Accuracy Assessment of Two GPS Fidelity Prediction Services in Urban Terrain

Andrew Moore
Dynamic Systems and Controls Branch
NASA Langley Research Center
Hampton, Virginia USA
andrew.j.moore@nasa.gov

Julian Gutierrez
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia USA

Evan Dill
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia USA

Michael Logan
Aeronautics Systems Engineering Branch
NASA Langley Research Center
Hampton, Virginia USA

J. Sloan Glover
Analytical Mechanics Associates
Hampton, Virginia USA

Steven Young
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia USA

Nathan Hoege
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia USA

***Abstract*—Low altitude flight in urban areas is susceptible to degraded GNSS-based navigation system performance due to terrain interference with radio signals from orbital positioning satellites. Predictive navigation performance fidelity tools are needed a) in preflight planning to assist in the creation of safe flight paths and b) in-flight to provide contingency management agents with the navigation risk of proximal flight corridors.**

Two navigation fidelity prediction services are validated by comparison with over 6000 readings from GNSS sensors collected along a five-mile path through urban areas of Corpus Christi, Texas, on three dates in 2022. Predictions are based on satellite line of sight through 3D terrain data collected in 2018. Each service predicts a set of navigation fidelity metrics over a user-specified time period. One metric estimated by both is the number of visible satellites. A direct comparison of the number of predicted visible satellites with the number sensed by the receiver is used to validate the prediction services. Results show an exact match in the number of predicted satellites for 60% of the measurements, and a match within +/- 4 satellites for 95% of the measurements. As expected, agreement improves away from vertical blocking terrain. Most cases of mismatch are due a lower predicted count than measured (false negatives), and can be accounted for by receiver pickup of stray signals caused by multipath propagation. About 10% of mismatches are false positives and are mostly accounted for by foliage effects. The two services predict visibility of the same set of satellites 80% of the time, differ by two or less satellites 95% of the time, and can compute predictions for one hour of observations in one minute or less. Validation is analyzed statistically and in detailed case studies of selected observation times.

The prediction services validated in this study run fast enough for preflight safety planning. The more stringent challenge of in-flight navigation fidelity prediction for contingency management requires both a speedup of the current level of modeling and equally fast stray signal modeling.

Keywords—GNSS, aviation, UAV, UAM, navigation

I. Introduction

Low altitude flight in urban areas is susceptible to degraded GNSS-based navigation system performance due to terrain interference with radio signals from orbital positioning satellites. Interference effects can range from simple blockage to complex reflection, diffraction, and attenuation.

NASA's System Wide Safety Project is developing a range of safety assurance services to support emerging low altitude urban flight operations. This involves identifying risks which information services may delineate (or mitigate), developing service prototypes, integrating them, and testing them in flight experiments. This R&D is intended to advance the concept of In-time Aviation Safety Management Systems (IASMS) [1]. One identified service type predicts navigation performance fidelity for use a) during preflight planning to assist in the creation of safe flight paths and b) in-flight to provide contingency management agents with the navigation risk of proximal flight corridors. Such a service should combine satellite orbital position and a digital surface model of the flight volume to determine satellite visibility along candidate flight paths and thereby estimate position quality of GPS-based navigation.

Two navigation fidelity prediction tools are tested in this study [2][3]. Their predictions are compared with readings from two GNSS sensors collected in Corpus Christi, Texas, on three dates in 2022. There is broad agreement between the predictions and measurements, and between the two prediction methods.

II. Objectives

This detailed comparison of sensor readings with predicted navigation fidelity advances three research objectives:

a. demonstrate the feasibility of rapid low-altitude navigation fidelity prediction,

b. uncover limitations to high fidelity prediction arising from the mapping, computational, and measurement methods employed, and
c. identify technology improvements required to improve the accuracy of navigation fidelity prediction.

Validation of two independently developed navigation fidelity predictors on common data allows for cross-validation between the two prediction services. Rationalization of any systematic differences between the two prediction methodologies was also an objective of this study.

I. Methodology

A. Prediction service.

Two navigation prediction services were evaluated, the Geometric Assessment of Positioning Systems (GAPS) [2] and the Corridor Assessment of Positioning Systems (CAPS) [3]. These services take different approaches to reduce the massive volume of 3D terrain data. GAPS projects from the satellite position to all locations within a volume and computes the shadows cast by 3D terrain features interposed between each satellite and each location (Fig. 1, left). In contrast, CAPS organizes 3D lidar data of a flight range into a spatial octree, computes rays from a receiver location to each satellite position, and finds intersections of each ray with voxels of the octree (Fig 1, right).

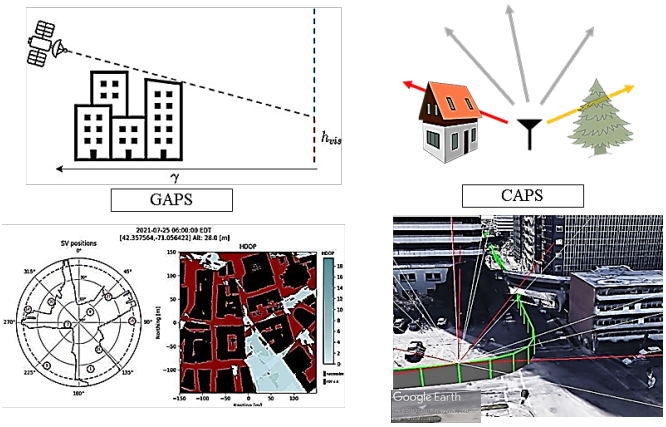

Fig. 1. Prediction methods under test. The GAPS prediction service computes a set of visible satellites by casting shadows (top left) of the terrain along all azimuth directions γ to find satellites above the minimum visibility height h_{vis}. The result (bottom left) is a visibility skyplot with overlaid cityscape and a map of horizontal dilution of precision (HDOP, blue scale bar). In the map, buildings are colored black and areas with less than 4 visible satellites are colored red. Regions of). Right: The CAPS prediction service casts rays to satellites and computes the total path obstruction along each ray (top right). CAPS results at one time sample of the June 6, 2022 test (bottom right) include unobstructed rays (grey) and rays obstructed by buildings (red). ©Graphics: NASA, NOAA, USDA; @Map data: Google, SIO, US Navy, NGA, GEBCO, Landsat/Copernicus

