OEM4 Inertial: A Tightly Integrated Decentralised Inertial/GPS Navigation System

Tom Ford, ***NovAtel Inc.***
Janet Neumann, ***NovAtel Inc.***
Pat Fenton, ***NovAtel Inc.***
Mike Bobye, ***NovAtel Inc.***
Jason Hamilton, ***NovAtel Inc.***

BIOGRAPHIES

Tom Ford is a GPS specialist at NovAtel Inc.. He has a BMath degree from the University of Waterloo (1975) and a BSc in survey science from the University of Toronto (1981). He became involved with inertial and GPS technologies at Sheltech and Nortech surveys in 1981. He became a member of the original group of GPS developers for NovAtel Inc. in 1989. He has helped develop many of the core tracking, positioning and attitude determination technologies at NovAtel Inc. His current focus is the integration of GPS other supplementary systems, especially INS.

Janet Brown Neumann obtained a BSEE from the University of Kansas in 1978 and an MSEE from Iowa State University in 1981. She has been active in many different areas of GPS software and algorithm development since 1983, with her recent focus being on carrier phase positioning methods and software and integration of GPS with other positioning methods. She has been working in the NovAtel GPS group since 1990.

Pat Fenton is the Vice-president in charge of research at NovAtel Inc.. He graduated from the University of Calgary in Survey Engineering in 1981. He became involved in the GPS field at Sheltech and Nortech surveys before becoming a member of the original group of GPS developers for NovAtel Communications (now NovAtel Inc.). He has been the primary designer on all of the NovAtel Inc. GPS receivers and is the inventor of the narrow correlator.

Mike Bobye is a member of the GPS research group at NovAtel Inc. where he has worked as a geomatics engineer since graduating with a BSc in Geomatics Engineering from the University of Calgary in 1999.

Jason Hamilton is a member of the GPS OEM development group at NovAtel Inc. where he has worked as an geomatics engineer since graduating with a BSc in Geomatics Engineering from the University of Calgary in 1998.

ABSTRACT

The synergy between GPS and inertial navigation has been well known in the industry since the inception of GPS. In theory, the continuity of the inertial system can both fill in positioning gaps left by GPS satellite outages and reduce the effect of high frequency GPS errors, while the unbiased nature of the GPS signals can limit the size of the low frequency errors in the inertial system. NovAtel Inc. embarked on a development process more than 2 years ago in order to build a prototype GPS/INS integrated system designed to take advantage of the complementary nature of the two systems.

The components of the prototype system were a Honeywell HG1700 IMU and a NovAtel Inc. OEM4 GPS receiver. The objective of the development was to provide a tightly integrated system at reasonable cost which could give positioning continuously at a 10 cm level provided the GPS signal outages were of short duration. The approach taken during the system development was to take advantage of the existing GPS navigation algorithms and supplement these with a set of inertial algorithms and to use these in a decentralised process that could run on the target processor on the OEM4 board. The result of this development is a modular system which fulfils the accuracy requirements noted and resides entirely on the NovAtel Inc. OEM4 receiver.

A decentralised filter has both advantages and disadvantages compared to a centralised filter. The advantages of simplicity, modularity, component size and independence, are offset to a certain extent by the requirement that at least 4 satellite observations are needed to generate a GPS solution that can be used to control the inertial errors. Another drawback of the decentralised approach is related to eliminating the use of data with poor integrity. In addition, a decentralised filter does not incorporate integer ambiguity states with the states related specifically to the inertial system, so another method of aiding the ambiguity resolution process needed to be devised. Solutions to these problems are described.

In an inertial system, the measurements of the system are used to provide the coefficients of the differential equation set which describes the dynamics of the system errors.

Noise on these measurements tricks the Kalman filter into a condition of “false observability”, in which theoretically unobservable states experience incorrect reductions in standard deviation. This condition was investigated and a solution to this problem was implemented.

In this paper, the authors propose to describe the system architecture, the particulars of the system development noted above and to provide test results which demonstrates the system performance in various environments.

INTRODUCTION

In areas where the satellite coverage is restricted for short periods of time, a serious shortcoming of a GPS only system in the unavailability of position and velocity data during those periods. In addition, the lack of GPS satellites for short periods of time can cause a serious degradation in the type of position information available even when satellites are visible. If the periods of visibility are too short, then the system will loose RTK availability entirely. For some applications the lack of accurate and continuous position information precludes the use of GPS as a navigation tool.

In late 1998 NovAtel Inc. identified an opportunity to improve the navigation capabilities of a GPS only system by supplementing it with some kind of medium grade inertial unit. The idea of the supplementary system was to bridge gaps in GPS coverage, and to seed the resolution process in the RTK filter in order to help the filter resolve ambiguities faster. The inertial measuring unit (IMU) chosen for the integration is a Honeywell HG1700 medium grade strapdown inertial system. The modified NovAtel Inc. OEM4 GPS receiver can use either the AG11 or AG17 models of the HG1700. The AG11 unit has a 1 degree/gyro bias uncertainty, and a 1 millig accelerometer bias uncertainty, while the lower cost AG17 has a 10 degree/gyro bias uncertainty, and a 3 millig accelerometer bias uncertainty. The integrated unit consists of a modified OEM4 GPS receiver, an IMU and a “Pizza Box” data recorder and a post mission package called “BlackDiamond”. So the user has the option of generating integrated results in either real time or post mission.

The integration has been a success. The real time unit has a number of notable features. It can produce position, velocity and attitude at rates up to 100 Hz. It has enhanced pseudorange reacquisition capability (2 seconds reacquisition after a brief outage). The time to ambiguity resolution after a brief outage has been reduced to 10 seconds. The position error after a 10 second outage is 10 cm at the 1 sigma level. The attitude accuracy in a dynamic setting in the continuous presence of GPS is 0.013 degree for roll or pitch and 0.04 degrees for azimuth. The velocity error if GPS is available is 0.007 m/sec. The post mission software (BlackDiamond) has most of the capability of the real time software with the exception of the signal reacquisition capability. BlackDiamond has the capability of estimating the IMU to GPS antenna offset. It can also accept an eccentric offset (for example from the IMU to a camera focal point) which can be applied to the normal output generated by the IMU. Finally, the BlackDiamond CPU load is not an issue, so it can output any combination of the various possible output items in any of 4 different reference frames.

