Ground Equipment for the Support of Packet Telemetry and Telecommand

Wolfgang Hell

Abstract - *This paper describes ground equipment for packet telemetry and telecommand which has been recently developed by industry for the European Space Agency (ESA). The architectural concept for this type of equipment is outlined and the actual implementation is presented. Focus is put on issues related to cross support and telescience as far as they affect the design of the interfaces to the users of the services provided by the equipment and to the management entities in charge of equipment control and monitoring.*

Introduction

This paper describes the telemetry and telecommand sub-systems which have recently been developed by European industry on behalf of the European Space Agency (ESA) and are presently being deployed to the Agency's ground station network. On the one hand, the design of these subsystems has been driven by ESA's Packet Telemetry and Telecommand standards (*PSS-04-106, PSS-04-107*) which in turn have been derived from the related "Blue Books" produced by Panel 1 of the Consultative Committee for Space Data Systems (CCSDS) (*CCSDS 102, CCSDS 201, CCSDS 202, CCSDS 203*). These standards in essence determine the functionality to be provided by the sub-systems. On the other hand, although final results are not yet available, also Panel 3 activities aiming at a standardisation of the services made available by ground segment entities have been taken into account. Specifically in the design of the subsystem interfaces care has been taken to cleanly separate services accessible to users from management issues. This approach and usage of the full OSI protocol suite will facilitate cross support between space flight agencies. In order to ensure appropriate growth potential and life time of the architecture, the design took also into account CCSDS work on Advanced Orbiting Systems (AOS) (*CCSDS 701*).

From this comprehensive set of requirements initially an "ideal" architecture has been derived which subsequently has been modified to accommodate constraints in terms of available hardware and software as well as cost.

The "Back-end" Architectural Concept

In ESA terminology, the Back-end of a ground station encompasses all equipment connecting the intermediate frequency equipment of the front-end to the ground communication network in order to provide the remote users (normally control centres) with the services required to operate the spacecraft and to acquire payload data. Set-up and monitoring of the back-end subsystems is done via a "management" interface. Figure 1 presents the back-end's context.

Wolfgang Hell is member of the Technology Development Division in the Stations and Communications Engineering Department of the European Space Operations Centre, Robert-Bosch-Str. 5, D-64293 Darmstadt, Germany
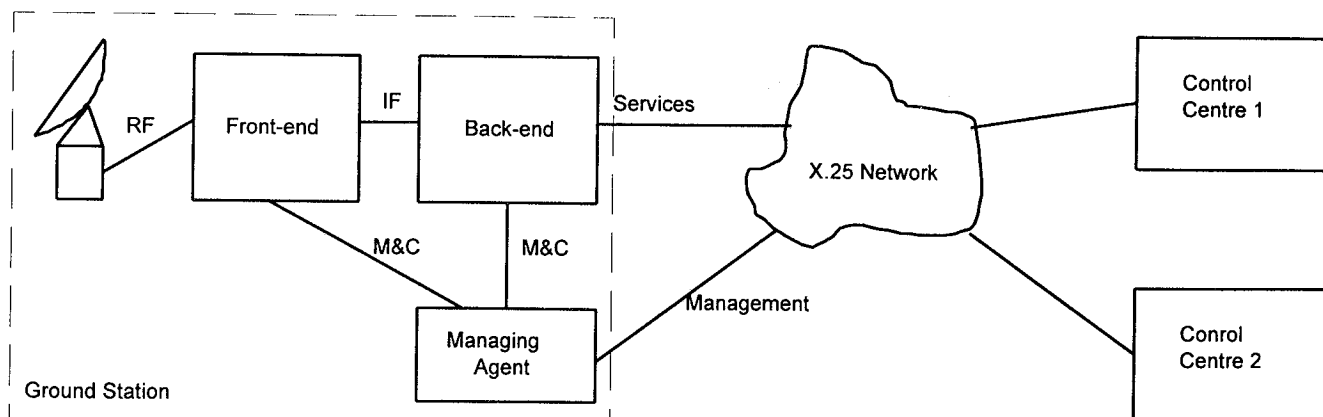
Figure 1: Back-end Context

On the return link, the back-end receives the demodulated symbol stream from the front-end which is either immediately processed and forwarded to the user (on-line service) or stored in the back-end and forwarded later on user request (off-line). Regarding the forward link, the back-end receives data from the user, i.e. the control centre, via the communication network and, depending on the user request, either forwards them immediately to the front-end (on-line) or stores them for uplink at the specified time (off-line service). Conventional telecommanding can be considered as a special type of on-line forward link service.

The back-end shall support the space data standards for spacecraft operation listed below:

- PCM Telemetry and Telecommand (*PSS-46, PSS-45*),
- Packet Telemetry and Telecommand (*PSS-04-106, PSS-04-107*),
- a subset of AOS (*CCSDS 701*).

A prime design objective has been to relieve the users from the need to be aware of configuration details which are internal to the back-end. In other words, the user should only need to know about the service requested, but not how the back-end manages to provide this service and how the required availability figure (e.g. by means of redundancy) is attained. These design goals have led to the internal back-end structure depicted in figure 2.

