INTEGRITY

Positioning Beyond Accuracy

ALISTAIR ADAMS
SWIFT NAVIGATION

Despite being a critical component of the autonomous sensor suite, precision GNSS is not yet a widely adopted sensor in ADAS (Advanced Driver Assistance Systems). This may be due to legacy trust issues in the accuracy of the GNSS signal. However, there are several applications where GNSS can add significant value. In this paper, we explore use cases where sensors are used and explain how trust in the GNSS signal is a key driver in the acceptance of its use in making ADAS and autonomous vehicles safer.

Geofencing ADAS and Autonomous Operation

"Are we there yet?" The sound every parent dreads from a bored kid in the backseat on a road trip is echoed by those who in 2015 spoke of autonomy as something that would be completed by 2020. While huge strides have been made, almost everyone limits operations to specific roads using geofencing. Most of the time, everyday GNSS accuracy is sufficient but sometimes adjacent roads and GNSS positioning errors can mistakenly allow or disallow operation on a particular road. When it comes to vehicular safety, "most of the time" is not good enough and trust in the position is critical.

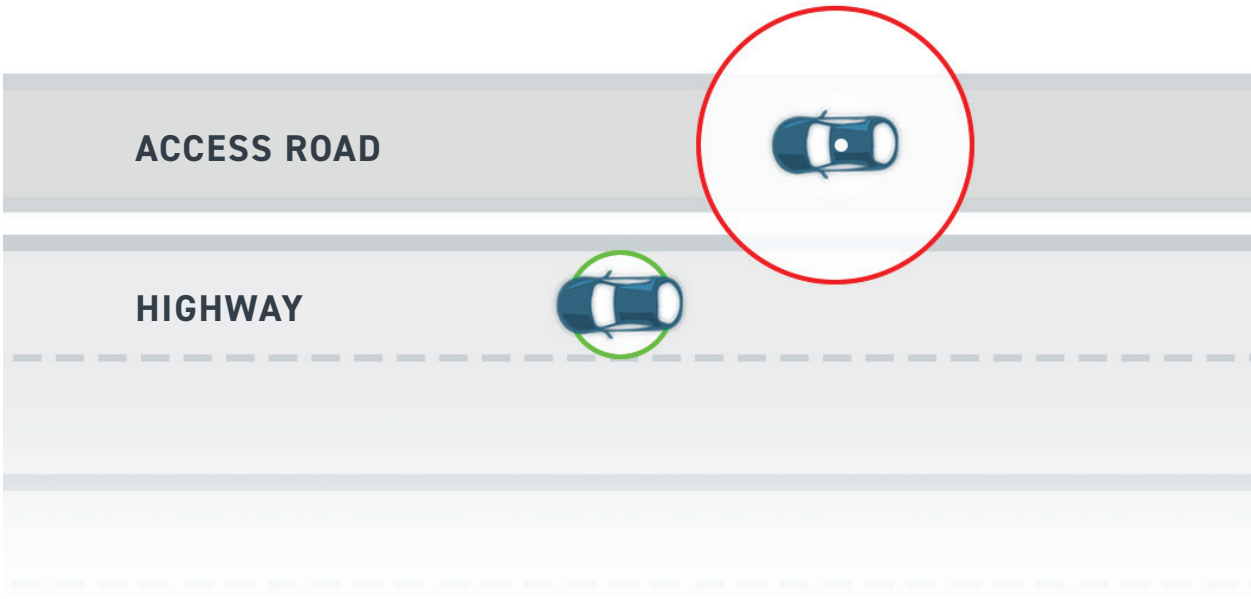

FIGURE 1
An access road next to a highway could result in the system placing the vehicle on the wrong road

Localization

Localization is the process of figuring out where we are in our environment. Cameras and LiDAR feed information into a perception system that is then matched with a high-definition map, containing information not normally found in standard consumer maps. Typically, these maps are accurate to 10 cm or better and define the location of curbs, lane markers, traffic signs and other road furniture. The localization algorithm takes the information from the perception system and searches the map to find the best fit for where the vehicle might be. This approach is computationally expensive but can be reduced by using a trusted, accurate, absolute GNSS position to provide a better starting point that considerably constrains the search space thus allowing the use of lower-cost computer hardware and even lower-cost sensors.

