Analysis of Sensors and Techniques for the Design of a Robust Personal Navigator

Dale Arden
Dale Arden Consulting
R.R. #3
Almonte, Ontario, Canada K0A 1A0

dale.arden@drdc-rddc.gc.ca

Jeff Bird
Defence R&D Canada - Ottawa
3701 Carling Ave
Ottawa, Ontario, Canada K1A 0Z4

jeff.bird@drdc-rddc.gc.ca

ABSTRACT

GPS has become ubiquitous on the modern battleground. However, current asymmetric threats often force our dismounted soldiers into combat in urban, indoor, subterranean or other environments where GPS signals are blocked or reflected or attenuated or jammed, resulting in the denial of adequate positioning performance. It is becoming increasingly clear to planners and commanders that precise location determination of these soldiers is critical to mission success. To fulfil the requirement of providing precise locations in these difficult environments, new systems and techniques are required.

This paper describes tests that have been conducted at DRDC Ottawa to help determine sensor configurations that are most likely to meet the goal of precise location determination in the absence of GPS. A Robust Personal Navigator Prototype (RPN-P) has been assembled to provide data for testing. The RPN-P includes MEMS IMUs, GPS receivers, a digital compass, a radar speed sensor, and the computer hardware and software needed for data collection.

1.0 INTRODUCTION

GPS has become ubiquitous on the modern battleground. GPS receivers are guiding aircraft to safe carrier landings, steering projectiles precisely to their targets, providing timing control to communication networks, guiding dismounted soldiers through unfamiliar terrain, and on and on. In a benign environment, GPS does all it is tasked to do, accurately and efficiently. However, often our men and women in uniform are required to do their jobs in environments that are hostile not only to them but also to their GPS equipment. The threats to GPS may be intentional (jamming) or inherent (unintentional signal interference, or signal attenuation or blockage by foliage, terrain, or buildings).

Arden, D.; Bird, J. (2007) Analysis of Sensors and Techniques for the Design of a Robust Personal Navigator. In *Military Capabilities Enabled by Advances in Navigation Sensors* (pp. 16-1 – 16-14). Meeting Proceedings RTO-MP-SET-104, Paper 16. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int.

GPS limitations have long been recognised. High-value assets like ships, aircraft and tanks carry Inertial Navigation Systems (INS) that can maintain required navigation and guidance performance through periods of GPS outage. When GPS is available and reliable, it is used to calibrate INS sensor errors. When GPS is lost or unreliable, the now-calibrated INS dead reckons until GPS recovers. Additional sensors often aid the INS to further slow error growth. Traditionally, height (such as a baro-altimeter on an aircraft) and velocity (like a ship's speed log) sensors have been the preferred extra-GPS aids.

Current asymmetric threats often force our dismounted soldiers into combat in urban, indoor, subterranean or other difficult environments. It is becoming increasingly clear to planners and commanders that precise location determination of these soldiers is critical to mission success. However, these environments are antithetical to reliable GPS coverage. Buildings block and reflect GPS signals; walls attenuate or block them; jammers prevent the receiver from locking onto them. To fulfil the requirement of providing precise locations, new systems and techniques are required.

In the last few years, miniature gyroscopes and accelerometers have been developed using MEMS (Micro-Electro Mechanical System) techniques. These MEMS inertial sensors have been combined into MEMS Inertial Measurement Units (IMU), producing devices that are small and efficient enough to be used in an integrated navigation system for a dismounted soldier. Such an integrated system is conceptually similar to that found on high-value assets: GPS and other sensors calibrate IMU errors when they are available; after GPS is lost, the calibrated IMU (with the help of any remaining sensors) dead reckons until GPS recovers. The real problem lies in the magnitudes of the MEMS sensors errors: they are currently orders of magnitudes larger than traditional inertial sensors. Even a well-calibrated MEMS IMU, on its own, will drift away from its true position very quickly. After GPS is lost, maintaining adequate position accuracy depends on the other aiding information.

Thus, while a MEMS personal navigator is similar in concept to a conventional integrated system, the implementation is very different, with different aiding information, different dynamics, different human factors constraints, and so on.

This paper describes tests that have been conducted at DRDC Ottawa to determine sensor configurations that are most likely to meet the goal of precise location determination in the absence of GPS. A prototype system (the Robust Personal Navigator Prototype or RPN-P) has been assembled at DRDC Ottawa. It is comprised of MEMS IMUs, GPS receivers, a digital compass, a radar speed sensor, and the computer hardware and software needed for data collection. Lacking in this configuration is an independent height sensor, such as a baro-altimeter.

Personal navigators typically rely on the walking motion of the user to provide measurements to reduce system error growth. Information from the inertial sensors can be used to simply detect steps, with a step prediction algorithm providing the aiding information. Alternatively, with an IMU mounted on the user's boot, a zero velocity update at each boot strike can be used. Or, stride length and direction can be calculated using a boot-mounted IMU. This information can be used as a constraint on error growth. The RPN-P currently includes two different IMUs, one mounted near the user's centre of gravity and a second that is boot mounted.

