

EGNOS CPF Check Set – Stand Alone Performance

R. Wolf, U. Rossbach , W. Werner
IfEN Gesellschaft für Satellitennavigation mbH (IfEN GmbH)
S. Lannelongue, J. Michelon
Alcatel Space Industries
P. Thomas
Logica Ltd.

BIOGRAPHIES

Robert Wolf holds a degree in Aerospace Engineering from the Technical University of Munich and a Ph.D. from University FAF Munich. Since 1995 he worked in the field of hybrid GPS / INS navigation and later on precise orbit estimation and integrity determination. He joined IfEN GmbH in 1999 as a systems engineer. Between 1999 and 2002 he was working on the development and optimisation of the EGNOS CPF Check Set algorithms, since 2001 as the technical responsible.

Wolfgang Werner holds a diploma in Computer Science from the University of Technology in Munich and a Ph.D. from the University FAF Munich. Since 1999 he is Technical Director of IfEN GmbH. Having been responsible for the EGNOS Independent Check Set algorithm development until end of 2000, he is currently working for Galileo integrity.

Udo Rossbach holds a Ph.D. in surveying engineering from the University of Federal Armed Forces Munich, Germany. Since 1999 he is working as a systems engineer with IfEN GmbH. There he is involved in the development of GPS/GLONASS/EGNOS/Galileo algorithms and software.

Stephane Lannelongue received in 1996 a M. Sc. engineer degree in electronics and signal processing from the Ecole Nationale Supérieure d'Ingénieurs Electriciens de Grenoble (ENSIEG/France). From 1997 to 2000 he worked in the field of satellite navigation at the European Space Agency (ESA/ ESTEC). He joined Alcatel Space Industries in 2000 as a system engineer on EGNOS and Galileo.

Julien Michelon graduated from the Ecole Nationale de l'Aéronautique et de l'Esace, ENSAE in Toulouse, France, in 1995, and holds a pre-thesis degree in Applied Mathematics from the Université Paul Sabatier in Toulouse, France. He joined the EGNOS CPF team at Alcatel Space Industries in 2001 to collaborate on the Check Set algorithms prototyping and porting to

operational implementations, and is currently involved in the RELY EGNOS project.

Pete Thomas is a principal consultant in Space and Defence Division of Logica UK and is the algorithm implementation team leader. He has more than 15 years experience of the implementation of scientifically complex systems within the space industry. Pete has worked on the satellite control and data processing facilities for several major European satellite programmes, mainly on earth observation applications. He has a PhD in geodesy from the University of Nottingham.

ABSTRACT

One of the main objectives of the European Geostationary Navigation Overlay System (EGNOS) is the protection of the user by offering not only position accuracy but also service integrity. For this reason the Central Processing Facility (CPF) is divided into two sets: the Processing Set (PS) and the Independent Check Set (CS). While the task of the PS is to generate ionosphere as well as ephemeris and clock corrections with corresponding error bounds, the CS is in charge of verifying the correctness of this information by checking the whole set of NOF (Navigation Overlay Frame) messages independently. The main checks, satellite correction (UDRE) Check and ionospheric correction (GIVE) Check, follow total different strategies. While the UDRE check applies the corrections and monitors the residual error, the GIVE check re-computes the correction and compares it to what has been broadcast to the user.

The CS is neither allowed to use any data produced by the Processing Set, nor to use the same input, i.e. measurement data. Many processing steps done by the PS have to be redone by the Check Set, using different algorithms to avoid common failure modes. This includes independent estimation of the RIMS inter-frequency biases and clock errors as well as estimating a Check Set internal ionospheric model.

The actual integrity checks are just the last (and not even most demanding) step in a long processing chain. The observations used to evaluate the NOF content can themselves be affected by errors and therefore have to be carefully monitored before delivered to the checks.

A major challenge during the Check Set Prototype development by IfEN has been the requirement of providing high integrity and being at the same time robust against data gaps, corrupted observations and other feared events at the input. Therefore, an observation has to pass many detection and exclusion stages from reception of raw data to a fully validated smoothed pseudo range. This includes mechanisms for the single line of sight, as well as cross checks for common events per satellite and per RIMS.

The Check Set is intended to ensure integrity of the EGNOS message. Although the current design can be tuned to provide nearly arbitrary levels of integrity & continuity, the price would be a significant reduction of availability. Who wants a system preserving integrity, but being not available most of the time? The current tuning tries to balance all requirements to provide just enough integrity & continuity to meet requirements, while keeping availability as high as possible.

The Check Set Prototype has not only undergone intensive Stand Alone Testing, but also test together with the prototype Processing Set on the so called integrated platform. These test have manly been conducted by the CPF contractor Alcatel. Those results (included in this paper) demonstrate the real added value of the Check Set in case of Processing Set failure.

