Development, Implementation, and Testing of a Prototype LAAS Architecture

Boris S. Pervan, Samuel P. Pullen, David G. Lawrence, Konstantin Gromov,
Jock Christie, Guttorm Opshaug, Vicky Lu, Ping-Ya Ko, Per Enge,
and Bradford W. Parkinson

Department of Aeronautics and Astronautics, Stanford University
HEPL/Gravity Probe B, Stanford, CA, 94305-4085, USA

Abstract: A Local Area Augmentation System (LAAS) architecture has been developed to provide satellite navigation performance compliant with the stringent requirements for aircraft precision approach and landing. Code and carrier phase measurements from ground-based airport pseudolites (APLs) located at each end of the approach runway are used to improve vertical performance and thus increase navigation availability. In addition, a comprehensive monitoring architecture is introduced to ensure navigation integrity. Detection and deletion of ground failures is performed both on the ground and at the aircraft to ensure the high levels of integrity and availability required.

To verify the performance of the proposed architecture, a series of FAA-sponsored flight trials will be performed on a Boeing 727 at Moffett Field, California in May 1997. APL signal quality will be compared using a fuselage-top-mounted antenna (also used for satellite ranging) and a forward-looking nose-mounted antenna (located within the 727 radome). The ground and air components of the prototype architecture implemented for these tests are described in detail.

1.0 INTRODUCTION

The Local Area Augmentation System (LAAS) is a ground-based differential satellite navigation system to be implemented by the Federal Aviation Administration (FAA) to provide the means for aircraft precision approach using satellite navigation. LAAS has two primary goals. The first is to provide Category I service for those airports that are not covered by the FAA's Wide Area Augmentation System (WAAS). The second is to provide Category II and Category III performance where required [1]. The requirements for the LAAS Signal-in-Space (SIS), as documented in the LAAS Operational Requirements Document [2] and subsequently modified by the FAA and Boeing Commercial Airplane Group [3], are based on derived Instrument Landing System specifications [4].

At Stanford University, an ongoing effort is focused on the research, development, implementation, and testing of LAAS architectures and architecture components. Previously, the Integrity Beacon Landing System (IBLS) was conceived and developed at Stanford to provide centimeter-level accuracy with high integrity [5]. This system combined differential carrier phase measurements with low-power ground-based pseudolites placed under the aircraft approach path – nominally at the ILS Middle Marker site. In this application, the large geometry change resulting from pseudolite overflight provides the observability for carrier phase cycle ambiguity estimation. While the performance of IBLS is exceptional, concern existed over the existence of critical equipment off the airport property. The FAA thus requested that Stanford constrain its LAAS development activities to architectures wholly within airport boundaries. The architecture described in this paper is the response to this directive.

The significant elements of the proposed baseline LAAS architecture are shown in Figure 1. Multiple reference receiver sites are included to provide both the redundancy for reference receiver fault detection and removal and the means for reducing the net effect of low-frequency code multipath at the reference station. The former function is accomplished by a combination of ground-based and airborne integrity monitoring components. The latter is achieved by effectively averaging pseudorange measurements from widely separated reference receivers.

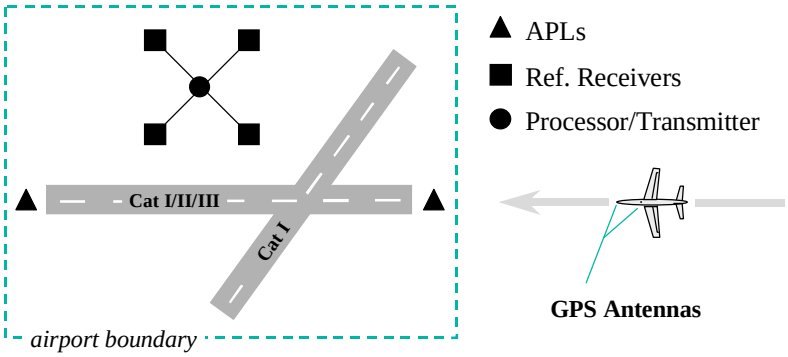

Figure 1: Baseline LAAS Architecture

Airport pseudolites (APLs) are introduced to augment LAAS availability by improving navigation accuracy, integrity, and continuity when the satellite geometry is deficient. Particular attention is placed on APL siting in the proposed architecture. This is especially true for Category III applications, where the 'intrack' configuration shown in Figure 1 is chosen to maximize vertical performance subject to the on-airport siting constraint. Because of the siting constraint, high-power (pulsed) pseudolites are used to ensure signal reception at the aircraft during the approach. A forward-looking APL antenna is present on the aircraft (in addition to a standard fuselage-top-mounted satellite antenna) to provide the capability for direct line-of-sight APL signal reception.

At the aircraft, code and carrier phase measurements are processed optimally to accommodate all available information. Of particular significance is the relative change in APL (and satellite) geometry during the course of the approach, which provides observability of the carrier cycle ambiguities. However, because of the limited geometry change available with APLs located within airport boundaries, a purely carrier phase solution (as achieved by IBLS) is not possible. For this reason, code phase measurements play an increased role. In addition, the airborne measurement processing system utilizes a new "multiple hypothesis" approach toward reference receiver measurement verification and selection.

This paper details the various aspects of LAAS architecture design and prototyping. Ground and airborne flight test implementations in preparation for flight trials on an FAA Boeing 727 and a NASA Beechcraft King Air are described.

2.0 ARCHITECTURE DESIGN

A summary of the architecture design is given below. Emphasis has been placed on the elements of the proposed LAAS architecture that are distinct from more traditional Differential Global Positioning System (DGPS) implementations. It is assumed that the only the GPS Standard Positioning Service (SPS) is guaranteed to be available. In addition, the selection of datalink broadcast frequency and modulation scheme will not be addressed in this paper. Where pertinent, it is assumed that LAAS will implement a VHF datalink compatible with Special Category I (SCAT I) installations.