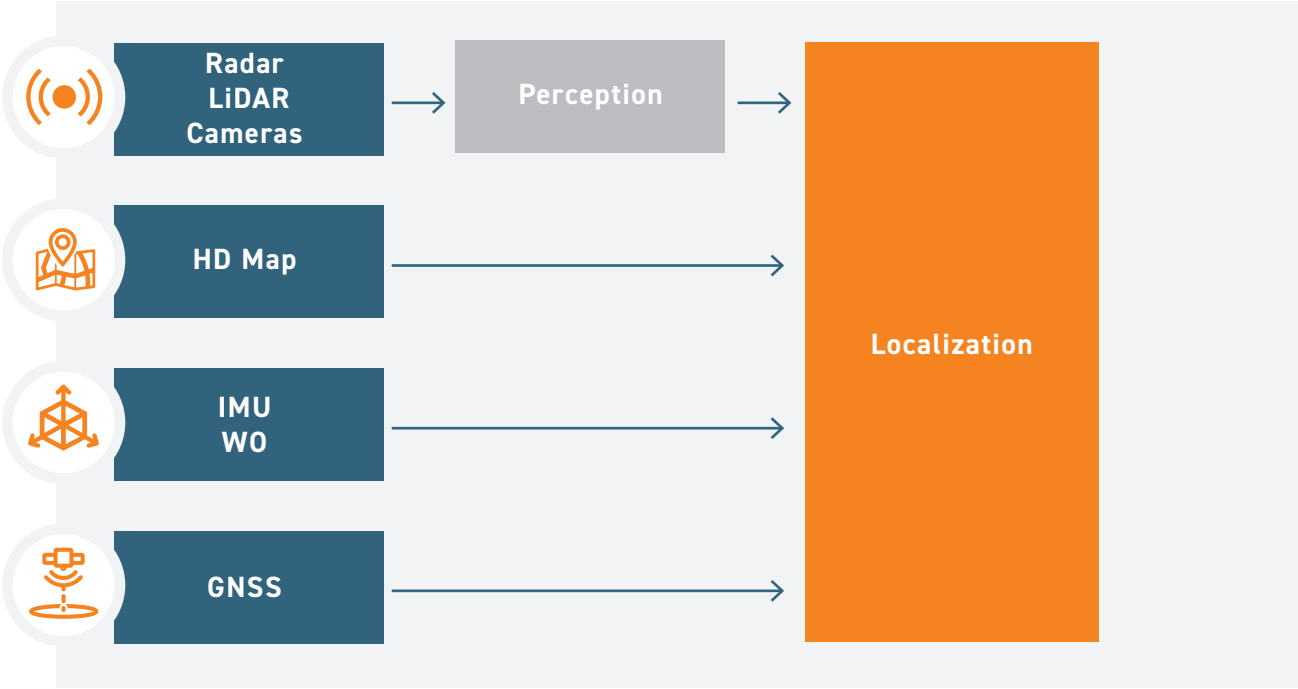

Feature Sparse Environments

Localization—based on an HD (high-definition) map with cameras and LiDAR—works great in urban environments with many features to key off. Feature sparse environments prove more challenging. A trusted GNSS position saves the day for lane-keeping operations, especially when lane markings are worn or in poor weather, such as heavy rain, fog and snow, where cameras and LiDAR are compromised. GNSS allows for faster, more accurate localization and lower fidelity maps.

