DISCRETE EVENT SIMULATION OF THE EOS-AM1 SCIENCE FORMATTING EQUIPMENT

Bruce O. Barringer
Fairchild Space and Defense Company
Germantown, Maryland

ABSTRACT

Fairchild is presently developing a high-rate telemetry collection and formatting component for one of NASA's Mission to Planet Earth's key missions. Because of the complexity and new technology involved, discrete event simulation tools have played a key role in the development process. This paper serves as a brief introduction to this component and to the model developed with the simulation tools.

KEY WORDS

Simulation, Modeling, Science Formatting Equipment, CCSDS

INTRODUCTION

In late 1992, Fairchild was awarded a contract to develop and manufacture the high-rate Science Formatting Equipment (SFE) to be included on NASA's EOS-AM1 spacecraft scheduled for launch in 1998. The SFE is responsible for collecting all of the spacecraft's science instrument data, formatting the data into frames compatible with the Consultative Committee for Space Data Systems (CCSDS) recommendations, and routing the formatted data to selected RF transmitters or solid state recorders.

Traditional telemetry formatting equipment, like that presently flying on the Fairchild built TOPEX spacecraft, has used synchronous, fixed size data structures to frame downlinked telemetry that is collected at fairly modest data rates. The timing within this equipment is fixed and easily analyzed. The EOS-AM1 spacecraft instruments can generate data packets at aggregate rates exceeding 100 Mbps and the downlinked telemetry is required to adhere to the CCSDS recommendations. CCSDS dictates the use of asynchronous variable length data structures making the analysis much more difficult.

Because of these complexities, Fairchild has chosen to implement and maintain a detailed discrete event simulation (model) as an integral part of the SFE design process. The results obtained from this model have been crucial in mitigating the risks associated with the hardware and software development.

OVERVIEW OF THE SFE

To understand how the model generates useful results one must first know a little about how the SFE operates. A block diagram of the SFE is shown below in Figure 1.