2.1 Airborne Measurement Processing

In the proposed architecture, reference receiver code phase pseudorange (PR) and carrier phase (CP) measurements (or, equivalently, measurement corrections) for each visible satellite (SV) and APL are transmitted directly to the aircraft. At the aircraft, these measurements are processed as indicated in the functional block diagram given in Figure 2. It is assumed that the datalink bandwidth is available to broadcast reference CP and PR measurements at 2 Hz rate. Generally, the airborne GPS receiver may be sampled at much higher rates, although the significance of the measurements will be limited by the tracking loop bandwidth. In order to achieve regular time synchronization between reference and airborne receiver measurements, the airborne measurement rate must be a multiple of 2 Hz. In the baseline architecture, a rate of 4 Hz is assumed. Note also that due to datalink latency, it is necessary that airborne measurements be stored for future time alignment with broadcast reference measurements.

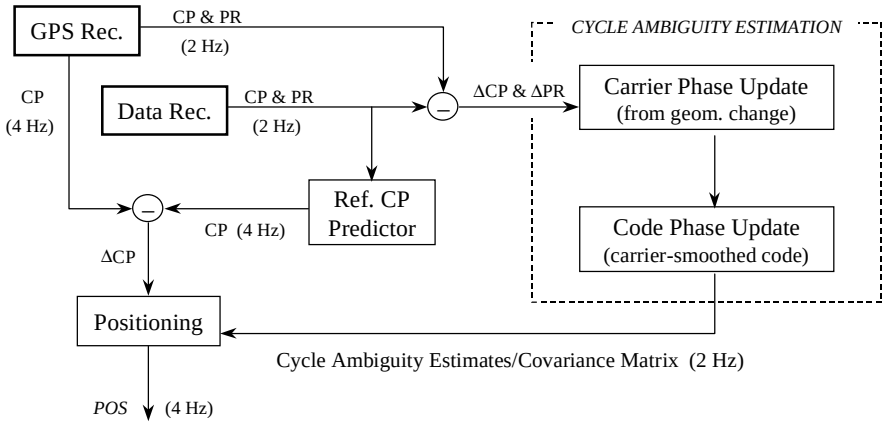

Figure 2: Nominal Measurement Processing

When a set of reference station measurements is received at the aircraft, single difference (aircraft minus reference) observables are formed for the carrier phase (ΔCP) and pseudorange (ΔPR). These time-delayed observables provide the measurement basis for real-time, sequential cycle ambiguity estimation. Two separate processes, indicated in Figure 2, are implemented in this regard:

Carrier Phase Update. Carrier measurements (ΔCP) are used directly for motion-based estimation of the cycle ambiguities. Although satellite motion and redundancy (required for cycle ambiguity observability) is limited during a typical approach, the addition of APLs helps ensure that minimum observability conditions are met. With specific APL ground configurations, the implementation of motion-based cycle ambiguity estimation leads to substantial improvements in vertical performance, and therefore, LAAS operational availability (see Section 2.3.3).

Code Phase Update. This implementation is simply a variation of carrier-smoothed code. The differences between ΔCP and ΔPR measurements are averaged over time to provide a source of carrier phase cycle ambiguity observability independent from the carrier phase update above. Since the predominant low-frequency components of reference station PR multipath error cannot be averaged out over the time scale of a typical approach, a bias state on ΔPR is maintained. In effect, a lower bound (typically 20-25 cm) is established on the standard deviation of cycle ambiguity estimate error using the code phase update.

The advantage of implementing carrier-smoothed code using single differences at the aircraft is that ionospheric divergence is irrelevant. As a result, longer smoothing durations may be used than if raw measurements were smoothed by the ground and aircraft separately. If pre-smoothed measurements must be transmitted to the aircraft (for compatibility with existing SCAT-I formats), the code phase update can also be implemented directly, without averaging, at the aircraft.

In order to generate timely (4 Hz) position fixes at the aircraft, stored reference CP measurements are projected to the current time using a quadratic fit (constant acceleration Selective Availability model) over the last seven measurements [6]. The resulting projected reference CP measurement is subtracted from the actual aircraft CP measurement to generate a projected ΔCP at 4 Hz. The current cycle ambiguity estimates are then subtracted from the projected ΔCP measurements, and a weighted least squares position fix is performed based on the sum of the cycle ambiguity error covariance matrix (an output of the cycle ambiguity estimation process) and the diagonal ΔCP measurement error covariance matrix. This processing architecture is an optimal implementation in the sense that all information available in GPS (SPS) and APL measurements is extracted. A detailed description of these algorithms is given in [6].

2.2 Integrity Monitoring

2.2.1 Monitoring Responsibility: The distributed nature of LAAS integrity threats among ground, space, and airborne suggests that it would be best to divide integrity monitoring between the ground and airborne segments to maximize fault observability [7]. However, it is desirable that LAAS fit into an "ILS-lookalike" format that places all of the responsibility for the integrity of the Signal-in-Space (SIS, encompassing ground and space segments) in the ground segment. This means that, in addition to verifying the integrity of the APLs and the datalink, the ground must detect all space segment failures that are not directly mitigated by DGPS (i.e., those viewed differently between ground and aircraft) and must detect and isolate (if possible) any reference receiver failures. This is difficult because not all of these are directly observable on the ground and because ground measurements may not be accurate enough to detect all errors that could be hazardous in the air. The aircraft must then assume that all received corrections have errors equal to the threshold with which they were screened on the ground, and the resulting likelihood of remaining inside the integrity alert limits is lower. Instead, if all measurements (or corrections) from each reference receiver are broadcast, the aircraft can use the actual measured differences to give the resulting error under the hypothesis that a given receiver had failed. This will always give improved availability, and the availability difference will increase with reference receiver noise. The aircraft can also attempt to achieve the required probability of HMI with a subset of the available satellites or ground measurements if the "all-in-view" calculation is unsatisfactory.

