Mars Express TT&C Contingency Operations – Hoping for the Best, Preparing for the Worst

Thomas Ormston[1]
VEGA Space GmbH, Darmstadt, Germany

James Godfrey[2]
European Space Agency, Darmstadt, Germany

and

Olivier Reboud[3]
SciSys, Darmstadt, Germany

The Telemetry, Tracking and Commanding (TT&C) system on any spacecraft is critical to operations as it is the only thing linking us to the spacecraft from ground. Therefore the response to a failure in the TT&C system especially requires a high level of autonomy and preparation both on board and on ground given that any failure could temporarily or permanently sever the link to the spacecraft and cut-off any opportunity to recover or troubleshoot from ground.

There are two different kinds of failure that can occur in any TT&C system – a failure on the receive chain (uplink) or a failure on the transmit chain (downlink). The response and operational approach for each kind of failure are very different and require a customized solution for each. In both cases the traditional method of analyzing telemetry, determining a response and recovering with telecommands cannot work completely.

In the case of an uplink failure it can be assumed that the spacecraft would keep transmitting normally and so diagnostics can be performed on the telemetry being downlinked. However, without on-board autonomy there is no way that the spacecraft can be commanded out of this failed situation if the receiver is not working. If a downlink failure occurs then the opposite is true – it can be assumed that commanding is functioning normally and recovery telecommands can be sent in the blind. However, there is no telemetry from which the problem can be diagnosed and the response determined. In fact there is no knowledge at all that the spacecraft is even there or listening.

Recently, as Mars Express ages, a project has been conducted to analyze, tune and augment our procedures, tools and autonomy to deal with a failure in the TT&C subsystem. This paper will detail the TT&C subsystem on Mars Express to highlight the complexity of the task. The multiple different failure modes considered will be covered to show where problems could occur. The paper will then go on to discuss what on-board autonomy is in place, or is under development, to cope with any of these problems.

[1] Mars Express Spacecraft Operations Engineer, HSO-OPM, ESOC, Robert-Bosch-Strasse 5, 64293 Darmstadt, Germany
[2] Mars Express Deputy Spacecraft Operations Manager, HSO-OPM, ESOC, Robert-Bosch-Strasse 5, 64293 Darmstadt, Germany
[3] Mars Express Spacecraft Operations Engineer, HSO-OPM, ESOC, Robert-Bosch-Strasse 5, 64293 Darmstadt, Germany

This description will show why we believe that we are better prepared than ever for a failure that we hope will never occur. With this knowledge we can demonstrate that should a unit in the Mars Express TT&C subsystem fail, we have a rapid, efficient and safe response ready to go.

I. Introduction

THIS paper will detail the protections and procedures that are in place on the Mars Express spacecraft to mitigate the impact of a failure in the Telemetry, Tracking and Commanding (TT&C) subsystem. These protections are considered particularly critical because without a working TT&C system, communication with the spacecraft is not possible and therefore ground recovery of the problem could be impacted.

In this paper, a description of the Mars Express TT&C subsystem will be covered as background to the risks that need to be mitigated. A general description will then be given to the different types of failure that can occur in the system (uplink or downlink) and the possible sources of these failures. The remainder of the paper will be dedicated to both the autonomous on-board response we have put in place to mitigate the impact of these failures and the ground procedures that have been developed to deal with any kind of failure and ensure the safe resumption of the mission.

II. The Mars Express TT&C Subsystem

The Mars Express TT&C subsystem is a dual-band redundant transponder, connected to three antennas. The following diagram shows a block representation of the subsystem.