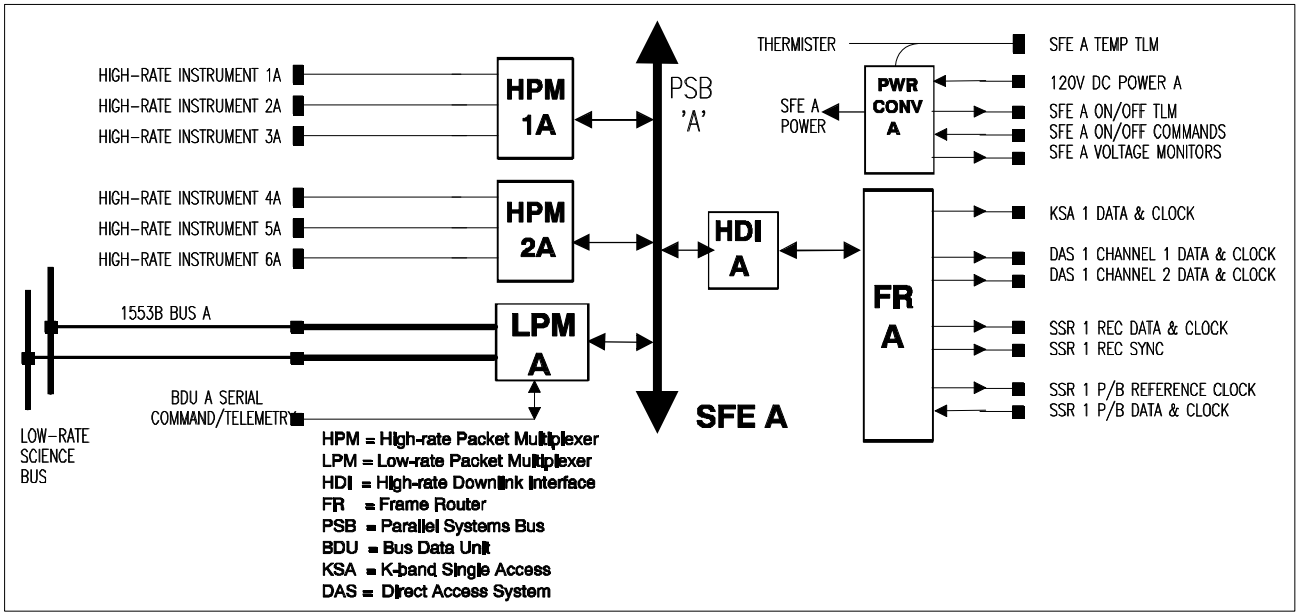

Figure 1: Block Diagram of the SFE

This block diagram functionally shows half of a redundant configuration. On the input (left-hand) side of the figure, it can be seen that the SFE interfaces to 2 types of instruments - high-rate instruments and low-rate instruments. High-rate instruments utilize dedicated, point-to-point electrical links to output their packetized data to a High-rate Packet Multiplexer (HPM) card at rates up to 50 Mbps. Each HPM card can service up to 3 such interfaces. Low-rate instruments utilize a 1553B serial data bus to transport their packets to the SFE's Low-rate Packet Multiplexer (LPM) card. The LPM card also provides the Command & Telemetry interface to the spacecraft computer. As implied by the name, the Packet Multiplexer cards multiplex the received instrument data packets into fixed-length, CCSDS data structures called Virtual Channel Data Units (VCDUs) that are routed to the High-rate Downlink Interface (HDI) card via a backplane. The HDI prepends header information and appends Reed-Solomon check bits to the frame to form Channel Access Data Units

(CADUs) which are then routed to 1 or more RF transmitter or Solid State Recorder outputs via the Frame Router (FR) card.

As mentioned above, the HPM, LPM, and HDI cards utilize a backplane interface to communicate with one another and to transport data. The industry standard Multibus II Parallel System Bus (PSB) was chosen to serve this function. The 32-bit wide PSB has a gross data transfer capability of 320 Mbps. One of the key reasons for choosing the PSB was the availability of a very efficient, low overhead transport mode called solicited messaging. It is this mode that is used to transport the frames between the SFE's Packet Multiplexers and the HDI cards.

As might be gathered from the discussion above, one of the more challenging design efforts was the HPM card. A block diagram of this card is shown below in Figure 2.