The remainder of this paper will describe the RPN-P and its component sensors, provide an analysis of different sensor data combinations (to judge sensor suitability for inclusion in an integrating Kalman filter), and present some conclusions derived from the data analysis as well as a preliminary road map for the way ahead.

The next step in the RPN-P evolution will see the integration of its sensors into a Kalman filter. DRDC Ottawa has Kalman filter software that has been used in a number of applications over the years. The current version is known as the EGIM – the Extensible GPS Integrated MEMS System. This is a loosely coupled implementation using GPS positions and velocities (along with data from other sensors) to control IMU navigation errors and to calibrate IMU sensor errors. To be used in the RPN-P, EGIM will have to be modified to add new sensors (e.g. the radar speed sensor and new IMUs), and additional measurement models will have to be developed and implemented (speed measurements, stride length estimation or calculation, etc.). At some point, a baro-altimeter will be acquired and added to the RPN-P and EGIM sensor suite. Height information is required to control INS height instability and to provide three-dimensional navigation in the absence of GPS height information.

2.0 DATA COLLECTION

Sensor data was collected using the RPN-P on the Shirleys Bay campus that includes DRDC Ottawa. Runs were made in different environments: along a loop with few buildings and thus good GPS satellite visibility and other loops through and around various buildings on the Shirleys Bay campus where satellite blockage and multipath significantly affected GPS performance.

2.1 The Personal Navigator Prototype

Development of the DRDC Ottawa Robust Personal Navigator is in its very early stages. The configuration used for the following results is shown in Figure 1. It is comprised of

- An ISIS IMU from Inertial Science, Inc. of Newbury, California,
- An Inertia-Link IMU from Microstrain, Inc. of Williston, Vermont,
- A DAGR military GPS receiver from Rockwell Collins, Inc. of Cedar Rapids, Iowa,
- A BU-353 high-sensitivity civilian receiver from GlobalSat Technology Corp. of Taipei, Taiwan,
- An HMR3500 digital compass from Honeywell Aerospace, and
- A DRS1000 radar speed sensor from KMH Engineering of Orem Utah.

Most sensors are mounted on a plywood base that is in turn attached to a modified backpack. The BU-353 is mounted on a short mast attached to the backpack frame, and the Inertia-Link was velcroed to the user's foot. The user wears the pack backwards, as a "frontpack". This allows him to monitor data collection software and sensor displays.

The computer is an Ultra-Mobile Personal Computer (UMPC), a model i7210 from tabletkiosk™ of Torrance, California. As stated on the manufacturer's website (www.tabletkosk.com) it "features a compact, lightweight design that packs the power and compatibility of a traditional PC and runs Microsoft® Windows® XP Tablet Edition 2005". This model has a 1 GHz Intel Pentium M® CPU with 1 GB of memory and a 60GB hard drive. Software developed on and compiled for a desktop PC can be run directly on the UMPC. At present, it is used for data collection. As the project progresses, the UMPC will be used to run the EGIM Kalman filter in real time. The 7" TFT LCD serves as the user display.