Computing hardware for the two prediction services was as follows. GAPS was executed on a machine with two Intel Xeon Silver 4216 CPUs, each with 16 cores/32 threads (hyperthreading enabled), 128 GB of RAM and an A6000 NVIDIA™,[1] GPU. CAPS was executed on a machine with two 8-core 2.6 GHz Xeon E5-2670 CPUs (32 total threads) and 128GB of RAM. Python code performed high-level setup and some low-intensity calculations (e.g. satellite propagation) for both services. GAPS' geometry calculations were computed with a combination of C++ and CUDA code. CAPS' geometry calculations were accomplished with custom C++ raycasting code and octree functions from Point Cloud Library [4] (version 1.7.2) .

Systematic differences between the two prediction methodologies were minimized by verifying spatial registration of terrain data inputs and by using identical coordinate systems for GNSS satellite orbits (Earth Centered Earth Fixed, or ECEF) and terrain data (World Geodetic System 1984, or WGS84).

B. Terrain data

The two prediction services used a common terrain survey, an aerial lidar scan conducted for the United States Geological Survey (USGS) in 2018 with an average point spacing of 0.6 meters [5]. The CAPS engine operates on point clouds; to improve point cloud sampling of the sides of buildings, additional points were added to the raw point cloud at 0.5m intervals between ground and the maximum survey altitude at each location. The GAPS engine operates on Digital Surface Models (DSMs) representing the maximum elevation captured by the LIDAR survey for each 1x1 m^2 tile. The USGS-generated geotiff files containing the DSMs were used with no modification. Spatial registration of the point cloud and tiff representations were visually verified in ArcGIS Pro (Figure 2, right). To allow direct comparison of the two prediction methods, a common grid size of one square meter (latitude by longitude) was used.

C. Data collection

Real-time geolocation measurements were conducted on three days at a different time each day to assure that the geometry of the orbital constellation varied between measurements. The measurement receivers were mounted on the roof of a vehicle (Fig. 3) and driven over approximately 5-mile paths that started and ended at a low-density, unobstructed area near the bayfront and coursed into the dense urban canyons of the Corpus Christi highlands (Fig. 2, left). The collection time was approximately one hour on the first two days and approximately 30 minutes on the third day. Urban terrain obstructed direct line of sight to a significant fraction of GPS Space Vehicles (SVs) for approximately 30 minutes on each measurement day.

D. Raw position data

Two independent receivers recorded position and associated satellite signals at one second intervals while the host platform traversed downtown Corpus Christi, TX:

1) a u-Blox Zed-F9P receiver configured to receive L1 and L2 band signals from the US GPS and Russian Glonass constellations.
2) a ruggedized VN-200 antenna configured to receive L1 signals from the US GPS constellation and produce a

[1] The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

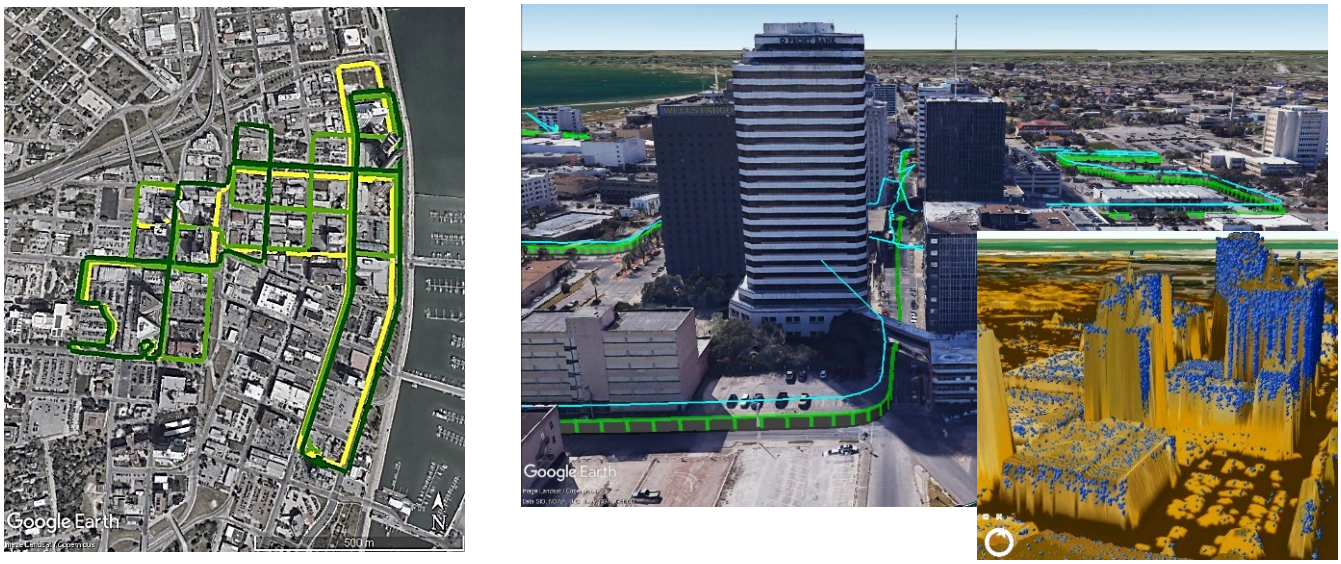

Fig. 2. Maps and geodata for Corpus Christi, Texas, used for computations. Left: Measurement path through flood plain and highlands on three dates in 2022 (June 6, dark green; July 28, middle green; August 27, yellow). Right background: View south along Carancahua Street showing the ground truth derived from the local elevation (green) and the position computed from receiver data on June 6 (cyan). As the cyan computed position trace indicates, navigation is severely degraded in the Carancahua Street urban canyon. Right foreground: Registration of lidar (blue dots) and TIF (brown hues) terrain representations used by the two services under test. @Map data: Esri/ArcGIS and Google, SIO, US Navy, NGA, GEBCO, Landsat/Copernicus

tightly integrated GPS/Inertial Measurement Unit (IMU) solution.