The software design architecture needed to support the performance goals for the system within the cost constraints imposed on the system development and maintenance. To facilitate both these objectives, a decentralised filter design approach was taken. This has two disadvantages and two advantages compared to a centralised filter development. The disadvantages include additional measurement correlation under some circumstances in the inertial Kalman filter, and the requirement that at least 4 satellites be present before an update is possible. The advantages are that the filter sizes are smaller since ambiguity and clock states are not required for the inertial filter, and that the total software architecture is more modular. The last advantage allows incremental testing and debugging during development and a reduced maintenance load after development is completed.

Although the Kalman filter used to process the inertial and GPS measurements is not a centralised filter, the designers have taken steps to ensure that the system is as tightly integrated as possible. GPS position measurements control the error growth of the inertial Kalman filter. The inertial measurements feedback to the GPS RTK filter to help it resolve ambiguities, and to the GPS pseudorange reacquisition process to help it reacquire pseudorange tracking earlier.

PERFORMANCE

The system performance [9] elements of note are position, velocity, attitude and their associated variance-covariance matrices. The position accuracy is shown in table 1.

Table 1: OEM4/AG11 Position Accuracy with GPS

GPS Position Type	Accuracy (1 sigma)
Stand Alone	0.5 to 2 m
Code Differential	0.25 to 1 m
RT-20 (Carrier Float)	0.05 to 1 m
RT-2 (Carrier Fixed Integer)	0.02 m
Post Processed	0.02 to 2 m

The velocity and attitude accuracy is shown in table 2, generated under the assumptions that the system had done a coarse alignment and RTK measurements have been available while the system has been moving for at least 1 minute. Both tables show AG11 results when GPS measurements are available.

Table 2: OEM4/AG11 Velocity And Attitude Accuracy with GPS

Item	Accuracy (1 sigma)
Velocity	0.007 m/sec
Roll	0.013 deg
Pitch	0.013 deg
Azimuth	0.04 deg

The system performance over short GPS outages will degrade according to the system errors at the time of the outage [10] and according to system noise [11]. The velocity errors will increase linearly as a function of attitude and accelerometer bias errors. The attitude errors will increase linearly as a function of the unmodeled gyro bias error. The position error is a quadratic function of accelerometer bias and attitude errors. If the system has had GPS RTK measurements available for some length of time, then the attitude and accelerometer biases will be more accurate, so the resulting position error in the momentary absence of GPS measurements will be less. The bias errors are estimated by the system. After some time, an accelerometer bias will be known to 0.1 milli-g compared to 1 milli-g with no estimation. The open loop position errors will vary depending on the knowledge of the accelerometer bias. An Ag11 system will show the open loop performance over time seen in the following table 3.

Table 3: OEM4/AG11 Open Loop Position Errors (RMS) from all sources

Time/Calibrated	No	Yes
5 sec	0.12 m	0.04 m
10 sec	0.50 m	0.11 m
15 sec	1.08 m	0.20 m
20 sec	1.98 m	0.34 m
25 sec	3.0 m	0.5 m
30 sec	4.5 m	0.7 m

SYSTEM ARCHITECTURE

The knee bone’s connected to the thigh bone,

The system has three main components in addition to the post mission software package. These are the inertial, GPS, and data collection components. The inertial sub-system is a Honeywell HG1700 tactical grade IMU. The GPS sub-system is a NovAtel Inc. OEM4 L1/L2 GPS receiver with software modified to process both GPS and inertial measurements. The data collection component (Pizza Box) time tags the inertial measurements with a GPS time and saves both measurements and processed data generated by the GPS sub-system.

The Honeywell IMU is a strapdown inertial measuring unit that uses a triad of accelerometers and ring laser gyros mounted orthogonally inside a compact 15 cm high by 15 cm diameter cylindrical case to measure specific forces and angular increments experienced in the unit’s body frame. Internally, delta velocity and delta angles are sampled at 600 Hz. From these, coning and sculling compensations are generated and applied to accumulated delta velocities and angles that we incorporate in our navigation software at a 100 Hz rate. The IMU can be either an AG11 (1 milli-g accelerometer bias, 1 degree/hr gyro drift) or AG17 (3 milli-g accelerometer bias, 10 degree/hr gyro drift).

The GPS sub-system is a NovAtel Inc. OEM4 dual frequency receiver modified to incorporate inertial measurements in its navigation solution. It can provide L1/L2 pseudorange and carrier measurements capable of single point, pseudorange differential and carrier based differential positioning at a 20 Hz rate. When the inertial measurements are available it time tags these with a GPS time accurate to 100 microseconds and then uses a decentralised approach to generate a blended GPS/INS solution.

The time synchronisation [13] depends on a counter within the MINOS4 correlation chip that is slaved to GPS time and has a resolution of 1 micro second. An interrupt service routine (ISR) with a high priority waits for incoming data on the serial port linked to the HG 1700. The first byte from the HG1700 causes the ISR to read the counter in the MINOS4 chip, and then uses the count to reconstruct the GPS time tag for the incoming inertial message.

The steady state process is described with reference to Diagram 1 as follows: The inertial measurements are collected and time tagged in the IMU task. Then they are sent to and processed in the INS task at a 100 Hz rate to generate position, velocity and attitude. These are available for logging to the user, although there is a limitation on the amount of 100 Hz data that can be sent at once (one of raw, position etc.). Every time a 1 second boundary is crossed, an interpolated copy of the system components is generated. The position is sent to the RTKMATCHED task in case it needs to be used for a once only per resolution floating ambiguity filter initialization. If the system determines it is stationary, it signals to the INS KALMAN task to do a zero velocity update (ZUPT). If not, it waits until a position is available from one of the GPS filter tasks as determined by the BESTPOS task (but originating in either PSR POSITION FILTER task or RTKFAST task) and uses this position in INS KALMAN task to do a position update. After the update, the state is propagated to the current time, applied to the system at the current time and reset.