The various Functional Units (FU) presented in figure 2 are physically interconnected by means of a dual ring (i.e. redundant) FDDI (Fibre Distributed Data Interface) Local Area Network which allows to set up so-called Functional Processing Chains (FPC) by establishing virtual circuits between the Functional Units as required for the provision of the service enabled by station management. For the Storage FU no redundancy is shown, as it is assumed that the required availability and performance for concurrent provision of multiple service instances is attained by means of internal redundancy (e.g. a RAID (Redundant Array of Inexpensive Disks) architecture). FUs drawn with dashed lines are not part of the initial implementation presented below, elements drawn with dash-dotted lines are only partly implemented. The concept depicted in figure 2 provides for a high degree of scalability in terms of attainable throughput (limited by the FDDI bandwidth) and availability. Typical station integration

problems where a high level of redundancy entails complex cabling and switching units which in turn have a negative effect on the actual availability are avoided.
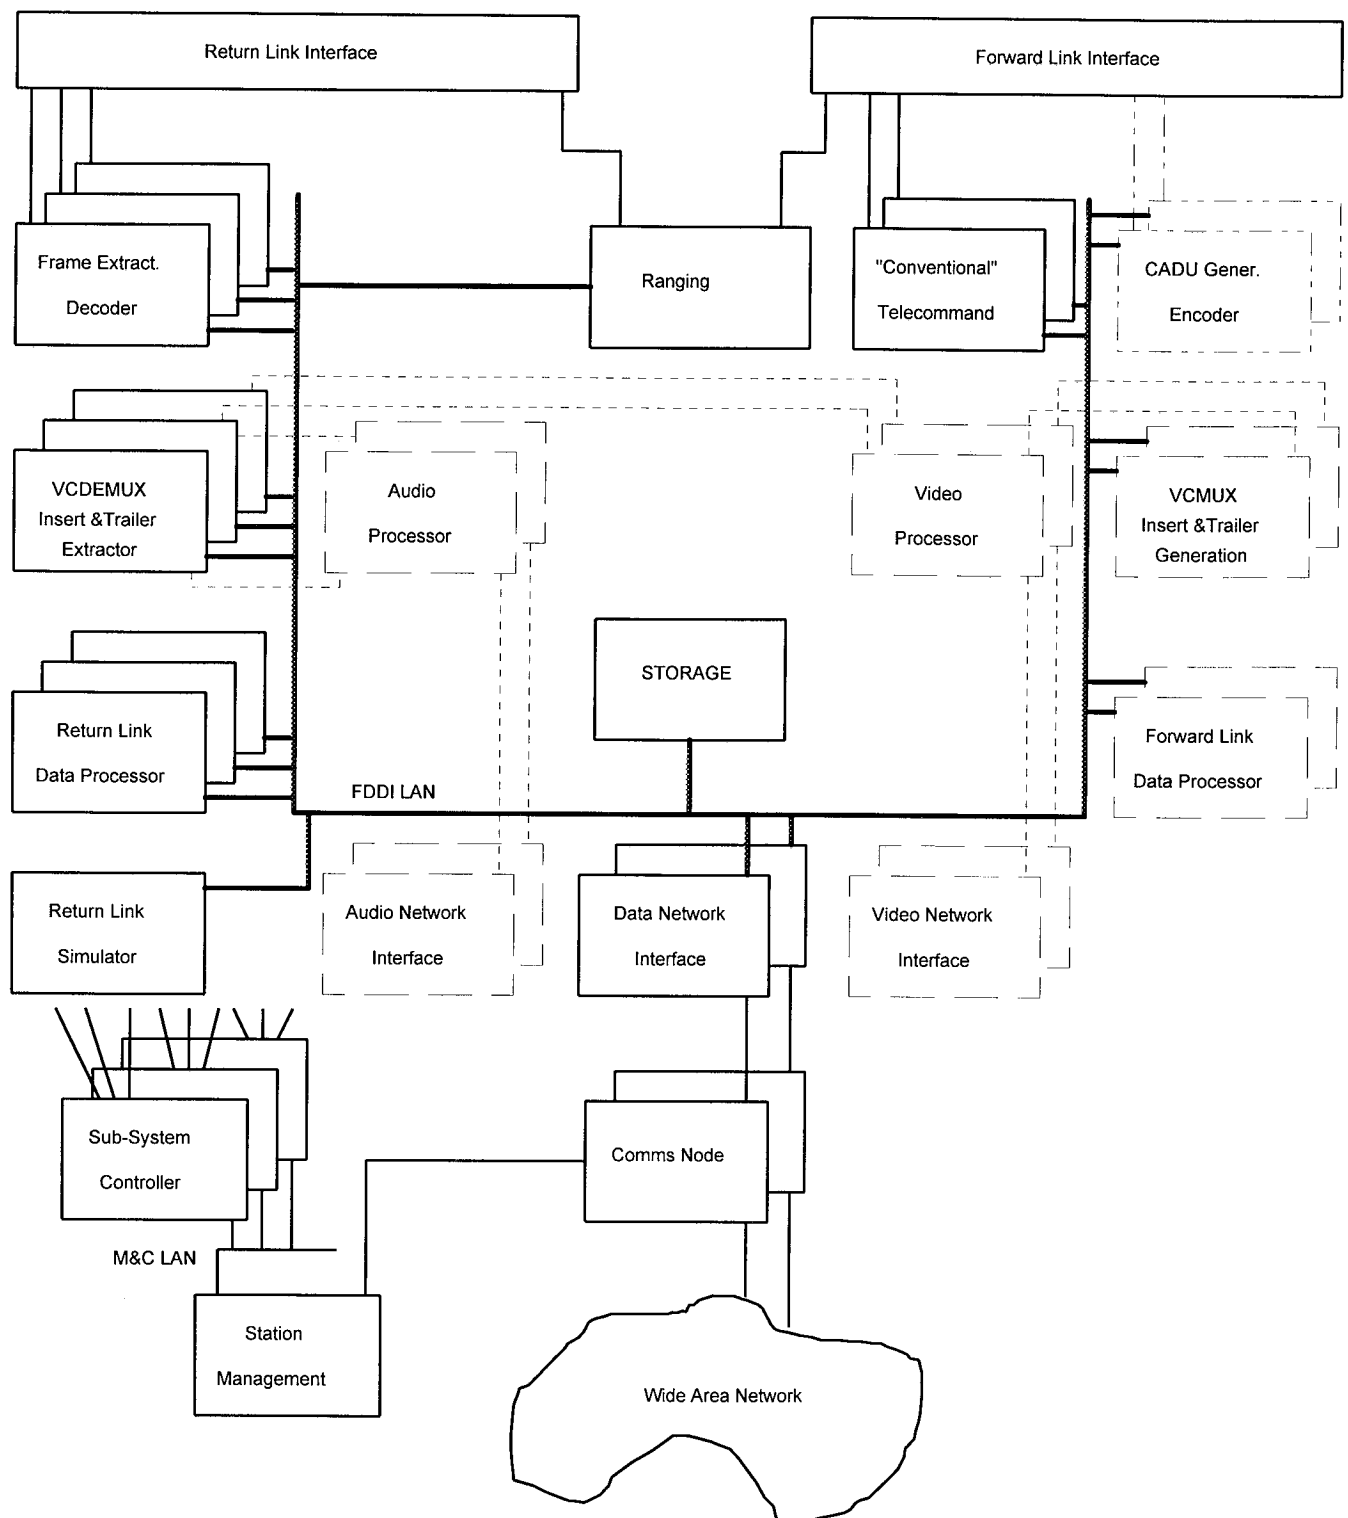
Figure 2: Back-end Architectural Concept