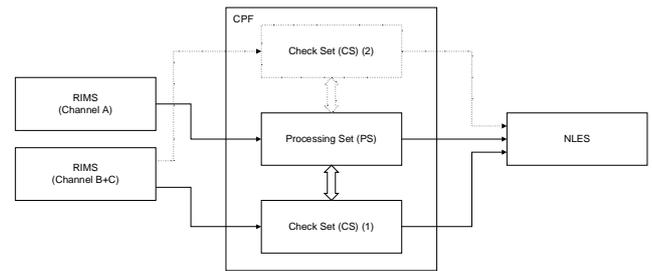
While IfEN has finished algorithm design and prototyping work, Logica is currently implementing the operational EGNOS Check Set. The last part of the paper describes the industrialisation of the complex integrity algorithms, which have been independently specified and prototyped, into a real-time, high availability service. The design of the system, in particular the partitioning of the overlapping augmentation footprints into semi-independent processing 'lanes' will be described. The performance constraints on the system will be outlined, and the operational solution will be described.

The validation of the safety-critical Check Set will be based on scenarios that collectively cover all nominal and foreseen extreme operating conditions. An outline of the validation tests will be given, and initial performance results for the operational software on the target hardware will be presented.

CHECK SET HIGH LEVEL OVERVIEW

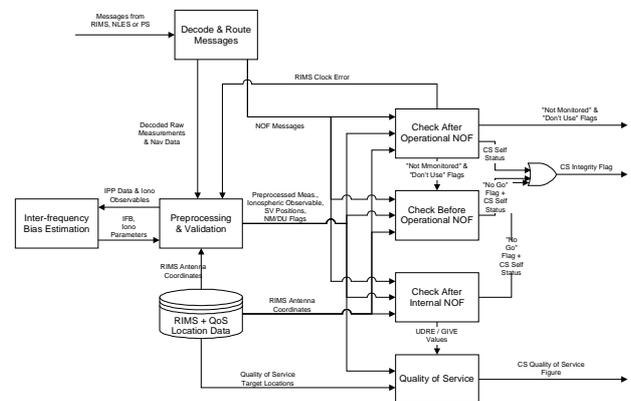
Figure 1 shows the CFP Check Set in the context of the EGNOS system. The CPF receives data from two independent sets of RIMS (Ranging and Integrity Monitoring Stations). One set is fed to the Processing Set (PS), one to the Check Set (CS). The output of PS and CS is send to the NLES (Navigation Land Earth Station) performing the uplink to the EGNOS GEO.

Figure 1: CPF CS context



The first level of Check Set functionality is depicted in Figure 2. The CS consists basically of 5 major components Check After (CKA), Check Before (CKB), Pre-Processing (PRE), Quality of Service (QOS). The CKB performs logical checks on the messages sent to the NLES, while the CKA performs statistical checks on the messages, which have already been received by the user. Both share a common Pre-Processing.

Figure 2: Check Set functional overview



The processing chain PRE-CKA is therefore mainly responsible for the performance of the CS in terms of continuity and integrity.

Note, that there exist up to three Check After and three Check Before of the operational NOF lanes. Each one is dedicated to one specific GEO. However, the CS is designed to be expandable to up to 6 operational GEO lanes.

The two main core checks in the Check After are the UDRE check and the GIVE check. In the following sections, the performance of these two will be shown.

ANALYSIS METHOD FOR STAND ALONE TESTS

To obtain statistically significant numbers for false alarm rate and missed detection probabilities, all types of errors in magnitude and direction have to occur. Basically there are two possibilities, manipulating the observations or manipulating the corrections. As the observations are subject to filtering processes, the errors have been triggered by manipulating the NOF corrections directly

A false alarm in the UDRE check is defined as an alarm raised although the true SREW is smaller than the HMI limit defined as

$$HMI := SREW > 1.62 \times UDRE$$

which is corresponding to a limit of 5.33σ . As the alarm curve will be a smooth, double curved figure it can not step from zero at the left of the HMI limit to 100% at the right of the limit. Therefore, a standardized weighting function has been used to derive the probability of false alarm.

$$pPfa = \frac{\sum_{i=0}^{15} \left(2 \cdot W_i \cdot \frac{N_{FA,i}}{N_{FA,i} + N_{CN,i}} \right)}{\sum_{i=0}^{15} W_i}$$

FA = False Alarm
 CN = correct non alarm

The weighting function has been derived from a Normal distribution with zero mean and a standard deviation of $1/3.29$.

$$W_i = P\left(\frac{SREW}{UDRE} < \frac{i+1}{10}\right) - P\left(\frac{SREW}{UDRE} < \frac{i}{10}\right)$$

The factor 2 comes from the fact that the SREW is always positive. The standard deviation has been chosen $1/3.29$, to ensure that the sum of weights is 0.999 which corresponds to the definitions of UDRE / GIVE as error bounds with a 99.9% probability. Note that we do not at all expect the error to be Gaussian! This choice is just driven by the philosophy that an alarm on an error close to zero ($i=0$) should have a higher weight (i.e. be more punished) than an alarm in class 15, which is very close to the HMI limit.