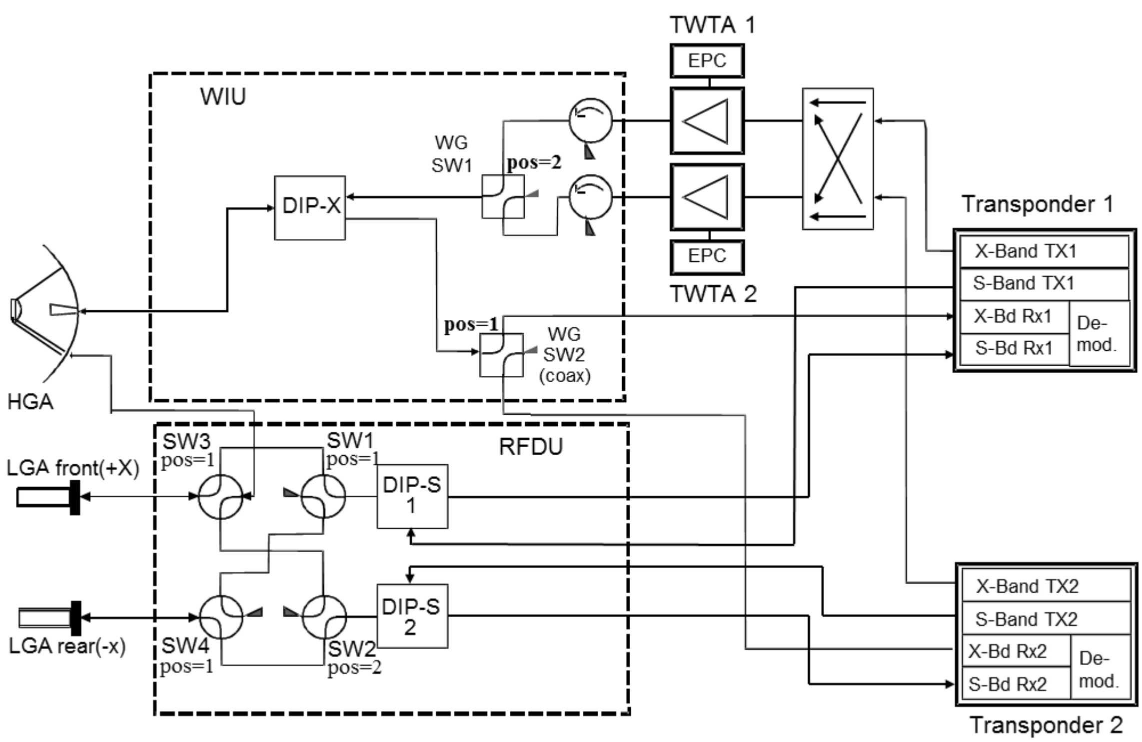

Figure 1. Block Diagram of Mars Express TT&C System

The following sub-sections describe each of the elements of the subsystem.

A. Transponders/Amplifiers

There are two identical transponders and amplifiers, to provide redundancy in the system. Each transponder consists of the following components:

1. X-Band Transmitter

Each transponder contains an X-Band transmitter that takes an incoming modulated reference signal and converts it to the Mars Express X-Band downlink frequency of 8420 MHz. The resulting X-Band signal is then passed out of the transponder, to the 3dB Hybrid Module.

The X-Band transmitter modules are self-contained, including their own power supply and phase modulator. Input information is in the form of the input signal, from the spacecraft data handling system, and a coherency signal which can be enabled or disabled by telecommand. If the coherency signal is enabled and the receiver is locked then the X-Band output frequency will be the appropriate turnaround ratio of the locked receive frequency. There is also a common TCXO module in each transponder that provides the frequency signal if coherency is disabled or no uplink signal is available.

2. S-Band Transmitter

Each transponder also contains an S-Band transmitter, transmitting at 2296 MHz, that is configured in a similar way to the X-Band transmitter. The major difference, aside from the frequency, is that the S-Band transmitters also contain a solid state amplifier that provides an amplification of the output signal, resulting in an output power of 37 dBm. This signal then requires no further amplification before transmission to an antenna. The S-Band transmitter output is passed directly to the Radio Frequency Distribution Unit (RFDU).

Aside from this difference, the S-Band transmitters are self-contained, as with the X-Band transmitters and receive the same input signals. Both take the input signal from the data handling system and so it is not possible to independently set X-Band and S-Band downlink bit rates. It is however possible to set differing ranging and telemetry modulation indices, or indeed disable the ranging and/or telemetry completely on a per-transmitter basis. It is entirely possible (and normal) to have both transmitters powered simultaneously, if the spacecraft power budget allows.

3. Dual-Band Receiver

The final part of the transponder units themselves is the dual-band receiver in each. These receivers are fed from the Waveguide Interface Unit (WIU), providing an X-Band signal, and the Radio Frequency Distribution Unit (RFDU), providing an S-Band signal. Each receiver takes the form of two receiver front ends, one X-Band, one S-Band and a common frequency receiver and demodulation stage. Each of the front ends is responsible for receiving the input frequency and then downconverting it to the intermediate frequency for further processing by the receiver. Only one front end can be selected per receiver at any one time.

Once the signal has been downconverted, the common receiver portion uses a hardware state machine to lock on the input signal, in either wideband or narrowband, depending on the strength of the incoming signal. The first stage of the lock is to detect a signal and track the incoming carrier. From this carrier a sub-carrier lock is then acquired and finally the data on this subcarrier is demodulated.

4. 3dB Hybrid Module

Both X-Band transmitters are connected to a 3dB Hybrid Module which splits the signal from each and routes each half to each of the two Travelling Wave Tube Amplifiers (TWTAs). The 3dB Hybrid Module is a passive device which ensures that both X-Band transmitters are constantly connected to both X-Band amplifiers.

5. Travelling Wave Tube Amplifier (TWTA)

There are two TWTAs, each separate units distinct from the rest of the transponder. The connection between these amplifiers and the transponder are by way of the 3dB Hybrid Module. The TWTAs have their own power supply and can be commanded on and off separately from the rest of the transmit chain.