Readings from receiver 2 were used to verify that receiver 1 tracked ground truth position.

E. Reference position and valid satellite list

Because the receivers are operating in a GPS-degraded environment for considerable intervals, the ground truth measurement location is inherently uncertain at times. Inspection of the ground tracks from these receivers with various post-processing treatments showed that several seconds of temporal smoothing of receiver 1 output via the u-Blox u-Center utility most closely tracked the actual lateral position. Vertical position error in this derived track was severe at some locations, and so the reference altitude was set to 2 meters above ground.

Receiver 1 output was also post-processed with RTKLIB (version 2.4.3/demo5/b31 [6]) to provide, at each one second measurement time,

a) a list of satellites with detectable signals,
b) the subset of SVs with sufficient signal validity to be useful for position determination (signal observed within one second and no detected cycle slips), and
c) receiver geolocation estimated from the instantaneous signals.

The cyan trace in Fig 2. (right background image) illustrates the instantaneous computed geolocation on one measurement date. At locations away from obstructing terrain, the computed position smoothly tracks the ground truth position (green trace), but in the urban canyon along Carancahua Street, it deviates severely from the ground truth geolocation.

Satellites below a typical navigational mask angle elevation of 15 degrees were not considered. Each timestamped list of valid satellites at 15 degrees or higher was retained for comparison with predictions.

II. Results

A. Broad agreement between prediction and measurement

There is high correspondence between the predicted count of unobstructed satellites and the measured count of valid satellites. Viewed as a time series (Fig. 4), the correlation is remarkably good, as the predicted and measured valid satellite counts rise and falls in near synchrony. Time is represented as distance driven in this plot (to condense intervals when the vehicle was stopped at traffic lights). The top series shows the predicted and observed SV counts for both prediction services for both US GPS and GLONASS constellations while the middle series shows these measures for the US GPS constellation only on June 6, 2022. (Note: most of the remaining results presented are for US GPS constellation only, to reduce clutter.)

The bottom in Fig. 4 shows the topographic elevation (black trace) and ground truth elevation used by the two prediction services (blue and green traces). A local indicator of building

Fig. 3. Receiver mounts on the test vehicle. The measurement receiver is on the left. A UAV body with a navigational receiver is on the right.

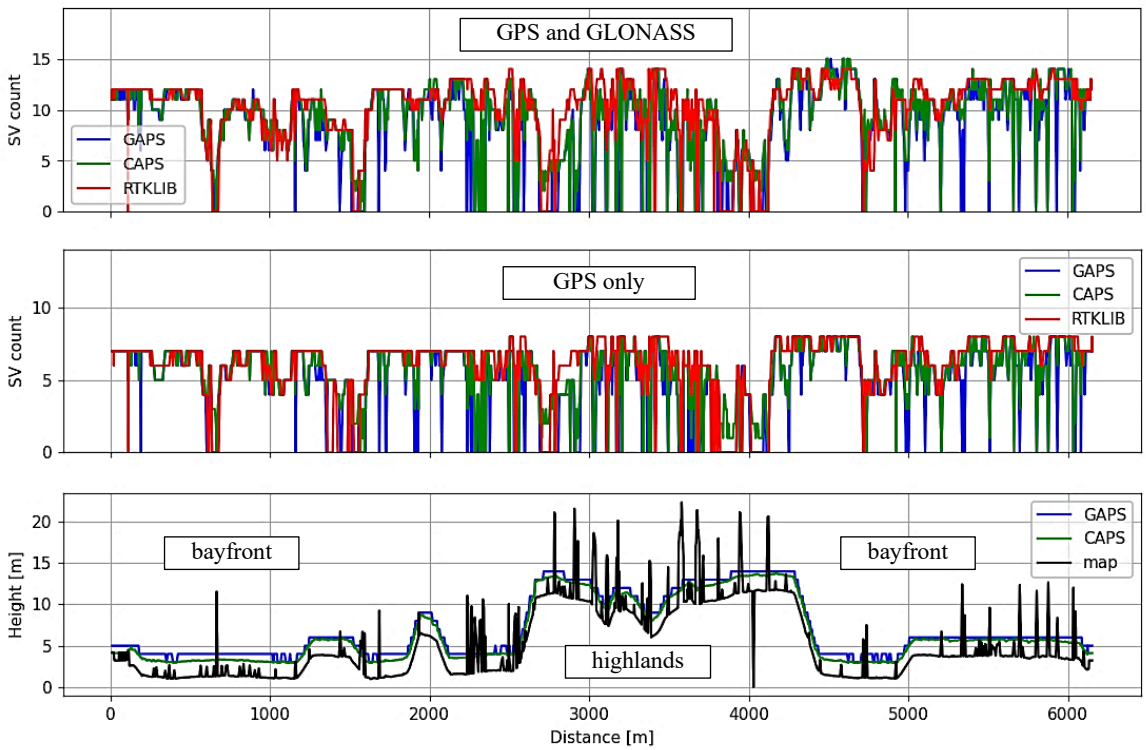

Fig. 4. Sequence of predicted and observed SV counts (top: US GPS and GLONASS, middle: US GPS) and topographic elevation (bottom) for the June 6 collection date. Time is represented as distance along the path through Corpus Christi. The delta elevation between the map and the CAPS/GAPS ground altitude is the vehicle height (2m). Degraded navigation is severe in the urban canyons of the highlands, traversed in the middle portion of this series.

height is also plotted in the black trace. Derived as a simple average of terrain height across the 3x3 m^2 area centered on the prediction location, this coarse neighborhood height indicator rises above ground most frequently in the urban canyons of the Corpus Christi highlands.