Diagram 1 GPS SubSystem Software Architecture

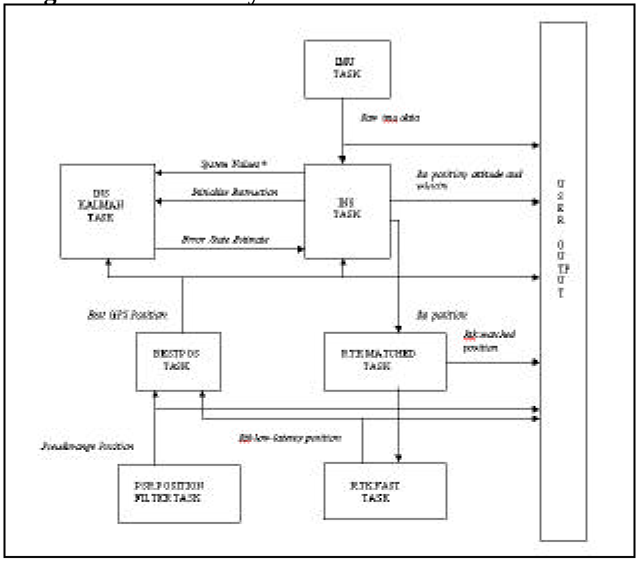

A detailed description of the inertial processing follows.

INERTIAL PROCESSING

The inertial/GPS integration software consists of 4 functional sections, namely a type/frame sensor section, a coarse alignment section, a mechanisation section and a Kalman filter. The type/frame sensor and coarse alignment sections are executed sequentially during a stationary period at the beginning of every mission. The mechanisation section is executed once every 10 msec, and the Kalman filter section is typically executed once per second, although this can vary depending on the availability of GPS position measurements and ZUPTS, and the covariance propagation portion of the Kalman filter takes place on ½ second boundaries.

The Kalman filter has 15 states in the real time process and 18 states in the post mission process. The real time states are ECEF position, velocity, attitude Euler angles wrto the ECEF frame and gyro and accelerometer biases wrto the body frame. The post mission states include these plus an estimate of the IMU to GPS antenna offset. Although states used to estimate both gyro and accelerometer scaling errors are not included, process noise tailored to account for the effects of these errors in a kinematic environment is applied to the velocity and attitude states. The details of the Kalman filter are documented in [15]. The basic 15 state filter was derived from both [3] and [4].

The mechanisation and Kalman filter sections are described in detail in [15] so these descriptions won't be repeated here except to note that the angular integration is based on a quaternion formulation and that the gravity model is parameterised in the ECEF frame.

FRAME DETECTION and COARSE ALIGNMENT

The AG11 and AG17 have different scaling (by a factor of 2) on the lsb of the accelerometer outputs. This is used to detect the type of IMU used in the system. Depending on the model detected, the software will modify the scale factor used, the attitude uncertainties after the coarse alignment, and the level of process noise applied to the covariance matrix during the Kalman propagation.

The coarse alignment estimates Euler angles in order to determine the attitude of the system. The procedure follows Britting Pg198 [2], and requires that the system's designated y axis not be aligned with the gravity vector. During the frame detection, the system will do a comparison of the magnitudes of the nominal x, y and z axis accelerations to see which is closest in magnitude to that of the gravity vector, and then it will reassign the nominal axis of the body frame so that the body frame axis are right handed with the z axis pointing up. The type and frame identification takes 5 seconds.

The accuracy of the alignment over time is dependent on the level of IMU measurement noise (as well as the gyro biases) so the noise level on the measurements must be known to determine the attitude accuracy after the coarse alignment. The noise level on the measurements is estimated during the frame identification process, and is used to determine the initial attitude accuracy to compute an appropriate level process noise during the Kalman filter propagation.

The coarse alignment takes up to 55 seconds, but if the system starts moving before the alignment is complete, the software will begin the navigation phase automatically.

FINE ALIGNMENT

One minute after the first GPS position becomes available to the system, it switches from coarse alignment to navigation mode. If the system remains stationary, there is very little physical impetus that allows the azimuth to become observable. To distinguish this from the navigation mode in the presence of motion, in which the azimuth does become observable, this phase of the process is called fine alignment. If there is noise on the accelerometer measurements, a condition of false observability occurs in which the azimuth standard deviation becomes much smaller than it should. The reason this happens is that the dynamics of the system links the velocity error rates to the errors in attitude states via a skew symmetric matrix of specific forces. If the specific forces have noise on them, then cross terms in the covariance matrix describing the correlation between the velocity and attitude states will vary according to the random walk of the accumulated measurement noise. The same condition also occurs for the same reason in the other attitude states and the bias states. This is documented in more detail in [15], but is interesting enough to repeat here. The following figure 1 compares the azimuth standard deviation computed over time on the bench with real data from the AG11, with the standard deviation computed with the same data, but with the noise on the measurements removed.

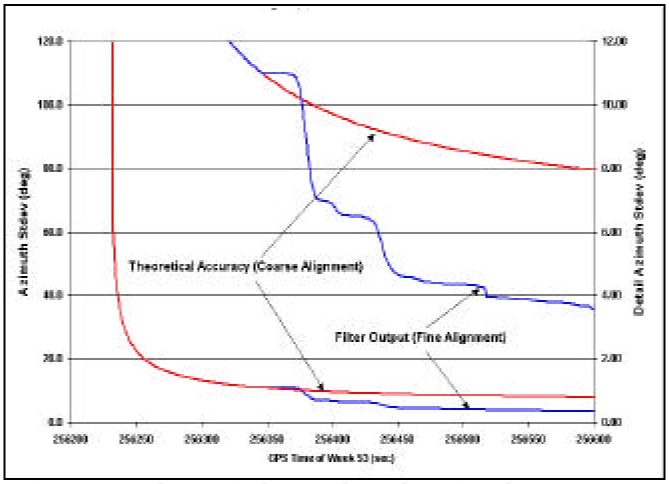

Figure 1: Theoretical vs Filtered Azimuth Accuracy Estimate

The noisy system computes an azimuth standard deviation of 3.7 degrees over a 370 second alignment period. This contrasts with a standard deviation of 8 degrees in the noise free system. In order to mitigate the effect of this noise, the measurements are pre-filtered during fine alignment. The figure 2 below shows the difference in azimuth and its standard deviation with and without measurement noise reduction during the fine alignment.

