in geodetic coordinates. Map matching entails feature extraction from observations in a Vehicle-Centered frame and data association of map data stored in the Geodetic frame, including road geometry. Radio SLAM entails estimation of transmitters and virtual transmitters together with the vehicle itself in a SLAM frame which has its origin as an arbitrary starting epoch of the Vehicle frame.

TO INCREASE PERFORMANCE FOR RELATIVE LOCALIZATION ESTIMATION, INVESTMENT IN THE OTHER SENSORS IS FEASIBLE, BUT EVEN A MILLION-DOLLAR STRATEGIC GRADE IMU CANNOT OBSERVE GLOBAL POSITION.

The localization function is responsible for providing road information together with the state estimates, and therefore the estimates should be in the same frame as the map data, and our first requirement enables us to perform absolute localization in the absence of GNSS. **Requirement 1: Localization estimates shall be generated in a well-defined geodetic reference frame.**

The three categories of cellular localization have different relationships with the reference frames. Machine learning methods can have labels assigned in any coordinate system, but nominally it is logical to use a system directly transferable to the geodetic system realization. In this case, the performance bottleneck for anchoring in geodetic coordinates is the labeling noise for the surveyor collecting fingerprints. Multilateration and triangulation have the inverse limitation: If transmitter antenna array phase center is not accurately surveyed in a global reference frame (as well as orientation, for angular methods), subsequent vehicular estimates of position will be biased proportionally with the offset. Fortunately, the data association problem is trivial if transmitters have unique IDs.

Multipath SLAM has the most complicated relationship with geodetic coordinates. In the classic formulation of a SLAM problem, an agent is placed in an unknown environment and aims to localize itself in relation to landmarks corresponding to sensor observations, seeking to close the loop. This is an inherently different problem than working from a map database known a priori. A vehicle coordinate system at a certain time point is initialized to generate a new coordinate system; a few car lengths of driven distance previously on the road, in the case of Figure 4.

From the perspective of the localization system of Figure 2, this is a similar function to inertial measurement and odometry, simultaneous with landmarks that are not consequential in the manner that other references like lane geometry are. To increase performance for relative localization estimation, investment in the other sensors is feasible, but even a million-dollar strategic grade IMU cannot observe global position.

Anchoring in global coordinates for SLAM can be done at the receiver, in the propagation medium or the transmitter. Presuming knowledge of the global receiver position and orientation upon initialization is merely augmented proprioception. Assuming knowledge of relevant scattering objects [15] to anchor through the propagation medium would entail a dramatic expansion of the map database beyond the scope of road-level data. Anchoring through transmitters may hold more promise, even if virtual anchors lack a unique global identifier with a bijective mapping to their sender ID; they appear as time-delayed copies of the original transmitter. Additional parameterization of virtual transmitters, including angles of departure and relative time delay to the original transmitter, may be a viable strategy for global identification.

Requirement Area 2: Tail Risks

Hardware and software failures can lead to hazards at a vehicle level, as the example of battery overheating case shows. Certain failure modes (voltage supply, for example) are common across electronic systems, but identification of new fault modes unique to cellular localization will be necessary.

GNSS localization integrity in aviation provides a procedural template for how to perform systematic analysis of risks. Certain GNSS fault modes can be disregarded; rapid changes in ionospheric electron content are not a likely error source for a cellular system, for example. Other faults will be introduced that might be especially difficult to quantify using cellular base stations that do not provide the same integrity guarantees as navigational satellites, though future cellular systems may evolve in that direction. Receiver hardware will also be subjected to a new level of complexity. Something as trivial as oscillators or amplifiers failing to operate within their specification, for example, may lead to potentially dangerous operation.

Overdetermined systems permit for integrity checking, as decades of development in receiver autonomous integrity monitoring for GNSS receivers have shown [11]. Multiple-input, multiple-output (MIMO) systems offer new possibilities for monitoring, including diversified channel estimation, consistency checks across estimated parameters (angles and ranges), consistency across multiple frequencies, or consistency across systems (inertial and cellular, camera, and cellular).

