Nationaal Lucht- en Ruimtevaartlaboratorium

National Aerospace Laboratory NLR

NLR TP 96625

Simulation of a cell switched network for the control of a switch matrix in a high-speed avionics network

D.A. Aupers, G.J. Heerink and S. Wellink

DOCUMENT CONTROL SHEET

	ORIGINATOR'S REF.		SECURITY CLASS.
	NLR TP 96625 U		**Unclassified**

ORIGINATOR
National Aerospace Laboratory NLR, Amsterdam, The Netherlands

TITLE
Simulation of a cell switched network for the control of a switch matrix in a high-speed avionics network

PRESENTED AT
the AGARD MSP 6th Symposium on "Advanced architectures for Aerospace Mission Systems", held in Istanbul, Turkey from 14 October to 17 October 1996.

AUTHORS	DATE	pp	ref
D.A. Aupers, G.J. Heerink and S. Wellink	**961023**	**12**	**2**

DESCRIPTORS

Architecture (computers)
Avionics
Channels (data transmission)
Computer networks
Data tranfer (computers)
Message processing
Network control
Optical switching
Packet sswitching
Radar data
Systems simulation
Wideband communication

ABSTRACT
This paper describes the research and experiments carried out by the National Aerospace Laboratory (NLR) in the field of high-speed interconnection systems for modular avionics. The research has been carried out in the EUCLID/RTP4.1-framework.
The avionics network that was modelled and simulated was an optical switch matrix under control of a cell switched network. The optical switch matrix offers the avionics system circuit-switch, uni-directional, point-to-point connections. A bandwidth of 2 Gbps is projected. The main purpose of the matrix is to connect sensors producing high data rates, such as an attack radar in fighter aircraft, with the core avionics processing cluster.
The cell switched network - in this case Asynchronous Transfer Mode (ATM) - controls the optical switch matrix and provides data transfer at lower data rates, file transfer, and status messages. The simulation model operated ATM at 149 and 622 Mbps.
The primary objective of our research was to assess ATM as a data link layer for a control and message network in an avionics data network. The computer-based tool to model the network was SES/Workbench.

Contents

(12 pages in total)

Simulation of a cell switched network for the control of a switch matrix in a high-speed avionics network

David A. Aupers, Gerald J. Heerink, and Steven Wellink
National Aerospace Laboratory (NLR)
P.O. Box 90502, 1006 BM Amsterdam, Netherlands
aupers@nlr.nl
http://www.nlr.nl

1 ABSTRACT

This paper describes the research and experiments carried out by the National Aerospace Laboratory (NLR) in the field of high-speed interconnection systems for modular avionics. The research has been carried out in the EUCLID/RTP4.1-framework.

The avionics network that was modelled and simulated was an optical switch matrix under control of a cell switched network. The optical switch matrix offers the avionics system circuit-switched, uni-directional, point-to-point connections. A bandwidth of 2 Gbps is projected. The main purpose of the matrix is to connect sensors producing high data rates, such as an attack radar in fighter aircraft, with the core avionics processing cluster.

The cell switched network - in this case Asynchronous Transfer Mode (ATM) - controls the optical switch matrix and provides data transfer at lower data rates, file transfer, and status messages. The simulation model operated ATM at 149 and 622 Mbps.

The primary objective of our research was to assess ATM as a data link layer for a control and message network in an avionics data network. The computer-based tool to model the network was SES/Workbench.

2 ABBREVIATIONS

AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ATM	Asynchronous Transfer Mode
B	Byte
B-ISDN	Broadband ISDN
CBR	Constant Bit Rate
CCITT	Consultative Committee for International Telegraphy and Telephony
CMN	Control and Message Network
DMA	Direct Memory Access/Addressing
EO	Electro-Optical
EUCLID	European Co-operation for Long term In Defence
Gbps	Giga bits per second
Hz	Hertz
ISDN	Integrated Services Digital Network
ITU-T	International Telecommunications Union, Telecom Standards Sector
kB	kilo Byte
LAN	Local Area Network
LCE	Link Control Element
Mbps	Mega bits per second
NLR	National Aerospace Laboratory
OSM	Optical Switch Matrix
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RF	Radio Frequency
RISC	Reduced Instruction Set Chip
RTP	Research and Technology Programme
SCI	Scaleable Coherent Interface
SDH	Synchronous Digital Hierarchy
SES	Scientific & Engineering Software Inc.
STM	Synchronous Transfer Module
VBR	Variable Bit Rate
VC	Virtual Circuit
WAN	Wide-Area Network
WEAG	Western European Armament Group

3 INTRODUCTION

This paper describes the experiments and the results of research in the field of high-speed interconnection systems for modular avionics. This research has been carried out in the framework of EUCLID RTP 4.1.