Navigation quality is generally good on the bayfront flood plain (beginning and ending third of recording) but navigational dropouts are common in the highland urban canyons (middle third of the recording).

B. Correlation of SV counts

Observed and predicted SV counts for all measurement times and all measurement dates are compared in Fig. 5 as confusion matrices (top) and histograms (bottom). In the confusion matrix plot format, exact matches are counted in diagonal cells. In plots comparing the observed count versus the predicted count (left: GAPS; middle: CAPS) of Fig. 5, nonzero counts below the diagonal represent false negative predictions (more satellites were observed than predicted) while nonzero counts above the diagonal represent false positive predictions (more SVs were predicted than observed).

The asymmetric prevalence of false negatives is thought to be due to radiofrequency effects such as attenuation, diffraction, and multipath signal reception (see *Discussion* below). Both predictors are conservative, as any obstruction along an infinitely thin ray to the SV is considered a complete blockage and there is no modelling of signal reflections. In contrast, the receiver can detect signals from SVs that are not in line of sight, as gigahertz-range radio waves diffract around the edges of blocking terrain and reflect from the ground and nearby structures. The measurement receiver duly reported all detected SV IDs, while the prediction methods only reported possible direct line of sight signals.

The two prediction services under evaluation report highly correlated SV counts (Fig. 5, right) with little asymmetry in the confusion plot (top) and histogram (bottom). This result indicates that there are only minor systematic differences between two prediction methods.

C. Prediction of visible SV IDs

The coarse validation provided by comparing SV counts at each time with the observed count can be further refined by considering the specific SV IDs that comprise the count. Prediction success is illustrated in two ways, as a histogram (Fig. 6) and as a mismatch count (Fig. 7). A ranking of the prevalence of predictions according to the number of mismatched SV IDs shows that a perfect match with observation list is most frequent (Fig. 6, top). Both the GAPS and CAPS predictions match observation exactly 60% of the time. Instances of mismatch by one satellite occurs 18% of the time, instances of mismatch by two satellites occurs 10% of the time, and instances of greater mismatch occur with a prevalence that monotonically decreases with increasing mismatch count.

The list of visible SVs predicted by the two services match more than 80% of the time. Instances of mismatch prevalence decline monotonically with mismatch magnitude. As for the SV count results in Fig. 4 and Fig. 5, this test shows close agreement between the GAPS and CAPS predictions.

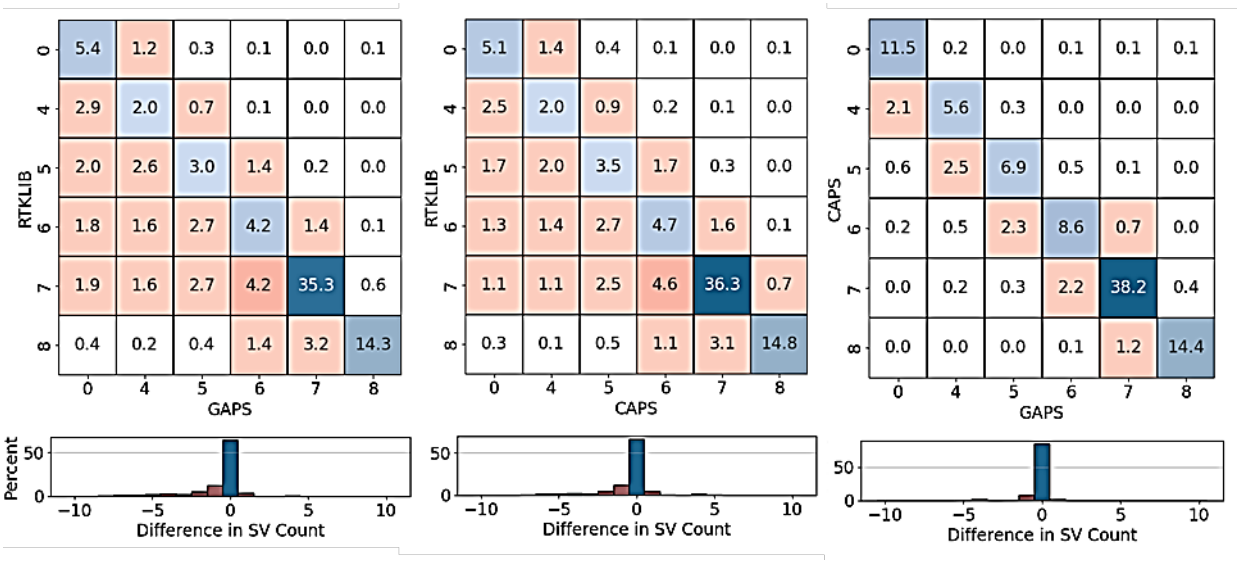

Fig. 5. Confusion matrix (top) and histogram (bottom) representation of match percentage between the predicted vs. the observed SV counts and between prediction methods across all dates. Each confusion matrix cell is labeled with its percentage prevalence. Left: GAPS vs. RTKLIB. Middle: CAPS vs. RTKLIB. Right: GAPS vs. CAPS.