Actual implementation

Unfortunately, in practice the above outlined concept had to be somewhat diluted since the actual implementation has been constrained by the availability of suitable hardware (and software) and cost. Furthermore, for some interfaces compatibility with existing installations had to be ensured such that the newly developed equipment could serve as an in situ replacement of existing, but obsolescent equipment. Figure 3 shows the resulting block diagram. It is intended to upgrade the subsystems in an evolutionary process to the concept outlined in figure 2.

As regards the return link handling part, the actual implementation does not (yet) support audio and video as defined for the AOS environment (components drawn with dashed lines in figure 2). For reasons of interface compatibility with existing stations, the Frame Extractor/Decoder (FED) functional units do not yet support the FDDI interface, but deliver the synchronised and decoded transfer frames to the VCDEMUX units via IEEE-488 point-to-point connections. Similarly, the VCDEMUX units provide the extracted Command Link Control Word (CLCW) messages via dedicated serial links rather than the FDDI LAN to the telecommand encoders. The Storage system had to be split into individual storage processors as a dual port RAID architecture was found to be by far too costly for the time being. This limitation unfortunately implies that in order to achieve redundant storage of telemetry, the data have to be duplicated by the VCDEMUX and sent twice through the FDDI network since the protocol (see below) does not allow "broadcasting".

The forward link handling part has been confined to the "conventional" (i.e. PCM and packet) telecommand function. The AOS type of forward link is not (yet) supported. A limited capability to generate and encode CLTUs for the forward link is (for test purposes only) available in form of the return link simulator and the encoding capability of the FED units. Therefore, the CADU (Channel Access Data Unit) Generator/Encoder units are drawn in figure 2 with dash-dotted lines.

Although the FDDI network provides for 100 Mbps bandwidth, the sustained throughput on a virtual circuit connecting two functional units was found to be less than 6.4 Mbps even with large block sizes. This is in part due to the fact that off-the-shelf only TCP/IP as protocol is available. This protocol is designed for a very unreliable network service and therefore introduces considerable overhead which in an FDDI environment is superfluous. The problem is further aggravated by the fact that the bit manipulations required for the TCP/IP error control field calculation are carried out on the CPUs of the connected units. From the introduction of a suitable "light weight" protocol such as XTP one can expect a substantial throughput improvement, in particular when in addition to point-to-point connections "broadcasting" is supported.

The Return Link Protocol Handling System (R-PHS)

Since the R-PHS must be usable as an in situ replacement for the presently deployed telemetry system, it must emulate the existing equipment as regards the PCM telemetry standard in order to ensure continuing support to on-going missions without any impact on the control centre system. For PCM

