Flight Test Evaluation of a Prototype Local Area Augmentation System Architecture

Boris Pervan, David Lawrence, Konstantin Gromov, Guttorm Opshaug, Jock Christie,
Ping-Ya Ko, Alexander Mitelman, Sam Pullen, Per Enge, and Bradford 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 alternative 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 optimally processed at the aircraft to improve vertical performance and thus increase navigation availability. In addition, a new integrity monitoring architecture is introduced to provide the tightest achievable protection limits with respect to LAAS reference receiver failures.

To demonstrate that the notional LAAS architecture is realizable, a prototype system was implemented at Moffett Federal Airfield in California for flight testing on a NASA Beechcraft King Air. In this paper, the ground and air components of the prototype architecture implementation are described. The real-time experimental results from King Air flight trials performed in September 1997 show exceptional navigation performance with the APL architecture, including a 95% vertical navigation error of 0.46 m.

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. In this regard, special attention has been devoted in three core areas: measurement processing, airport pseudolites (APLs), and integrity monitoring.

Multiple reference receiver sites are included in the LAAS architecture 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 latter is achieved by effectively averaging pseudorange measurements from widely separated reference receivers.)

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 a runway-'intrack' configuration is chosen to maximize vertical performance subject to on-airport siting constraints. Because of the siting constraint, high-power (pulsed) pseudolites are used to ensure signal reception at the aircraft during the approach.

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 supplementary to that effectively gained through traditional carrier-smoothed code processing. In addition, the airborne measurement processing system utilizes a new 'multiple hypothesis' approach toward protection limit computation with respect to reference receiver failures.

DESIGN

A summary description of the core elements of the proposed LAAS architecture is given in this section.

Measurement Processing

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 1.

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 (due to data broadcast latency) observables provide the measurement basis for sequential cycle ambiguity estimation. Two separate processes, indicated in Figure 1, are implemented in this regard:

Carrier Phase Update. Carrier measurements (ΔCP) are used directly for motion-based estimation of the cycle ambiguities. As will be discussed later, with specific APL ground configurations, motion-based cycle ambiguity estimation can provide substantial improvement in vertical performance.

Code Phase Update. This implementation is essentially an equivalent 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.

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. However, 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 addition a tropospheric refractivity state is included to accommodate differential tropospheric model error. Both of these additional states are observable using the carrier phase update with the APL configuration to be described shortly

In order to generate real-time position fixes at the aircraft, stored reference CP measurements are projected to the current time using a quadratic fit (constant range acceleration) over the last seven measurements [5]. The resulting projected reference CP measurement is subtracted from the actual aircraft CP measurement to generate a projected ΔCP. 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 [5].