A multidisciplinary approach is necessary to identify faults at a sensor level, as the weakest link in the chain will dominate the tail errors for the final localization estimates, expressed mathematically as a convolution of probability density functions. Both the faults themselves and areas for monitoring and mitigation flow naturally from the fault tree. This work is fundamental for development of a safety-critical system [12], and we should therefore include this procedural matter as an explicit requirement: **Requirement 2: Fault tree analysis shall be used to identify and monitor faults and mitigate risks.**

Ultimately, the goal in localization integrity is to establish quantified error overbounds by determining the frequency at which faults are expected to occur. Fault monitors can be effective at reducing the frequency at which undetected faults are encountered and can reduce the risk of subsequent vehicle-level hazards by orders of magnitude. After decades of research in aviation, the vertical position error overbound was reduced to 10 m with 10^{-7} failures/h integrity risk [8], but a position overbound of 10 m may not be a sufficiently powerful tool for an automotive use case. This leads to perhaps the most difficult requirement, but the most useful for the localization function. **Requirement 3: Localization estimates shall provide defined instantaneous upper error bounds.**

With this end goal, it is useful to start considering the placement of cellular localization in the architecture of Figure 2. Designing for accuracy and integrity are often in direct opposition, as choices that might lead to better accuracy, such as tighter filter tuning, may also lead to infrequent extreme errors. This is why the proposed architecture in Figure 2 does not use tighter feedback loops or permit downstream sharing of proprioception sensor data, to prevent dependent failures among systems (specifically common cause failures, in ISO 26262 taxonomy).

It is important to keep in mind that only some hazards in an automotive environment are a predictable function of absolute location and orientation, as collision risks are largely derived from dynamic objects. Drifting outside of a lane (see, e.g., [6]) can be hazardous, but because map data are by definition historical, AD vehicles are vulnerable to dynamic environmental changes to the semantic labels or physical geometry such as construction [7], or (less frequently) earthquakes. An important consequence of this is that to a much larger extent than aviation, absolute location certainty can only contribute to, but not satisfy, safety goals consequential to the localization system. Redundant systems are necessary, and perception and localization are necessarily interwoven. If lanes are not visible, e.g., in snowy conditions or road construction, then driving within historically defined lane boundaries is hazardous.

An illustrative example of a single AD hazard from a HARA analysis is shown in Figure 5. The vehicle-level hazard propagates down to functional safety requirements on a sensor level. The benefit of cellular localization for AD is the blue AND gate providing diversified redundancy. Satellite and cellular localization can work in a redundant fashion when both are available, but either can protect against hazards on an individual basis, and with cellular localization, absolute localization is possible in environments where GNSS is completely unavailable.