Analysis of mismatch across location and elevation can reveal some of the underlying causes of differences between prediction and observation. Fig. 7 combines analysis along those two dimensions in a single plot for the GAPS results, with projections along the time axis at top and along the elevation axis at right. (CAPS results, not shown, are very similar). The central plot shows the change in SV elevation with time on the three validation dates for all 6175 experimental times. Different colors are assigned to individual SVs and their ID and instance count is inset at top right.

Experimental dates and times were chosen with an aim of validating navigation quality prediction with no repetition of constellation geometries; the diversity of SV IDs and elevations across dates evident in Fig. 7 shows that this goal was achieved.

Projection along elevation (Fig. 7, right) shows that mismatch magnitude is skewed toward low-elevation SVs. This is the expected result for urban canyons: signals from SVs directly overhead would not be obstructed by terrain; signals from SVs at high elevation are less likely to be obstructed but may be diffracted and reflected from the sides of buildings; and signals from SVs at low elevation are most likely to be obstructed, diffracted, and reflected.

Projection along time (Fig. 7, top) reveals instances amenable to individual analysis. Three such instances were chosen near the time indicated by the block arrow labeled '*' on July 28 and are shown below: one false positive instance with high mismatch (Fig. 8, foreground), one instance of exact match with few obstructions (Fig. 8, background), and one instance of exact match with many obstructions (Fig. 9).

D. Detailed validation of three cases

The false positive result with the highest mismatch amongst those plotted at the top of Fig. 7 is shown in the foreground of Fig. 8. At this location on Lipan Street adjacent to the Corpus Christ Cathedral, the canopies of tall trees form a wall of foliage to the north and overhang the street. At this experimental time, only one SV was visible directly south (gold ray of CAPS image) while the remaining SVs to the north and east were obstructed by foliage (red rays of CAPS image, dark dot of GAPS image). Preliminary results [3] show that the canopy depth should attenuate but not block signal reception, but both prediction services conservatively model any obstruction as completely blocking.

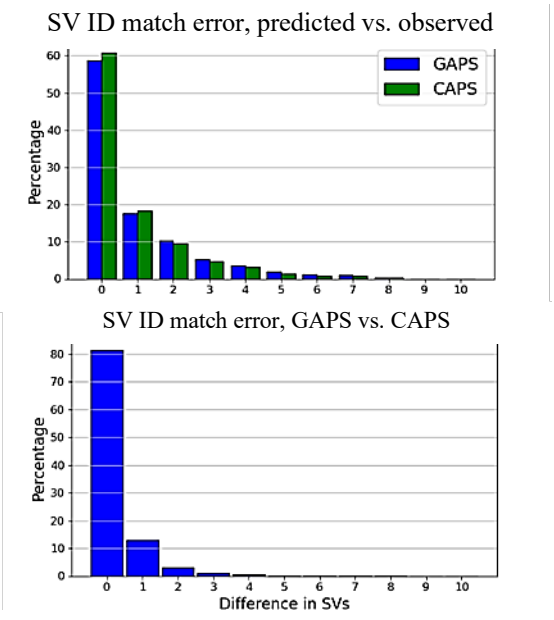

Fig. 6. Top: Relative frequency of the magnitude of SV ID mismatch, all observations on all days. The predicted constellation (GAPS, blue bars and CAPS, green bars) matches observations exactly 60% of the time and differ by four or fewer satellites 95% of the time. Bottom: The two services under evaluation predict the same constellation 80% of the time and differ by two or less SV IDs 95% of the time. This measure is more sensitive than comparison by SV count but shows similar characteristics: broad concurrence between predicted and observed satellites and close agreement between prediction methods.