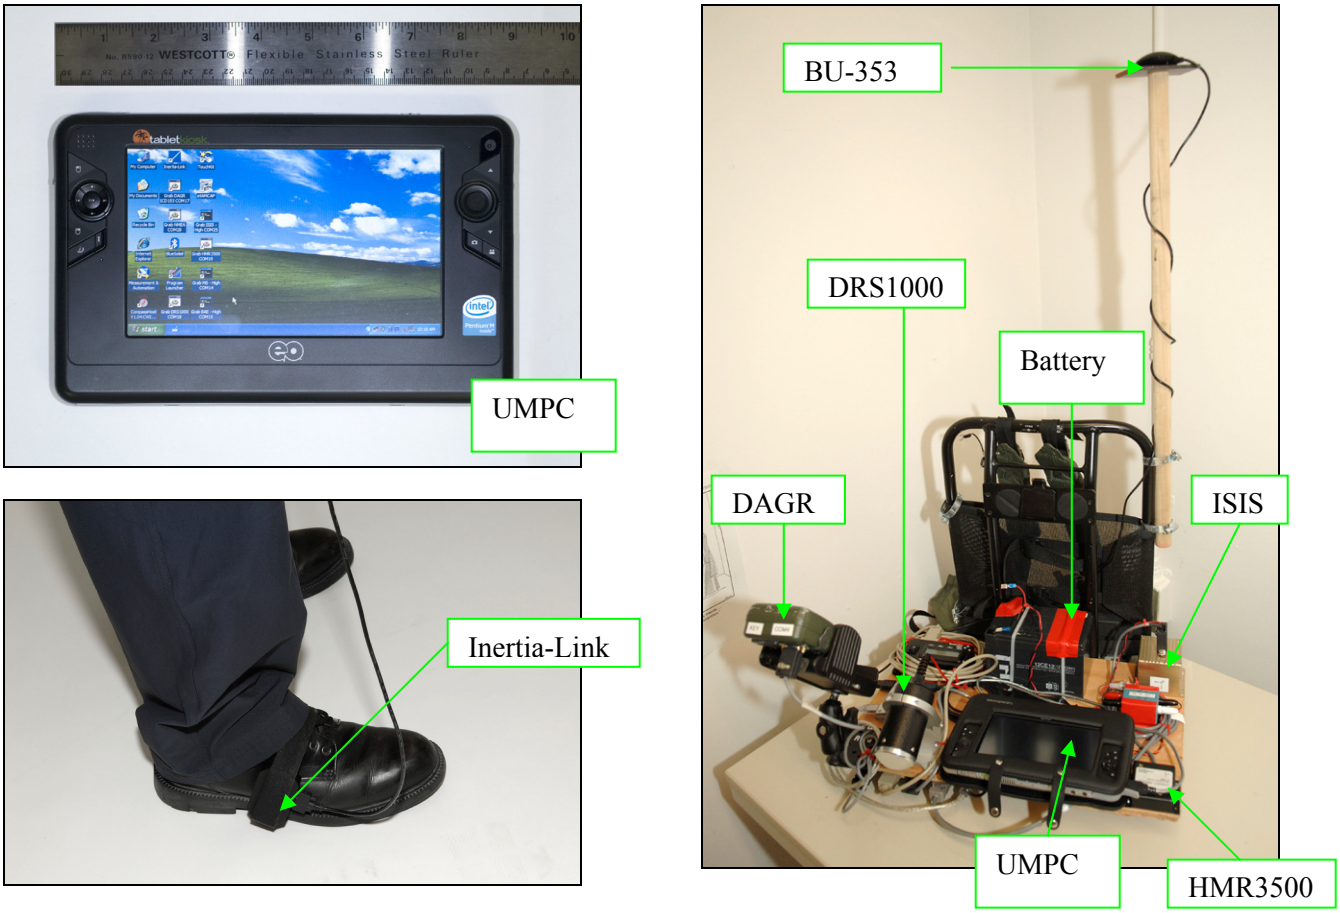

Figure 1: Robust Personal Navigator Prototype

Figure 1 shows a close-up of the UMPC, the RPN-P with the Inertia-Link mounted onto the user's shoe and all other hardware mounted on the frontpack.

Figure 2 is a block diagram of the RPN-P showing hardware components with data and power connections. The ISIS IMU, DAGR, DRS1000 and HMR3500 output data over RS-232 serial communication ports (shown in red). Each is connected to a USB / serial adapter which is in turn connected to a USB hub. The BU-353 and Inertia-Link connect directly to the USB hub. All USB connections are shown in blue. A central 12V DC lead-acid battery powers the ISIS IMU, the HMR3500 compass (after being transformed to 5V DC – not shown) and the DRS1000 speed sensor. The computer and DAGR operate on their own internal batteries, while the Inertia-Link and BU-353 are powered by the computer battery via their USB connections.