Each TWTA provides an amplification of the input X-Band signal, resulting in an output power passing into the Waveguide Interface Unit (WIU) of 48.4 dBm. There is no explicit configuration of the TWTAs and they can be safely left on even when no X-Band transmitter “drive” is present. Only one TWTA can be connected to an antenna at any one time, the other to a load. There is a major risk to spacecraft safety if a powered TWTA were to be connected to the load. In the worst case of powering into the load, the WIU maximum qualification temperature would be reached in 44 seconds. An autonomous monitoring constantly checks the position of the switch and status of the TWTAs to ensure this does not happen.

B. Signal Distribution

The uplink and downlink signals are distributed between the three antennas and the transponders by means of two systems, one for X-Band and one for S-Band.

1. *Waveguide Interface Unit (WIU)*

The Waveguide Interface Unit (WIU) provides the interface to the High Gain Antenna for the X-band signals; the HGA is the only antenna available for X-band uplink and downlink. The WIU has two coaxial switches: Switch-1 connects the output of one of the two TWTAs to the HGA and the other to an RF load; Switch-2 connects the HGA to one of the two X-band receivers. It is the responsibility of ground commanding to ensure that the selected receiver is also configured for X-Band reception.

An X-band diplexer routes the received signal from the HGA to the WIU Switch-2 and the transmit signal from the output of Switch-1 to the HGA.

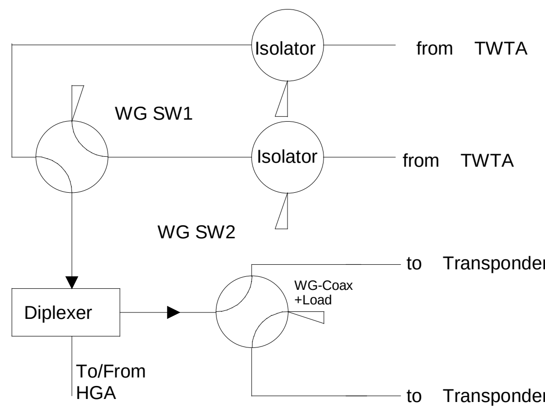

Figure 2. Mars Express WIU Block Diagram

2. *Radio Frequency Distribution Unit (RFDU)*

The Radio Frequency Distribution Unit (RFDU) performs a similar function to the WIU but for S-Band signals. In this case it handles the connection of each of the two transponders to one of the three S-Band antennas; two low gain and the high gain antenna. In order to perform this task, the RFDU has four RF switches, each of which can be moved independently from the others. The various combinations of positions for the switches allow any antenna to be connected to any transponder.

S-band diplexers are placed between the RFDU switch assembly and the two S-band transponders. The S-band uplink and downlink signals share a common route through the RFDU assembly from these diplexers to the antenna such that each S-band transponder that is connected to a given antenna for downlink signal is also connected to that antenna for uplink and vice versa.

Switches 1, 2 and 4 have RF loads connected to them, meaning that there are certain "forbidden" configurations where one of the transponders would be connected to an RF load. It is operational practice to never use these configurations, although there would be no immediate impact on spacecraft safety, even with the S-Band transmitter powered and radiating into a load.

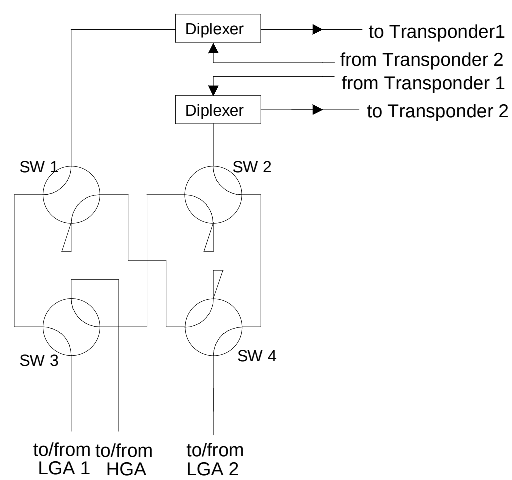

Figure 3. Mars Express RFDU Block Diagram

C. Antennas

There are three antennas on the spacecraft – one High Gain antenna and two Low Gain antennas.

1. *High Gain Antenna (HGA)*

The High Gain Antenna (HGA) is a 1.65 m diameter dual center-fed Cassegrain system consisting of a centered paraboloid main reflector and a hyperboloid dichroic sub-reflector. The antenna is fixed to the body of the spacecraft. The antenna has a 3dB half-cone angle of 0.65º in X-Band and 2.65º in S-Band and therefore must be accurately pointed to Earth by adjusting the spacecraft attitude before communications can be performed. The antenna is connected to the transponders by way of both the WIU and RFDU.