216
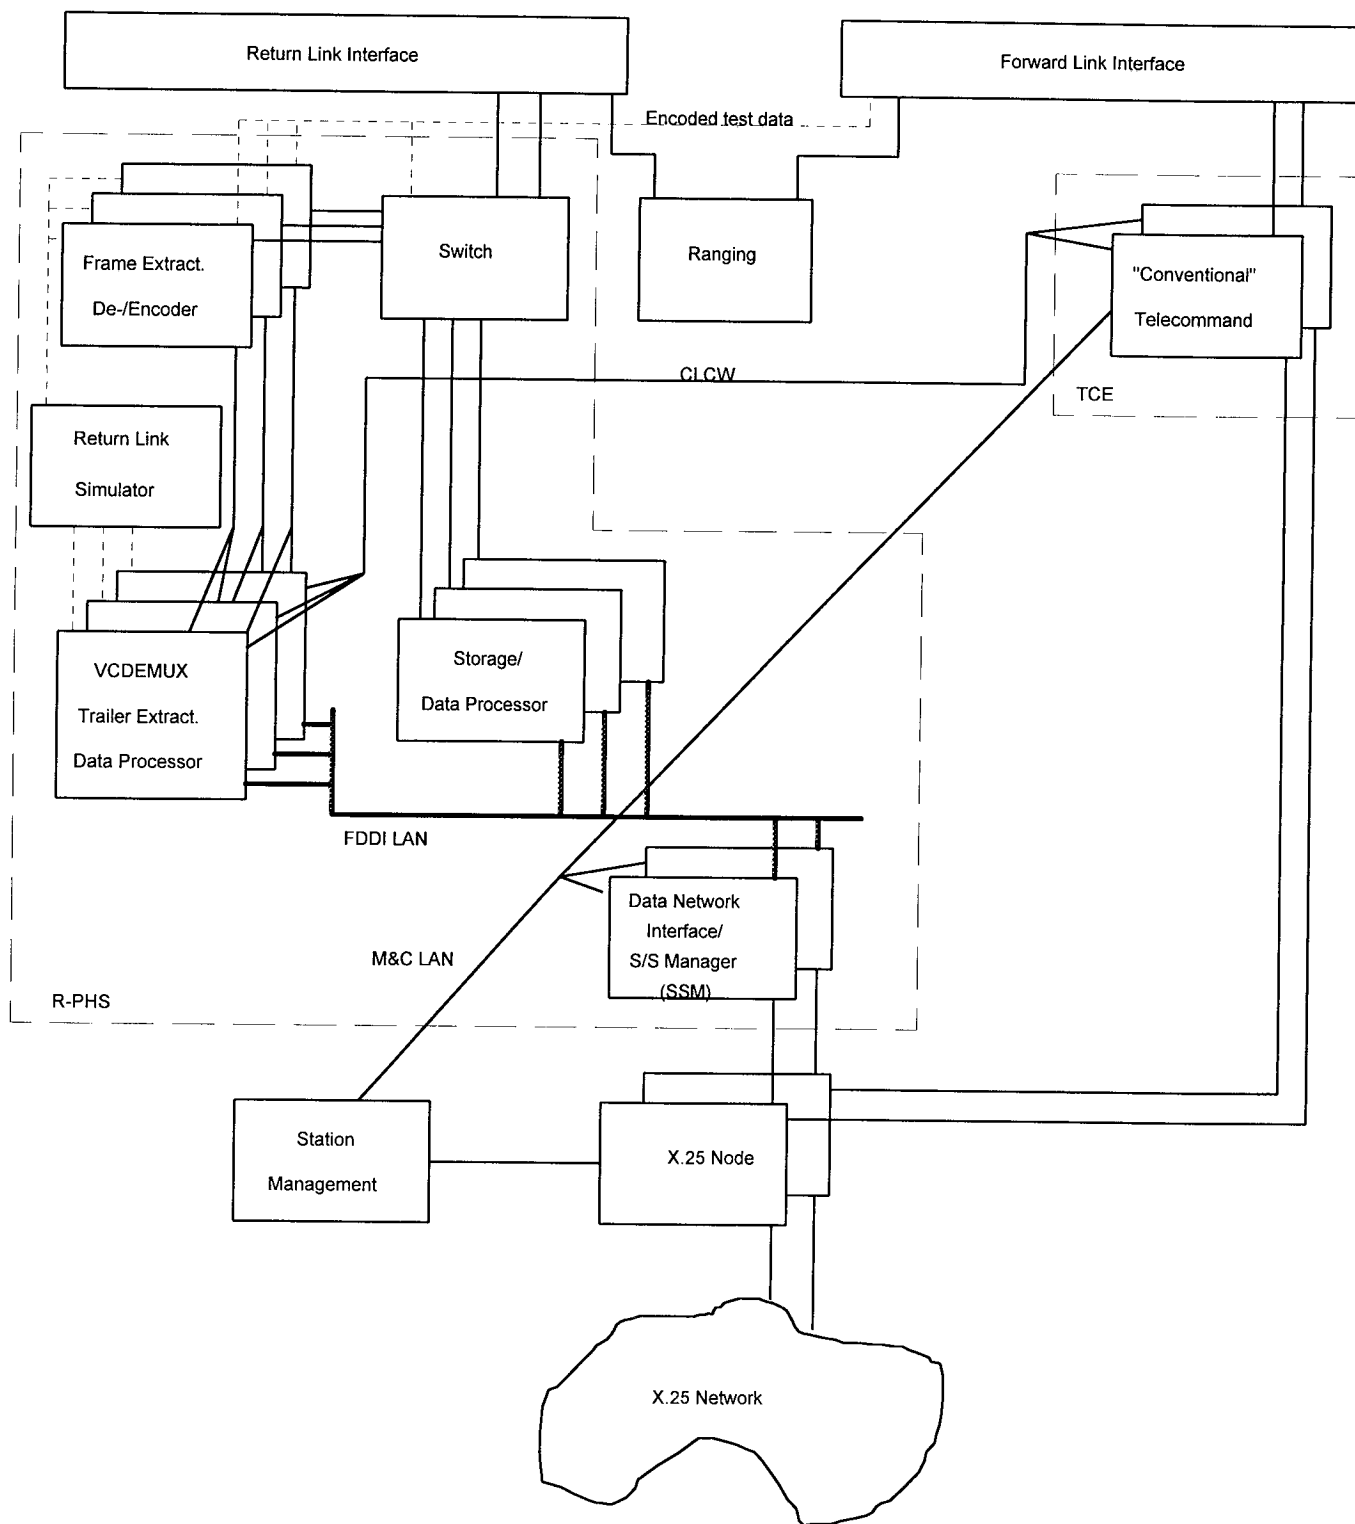
Figure 3: Back-end Implementation

mode, the R-PHS thus supports the private "ESA Message Protocol" built directly on top of the X.25 network layer and provides both on-line and off-line telemetry data delivery.

For the services related to packet and AOS telemetry, a prime objective of the project has been to develop a system which due to the services and protocols provided facilitates telescience and cross support. Application layer (i.e. layer 7 in the OSI reference model) protocols such as FTAM and CMIP will give the least problems in terms of interoperability between the R-PHS and the user which will generally require inter-operating of computer systems from different vendors. However, in most cases the telemetry system is connected to the user via an X.25 network which provides limited bandwidth only. Therefore, the efficiency of the bandwidth usage is a critical issue in particular for the transfer of (high-volume) payload telemetry, as long as high-speed wide area network technology like Frame Relay or ATM (Asynchronous Transfer Mode) are not yet widely available. For the R-PHS the link to the user has been split into a control channel, on which the layer 7 Common Management Information Protocol (CMIP) is used, and a data channel for bulk data transfer, on which for efficiency reasons the OSI protocol suite has been limited to session layer.