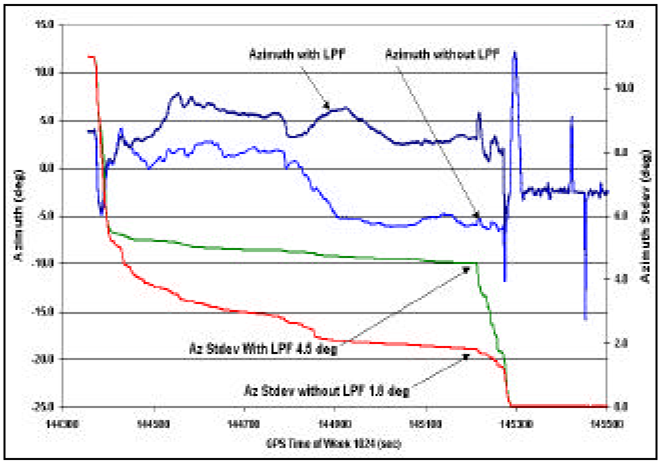

Figure 2: Effect of Measurement noise Filtering

The measurement pre-filtering causes the azimuth estimate to be different by 10 degrees and the standard deviation to be larger by 2.7 degrees (1.8 to 4.5 degrees) over a fine alignment period of approximately 1000 seconds. This increase in standard deviation allows the system to adapt once the system actually moves.

CENTRALISED VS DECENTRALISED

As mentioned earlier, the GPS and inertial processing is carried out in two separate but interacting filters. Together, they constitute a decentralised filter process. The decision to proceed with this design rather than one that incorporated a centralised filter appears to be at first glance a questionable one. A centralised design has a number of advantages. The NovAtel Inc. decentralised filter design requires that at least 4 satellites be available before an inertial update can be made, whereas a centralised filter can accept observations from just one satellite to control its state error growth. A decentralised filter can have measurement correlation arising from the filter supplying the measurements, whereas a centralised filter only has measurement correlation arising from low frequency error sources on the raw measurements. Both advantages are real. The first is partially mitigated by the fact that normally periods of satellite outages in which less than 4 satellites are available are infrequent. The second by the fact that in the NovAtel Inc. GPS filters, the time periods when the filters increase measurement correlation is only during the initial periods of floating ambiguity resolution, and this is relatively short.

Another advantage of a centralised filter is that there is only one filter. All the states are in one place, so the GPS system is naturally enhanced by the effects of the inertial system. This shortcoming of the decentralised filter is overcome with appropriate feedback from the inertial to the GPS filter.

Therefore, many of the disadvantages of a decentralised filter compared to a centralised filter can be overcome to some extent. At the same time, the decentralised filter has distinct advantages over the centralised filter, especially in the NovAtel Inc. development context. Some of these are noted below.

NovAtel Inc. has legacy software which is up to 20 years old. The original pseudorange filter was developed by some of the authors of this paper in the mid 1980s. The RTK software [8][9] has been developed over the last 8 years. All of this has had many man years of testing, and its performance is very stable because of this. Since a decentralised filter would start with a stable GPS core, while a centralised filter would not, this is a strong advantage for the case for a decentralised filter.

To have the same modelling functionality as the current system, a centralised filter incorporating both GPS and inertial states would have all of the inertial states (position, velocity, attitude, gyro biases, accelerometer biases, possibly antenna offset) in addition to the uniquely GPS states (clock, clock rate and ambiguities). This is a filter with a variable number of states between 20 and 30, corresponding to 0 and 10 ambiguity states. Without any optimization, a 25 state filter will require 31250 multiplies during the propagation stage and 179025 multiplies for 6 pseudorange and 5 double difference carrier measurements. In the current design with 18 states, there are 5832 multiplies per propagation and less than 6000 multiplies in a 3 dimensional position update. This is a significant advantage for the decentralised filter side, particularly given that all of the computation takes place on the OEM4 processor board.

Finally, the development platform is an OEM4 receiver, which was released from formal testing in late 2000. The target date to complete an integration was the summer of 2001, so to develop and test a centralised filter for a new platform in that time frame would have included undue risk. As a development strategy, a decentralised filter is safer because the members of the development team can rely on a smaller area of domain knowledge as the scope of the problems they have to deal with is narrower.

For these reasons, a decentralised filter development was chosen. Although the Kalman filter design is not centralised, the system is in some respects tightly coupled, and this coupling is the subject of the next section.

GPS INS INTERFACE

The GPS filters and the GPS/INS filters share position and position covariance information. The GPS filters send the best available position and covariance to help the GPS/INS filter control its state errors. The INS/GPS filter sends its position and covariance to the GPS RTK floating ambiguity filter to aid it in ambiguity resolution. The INS/GPS filter also sends position and velocity to the pseudorange reacquisition routine.

The GPS filters include a pseudorange least squares position estimator, a pseudorange/carrier RTK floating ambiguity filter and a carrier based RTK fixed ambiguity filter. The position and covariance from these are assembled and a best position and associated covariance is chosen based on a minimum covariance trace criteria. This is passed to the GPS/INS filter to use as a position update. Each of the 4 position types have distinctive accuracy and epoch to epoch correlation. The pseudorange filter errors are dominated by multipath in differential mode and multipath plus ionospheric errors in single point mode. Therefore the frequency of the error variation is governed by these error sources. The floating ambiguity RTK filter is dominated by multipath, but the time correlation of the errors is dependent on the nature of the floating ambiguity convergence. The position types and the characteristics of their errors are summarized in the following table 4.

Table 4: GPS Measurement Error Characteristics

Type	Accuracy (m)	Time Constant	Error source
Single Point PSR	3	5 min	Iono/MP (PSR)
Differential PSR	1	3 min	MP PSR
Carrier Float	1 to 0.20	3 min	MP/Conv
Carrier Fixed	0.02	3 min	MP (car)

The multipath error (MP) related time constants are reduced significantly when the system is moving, down to about 5 seconds except for the carrier float solution type which has time correlated errors as a result of ambiguity convergence errors. The pseudorange position estimators doesn't increase the measurement correlation in of itself because it is a single epoch least squares process, rather than a Kalman filter with some time history from clock rate or velocity states. The carrier fixed position type also has long time constant errors associated with carrier multipath, but the amplitude of these errors is smaller than the noise level of the accumulated inertial measurements, so they can be ignored. Time correlated measurement noise generates a modelling error in the filter if it is not taken into account. The measurements are de-weighted for 9 out of 10 observations when the system is stationary or if the system is in carrier float mode. Otherwise, the covariance matrices provided by the GPS position filters are used directly in the Kalman update.