3.1 Project background

The European Co-operation for Long term In Defence (EUCLID) Research and Technology Programme 4.1 "Modular Avionics Harmonisation Study" identified and researched the technologies available in Europe for the development of future avionics systems architectures. The programme was a joint effort of 27 companies in 6 European nations: France, Germany, United Kingdom, Spain, Italy, and the Netherlands. The consortium consisted of most European airframe manufacturers and equipment suppliers. EUCLID is a programme of the Western European Armament Group (WEAG).

The in-service time frame of the envisioned avionics systems was 2005-2010. The target programme can either be a retrofit of an existing aircraft or the development of a new aircraft. The types of activities in the programme involved definitions, specifications, surveys, simulations, and laboratory demonstrations.

The areas in which the National Aerospace Laboratory (NLR) was involved covered the following topics:

- high-speed interconnection systems;
- digital signal processing;
- fault-tolerance;
- component and rack cooling;
- system development tools.

This paper focuses on our activities in the field of high-speed interconnection systems, the modelling and simulation thereof in particular.

3.2 Modular avionics architecture

The avionics architecture defined in the programme formed the basis for the simulation model. The core avionics architecture consists of ten functional areas and a unified data network interconnecting the functional areas. The ten functional areas are vehicle control, crew interface control, mission control, systems control, data base control, RF, EO, image analysis, image generator, and acoustics. Each functional area hosts a group of related functions to optimise the traffic across the network. Reference 1 describes in detail the rationale for the division of the core avionics into functional areas.

The following modules are the building blocks for the functional areas: data processing, signal processing, image processing, graphics processing, and memory modules. With the continuous increase in performance of processing devices, it is likely that eventually all processing takes place on generic processing modules.

Table 1 on page 16-8 shows the expected data traffic categories and their characteristics (Ref. 2). These categories and characteristics formed the basis for the workload for the simulations.

Analysis of the data traffic shows that the avionics network shall support three basic types of transmissions:
1. sustained, large amounts of data;
2. bursty, medium sized amounts;
3. short, but time critical messages.

To be able to service this variety of transmission types, a dual network approach was chosen. The dual network is called the 'Matrix Switched Network' (MSN). The MSN provides:
- a connection-oriented data transfer network for sustained, large amounts of data, typically originating from sensors;
- a control and message network to control access to the data transfer network and to facilitate transmission of bursty, medium sized amounts of data.

For the control and message network, the following protocols have been evaluated: 1553, FDDI, ATM, and SCI. ATM came out as most promising candidate, closely followed by SCI.

Because of the limited amount of resources we were able to model one type of protocol. For several reasons we decided to go for ATM:
- ATM came out of the evaluation as most suitable;
- ATM-technology is available on the market;
- there are several commercial as well as academic models available.

3.3 Objectives of the modelling and simulation

Our research involved the modelling and simulation of a typical functional area with the following three objectives:
1. Development of a model of the core avionics architecture defined in the programme.
2. Performance modelling of the avionics architecture model.
3. Assessment of ATM as data link layer for a control and message network for an avionics data network.

4 DESCRIPTION OF THE NETWORK MODEL

Before explaining how an ATM network can be used as a data link layer for a control and message network, an introduction to ATM networks will be given in section 4.1. Section 4.2 explains how an ATM network can be used as a basis for the control and message network. Section 4.3 describes the limitations of the model. Section 4.4 describes the simulation tool SES/Workbench briefly.

4.1 Introduction to ATM

In the mid-1980s when the ISDN standard was being developed, the CCITT began working on the successor of ISDN; it was acknowledged that ISDN would not offer enough bandwidth in the future. This successor is known as Broadband ISDN (B-ISDN). One key objective was to develop a technology that would allow for efficient transport of all kinds of traffic (bursty and isochronous). Further, the new technology should support future speeds of several Gigabits per second (Gbps). In 1988 the CCITT decided to base the development of B-ISDN on ATM which was formalised in the late 1980s. B-ISDN became one of the services that can use ATM technology.