V2X

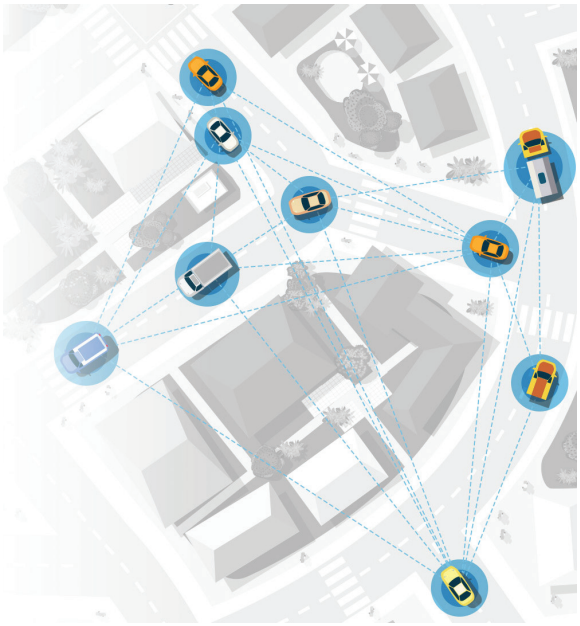

Vehicle-to-Vehicle, Vehicle-to-Infrastructure and Vehicle-to-Pedestrian fall under the term V2X (Vehicle-to-Everything). At its base are messages announcing, "here I am", "this is what I see" and "this is what I'm doing or going to do". Using such messages, other participants in the V2X community can adjust their behavior. Inherent in these messages is a common absolute positioning frame. The 5G Automotive Association (5GAA) applies a set of positioning requirements to different use cases for V2X with the highest accuracy requested as three sigmas. This means that 99.7% of the time the accuracy requirement is met but also that three times out of 1000, it could be wrong, threatening the reasonable safety tolerance. Swift believes that much higher levels of trust are required for safe V2X use cases.

Precise Mapping

Precise mapping HD maps are emerging as a key part of ADAS and autonomous solutions. HD maps are expensive to create and maintain, largely done by driving and re-driving with expensive surveying vehicles. Current estimates[1] show that just 2% of the US roads are mapped to HD levels. The concept of self-healing-maps has been around for about a decade but the amount of vehicles capable of contributing to the healing remains very small. If corrections from many vehicles are to be stitched together, then it's going to be a lot easier if there's a common, accurate anchor—an anchor that a trusted precise absolute GNSS position can provide.

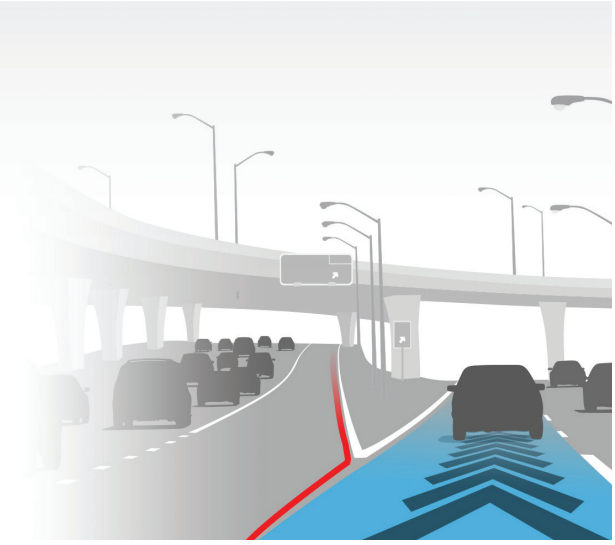

Integrity

All these use cases have one thing in common; a need to trust the GNSS position.

The GNSS world uses the term Integrity which is described as:

> A measure of **trust** that can be placed in the **correctness** of the information supplied by the system. Provides **timely warnings** when the information **should not be used**.

History of Integrity

The concept of GNSS Integrity began in the 1990s for airlines. Various algorithms were iterated over time but in general, sub-meter threats—that is errors that contributed less than a meter in error—were ignored since these were too small to be considered significant to this industry. An understandable decision given that aviation designs with an Alert Limit of 35 meters.

Automotive is understandably much more stringent, with alert limits in the 1-3 meter range. Automotive also must deal with a lot more multipath in the environment than aircraft, which have the benefit of being in the sky away from anything or landing on airstrips with a lot of space around them. Ground vehicles are impacted by signals reflected from adjacent vehicles and buildings. Automotive also uses lower-cost antennas, further adding threats to the integrity of the position.