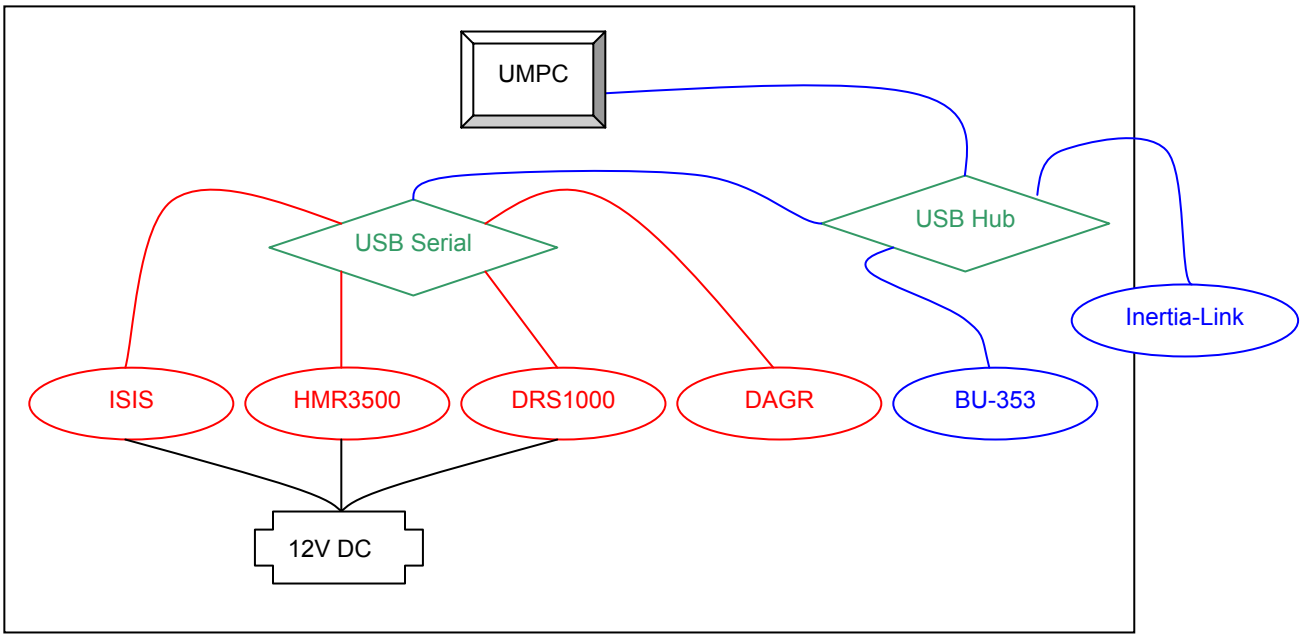

Figure 2: RPN-P Block Diagram

2.2 The Shirleys Bay Campus

Figure 3: Test Tracks at Shirleys Bay

Figure 3 is an aerial view of the Shirleys Bay campus. There are two data collection loops shown. Track 1 is the longer, solid red loop that is mostly away from buildings, where GPS satellite visibility is good. Track 2 is the shorter, dashed green loop that is in and amongst various one to three storied buildings, where signal blockage and multipath simulate an urban canyon environment.

All runs started near the upper right corner of the green track. Both tracks followed a generally counter-clockwise direction, with the red track proceeding from the start point in an easterly direction and the green track proceeding from the start point in a westerly direction. North is up in all figures.

3.0 DATA ANALYSIS

This section is designed to illustrate the usefulness of non-GPS data in controlling IMU position errors in the EGIM Kalman filter. In general, data from the IMU(s), compass and speed sensor are pre-processed and then compared to a GPS reference. The section begins with an analysis of the comparative performance of the two GPS receivers.

Horizontal position performance is demonstrated by superimposing horizontal tracks on aerial photos of the site. Note that the registration of the aerial photos relative to the WGS-84 reference frame is approximate. Scale can be estimated by the vehicles seen in the parking lots in the aerial photos.

3.1 GPS Data

The RPN-P includes two GPS receivers, a dual-frequency (L1 and L2) military DAGR and a civilian (C/A code, L1 only) BU-353. The DAGR was keyed (using encrypted Y-code signals) for all tests. It was tested with the external antenna (co-located with the BU-353) and with the internal antenna only. The BU-353 is a high-sensitivity integrated receiver / antenna built around a SiRF™ Star III High Performance GPS chipset that is WAAS (and EGNOS) enabled. It was located on the mast shown in Figure 1.

The following track plots show the DAGR and BU-353 results superimposed onto a background aerial photo. In all cases, the DAGR track is magenta and the BU-353 track is cyan. The tracks shown in each figure were collected simultaneously. The track plots are accompanied by corresponding plots of height for each receiver (using the same colour scheme).

Figure 4: Track 1 - DAGR (Keyed, Int Ant) vs BU-353

Figure 4 (horizontal track and height) shows the only run taken around Track 1. Only a representative section of the track is shown. With clear-sky satellite visibility, this is intended as a reference for the following "urban canyon" runs.

Differences between the horizontal and vertical positions both average 2 metres. There is one 30-second anomaly (originating with the DAGR data, as seen in the height plot of Figure 4) where differences jump to 9 and 11 metres respectively. The cause of the DAGR's difficulties at this time is not clear.

The next two tests were conducted over Track 2, the simulated urban canyon. The buildings through which the tests were run include single-story steel-sided and three-story brick-sided structures. Away from the buildings, the DAGR and BU-353 tracks are generally very similar. In amongst the buildings, the most interesting differences are found, so the subsequent plots show only that part of Track 2 that is affected by the buildings.

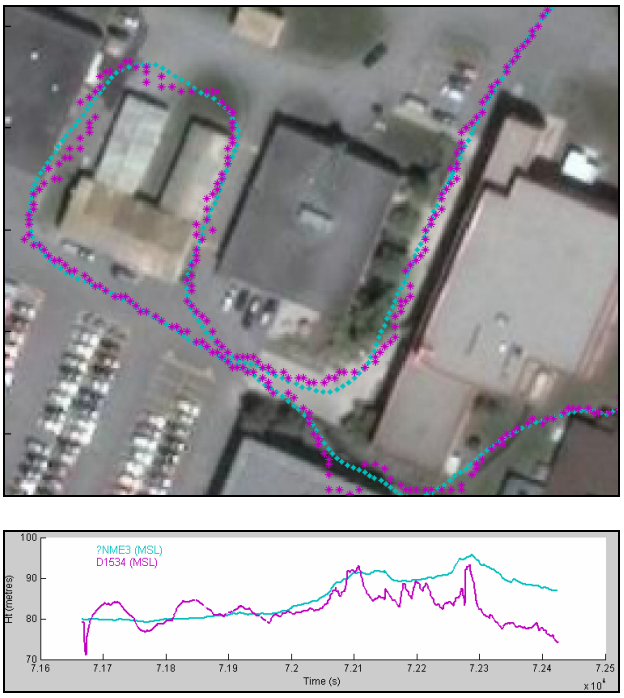

Figure 5: Track 2 - DAGR (Keyed, Int Ant) vs BU-353

Figure 6: Track 2 - DAGR (Keyed, Ext Ant) vs BU-353

For the first Track 2 test, shown in Figure 5 (horizontal track and height), the DAGR only had access to signals via its internal antenna. Position differences relative to the BU-353 average 2 metres (horizontal) and 4 metres (vertical).