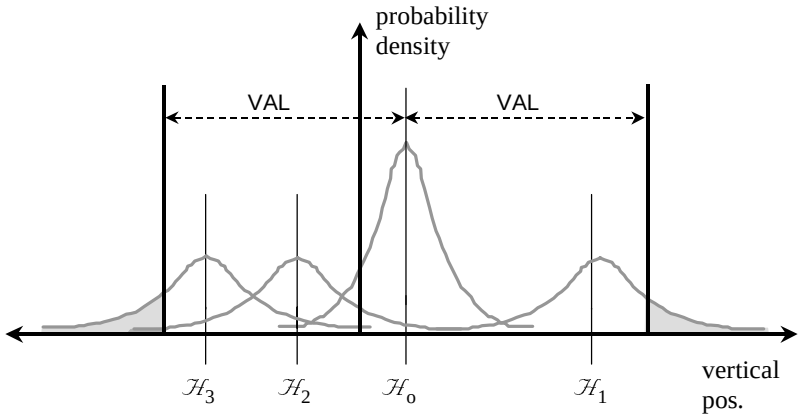

Figure 3: Multiple Failure Hypotheses

2.2.2 Airborne Multiple Hypothesis Approach: LAAS ground segment failures are mitigated at the aircraft through direct evaluation of integrity risk. In this context, integrity risk is defined as the likelihood that the position estimate exceeds a pre-specified alert limit. The computation of integrity risk considers all single reference receiver failure hypotheses and the no-failure hypothesis taken together. This approach for vertical position integrity monitoring is graphically illustrated in Figure 3 for the three-receiver case. The four Gaussian curves in the figure are aircraft position probability density functions (pdfs) corresponding to the three reference receiver failure hypotheses ($\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{H}_3$) and the no-failure hypothesis ($\mathcal{H}_0$). The mean of the $\mathcal{H}_0$ pdf is simply the least-squares vertical position estimate using data

from all three reference stations, while the mean of the $\mathcal{H}_1$ pdf is the least-squares vertical estimate using information from only reference receivers 2 and 3. The means of the $\mathcal{H}_2$ and $\mathcal{H}_3$ pdfs are obtained in an analogous manner. The prominence of the $\mathcal{H}_0$ pdf relative to the three failure pdfs signifies the relative prior probabilities of the associated hypotheses; in general the prior probability of failure is much smaller than one. In the case of LAAS, the prior probability of reference station failure at a given measurement epoch must be smaller than 10^{-5} to ensure that the likelihood of simultaneous multiple reference receiver failures is negligible with respect to the integrity risk requirement of 10^{-9} for Category III. The computed integrity risk is then simply the sum of the pdf 'tail' areas weighted by the prior probability of each hypothesis. More detail on this approach to LAAS integrity monitoring will be given in a forthcoming paper.

2.2.3 Ground Screening: As discussed previously, the ground segment is capable of detecting and removing satellite and reference receiver signals that exceed thresholds dictated by ground measurement noise. Ground screening eliminates the vast majority of hazardous errors before they are broadcast to the aircraft. Also, failures that require repair are flagged promptly to minimize the equipment unavailability that results. Finally, screening ensures that the prior probabilities assumed in the airborne HMI probability calculation are attained.

The first element of ground screening involves signal quality monitoring (SQM) of individual measurements in each reference receiver. This step weeds out code or carrier measurements with anomalously low C/N_o, large range rates and/or accelerations that may be due to erratic clocks, and large code corrections that suggest a satellite ephemeris error. The measurements are differenced with respect to a GPS signal generator to cancel out individual receiver clocks. They are then projected to a common 'virtual' reference point by removing the projection of the known baseline vector from this point to each actual antenna location along the satellite line-of-sight unit vector. The surviving n measurements for each satellite are transformed into an $n-1$ length parity vector whose norm provides a test statistic that is compared to a 7σ detection threshold. Once detection occurs, it may be possible to identify and remove the responsible reference receiver channel assuming that the probability of more than one reference receiver failing simultaneously is negligible [7]. If a failure is isolated on a particular satellite and reference receiver, the measurements for that satellite/receiver channel combination are not broadcast until the test statistic passes or for a specified minimum recovery time. If a failure is detected, but the probabilistic isolation condition is not met for any reference receiver, the satellite measurement is declared unusable on all reference stations. Finally, if the isolation condition is met on one reference station for more than one satellite during any recovery-time interval, all measurements from the affected receiver are declared unusable for an additional recovery interval. Finally, a maintenance alert function disables reference receivers that remain unusable for a substantial period (on the order of 5 minutes) until they can be checked and repaired if necessary (see [7] for details).

For the flight tests described in Section 4.0, ground screening is used as a diagnostic tool only. All measurements (including ones flagged as bad on the ground) are transmitted to the aircraft, which then performs detection and isolation in real-time. These results are compared to the ground screening outputs in post-processing to correlate any integrity flags that arise and to compare the sensitivity of screening on the ground and in the aircraft.

2.3 Airport Pseudolites

2.3.1 Unaugmented LAAS Availability. The availability of LAAS for precision approach and landing is primarily an economic consideration. For an approach to be available, the accuracy, continuity, and integrity requirements must all be met before the approach commences. *Service availability* is the long-term average that LAAS will be available for a given category of service at a given location. Service availability of current ILS approaches is generally in the range of 0.998 – 0.9999 over an airport and is bounded by the outage probability of the common ground equipment that supports all ILS approaches [8]. However, LAAS approaches are dependent on satellite geometry that is repeatable from day to day but is subject to satellite outages that could affect availability at several locations in a given region. *Operational availability* reflects this by examining the duration of the longest outage (period without available approaches) for a given number of healthy satellites (out of the 24 in the nominal constellation). Outages longer than 10-20 minutes are significant in terms of the inconvenience posed to air traffic control and to individual aircraft [9].

For operational reasons (e.g., issue of NOTAMs to departing aircraft), availability must be predicted well before it is actually determined at the start of an approach. GPS geometry simulation can predict (given the number of satellites predicted to be healthy) the satellite geometry available at any future time. Given this, the airborne multiple-hypothesis calculation is made assuming that the aircraft is at the decision height to determine if the probability-of-HMI condition will be met. Since protection limits and the probability of HMI scale directly with ranging error, the ability to limit errors under normal conditions is vital to obtaining high availability. Recent flight-test results analyzed in [10] provide an "optimistic" elevation-dependent range error model assumes improved ground multipath mitigation as a prerequisite for high availability. This model is used in the availability analyses to follow.