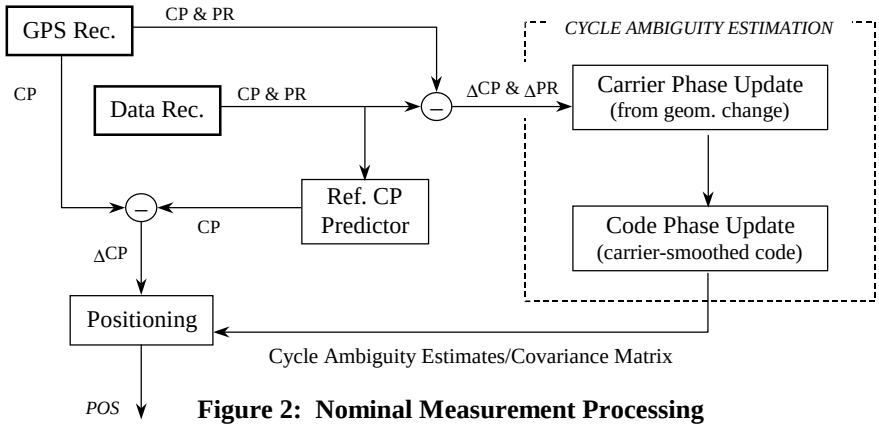

Figure 2: Nominal Measurement Processing

Figure 1: Nominal Airborne Measurment Processing Architecture

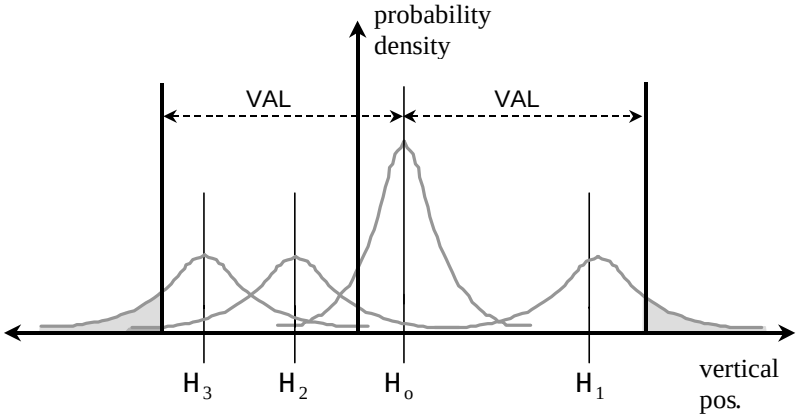

Figure 2: Multiple Hypothesis Integrity Monitor Approach

Integrity Monitoring

The integrity risk and protection limit associated with ground segment failures are computed in real-time at the aircraft. 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 2 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 (H_1, H_2, and H_3) and the no-failure hypothesis (H_o). The mean of the H_o pdf is simply the least-squares vertical position estimate using data from all three reference stations, while the mean of the H_1 pdf is the least-squares vertical estimate using information from only reference receivers 2 and 3. The means of the H_2 and H_3 pdfs are obtained in an analogous manner. The prominence of the H_o 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.

Airport Pseudolites

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 [7]. Specific APL signal standards for LAAS are currently under development by the Radio Technical Committee on Aeronautics (RTCA) Special Committee 159 Working Group 4A. For the present work, an

Figure 3: IN200 Pseudolite at Moffett

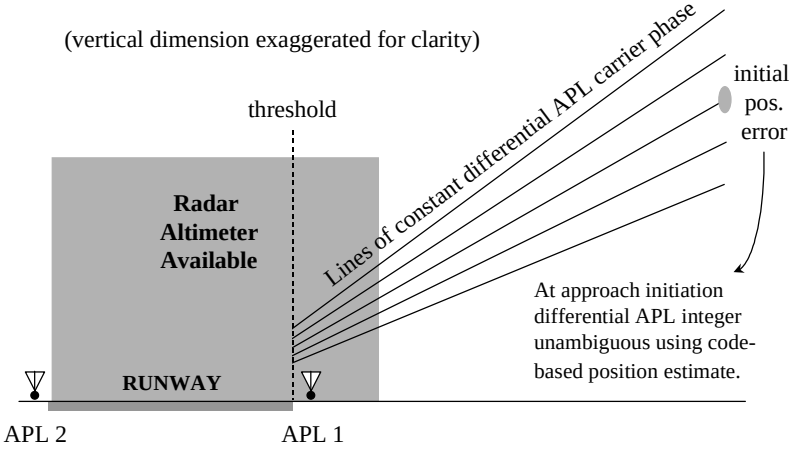

Figure 4: Intrack APL Concept

IntegriNautics IN-200 L1 pulsed C/A code pseudolite (see Figure 3) has been implemented as part of the architecture proposed here.

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.

In general, when APLs are used, motion-based cycle ambiguity estimation becomes possible. In practice, however, the performance obtained will depend strongly on the location of the APLs relative to the final approach path. Specifically, when the APLs are implemented in a runway-'intrack' configuration [8], the resulting geometric observability can provide significant improvement in vertical navigation performance. In this arrangement, for a selected Category III runway, one APL is situated at each runway end, along the runway centerline. Note that the APL at the departure end of the runway can also be used to support Category III rollout operations.

The basis for vertical performance improvement using intrack APLs is illustrated in Figure 4. 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 the same manner as satellite measurements).