2. *Low Gain Antenna (LGA)*

The spacecraft has two Low Gain Antennas (LGAs), one on each side of the spacecraft. The two LGAs are otherwise identical – hemispherical quadrifilar antennas with a quasi-omni-directional performance. The LGAs are only designed to receive and transmit S-Band signals and so are only connected by way of the RFDU. Their principle role was to provide communications immediately after launch where spacecraft attitude could not be guaranteed. In their current role they provide a method of emergency communications with the spacecraft in the case where attitude control has been lost and the spacecraft is not pointing, or cannot point, the HGA towards Earth.

III. Possible Failure Cases

There are two overarching types of failure possible in any TT&C subsystem – a failure in the reception capability or a failure in the transmission capability. These could of course happen together, although by design the two chains are isolated in such a way as to minimize this possibility. There are also several different levels of severity of each type of failure, depending on the source of the failure.

The following will summarize some of the possible failure types in each case along with a description of the symptoms. This is not designed to be an exhaustive list but rather a guide to some examples of what needs to be dealt with when handling a failure. Double failures or failures thought highly unlikely (i.e. catastrophic mechanical failure) are also discounted in this paper, although the responses would be just as valid.

A. Uplink Chain Failure (Failure to receive a signal from ground)

Failure in the uplink chain is in many ways far more serious than a failure in the downlink chain. The major positive of this type of failure, with respect to a downlink chain failure, is that it can be assumed that data about the spacecraft status will still be received from the downlink chain (assuming the failure has only impacted the uplink). However, even if sufficient telemetry can be received to diagnose the problem then reaction will be limited by the fact that no recovery commands can be sent to the spacecraft.

In the case of Mars Express, the commands to enable the downlink chain are almost always loaded into the on-board timeline in advance, and therefore a downlink signal would be received, even in case of an uplink chain failure. In cases where the commands are not loaded in advance, or where the on-board timeline has expired, an uplink chain failure can be considerably more difficult to diagnose and distinguish from a downlink chain failure, as no downlink signal would be observed.

The chances for a failure in the Mars Express uplink chain are limited to three broad cases:

1. **Switch Failure**
   A failure in either a WIU (X-Band) or RFDU (S-Band) switch could result in the path from the antenna to the receiver being cut. This precise failure occurred on another ESA spacecraft, XMM-Newton, in October 2008. In this case a switch similar to the RFDU switches on Mars Express was stuck in the middle position and did not properly route the RF signal to any receiver. A similar failure could occur on Mars Express, although operationally we usually have one receiver connected via the WIU and the other via the RFDU to mitigate the impact of any single switch failure. Commands to change switch positions are always repeated to reduce the probability of an intermittent failure like that on XMM-Newton.
2. **Receiver Failure**
   A failure in any part of the receiver module itself would of course impact the ability of Mars Express to receive the uplinked signal. As described in the previous section, there are multiple parts to the Mars Express receiver and failure of any of the parts could completely or partially impact the ability to receive a ground station signal. Failure of the in-use front-end would completely impact the receive ability, whereas a failure later in the chain could result in partial lock being acquired but without the ability to decode bits.
3. **Configuration Failure**
   An incorrect configuration of the receiver by ground could also result in an inability to uplink to the spacecraft. In the case of selecting the wrong front end for the signal being received this could mean a complete loss of the ability to command the spacecraft. Changing the bitrate that the receiver is working on could also impact the ability to command, although by cycling through the bitrates with the ground station a link could be re-established. Wrong configuration of the switch positions in the WIU or RFDU from ground could also result in a complete failure in the ability to uplink commands to the spacecraft, in the case where a receiver is connected to an RF load.

It is these three broad groups of failures that drive the response that would be required to restore operations of the spacecraft. This may be guided from ground in case of a partial failure but the main characteristic of an uplink chain failure is that it is highly unlikely that the ground will be able to command the spacecraft.

Initial mitigation is provided on Mars Express by having two receivers permanently powered and connected to different antennas – one the prime X-Band HGA link and the other an S-Band LGA link. The latter is our emergency link in case of a prime uplink chain failure. However, the large Earth-Mars distance and low gain of the LGA mean that this link cannot always be guaranteed.

B. Downlink Chain Failure (Failure to transmit signals to ground)

Failure in the downlink chain presents an entirely different set of problems to failure in the uplink chain. In the worst case of a downlink chain failure, no signal or data from the spacecraft would be available at all. In the best case a signal may be available but without data modulated onto it. This latter case is somewhat easier to deal with because at least the spacecraft can be tracked, and indeed much information can be gained even from carrier transmission as it shows that the spacecraft is properly pointed and that commanding on-board is working.

The worst case of a complete loss of signal is considerably more difficult to deal with because it cannot be said with any certainty that the failure lies within the TT&C system. Poor spacecraft pointing, a power or commanding failure or any number of system level problems could lead to the inability to downlink a signal. However, in terms of recovery, the safest assumption is that the lack of signal reception is indeed down to a failure in the downlink chain of the spacecraft. If no progress is made on recovery then further investigation into more serious causes can be carried out.