ATM is a relatively new method to transport information. Two classical ways of transporting information are:
- Circuit switching: requires a circuit to be established prior to transport of data. Resources in the network stay reserved until the connection is torn down. Circuit switching is well suited for isochronous traffic. ISDN and the classical telephone network are examples of the use of circuit switching.
- Packet switching: suitable for bursty data transmission and unsuitable for isochronous applications. It is more efficient than circuit switching, because network resources are only used when traffic is present. Packet switching is used in LAN environments.

ATM is a cell switching technique. Cells are small, fixed-length packets of 53 Bytes that are switched to their destination by the hardware in network nodes (ATM switches). Cells can carry data from arbitrary applications (isochronous as well as bursty). ATM systems are connected to ATM switches by a dedicated link; there is no shared medium like in LANs. This means that distinct pairs of ATM systems can communicate at full wire speed with each other (if the switch has enough switching capacity). A switch can be equipped with different types (speeds) of ATM ports; this way a server on ATM can have a faster connection to the ATM network than its clients.
Before data can be transported, a Virtual Circuit (VC) has to be established between the two end-points that wish to communicate (ATM is connection-oriented). An application can negotiate a QoS required for its VC. An ATM system

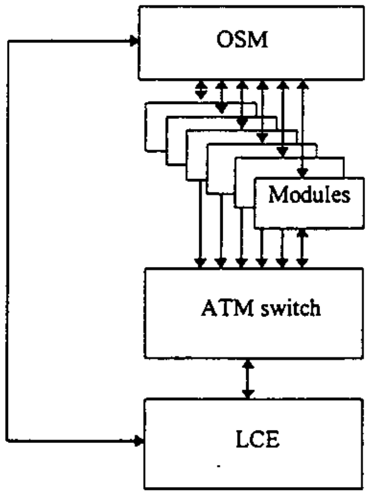

Figure 1 ATM as a control and message network

typically may use up to several thousands of Virtual Circuits simultaneously to different other ATM systems.

ATM supports four classes of traffic. Ordered in a decreasing priority the traffic classes are:

Class A Constant Bit Rate (CBR), connection-oriented, synchronous traffic (uncompressed voice or video)
Class B Variable Bit Rate (VBR), connection-oriented, synchronous traffic (compressed voice or video)
Class C Variable Bit Rate (VBR), connection-oriented, asynchronous traffic (X.25, Frame Relay)
Class D Available Bit Rate (ABR), connectionless, packet data (LAN traffic)

ATM is scaleable regarding both bandwidth and topology. Speeds are supported from 2 Megabits up to several Gigabits per second. ATM is often run over a physical layer consisting of one of the standards from the SDH hierarchy of optical standards. The hierarchy ranges from STM-1 (155.52 Mbps) up to STM-16 (2.4 Gbps) while even faster standards are being developed.

ATM is suitable in LAN as well as in WAN environments. LAN and WAN connections differ regarding available bandwidth. That is why congestion and flow control are important issues in large ATM networks. The ATM Forum and the ITU-T (former CCITT) are currently working on standards to address these issues.

4.2 The model and its traffic

Figure 1 shows schematically how an ATM network can be used as the control and message network for the OSM.

The functional area that is modelled contains 6 modules that are all connected to both the OSM network via optical links and the CMN network (in this case implemented by an ATM network) via an ATM network interface to which an optical or electric link is attached. The OSM controller (LCE) is also connected to the CMN network.

Note that in this set-up, modules can not only communicate with the LCE, but also directly with each other by using a direct ATM virtual circuit between them without bothering the LCE.

Four kinds of traffic will be simulated in the model. These will be explained in the following sections.

4.2.1 Commands between modules and LCE

Each module can issue commands to the LCE to set up or tear down an OSM connection with other modules. The time between the transmission of the request and the moment at which transmission of data on the OSM connection can start, is called the link time. For the so-called unlink time (for tearing down a connection) a similar definition is valid. A driving requirement was that the (un)link time had to be less than 50 μs.