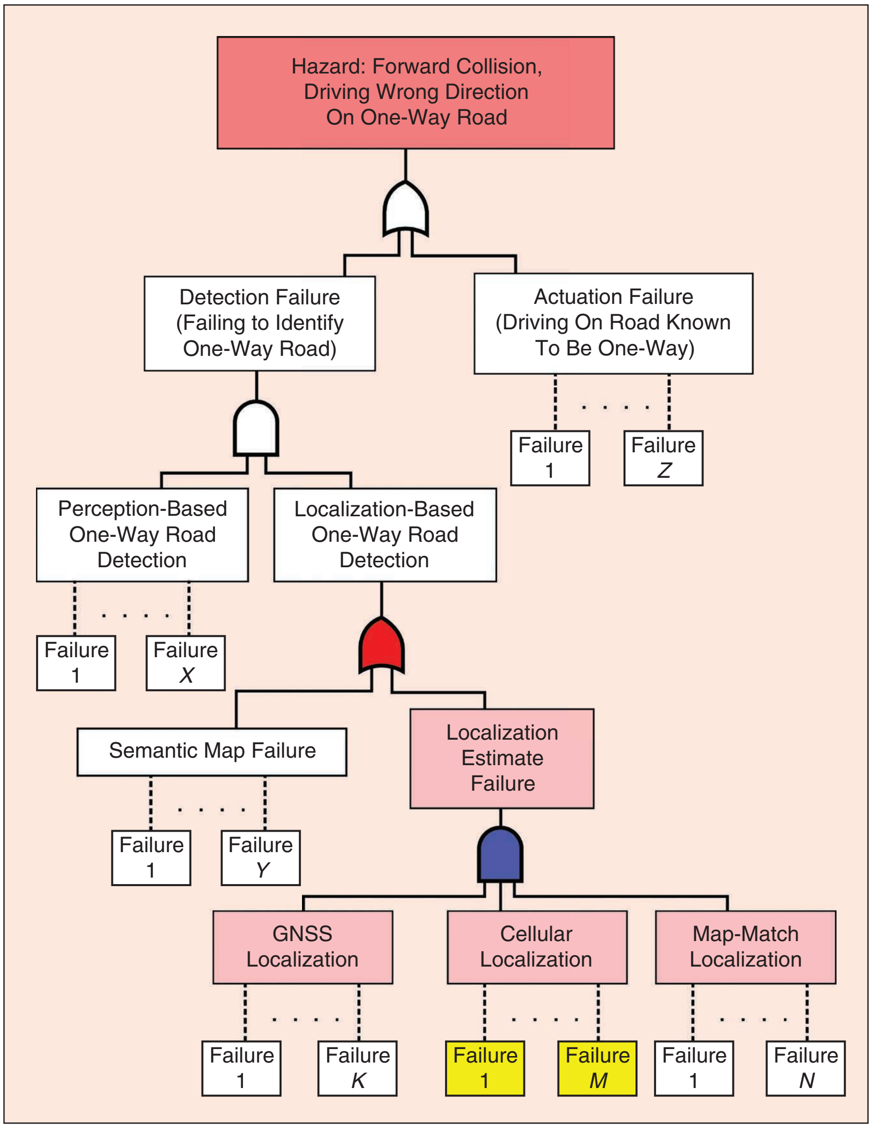

FIGURE 5 Part of a fault tree for a functional safety concept. AND gates allow for a multiplication of risk probabilities (reduction in risk), in the absence of dependent faults. OR gates result in a summation of risk. Note that the OR gate between absolute localization and map data means that the map is a bottleneck for any safety goal met with the help of absolute localization. Cellular localization Faults (yellow) are unexplored in literature. GNSS or Cellular blocks can include fusion with proprioception.

Requirement Area 3: Wave Propagation

Far-field propagation is typically modeled as a sum of plane waves and, for some more advanced propagation models, an additional diffuse multipath component term

IT IS IMPORTANT TO KEEP IN MIND THAT ONLY SOME HAZARDS IN AN AUTOMOTIVE ENVIRONMENT ARE A PREDICTABLE FUNCTION OF ABSOLUTE LOCATION AND ORIENTATION, AS COLLISION RISKS ARE LARGELY DERIVED FROM DYNAMIC OBJECTS.

to capture components that are too strong collectively to ignore when estimating plane waves but not resembling a plane wave sufficiently (diffuse waves in an urban canyon, for example) to estimate individually [15]. Single-bounce reflections from the buildings shown in Figure 4 might result in estimable plane waves, and diffraction over rooftops or scattering might contribute to diffuse multipath, but the distinction is dependent on physical geometry, measurement aperture and signal processing capabilities.

A physical model well representative of the physical reality of wave propagation allows for more precise localization estimates. Multilateration methods, whether based on clock synchronization or time-of-flight, typically presume that the first (and usually strongest) arriving wave is a direct path between the transmitter and receiver. They are therefore subject to ranging biasing because of multipath and blocked line-of-sight. Even if larger bandwidths offer improved time resolution, the very properties that have driven the MIMO revolution in wireless communication pose challenges for multilateration. In MIMO communication, the number of linearly independent paths from transmit array to receive array (which can increase the rank of the channel matrix) allows for linear scaling of data rate. Multipath is considered beneficial because it serves both to enable greater physical coverage and to achieve higher data transmission rates. This is a strong contrast to GNSS, in which signals and hardware are designed to push reality to fit the multilateration model better, with polarized antennas and signal design that explicitly aims to physically suppress multipath, leaving the first and strongest signal if it is present. Scientific measurements in geodesy even employ elaborate choke ring structures for multipath suppression.