In many environments, the system experiences severe multipath errors (urban canyons are one example where the predominant or even only signal may be a reflected one). To prevent the positions generated with these from corrupting the inertial system parameters via the Kalman update, a six sigma bound is placed on the innovation before it is used to update the inertial filter.

The inertial position and velocity output is used to help the GPS receiver reacquire satellite signals. The satellite reacquisition logic requires the instantaneous code and doppler of the satellite signal as seen by the GPS receiver. The inertial position in conjunction with the receiver clock offset and satellite position is used to generate the theoretical pseudorange to the satellite. The inertial velocity is projected onto the line of sight vector to the satellite and the resulting line of sight velocity is combined with the receiver clock offset rate and satellite motion to generate a theoretical doppler rate for the satellite. To test the effectiveness of this reacquisition method over the current method using the propagated GPS position and last estimated velocity, an OEM4 and OEM4INS systems were set up together on the bench and each were issued commands which caused all the satellites to drop lock for varying lengths of time, from 4 to 120 seconds. Every loss of lock time interval was repeated 10 times, which depending on the number of satellites in view caused the two systems to initiate between 70 and 100 reacquisitions per time interval without satellite signals. The results of this experiment are shown in table 5 below. The OEM4INS system has slightly faster reacquisition times for short outages, but equivalent reacquisition times when outages exceed 60 seconds. This is because the GPS clock offset and dynamics are required for fast signal reacquisition, and the clock model is discarded after 60 seconds without GPS position fixes.

Table 5: Signal Reacquisition Times (Static)

Time Outage	OEM4	OEM4INS	Samples
4 sec	2.3	1.0	90
6 sec	2.5	1.1	90
10 sec	2.7	1.5	90
20 sec	3.0	1.8	90
30 sec	2.8	1.7	100
60 sec	3.1	4.4	95
120 sec	3.2	3.8	74

Statistics were also generated when the system is moving, and these results (times to L1 reacquisition) are shown in the following Tables 5a and 5b for 5, 10, 15 and 20 seconds. Table 5a shows the mean time for a single reacquisition. Table 5b shows the average time it takes to reacquire 80% of the signals. This is important because in order to fix integer ambiguities, six or more satellites are generally required.

Table 5a: Mean Signal Reacquisition Times (Kinematic)

Time Outage	OEM4	OEM4INS	Samples
5 sec	3.73	1.79	210
10 sec	4.15	1.60	145
15 sec	3.40	1.45	140
20 sec	3.40	1.07	140

Table 5b: Mean Signal Reacquisition Times for 80% of satellites (Kinematic)

Time Outage	OEM4	OEM4INS	Samples
5 sec	5.25	1.50	210
10 sec	6.15	1.65	145
15 sec	4.60	1.65	140
20 sec	4.90	1.65	140

The RTK process in the NovAtel Inc. receiver consists of a floating ambiguity filter and a fixed ambiguity filter. When the floating ambiguity filter position becomes accurate enough, it is used to initialise a search space for the fixed ambiguity filter. This process usually takes about 40 seconds. Once the search begins, the ambiguities are resolved fairly quickly and fixed ambiguities are typically generated in less than one minute.

In order to fix ambiguities more quickly with the help of the inertial positions, one of two approaches is possible with the decentralised filter. If four satellites are available, then the floating ambiguity position and its associated covariance matrix can be set by the inertial position and covariance matrix. Then the normal double difference carrier updates will cause the assigned position uncertainty to projected onto the variance elements of the floating filters ambiguity states.

Alternatively, if less than four satellites are available, the ambiguities can be set directly from the difference between the double difference phase measurement and the double differences of the theoretical ranges. The inertial covariance matrix can be explicitly propagated to generate ambiguity variances according to:

$$\sigma^2 = H\ P\ H^T$$

Where

H = Single difference of direction cosine vector for satellites i and j.

$H^{ij} = [\Delta x^i / R^i - \Delta x^j / R^j, \Delta y^i / R^i - \Delta y^j / R^j, \Delta z^i / R^i - \Delta z^j / R^j]$

R = Theoretical range to satellite

P = Inertial position covariance matrix

The method describing the individual ambiguity initialization allows the decentralised filter to assume one of the advantages a centralised filter would have. In some instances the advantage of single ambiguity initialisation over group ambiguity initialisation is small because many outages are caused when the vehicle goes under bridges. Often in these cases 4 satellites or more become available at the same time.

In order to quantify the improvement in resolution times, a test was designed. (In this test a software load was used which limited the ambiguity resolution aiding to the case in which at least four satellites are available.) One OEM4INS unit and two OEM4 receivers with RTK capabilities are connected to the same GPS antenna mounted on a van. All GPS receivers obtain differential measurements from the same GPS base station via a 900 MHz spread spectrum radio. All systems are allowed to resolve integer ambiguities while the van is stationary, then the van is driven on a route that allows for the reception of both continuous uninterrupted satellite and differential observations. After 200 seconds a command is issued to the OEM4INS and one of the standard OEM4 receivers to drop all the satellites for 5 seconds. Then both reset systems are allowed to reacquire satellites and fixed ambiguity resolution. Then the experiment is repeated 9 more times. Then the reset time is extended to 10 seconds and the 10 reset experiment is repeated. Then a 15 second reset test is initiated, followed by a 20 second reset test. In post mission analysis, the sum of the reacquisition and resolution times are compared for the OEM4INS and OEM4 systems, and the reliability of the resolutions is verified for both with the control generated with the third RTK system on the vehicle. The ambiguity resolution times for the integrated system (OEM4INS) and for the standard OEM4 are shown in the following Table 6.