The second Track 2 test, with the DAGR using an external antenna co-located with the BU-353, is shown in Figure 6. Horizontal position differences average 2 metres, while vertical errors average 7 metres.

Note the large DAGR excursions visible in both Track 2 plots at the extreme bottom of the plots near the right corner. This occurs during the approach to the tallest, narrowest urban canyon portion of the track. The last spike in DAGR height in each plot coincides with the track excursions. This error behaviour along with the DAGR status information suggests that multipath is the cause.

3.2 IMU Data

The RPN-P has two IMUs that are somewhat similar in performance. However, the mounting arrangements are decidedly different: the ISIS is mounted on the RPN-P sensor platform, at waist height, with relatively low dynamics; the Inertia-Link is boot-mounted, experiencing high-dynamics as the user walks.

A rigorous, standalone strapdown navigator algorithm cannot be applied to these units because of their high error growth rates. The EGIM software uses closed loop error control to continuously correct IMU errors. To get a rough feel for the performance levels, the gyros were (mostly) removed from consideration. Accelerations were integrated twice to calculate distance travelled. The results were compared to GPS-based distances.

3.2.1 Data Processing

To calculate distance travelled using IMU accelerometers, all non-motion effects must be removed from the accelerometer measurements. A simple process is used to do this. When the IMU is stationary, the accelerometers sense only gravity and instrument biases. Assuming gravity and sensor biases are constant over the course of any test, they were calibrated out of the measurements as follows.

First, at each IMU output epoch, the magnitude of the output acceleration vector was calculated. Then, periods of time when the IMU was stationary were identified. All acceleration magnitude values over these motion-free periods of time were averaged. This average is assumed to include all acceleration biases – sensor and gravity. It acts as the calibration correction for all other accelerations.

By subtracting the calibration value from every acceleration magnitude, values that are largely free of biases are produced. Errors in this calibration technique arise from movement during calibration or from changing accelerometer biases. Gravity variations are insignificant.

Note that the accelerations from individual accelerometers cannot be calibrated due to the significant changes in the gravity signal visible to each accelerometer during a step cycle, as its attitude changes.

Integration of an acceleration vector time series will generate changes in velocity. For example, a boot-mounted IMU will start at zero velocity. As a step is initiated, there is an acceleration forward and up. As the step is concluded, accelerations are backward and down, ending with another zero velocity state between steps. With acceleration magnitude, there is no direction information. Simply integrating the acceleration magnitudes will result in a “speed” that increases under positive as well as negative acceleration. In our example, the calculated speed will increase as the step begins (as it should), but speed will also increase as the step ends (not as it should). At the end of a step, the “speed” calculated in this way will not be zero.

However, with some empirical massaging, a reasonable estimate of total distance travelled can be produced. The procedures used for each IMU are detailed next.

3.2.2 Inertia-Link IMU

Recall that this IMU was boot-mounted – we can be confident that the IMU is stationary at each step when the boot is planted on the ground. Errors in the distance calculations can be reduced by identifying the times when the IMU is stationary and forcing the speed to be zero – much like a ZUPT (zero velocity update) often used in Kalman filtering.

Inertia-Link velocity increments were collected. The velocity increments are the time integral of the unit's compensated acceleration measurements taken over 0.01 second intervals. The Inertia-Link velocity increments are described in the documentation as metres per second per 0.01 seconds. In fact, they are provided as 9.80665 metres per second per 0.01 seconds (or g * s).

To help identify zero speed, the bias-corrected velocity increment magnitude values were lightly filtered (a simple moving average of five samples). Figure 7 is a plot of a representative 5-second sample of the bias-corrected velocity increment magnitude values (calculated using the method described above) as well as the filtered values.