The good news in case of a simple failure in the downlink chain is that the spacecraft should otherwise be functioning normally – indeed the ability to downlink is not critical to safety of the spacecraft as a system. In this case the receiver side of the transponder should be functioning normally and so commands to recover the anomaly can be sent in the blind.

The chances for a failure in the Mars Express downlink chain have four broad root cases:

1. **Switch Failure**
   As in the uplink case, a failure in either a WIU (X-Band) or RFDU (S-Band) switch could result in the path from the transmitter to the antenna being cut. In case of the X-Band switch getting stuck this could dissipate a large amount of energy into the switch or RF load, with potentially catastrophic thermal implications. In the case of S-Band it would not cause any catastrophic failure but would certainly greatly attenuate if not completely block the downlink signal.
2. **Transmitter Failure**
   A failure in any part of the transmitter module (either the X-Band or the S-Band module) could cause a partial or complete failure of the downlink chain. Depending on which part of the transmitter fails will dictate the exact type of failure which will be observed, with the most likely being a complete lack of output signal. Other possibilities would include an attenuated output or a failure to modulate telemetry on the output. A problem with the TCXO could also cause a drift or change in the frequency of the downlink signal.
3. **Amplifier Failure (X-Band Only)**
   The S-Band amplifier is part of the S-Band transmitter, but in the case of X-Band there is a separate Travelling Wave Tube Amplifier (TWTA). The TWTA is a delicate piece of high voltage equipment and so the possibility of failure is considered high compared with other TT&C units. The Ulysses mission is an example of in-flight failure of X-Band TWTAs, with its first malfunctioning in 2003 and the remaining TWTA failing in 2008 – ultimately leading to the end of the mission. A failure of the TWTA would likely be immediate and catastrophic, resulting in a complete loss of X-Band downlink signal.
4. **Configuration Failure**
   As with the uplink case, a failure in the downlink configuration could cause a failure of the downlink signal reception. This could be a simple failure, such as a wrong bitrate or downlink band. In both of these cases the ground station would likely be able to detect and correct for the issue. In case of the modulation index being set wrongly, there would be the potential for too much power to be removed from the carrier. In this case it could become hard to distinguish the carrier signal from the noise.

These four types of failure guide the recovery of a problem with the Mars Express downlink. As mentioned though, it may be challenging to determine exactly what has occurred because there will likely be no information being transmitted by the spacecraft to aid the diagnosis.

Mars Express does have separate transmitters for each band, so it is possible to attempt downlink using the alternate downlink band. However, the link budget only supports high bit rates in X-Band, so before any attempt the downlink rate would have to be adjusted accordingly if telemetry was required.

IV. Recovery from an Uplink Chain Failure

The recovery of an uplink chain failure must rely heavily on on-board autonomy, because by definition the ability of the ground to control the recovery has been impaired. There are two stages to the autonomous response to a downlink failure – the detection of a failure and the recovery from that failure.

A. Autonomous Detection of an Uplink Chain Failure

The detection of the root cause of an uplink chain failure is complex. This is because, as described in the previous section, there are several different ways in which the uplink chain can fail. Ultimately it would be too difficult and error prone for the on-board software to monitor different health statuses of the different parts of the receive chain in order to detect an error. Indeed even if that were possible, it could still miss errors such as configuration failures.

The software instead uses a simple system-level approach to detecting a failure in the uplink chain – it works based on when the last command was received from ground. There is a timer on the spacecraft, which is restarted every time the spacecraft receives a command from the ground. There is a configurable limit in the on-board software which, when reached, will trigger the recovery action. This approach is both very simple and very effective at detecting a problem with the uplink chain and covers all three of the identified cases in the previous section. It even covers the case of a configuration failure. The reason it is so effective is that the end result of all of the different failures is the same – the failure to receive commands. Therefore if commands aren't received then there is a problem somewhere in the system.

The timeout is variable and is nominally set on Mars Express to be four days. This value is a tradeoff between avoiding false triggering of the timeout and not delaying the recovery action too far. In case of avoiding false triggering, the maximum allowed duration between passes on Mars Express is 28 hours. Therefore in case of a ground station problem the period without any ground commanding could reach 56 hours. By adding margin to this value, and time to prepare for the recovery action, the value of four days is reached. Any longer would result in an unacceptable delay in recovery. As mentioned, no recovery is possible before the autonomous reaction is triggered and so longer timeouts would just waste time before the mission could be restarted.

There is one exception to the four day rule – this is during solar conjunctions. This occurs every two years when Mars passes behind the Sun as seen from Earth. This can cause a complete loss in RF communications for up to a month. Therefore at entry into this season the timeout is set to 30 days. Every two days the timeout is reduced by two days if ground commands can be received, until it returns to its normal four day setting at the end of conjunction.

