

REAL-TIME TELEMETRY DATA INTERFACE TO GRAPHICS WORKSTATION

Amy Sidorovich
NAWCAD
Telemetry Division
Atlantic Ranges and Facilities
Patuxent River, Maryland 20670

KEY WORDS

Data Interface, Replicated Memory Network

ABSTRACT

The demand for additional computing power and more sophisticated graphics displays to strengthen real-time flight testing prompted the Real-time Systems Team to turn to graphics workstations. In order to drive graphics displays with real-time data, the questions became, "What interface to use?" and "How to integrate workstations into our existing telemetry processing system?". This paper discusses the interface and integration of graphics workstations to the Real-time Telemetry Processing System III (RTPS III).

INTRODUCTION

The task of integrating graphics workstations into an existing telemetry ground station is a challenge that is important to the growth of real-time flight testing. By adding new graphics displays, safety of flight data output will be increased and a more flexible graphics environment will be provided to the test engineer. With that challenge comes the issue of determining the data interface that will provide the most capability without disturbing the current system.

The Real-time Telemetry Processing System III (RTPS III) can support six concurrent real-time flights, providing a variety of data output devices in each of the six Project Engineer Stations (PES). With the onset of new aircraft projects, the decision was made to enhance each PES with the addition of two SGI Indigo2 Extreme workstations. Through investigation, it was determined that a replicated memory interface would best fit the current system configuration and provide the needed information to the graphics workstations. In this case, Systran's SCRAMNet Network was used.

The purpose of this paper is to discuss the integration of workstations to an existing telemetry processing system. There are three major facets discussed: determine the best interface to the workstation, integrate the workstations, ensure data integrity and throughput. The success of this integration will be demonstrated by graphics applications on the workstation driven with real-time telemetry and calculated data.