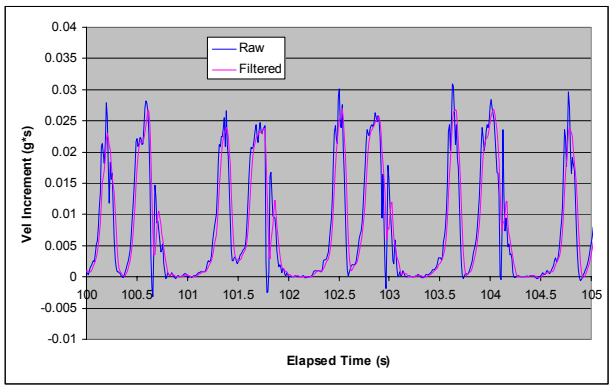

Figure 7: Inertia-Link - Example Velocity Increments

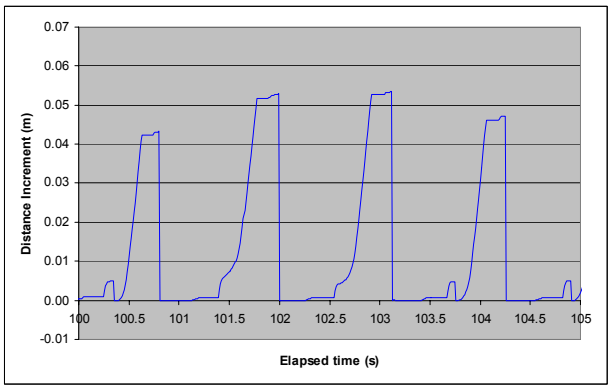

Figure 8: Inertia-Link - Example Distance Increments

An inspection of the data indicated that zero speed could be assumed when the filtered, corrected magnitude values were less than 0.0005 g*s (about 0.005 m/s). Whenever the speed was assumed non-zero, the distance increment (distance travelled per 0.01 seconds) was estimated by summing the filtered, corrected velocity increment magnitude values. Whenever the speed was assumed zero, the distance increment was set to zero. Figure 8 is a plot of a representative 5-second sample of these distance increment values. Every cycle covers about 2 metres (e.g. 0.05 metres per 0.01 seconds times 0.4 seconds). As an example, the cycle ending at 102 seconds covered a distance of 1.88 metres. To complete the distance calculation, all distance increments were summed.

By carefully inspecting the IMU data, it was determined that the output of the "pitch" gyro could be used to eliminate the decelerations at the end of a step cycle. This is a rough approximation that, as a by-product, eliminates motion at the start of a step, as the foot rotates about the toe, and before the toe lifts off the ground.

The distance results for the run are shown in Figure 9: the raw IMU distance is shown in blue line, the pitch-corrected in pink and a BU-353 GPS reference in yellow. Note that after 375 seconds, the GPS reference became unreliable as the system entered a steel-shelled building: later results were not plotted.

As expected, the dead reckoning-like Inertia-Link distance drifts away from the GPS reference distance with time, the distance computed using raw IMU velocity increment magnitudes at a higher rate than the pitch-filtered IMU distance. The errors in the former increase at a fairly steady pace, reaching 230 metres by the end of the run. The errors in the latter reach their maximum of about 40 metres at 150 seconds, ending at an insignificant 5 metres.

These results indicate a good deal of promise in techniques that use step-detection and zero velocity updates using a boot-mounted IMU.

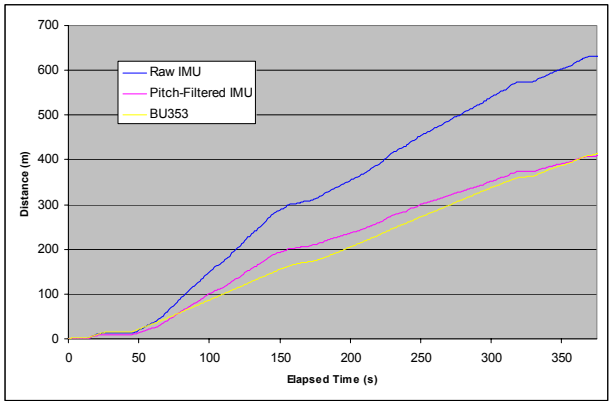

Figure 9: Inertia-Link - Distance Travelled vs GPS

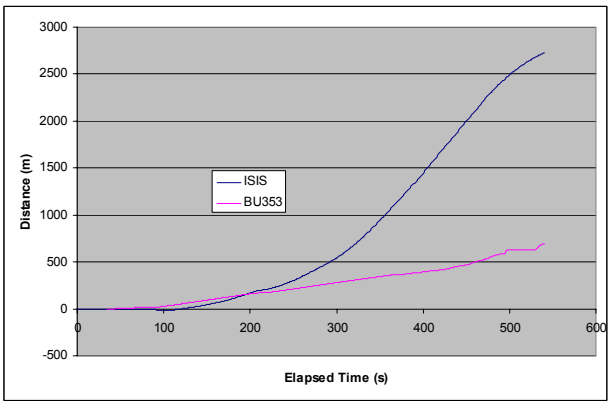

Figure 10: ISIS - Distance Travelled vs GPS

3.2.3 ISIS IMU

The ISIS IMU was mounted on the RPN-P platform. This location has some advantages in a standard Kalman filter navigation solution but has major disadvantages when attempting to integrate the outputs independently. The primary disadvantage is the lack of the frequent, identifiable zero velocity intervals that were critical to limiting the velocity error growth and thus position error growth rates in the Inertia-Link. Figure 10 is a plot of distances computed by integrating ISIS-provided accelerations (along with the corresponding BU-353 GPS reference). With uncontrolled IMU speed error growth, distance errors grow at a much higher rate.