The typical vertical performance improvement obtained with the use of intrack APLs is shown in Figure 5. In this figure, the vertical 2s performance computed for a typical 6 SV geometry with and without APLs. Note that when the APLs are added in the intrack configuration, but processed only using carrier-smoothed code, performance

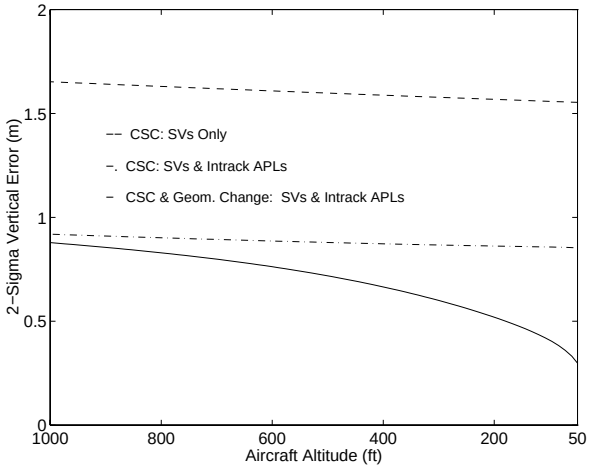

Figure 5: Intrack APL Performance Example

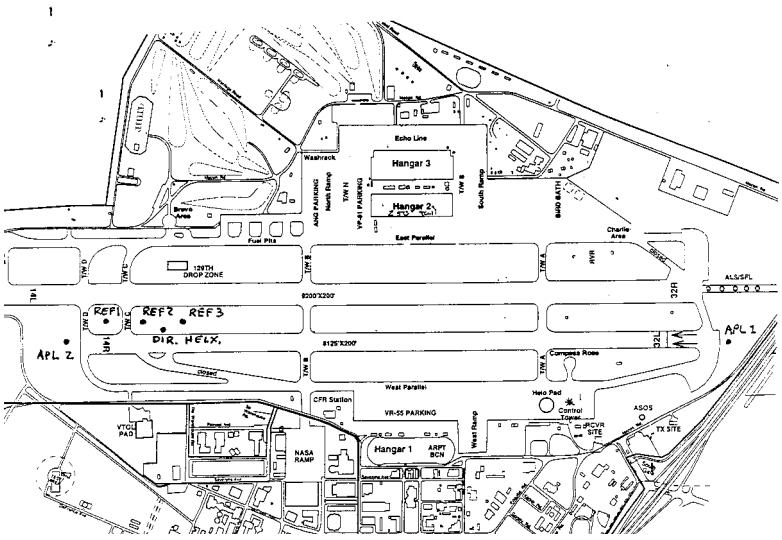

Figure 6: Moffett Field Ground Configuration

improves throughout the approach due to the improved geometry (vertical dilution of precision). However, when motion-based cycle ambiguity estimation is incorporated, vertical performance shows a further steady performance improvement during the course of the approach.

PROTOTYPE DEVELOPMENT

A prototype system has been implemented at Moffett Federal Airfield in California and is described below.

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 11 and 12) 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. 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 reference configuration. These antennas, located as indicated in Figure 6, are mounted approximately 11 ft above the ground to provide a relatively clear line of sight to both APLs.

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 7).

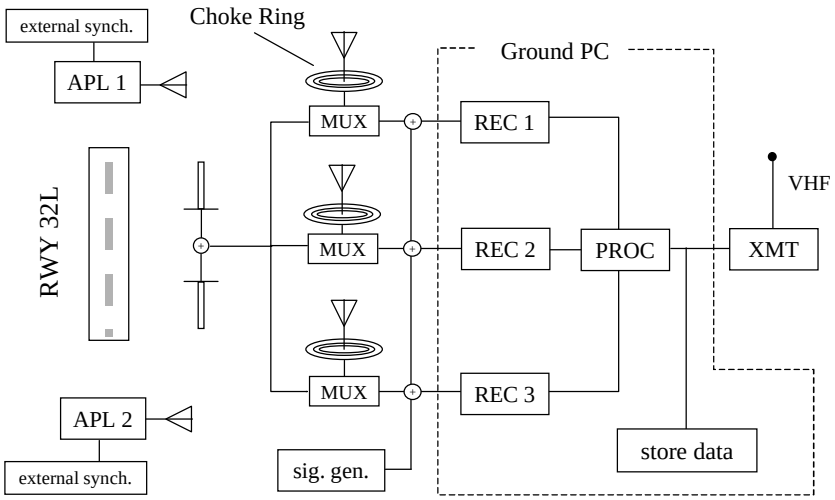

Figure 7: Ground System Functional Diagram

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 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.