Table 6: Mean Ambiguity Resolution Times

Time Outage	OEM4	OEM4INS	Samples
5 sec	64.8	20.3	10
5 sec *	143.3	29.7	10
10 sec	68.3	23.9	10
15 sec	66.4	26.6	10
20 sec	57.9	32.1	10

* This was a bad geometry case. In 3 out of 10 resolutions the OEM4 did not resolve after 200 seconds, so in fact the OEM4 means would have been even worse if the test extended to the OEM4 resolution in every attempt. Comparisons with control generated from a third receiver that had no signal outages verified that the resolutions on both test receivers were correct.

PERFORMANCE TESTS

In addition to the specific test results tailored to show the performance of the alignment logic, the pseudorange reacquisition and ambiguity resolution functions, some test results are included to show the type of degradation over time when no GPS position observations are available. The results are all taken in vehicles that make frequent stops where zero velocity updates (ZUPTS) can be initiated.

The first test results show the position degradation in the system that occurs when GPS measurements are unavailable intermittently for 30 seconds followed by 100 seconds of GPS availability. The positions are generated from real time inertial data collected on an "L" shaped route north of Calgary where the satellite visibility is almost continuous. The longest base to remote distance is 10 kilometres. Once the system has estimated its gyro and accelerometer biases, the GPS updates are not used for 30 seconds, then are used for 100 seconds. The GPS RTK positions are continuously available and are accurate to 2 centimetres or less, so the INS errors can be reliably computed from this GPS control. The open loop errors are seen on the accompanying figures 3 and 4. The error growth is quite close to the predicted growth seen in table 3.

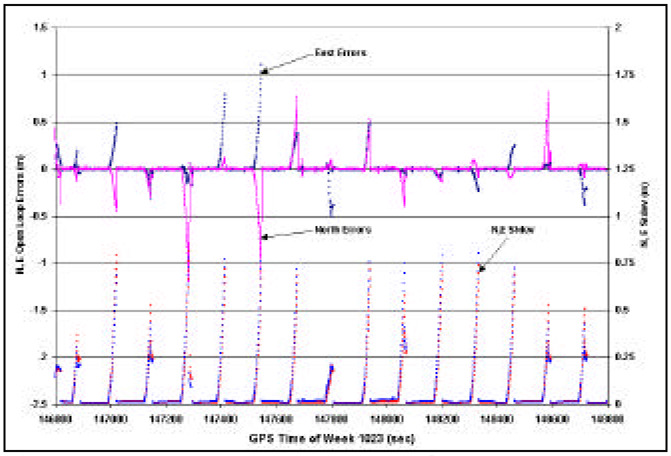

Figure 3: 30 sec Open Loop Error Balzac Aug 16, 1999

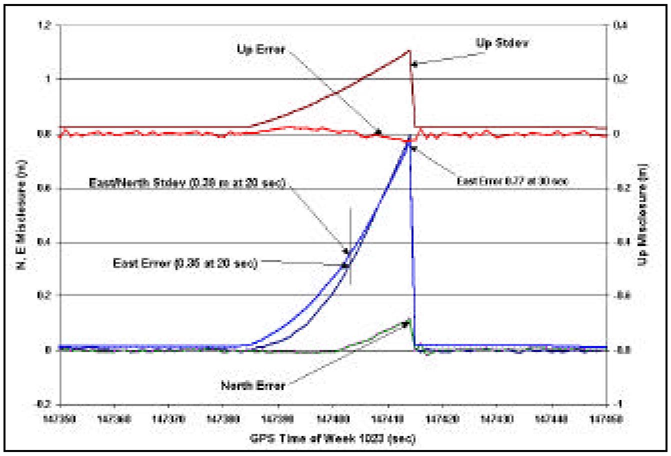

Figure 4: Detail 30 sec Open Loop Error

Figure 5 shows the ratio of the errors generated by comparing the inertial solution with the RTK results computed with GPS data.

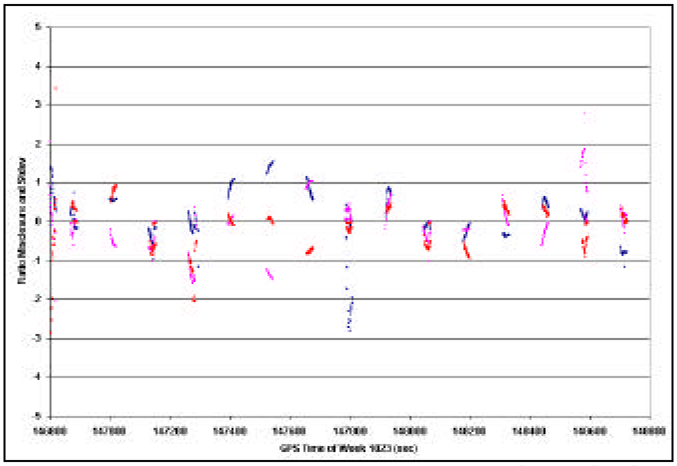

Figure 5: Ratio Open Loop Error to Reported Standard Deviations

Only the standard deviations in excess of 0.1 metres were used, which corresponds roughly to open loop time outages of more than 10 seconds on average. Frequency analysis on this data shows that the standard deviations for the typical outage cases are pessimistic by a factor of 0.67 (too large by that factor), but there are a significant number of larger errors in the tails of the distribution, so the reported standard deviations give a closer representation of the larger magnitude errors.

The post mission software has the capability of processing data with or without ZUPTS. Anecdotal evidence suggests that ZUPTS are effective in controlling the error growth in the OEM4INS system. The effect of ZUPTS on the Balzac data during different lengths of GPS outages can be seen in the following tables 7 and 8.

Table 7: Error growth with ZUPTS

Outage	E RMS	N RMS	H RMS	PTS
30 sec	0.22 m	0.29 m	0.16 m	599
60 sec	0.68 m	0.73 m	0.18 m	890
120 sec	0.41 m	0.39 m	0.34 m	1485
240 sec	0.44 m	0.43 m	0.53 m	1890

Table 8: Error growth without ZUPTS

Outage	E RMS	N RMS	H RMS	PTS
30 sec	0.29 m	0.35 m	0.17 m	370
60 sec	0.72 m	0.75 m	0.19 m	860
120 sec	0.53 m	0.49 m	0.70 m	1051
240 sec	0.56 m	0.73 m	1.06 m	1411

So it is clear, that in this data set, the use of ZUPTS significantly improves the performance, especially when the length of the outage is 4 minutes. The position errors and associated standard deviations are shown in Figures 6 and 7 below for the 4 minute case.