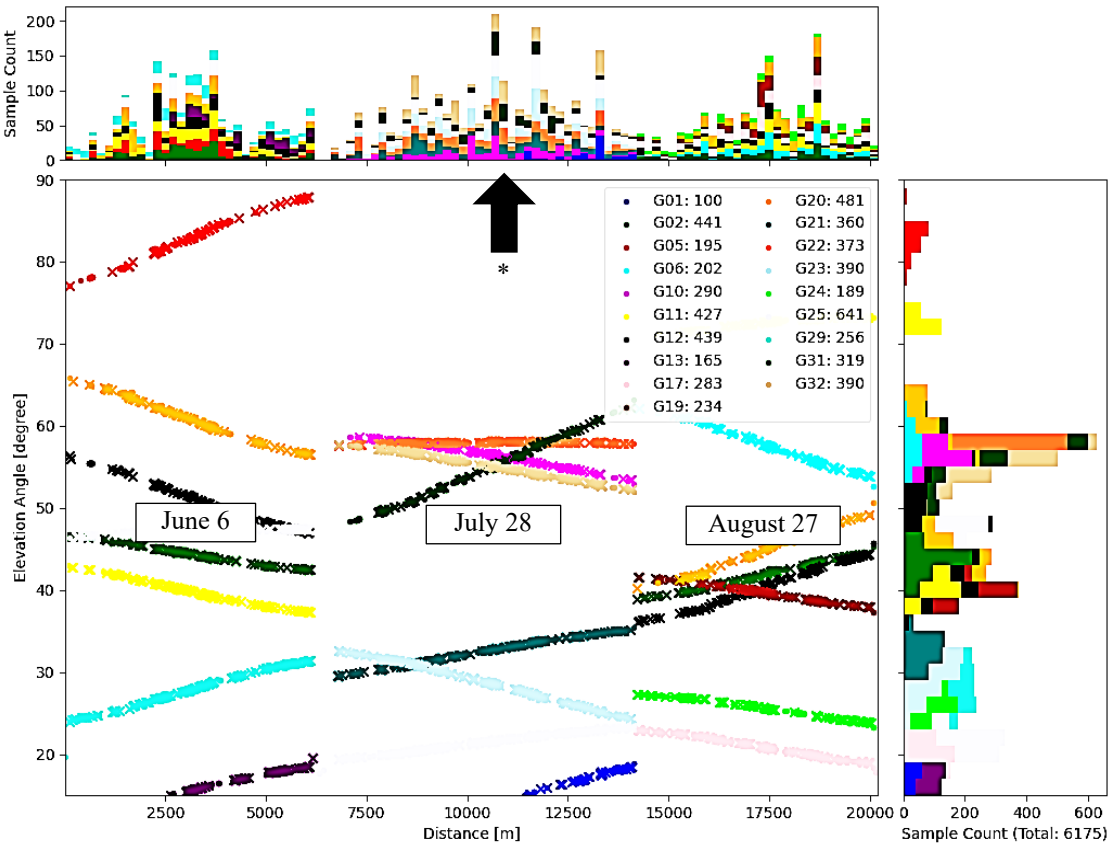

Fig. 7. Satellite elevation above horizon of instances of observed SV IDs that were not predicted by GAPS for all three measurement dates. Time is represented as distance driven on the x-axis of the central plot, to collapse time spent at stop lights. Mismatches between prediction and observation arise from non-line-of sight (NLOS) satellite signals detected by the receiver. Projection of the central plot by elevation angle (right) shows that stray signal pickup is skewed toward low elevations while projection by sample sequence (top) shows that high mismatch counts are clustered to certain map locations. Three cases of prediction performance on July 28 (near time with block arrow labeled *) are illustrated in Fig. 8 and Fig. 9 below. The corresponding plot for CAPs is nearly identical.

An instance of an exact SV ID match in an unobstructed location measured one minute later is shown in the background of Fig. 8. All lines of sight to SVs above 15 degrees elevation are clear (gold ray of CAPS image, light dot of GAPS image).

Lastly, verified predictions of complete navigation loss deep in an urban canyon is shown in Fig. 9. The sensor reported less than 4 valid satellites at this location (dark dot indicated by arrow in 'RTKLIB' image, top left). The skyplot (top right) produced by the GAPS service includes the bounds of obstruction (blue trace) and shows that the narrow slot of visible sky contains only three SVs above 15 degrees of elevation. The corresponding GAPS SV visibility prediction (bottom left) matches the list of observed SVs exactly, as does the CAPS prediction (bottom right).

III. Discussion

A. *Limitations to navigation fidelity prediction*

The following primary limitations to high fidelity navigation quality prediction with these methods were identified.

- **Terrain map currency.** The lidar survey used in this validation study is three years old, and the urban landscape of Corpus Christi is somewhat changed due to construction and demolition of buildings.
- **Multipath.** Radiofrequency (RF) signals reflect from buildings and foliage, therefore, some satellites are sensed but invalid (as judged by RTKLIB). Additionally, diffraction/lensing around building edges occurs at GPS radio frequencies. Since the prediction methods compute direct line of sight (no reflection and no diffraction), they cannot be expected to fully agree with the measured signal set in highly degraded areas. However, we found that the intrinsic approximation of direct line of sight prediction is useful and models most of the navigation loss.
- **Computation time.** Both prediction engines are designed to perform rapidly to advance progress toward real-time navigation quality estimation. The execution time needed to produce predictions of navigation quality at all measurement times of one of the validation measurement days for this terrain dataset is less than one minute for both GAPS and CAPS. For GAPS, this required sampling SV orbits at ten second time intervals and using the most recent orbital point for a given prediction (worst-case angular error < 0.05 degrees). For CAPS, this required limiting predictions to points along the ground truth path. These execution times are adequate for preflight planning, but faster (3 seconds or less) execution time is needed for in-flight contingency management.

