MONITORING OF GLOBAL NAVIGATION SATELLITE SYSTEMS

P. Bernedo*, T.J. Martin-Mur, J.M. Dow, C. Garcia Salinas*

ESA/European Space Operations Centre
Darmstadt, Germany
*GMV at ESOC

ABSTRACT

ESOC has been operating a GPS Tracking and Data Analysis Facility (GPS-TDAF) since June 1992, using it for its participation in the International GPS Service for Geodynamics. The GPS-TDAF operates GPS receivers located at ground stations in Kiruna, Kourou, Malindi, Maspalomas, Perth and Villafranca. Data from these stations is downloaded daily to the control centre in ESOC, where it is validated, formatted and transmitted to other IGS (International Service for Geodynamics) centres. The operations are automated, and the data have so far been available the day after the measurements are taken.

The existing infrastructure of receivers and communication lines has been upgraded to support automatic real time operations of the GPS-TDAF. The real time GPS-TDAF will evolve into a real time monitor for Global Navigation Satellite Systems (GNSS) providing two main services: support for ESA missions that will use GPS for navigation, and support the development of European satellite based navigation systems.

INTRODUCTION

Global Navigation Satellite Systems (GNSS) are increasingly used for navigation of spacecraft and they are being proposed as the future primary means of navigation for safety critical applications like aircraft navigation. Most uses of GPS for spacecraft navigation have been done in postprocessing mode for analysis of the mission's data. In future missions GNSS data will be used for the navigation and control of the spacecraft. Errors in the GNSS signals could affect the success of those missions, if they are not detected and a warning is not given within a short time. The Real Time Monitoring of GNSS is an evolution of existing infrastructure to support future needs of ESA in the field of GPS-based navigation.

The existing GPS Tracking and Data Analysis Facility (GPS-TDAF) in ESOC performs the following functions:

- Control and monitoring of remote stations.
- Retrieval of data from remote stations.
- Preprocessing of data from remote stations.
- Delivery of preprocessed data to other IGS (International GPS Service for Geodynamics) centres.
- Processing of data.
- Delivery of processed products to other IGS centres.

The GPS-TDAF operates automatically using existing workstations and communication links. The links are data packet networks where available or the public telephone network. The existing GPS-TDAF infrastructure has a proven record of performance and reliability.

The extension of the GPS-TDAF to real time activities is a natural evolution of the system. The following principles guided the design of the real time facility:

1. Reuse of existing hardware.
2. Reuse of existing communication links.
3. Reuse of existing software.
4. Uninterrupted support for on-going operations.
5. Maximum compatibility with the current system.

Existing hardware and communication links fully support the required features for a real time extension. No need has been identified that is not covered by existing hardware. The existing software can be partially reused for some functions. Entirely new functions required the development of new specific software.

It is vital not to disrupt the existing operations of the GPS-TDAF. Data collected from ESA's stations are used by other IGS centres and from many other users accessing the IGS data. These data are updated daily implying that interruptions of the current service may not last longer than one day. Once the system has been deployed the previous operating procedures of the GPS-TDAF must still be allowed. This is only possible if the data brought by the GPS-TDAF from the remote stations is also available from the real time extension.

SYSTEM OVERVIEW

The system uses the existing hardware and communication links. Sun workstations operating under Solaris are used in the control centre in ESOC. Remote stations have a MiniRogue/TurboRogue GPS receiver with dual frequency capability and a personal computer running MS-DOS. Remote station computers are connected to modems for the appropriate type of data link: either a telephone voice line or a packet switched data network. The remote stations operate controlled by the control centre: they only transmit data when commanded by the control centre.

The hardware used by the existing system is able to support the real time extension. However, no claims are made that it may be the optimal selection if the real time project were to start from scratch. The existing computing hardware for the central control facility in ESOC uses Sun SPARC workstations. These