B. Recovery Action after Uplink Chain Failure Detection

After the mechanism described above detects a failure in the uplink chain, it triggers another program in on-board software that performs the recovery action. Unfortunately, the drawback of the system level approach to detection means that a specific unit cannot be identified by the recovery software as cause for the failure. Therefore the software makes assumptions about the intended command link and assumes that it has failed. It then makes a series of reconfigurations to redundant units, continuing until one restores the link and the ground can intervene. There are two stages to the reconfiguration:

1. Initial reconfigurations

The initial reconfigurations differ depending on whether an X-Band receiver is in use or not. If one is then the system assumes that X-Band is the primary command path and switches to the redundant X-Band receiver, moving the WIU switch to select this new receiver. It then waits for a configurable period. Originally this was set to 18 hours, although given the previous statements about pass scheduling for Mars Express this was increased to approximately 4 days to give the ground time to recover and intervene. In case intervention is not possible, due to a more serious issue with the X-Band (such as a stuck WIU switch), the next reconfiguration will reconfigure both receivers to S-Band, which in normal operational usage means one connected to the HGA and the other connected to an LGA. There is then another wait period where the ground can try and contact the spacecraft in S-Band and respond to the failure.

If the initial configuration has both receivers configured for S-Band then the system does not know which should be the prime receiver. Therefore every four hours it swaps between the two receivers, moving the RFDU switches to perform the swap. This should overcome any of the problems mentioned in the failure cases above. Ground telemetry monitoring allows the ground to determine when each switch in the cycle will take place.

2. Ultimate Dual-LGA Configuration

If all of the initial reconfigurations fail to restore a working command link then the spacecraft takes more drastic action, taking the form of the Ultimate dual-LGA (ULGA) mode. This mode is entered by the spacecraft triggering a safe mode and then returning from this safe mode with both of the receivers configured in S-Band, each connected to one of the Low Gain Antennas (LGAs). This reaction should restore receive capability to the spacecraft even in cases resulting from a wider system issue than a pure TT&C failure. Any simple TT&C failure should be solved by the Initial Reconfigurations above. The safe mode triggered by ULGA entry would reset and reconfigure the whole spacecraft system, which should solve any further system failures.

The major issue with entry into ULGA mode is the link budget necessary to support an uplink through the low gain antenna. It is likely, unless Earth and Mars are close, that a spacecraft emergency would need to be declared in order to obtain the largest and most powerful stations possible to contact the spacecraft (likely a NASA DSN 70m station). Even in this case it would be questionable whether useful commanding with the spacecraft could be established and highly unlikely that downlinked telemetry would be visible until HGA communications could be restored.

The initial timeout before entering ULGA mode was originally set to only 18 hours, however we are in the process of changing this to be closer to 70 days. The reason for this is the difficulties that would likely be encountered when recovering the spacecraft from ULGA mode. The particular case that this timeout is dimensioned for (with margin) is the solar conjunction period on Mars Express, where contact could be lost for up to 6 weeks. If the uplink failure reaction were to trigger immediately before this period then we would want to wait until commanding was again stable before attempting recovery and ULGA would severely hamper this. On top of the 6 week period an extended window of four weeks was allowed to gather the necessary ground resources and scheduling considerations for recovering from an ULGA triggering.

V. Recovery from a Downlink Chain Failure

The recovery of a downlink chain failure originally relied far less on spacecraft autonomy, as by definition the ground would be able to intervene, even if the results of the intervention may not be visible initially. This approach was as designed into the spacecraft, because a downlink failure does not represent any system-level impact to the safety of the spacecraft. Of course it impairs the ability to monitor the spacecraft activities but there are no safety implications and recovery can proceed without hindrance, albeit in the blind.

The on-board detection of a downlink chain failure presents similar challenges to detection of an uplink chain failure. These challenges relate to the definition of exactly what telemetry would indicate a problem in the downlink chain. However, the complication is that no simple system level metric is available on board. Nonetheless, there are several advantages to having autonomous recovery of the downlink chain in case of failure:

- *Additional protection of units against transient problems*
  At present the only on-board protection against transient problems is the over-current protection provided by the X-Tx and TWTA LCLs. Additional protection might be able to detect and respond to problems earlier and thus minimize any damage caused to the hardware
- *'Seamless' recovery*
  If a failure of the transmitter chain were to occur now the ground would only be aware of the problem due to the lack of signal. Recovery would then take time while attempts would be made to switch on firstly the prime and (if unsuccessful) the redundant chain. On-board protection would provide recovery of the downlink signal after only the few minutes that it takes for the switch over place and the ground would be immediately aware via out of limit and event alarms that a failure had occurred.
- *Minimal loss of data*
  Due to the swiftness of the recovery the loss of science and housekeeping telemetry would be kept to a minimum.