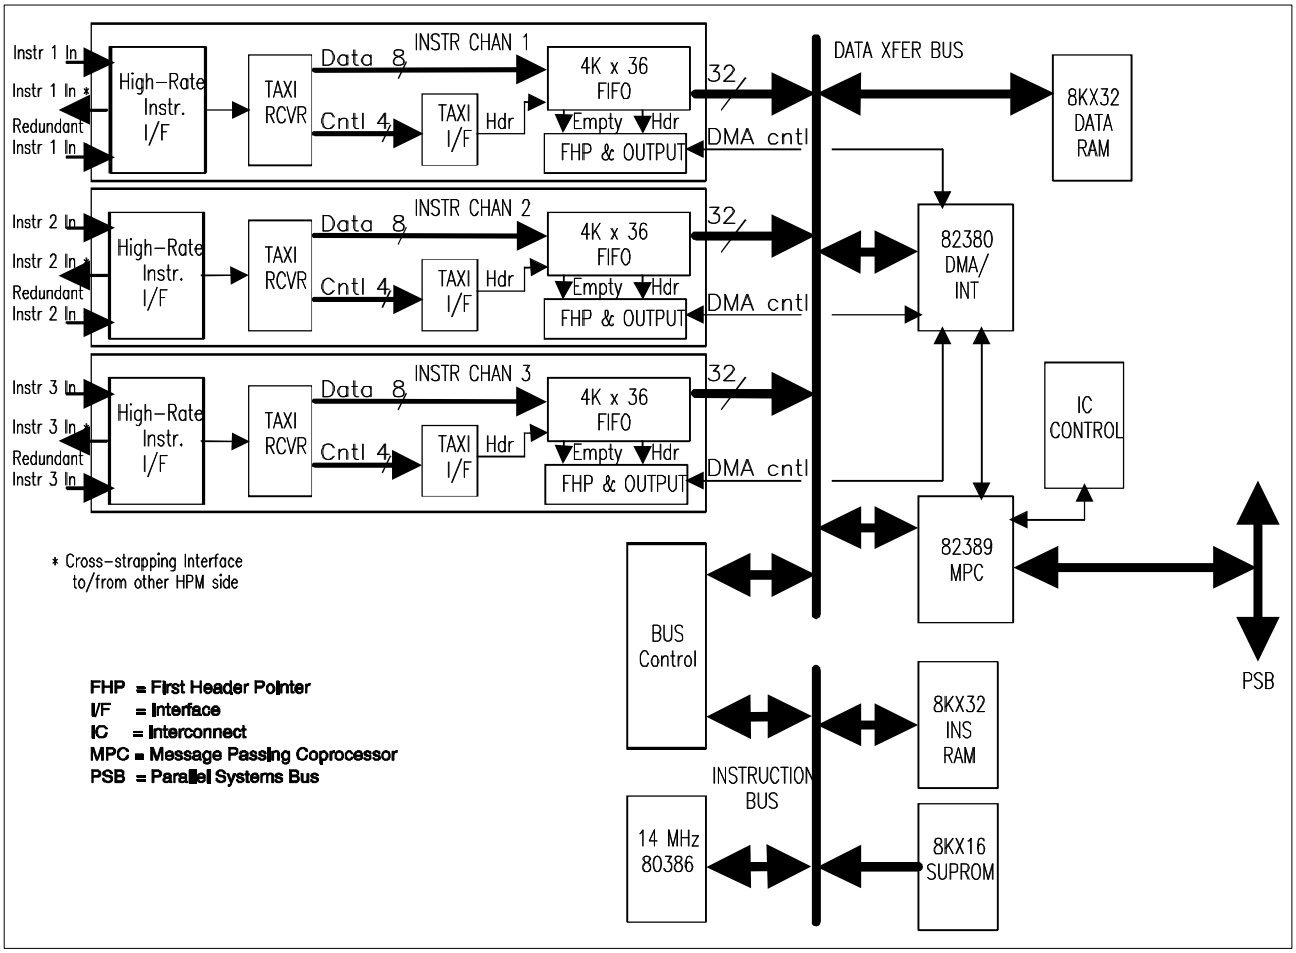

Figure 2: Block Diagram of the HPM

The HPM consists of 3 sections: a high-rate instrument interface section that services 3 instruments, a Direct Memory Access (DMA) based data transfer section, and the

microprocessor based control section. As can be seen in this figure, the design is based around the Intel 80386 microprocessor and related family parts.

OVERVIEW OF THE MODEL

Early in the design process it was determined that 2 areas needing significant attention were the utilization of the PSB and the utilization of the DATA XFER BUS (DXB) on the HPM card (see Figure 2). To minimize hardware costs it was desirable to limit the speed of the hardware to 14 MHz. For similar reasons, it was decided to utilize the PSB in a relatively simplistic (i.e., bandwidth inefficient) fashion by only allowing a single solicited message transfer at a time. To verify that the SFE would work properly with these constraints it was necessary to construct a detailed model with focus on these 2 areas.

To assist with this effort, the NETWORK II.5 discrete event simulation tool from CACI Products Company was utilized. This tool allows one to build models based around hardware elements (Processing Elements, Transfer Devices, and Storage Devices) and software elements (Instructions, Modules, Files, Messages, and Semaphores) that are event driven. A block diagram of the hardware configuration of the constructed model is shown below in Figure 3. The hardware model consists of 20 Processing Elements, 2 Storage Devices, and 5 Transfer Devices. A total of 55 software Modules are distributed amongst the 20 Processing Elements to define the functionality of the system.

Since an area of focus in the model is the HPM card, the model is more detailed for this card. For example, in Figure 3 the functions of each of the HPM's are actually represented by 6 Processing Elements (HPM #1, PE 1, PE 3, I1, I2, and I3), one explicit Transfer Device (BUS1), and a Storage Device (RAM1). In this model IN1 - IN6 represent inputs from the 6 high-rate instruments. The condition modeled is where the instruments always generated maximum length packets which is anticipated to be the worst case.

One of the key purposes of the simulation was to measure the loading on the PSB and DXB busses of the HPM cards (BUS1 and BUS2 in Figure 3) as a function of instrument data rate. NETWORK II.5 provides a number of tools, both graphical and tabular, to present simulation results.

In Figure 4 below the PSB loading is shown for a 20 msec run with all the instruments setup to generate data at their maximum allocated rates (an aggregate of about 109.4 Mbps). At these rates the loading on the PSB averages 67%.