Several high-level protocols have been considered for accomplishing a reliable connection set-up. To minimise the link time, the protocol in Figure 2 was chosen. The protocol works the following way:

Suppose module A wants to set up an OSM connection with module B. Module A sends a connection request to the LCE. After receiving the request, the LCE checks whether module B is available for the requested connection. If not, the LCE sends a negative response to module A. If module B is available, the LCE sends a message to module B to inform about the OSM connection that is about to be activated. At the same time the LCE sends a positive response to module A and starts setting up the OSM connection. When module B receives the message from the LCE, it sends a (positive) acknowledgement to module A. When module A has received positive messages from both the LCE and module B, it may start transmitting data via the OSM connection.

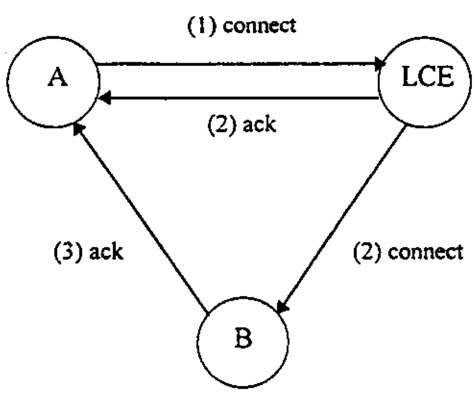

Figure 2 OSM-command Protocol

For the simulation it has been assumed that 3 modules maintain the same, static OSM connection configuration (e.g. continuous high bandwidth demanding sensor processing). The remaining 3 modules are resources for which is competed. They randomly issue OSM-commands. These requests are exponentially distributed with a mean of 10 Hz. In a real network the mean OSM-command release rate is probably lower than 10 Hz. Because the link time of an OSM-connection is only worthwhile when relatively large amounts of data have to be transported. The message length is 25 B (fits into 1 ATM cell). Because of the need to minimise the link time, this traffic is assigned to the VBR traffic class and not to the low priority ABR class..

4.2.2 Status (synchronisation) traffic

It is assumed that each module periodically sends synchronisation or status data to the LCE for configuration management purposes. These are small messages (25 B) that fit into 1 ATM cell. They are generated with a triangular distribution with a mean of 1 ms (1000 Hz), a minimum of 0.8 ms and a maximum of 1.2 ms. This traffic is assigned to the ATM ABR traffic class.

4.2.3 Control/data traffic between modules

Applications use a higher layer protocol to synchronise their activities and to exchange information. This dynamic behaviour depends on the functionality and implementation of the modules. The dynamic behaviour is modelled by the following random parameters:

- message size;
- transmission interval;
- source module;
- destination module.

The size of the messages is uniformly distributed between 1 kB and 16 kB. The messages are generated with an exponential distribution with a mean of 5 ms (200 Hz). The source and destination modules are chosen according to a uniform distribution. This traffic uses the ATM ABR traffic class.

4.2.4 File transfer traffic between modules

Modules can exchange certain amounts of data for which it is not effective to request an OSM connection or when the desired OSM connection is unavailable. This data can be transported by means of a file transfer using the CMN. This results in a burst of maximum sized packets between two modules. The dynamic behaviour is modelled by the following random parameters:

- message burst size;
- source module;
- destination module.

The modules are chosen according to a uniform distribution, just like the file-size (between 64 kB and 192 kB). A file burst is generated every 0.1 s (10 Hz) and is assigned to the ATM ABR traffic class.

4.3 Abstractions and limitations of the model

This section describes limitations and abstractions of the model when compared to a possible real world implementation.

1. *Only Permanent VCs are used in the model.*
   The process of dynamically (on demand) setting up an SVC (Switched VC) can take milliseconds in a real ATM network. In the avionics system being modelled, such a delay is intolerable. Hence, it is assumed only PVCs (Permanent VC) are used. In a real implementation these can be set up automatically during system initialisation. As a consequence an ATM node can start transmitting data immediately; it is not necessary to set up a VC first.

2. *All links have the same bandwidth.*
   In an ATM network it is possible for a node that will receive/transmit more data than other nodes to have a higher capacity network-connection. Since in the model the traffic is fairly well distributed, an optimisation like this is not used. Simulations are run for ATM networks based on SDH STM-1 and SDH STM-4.