HISTORY

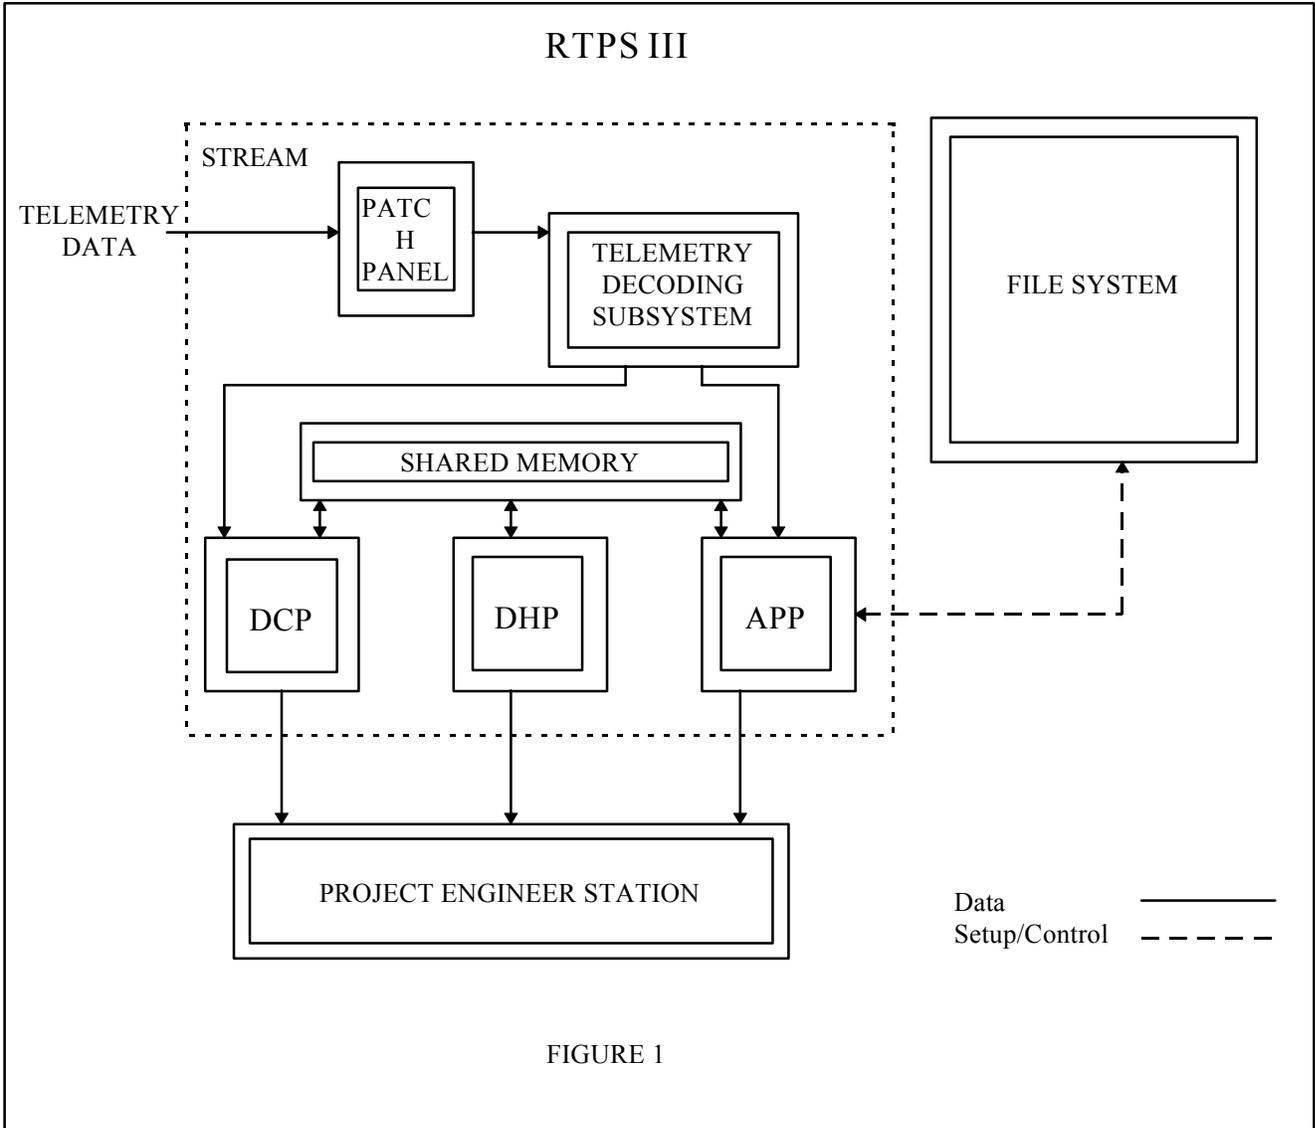
RTPS III was constructed by Computer Sciences Corporation of Lompoc, California and was delivered to Patuxent River, Maryland in 1988. The system was based on a government developed prototype system. It is made up of six independent systems (streams) configured with a common file system. Each stream consist of two Encore 32/67 minicomputers serving as a data channel preprocessor (DCP) and display host processor (DHP), an Encore 32/97 superminicomputer serving as applications processor (APP), an Aydin SG2000 based telemetry decoding subsystem, two Adage graphics processors, a vector processor and strip chart recorders. The three Encore processors are connected via a shared memory implementation. The applications processor on each stream is connected to the file system via a Hyperchannel A470 Network Adapter. (See Figure 1)

THE DECISION

Deciding what interface to use is dependent on many factors. To begin, the reasoning behind the addition of two workstations per stream was not only to provide additional graphics capability, but also to provide some additional computing capability. Therefore, the ability to have all incoming data available on the workstations and provided at the current throughput rate would be ideal. Since the interface to be implemented will go between the applications processor and the workstations, the decision would then also depend on the hardware of these computers and their common interface capabilities.

Data

A description of the data to be transferred will help illustrate the path that was chosen. RTPS III has the ability to process 2048 telemetry measurements through the telemetry decoding subsystem's Current Value Table (CVT) path. This data is passed directly to the applications processor and stored in a shared memory area referred to as the CVT. In addition, measurements that can be calculated from the telemetry data are computed and stored in the CVT. There are 512 memory locations available for storing calculated data, giving a total of 2560 32 bit memory locations in the CVT. The maximum throughput of the telemetry data is 250,000 32 bit words per second, (1 Mbyte/sec), aggregate.