Accuracy vs. Integrity

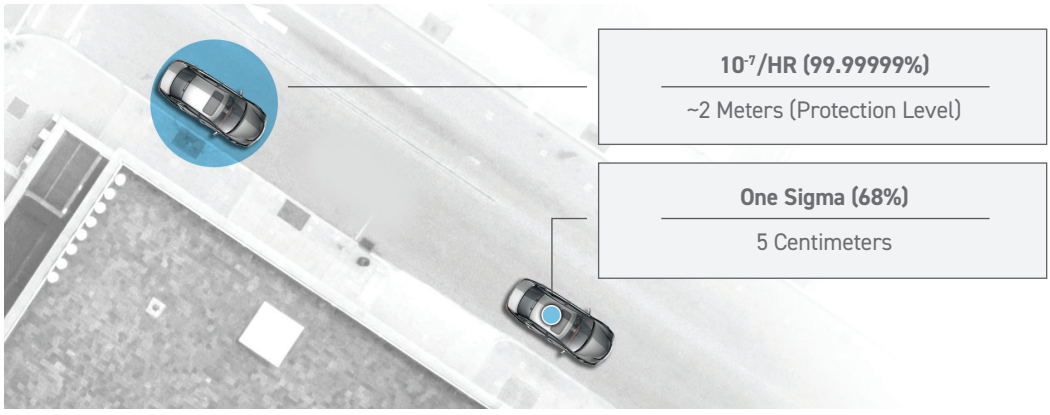

In GNSS terms, accuracy is often quoted in terms of standard deviations. This assumes a Gaussian distribution (bell curve) that works until the long tail is reached. A typical headline figure of one sigma (68%) may be quoted with five centimeters—well within

the realms of automotive grade and cost-sensitive components. But 68% means that one out of every three times it is more than five centimeters and it can occasionally be tens of meters off, especially in cities. Integrity factors in all the sources of errors, eliminates erroneous satellite signals and provides a much more certain position at a larger confidence circle. The error is referred to as a *Protection Level* in the integrity terms. The probability that the real error exceeds the Protection Level is called the *Target Integrity Risk* (TIR), with the goal to make that probability as small as possible, down to 10^{-7} per hour. An *Alert Limit* defines the maximum allowable protection level at the system level before declaring that the GNSS System is unavailable. The diagram below shows the relationships between Protection Level (PL), Alert Limit (AL) and the Position Error (PE). The PL grows and shrinks depending on the GNSS signal environment. In the diagram on the right it has exceeded the Alert Limit and the system is considered *unavailable*.