3. *Physical layer overhead not modelled properly.*
   In an ATM network based on SDH there is some overhead at the physical layer. On an STM-1 (155.52 Mbps) trunk every 27th cell is needed for that overhead limiting the available bandwidth to 149.76 Mbps. This is modelled assigning an overall available bandwidth of 149.76 Mbps to the ATM trunks; in stead of reserving every 27th cell. For STM-4 (622.08 Mbps) every 108th cell is not available, resulting in an available bandwidth of 616.32 Mbps.

4. *ATM interface processing overhead not modelled.*
   Of course, some processing needs to be performed at an ATM interface. ATM Adaptation Layer (AAL) headers/trailers must be added or removed. Packets of data have to be segmented/reassembled to/from cells. In state-of-the-art ATM adapters dedicated hardware is used to obtain a minimum latency (64 bit RISC processors, DMA, etc.). Data is transferred from/to the host memory while the cells of a packet are being transmitted/received to/from the ATM network. Latency introduced by a carefully designed ATM adapter is small when compared to the total latency of transferring a message through the ATM network.

5. *Higher layer protocols are not modelled.*
   The objective of the simulation was to focus on the ATM level of the CMN. Because of this, no higher layer protocols have been modelled. As a consequence no higher layer protocol headers have been taken into account when decreasing the maximum packet size during consecutive simulation

runs. As another, more serious consequence, no flow control is available. This means that all data for a file transfer enters the ATM-interface of a module as maximum-sized packets simultaneously.

6. *VC and their QoS parameters are not modelled.* In a real ATM network all data offered to an ATM network interface must be transmitted on a pre-established VC while respecting the QoS parameters that were agreed upon during the VC set-up ('traffic shaping'). The ATM model did not have options to specify other QoS parameters than the ATM traffic class. All data offered to an ATM network interface is transmitted as fast as possible (at the speed of the trunk connected to it). This is slightly worse than in the real world and increases the probability of cell loss in the switch.

7. *Packet transmission is considered un-interruptable.* It is desirable that the transmission of a (possibly large) low priority packet (e.g. ABR) is interrupted because a higher priority message (e.g. VBR) is offered for transmission to the ATM interface. Depending on the implementation of an ATM adapter and whether (and in what way) it enforces traffic shaping; this may be possible in a real ATM adapter. It is not included in the model. As a result the latency of high priority messages depends on the maximum packet size for lower priority messages.

4.4 Short description of the simulation environment

SES/Workbench is a graphically oriented general-purpose simulation language that contains features for modelling computer systems and communication networks. The graphical interface allows users to build and represent designs pictorially. The major building blocks are:

- Nodes
- Arcs
- Transactions

There four basic types of nodes:

Resource management nodes
Resource management nodes create, allocate and release the resources used by transactions. The resources may be processors, memory, communication links, busses and system processes.

Transaction flow control nodes
Transaction flow control nodes create, destroy, and alter flow of transactions through the system.

Sub-model management nodes
Sub-model management nodes allow a model to be developed and specified as a hierarchical collection of sub-models.

Miscellaneous nodes
Such as, user-defined nodes.

A collection of building blocks represents system components, processors, resources, transaction flows, and others. To build a model one defines a transaction that corresponds to a message. After that, a directed graph consisting of nodes and arcs is created. This is done by placing icons on the display that can be connected by arcs. The arcs and nodes describe how transactions flow through the model. In this way it is easy to create and view complex models.

SES/Workbench provides real-time animation that displays transactions flowing through the model and shows events that occur at nodes. Simulation results can be displayed in both numerical and graphical formats, either during the simulation or after it has been completed.

Several statistical functions, built-in probability density functions, and queuing disciplines are available.

5 EXPERIMENTS

5.1 Measurements

For the experiments the following statistics were measured:

- mean ATM-utilisation for:
    - network interfaces/links;
    - the ATM-switch;
- ATM-switch lost-cells;
- mean OSM-command response-time (including acknowledgements);
- mean status-message response-time (including acknowledgements);
- mean file-burst response-time (including acknowledgements).