Simulations of GPS geometries with a selected number of unhealthy satellites are used to determine service and operational availability. These simulations cover six large U.S. airports with Category III ILSs: Los Angeles (LAX), Chicago (ORD), Dallas-Fort Worth (DFW), New York/Kennedy (JFK), Atlanta (ATL), and Seattle/Tacoma (SEA). In all cases, 3 operating reference receivers are assumed. For each user and combination of unhealthy satellites, one day of GPS geometries is generated at 5-minute intervals, and a predictive availability check is made for each geometry and user location. Representative vertical and lateral alert limits (VAL and LAL) for navigation integrity are given in [11]. For Category I, VAL is 7.6 meters and LAL is 10.5 meters. For Category III, these numbers tighten to 3.81 meters and 6.0 meters respectively. Overall service availability and maximum outage duration are computed for each location once all samples are complete. Note that each service availability result can be weighted by an assumed likelihood of having the assumed number of healthy satellites (see [12]) to get an overall service availability estimate.

Airport	CATEGORY I		CATEGORY III	
	Serv. Avail.	Max. Out. (min)	Serv. Avail.	Max. Out. (min)
LAX	0.9977	35	0.9905	85
ORD	0.9993	20	0.9884	65
DFW	0.9992	25	0.9871	50
JFK	0.9989	25	0.9902	60
ATL	0.9991	50	0.9870	90
SEA	0.9992	25	0.9871	55
Overall	0.9989	50	0.9884	90

Table 1: LAAS Availability for 22 Healthy SVs and no APLs

Without APL augmentation, the Boeing optimistic error model provides essentially no outages (service availability ~ 1) for either Cat. I or Cat. III when all 24 GPS satellites are healthy; thus the actual availability obtained will be dictated by ground equipment outages (it should be above 0.999, see [8]). For 23 healthy satellites, satellite service availability is better than 0.999 for Cat. I and 0.998 for Cat. III at all locations tested, and the longest outage is less than 25 minutes long. Table 1 contains the same results for 22 healthy GPS satellites. Not only does the achievable service availability drop significantly as more satellites become unhealthy, the longest outage duration jumps precipitously to 50 minutes for Cat. I and 90 minutes for Cat. III. This worst-case duration increases to 255 minutes for Cat. III in the unlikely but plausible 21-satellite case (service availability drops to 0.974). The use of APL's is intended to shorten these undesirable outages by making more satellite geometries available through improved accuracy performance.

2.3.2 APL Signal Format: Because of the significant outages in LAAS availability with the existing GPS constellation, augmentation by ranging measurements from airport pseudolites (APLs) is proposed. For an APL ranging signal to be useful for LAAS, APL signals strong enough to be tracked at the initiation of approach must not jam weaker satellite signals as the aircraft approaches the APL. This effect, known as the "near/far problem," may be mitigated by a number of APL signal design approaches, the most practical of which involves on/off pulse modulation of the APL signal [13]. Specific APL signal standards for LAAS are currently under development by the Radio Technical Committee on Aeronautics (RTCA) Special Committee 159 Working Group 4A. In the meantime, an IntegriNautics IN-200 L1 pulsed C/A code pseudolite (see Figure 4) has been implemented as part of the architecture proposed here.

2.3.3 APL Layout: Given an acceptable pseudolite design, it is necessary to define precisely how they will be used (e.g., how many and where located) within the LAAS architecture. APL siting, in particular, will be critical to LAAS performance and must be carefully considered with regard to airport boundaries, approach visibility, and low-multipath constraints. One APL mounting location that would provide an unobstructed view of all approaches is atop the airport tower. In this case, however, pseudolite ranging would be restricted to altitudes above the tower height since APL transmissions at negative elevation angles must be suppressed in order to mitigate APL transmit multipath. As a result, a tower APL that could support Category I or II applications down to a decision height (DH) of 200 or possibly 100 ft (depending on the height of the tower) could not support Category III requirements through touchdown and rollout. Furthermore, when only one APL is used, LAAS performance is improved, (see Table 2 in Section 2.3.4) but undesirably long outages still remain for Category III. Multiple pseudolites will thus be required for large airports with the highest availability requirements.

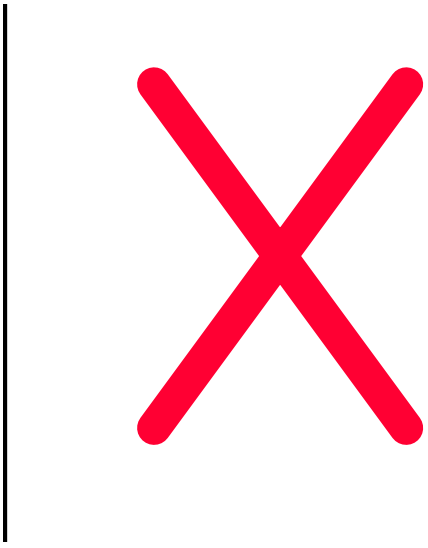

Figure 4: IN-200 Pseudolite

When two or more APLs are considered, the performance contribution due to cycle ambiguity estimation can be significant. In practice, the performance obtained will depend on the location of the APLs relative to the final approach path. It has been suggested that a vertically stacked APL configuration (implemented using a tall antenna tower) would provide a continuous gain in vertical performance as the aircraft approaches the runway [14]. In this case, performance would also improve with progressively taller towers. It has also been proposed that equivalent performance can be achieved on a per-runway basis without excessive tower heights through the implementation of an "intrack" APL configuration. In this arrangement, for a selected Category III runway, one APL is situated at each runway end, along the runway centerline. The resulting vertical performance is equivalent to using a stacked APL with an antenna height in excess of 500 ft – clearly unattainable in a practical airport environment. Note that the APL at the departure end of the runway can also be used to support Category III rollout operations.