The probability of missed detection has been computed without a weighting function, just by dividing the number of non-alarmed HMI by the total number of HMI events

$$Pmd = \frac{\sum_{i=16}^{30} N_{MD_i}}{\sum_{i=16}^{30} N_{MD_i} + \sum_{i=16}^{30} N_{CD_i}}$$

MD = missed detection
 CD = correct detection

Note that this approach is quit pessimistic, because it weights an HMI of a few millimetres as high as an HMI of a several meters!

CHECK SET STAND ALONE PERFORMANCE

UDRE Check

shows the alarm curve of the UDRE Check under "nominal" conditions, i.e. nominal with respect to observation noise, cycle slips etc. but nevertheless non-nominal with respect to errors to detect!

UDREs generated by the Input Driver were based on the "nominal PS behaviour" assumption wrt. UDRE profile, i.e. the UDRE assumed for a given observation geometry. The corrections were corrupted in a random way to produce over as well as under bounds.

Figure 5: UDRE Check Alarm Curve

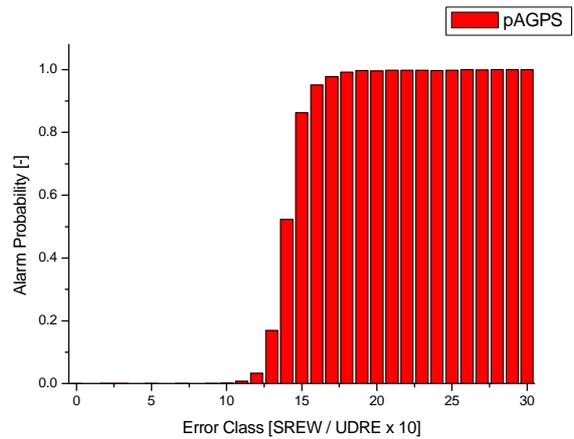


Figure 6: P(MD) vs. UDRE for dif. num. of Measurements

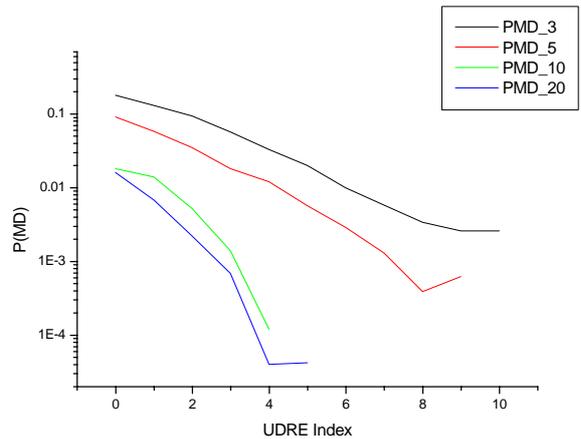


Figure 5 has been created from the event table previously shown. For example it reveals that an error, which is between 1.62 and 1.7 times the broadcast UDRE will be alarmed with a probability of 95%.

The alarm curve is the representation of the performance taking no assumptions, but the requirements in terms of missed detection and false alarm are given as single

numbers. With the equations given, this yields mean values for P(FA) and P(MD)

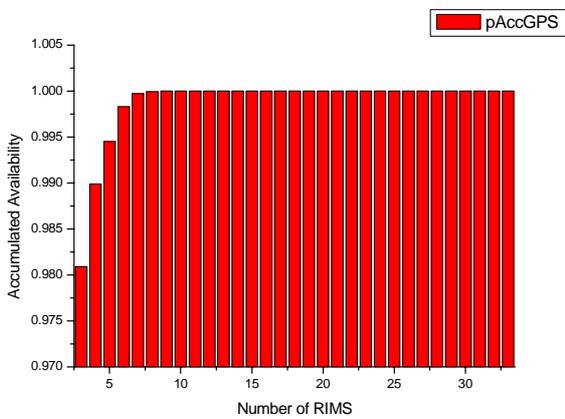
$$P(FA) = 7.16E-05$$

$$P(MD) = 4.75E-03$$

for the „nominal“ UDRE profile. **Figure 6** shows the alarm probability for a given UDRE depending on number of observations.