5.2 Parameters

To investigate the model, the following parameters were varied:

• ATM-bandwidth	149 Mbps, 616 Mbps
• Workload	nominal, high (5 times nominal)
• Maximum packet-size	100%, 40%, 12.5%, 6.25% of control/data traffic and file-burst maximum packet-size

ATM-bandwidth and workload parameters were used to vary the traffic and stress of the ATM-network. The nominal workload described in section 4.2 results in a mean control and data traffic load of 13.67 Mbps and a mean file-burst load of 10 Mbps. The high workload approximately produces 5 times more traffic than the nominal workload: a mean control/data traffic load of 68.36 Mbps and a mean file-burst load of 50 Mbps. The high load control/data traffic is created by increasing the release-rate. The distribution in time of the extra packets is uniformly. The high load file-burst is created by increasing the range, from which the size of the file-burst is uniformly chosen, from [64 kB, 192 kB] to [320 kB, 960 kB]. The resulting extra file-burst traffic enters the network simultaneously.

As explained in sections 4.2 and 4.3, the larger the allowed packet-sizes, the higher the latencies of messages can become. For this reason, the maximum packet-size is varied during the experiments. For file-burst messages a smaller packet-size will result in more (but smaller) packets being generated simultaneously. The model does not include the effect of additional overhead needed to reassemble messages from multiple smaller packets. The OSM-commands and status-messages are not varied. As described in section 0 and 4.2.2 they always occupy 1 ATM cell.

5.3 Experiments

The following groups of experiments will be described:

- Reference experiments with single messages;
- experiments with an ATM-bandwidth of 149 Mbps during nominal operation;
- experiments with an ATM-bandwidth of 616 during nominal operation and operation under high loads.

All results are mean values over the complete simulation period and times are reported in µs.

5.3.1 Reference experiments

The response-times of an OSM-command message, a single-cell status-message and file-burst messages were determined, while no other messages were in the network. Note that all measured response-times include the latencies of acknowledgements being sent back to the source of the original messages. The service times in the involved modules and LCE were fixed to their mean service times (10 µs).

	149 Mbps ATM	616 Mbps ATM
Status-message	33.2	15.6
OSM-command	54.8	28.4
64 kB file-burst	3900	955
640 kB file-burst	42600	10400

The single-cell status-message response-time includes the following latencies:

- Four cell transmissions (message and acknowledgement) from the network interface to the ATM-switch and vice versa. This is the time needed to put bits of a cell on a link.
- Four link propagation-delays (2 for the message and 2 for the acknowledgement) from the source network interface to the ATM-switch and from the ATM-switch to the destination network interface. For the experiments all links had a length of 1 meter.
- Two switch delays (message and acknowledgement). This is the time to move the cell through the switch-fabric from the input port to the output port.
- One service time (10 µs) needed in the destination module to produce the acknowledgement.

The table shows that the 50 µs OSM (un)link-time requirement cannot be achieved with the 149 Mbps ATM-bandwidth.

5.3.2 149 Mbps experiment

Description:

OSM-command, status-message and file-burst mean response-times, during nominal load, 149 Mbps ATM, a 20 second simulated time and triangular distributed services times for modules and the LCE with a minimum of 5 µs, a mean of 10 µs and a maximum of 15 µs.

Parameters:

packet-size (in percentages).

Results:

Size	OSM-command	Status	File-burst
100	147.6	167.2	14678
40	81.4	125.4	9674
12.5	66.5	101.5	8408
6.25	59.1	89.9	8507

Figure 3 shows the measured response times of the OSM-commands and status messages with the ATM network opating at 149 Mbps and with a nominal workload.

Conclusions from the measurements:

- smaller packet sizes reduce the response-times;
- the improvement of the mean file-burst response-time with smaller packet-sizes is because of a decrease in packet-size of the control/data traffic. This decrease causes the control/data traffic load to be more uniformly distributed in the functional area (space) and in time;

The mean OSM-command response time during high load, 149 Mbps ATM with 100% packet size was 400 µs. (Because this result is far from the desired 50 µs, further experiments were concentrated on 616 Mbps ATM-bandwidth experiments.)

The following ATM-network statistics were measured:

Utilisation of:	Nominal load	High load
ATM switch	1.5%	6.2%
Module Net Interface	3.6%	15.5%
LCE link	1.8%	1.8%

The experiments showed that when occurrences of file-bursts overlap in time and space (to the same destination-module), cell-loss can occur even during nominal load. (In such a case the ATM-switch output buffer to the involved destination-module is easily congested.)