Other Requirements

In addition to the need for 1 Mbyte/sec data throughput, there are several other requirements that the chosen data interface would have to meet. Due to the layout of RTPS III there would be a maximum distance of 120 feet between the APP and the workstations. This infers that the cable length specification of the interface would have to meet this distance. Since the purpose of the workstation addition is to provide more capability, the minimization of CPU overhead on the APP and the workstations is important. Also important is the minimization of cost and development. Finding solutions that take advantage of existing hardware and software available on each platform (Encore 32/97 and SGI Indigo2) could support this effort.

Modes of Transfer

There were two modes of transfer explored: program controlled data transfer and replicated memory data transfer. For the program controlled data transfer, ethernet was investigated. And, as mentioned before, the SCRAMNet communications network was investigated as a replicated memory data transfer solution. An examination of cost, the amount of development required, the capability each would provide and the effect each would have on the current system was applied to each.

Ethernet was considered because there were ethernet capabilities available for both the workstation and the applications processor. The workstation has a built in ethernet port and development had been done in the past using an ethernet interface card in the APP. The major hardware cost in this scenario would be the purchase of an ethernet card for each of the six applications processors. There are no cable length limitations, given ethernet can cover distances up to 1000 meters. There would be software needed on both ends to handle the transfer of data. This software would require time to develop and, more importantly, it would require a portion of the CPU bandwidth on each computer. Data transfer rates for ethernet, at 10 Mbits/sec (1.25 Mbytes/sec), could just meet the CVT throughput requirement, theoretically. However, due to the overhead of the protocol, ethernet could never achieve the 1 Mbyte/sec wanted. Its most effective use would be to transfer only the subset of data needed to drive low rate graphics displays.

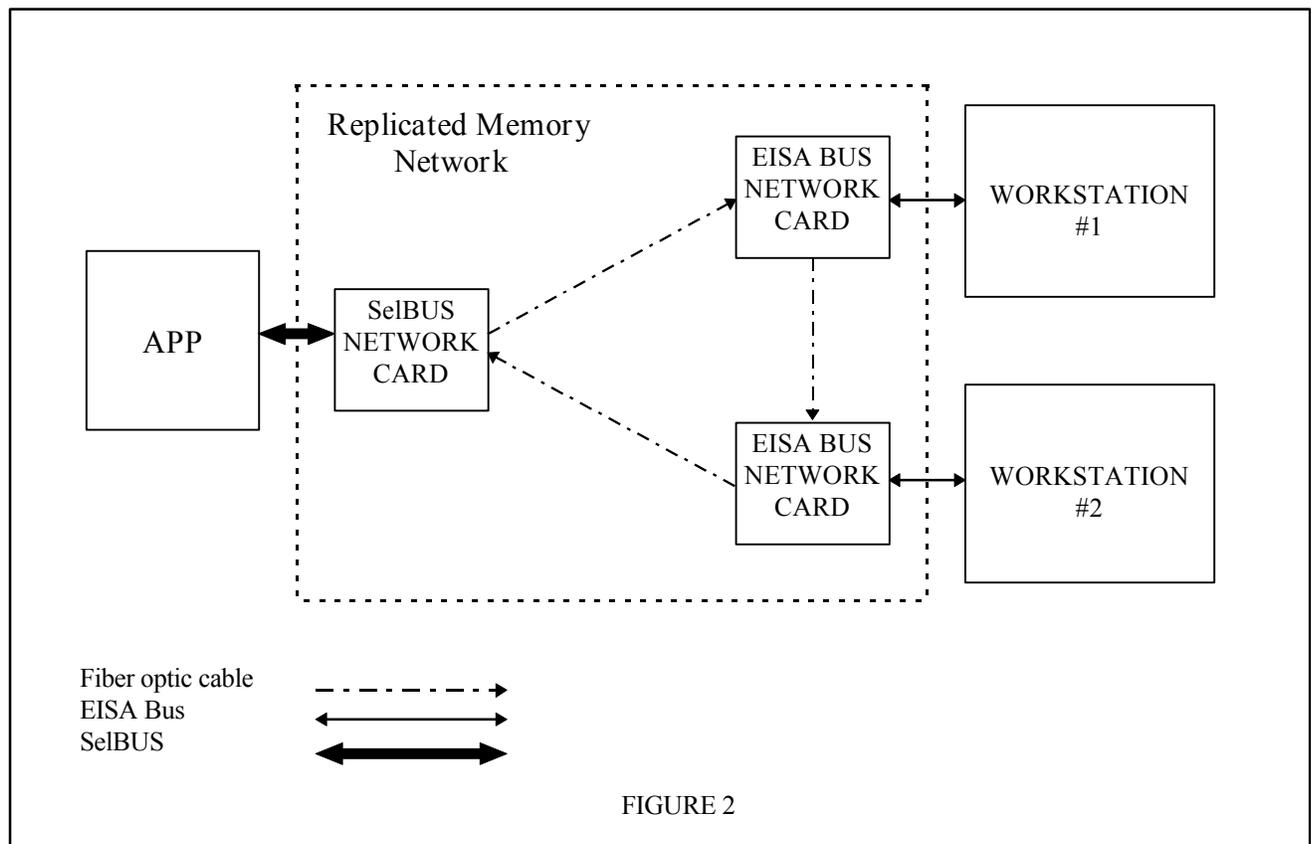
The promise of all the CVT data being transferred at its aggregate rate was enticing; and, the replicated memory option promised just that. The architecture of the two types of computers played more of a role in this situation. The APP has a SelBUS and the workstations are designed with an EISA bus. The solution was dependent on finding a network that could connect these two architectures. The SCRAMNet network provided the connection needed. Hardware costs in this situation would involve the purchase of one network card for each computer connected to the ring. One card plugs into the SelBUS of the applications processor and one card plugs into the EISA bus of each workstation. This would be a total of three cards per stream, for six streams. Although this would be the most expensive solution, there would be no software required to facilitate the transfer of data, which would save development time and preserve CPU bandwidth. There are no cable length conflicts since the network supports up to 1000 feet between nodes. The data throughput offered by this solution is 6.5 Mbytes/sec, which is more than enough to meet the 1 Mbyte/sec throughput requirement and provide room for expansion.