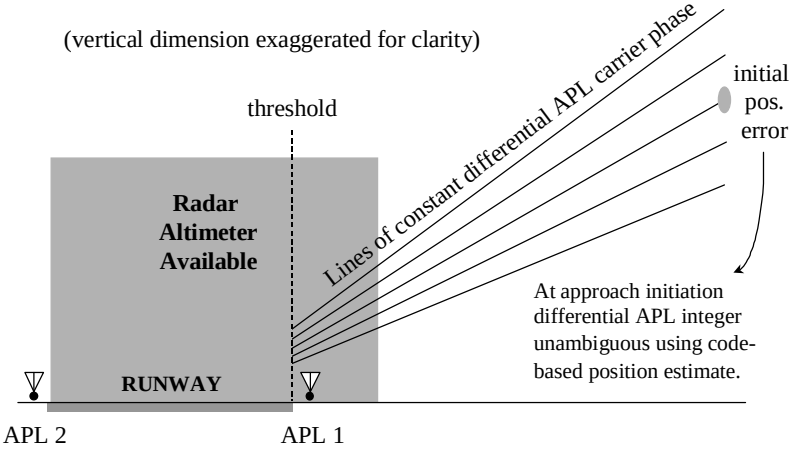

Figure 5: Intrack APL Performance Improvement

The basis for vertical performance improvement using intrack APLs is illustrated in Figure 5. This diagram shows a cross-section of the contours of constant double-difference carrier phase (single-difference phase differenced between the two APLs). In actuality, the contours are hyperboloids in three dimensions with foci at the APL locations. Because the contours become denser as the runway is approached, the APL differential phase becomes more sensitive to vertical position deviations. These differential phase contours represent *relative* lines of position because the differential cycle ambiguity is unknown. However, at the initiation of an approach, the wide spacing of the contours permits the effective resolution of this ambiguity using a vertical position computed directly from carrier-smoothed code. While the illustration in Figure 5 helps explain the vertical performance improvement due to intrack APLs, in practice, code and carrier measurements are simply processed as described earlier in Section 2.1 (in the same manner as satellite measurements).

In order to extract maximum vertical performance benefit from an intrack APL configuration, it is necessary that the aircraft flight path be along the centerline during the course of the entire approach. Obviously, this is not realistic; however, for safety reasons, the aircraft deviation from the nominal flight path must not be too large. One proposed "inner tunnel" requirement [15] stipulates that 95% of the approach trajectory must lie within a specified containment surface. Given this constraint, the worst-case intrack-APL performance will result when the aircraft remains at the outer-lower edge of the tunnel from 1500 ft to 175 ft in altitude (the first 95% of the approach) and outside the tunnel thereafter. It is assumed that during this last 5% of the approach, vertical performance does not improve. Under these worst-case conditions, the effective vertical dilution of precision (VDOP) is lowered to approximately 0.55 times VDOP for case where the two APLs are used as supplementary ranging sources but cycle ambiguity estimation is not implemented.

2.3.4 Predicted performance improvements with APLs: APLs at a given user location provide additional ranging signals from fixed locations (broadcast by LAAS) that can be used as if they were additional GPS satellites. Intrack APL's are specifically sited to provide substantial accuracy improvement as explained above. Even if intrack APLs are not provided to a given airport, the performance improvement obtained from APLs depends greatly on their siting with respect to the landing runway and to local obstructions. To augment the availability simulations introduced in Section 2.3.1, APLs are assumed to be located relative to the primary Category III approach at each airport. A single APL is assumed to be located near the runway threshold for that approach, while intrack APLs add another APL on the opposite end of the runway.

Tables 2 and 3 show the resulting LAAS availability with 22 healthy satellites for the

Airport	CATEGORY I		CATEGORY III	
	Serv. Avail.	Max. Out. (min)	Serv. Avail.	Max. Out. (min)
LAX	0.9999	20	0.9993	75
ORD	1.0000	15	0.9995	30
DFW	0.9999	10	0.9973	45
JFK	0.9996	15	0.9968	35
ATL	1.0000	0	0.9994	40
SEA	1.0000	10	0.9995	50
Overall	0.9999	20	0.9986	75

Table 2: LAAS Availability for 22 Healthy SVs and One APL

single-APL and intrack-APL augmentations respectively. There is a substantial improvement for both Cat. I and Cat. III when one APL is added. Not only does (satellite) service availability increase to almost 0.999 for Cat. III, the maximum outage duration drops by 30 (Cat. I) and 20 minutes (Cat. III) from the no-APL results in Table 1. Note that availability does vary from site to site based on the degree to which the single APL location benefits the local satellite geometry. For each airport, a better site could probably be found than the one used in this simulation.

Airport	CATEGORY I		CATEGORY III	
	Serv. Avail.	Max. Out. (min)	Serv. Avail.	Max. Out. (min)
LAX	1.0000	10	0.9999	35
ORD	1.0000	5	0.9999	20
DFW	1.0000	5	0.9997	20
JFK	1.0000	5	0.9980	20
ATL	1.0000	0	1.0000	0
SEA	1.0000	0	0.9999	20
Overall	1.0000	10	0.9997	35

Table 3: LAAS Availability for 22 Healthy SVs and Intrack APLs

The intrack APL results in Table 3 show a substantial improvement beyond the single-APL augmentation. Service availability for Cat. I is almost one, and the few outages noticed are all in the range of 5-10 minutes. For Cat. III, service availability is well above 0.999, and the worst-case outage of 35 minutes at LAX is much less operationally hazardous than the 90-minute outages possible without APLs. Intrack APLs become even more important if only 21 satellites are healthy (not tabulated), as they limit the longest outage duration for Cat. III to 75 minutes. Category III service availability exceeds 0.999 even for the 21-satellite case.

3.0 PROTOTYPE DEVELOPMENT

In order to validate the projected performance the LAAS architecture detailed above and to assess the practical issues in fielding a LAAS architecture in general, a prototype system has been installed at Moffett Federal Airfield in California. This ground system has been implemented in preparation for upcoming flight trials on an National Aeronautics and Space Administration (NASA) Beechcraft King Air at the end of April 1997 and on an FAA Boeing 727 in mid-May 1997.