The Check Set needs a steep alarm curve to provide good integrity (small P(MD)) and continuity (small P(FA)), but the availability has to be kept as high as possible. **Figure 7** indicates that for satellites being visible to 3 or more RIMS, an availability of 98% is guaranteed.

Figure 7: Accumulated Availability of UDRE Check vs. Number of Visible RIMS



GIVE Check

The GIVE Check is the second big core check in the Check Set. I has to verify that the ionospheric error remaining after application of the broadcast GIVD is bounded by the $1.62 \times GIVE$ with a probability of 99.99999%. **Figure 8** shows the alarm curve of the GIVE Check.

Figure 8: GIVE Check Alarm Curve

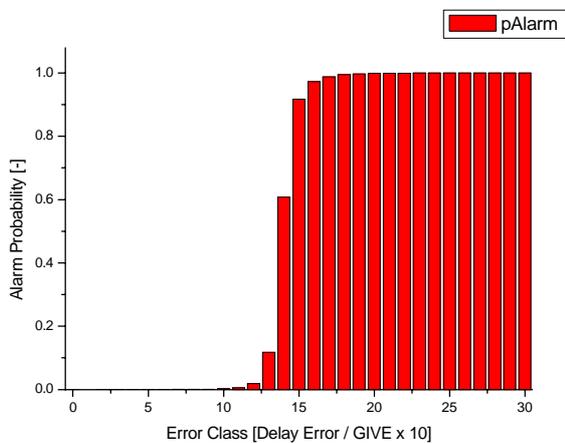


Figure 9: Geographical Distribution of GIVE Missed Detection Performance

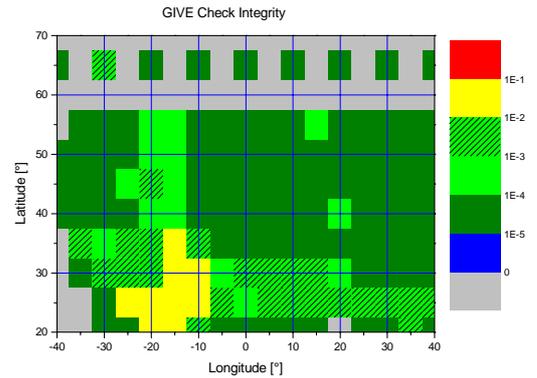


Figure 10: Geographical Distribution of GIVE False Alarm Performance

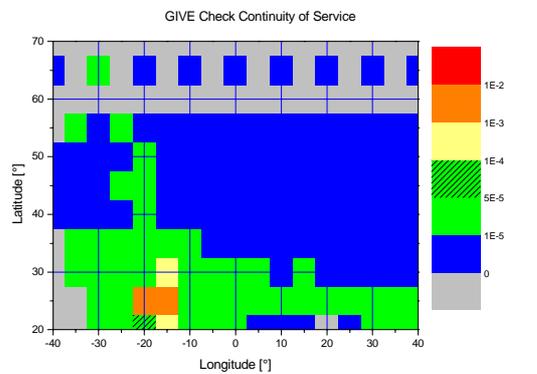
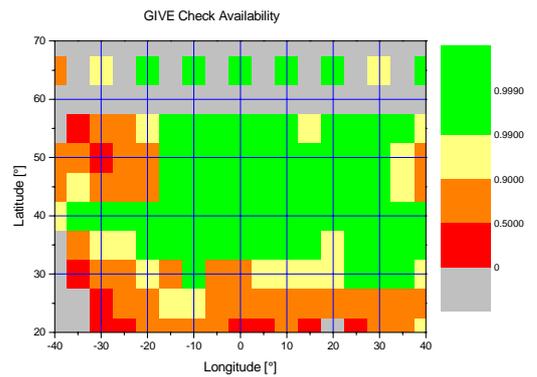


Figure 11: Geographical Distribution of GIVE check Availability



The weighted false alarm and (not weighted) missed detection numbers for the CIVE Check are

$$P(FA) = 1.4E-05$$

$$P(MD) = 1.4E-03$$

These again are mean numbers over the entire service area (ECAC). Especially for the GIVE Check, the spatial distribution is very interesting. For this reason, missed detection probability (**Figure 9**), false alarm rate (**Figure 10**) and availability (**Figure 11**) has been computed per IGP and displayed as colour coded plot over the longitude / latitude grid of ECAC. The performance over the service area is very good, except one small area over the Azores. This is slightly degraded due to high ionospheric gradient near the equator. Note that the plots below assume the Processing set to perform perfectly and compute very low GIVE values. On the integrated platform it showed that the Processing Set has a problem with the same region too and produces therefore higher GIVES. The overall CPF integrity and continuity can therefore be kept even there.