There were several other interface options considered of both the program controlled type and replicated memory type. However, these options were never fully investigated since they could not meet some of the basic requirements: hardware compatibility, device drivers for the workstation, agreeable cost.

From this investigation, it was found that program controlled data transfer would not satisfy the data throughput requirements. There is too much overhead due to the software involved to transfer data. Replicated memory was found to meet all of the requirements and did not hinder the existing system in any way. Therefore, it was decided that replicated memory would provide the best solution.

THE INTEGRATION

The replicated, shared-memory network consists of one card that resides on the SelBUS of the applications processor, one card that resides on the EISA bus of each workstation and fiber optic cables connecting each node into a ring. (See Figure 2) The transfer of data is generated from the applications processor to the workstations. Although data could be passed from the workstations to the applications processor, there was no immediate need to do so and development was pursued with only one node in the ring writing across the network. Since only one node would be transmitting, the common network protocol between the two types of network cards, Burst mode, would allow up to 6.5 Mbytes/sec data throughput.



Hardware

The applications processor network card resides on the SelBUS; however, the APP does not recognize the card in either the hardware or the software sense. This eliminates any load on the current system. There is no memory resident on the card; yet, it has the ability to monitor the memory write transfer activity on the bus. Any writes to a specified address range in memory on the APP can be recognized, then that same data is transmitted to the replicated memory network. Via hardware control switches, the base address and size of the memory that is to be replicated from the APP's memory to the network was established.

This range of memory to be replicated included the CVT and also a measurement description table (MDT). (See Figure 3) The MDT is a table of static information describing each measurement in the CVT. It is written once at system initialization.