3.1 Moffett Ground Segment

An IntegriNautics IN-200 pseudolite is mounted at each end of Moffett runway 32L along the runway centerline as indicated in Figure 6. Each pseudolite transmits L1 C/A code signals (PRNs 12 and 13) pulsed with the RTCM-104 modulation format. The pulse windows are 93 μs (1/11 C/A code epoch) in duration, and the pseudorandom pulse sequence is repeated every 200 ms. The pulse sequences are synchronized to a fixed time offset from the GPS second. GPS timing (one pulse per second, or 1 PPS) is provided by an external Motorola Encore 8-channel GPS receiver (see the background of Figure 4). The pulse time offsets have been selected reduce the number of pulse 'collisions' between the two APLs. This is especially important during the final phase of an aircraft approach, when the APL 1 power is large enough to suppress an APL 2 signal with synchronous pulses.

Reception of the APL signals at the reference antennas is necessary due to the need for differential corrections. For the intrack configuration, this presents a problem due to the great distance separating the reference site from the APLs (APL 1 is over 7000 ft away from the nearest reference antenna). Because of the large lengths required and the need to cross several operational taxiways, the use of coaxial cable was precluded. In addition, preliminary tests indicated that direct reception of the APL signal through the reference antennas was not possible due to the large ground losses and the attenuating effect of the choke ring on low-elevation signals. As a result, two directional helix antennas, one pointed at each APL, were incorporated into the

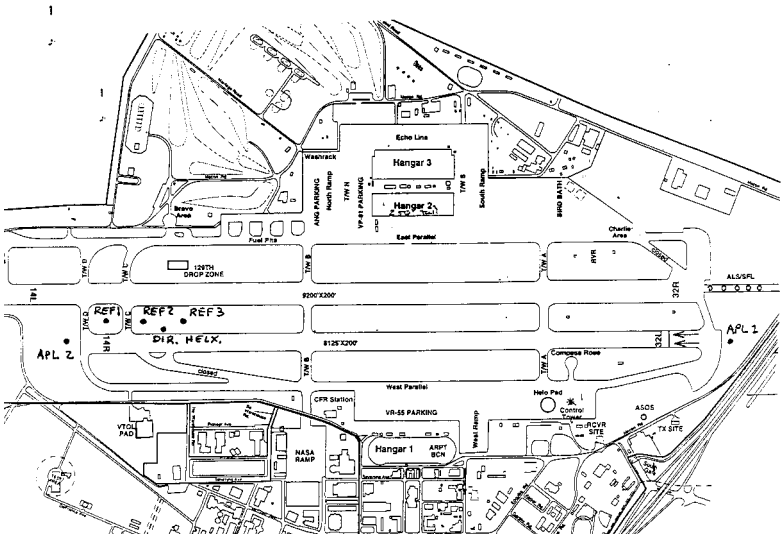

Figure 6: Moffett Field System Layout

reference configuration. These antennas, located as shown in Figure 6, are mounted approximately 11 ft above the ground to provide a relatively clear line of sight to both APLs (see Figure 7).

A straightforward summation of APL signals from the helix and satellite signals from the choke rings is inappropriate because of the possibility of receiving low-elevation satellites through the helix and choke ring simultaneously, resulting in a large multipath-like effect. Instead, an *antenna multiplexing switch* is used in place of a summing junction. The switch is timed in synchronization with the APL pulses to ensure that when the signal pulse from either APL is on, the directional helix connection is engaged. During this short time interval (2 x 93 µs each millisecond), any satellite signals received through the helix are suppressed by the high-power APL signal. When the signal pulses from both APLs are off (over 80% of the time) the choke ring antennas are connected. Synchronization is achieved using GPS time pulses supplied by a reference receiver and accounting for the constant transmission delay from the APLs (~7 µs for APL1). Three switches, one allocated to each antenna/reference receiver pair, are implemented in parallel (see Figure 8).

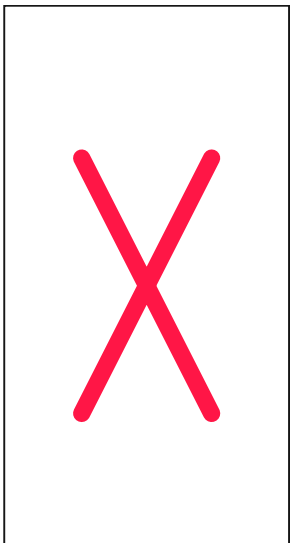

Figure 7: Helix Antenna

Before the satellite and APL signals enter the reference receivers, they are combined with the output of a Welnavigate GS100 single-channel GPS signal generator (PRN 32). This makes it easy to remove receiver clock biases (for ground measurement screening and reference correction generation) but is not a necessary element of the ground architecture, as one can directly compute and remove the receiver clock bias from each range measurement. The LAAS prototype reference station includes three NovAtel 3951R GPSCard Receivers (one dedicated to each reference antenna) residing in a single Pentium PC. The reference station rack incorporating all of these elements is shown in the foreground of Figure 9. The four RF inputs to the rack originate from the three choke ring antennas and the combined output of the two directional helix antennas. The only additional input to the reference station rack is 120 V A/C power. The single reference station output is the data to be transmitted to the aircraft via a Pacific Crest VHF data transmitter sited nearby.