CHECK SET PERFORMANCE ON INTEGRATED PLATFORM

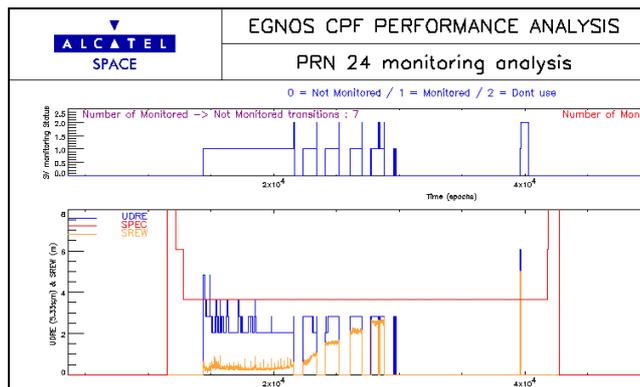
The purpose of the Check Set is to be integrated in the EGNOS CPF together with CPF Processing Set. Once integrated it is difficult to characterise Check Set performance since in order to see Check Set impact it is necessary that the Processing fails to meet its specifications. This does not happen in nominal conditions, however during CPF experimentation extreme scenario (so called feared event scenarios) were played allowing to identify some limitation of the Processing set. Such limitations are under corrections are not representative of the operational processing set. However those intermediate results allow to identify the real added value of the Check Set in case of Processing Set failure.

Satellite L1/L2 HW bias Feared event

One good example are the results observed when simulating L1/L2 HW biases jumps at satellite level. The upper curve of **Figure 12** represents the status of the satellite at CPF output (0=NM, 1=USE, 2=DU). The second curve represents the true SREW (bright) and the UDRE scaled at 5 sigma (dark). L1/L2 HW biases jump are simulated as a four steps function of 50 cm. First what can be seen is that each jump is detected by the Check Set and a DONT USE is raised at CPF output on the faulty satellite. At Check Set level, when such event is detected satellite filters are reinitialised and the L1/L2 HW bias is estimated again. That is why the satellite remains Not Monitored for some minutes after the alarm.

In the mean time it can be seen that the real SREW is slowly degrading. This is due to the fact the Processing Set has not detected anything and still work with a no longer valid bias estimation. After the four step the SREW still degrades and gets close to the UDRE threshold. Before that the threshold is hit the Check Set detects successfully the problem and flag the satellite DO NOT USE again.

Figure 12: CPF behaviour vs SV L1/L2 HW bias jump

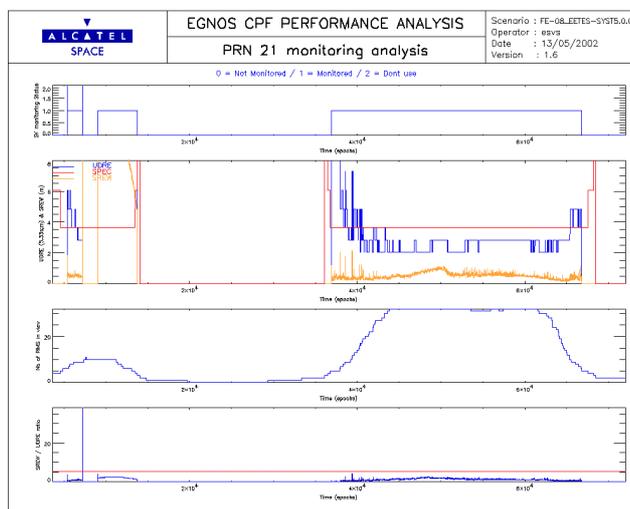


This demonstrates that Check Set performs a good estimation of the SREW on an integrated platform and can detect inconsistency between SREW and UDRE

Satellite Clock Bias Feared Event

Figure 14 shows CPF behaviour when confronted to a satellite clock jump. It focuses on PRN21 affected by the jump. It can be seen that when the jump occurs a “UDRE (scaled at 5.33 sigma) not bounding the SREW” arises instantaneously. This is quite normal since the information used at user level were computed by the CPF before the feared event and expecting a nominal behaviour of the satellite.