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 is designated as 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.

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 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.

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. 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.

Airborne Segment

A NASA Beechcraft King Air test aircraft (see background of Figure 3) was equipped to support a series of flight executed in September 1997. The aircraft was also pre-equipped with multiple GPS antennas. For these flight trials, an existing antenna on top of the aircraft tail was chosen because its favorable visibility with respect to the ground-based APLs on final approach.

The airborne implementation includes a NovAtel 3951R GPS receiver residing in a Pentium PC. Raw code and carrier phase measurements from the GPS receiver and data received from the ground are processed as discussed in the earlier sections of this paper. Real-time position fixes and integrity flags (1 Hz) and all data is stored in the PC for post-flight evaluation.

FLIGHT TEST

In total, 35 approaches were executed at Moffett Field in the NASA King Air between 10-12 September 1997. Of these approaches, 28 were performed using intrack APLs. The true trajectory of the aircraft during these approaches was obtained (in post processing) to within a few centimeters using kinematic carrier phase GPS, with cycle ambiguities obtained during a preflight static survey. The truth trajectory was also validated through another static survey performed after each flight. In real-time, aircraft position fixes and vertical/horizontal protection limits were computed and stored.

Figure 8 is a plot of lateral (cross-track) real-time position error as a function of aircraft altitude for the 28 approaches with APLs. The result is consistently small ($|\mu|+2\sigma$ = 0.40 m) and slowly-varying (bias-like) in each of the approaches. This slowly-varying error structure is attributed to the very low frequency components of the code phase multipath at the reference stations.

Figure 9 shows the vertical navigation sensor error as a function of aircraft altitude for the 28 APL approaches. As expected, the vertical position error becomes smaller as the aircraft approaches the threshold (50 ft altitude). Between 50 and 200 ft altitude, $|\mu|+2\sigma$ = 0.74 m.

In the results of Figure 9, it was noted that the dispersion of the vertical error at altitudes above 300 ft was larger than expected. A detailed analysis of the measurement error structure in the collected raw data revealed that the reason for this effect was APL pseudorange measurements noisier than the SV pseudorange measurements. This was not unexpected given the low (negative) APL elevation angles as seen from the aircraft and the inevitable multipath associated with the RF crosslink in the ground system. Based on the collected flight data, the APL