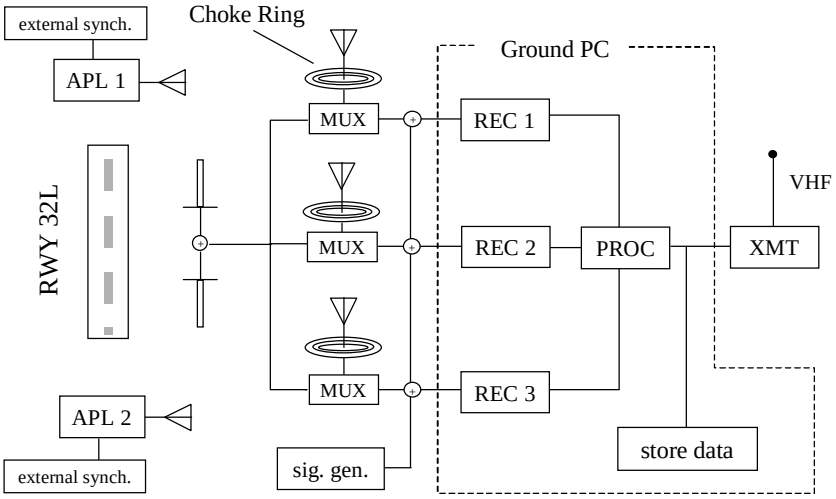

Figure 8: Ground System Implementation

The primary function of the ground processor is to collect and format the data so that the airborne processor is able to reconstruct the raw code and carrier-phase measurements for each satellite and reference receiver. At Moffett, REF 1 selected to be the virtual reference point to which all measurements are projected, after which calibrated antenna cable biases from each receiver are removed. Code and carrier measurements from all channels tracking satellites and APLs are then differenced from the GPS signal generator measurements to remove the receiver clock bias. It is then possible to directly compare measurements from all three receivers for any given satellite or APL (see Section 2.2.3).

Finally, in order to reduce the amount of data transmitted to the aircraft, raw code and carrier measurements are sent only for REF 1, and measurement deviations from REF 1 are sent for REF 2 and REF 3. It is possible to further reduce the amount of data transmitted by removing a computed range (based on the

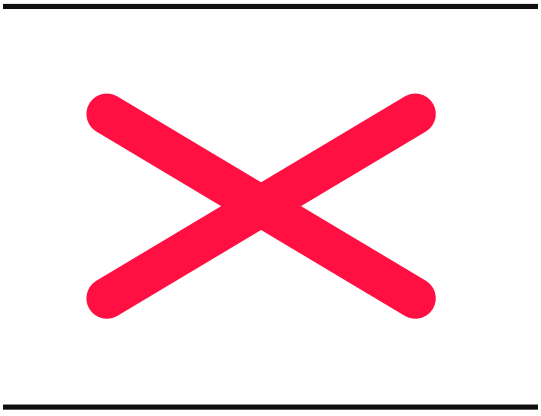

Figure 9: Ground Reference Station

ephemeris) to each satellite on REF 1 and transmitting the resulting correction instead. However, this requires the definition of a protocol to ensure that the aircraft uses the same ephemeris-based range to reconstitute the original measurement; this is not implemented at present. A one-byte phase-valid counter is also transmitted to the aircraft for each channel on each receiver. This counter is incremented when the receiver outputs a loss-of-lock flag to the processor or when the REF 2 and REF 3 measurement deviations exceed the allocated dynamic range. The use of a counter instead of a single-bit valid/invalid flag ensures that the aircraft will be alerted of ground receiver cycle slip even in the event of a temporary datalink dropout.

Cable bias variations due to thermal expansion (associated with ambient temperature changes) have been observed in ground tests performed to date. Because the thermal expansion effect is proportional to the differential cable length, the cable bias variations are significant when the cable length difference is large. In the Moffett implementation, all cable lengths have been cut to the same length within a tolerance of 5 m in order to ensure that the differential carrier phase datalink dynamic range constraint (currently set at 0.5 L1 cycles) is not exceeded under varied weather conditions.

The redundancy in the reference antennas and receivers is required to support the detection and removal of reference receiver measurement failures, although this function is actually implemented at the aircraft (see Section 2.2.3). In an operational LAAS system, it is also likely that redundant processors and data transmitters will be included. At this stage of prototype development, these elements have not been incorporated. However, redundancy management for these future elements will be simpler in the sense that no noise is present and elementary voting and on/off checks can be employed.

3.2 Boeing 727 Setup

An FAA Boeing 727 has been equipped to support flight trials scheduled for mid-May 1997. The test objectives are:

- Assess APL signal quality on a commercial transport class aircraft using fuselage-top-mounted and forward looking (nose-mounted) antennas.
- Verify the performance of the intrack APL architecture.
- Provide a preliminary evaluation of the prototype LAAS integrity monitor performance.

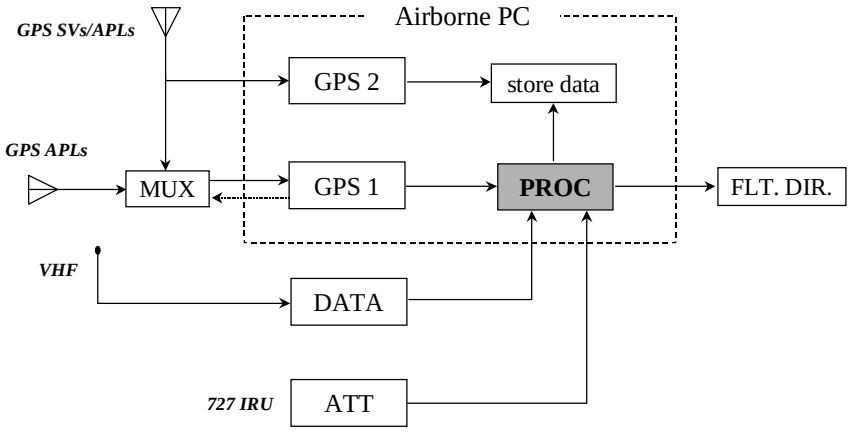

Figure 10: Airborne Implementation

The airborne implementation is shown in Figure 10. Two NovAtel 3951R GPS receivers reside in a Pentium PC. Raw code and carrier phase measurements from the primary GPS receiver (GPS 1) and data received from the ground are processed as discussed in the earlier sections of this paper. The real-time output to the flight director consists of 4 Hz position fixes and vertical and horizontal integrity flags. However, all raw data is stored for post-processing.