Before a user can connect to the R-PHS, the system must be configured via the Sub-System Manager (SSM) such that the Functional Processing Chains providing the services to be granted to the user are established by connecting the various functional units in the appropriate way. The correct functioning of the established chains is checked by means of built-in test facilities and, in case a functional unit is found to be faulty, alternative chains excluding the defective unit are set up by the SSM. The Data Network Interface units are notified of the final set-up in terms of service providers and user access rights such that incoming requests can be validated and routed to the unit providing the service requested by the user. In this way, the user is relieved from the need to be aware of the internal R-PHS set-up. The R-PHS appears to the user as a telemetry server which can be accessed by using the network addresses of the DNI units. Which actual Functional Processing Chain (i.e. which physical FUs) provides the requested service is transparent to the user.

As for PCM telemetry, for packet and AOS telemetry the R-PHS provides on-line and "off-line" services. In on-line mode the data are delivered without flow control, but with overflow management. When the volume of data requested by the user exceeds the available bandwidth of the connection to the user, data are discarded in a controlled way such that minimum size blocks of contiguous (as far as successfully reconstructed from the incoming symbol stream) telemetry are delivered. The block size is user selectable. In addition, a user controlled release timer warrants a worst case latency of the telemetry delivery.

In terms of data selection, the R-PHS supports these options:

- Space Link Channel SLC (i.e. all frames)
- Master Channel MC (i.e. all (good) frames of the specified S/C ID and Version ID)
- Virtual Channel VC (i.e. all good frames of the specified VC in the specified MC)
- MC Secondary Header (i.e. the Primary and Secondary Headers extracted from the frames received on the specified MC; Packet Telemetry only)
- VC Secondary Header (i.e. the Primary and Secondary Headers extracted from the frames received on the specified MC/VC; Packet Telemetry only)

- VC Access (i.e. the VCDU Data Zone extracted from the frames received on the specified MC/VC; AOS only)
- VC Bitstream (i.e. the Data Field Status (Packet Telemetry only) and the Data Field/Data Unit Zone (without fill data) extracted from the frames of the specified MC/VC)
- MC Control Field (i.e. the Operational Control Field extracted from the frames received on the specified MC)
- VC Control Field (i.e. the Operational Control Field extracted from the frames received on the specified MC/VC)
- Source/Path Packets (i.e. the reconstructed source/path packets with the specified AP-IDs received on the specified MC/VC; the AP-ID list can be modified on-line; synchronisation markers inserted into the data transferred to the user indicate when the new selection has become effective)
- Time Calibration (i.e. the Time Calibration Packet constructed for the specified VC; Packet Telemetry only)
- Space Link Status (on user request, the R-PHS monitors the space link status and reports any status changes detected)

The other service class is the so-called "immediate data access" (IDA), which as opposed to the on-line service delivers the selected data with full flow control. This means that the selected data will be delivered to the user as fast as the available link bandwidth allows, but due to the applied flow control no data will be lost. This service class is not called "off-line", since it allows the user in a single selection not only to request data already stored by the R-PHS, but even telemetry still to be acquired. This means that, available communications bandwidth permitting, the IDA service class can also be used for near real-time telemetry delivery, in case flow control is essential.

The data selection options are mostly identical to the on-line service class with the following exceptions:

- Information on the presently stored telemetry can be retrieved in the form of directories
- Data selections can be further refined by specifying
  - the start and end time and or counter range
  - the start time or counter and the number of data units to be delivered
- List of AP-IDs can only be changed when a data transfer invoked earlier has been terminated
- Space link status reports are not available

If the real-time telemetry received by the R-PHS contains also a Virtual Channel conveying so-called tape dump data, the transfer frames of this virtual channel will initially be stored as any other VC. Under control of the SSM, the data zones of the transfer frames of the tape dump VC are extracted and, if required, in reverse order, serialised and forwarded to a Frame Extractor Decoder (FED) unit which performs frame synchronisation and decoding. As real time telemetry, the annotated frames are forwarded via the VCDEMUX to the Storage unit which stores them in the "tape-dump" directory rather than the real-time directory. By means of the IDA services the user has access both to real-time and tape-dump telemetry by selecting the appropriate directory.

The Telecommand Encoder (TCE)

The "conventional" telecommand function is implemented in a single physical unit encompassing the communication network interface for connection to the user, the telecommand engine proper and the PSK modulator. As opposed to the R-PHS, the Functional Processing Chain providing the telecommand service cannot be dynamically built from a pool of functional units interconnected by virtual circuits over a LAN. Therefore, for the telecommand function the "server" concept has not yet been implemented. By selecting the network address, the user also specifies implicitly the physical resources which will provide the requested service.

Otherwise, as regards the protocols, the same considerations as presented for the R-PHS have been applied. Since the present modulation standard limits the maximum throughput on the space link for telecommands to 4 kbps, and, in general, compared to telemetry, the throughput requirements are moderate, for telecommanding a single seven layer protocol architecture (i.e. not split into control and data channel) is used.

The TCE provides three types of services:

- the PCM Telecommand service,
- the Packet Telecommand service,
- the Physical Layer Interface service.

Before a user connects to the TCE, the service to be provided is enabled through the management interface and the access rights for the user(s) are set.

The PCM telecommand service has been implemented in order to ensure the continuation of support to on-going missions, whenever the new TCE has to be installed as an in situ replacement of the equipment presently deployed. To avoid the need for any modifications on the control centre side, the related private ESA message protocol implemented on top of the X.25 network layer has been implemented for accessing the PCM telecommand service.