The rich scattering environments in which GNSS is subject to range biasing are correlated with those in which cellular transmitters will also experience blocked line-of-sight signals and heavy multipath, such as urban canyons or parking garages. Mitigation strategies suggested for GNSS have included improved antenna polarization, map consistency checks or other cross sensor utilization (see [11] Table 2). MIMO systems should be able to move from mitigation

TABLE 1 ***AD function requirements, research directions, and benefits.***

Requirement	*Function Benefits*	*Possible Research Directions*
Localization estimates shall be generated in a well-defined geodetic reference frame	■ Geographical expansion of AD operational design domain to GNSS-denied areas ■ Relaxation of GNSS integrity requirements when cellular localization is available ■ Relaxation of bias drift specification for IMUs ■ Absolute localization reference sources that move seismically with other items in the map database	■ Surveying and parameterization of transmitters and virtual transmitters in a global frame ■ Cross-sensor perception to correlate cellular localization references with observations from other sensors ■ Cellular fusion with proprioception
Fault tree analysis shall be used to identify and monitor faults and mitigate risks	■ Procedural compliance with functional safety	■ Identification of localization faults related to antennas, radio electronics, and channels ■ Integrity monitoring techniques for different observation models
Localization estimates shall provide defined instantaneous upper error bounds	■ Quantitative contribution to the functional safety concept ■ Reduced search space and higher confidence for map-matching	■ Novel statistical overbounding approaches ■ Scalable data collection hardware and software platforms for wireless systems
The channel estimation shall resolve multipath components	■ Absolute localization in urban canyons, parking garages, indoors ■ Decorrelation of GNSS and cellular localization failures	■ Propagation models providing more granular resolution of diffuse components ■ Computationally efficient and redundant channel estimation
Position, heading, and velocity estimates shall be generated	■ Stationary orientation estimates independent of inertial measurement ■ Odometry-independent source of velocity estimates ■ GNSS spoofing detection	■ Ego-velocity estimation from Doppler observation ■ Dynamic receive phase calibration and antenna pattern characterization

to utilization. With this in mind, we can formulate our next requirement to ensure performance works well when it matters most. **Requirement 4: The channel estimation shall resolve multipath components.**

One promising opportunity afforded by MIMO systems is the ability to both steer and estimate angles of departures and angles of arrival. This has potential benefits for integrity, but it also affords opportunities for velocity and orientation estimation. Multi-antenna GNSS receivers (common in construction or agriculture, but not on passenger vehicles due to cost and packaging) can provide the ability to perform static orientation estimation, but cellular systems can estimate angles at both sender and receiver (and are not subject to International Traffic in Arms Regulation in the United States, as some multiantenna GNSS receivers are). Three-dimensional receiver orientation estimation is possible if the physical geometry of the array permits angular resolution in both azimuth and elevation, so yaw, pitch, and roll angles can be estimated without inertial measurement of gravity or magnetic measurement of compass direction. For a vehicle that is highly constrained in both pitch and roll, heading estimation is typically a primary constraining factor for localization system performance.

Similarly, the ability to perform Doppler estimation allows for an independent estimation of velocity. Dynamic objects in the environment such as other vehicles result in multipath Doppler contributions that are additive to ego-vehicle velocity vector, but reflections from a static transmitter off of static objects in the environment will have Doppler contributions derived only from ego-vehicle velocity and frequency mismatch. With the full dynamism of the sensor in view, we can formulate our final requirement. **Requirement 5: Position, heading, and velocity estimates shall be generated.**

Conclusions

The five requirements introduced for safety-critical cellular localization for AD, together with a summary of the functional benefits and research directions for realization, are listed in Table 1. We contend that a cellular localization sensor that can provide position, heading and velocity estimates with error overbounds in an external reference frame is a powerful tool for autonomous drive systems.

Safety-critical applications provide a unique set of challenges, as a long history of integrity systems for satellite navigation in aviation has shown. The natural affinity between advanced communication technologies and localization should make such solutions feasible to implement on a wide scale.

Acknowledgment

This work was supported by the Swedish Innovation Agency VINNOVA through the MIMO-PAD Project (Reference 2018–05000). The author would like to thank Max Boerboom at Volvo Car Corporation for his assistance in rendering Figure 4.