Although it is desirable for an operational LAAS system to utilize a single antenna for both satellites and APLs, the quality of an APL signal received through an antenna mounted on top of an aircraft fuselage is questionable because a direct line-of-sight does not exist. For this reason, the 727 has been equipped with both a forward-looking antenna mounted within its radome and another antenna mounted on the top of the fuselage. The input to the primary receiver is multiplexed between the top and nose-mounted antennas in accordance with the APL pulse protocol. The nose antenna is engaged when the APLs are pulse-on and the top antenna is engaged for satellite tracking when the APLs are pulse-off. Synchronization is assured by providing a 1 PPS output from the GPS receiver and a computed range delay correction from the PC to the multiplexing switch. A secondary GPS receiver is dedicated solely to the top mounted antenna to provide the basis for APL signal quality comparison. Data from this receiver is collected and stored for post-processing only; it is not processed in real-time for aircraft navigation.

In preparation for the 727 flight trials, the airborne equipment has been installed in a NASA Beechcraft King Air (Figure 11) for a series of preliminary flights that are being executed at the time of this writing. The existing King Air experimental setup is essentially identical to what will be used in the 727 testing. Although a nose-mounted antenna does not exist on this aircraft, a number of the multiple existing GPS antenna locations are particularly useful. These include

mounts on the top of the fuselage, the top of the tail, and the fuselage belly. The results of these initial flight tests are intended to demonstrate airborne system preparedness for the 727 flight trials.

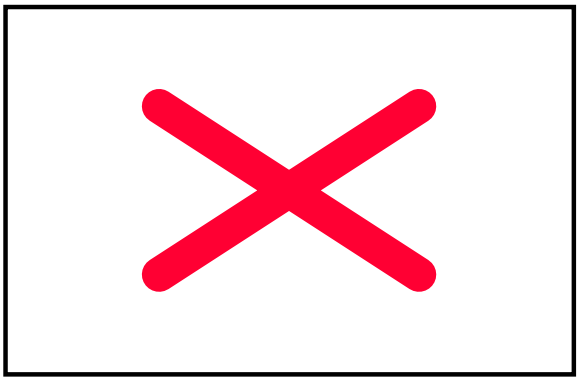

Figure 11: NASA King Air Test Aircraft

4.0 CONCLUSION

A LAAS architecture has been designed with the goal of meeting the stringent navigation requirements for aircraft precision approach and landing. Elements of particular interest in the proposed system include:

- The use of airport pseudolites for availability augmentation, including use of an intrack configuration to improve vertical performance for Category III applications.
- An optimal measurement processing architecture incorporating all information provided by the satellites and pseudolites, including code, carrier, and signal line-of-site motion during the approach.
- A new airborne position domain integrity monitoring architecture.

The proposed LAAS system has been developed and implemented in preparation for upcoming flight trials on an FAA 727 to be performed at Moffett Federal Airfield in California.

5.0 REFERENCES

[1] R. Swider, "Recommended LAAS Architecture," Presented to RTCA SC-159 WG-4, Anaheim, CA, Nov. 12, 1996.

[2] *Operational Requirements Document: Local Area Augmentation System*. Satellite Navigation Program Office, Federal Aviation Administration, February 28, 1995.

[3] R. Swider, D. Miller, A Shakarian, S. Flannigan, and J. P. Fernow, "Development of Local Area Augmentation System Requirements," *Proceedings of ION GPS-96*, Kansas City , MO, September 17-20, 1996.

[4] International Civil Aviation Organization (ICAO), *International Standards, Recommended Practices and Procedures for Air Navigation Services -- Annex 10*, April 1985.

[5] C. E. Cohen, B. S. Pervan, H. S. Cobb, D. G. Lawrence, J. D. Powell, and B. W. Parkinson, "Real-Time Cycle Ambiguity Resolution using a Pseudolite for Precision Landing of Aircraft with GPS," *Proceedings of DSNS-93*, Amsterdam, The Netherlands, March 30 - April 2, 1993.

[6] D. G. Lawrence, B. S. Pervan, C. E. Cohen, H. S. Cobb, J. D. Powell, and B. W. Parkinson, "Real-Time Architecture for Kinematic GPS Applied to the Integrity Beacon Landing System," *Proceedings f the ION 51st Annual Meeting*, Colorado Springs, Colorado, June 1995.

[7] S. Pullen, B. Pervan, P. Enge, and B. Parkinson, "A Comprehensive Integrity Verification Architecture for On-Airport LAAS Category III Precision Landing," *Proceedings of ION GPS-96*, Kansas City, MO., Sept. 17-20, 1996, pp. 1623-1634.

[8] W. Hundley and M. Sams, "LAAS Ground Equipment Availability", *Proceedings of ION NTM-97*. Santa Monica, CA., Jan. 14-16, 1997, pp. 463-474.

[9] B. DeCleene, Operational Availability Requirements Discussion, LAAS Architecture Review Committee Meeting, October 4, 1996.

[10] T. Murphy, "Integrity Monitoring: Position Domain vs. Range Domain". RTCA WG-4A Meeting, Athens, OH., December 6, 1996.

[11] "Appendix: Rationale for Requirements", RTCA WG-4A Meeting, MT-282, Washington, D.C., April 8, 1997 (Draft).

[12] W. Phlong and B. Elrod, "Availability Characteristics of GPS and Augmentation Alternatives", *Navigation*, Vol. 40, No. 4, Winter 1993-94, pp. 409-428.

[13] H. S. Cobb, C. E. Cohen, and B. W. Parkinson, "Theory and Design of Pseudolites," *Proceedings of the ION National Technical Meeting*, San Diego, CA., January 1994.

[14] D. Lawrence, S. Cobb, B. Pervan, G. Opshaug, P. Enge, J. D. Powell, and B. Parkinson, "Performance Evaluation of On-Airport Local Area Augmentation System Architectures," *Proceedings of ION GPS-96*, Kansas City, MO., Sept. 17-20, 1996, pp. 1623-1634.

[15] R. J. Kelly and J. M. Davis, "Required Navigation Performance (RNP) for Precision Approach and Landing with GNSS Application," *Navigation*, Vol. 41, No. 1, Spring 1994.