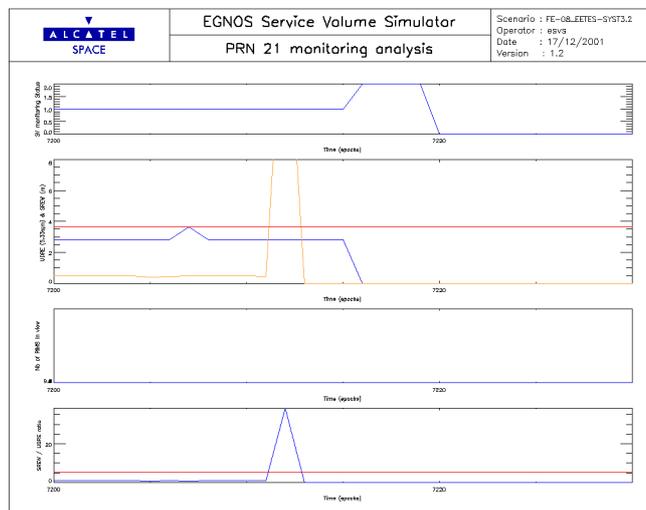
Figure 13: CPF behaviour vs SV Clock bias jump



When zoomed on the event (see **Figure 14**, it can be seen that the UDRE does not bound the SREW only during 4 epochs. This delay is the time simulated on the prototype platform between the generation of the message and the reception by the user. Therefore what can be said is that the UDRE fails in bounding the SREW for a few epoch but this is correctly detected by the CPF and in particular by the Check Set. This allows an alarm to be broadcast

within the Time To Alarm. This is illustrated by the SV status (upper curve of the figure). The non bounding situation does therefore not lead to a UDRE integrity failure.

Figure 14: Zoom on CPF behaviour vs SV Clock bias jump



INDUSTRIALISATION OF THE EGNOS CHECK SET

1 Purpose

Logica's role is the implementation of the Check Set algorithms on the target hardware respecting the EGNOS real-time constraints. This work is being performed under a subcontract to Alcatel Space Industries. The approach to the implementation and validation of the algorithms on the operational Check Set is that the scientific content of the algorithms has been already validated during the algorithm development phase; the validation of the software implementation can include a comparison of the behaviour of the operational software against the prototype results used to validate the algorithm specifications.

2 Implementation solution

The software being developed for the CPF Check Set consists of the following items:

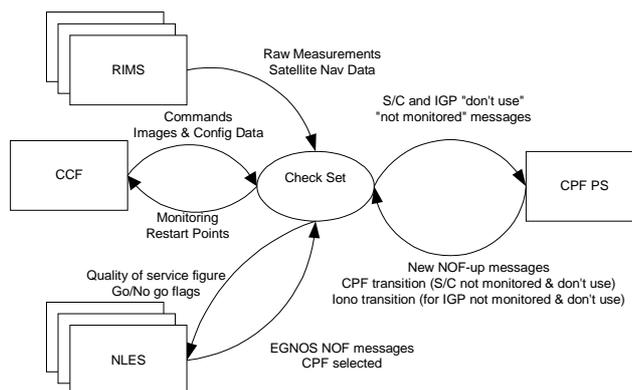
- Boot Loader to perform initial power up tests when the set is energised, and perform minimal functions to support the transition to the application software
- Real-time monitoring and control (RTMC) Application Software to manage all external interfaces and to route information between the Integrity Software component, the CPF Processing Set and other external facilities
- Integrity Software consisting of the four algorithm functions Preprocessing and Validation, Check After, Check Before, Quality of Service. In addition to the five algorithms,

there is also a central control class 'Integrity Control', which controls the timing of the algorithms and the interfaces with the RTMC Slave Application

The operational software is implemented almost entirely in ANSIC.

The external interfaces to the Check Set are illustrated in the figure below.

Figure 15: Check Set External Interfaces



The interface to the RIMS is one-way. Raw RIMS measurements and satellite navigation messages received are transmitted each second. All data are received by the RTMC subset, which rejects messages with invalid CRC, invalid timestamps or from RIMS that are not configured for this Check Set (e.g. RIMS A messages). Filtering of messages for content and consistency takes place within the Preprocessing and Validation function.

The interface between the Check Set and the CPF Processing Set (PS) being checked is two-way. Clearly, part of the interface must be the new NOF-up messages from the CPF PS that is to be checked. The Check Set reports to the PS on the status of the satellites and the ionospheric grid points (IGPs) used in computing the NOF. Each of these may be flagged as either "Not monitored" or "Don't use"

- Not monitored means that the CS has not received enough data to verify that a navigation message computed using the satellite or grid point in question meets the integrity requirements.
- Don't use means that the CS has detected a problem with either the satellite signal or one of the satellite or ionospheric corrections.

In both cases, the PS needs to be informed in order to incorporate the status information in its next generated NOF message.

The new NOF-up message will be checked for format and consistency with the previously transmitted messages.

The CS will send a “Go” flag to the NLES if the checks are passed and a “No go” check if the NOF-up message is unacceptable.

It is important to distinguish between the “Don’t use”, “Not monitored” and “No go” flags. The former two indicate that the CS has correctly detected and responded to external problems. The latter indicates that there is some error in the CPF (PS or CS) itself and, if the CPF generating it is the operational one, the NLES responds to receipt of this flag by initiating a switchover to another CPF.