5.3.3 616 Mbps experiments

Description:

OSM-command, status-message and file-burst mean response-times, during nominal and high load, 616 Mbps ATM, a 20 second of simulated time and triangular distributed services times for modules and the LCE with a minimum of 5 µs, a mean of 10 µs and a maximum of 15 µs.

Parameters:

packet-size (in percentages);

load (nominal or high).

Results:

Nominal load:

Size	OSM-command	Status	File-burst
100	36.2	23.8	2863
40	29.0	20.5	2367
12.5	29.2	19.9	2069
6.25	28.5	19.5	2024

Figure 4 shows the measured response times of the OSM-commands and status messages with the ATM network opating at 616 Mbps and with a nominal workload.

High load:

Size	OSM-command	Status	File-burst
100	50.2	150.5	12018
40	39.9	91.9	9911
12.5	32.6	150.6	12482
6.25	30.7	101.5	11109

Figure 5 shows the measured response times of the OSM-commands and status messages with the ATM network opating at 616 Mbps and with a high workload.

Conclusions from the measurements:

- With the high load, 94% of the OSM-commands are processed within 50 µs. (This statistic is not shown in the tables).
- With the nominal load the packet size has only minor influence on the response-times, because packets that block the network interface of a sending module (or the ATM-switch) for packets that follow, are served 4 times faster in the ATM-network than during the 149 Mbps experiments;
- The irregular shape of the graph of the high load file-burst response-time may be caused by file bursts that overlap in time and/or space. One of the following scenarios might have occurred:
    1. Bursts overlap in time and are transmitted from the same module. Because all bursts have the same priority, the second burst is delayed until the first burst has been transmitted;
    2. Bursts overlap in time and are transmitted from different modules, but to the same module. This causes both extra delays and cell-loss in the ATM-switch. During the experiments with nominal load, almost no cell-loss occurred.

Because of the relatively short simulated time of 20 seconds an occurrence of one of these scenarios has a large impact on the shape of the graph. Inspection of the collected statistics showed that the experiments that were responsible for the peaks in the high load graphs suffered from severe cell-loss when compared to the other experiments in the same graph. This may indicate that scenario (2) is responsible.

- The irregular shape of the graph of the high load status-message experiments is similar to the shape of the high load file-burst experiments. The status-messages suffer from the file bursts the most, because status messages have the same priority as file bursts, while the OSM-commands have a higher priority.

The following ATM-network statistics were measured:

Utilization of:	Nominal load	High load
ATM switch	0.4%	1.5%
Module Net Interface	0.9%	4.0%
LCE link	0.4%	0.4%

6 CONCLUSIONS AND RECOMMENDATIONS

From the simulation experiments the following can be concluded.

1. During high load with a maximum packet-size of 64 kB, 94% of the OSM-commands (link or unlink commands) are processed within 50 µs when:
    - at least a 616 Mbps bandwidth is used for the ATM-links and network interfaces;
    - an ATM-switch is used with a 9.86112 Gbps aggregate bandwidth, with a typical switch fabric latency of about 5 µs;
    - the module and LCE service times approximate a triangular distribution with a minimum, maximum and mean of respectively 5 µs, 15 µs, and 10 µs;
    - OSM-commands have priority over other packets in the ATM network interface and other cells in the ATM-switch.

2. Because the status traffic introduces only minor workload (thus minor latency for lower priority messages) and it is important that status-messages have a low latency, it is recommended that these messages have a high ATM-priority like the OSM-commands.

3. It is recommended that both the ATM network interfaces and the ATM-switch have separate output-buffers for cells with different priorities to make the latency of high priority messages independent of the packet-size of lower priority messages.

4. Traffic-bursts such as simulated in the model should be suppressed or controlled to prevent:
    - that the network interface of a module is blocked for other traffic;
    - that the packets of a burst are transmitted one-after-the-other without gaps, causing severe load-peaks. For this purpose higher layer protocols could be used, that apply flow control, e.g. sliding-window mechanisms, and at the ATM-layer Virtual Circuits for each traffic-stream with properly configured QoS-parameters to enforce traffic shaping.