Therefore, on-board autonomy is planned for introduction on Mars Express, to cope with a failure in the downlink chain.

A. Autonomous Detection of a Downlink Chain Failure

The detection of a failure in the downlink chain must be based upon monitoring spacecraft telemetry parameters, as the spacecraft has no knowledge of the last time data was received on ground (i.e. nothing analogous to the system-level uplink case). It was decided that the parameters that would be most indicative of a failure of an X-band transmitter, or a TWTA, would be the RF output power and the TWTA power supply (EPC) LCL current respectively. Unfortunately limit values had not been provided for either of these parameters by the spacecraft manufacturer. It is therefore necessary to determine useful values for the monitoring limits on these parameters without having observed an actual failure. We therefore have to rely upon the information that is available to us, which are the parameter calibration curves supplied by Industry and the TM history from the spacecraft.

Firstly, we consider the RF output power for monitoring the health of the X-band transmitter. It is known that the calibration curves are polynomials. However only four points were provided by Industry, representing the minimum and maximum values of the curve and two points that delimit a linear fit to the central section. This linear section brackets the range of values observed in the spacecraft telemetry when the X-Band transmitters are operating. This was therefore assumed to be the nominal operating range and the monitoring values were selected to be one raw value outside of this linear section.

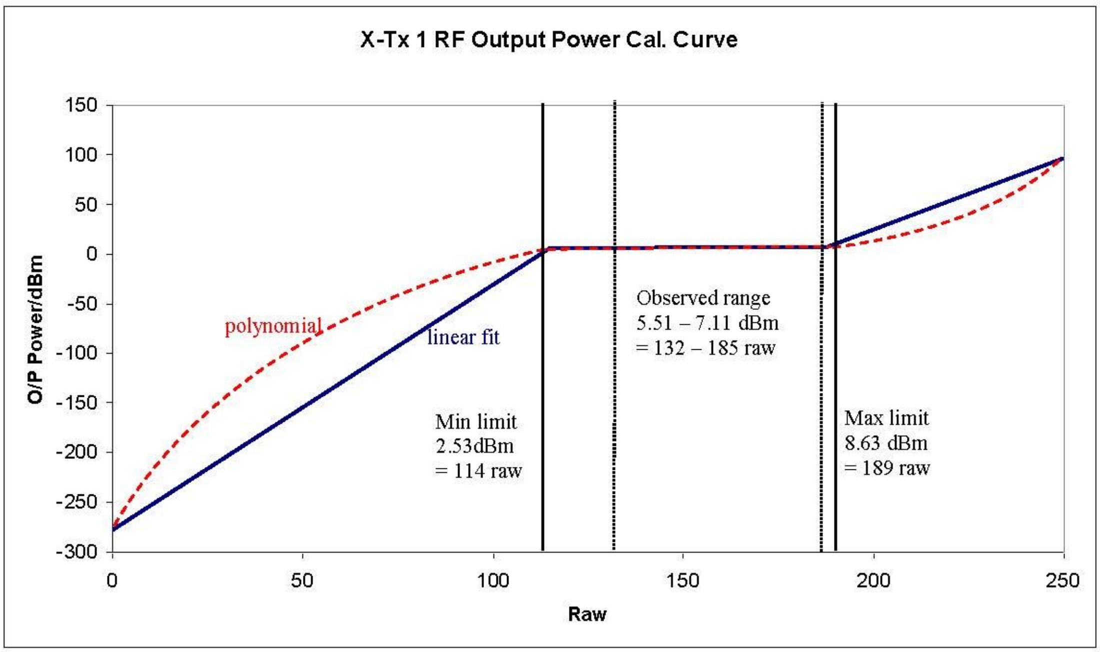

Figure 4. X-Tx-1 Output Power Calibration Curve and Limits

Secondly we consider the TWTA LCL Current Limit selection. The observed range of current consumption of the EPCs are 1.6 +/- 0.2 A if it is on with the transmitter off ('Drive-Off') and 4.2 +/- 0.2 A when on with the Tx on ('Drive-On'). When the EPCs are off the current consumption is less than 0.3 A. Instantaneous values outside these ranges have also been recorded (in all cases <4.4 A). However, in all instances these occur at transitions between Off, Drive-Off and Drive-On states and last for one TM sample only.

The trigger levels chosen for monitoring are therefore Min. = 1.0 A, Max. = 5.0 A. These values comfortably bracket the values observed in TM, both for the Drive-On and Drive-Off cases, although the monitorings would be disabled during Drive-Off and must be disabled during full off.

These parameters are each sampled every 16 seconds in the most commonly used TM mode. Both monitorings will be set to trigger after 5 consecutive out of limit samples i.e. between 64 and 80 seconds after the first excursion. This generous delay has been chosen because avoiding false triggering is considered more important than an immediate response. This triggering is handled by an on-board monitoring function built into the spacecraft software. When triggered this can then execute a recovery action, as described below.