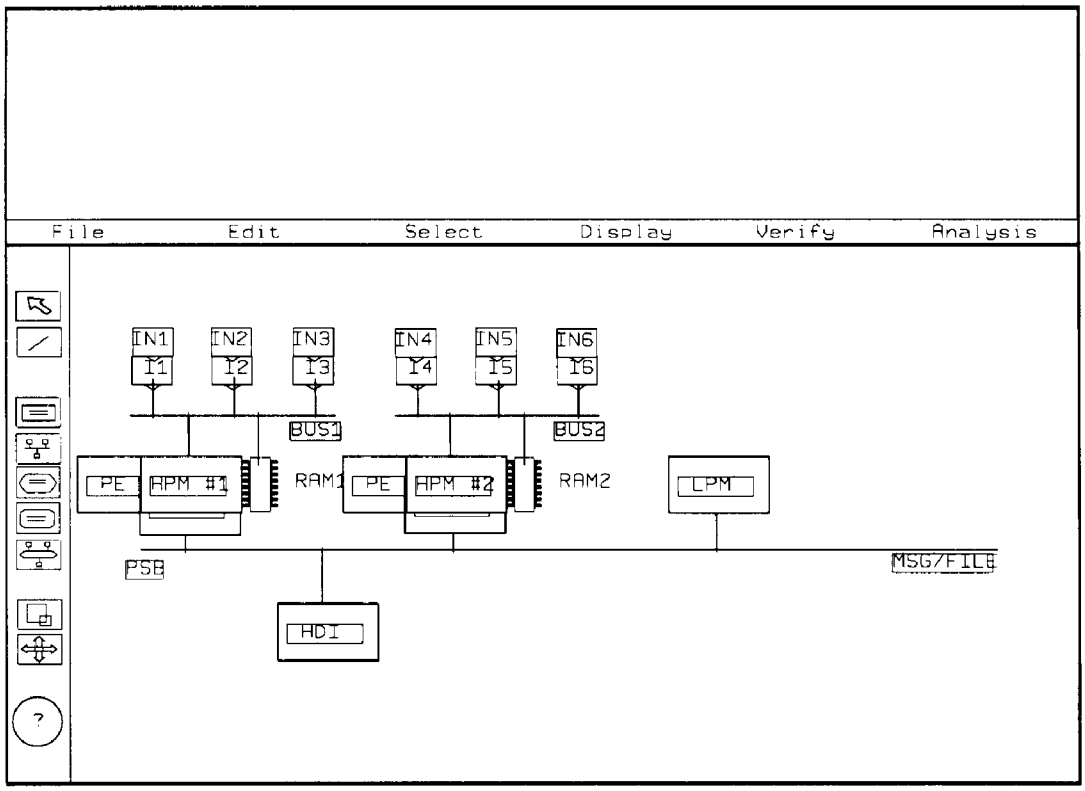

Figure 3: Simulation Hardware Block Diagram (NETWORK II.5 generated)