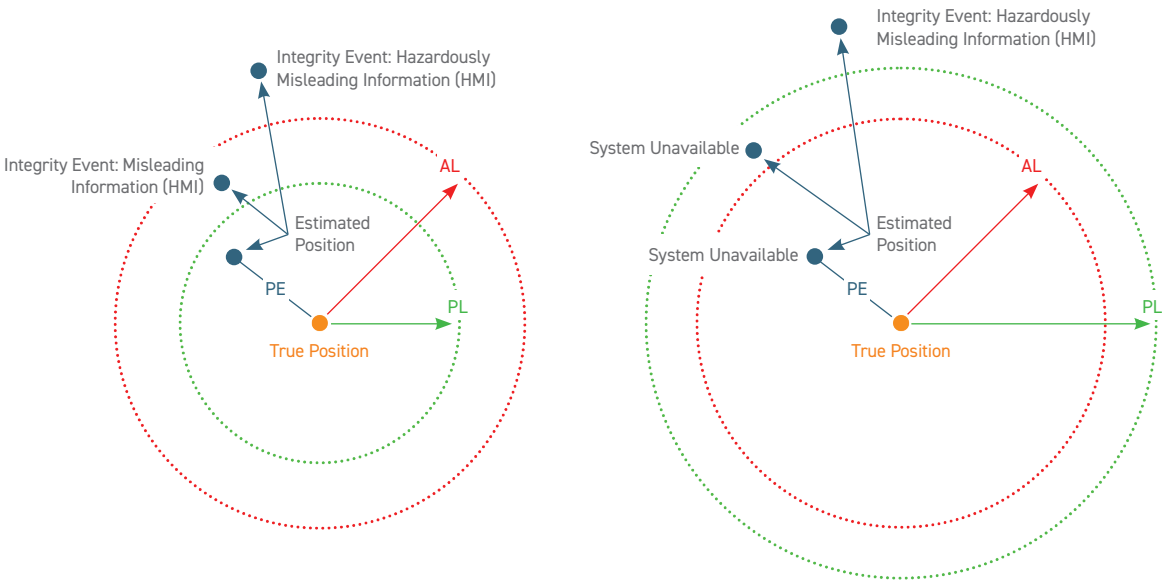

GNSS Positioning and Sources of Error

GNSS satellites are continuously transmitting a signal. A GNSS receiver estimates how long it has taken for the signals to reach it and calculates the location of the satellite in the sky using known orbit data. This duration—when converted to distance between transmitter and receiver—is called the pseudorange. Simply put, the distance from the satellite is the speed of light multiplied by the time taken.

If only calculating position were that easy. While the receiver gets updates every two hours of the satellite orbit—known as the ephemeris—during that time the satellite is pushed and pulled by solar winds and gravity, which can result in errors of upwards of two meters. The clocks on the satellite can also drift, so one doesn't know exactly how long the signal took to reach the receiver, also contributing to up to two meters of error. The ionosphere—ionized by solar radiation—delays GNSS signals adding up to another three meters of error while the wet troposphere adds another meter. Finally, multipath in cities can result in up to as much as 100 meters of error though newer receivers detect and eliminate a lot of multipath.

As the precision and accuracy increases, smaller contributions to errors such as the antenna orientation, signal deformation and thermal effects also start to have an impact.

Correcting for Errors and Integrity

To correct for these variety of errors, Swift has deployed—and utilizes—a network of ground-based reference stations to monitor the signals from the GNSS satellites. These signals are uploaded to Swift's Skylark™ Precise Positioning Service, which has models built to correct all the errors, including global errors like clocks and orbits. Others are more local, such as the ionospheric and tropospheric delays. These corrections are sent to the GNSS receiver running a positioning engine such as Swift's Starling®.

For the positioning engine to calculate a protection level, it needs to know the error distribution of these corrections. This is what the integrity chain in Swift's Skylark does. It assigns bounds to each correction and then cross-checks them against known reference stations to ensure accuracy. It also monitors the satellite network and atmosphere for faults which could impact the ability for GNSS positioning systems to operate accurately and safely. These correction bounds are transmitted to the positioning engine, which uses them to calculate protection levels. Skylark also sends flags to indicate if a satellite should not be used.

In addition to the errors already mentioned which are inherent to satellite positioning, other threats to a GNSS signal exist. These threats occur infrequently but left undetected could cause a positioning solution to be unsafe.