B. Recovery Action after Downlink Chain Failure Detection

In response to a failure only the unit that is detected as failed will be switched to the redundant unit i.e. if a TWTA fails it will be switched over to the redundant unit, but the X-band transmitter will not be switched. This minimal response is similar to the approach that is followed when recovery of such a failure is performed manually from the ground, as it makes it simpler to identify, isolate and recover from the failure.

In order to implement this, a suite of On-Board Control Procedures (OBCPs) are under development that will allow automatic re-configuration of the S- and X-band elements of the TT&C subsystem. Two OBCPs to switch the X-band transmitter on (OBCP 6600) and off (OBCP 6610) have already been implemented on board the spacecraft and are now routinely used for transmitter switching and configuration. They reduce the commanding volume for routine operations, which is extremely useful as a new scheme for operating Mars Express in response to intermittent anomalies on the SSMM has placed severe restrictions on the number of commands that can be stored on-board. In addition they increase safety as only safe configurations and operations are allowed by the OBCPs, and this will help to streamline recoveries in case of Safe Mode, or TT&C failures.

When one of the downlink failure monitorings does trigger it will call another OBCP (OBCP 6650, still under development) that will manage the response to the failure by switching over the unit that has been detected as failed. To avoid 'illegal' configurations (e.g. both TWTAs powered) during the reconfiguration process the OBCP also needs to know what the configuration of the non-failed unit is. That is if a TWTA has failed it also needs to know which X-Tx is in use in order to produce the final configuration with the TWTA configured on the redundant unit and the X-Tx on the original unit. The reverse is also true. This information will not be communicated to the OBCP by the monitoring trigger itself, so the OBCP will have to interrogate the spacecraft data pool to ascertain the configuration of the non-failed unit. Once this is completed the recovery OBCP will call the already existing OBCP 6600 to select the desired X-Tx and TWTA configuration.

If this OBCP is triggered as a result of a failure of the X-Band Transmitter its execution may result in uplink and downlink being configured on separate chains. In this case, regardless of the configuration of the X-band transmitter, the downlink signal will be non-coherent, and ranging will not be possible. However, this is not seen as a major inconvenience and it will be easily remedied during the recovery operations either by swapping the X-band receiver chain, or if the failure was transient by switching the X-band receiver back to the failed unit. In any case, the recovery actions will be greatly improved by having an immediate return of a signal, rather than no signal at all.

VI. Conclusion

The Mars Express spacecraft has a relatively complex communications system that nonetheless provides good scope for recovery in case of a failure of any part of it. The dual-band, triple antenna set-up allows for a wide range of possible configurations but also gives a variety of different areas where a failure could impact the ability of the spacecraft to send and receive data. A failure of this kind is complicated by the fact that the spacecraft is at Mars and therefore the distance is considerably higher and link budget far less favorable than an Earth-orbit spacecraft.

Failure in any part of the uplink chain would result in a complete lack of ability to command the spacecraft and thus recover from the failure, despite the possibility that downlink information would be available. On Mars Express this is managed by on-board autonomy, which takes a system-level approach to failure detection – detecting that no ground commands have been received for a certain period of time. This in turn triggers an on-board recovery action which switches different receiver configurations until one works and the ground can intervene.

Failure in the downlink chain is not critical to spacecraft safety and therefore originally had no on-board autonomy covering it. However, there are numerous reasons why it would be preferable to have an on-board response in the case of a downlink failure too. Therefore a process is ongoing on Mars Express to implement an OBCP-based approach to recovery of a downlink chain failure, triggered by monitoring of on-board transmitter and amplifier health parameters.

These measures ensure that Mars Express is robust to any kind of failure in its TT&C equipment. By ensuring this, one of the most vital parts of our ability to control spacecraft – the link between Earth and space – can be relied upon to provide all the needs of the mission and its operators.

Appendix A
Acronym List

DIP	Diplexer
EPC	Electronic Power Conditioner
ESA	European Space Agency
HGA	High Gain Antenna
LCL	Latching Current Limiter
LGA	Low Gain Antenna
OBCP	On-Board Control Procedure
RF	Radio Frequency
RFDU	Radio Frequency Distribution Unit
Rx	Receiver
SW	Switch
TCXO	Temperature-Compensated Crystal Oscillator
TM	Telemetry
TT&C	Telemetry, Tracking & Commanding
TWTA	Travelling Wave Tube Amplifier
Tx	Transmitter
ULGA	Ultimate dual-LGA mode
WIU	Waveguide Interface Unit

Acknowledgments

The authors would like to thank the designers of the Mars Express spacecraft for the robust existing failure detection and recovery software upon which the response outlined in this paper has been based.