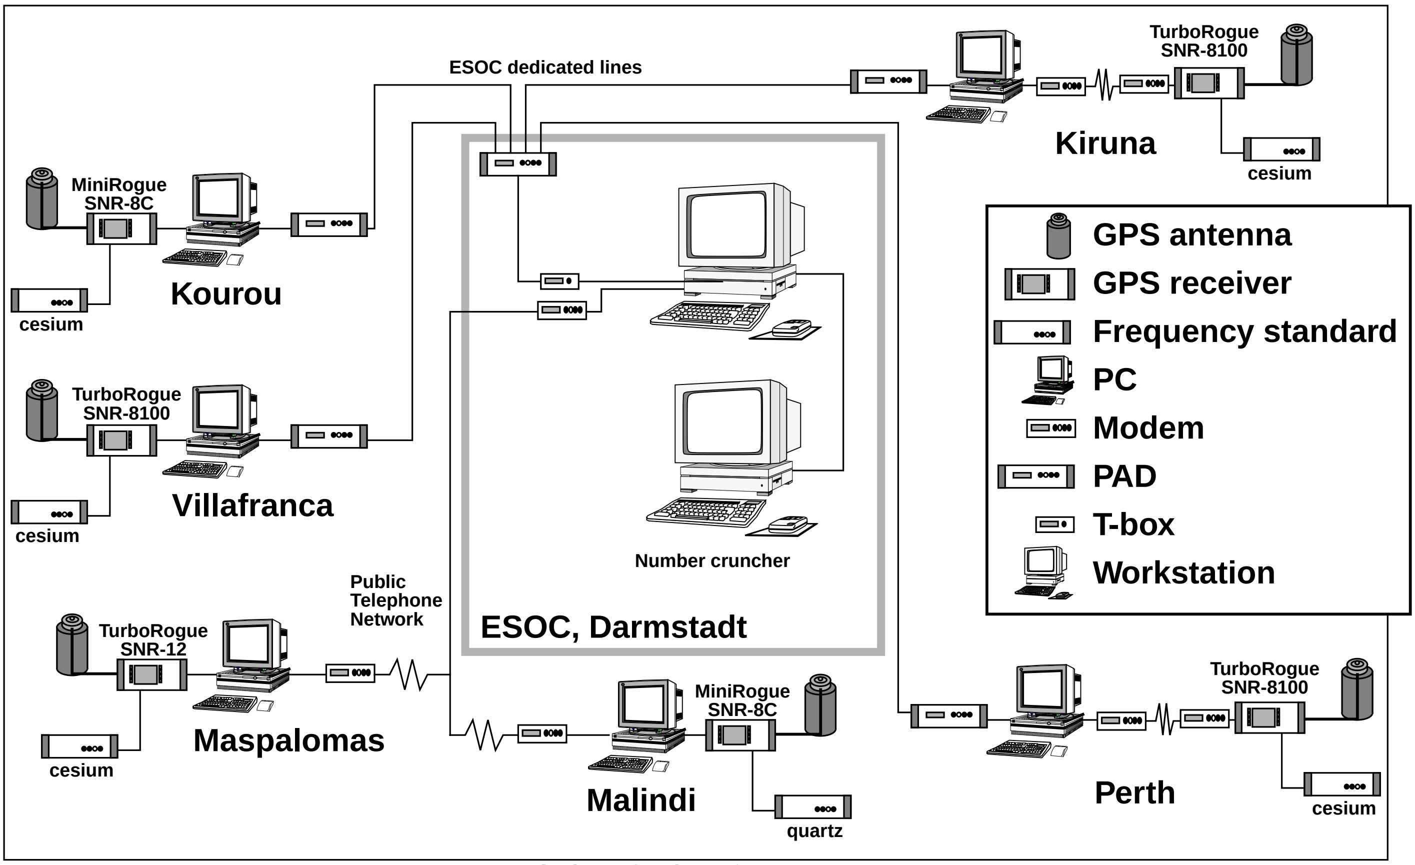

GPS-TDAF Overview

workstations are used for other tasks in addition to the GPS-TDAF, that require the processing power of the workstation. Given the low load imposed by the data transfer activities, almost any type of machine can support those activities. The choice of the platform is guided by other considerations like availability, reliability and cost.

The GPS receivers already installed in the remote stations were designed for geodetic measurements like those done for the IGS. The command set and performances are optimised for storage of measurements and postprocessing. Other types of receivers are now available that have been designed for other applications, like generation of real time differential corrections. Such a receiver would increase the capabilities of the current system, allowing better control of real time capabilities.

New software has been developed for both the remote stations and the central control facility. The software has been written in the C++ language, with sharing of source code for modules that perform similar functions in both the remote stations and the central facility.