The Check Set also generates a “Quality of Service” metric that is sent to the NLES. The purpose of this is to allow the NLES to decide which of the currently active CPFs is providing the best service and thus should be used operationally. The NLES informs the CPFs which of them is selected in its Feedback message.

The remaining part of the CS/NLES interface is the delta-NOF; that is the four EGNOS NOF messages that have been accepted but not yet received by the RIMS as part of the Signal in Space (SIS).

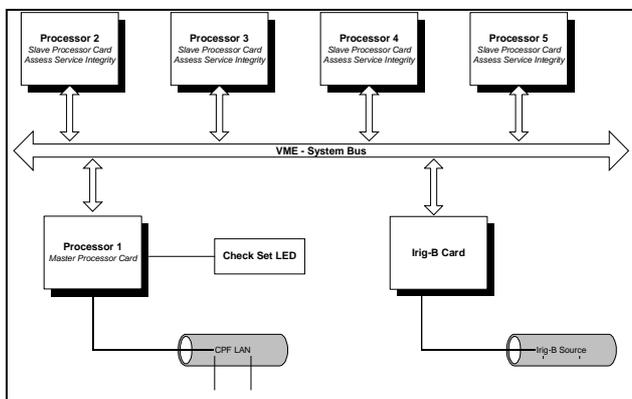
3 Hardware

This section describes the baseline hardware architecture of CPF Check Set. Each Check Set consists of the following items:

- One 19” VME enclosure with integrated power supply, containing a 17 slot VME64 extensions backplane.
- Five Concurrent (VP PMC P34) Intel Pentium 850 MHz CPU boards, each with 128 Mb DRAM, 16 Mb flash memory, 2 Mb of battery backed SRAM, EEPROM, Watchdog timer, and Ethernet.
- One Truetime VME SG2 Irig-B timing card.

This hardware configuration is illustrated below:

Figure 16: Check Set Hardware Configuration



All 5 processor cards, and the Irig-B card are located in the VME Chassis and connected to the VME backplane. The Master Processor Card is located in slot 1, and so takes on the role of VME Master (i.e. it has overall control over the VME backplane). This card is also connected to the CPF LAN via the Ethernet port on the front of the card, and the internal wiring of the chassis connects it to the I/O panel on the back of the chassis.

The IRIG-B Card is located in slot 11 on the backplane. This card provides the time signal for the set, and is connected to the external source via internal wiring from the front of the card to the I/O panel on the back of the chassis.

Slots 3, 5, 7 and 9 are occupied by the 4 Slave Processor Cards, which perform all of their communications across the VME backplane.

The remaining odd numbered slots in the chassis are reserved for future expansion, and will be filled with air baffle modules initially. Even numbered slots are reserved for cooling. This is required in order to ensure correct air flow through the chassis.

4 Key considerations

EGNOS uses the concept of processing lanes. There are two types of lane, operational and internal. An operational lane consists of the Processing Set processing required to generate a NOF for a single GEO, the Check Set processing required to check it and the dedicated NLES to uplink it. The check after processing for the lane in any instance of a Check Set accepts the NOF which was actually uplinked regardless of which of the CPFs actually generated it. The check before processing checks the locally generated NOF to ensure that it could be uplinked, even if the NLES will not do so.

The internal lane is an additional NOF which is not intended for uplink; there is therefore no check before in this lane. It is generated by the Processing Set and sent only to the local Check Set where it is subject to the check after processing. If the internal NOF passes this check, the Check Set then goes on to use the internal lane check after processing data output to compute a quality of service metric which is sent to the NLES to allow it to determine which CPF is producing the best results. If it fails, a message is sent to the NLES indicating that it should not use this CPF.

The Internal Lane running on a Check Set in Backup mode sets “NoGo” flag whenever the Check After marks a satellite or IGP as “Don’t Use”.

Each of the four processing lanes (i.e. three operational plus one internal) will be implemented on a separate Slave Processor Cards. At initialisation, the Integrity software on the card determines which lane it is to process from its location with respect to the Master Processor Card. In this configuration, the Pre&Val subset will be repeated on

each of the processor card. Whilst this may be inefficient from a processing load perspective it has a number of key advantages:

- It restricts the transfer of the large amounts of data between the Pre&Val and Check Before/After subsets to tasks on the same processor card thus minimising the inter-board data transfer over the VME backplane
- It leads to a broadly symmetrical architecture without the need for different control logic on each Slave Processor Card
- It simplifies the expansion mechanism when extra GEOs are added.

Thus, the three processor cards allocated to the operational lanes (i.e. each responsible for controlling and monitoring a single EGNOS GEO) will execute Preprocessing and Validation, Check After and Check Before.