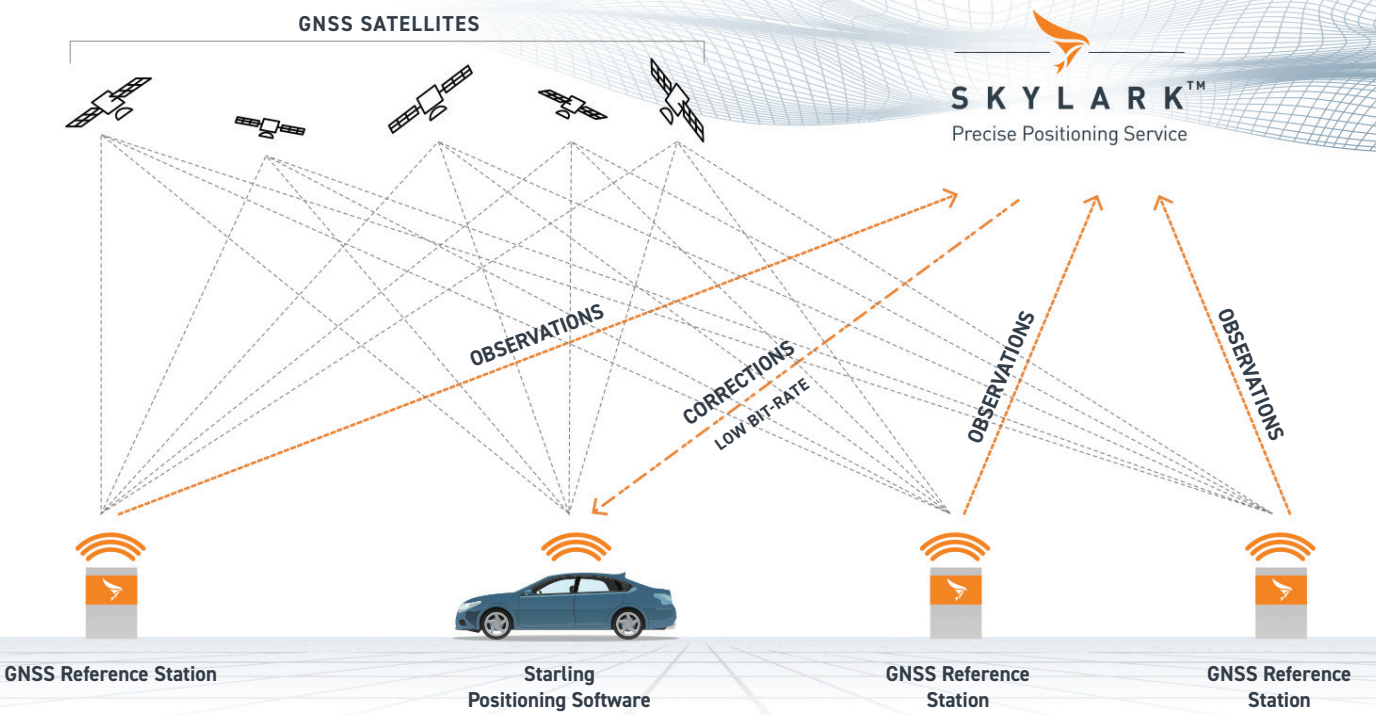

Swift's Skylark Checks for the Following:

Satellite Issues

- Satellite clock instabilities
- Signal anomalies
- Unscheduled maneuver
- Satellite inter-frequency bias
- Satellite orientation issues
- Satellite antenna gain change

Atmospheric Issues

- Ionospheric irregularities (that could be caused by solar storms, eruptions or earthquakes)
- Scintillation

Ground Segment

- Ephemeris issues (e.g., erroneous orbit/clock)

Swift's Starling Considers:

- Multipath
- Jamming
- Spoofing
- Cycle slips
- Galileo Binary Offset Carrier (BOC) side peak tracking

Example of Rapid Ionosphere Changes

In January 2022, a volcano erupted near Tonga. This event sent ripples through Earth's ionosphere, disturbing the Total Electron Count (TEC). Free electrons delay radio signals, so when there are rapid changes in the electron count, there are corresponding rapid changes in the signal delay that impact positioning accuracy. These ripples have the potential to impact GNSS corrections if not observed and corrected.