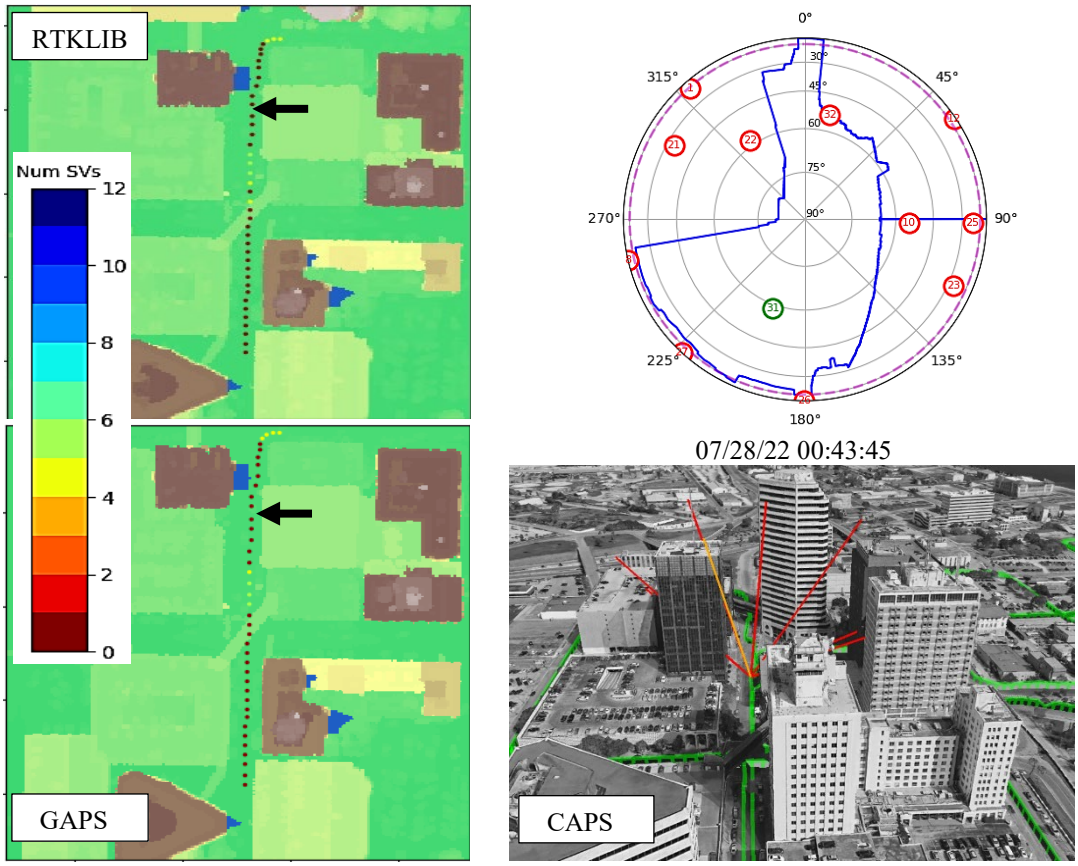

Fig. 9. Prediction of navigation loss in urban canyon from the July 28 experiment. As expected from the skyplot (top right), few satellites are available to the sensor at this location due to urban blockage, and post-processed sensor readings show complete loss of navigation (top left, arrow). GAPS (bottom left) and CAPS (bottom right) faithfully predict the set of visible and blocked satellites. The scale at left assigns color according to the numbers of SVs in the RTKLIB and GAPS images. @Map data: Google

E. Limitations of this validation method

The following primary limitations to the validation method were identified.

- **Ground truth reference (spatial).** The ground truth measurement location is inherently uncertain in highly degraded areas. While a reasonably well-behaved spatial ground truth was constructed from the measured receiver data, there were instances of lateral variation well outside the lane of the roadway. To minimize this limitation, ground truth data was manually inspected and the time intervals of such instances were excluded from the validation dataset.
- **Ground truth reference (satellite count).** The subset of valid satellites was obtained using the default validity tests of the post-processing tool (RTKLIB). The tools' default thresholds (e.g., signal strength, cycle slip parameters) may not accurately model the actual validity tests embedded in a particular GNSS receiver. Additionally, if RTKLIB cannot find the minimum number of satellites (four by default) needed for a position fix, it reports zero valid satellites. In contrast, the predictors will report the number of visible satellites even if it is less than four.
- **Foliage transparency.** In the few tree species studied to date [7], the foliage thickness that results in GNSS attenuation high enough to prevent detection by a receiver is 10m or more. This area of Corpus Christi is relatively free of foliage, and so treating foliage as completely opaque to GNSS signals was a known simplification of this validation study. The manually intensive [8] effort required to segregate foliage from buildings would reduce the mismatch frequency and could eliminate some false positives (such as the one shown in Fig. 8, in which a canopy much thinner than 10m is regarded as opaque).

F. Major technology gaps

Better terrain map currency and fast multipath modeling are needed to improve navigation fidelity prediction tools and to increase the accuracy of their validation.

Lidar surveys are slow and expensive and as a result are conducted infrequently. Typically, surveys of a given area are separated in time by 3 or more years. The satellite photogrammetry method of 3D terrain mapping is approaching the level of detail needed to perform physics-based GPS navigation prediction. In principle, satellite mapping could refresh terrain maps several times per year.

Extensive multipath computation that accounts for reflections and diffraction will improve prediction quality but will increase execution time. There is not currently a firm estimate of the added computational cost. The prediction services validated in this study, which compute direct line-of-sight only, differ from observations by four SVs or less 95% of