RELIABLE DATA LINKS

All information exchanged between the central control facility and the remote stations is divided into messages. The format of the messages has been specifically tailored to the needs and environment of the real time GPS-TDAF extension, although its guiding principles can be found in other formats for real time broadcast of GPS corrections like RTCA/DO-217 or RTCM SC-104 standards. The existing standards do not cover such things like transmission of remote station status or download of recorded data.

Each message has markers to signal the start and end of the message. These markers allow a listener of the data stream to synchronize with the message sequence, even though the transmitted data are all binary data. The contents of the message are protected against undetected transmission errors with a 32 bit CRC code that rejects corrupted messages. Each message has a message type number defining its contents and purpose.

A set of messages supporting the current application have been defined and implemented, with plenty of room for addition of new message types. The defined messages cover data transmission, control of the remote station and reporting of configuration parameters.

It is worthwhile to note that the selected message format rejects virtually all messages that are corrupted by the transmission lines. The primary property of the format is that it ensures that no wrong messages will be accepted, although messages may get lost. There is no protocol to make sure that a given message has been received. Instead, messages are sent even without knowing if the computer at the other end is still working. When acknowledgement of reception of a message is required, a matching “answer” message type has been defined. It is the responsibility of the machine receiving a message to react to a message type that requires acknowledgement. Supplying an answer message was not considered appropriate for this system. Data either arrives in time to the system, or it is not available. When supporting real time monitoring of the health of the constellation, we cannot delay the start of each processing step while waiting for retransmission of a message: it would take precious seconds from the time to alarm. This strategy is possible due to the high reliability of the data links: only very few messages are lost, while most of them get through the system without loss.

REMOTE STATIONS

The remote stations have a MiniRogue/TurboRogue receiver that is continuously tracking the visible GPS satellites and sending the stream of measurements to the controlling computer through a RS-232C serial link. The computer is a standard personal computer running the real time GPS-TDAF software.

Measurements from the receiver are decoded, checked for errors and stored in the computer. Both measurements and navigation data (ephemeris and almanacs) are stored in the computer’s hard disk. When the disk is full, the oldest measurements are replaced with the new ones. This results in always having the most recent set of data that fits into the hard disk. It is assumed that the data will either be read in real time by the control facility, or that they will be downloaded from the station routinely. The rate at which measurements are stored can be set from the control facility.

Measurements that have passed the error detection test are sent immediately to the control facility (these are the real time data). The current navigation data is sent periodically to the control centre. The control centre can enable and disable the transmission of the real time data. It can also control the period of time between two transmissions of navigation data. The purpose of the transmission of all current navigation data is twofold: (1) it makes the system robust against losses of messages, because the navigation data will be repeated, and (2) allows any processes within the control centre to acquire the navigation message without querying the remote station.

The described operation is the real time operation, where measurements are sent immediately. The remote station also accepts commands to download recorded data. This feature was introduced to support the existing IGS operations. It allows the retrieval of data for periods when the control centre was not listening to the remote station. The control centre may send commands to the remote station to retrieve the list of recorded data that is available and can select the download of specific hours. Downloading can proceed in parallel with the real time operations: the downloaded data is divided into small blocks and the blocks are sent when the communications line is idle. The real time data always takes precedence over the downloaded data, suppressing the risk of degrading the performance of the real time operations.

One of the missions of the remote station software is to translate the measurements coming from the receiver into a common format used throughout the system. The receivers currently in use send their measurements in the TurboAscii format. This format is specific to the receiver used (Turbo Rogue) and is optimised for postprocessing operations. The format used by the real time extension of the TDAF is designed to efficiently support the transmission of data in real time. The advantage of using a format independent of the receiver is that receiver types can be changed, by modifying only the software running in the remote station. This software is already built with a modular structure and the translation from the receiver format to an internal format is the first step in the processing of measurements. The remaining part of the remote station software is not affected by changes in the receiver.

DATA LINKS

Data are transmitted from the remote station to the control centre and commands are sent back to the remote station. All these messages flow through data links. The data links used in the system are of two types. Most ESA stations are linked to ESOC through data links that offer very reliable and fast communications. These links are the preferred option where available. Each remote station generates a low rate of data (about 300 bytes per second when sending measurements at one hertz), and can be accommodated in the data link together with other users. The result is that the real time monitoring data can be transmitted through the unused capacity of existing links, without requiring the allocation of any additional capacity for this purpose.