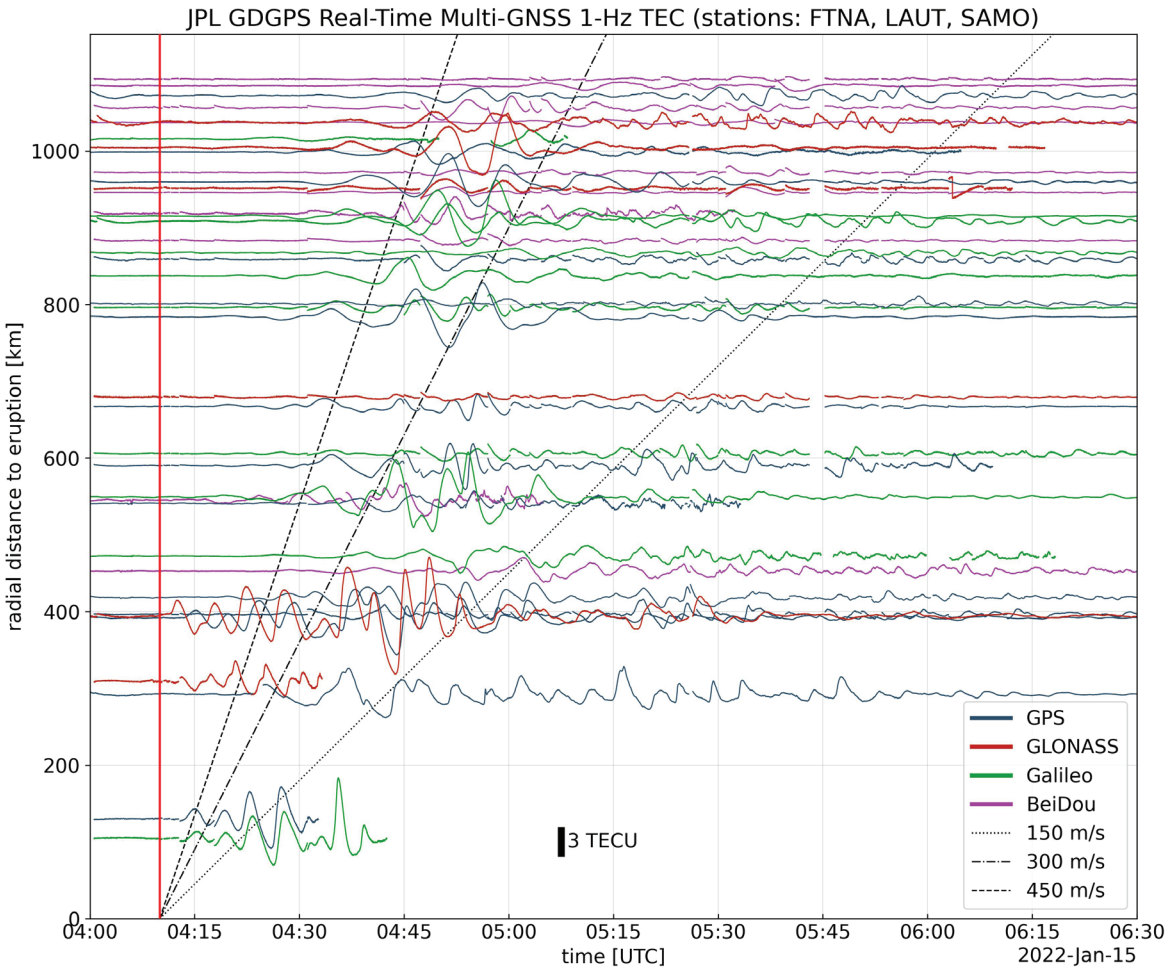

Credits: NASA/JPL-Caltech/GDGPS

The bar in the middle of the graph above indicates a 3 TECU amplitude. This equates to approximately 0.5m of delay, so some of these oscillations could contribute to a 1m variation in position accuracy. Swift has dimensioned its monitoring network to capture these events and other large TECU gradients in the ionosphere, reflected in the corrections sent to the positioning engine.

Integrity Needs Standards

With so many factors impacting automotive positioning, systems must utilize more sophisticated algorithms to meet integrity requirements. The probability of hardware and software faults also impact integrity risk. In the automotive world, these factors are handled by developing to the ISO 26262 standard, with the ASIL B(D) level being the typical requirement. ASIL B(D) is read as "ASIL B of D" which means that there will be one or more other ASIL B systems that combine to create an ASIL D—the most stringent safety case.

This paper has reviewed how integrity parameters are impacted by a corrections stream calculated in the cloud but the ISO 26262 standard specifically states: "ISO 26262 is the adaptation of IEC 61508 to comply with needs specific to the application sector of electrical and **electronic (E/E) systems within road vehicles**."[2] Because the cloud is not an electronic system within a road vehicle additional considerations need to be factored in.

Bringing Safety to the Cloud

There are currently no industry standards for cloud safety. Fortunately, many critical services are already running in the cloud, so best practices exist. One example is containerization, scalable services that spin up another container when loads increase or should one fail. Regional data centers handle geographic load balancing and failover should one center fail. These together allow a robust cloud architecture to be developed.