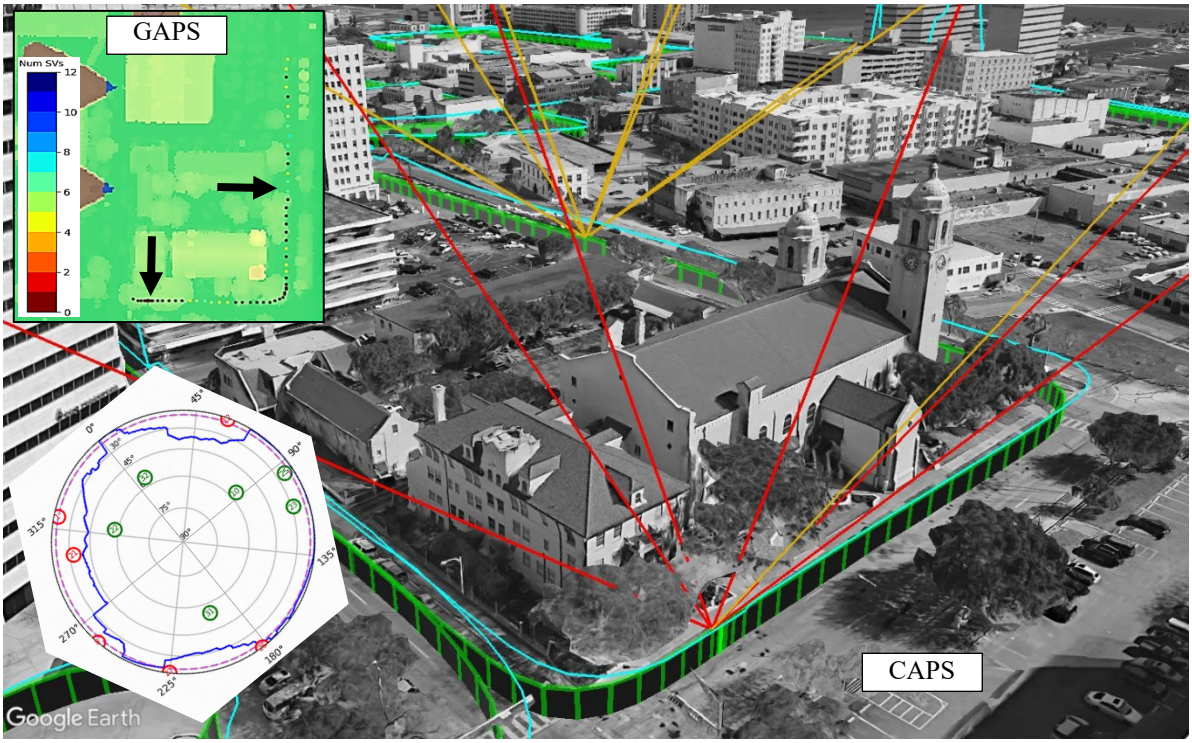

Fig. 8. Locations of a notable false negative (foreground) and an exact match (background) from the July 28 recording, near the time marked '*' in Fig. 7. The computed position from sensor readings (cyan trace) is stable on this segment of the measurement path, indicating good navigation quality. In the foreground location CAPS predicts visibility from only one visible SV (gold ray) with the rest blocked by terrain (red rays). This false negative arises from the tree canopy in the foreground location which is regarded as opaque. In the background location, the set of predicted and measured SV IDs match exactly. Top inset: congruent GAPS results at these locations (black arrows). Bottom inset: visibility skyplot computed by GAPS and oriented to align with North in the background image. @Map data: Google

the time. Even partial approximations to the true complexity of multipath GNSS signal propagation (e.g., single bounce raytracing) could materially narrow the difference between prediction and measurement. Predictions computed in three seconds or less with an accuracy of two SVs or less 95% of the time are desirable for low altitude navigation.

IV. Conclusions and Significance

To enable anticipated demand for new types of flight operations at low altitudes and in urban areas, new safety assurance challenges must be addressed proactively. One such challenge is associated with navigation system performance. Requirements have yet to be defined for some of these operations; however, GNSS is likely to play a key role. Predicting performance using methods such as described here can (a) aid in planning (e.g. flight profiles and scheduling, takeoff/landing site selection, and identifying needs for augmentation infrastructure); and (b) support in-flight risk mitigation (e.g., via trajectory prediction, contingency management and execution).

There is a small but growing number of research efforts to predict low-altitude navigation fidelity with realistic terrain models (see [2] for a review of this literature). To our knowledge, none have been validated to this level of rigor. To serve as safety or mission-critical services in the emerging infrastructure required for urban air mobility and low-altitude unmanned flight operations, exercises such as this are required to uncover technology gaps and overcome them. While some methodological limitations were identified in this study, none seem insurmountable, and agreement between prediction and measurement was remarkably close.

In urban canyons with degraded signals from positioning satellites [9][10], aviators and public safety officials need to know where and when flights can be conducted reliably and safely. The prediction services validated in this study run fast enough to address this preflight safety planning need and their incidence of false positive predictions is negligible. Their incidence of false negative predictions is unacceptably high, however; better modeling of stray signal reception is needed to narrow the gap between prediction and observation. The more stringent challenge of in-flight navigation fidelity prediction for contingency management requires a speedup of line-of-sight modeling — from tens of seconds to three seconds or less — and equally fast stray signal modeling.

V. Acknowledgements

We are grateful to Nick Rymer of the NASA Langley Research Center for assistance on receiver setup and Kyle Ellis and Misty Davies for program support. N. Hoege acknowledges the University Space Research Association for internship logistics. This work was performed with support from the NASA System Wide Safety Project, which is part of NASA's Aviation Operations and Safety Program.

VI. REFERENCES

[1] S. Young, E. Ancel, A. Moore, E. Dill, C. Quach, J. Foster, K. Darafsheh, K. Smalling, S. Vazquez, E. Evans et al., "Architecture and information requirements to assess and predict flight safety risks during highly autonomous urban flight operations," NASA Tech Memo NASA-TM-2020-220440. 2020.

[2] E. Dill, J. Gutierrez, S. Young, A. Moore, A. Sholz, E. Bates, K. Schmitt, J. Doughty, "A predictive GNSS Performance Monitor for Autonomous Air Vehicles in Urban Environments", Institute of Navigation (ION) Global Navigation Satellite Systems (GNSS+) Conference, September 2021.

[3] Moore, A., Schubert, M., Rymer, N., Villalobos, D., Glover, J., Ozturk, D., & Dill, E. (2022), Volume raycasting of GNSS signals through ground structure lidar for UAV navigational guidance and safety estimation. AIAA Scitech 2022 Forum. 2022.

[4] Rusu, Radu Bogdan, and Steve Cousins, "3d is here: Point cloud library (pcl)." In 2011 IEEE international conference on robotics and automation, pp. 1-4. IEEE, 2011.

[5] United States Geological Survey, "South Texas FEMA 2017 LiDAR Project Report," https://coast.noaa.gov/htdata/lidar1_z/geoid18/data/8841/supplemental/31013_TX_FEMA_R6_ProjectReport.pdf

[6] Everett, T., rtklibexplorer, RTKLIB Demo5, 2021 URL: https://github.com/rinex20/RTKLIB-demo5.

[7] A. Moore, N. Rymer, and J.S. Glover, "Predicting GPS Fidelity in Heavily Forested Areas." 2023 IEEE/ION Position, Location and Navigation Symposium (PLANS), Monterey, CA, 2023.

[8] A. Moore, M. Schubert, T. Fang, J. Smith, and N. Rymer, Bounding Methods for Heterogeneous Lidar-derived Navigational Geofences. NASA Tech Memo NASA-TM-2019-22399. 2019.

[9] Soloviev, A., Bruckner, D., van Graas, F. and Marti, L., 2007, January. Assessment of GPS signal quality in urban environments using deeply integrated GPS/IMU. In Proceedings of the 2007 National Technical Meeting of the Institute of Navigation (pp. 815-828).

[10] Kumar, R. and Petovello, M.G., 2014, September. A novel GNSS positioning technique for improved accuracy in urban canyon scenarios using 3D city model. In Proceedings of the 27th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2014) (pp. 2139-2148).