In stations where no packet switched data links are available, the public telephone network is used with data modems. The telephone network is not as reliable as the packet switched network and has higher operating costs.

Although these two types of data links are currently used, the system can use any data link that is available (eg. Internet), just by replacing the current modems with the appropriate modem for the new link.

The data links are one of the critical points of the system. The system depends on the capability of the data links to deliver a continuous stream of data without significant errors. Errors in the line will corrupt messages that will be rejected by the central facility. If the message carried measurements, they will be lost, unless the central control facility requests them to be sent. Usually there is not enough time to request the measurement and to receive it before the measurement for the next epoch arrives. Therefore, lost measurements are not retransmitted. Instead the next measurement is used.

When recording incoming data for postprocessing, not all measurements are required. It is usually enough to record a measurement every 10 or 30 seconds, even though the control centre receives a measurement every second. In this case, loss of one measurement has no impact on the recorded data.

The data links impose a latency in the system. Data links have different and varying latencies. Telephone connections have very small latencies, only the time required for the whole message to be transmitted by the modem, plus up to a few tenths of a second for the case of a telephone line passing through a geostationary satellite. Packet switched networks, may route each packet through a different path, and the actual time required to reach the destination varies. This time is usually short, but there are no guarantees that exceptional long delays will not occur.

Considering this variation in the latency, we must be prepared to expect the arrival of data from different stations at different times. There is no GPS receiver installed in the control facility, preventing the software in the control centre from having an accurate timing source. It can only guess the real GPS time, by estimating the latency of the transmission links. One method to do it is to send a request to the remote station and wait for the answer. The interval between the query and the answer is approximately twice the one way latency plus the latency in the station, which is negligible.

One issue that must be dealt with when processing data from multiple stations is when to consider that the data should have arrived. If we wait for all stations to send their data, it may happen that one of them is down, or that its data link has very long latency. Waiting for that station's data will degrade the performance of the system. It is better to set a maximum time interval between the arrival of the first station's data and the start of the processing for this epoch. Any data for this epoch that arrives after that interval will not be considered in the computations. The actual length of the waiting interval is a balance of the probability of missing some data because it arrives too late and the maximum time to raise an alarm if one of the GPS satellites goes wrong.

CONTROL CENTRE

The control centre, located in ESOC, controls the remote stations, manages the incoming data and distributes the real time measurements to users within ESOC. The control centre is implemented as a process running on a Sun workstation connected to modems for the data links. Each workstation in the current configuration can be attached to two modems (due to physical limitations). More than two remote stations can be accommodated just by using other workstations. This is not a problem because all workstations are linked by a local area network and all clients access the control centre through the network.

The control centre acts as a bridge between remote stations and users of the data: these users just see an incoming stream of measurements coming from a socket and are not concerned about the receiver type or the data link. Most processing of the measurements is carried out by client programs that connect to the control centre.

The control centre provides mechanisms for commanding the remote stations, downloading past measurements from the remote stations and retrieving the configuration of the remote stations. All these functions are accessible from a suitable user interface. The functions available from the user interface mimic the capabilities of the remote stations. For instance, the user may set the period between broadcasts of navigation data, stop the transmission of real time data or command a download. All users may view the parameters of the remote stations, but modifying them requires a password. Connecting to the real time data stream does not require any password, and the data is available to any computer connected to the ESOC LAN.