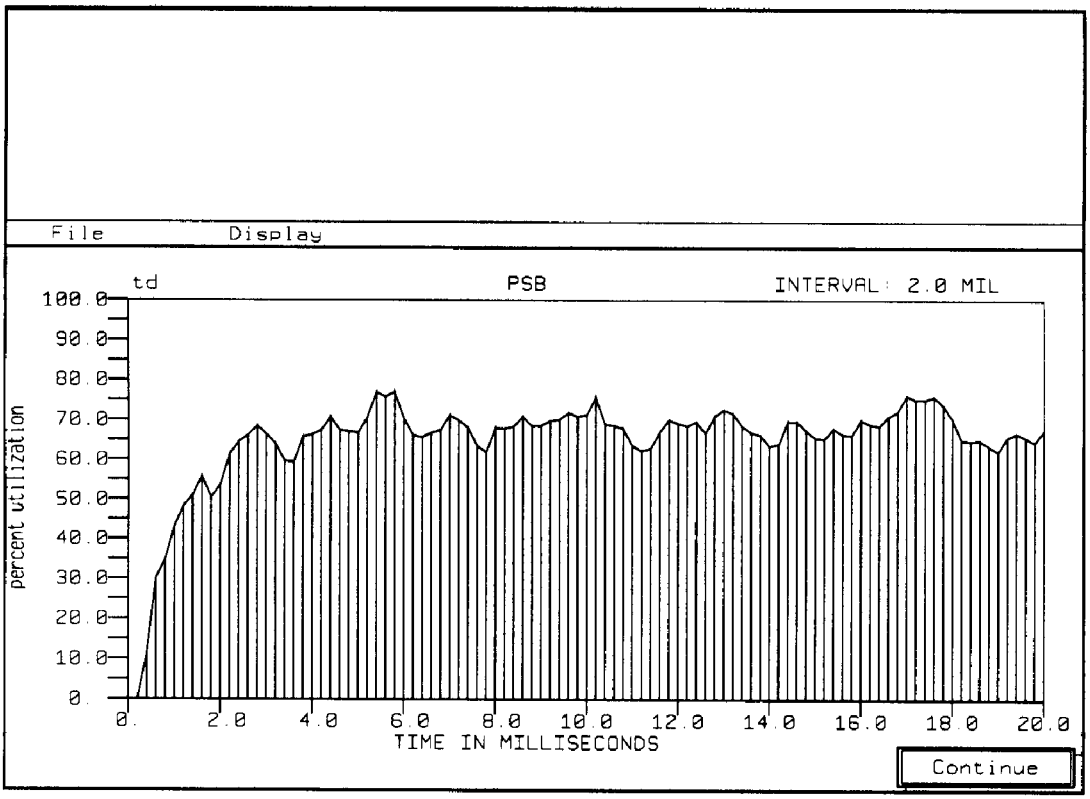

Figure 4: Loading on the PSB at the Maximum Allocated Instrument Data Rates (NETWORK II.5 generated)

In Figure 5 below a similar diagram is presented that shows the HPM1 DXB loading (the worst case of the 2 HPMs) for a 20 msec run with the 3 attached instruments setup to generate data at their maximum allocated rates (an aggregate of about 58.2 Mbps). At these rates the loading on this bus averages 83%.

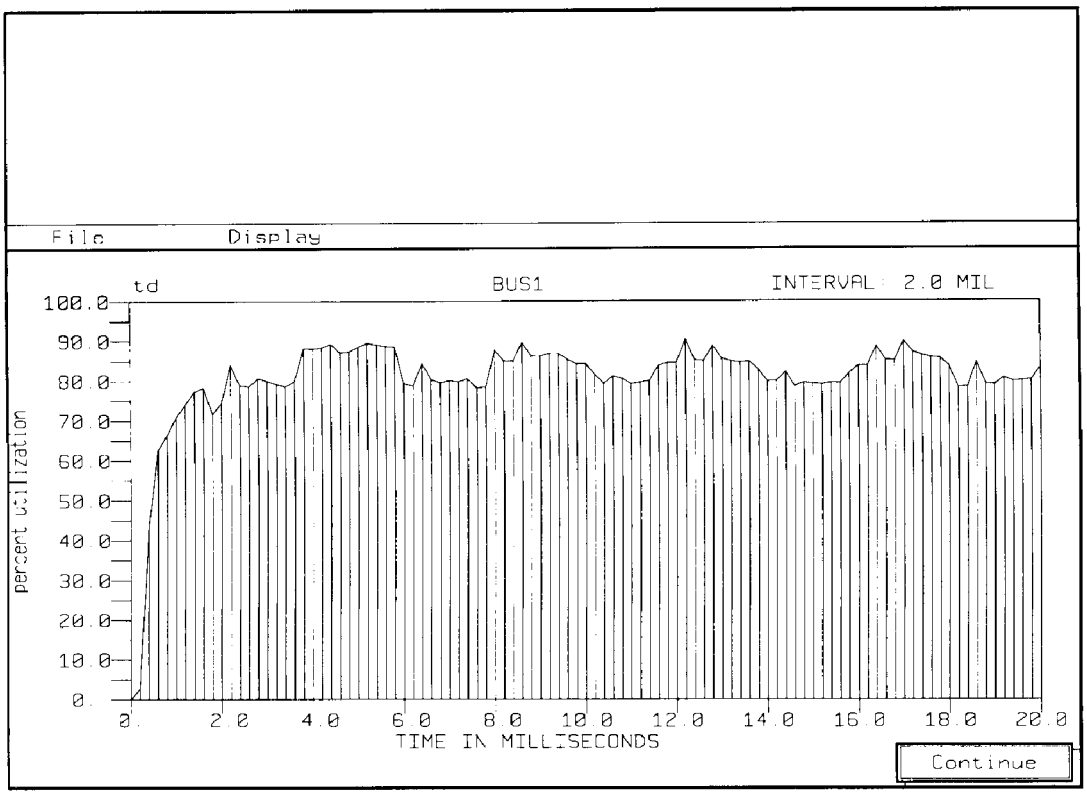

Figure 5: Loading on the HPM1 Local Bus at the Maximum Allocated Instrument Data Rates (NETWORK II.5 generated)

SIMULATION RESULTS

To determine the PSB loading curve (Figure 6 below), several simulation runs were made with the instrument data rates set to various values. Bus loading data was not collected for the first 10 msec of the run to allow the startup transient to settle out. The fourth run in the graph (labeled MAX) represents the loading at the allocated maximum aggregate instrument data rate of 109.4 Mbps. The other runs in this series scaled MAX upward or downward accordingly. The first value in the graph is 20% of MAX which is the approximate average aggregate instrument rate for EOS-AM1. At the average aggregate rate the PSB loading is about 16%.