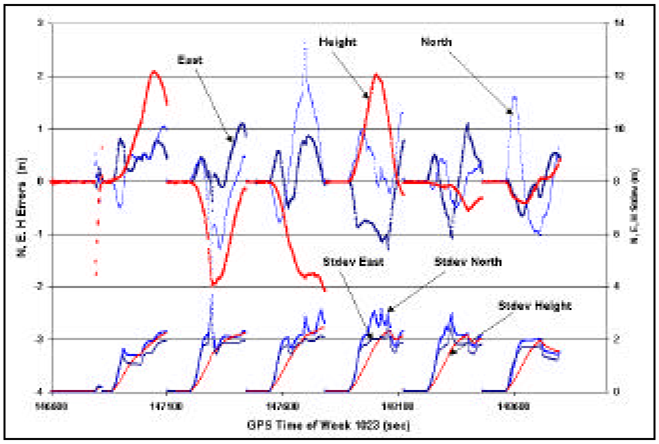

Figure 6: 240 Sec Open Loop, Without ZUPT Balzac

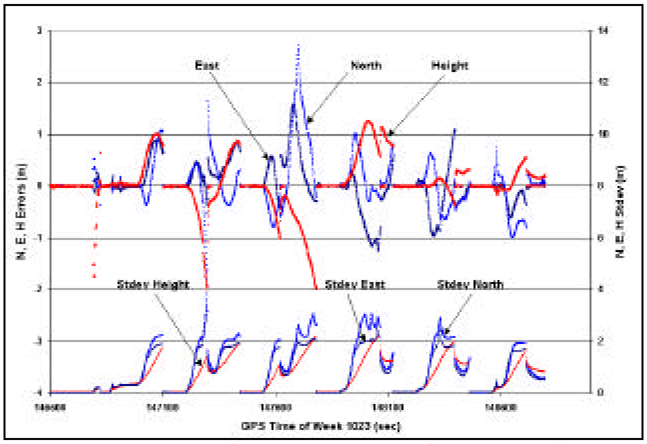

Figure 7: 240 second Open Loop, With ZUPT Balzac

The plots show the effect of changing dynamics on the growth of the errors, even when ZUPTS aren't used to dampen the inertial state error growth.

As a demonstration, some data through the urban canyons of downtown Calgary are processed with and without ZUPTS. The plan view of the positions from this data are shown in Figure 8.

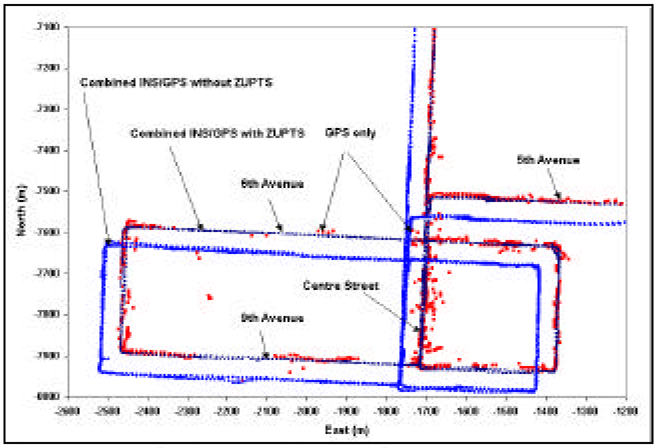

Figure 8: Real Time Positions with ZUPTS vs Post Mission without ZUPTS data collected Downtown Calgary, May 2001

The positions without ZUPTS have been offset by 50 metres. The OEM4INS positions have been aided with single point GPS. Generally, both sets of INS/GPS positions maintain a good agreement with the road plan except the non-ZUPT positions deviate somewhat in the south west corner of the plan. The reason this happens is that the standard deviations grow in the non-ZUPT case in the absence of GPS so that corrupted GPS data is accepted in the non-ZUPT case. But in the real time data process in which ZUPTS were used, the inertial position standard deviation growth was limited by the ZUPTS and as a result the corrupted GPS data was not used to update the inertial

Kalman filter. The difference in position between the ZUPT and non-ZUPT case is shown in Figure 9.

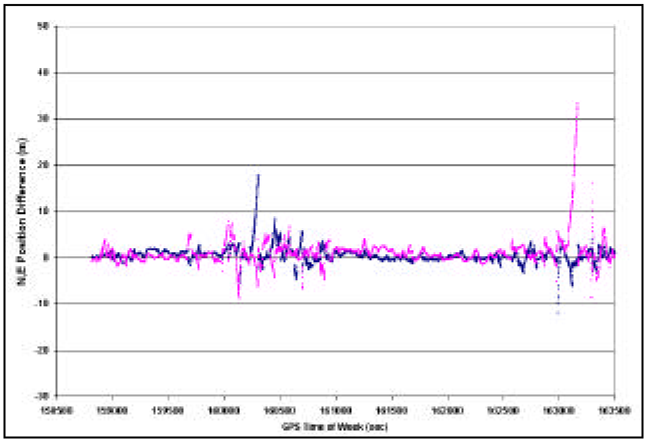

Figure 9: Urban canyon Position Difference ZUPT vs Non-ZUPT

The standard deviations of the inertial positions are shown in Figure 10.

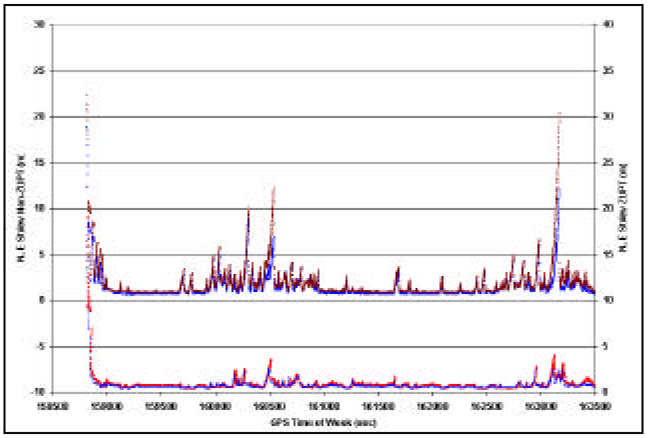

Figure 10: Urban canyon Position Standard deviations ZUPT and Non-ZUPT over time

The uncertainty around times 160500 and 163000 is greater for the non-ZUPT case, making that filter more susceptible to erroneous GPS positions. At the same time, the position discrepancy seen in Figure 8 is reflected in the standard deviations in the non-ZUPT case, indicating that the discrepancy is likely the result of inertial errors that occur when ZUPTS are not used.