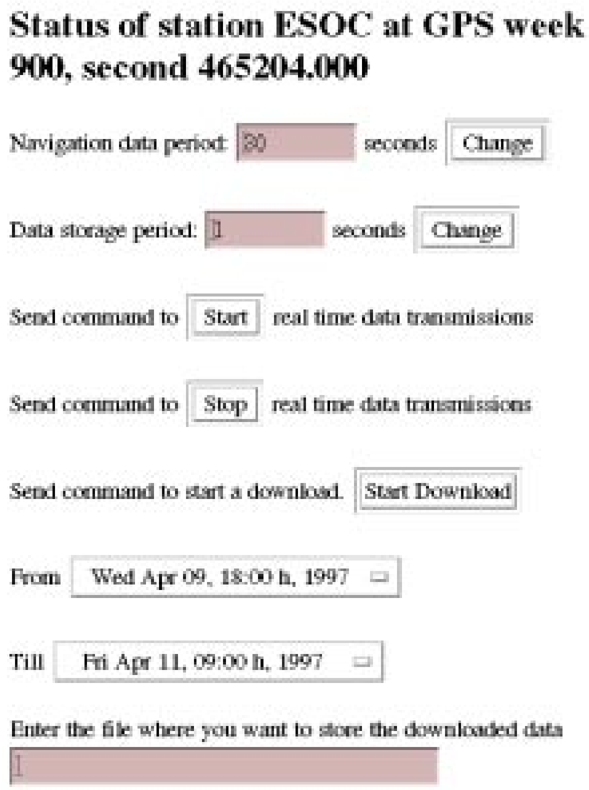

Control of remote station

The other interface for the control centre is the client interface. Client means any program that needs access to the real time data stream. The actual interface used is based on sockets, making it portable to virtually every network and operating system. Multiple clients can log to the control centre and receive the stream of real time data. This effectively turns the control centre into a system for broadcasting real time measurements to any interested workstations in the LAN. It is important to notice that the connection of users to the control facility has no impact on the remote stations themselves. All users get the data on a read only basis: a user can not modify the remote station's parameters to fit his own needs. This is only possible to authorized users with a password.

The client programs do the processing of the data brought by the control centre. They connect to the source of real time data from the control centre and process the data according to their purpose. Given that the system can support an unlimited number of clients, it is always possible to extend the processing performed, just by adding a new client. This architecture simplifies the maintenance of the system, because the control centre provides only the real time data infrastructure with little processing beyond formatting and storage.

Several types of client programs are always active. The simplest client is one that stores the received data in a file. The data will then be available for postprocessing without requiring a download. This file will also reflect the actual data used in the computation, which may be less than the data stored in the remote station due to loss of messages in the data link. Another useful client is one that keeps a log of all events and non-measurement messages that flow in the system.

Other clients of the control centre are the single station monitor (evaluation of the health of the constellation using data from one station only) and the multiple station monitor.

SINGLE STATION MONITORING

The single station monitor is a client that can evaluate the health of the GPS constellation using data from a single remote station. Other clients use data from several stations to conduct more comprehensive tests, but require data from all stations. The single station monitor was the first monitor to be implemented.

Anomalies in the GPS constellation can be detected in different processing steps. The first step is to process each channel's data to combine the observations in L1 and L2 to remove ionospheric effects. This is possible due to the availability of both L2 carrier phase and cross-correlated P (Y) code pseudorange. The resulting pseudorange and carrier are passed to a filter to remove the noise of the pseudorange. The filter's gain is limited to allow it to forget very old measurements. The filter allows us to detect big cycle slips and remove some observation noise for the following processing steps.

The next step is the decoding and verification of broadcast ephemerides. The ephemeris are compared against the predicted orbits generated within the framework of the contribution of ESOC to IGS. Predicted orbits are generated daily and provide a very accurate orbit that can be used as the reference of truth. Small discrepancies are usual between the predicted and broadcasted ephemeris due to the limited accuracy of the representation of the broadcasted ephemeris and the introduction of SA in the orbits. Larger discrepancies may be due to errors in the navigation message or a manoeuvre of the GPS satellite. Both events are detected by this test. Although a manoeuvre is not an error (the new ephemeris will reflect the correct orbit after the manoeuvre), it may have an impact if some users were using data from a predicted ephemeris, and it must be detected.

Once the ephemerides have been checked, they can be used to compute the expected pseudorange and compare it to the measured pseudorange. The result is a series of "raw corrections" to the pseudoranges that include the real correction and the clock bias of the remote station. Comparing these "raw corrections" with the previous "raw corrections" gives the "raw rate of change" for the correction of each satellite. The "raw rate of change" has the clock drift of the station plus the actual rate of change of the clock of the satellite.

Satellites whose corrections are out of predefined bounds are detected by taking the largest group of satellites whose "raw corrections" are not separated more than the predefined limit. All satellites belonging to this group are declared to have their corrections within reasonable bounds. All others are declared unhealthy due to excessive correction. A similar procedure is repeated for the "raw rate of change of corrections", and for the Doppler measurements.

Satellites that pass all checks (corrections, rates and doppler) are declared healthy. Otherwise an excessive correction or excessive rate of change of correction is declared.