The Packet Telecommand service, which can be accessed using the Common Management Information Protocol (CMIP), supports telescience applications by allowing payload control centres to connect in parallel with the flight agency's control centre. To safeguard the mission, only the flight agency's control centre has control over the telecommand session, i.e. the establishment and release of the radio link to the spacecraft. Only this control centre has access to the Bypass Control (BC) service and is allowed to send directives affecting the state of the telecommand protocol engine in the TCE. It also determines the uplink bandwidth allocation by specifying the MAP (Multiplexor Access Point) multiplexing scheme. Should an emergency require to do so, the flight agency's control centre can at any time lock out payload control centres. The TCE will only accept telecommand requests as long as the user access rights encompass the selected MAP and Application Identifier(s) (AP-ID). In order to facilitate cross support, the TCE implements, in addition to the ESA packet telecommand standard, also the CCSDS Blue Book, where the latter, in contrast to the ESA standard, allows Bypass-Control (BC) services to be supported without first terminating the Acceptance-Data (AD) service.

The Physical Layer Interface service is intended for support of spacecraft which adhere neither to the PCM nor to the Packet Telecommand standard. In this service, the TCE is practically transparent and enables uplinking of CLTUs (Command Link Transfer Unit) as submitted by the user. In addition, the user has control over the insertion of acquisition and idle sequences, where in case the TCE is requested to insert them automatically the bit patterns can be defined.

The Monitoring & Control Concept

The objective of the new M&C concept has been to establish a clean management hierarchy (ground station network, individual ground station, individual subsystem) to allow control of the services made available to users by the various ground segment entities. Furthermore, the implementation of the concept should exploit more advanced technology to replace the mostly IEEE-488 bus-based station internal M&C infrastructure which is cumbersome and expensive to maintain because of the lack of truly standardised protocols, data types, data presentation and bandwidth limitations.

Within the scope of this paper, it is not possible to present the entire M&C concept. What is presented below, addresses the sub-system related aspects of this concept.

Also in the area of ground station equipment, considerations of cost and user friendliness have led to the introduction of Graphical User Interfaces (GUI), replacing the expensive, individually designed front panels. The availability of powerful GUI building packages in the UNIX world resulted in the introduction of UNIX in embedded systems which traditionally had been built exclusively around real-time kernels. UNIX also made the LAN technology readily available which enabled to place the GUI infrastructure, which then is shared between individual sub-systems, at the stations operator's normal working position, relieving him from having to walk to the individual sub-system to control it.

This evolved environment also facilitates the introduction of a modern M&C infrastructure, where the expensive and cumbersome IEEE-488 infrastructure is replaced with the LAN and where due to UNIX suitable protocols available as off-the-shelf products can be introduced providing for truly standardised data exchange. Candidate protocols have been evaluated and CMIS/CMIP has been chosen. The initial implementation only uses a subset of the service elements provided by this protocol, in particular scoping and filtering are not used. In order to obtain good adaptation to the application and to avoid unnecessary complexity, privately defined managed objects (M&C-ID-0) are used rather than the Generic Managed Objects defined in ISO 10165. By means of these objects, a "conceptual" view of the sub-system which is then available to the managing entity can be modelled. These objects are briefly described below.

Any sub-system is assumed to arrange the M&C related resources in a tree structure in line with the hierarchy depicted in figure 4, where this tree has three levels: the sub-system proper, individual subsystem units, and function blocks (function blocks can however exist at sub-system level). These elements are not only structuring the resource tree, but they are also Managed Objects which support generic sub-system administration like state transitions from set-up to operable, control mode (local or remote) and the like. The tree structure determines the scope of visibility of resources. At any node within the tree only those resources which belong to that node or a branch below that node are visible.