Author Information

Russ Whiton (russell.whiton@volvocars.com) is a senior hardware engineer with Safe Vehicle Automation at Volvo Car Corporation, Gothenburg 405 31, Sweden, and a Ph.D. student at Lund University. His research interests include localization and wireless communications. He received his M.S. degree in electrical engineering from Lund University.

References

[1] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE Access*, vol. 8, pp. 58,443–58,469, 2020, doi: 10.1109/ACCESS.2020.2983149.

[2] Z. Kassas, J. Khalife, K. Shamaei, and J. Morales, "I hear, therefore I know where I am: Compensating for GNSS limitations with cellular signals," *IEEE Signal Process. Mag.*, vol. 34, no. 5, pp. 111–124, Sep. 2017, doi: 10.1109/MSP.2017.2715363.

[3] S. Dwivedi et al., "Positioning in 5G networks," *IEEE Commun. Mag.*, vol. 59, no. 11, pp. 38–44, Nov. 2021, doi: 10.1109/MCOM.011.2100091.

[4] S. Bartoletti et al., "Positioning and sensing for vehicular safety applications in 5G and beyond," *IEEE Commun. Mag.*, vol. 59, no. 11, pp. 15–21, Nov. 2021, doi: 10.1109/MCOM.011.2100339.

[5] M. Maaref and Z. M. Kassas, "Autonomous integrity monitoring for vehicular navigation with cellular signals of opportunity and an IMU," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 6, pp. 5586–5601, Jun. 2022, doi: 10.1109/TITS.2021.3055200.

[6] T. G. Reid et al., "Localization requirements for autonomous vehicles," *SAE Int. J. Connected Autom. Veh.*, vol. 2, no. 3, pp. 173–190, 2019, doi: 10.4271/12-02-03-0012.

[7] K. Jo, C. Kim, and M. Sunwoo, "Simultaneous localization and map change update for the high definition map-based autonomous driving car," *Sensors*, vol. 18, no. 9, p. 3145, 2018, doi: 10.3390/s18093145.

[8] M. Joerger and M. Spenko, "Towards navigation safety for autonomous cars," Inside GNSS. Accessed: Feb. 11, 2022. [Online]. Available: https://insidegnss.com/towards-navigation-safety-for-autonomous-cars/

[9] P. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd ed. Norwood, MA, USA: Artech House, Mar. 2013.

[10] N. Joubert, T. G. Reid, and F. Noble, "Developments in modern GNSS and its impact on autonomous vehicle architectures," in *Proc. 2020 IEEE Intell. Veh. Symp. (IV)*, pp. 2029–2036, doi: 10.1109/IV47402.2020.9304840.

[11] N. Zhu, J. Marais, D. Bétaille, and M. Berbineau, "GNSS position integrity in urban environments: A review of literature," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 9, pp. 2762–2778, Sep. 2018, doi: 10.1109/TITS.2017.2766768.

[12] O. Kirovskii and V. Gorelov, "Driver assistance systems: Analysis, tests and the safety case. ISO 26262 and ISO PAS 21448," in *Proc. IOP Conf. Ser., Mater. Sci. Eng.*, 2019, vol. 534, no. 1, p. 012019, doi: 10.1088/1757-899X/534/1/012019.

[13] J. A. del Peral-Rosado, R. Raulefs, J. A. López-Salcedo, and G. Seco-Granados, "Survey of cellular mobile radio localization methods: From 1G to 5G," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 2, pp. 1124–1148, 2nd Quart. 2018, doi: 10.1109/COMST.2017.2785181.

[14] D. Burghal, A. T. Ravi, V. Rao, A. A. Alghafis, and A. F. Molisch, "A comprehensive survey of machine learning based localization with wireless signals," 2020, *arXiv:2012.11171*.

[15] K. Witrisal et al., "High-accuracy localization for assisted living: 5G systems will turn multipath channels from foe to friend," *IEEE Signal Process. Mag.*, vol. 33, no. 2, pp. 59–70, Mar. 2016, doi: 10.1109/MSP.2015.2504328.

VT