CONCLUSIONS

NovAtel Inc. has successfully integrated a Honeywell HG1700 IMU with a NovAtel Inc. OEM4 GPS receiver. This system has been described in this paper.

The system is tightly coupled in the sense that all of the processing takes place on the CPU of the OEM4 receiver and that inertial positions are used to aid ambiguity resolution and with inertial velocities pseudorange reacquisition. The inertial Kalman filter is aided by GPS positions and ZUPTS.

The system is loosely coupled in the sense that the GPS and inertial filters are separate and share processed parameters rather than the alternative of having a system with one filter producing parameters from both GPS and inertial observations.

The design was chosen to limit the size of the Kalman filter and its associated complexity and throughput requirements as well as to take advantage of the legacy software developed over the years at NovAtel Inc.. The effect of correlated GPS measurement noise on the GPS positions used to update the inertial Kalman filter is reduced because of the GPS position generation methods and through selective deweighting of certain position types.

Noise on the inertial measurements creates a condition of "false observability" in which various states in the Kalman filter have reported accuracies that are optimistic by a factor of 2. The effect can be mitigated by pre-filtering the raw inertial measurements before they are used to generated the elements of the transition matrix of the inertial Kalman filter.

The inertial positions and velocity are used to reacquire GPS pseudorange measurements and this significantly reduces the time required for code reacquisition. For short outages (20 seconds or less), the L1 signal reacquisition times are reduced from 3.7 to 1.5 seconds when the system is moving. The average time to reacquire 80% of the signals is reduced from 5.2 to 1.6 seconds.

The inertial positions are used to seed the ambiguity resolution filter. For short signal outages and good geometry, the narrow lane resolution times (this includes signal reacquisition times) are reduced from 64 seconds to 26 seconds on average.

The open loop performance of the system has been examined under different lengths of time when GPS is unavailable. The error growth with and without ZUPTS is compared. Typically, the error growth is 10 cm in the first 10 seconds, and will double every 10 seconds thereafter when ZUPTS are not available.

Analysis of data shows that the open loop errors are reflected pessimistically when the time outage is small and more realistically as the GPS time outage increases.

Tests conducted in urban canyon show that the system performs well during real signal blockages and the signal degradations associated with urban canyons. During these tests, ZUPTS were used to advantage to limit the size of the error growth. This helps the immediate system performance through the damping of the state error growth in the Kalman filter as well as allowing the system to more capably reject corrupted GPS positions before they are used to update the inertial Kalman filter.

ACKNOWLEDGEMENTS

The following people at NovAtel Inc. made significant contributions to the system: Ian Williamson, Theo Smit. It should also be noted that research into inertial navigation has been underway in Calgary for many years, and so the authors have gained knowledge from many of those people that were involved in it: John Hagglund, Richard Wong and KP Schwarz plus all the graduate students [4] that worked under Professor Schwarz on Kingspad at the University of Calgary. Also thanks to Dr. Naser El-Sheimy and Prof. R.G. Brown for their informal contributions through discussions.

REFERENCES

[1] Military Avionics Division Honeywell Inc., "Specification No. DS34468-01"

[2] K. Britting, **Inertial Navigation Systems Analysis**, John Wiley and Sons Inc 1971.

[3] K.P. Schwarz, M. Wei, "INS/GPS Integration for Geodetic Applications (partial lecture notes for ENGO 623)" , Dept. of Geomatics Engineering The University of Calgary, Jan. 1997

[4] Dr. K.P. Schwarz, Hugh Martell, Dr. Ming Wei, Dr. Naser El-Sheimy, Dr. John Zhang, Dr. Ziewn Liu, Dr. Ahmed Mohamed, Dr. Jan Skloud, Dr. Craig Glennie, Dr. Alex Bruton, "KINGSPAD Source code", University of Calgary, 1999.

[5] G. Siouris, Aerospace Avionics Systems: A Modern Synthesis", Academic Press Inc. 1993

[6] A. Gelb et al, **Applied Optimal Estimation**, M.I.T Press, 1974

[7] R.G. Brown, P. Hwang, **Introduction to Random Signals and Applied Kalman Filtering 2nd Edition**, John Wiley & Sons Inc, 1983

[8] Ford, T.J. and J. Neumann , NovAtel's RT20 - "A Real Time Floating Ambiguity Positioning System", Proceedings of ION GPS '94, Salt Lake City, Utah, Sept. 20-23, 1994, The Institute of Navigation, Washington, D.C. pp. 1067.

[9] J. Neumann, A. Manz, T. Ford, O. Mulyk, "Test Results from a New 2 cm Real Time Kinematic GPS Positioning System", Proceedings of ION GPS '96, Kansas City, Mo., Sept. 17-20, 1996, The Institute of Navigation, Washington, D.C. pp. 183

[10] NovAtel Inc. Document OM-20000046 OEM4 Users Manual

[11] M. Kayton, W. Fried, **Avionics Navigation Systems**, John Wiley and Sons, New Yory, 1997

[12] J. Huddle, "Inertial Navigation System Error Model Considerations in Kalman Filtering Applications", Control and Dynamic Systems, Vol. 20 Academic Press 1983.

[13] I. Williamson, Private conversation, NovAtel Inc, May 28, 2001

[14] T. Ford, "Specification of Performance of the OEM4/HG1700 (AG11) Integrated System (Preliminary)", NovAtel Inc. internal document D03011 Rev. A, Sept. 2000

[15] T. Ford, J. Neumann, M. Bobye "OEM4 Inertial: An Inertial/GPS Navigation System on the OEM4 Receiver", Presented at the INTERNATIONAL SYMPOSIUM ON KINEMATIC SYSTEMS IN GEODESY, GEOMATICS AND NAVIGATION (KIS), Banff, Alberta, Sept. 2001