Associated with each node of the tree, different types of resources like Variable Lists (VL), and tasks may exist. These resources can be mapped to the different managed objects accessible through the management interface of the subsystem. This resource structure as well as the mapping to the managed objects is specified in the so-called Management Information Base (MIB) description file (*M&C-ID-1*), which is evaluated at start-up of the subsystem. If the particular site or the mission to be supported require a modification of the view presented to the managing entity, the MIB description file is updated accordingly. The generation of a modified view, i.e. a different mapping of resources to the managed objects, does not require any change to the software of the sub-system. At start-up, the MIB is built according to the MIB description file. Obviously, since the resource structure is implemented in the sub-system's application software, this part of the MIB description file is fixed.
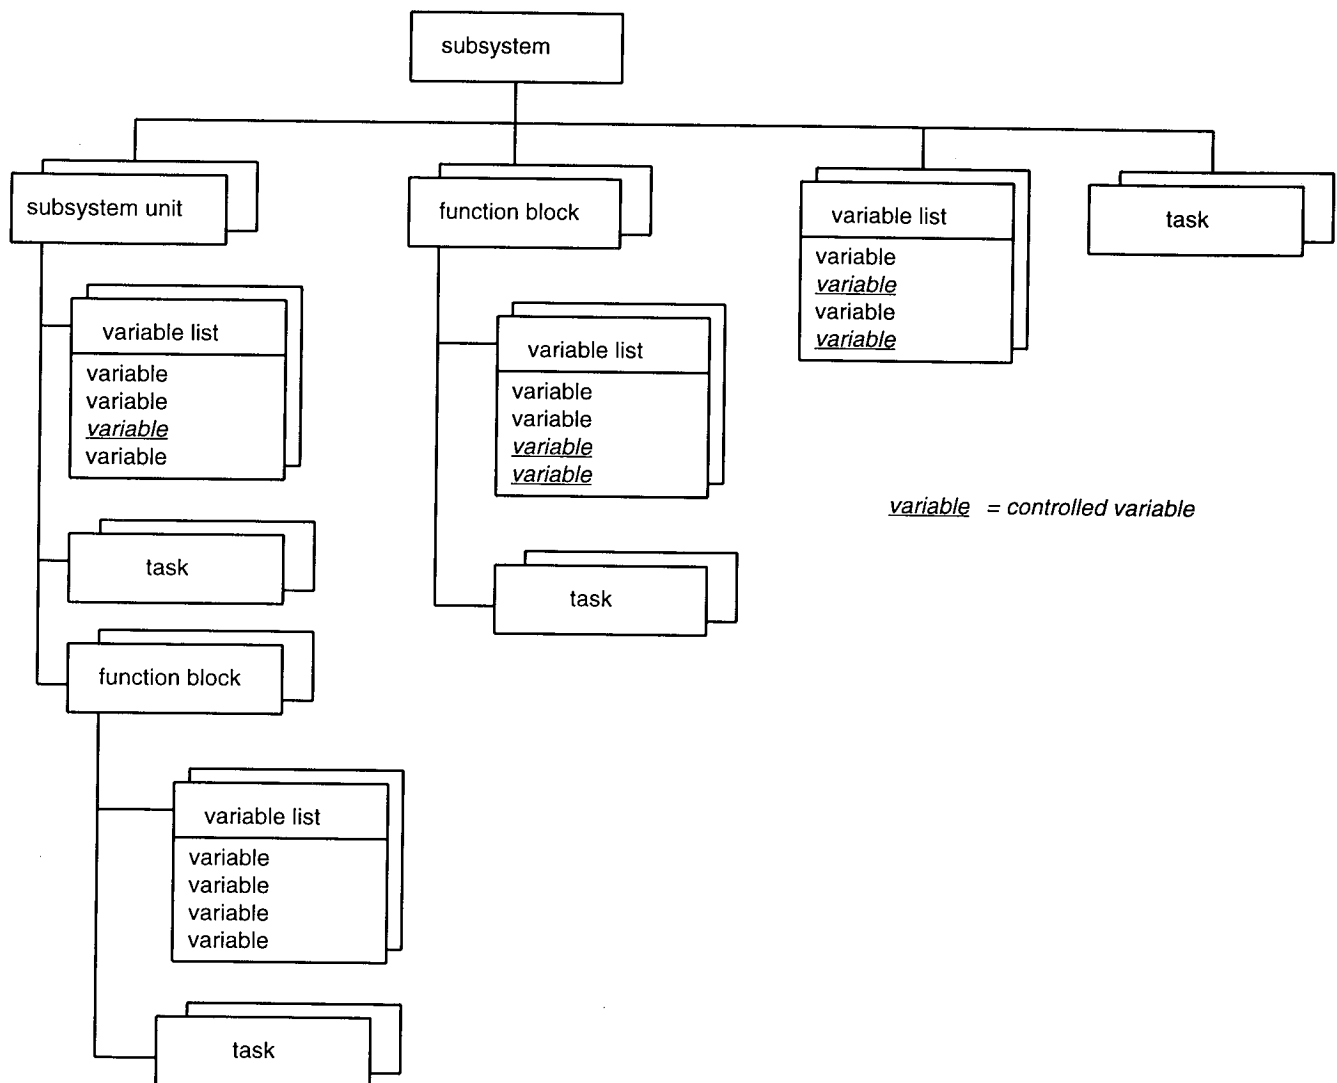
Figure 4: Sub-system Resource Tree Structure

The mechanism by which the resources are mapped to the Managed Objects accessible to external entities is illustrated by means of a very simple example in figure 5. The purpose of the Managed