3.3 Speed Sensor and Compass Data

When GPS data is degraded or unavailable, other information must be used to limit the growth rate of the IMU navigation errors. In the RPN-P, data from a radar speed sensor and a digital compass were collected. To judge the effectiveness of these data in the IMU aiding in an integrating Kalman filter, they were combined in an independent dead reckoning system. The process and results are described below.

Data from the Honeywell HMR3500 digital compass requires no processing. This unit contains three magnetometers and three accelerometers that are used to provide roll and pitch as well as azimuth, in almost any orientation.

The KMH DRS1000 radar speed sensor requires more attention. It is mounted on the RPN-P frontpack, facing forward, tilted 30º downward relative to the RPN-P platform (see Figure 1). It outputs pulse counts that are proportional to the measured line-of-sight speed. To get the required horizontal speed in metres per second, scale factor and pitch corrections have to be applied to the raw data. Since the RPN-P platform is not perfectly level while being worn, the tilt of the platform is required. The HMR3500, which is also mounted on the RPN-P platform, provides platform pitch that can be used with the fixed 30º downward pitch of the sensor relative to the platform to correct the DRS1000 speed to the horizontal.

After the DRS1000 speed data had been transformed to the horizontal and rescaled, they were combined with HMR3500 data to generate a horizontal track. At each DRS1000 measurement epoch, an east and north displacement was computed. Then the individual displacements were summed to generate a dead reckoning track. An example is provided in Figure 11. The leftmost plot (with all data uncorrected) shows an apparent rotation in the dead reckoned track relative to the GPS reference.

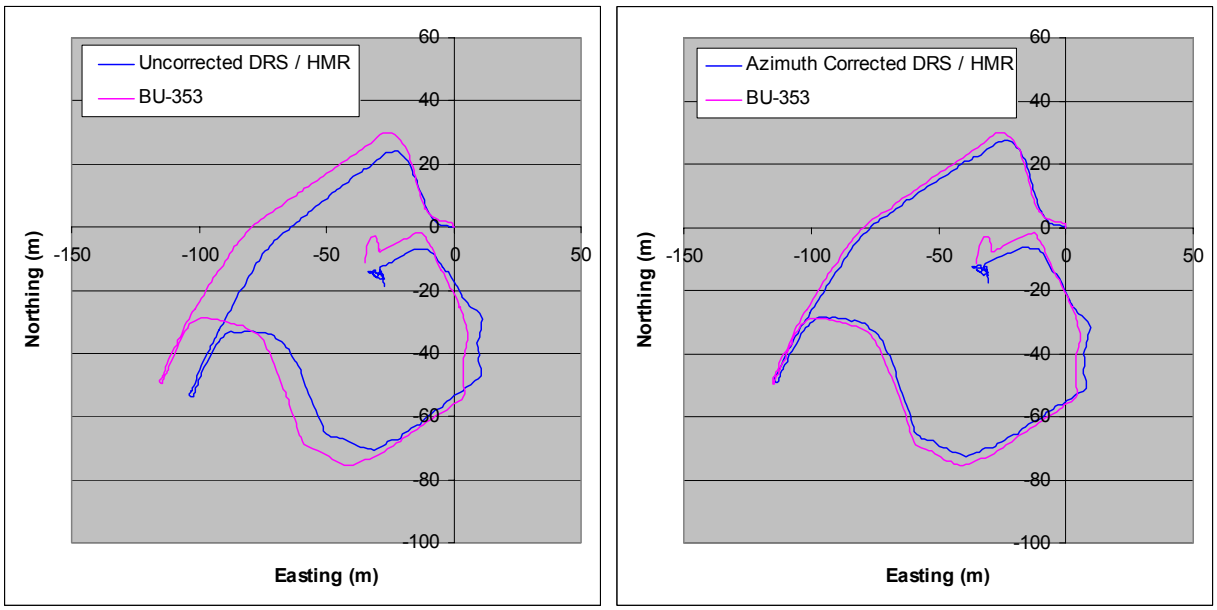

Figure 11: DRS1000 / HMR3500 Dead Reckoned Tracks

To check the HMR3500 azimuth, it was compared to a GPS reference from the BU-353. A comparison of the HMR3500 azimuth and the BU-353 Course-Over-Ground showed an average difference of 4.3 degrees (with a standard deviation of 10.6 degrees). The large standard deviation is due to azimuth periodicity resulting from hip motion as the user walks.

The bias can be attributed to compass calibration errors, errors in the magnetic variation, a compass misalignment relative to the RPN-P platform, or to a misalignment of the RPN-P relative to the direction of travel. Given that the compass was calibrated just before testing began (and reported a good result) and that the compass was mounted on the edge of the RPN-P platform (where it could be accurately aligned), it can be assumed that the 4.3 degree azimuth / COG difference is due to variation error or a misalignment of the RPN-P platform relative to the direction of travel. The RPN-P is strapped on using shoulder and hip straps. With no practical way to perfectly align the RPN-P with the direction of travel, a misalignment of several degrees is not unexpected.