The processing required of the internal processing lane is very slightly different in that it processes RIMS data for the superset of all EGNOS GEO and computes the quality of service figure. There is also no need to run the Check Before algorithm on the internal lane, since there is no NOF up message to Check. One consequence of this architecture is that the Quality of Service figure computed on the internal lane must be transmitted to the Operational Lanes for inclusion in the cyclic data messages sent to each NLES. This can easily be achieved using the RTMC VME message transmission software. Other than that, there is no inter-lane communications required during the critical algorithm processing cycle.

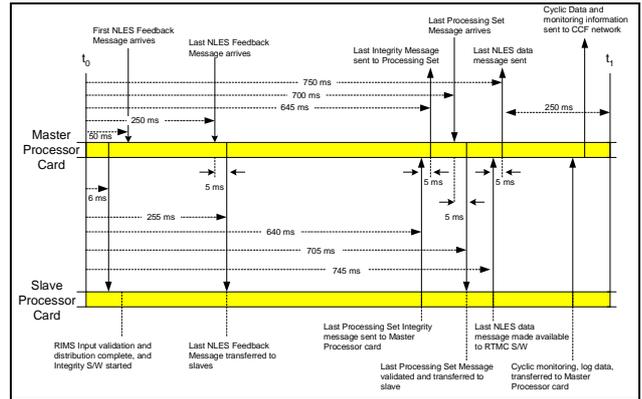
All of the sub-functions discussed thus far run on the standard once-per-second EGNOS cycle, under real-time constraints. However, the Interfrequency Bias Estimation (IFB) algorithm, which is notionally part of Pre&Val, is a computationally intensive process that does not need to run at the once-per-second rate. It is run at a frequency determined by a constant in the set of configuration parameters. The nominal frequency is once every 180s.

The IFB algorithm is run on the Master Processor Card and runs when the Master Processor Card would otherwise be idle. This gives it approximately 750ms of execution time per 1 second cycle.

5 Performance

The Check Set ‘Top of Cycle’, t_0 , is set to be at a configurable offset between 0 and 1000 ms from the GPS second boundary. This offset is initially expected to be set to around 745 ms.

Figure 17: Check Set Timing Diagram



EGNOS operates on a 1 second processing cycle with a data lifecycle of 4 seconds. The latest data will be generated at the CPF (including a check within the CS), sent to the NLES, uplinked to the GEOs, broadcast to the RIMS (and users), transmitted to the CPF, and checked again within the CS.

The key performance drivers can be derived from the above figure.

1. The “Not monitored/do not use” flags must be available not later than t_0+640 ms. This means that all data decoding, preprocessing and validation, and check after must be completed by this time.
2. The new NOF-up message is received from the processing set by t_0+705 ms. The “Go/no go” flags must be available not later than t_0+745 msec. This means that only 40 ms is available to the check before function.

Only 745 ms of each cycle is available to complete the Integrity Software because of the requirement to send the completed check set results (the Quality of Service value and “Go/No go” flag for each GEO) to the NLES by that time. The remaining 255 ms of each cycle will be used for RTMC activities.

The projected length of time taken to perform each major function is summarised in the table below.

Table 2: Check Set Real-time Performance

Processing Step	Elapsed Time (ms)
RIMS Data Transfer	6
RIMS preprocessing	10
Preprocessing and Validation	46
Check After	280
Check Before	15
Quality of Service	9

The time taken to complete the Check After is:

RIMS Data Transfer	6 ms
RIMS Preprocessing	10 ms
Preprocessing and Validation	46 ms
Check After Downlink Algorithm	280 ms
Total	342 ms

This is comfortably within the 640 ms limit. Also, the 15 ms execution time for check before is well within the 40 ms window available for that function.

These figures assume the current hardware baseline (i.e. 850 MHz cards), and have been computed by estimating the number of operations to be performed at each processing step (depending on number of RIMS, satellites etc.). The number of operations has then been converted to elapsed time using a conversion formula derived from experiments made on the target hardware.

CONCLUSION

The EGNOS CPF Check Set has proven to meet requirements in any sense. It is a mature product and a real added value to the overall EGNOS system.

ACKNOWLEDGMENTS

The algorithms and results presented in this paper have been developed within the frame of the EGNOS project, funded by the European Space Agency (ESA).

REFERENCES

- [1] *Minimum Operational Performance Standards For Global Positioning System/Wide Area Augmentation System Airborne Equipment*, RTCA Document No. RTCA/DO-229B, Prepared by RTCA Special Committee 159 (RTCA SC-159), RTCA, Inc., Washington, D.C., October 6, 1999.
- [2] *Algorithms and Performance of the EGNOS CPF Independent Check Set*, ION GPS-2000, Salt Lake City, Utah, Sept. 19-22, 2000.