The average of the "raw corrections" of the healthy satellites provides an estimation of the clock bias of the station, with the error due to selective availability added to it. Similarly, the average change of "raw corrections" between one epoch and the next (for satellites that are healthy in both epochs and have the same ephemeris) is an estimation of the change in the clock bias, with the addition of the effect of SA. The clock of the station is modelled as a bias and a drift, with the observables being the estimations computed with the average correction and the average change of corrections.

The estimated clock bias is then removed from each satellite's correction, yielding the final clock bias for the satellite. The estimated clock drift is removed from each satellite's rate of change of correction, yielding the final value of the rate of change. Errors in the clocks of the remote station can be detected by comparing the clock bias and clock drift against predefined bounds.

MULTIPLE STATION MONITORING

Using multiple stations will improve the constellation health evaluation capabilities of the system. There are several advantages in using multiple stations:

- The area of coverage is larger.
- It is possible to separate in real time orbit errors from clock errors.
- More accurate evaluation of errors is possible due to redundancies in the measurements.
- The reliability of the system is improved.
- More exhaustive verifications of the health of the monitoring stations and satellite constellation are possible.

There are no differences in the data distribution system described above, when used with several remote stations. The addition of remote stations requires just the addition of a new modem and a new computing process in the control centre. This is a minor change in the system. The main change is the installation of a new client program that takes data from central control facility and uses the multiple stations.

The multiple station monitoring capability is not yet implemented at the time of writing, although it will be part of the real time monitoring system. We describe here the major features of the multiple station monitoring system.

The first steps in the processing of multiple stations are similar to those described in the single station case. The input data for each channel of each station is filtered to remove observation noise. The same procedures used by the single station can be used in the multiple station case without modifications, except that instead of having a single indicator of the health of each satellite, there will be several indicators, one from each station tracking the satellite.

Several algorithms can be applied when multiple station data are available. It is possible to separate errors in the broadcast ephemeris from errors due to satellite clocks. Geometrical algorithms (snapshot methods), can determine the orbit and clock of each satellite that is observed by the network of stations. The computed orbit and clock are not very accurate when using pure geometrical methods, but these algorithms do not rely on past measurements and react immediately to anomalies in clocks or to satellite manoeuvres. Even if the computed orbits and clocks are not very accurate, the errors in pseudorange using the computed orbits and clocks may be quite small in some parts of the coverage area.

Kalman filtering is a combination of geometrical methods (for updating) and dynamic methods (for propagation). Depending on the accuracy of the propagation model the Kalman filter's performances can vary from pure geometric to dynamic methods. Real time constraints require that the propagation models are simple enough to be computed in real time, limiting the accuracy that can be achieved with the Kalman filter. It is however the best method to obtain real time accurate ephemeris and clocks. The Kalman filter technique can use partial measurements, for instance observations from only two stations, and still get good results.

When using multiple stations for real time orbit and clock determination, the issue of station clock errors becomes important. Batch methods benefit from the dynamic propagation of the orbit and can filter out measurement noise and compensate station clock errors. Real time methods can not count on the availability of fully dynamic methods (they take too long to compute) and using accurate clocks in the stations increases the performance of real time algorithms.

Kalman filters can include in the propagation model not only the dynamics of the satellite, but also the propagation of the clock offsets of the ground stations. This is possible only if atomic clocks are installed in the stations; quartz clocks are too unstable to get any useful propagation between measurements.

The set of remote stations has almost full coverage of the GPS constellation, but there are still some areas where the satellites can not be observed. Even in the covered area, the number of stations simultaneously seeing a satellite may be quite low, perhaps just one or two stations. In those cases we can revert to single station processing or we can use the Kalman filter. At the beginning the Kalman filter will just propagate the precomputed orbit, but the observation data will gradually refine the orbit. This approach provides a smooth transition from the case of observations from one station to many stations. This will be the normal case when satellites rise in the European sky after a period of no tracking above the Pacific ocean.

A similar issue is the effect of station failures on the computed orbit and clocks. The Kalman filter method allows graceful degradation of results when stations (or the links to the stations) suddenly fail.

CONCLUSION

The existing GPS-TDAF has been upgraded for real time operations. The real time system is fully compatible with the existing GPS-TDAF and was implemented reusing most of the existing resources. The resulting system opens a new field of applications for users that rely on GPS for their navigation.