The rightmost plot in Figure 11 shows the track after the HMR3500 azimuth was corrected by the estimated 4.3 degree bias. The results are remarkably accurate, with maximum errors over the 350-second span of the run of no more than 6 metres.

A speed sensor and a compass integrated into a Kalman filter integrated navigation system have the potential to provide valuable aiding information in the absence of GPS. However, there are some caveats that must be understood. These caveats relate to the significant physical differences between the RPN-P and a practical handheld navigator.

First of all, the radar speed sensor is too large for a practical system. It is possible that a more compact sensor can be found (maybe ultrasonic), but there remains the question of finding a location to mount the sensor. In theory, if its orientation relative to the user's track was known, if the ground was within range of the sensor, and if the orientation was not perpendicular to the track, the speed-over-ground (SOG) could be computed. The speed sensor mounted on the RPN-P was useful because it had a fixed orientation relative to the user and could be used to determine the user's SOG. A speed sensor mounted on a handheld device would have limited value as it could be pointed in any direction at any time – it could be pointed away from the ground or at an orientation that measures too small a component of SOG. So perhaps the speed sensor could be mounted apart from the handheld device, on a more stable platform.

Wherever the speed sensor is mounted, its orientation is required. A sophisticated compass like the HMR3500 mounted rigidly with the speed sensor could provide the complete orientation. Or the speed sensor could receive orientation information from a co-located IMU. Since the IMU needs the compass for initialisation and error control and the speed sensor needs its orientation, it would seem logical to mount the IMU, the compass, and the speed sensor all on the same rigid body. What would be an appropriate location? Maybe they could be mounted on the user's belt. If all three sensors were small enough, perhaps the user's boot or helmet could be used.

4.0 SUMMARY AND FUTURE WORK

To date, a "breadboard" prototype has been assembled consisting of a MEMS IMU, military and civilian GPS receivers, a digital compass, and a radar speed sensor. A second, boot-mounted MEMS IMU is included in the sensor suite. An Ultra-Mobile PC is included on the breadboard for data collection.

Several brief field trials have been carried out in different GPS environments. Data collected during these trials have been analysed and results presented. In urban canyon-like locations on the DRDC Ottawa campus, the high-sensitivity BU-353 civilian GPS receiver tends to perform better than a keyed DAGR military receiver. Nominal differences were measured in a more benign GPS environment. Steps were reliably detected using the boot-mounted IMU and step lengths could be accurately estimated.

The speed sensor and compass were combined into a simple dead reckoning system that was found to be quite accurate. This leads to valid questions regarding the need for the IMU(s), and given speed and heading sensors of adequate accuracy, there may be applications where such a dead reckoning system could be used. However, finding appropriate sensors (especially a speed sensor) may be difficult. Perhaps a boot-mounted IMU can serve as the "speed" sensor in a kind of hybrid system. Another limitation of a compass / speed sensor dead reckoning system is compass disturbance due to magnetic interference. Since gyros are immune to such effects, an IMU co-located with the compass can be used to detect magnetic anomalies that are affecting the compass.

Some of the other outstanding questions that will be examined:

- Keeping in mind lever arm effects, what combinations of IMUs at what locations are best? For example, if an appropriate speed sensor can be found, the boot-mounted IMU may be redundant. Can a boot-mounted IMU replace the speed sensor? Perhaps, the boot-mounted IMU can be used for navigation, eliminating the need for a central unit.

- What effect does body shielding have on GPS antenna performance? Will the antenna have to be mounted above the user (on a helmet, for example)? Lever arm effects will be an important consideration.
- The tests run for this paper showed that a modern civilian receiver performed better than a keyed military receiver, suffering perhaps from the long military acquisition lead times. Are two GPS receivers needed? A high-sensitivity civilian unit for optimum performance when multipath and signal attenuation errors predominate, and a keyed military receiver when jamming is present? With luck, there will be a receiver available in the near future that will work well in both environments.

To help address these questions, the Navigation Group at DRDC Ottawa has begun the work of incorporating all RPN-P into its existing Kalman filter navigation software (EGIM). This involves software modification and sensor error modelling. Part of this effort will include step detection and step length estimation using a boot-mounted IMU.

The search for better sensors will be an ongoing process. Initially, sensor performance will be the primary criterion. As development proceeds, and the list of possible configurations is narrowed to a small group, sensor form factor considerations will gain prominence in the selection process. This will inevitably be an iterative process as changes to smaller packaging affects performance. There is also a high likelihood that new sensors will added be to the existing suite, For example, reference [1] describes the possibility of using information from MANET (Mobile Ad Hoc Network) radios to aid an integrated personal navigator.

The ultimate goal - to have a robust navigator that operates adequately well in difficult GPS environments where satellite signals are attenuated, blocked, reflected, jammed etc. and has the potential to be operated as a handheld, personal device – is elusive but remains the objective.

5.0 REFERENCES

[1] Paul Labbé, Louise Lamont, Ying Ge and Dale Arden, “Cooperative blue-force tracking (BFT) and shared situation awareness (SA) in complex terrains”, in these proceedings.