As is shown in Figure 6, within the normal operating region, the PSB loading will be linearly related to the offered aggregate instrument data rate. This linear relationship breaks down at an offered instrument data rate greater than about 138 Mbps. At this point the PSB (as it is being simplistically utilized) can no longer transfer the data fast

enough to keep up with the inputs into the HPM cards. The symptom of this breakdown is that the RAM1 and RAM2 Storage Devices on the HPM cards begin to overflow. At the operating point immediately preceding overflow, the maximum utilization of the RAM storage that was observed was 79,616 bits - well less than the provided 266,144 bits. Likewise, this run exhibited a maximum FIFO utilization of 194 32-bit words - again well less than the provided 4,096 words. This confirms that the limiting element is the backplane and that all data queues are properly sized.

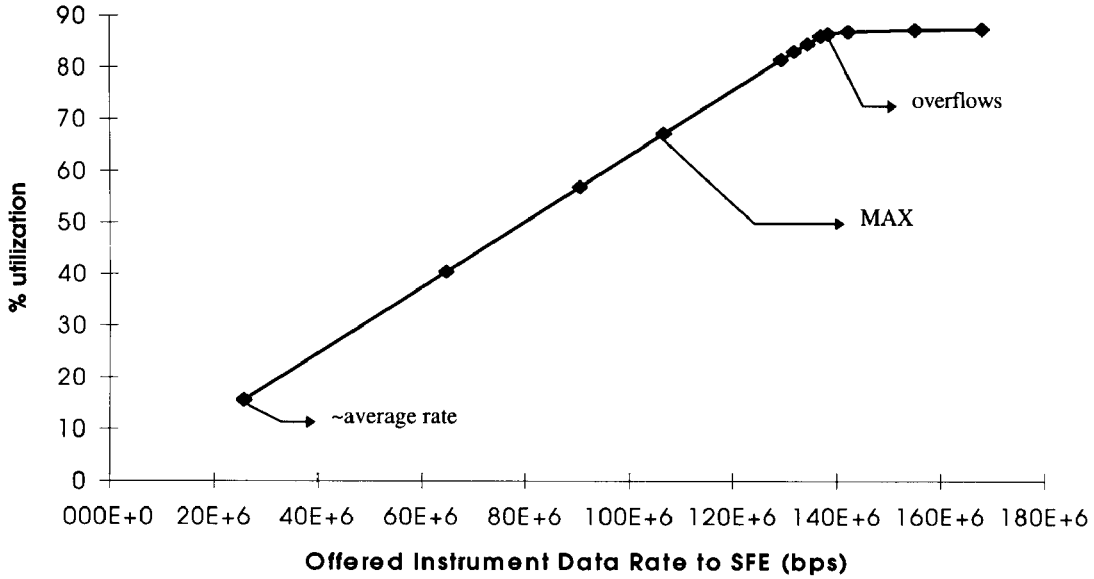

Figure 6: Loading on the PSB as a Function of Aggregate Instrument Data Rate

Similarly, to determine the HPM's DXB loading curve, several simulation runs were made with the attached instrument data rates set to various values. HPM1 was analyzed since it is expected to handle the most data. To decouple possible PSB interactions the HPM2 instruments were turned off (i.e., their data rates set to 0).

The HPM1 DXB loading results are shown in Figure 7 below. The fourth run in this graph (labeled MAXH) represents the loading at the HPM1 allocated maximum data rate of 58.2 Mbps (for the attached 3 instruments). The other runs in this series scaled MAXH upward or downward accordingly. The first value in the graph is 20% of MAXH which is the approximate average aggregate instrument rate expected for HPM1. At the average aggregate rate the DXB loading is about 21%.