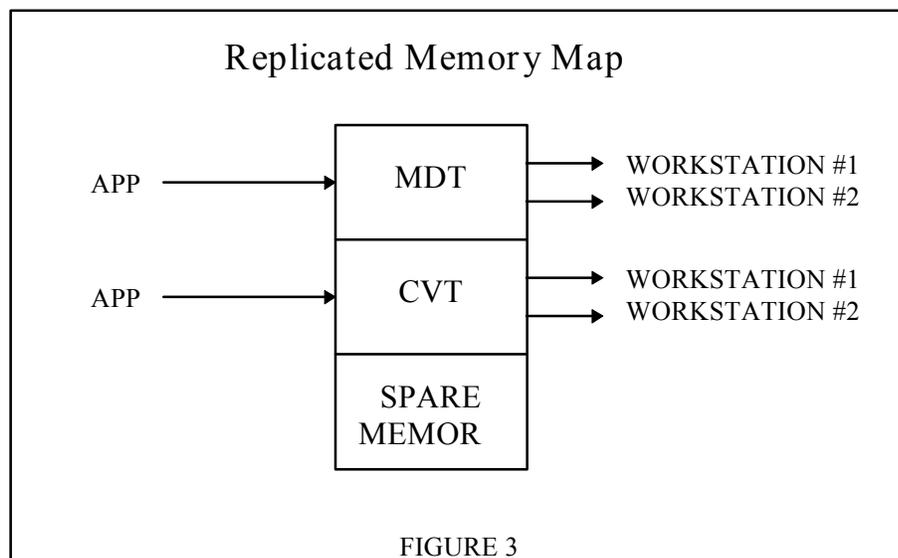


FIGURE 3

Each workstation network card in the ring has 2 Mbytes of memory (expandable up to 8 Mbytes) resident on the card where all network data is written. Therefore, both workstations on the ring have a copy of the MDT and a continuously updated copy of the CVT. Using provided software functions, the memory on the network card can be mapped onto the EISA bus. This makes access to the CVT and the MDT on the network card similar to access to local workstation memory. Initialization of this card is done using software functions provided by the vendor.

Software

On RTPS III there is software that allows the programmers to fetch data without an in depth knowledge of how and where the data is stored. A group of functions was developed on the workstation to provide the same ease of access to network data. These functions take care of mapping and unmapping the memory on the EISA network card. They provide a mechanism for fetching only a subset of measurement data from the CVT

needed for any one application; and, they provide floating point format conversion, if needed. Also, a function is provided that gives the programmer direct access to all of the variables defined in the MDT for any particular measurement.

THROUGHPUT AND INTEGRITY TESTS

There were three tests created to ensure the bandwidth and integrity of the replicated memory network: a no lost data samples test, a bandwidth test and a network memory access test. To ensure that data samples were not being lost during transmission, a test was devised to write once to every memory location being replicated and verify on the workstation end that the same value was received in the appropriate memory location. The result to this test was no lost data samples. The next test was to determine that the network would not be a bottleneck. Data was written to the CVT as fast as the applications processor's CPU would allow. This throughput rate was determined to be approximately 380,000 32 bit words per second (1.52 Mbytes/sec). No data was lost on the workstation end; as expected, given that the 1.52 Mbytes/sec test rate was below the network specified rate of 6.5 Mbytes/sec. The test showed that the network bandwidth would not be a limitation. The third test done determined the rate of access to the network card memory across the EISA bus. This access rate was approximately 1.44 Mbytes/sec. Since the CVT throughput was slower than this rate, there would be no hindrance to data access from the workstation. Each of these tests provided confidence that the interface would meet the demands of the CVT transfer rates. It should be noted, however, that the 1.44 Mbytes/sec memory access rate could be quickly outgrown if additional data buffers are added to the memory map.

CONCLUSION

Finally, the network ring is in place, software access to the CVT and the MDT has been established, data throughput and integrity have been verified and graphics applications have been written. The final test is to provide data to real-time applications running on the workstation. Using telemetry data recorded on tape from a previous flight, the graphics displays were successfully driven by both telemetry and calculated data; thus, the integration was complete.

Although the real-time data interface to the graphics displays was proven successful, there are still on going investigations to determine the effectiveness of using the graphics workstations for additional computing power. At the moment, the two major hurdles are accessing data across the EISA bus from the workstation and the inability to perform deterministic scheduling on the single CPU workstations.

Future plans include further investigation of the interrupt capability of the replicated memory network, the addition of an SGI Challenge (VME bus node) to the network and the addition of other data buffers to the memory map.

ACKNOWLEDGMENTS

The author would like to thank the Real-time Systems Team for their work and support. A special thanks to Thomas F. Hayes for his technical input.

REFERENCES

Hennessy, John L. and Patterson, David A., Computer Architecture A Quantitative Approach, Morgan Kaufmann Publishers Inc., 1990, pp. 526-528.

SCRAMNet-LX Network EISA Hardware Interface Reference , Systran Corporation, Dayton, OH, 1994.

SCRAMNet Network Reference Manual for the Model SE-1000 SelBUS Hardware Interface Node, Systran Corporation, Dayton, OH, 1993.