5. Because it is expected that the bandwidth of ATM-networks will be increased significantly in the near future and because in an ATM-network different types of data can be transferred with different QoS, it should be considered to transfer the high bandwidth data via the ATM-network as well. Because the OSM would no longer

be necessary, the complexity of the avionics network would be reduced significantly.

7 REFERENCES

1. A. Marchetto, "Modular Avionics System Architecture Definition in the EUCLID Research and Technology Programme 4.1: Methodology and Results", *presented at* "AGARD Mission Systems Panel 6th symposium on Advanced Architectures for Aerospace Mission Systems", 14-17 October 1996, Istanbul, Turkey.
2. Neue Avionikstruktur, Dok-Nr.: 1LN1-BT-0008, "Zusammenfassung der Ergebnisse", ESG, Issue 1, 16 July 1988.

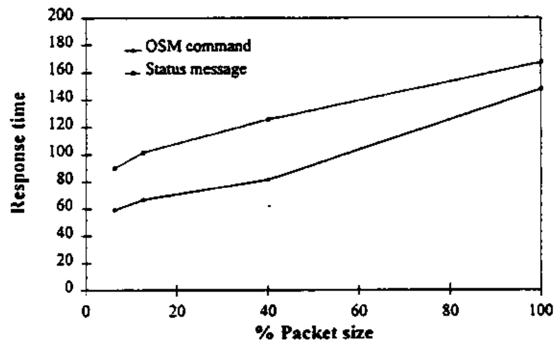

Figure 3 Measured response times with 149 Mbps ATM with nominal workload

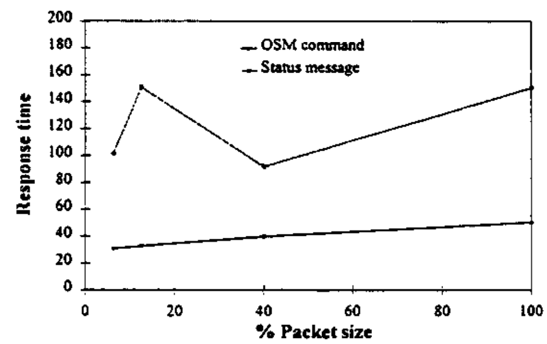

Figure 5 Measured response times with 616 Mbps ATM with high workload

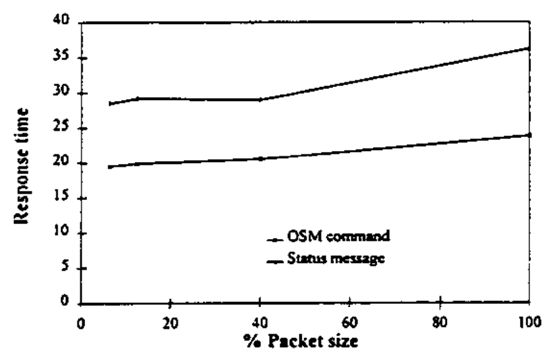

Figure 4 Measured response times with 616 Mbps ATM with nominal workload

Table 1 Data categories and characteristics

Parameter	Video	Fast sensor	Medium sensor	Slow sensor	Control/data	Sync	File transfer
Data rate - applications - frame length - rate	2 Gbps Video 25 Mbits 80 Hz	2 Gbps Radar 2 Mbits 1 kHz	750 Mbps Beam steer 64 kbits 10 kHz	250 Mbps E/O data 5 Mbits 50 Hz	< 1 Mbps Various 32 bits - 132 kbits 50 - 200 Hz	< 1 Mbps Various < 100 kbits 50 Hz - 1 kHz	< 1 Gbps Various 1 Mbit
Periodic/aperiodic	periodic	periodic	periodic	periodic	both	periodic	aperiodic
Persistence	10s of s	10s of s	10s of s	10s of s	10s of s	10s of s	message length
Latency - bit - frame	 10s of ms	 5 µs	 5 µs	 5 µs	 100 µs	 10 µs	 1 ms
Time tagging	no	yes	yes	yes	-	-	no
Topology	point/point	point/point	point/point	point/point	multipoint	multipoint	point/point
Delivery guarantee	Bit errors detected	Bit errors corrected	Bit errors corrected	Bit errors corrected	Frame acknowledgemt	Frame errors corrected	Frame acknowledgment