The HPM DXB circuitry is designed to operate more efficiently at higher data rates. Therefore, unlike the PSB, a linear relation between data rate and utilization is not expected. This behavior is verified in the graph below. It was found that the DXB could support an offered instrument load up to about 84 Mbps. At this point the DXB can no longer transfer the data fast enough to keep up with the inputs into the HPM

card. Within the model, the symptom of this breakdown is that the FIFO semaphores on the HPM card begin to overflow. At the run immediately preceding overflow, the maximum utilization of the RAM storage that was observed was 50,718 bits - well less than the provided 266,144 bits. Likewise, this run exhibited a maximum FIFO utilization of 447 32-bit words - again well less than the provided 4,096 words.

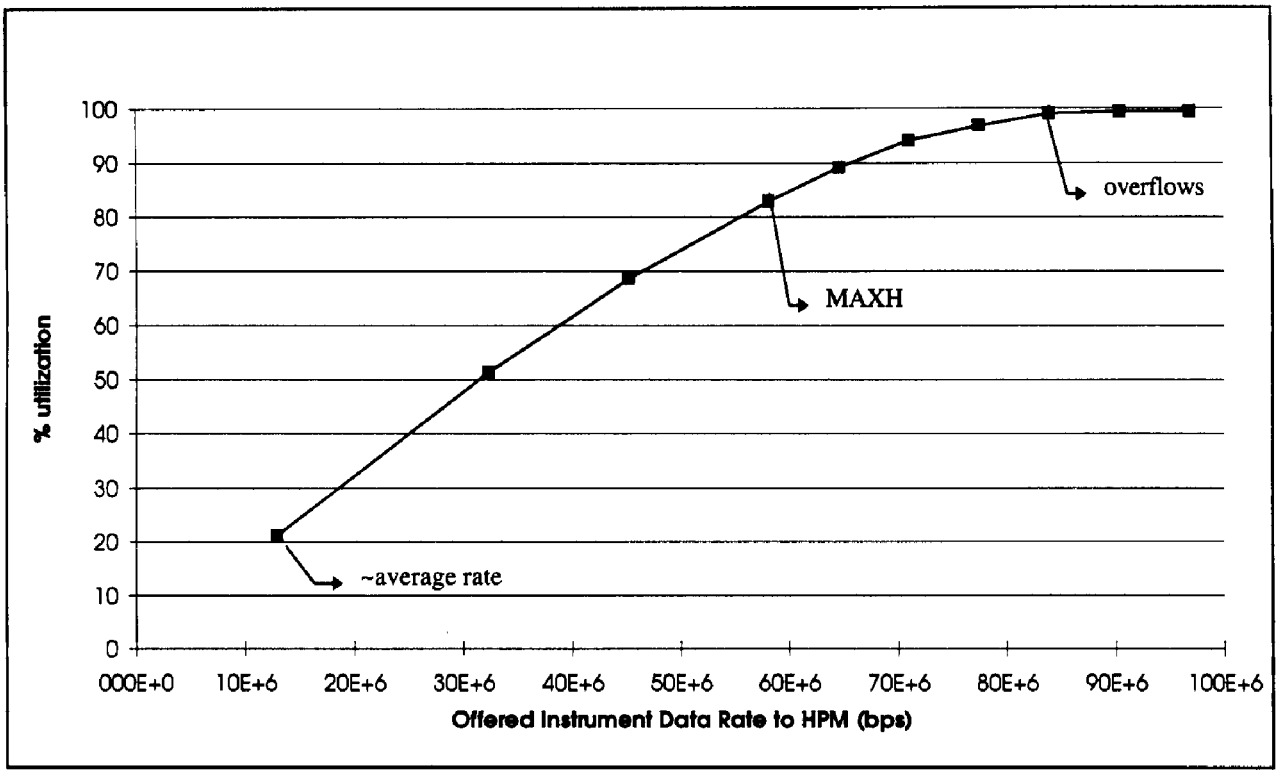

Figure 7: Loading on the HPM Local Bus as a Function of Instrument Data Rate

CONCLUSIONS

Using the simulation results discussed above (some of which were validated with a partial commercial prototype implementation), Fairchild was able to simplify the SFE design by operating at a lower frequency and by utilizing the interconnect bus in a simplified fashion. The results of the simulation predictions as compared to the design requirements are specified in the table below.

REQUIREMENT	AVERAGE (Mbps)	PEAK (Mbps)	PERFORMANCE (Mbps)
PSB Throughput	25.9	109.4	138.4
DXB Throughput	12.9	58.2	84.0

As can be seen by this table, plenty of margin exists in the design.