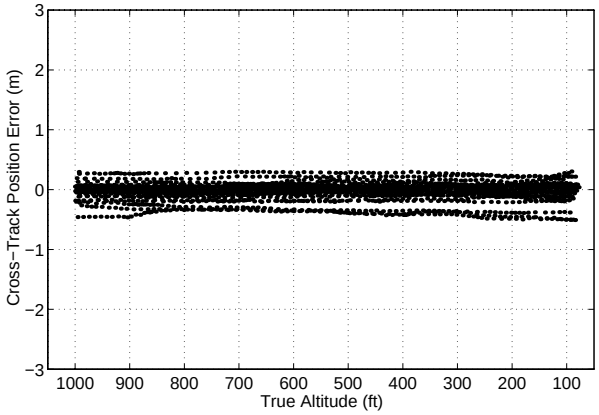

Figure 8: Real-Time Lateral Navigation Error

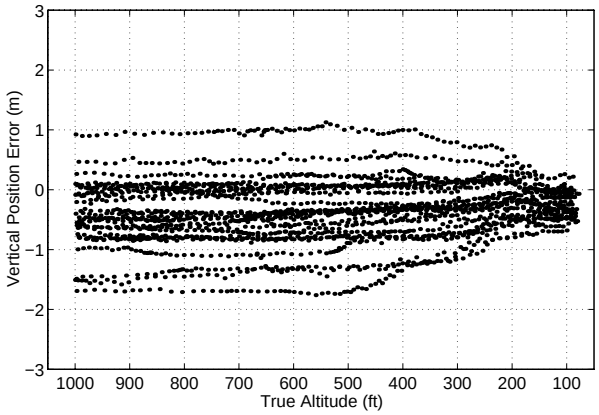

Figure 9: Real-Time Vertical Navigation Error

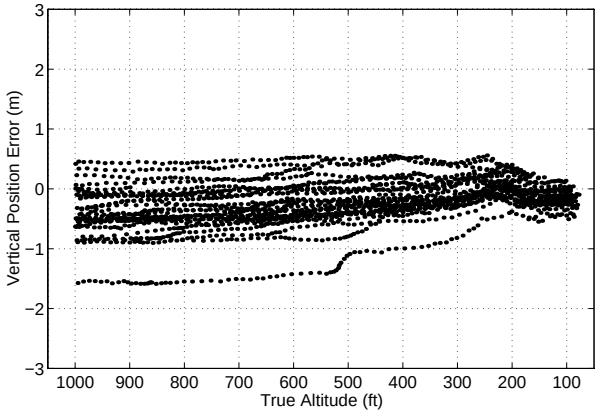

Figure 10: Played-Back Vertical Navigation Error

pseudorange error statistics were modified in the real-time flight code and the raw data was 'played-back' using this code. The results, given in Figure 10, show notably improved vertical performance .

Table 1 summarizes the position error statistics for the 28 APL approaches. In addition, the stored raw data was again 'played-back' for the 28 approaches to evaluate positioning performance using satellites only (i.e., no APLs). The dramatic improvement in vertical performance using the intrack APL architecture is clear in the results. Furthermore, it is significant that the actual intrack-APL performance is more closely characterized by the played-back results (using the updated APL pseudorange error statistics in the estimator) as this is the expected level of performance achievable in future flight testing of the architecture.

Finally Figures 11 and 12 show the vertical and lateral protection limits (VPL/LPL) computed in real-time for 14 of the APL approaches. It is clear in Figure 11, that VPL decreases as the aircraft approaches the runway down to approximately 150 ft altitude. At this point, the signal from APL 1 (at the approach and of the runway) is blocked by the aircraft fuselage. As a result, when the aircraft is lower than 150 ft, the VPL remains a constant level. Note also that LPL (Figure 11) is relatively unaffected by the APL geometry change during the course of the approach.

Nav. Error (m) ave. over 50-200m		Real-Time	Play-Back (SVs only)	Play-Back (SVs/APLs)
μ	LAT	-0.03	-0.03	-0.05
μ	VER	-0.22	-0.40	-0.12
σ	LAT	0.18	0.18	0.19
σ	VER	0.26	0.32	0.17
$\lvert\mu\rvert + 2\sigma$	LAT	0.40	0.39	0.43
$\lvert\mu\rvert + 2\sigma$	VER	0.74	1.05	0.46

Table 1: Flight Test Results Summary

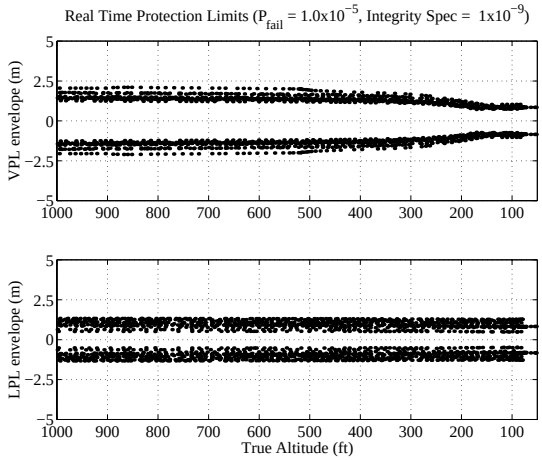

Figure 11: Real-Time VPL and LPL Results

CONCLUSION

A LAAS architecture has been designed, implemented and tested. The features of primary significance in this architecture are: airport pseudolites, an optimal measurement processing architecture, and an integrity monitoring system using multiple reference receivers.

A prototype LAAS system was implemented at Moffett Federal Airfield in California for flight tests using a NASA Ames King Air. In September 1997, 35 approaches were executed in total, 28 in the intrack APL architecture. The fight test results indicate that significant improvements in vertical performance are achievable using the intrack APL architecture (consistent with analytical predictions). In addition a real-time integrity monitoring system using multiple reference receivers was demonstrated.

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] 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 of the ION 51st Annual Meeting*, Colorado Springs, Colorado, June 1995.

[7] 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.

[8] 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.