But that is not enough. There can still be systemic algorithmic errors and simple software bugs to consider. To counteract that, Swift developed its cloud software using ASPICE and ISO 26262 standards, so while ISO 26262 was not designed for the cloud, Swift follows the spirit of the standard until applicable standards catch up. Swift works closely with cloud providers to bring safety standards for automotive to the cloud.

Cybersecurity

It doesn't matter how robust you design your software—or the hardware it runs on—if someone can access or hack the data in it, it isn't safe. Cybersecurity is critical to safe positioning.

One of the many attack vectors to protect is spoofing. With hundreds of reference stations per continent, it would be very difficult for someone to spoof all of them. In addition, each station is equipped with dual antennas—that immediately detect a spoofing attack—and multiple redundant pathways to get the data back to the cloud. On the vehicle, spoofing is detected by using the fusion engine to detect, and reject, erroneous signals. Ephemeris spoofing is handled by using ephemerides sent in the corrections stream from Skylark, reading ephemerides from many sources and rejecting any that may have been spoofed.

All corrections are digitally signed with key rotation, so the GNSS client knows they can be trusted and have not been altered. Encryption can be layered on for even more security. Data security is bolstered with penetration testing and the following of industry best practices for cloud-based services.

At Swift, many standards are followed, including:

- ISO 20000 (Information technology — Service management)
- ISO 27001 (Data Security)
- ISO 21434 (Automotive Cybersecurity)
- ISO 26262 (Automotive Functional Safety)
- ISO 21448 (Safety of the intended functionality)

Proving Integrity

Proving that the probability of an integrity failure is less than 10^{-7} per hour can be difficult. Simply using drive tests would require billions of miles, and even then, one would have to determine if a satellite error occurred during that time and if that error was detected and transmitted to the positioning engine. A diversity of testing methods is as important as quantity when proving a system.

Swift has a multi-pronged approach to testing Integrity. Conducting **road tests**, we not only measure position errors but also capture the radio frequency (RF) signals from the antennas and a host of outputs from the receivers. Using the captured RF and data, Swift can **replay** these signals to racks of devices in a **hardware-in-the-loop test** situation. Swift has hundreds of receivers running thousands of automated tests 24 hours a day, 7 days a week, 365 days a year, easily scaling to millions of monthly test hours. **Cloud-based software-in-the-loop** testing can replay a four-hour test drive in less than two minutes.

As previously noted, one can drive for years and not encounter some specific satellite faults. How can one be confident in detecting and flagging such faults using only road testing? This obscurity is addressed using fault injection. With **fault injection**, replay, software-in-the-loop and cloud-based software-in-the loop testing of all combinations of scenarios, single feared events or any number of multiple feared events can occur delivering confidence that these can be handled if encountered.

Swift works with NASA's Jet Propulsion Labs (JPL) and researches the literature to capture any event that's ever gone wrong with GNSS and positioning sensors and incorporates it into our testing.

Summary

Swift has built an end-to-end integrity system from the ground up that will allow automotive and other safety-critical applications to trust the GNSS position and understand the bounds of the position error for that trust. The system starts with the reference stations that monitor the GNSS constellations. It flows through the Skylark cloud infrastructure that processes signals captured by the reference stations, builds a model of the GNSS error components with integrity error bounds and distributes these to the positioning engines. Swift's Starling positioning engine works with the GNSS receiver, taking in the integrity corrections stream and outputs position with protection levels for the TIR required for a specific positioning need. Starling furthermore integrates inertial measurement and vehicle motion sensors such as odometry to add robustness and availability to the position estimate.

References:

1 Sonekar, Saket and Puttagunta, Sravan (2021, September 18), Are HD Maps a bottleneck for self-driving cars? LinkedIn, https://www.linkedin.com/pulse/hd-maps-bottleneck-self-driving-cars-hyperspec-ai/?trk=organization-update-content_share-article

2 International Organization for Standardization, ISO 26262, Road vehicles—Functional safety, www.iso.org, https://www.iso.org/obp/ui#iso:std:iso:26262:-8:dis:ed-2:v1:en