Objects "sub-system", "sub-system unit", and "function block" is explained above. The "monitored variable list" provides read access to subsystem internal variables, where the manager can choose to receive a report either cyclically, on change of (at least) one variable contained in the list, or only on request. The "controlled variable list" enables the manager to set subsystem variables to either the values specified in the request or to default. Any subset of the variables contained in the Managed Object is accessible. The "task" object is used to invoke, stop, or abort the execution of specific functions in the sub-system, where the object is used both to convey any arguments as well as for monitoring of the function execution. The "event handler" objects allow the detection of sub-system internal events and the automatic triggering of associated actions. The manager can switch on or off the complete event detection as well as the individual associated actions. The "log" object is used to copy the specified subset of the sub-system log into a "public" file store from which it can then be retrieved by the manager. As a future extension, it is intended to implement a Managed Object for the administration of sub-system schedules. Another set of Managed Objects is used for control of inter-subsystem communication. This feature is used e.g. by the TCE which connects to the Front-End Controller for checking the front-end status.
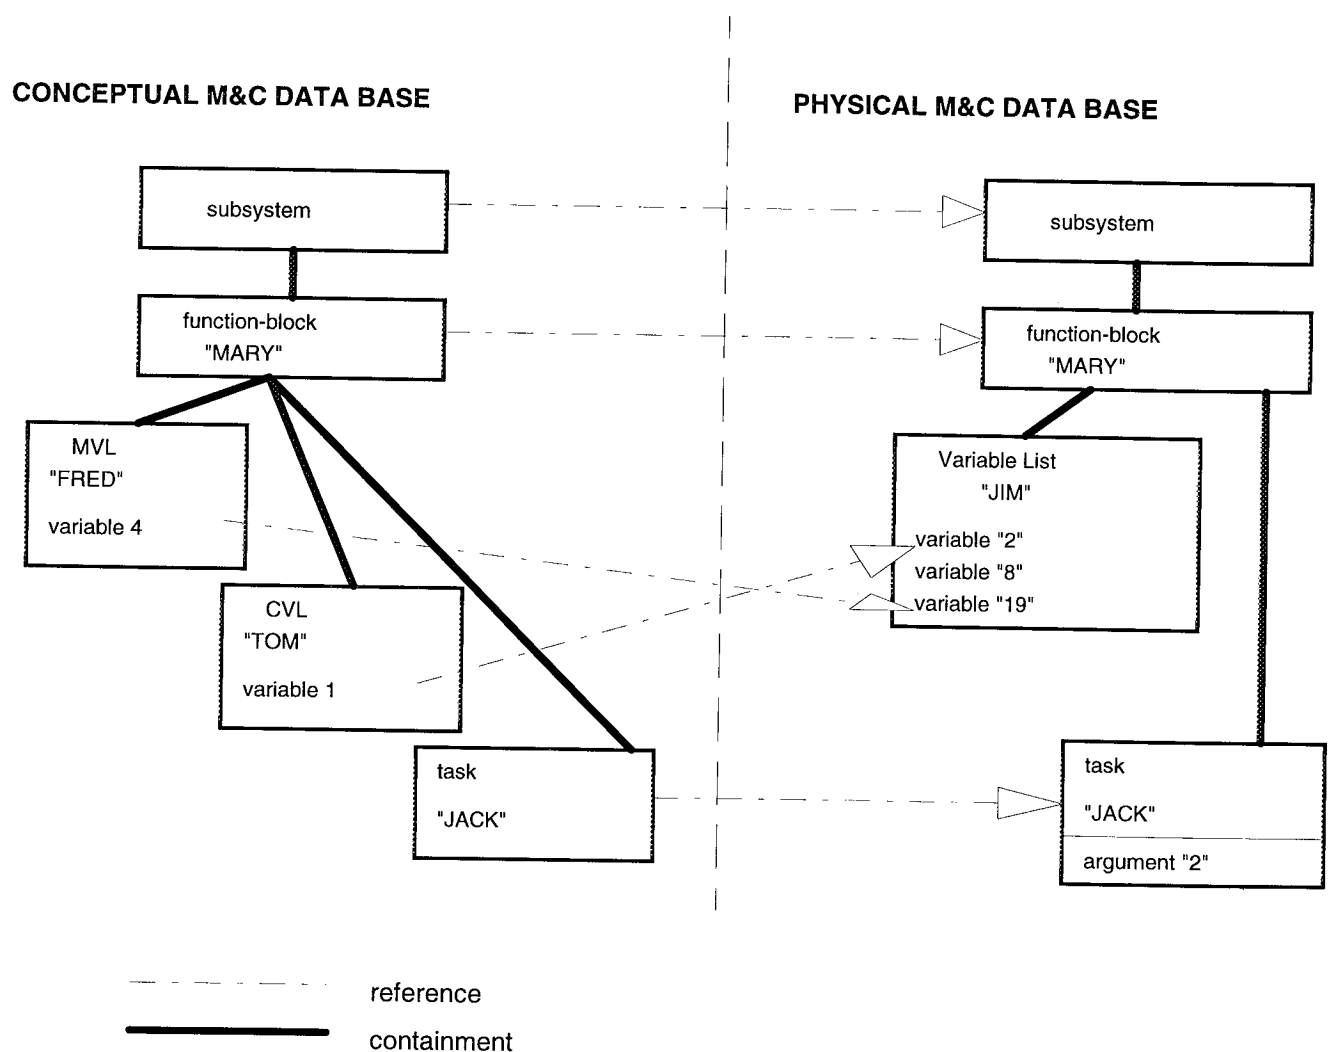
Figure 5: Mapping of Resources to Managed Objects

Conclusion and outlook

Starting from the architectural concept for the ground station back-end equipment, this paper has described the actual implementation as constrained by presently available hardware and software, cost and the need for backward compatibility. The growth path towards full implementation of the CCSDS AOS recommendations has been outlined.

The features designed into the equipment to facilitate cross-support and to promote telescience in terms of available services and management concept have been high-lighted.

Further system enhancements of the described sub-systems will be driven by mission needs. Hardware modifications will aim at getting closer to the architectural concept, in particular as regards the Frame Extractor/Decoder component. Mid-term extensions of functionality are expected in the area of further refined services resulting from the introduction of the Packet Utilisation Standard (PUS). Another activity which has already been started is the development of a "low-end" telemetry system which while retaining the functionality and user interface will be based on much simpler (and therefore cheaper) hardware. The considerably lower performance of this system is still sufficient to cover a wide range of TT&C applications such as geostationary communications satellites.

References

PSS-04-106   Packet Telemetry Standard, ESA PSS-04-106, Issue 1, January 1988
PSS-04-107   Packet Telecommand Standard, ESA PSS-04-107, Issue 2, April 1992
CCSDS 102   Packet Telemetry, CCSDS 102.0-B-2, Blue Book, January 1987
CCSDS 201   Telecommand: Part 1 Channel Service, Architectural Specifications, CCSDS 201.0-B-1, Blue Book, January 1987
CCSDS 202   Telecommand: Part 2 Data Routing Service, Architectural Specifications, CCSDS 202.0-B-1, Blue Book, January 1987
CCSDS 203   Telecommand: Part 3 Data Management Service, CCSDS 203.0-B-1, Blue Book, January 1987
CCSDS 701   Advanced Orbiting Systems, Networks and Data Links, Architectural Specification, CCSDS 701.0-B-1, October 1989
PSS-46   PCM Telemetry Standard, ESA PSS-46 (TTC.A.02), Issue 1, April 1978
PSS-45   PCM Telecommand Standard, ESA PSS-45 (TTC.A.01) Issue 1, April 1978
M&C-ID-0   M&C Interface between the Station Computer II and the Sub-System Controllers, ESA, M&C-ID-PTS-0000 (SCED/TDD/27318/PTS), Issue 1.3, July 1994
M&C-ID-1   ESTRACK Standard for Management Information Base Definition Files, ESA, M&C-ID-PTS-0001, Issue 1.1, July 1994