DATA DISTRIBUTION
In The
TELEMETRY GROUND STATION OF THE 1990's

O.J. (Jud) Strock
Principal Applications Engineer
Fairchild Weston Systems, Inc.
Sarasota, Florida

ABSTRACT

For as long as telemetry has been used in scientific research, users have asked for transmission of more data points at higher data frequencies. Now, the increased complexity of vehicles under test and the presence of data from one or several computer systems on a vehicle has further increased the rate and format complexity of a typical telemetry data stream.

To accommodate higher data rates and increased complexity, many telemetry ground stations use distributed processing techniques, typically employing a hardware preprocessor, a host computer, and one or more intelligent display stations. While this distribution of power potentially enhances data throughput rates, it imposes new demands on data distribution networks within the processing area, and the full potential of the processors cannot be met until these demands have been met.

This paper looks first at system architecture of the typical ground station, and how this architecture and telemetry data rate capability have progressed during the past 20 years. Then it looks at the data distribution requirement in a modern telemetry ground station, explores possible solutions to improve throughput rates, and describes a set of solutions for typical system applications in the next few years.

EVOLUTION OF GROUND STATION ARCHITECTURE

Since the first dedicated real-time computer was used in a telemetry station in 1969, telemetry system users have seen a constant beneficial evolution in ground station architecture.

In 1969, the typical ground station architecture was very simple. As Figure 1 shows, the general-purpose host computer received all telemetry data from the front-end

synchronizers, provided 100% of the processing power, and sent data to storage and to a limited number of simple display devices. The maximum data input rate was very low, because of computer speed limitations.

By 1979, many stations incorporated hardware compressor/preprocessors as shown in Figure 2; this freed the host computers for more display servicing, and more display devices could be driven. As the graph illustrates, possibly 40% of the processing power resided in the compressor/preprocessor, and a system was able to handle perhaps a ten-fold increase in input data rate.

Present-day systems have added a new element, as shown in Figure 3. Not only is the compressor/preprocessor more powerful, but also the displays have become microcomputer-powered workstations. With distributed power in the system, the host computer is now doing perhaps 30% or less of the total data processing. Input rate capability has increased rapidly as compression and preprocessing algorithms reduce data volume and processing requirements at the host.

It will be apparent from a close look at these graphics, as we knew already, that system computers are not getting smaller as we distribute system power; in fact, our demands on the typical system are increasing daily and it is necessary to use larger and larger computers even as we off-load as much of the power as possible.

Up to this point, the system designer has not been required to take any unusual steps to improve the data distribution buses and networks in the system. Rates can be met in most cases with standard techniques, and with readily-available hardware.

ARCHITECTURE OF THE NEAR FUTURE

We know from observations and can see from Figure 4 that the evolution outlined above continues. We will demand ever-increasing data input rates, and will continue to expect more and better displays of data in real time. Processing power will be moved from the host computer in both directions, forcing the system designer to put more power into the hardware preprocessor and to use higher-powered workstations for data processing, formatting, and display.

The preprocessor will become an even more prominent element in system architecture in the future than it is now as data rates increase, the number of data streams in a system increases, and formats become even more complex. A present-day result of the trend in preprocessors is the EMR 8715 Preprocessor, a modular device which can incorporate as many as 16 input data streams plus up to three independent time sources, and output preprocessed data via as many as eight independent ports. It can accept input data at an

aggregate rate as high as 2.4 megawords (32 bits) per second, examine every word, and output all data and identification tags, or it can output compressed data if programmed to do so.

The highest-performance front-end synchronizers of the future will still be stand-alone units, each with versatility, and each containing a microcomputer for setup, format storage, and self-diagnostics. However, lower-performance modular front-end synchronizers can be incorporated into a preprocessor as shown where space cannot accommodate the stand-alone units.

Similarly, even though the high-performance preprocessor incorporates a microcomputer module for housekeeping, it normally works with a stand-alone general-purpose host computer. This has many advantages in terms of flexibility and maintainability; even when the extra power of the preprocessor is available, most users will benefit from the versatility of a stand-alone computer. By design, the powerful preprocessor is intimately related to the host in both hardware and software functionality, even though they are generally different physical entities. If space is not available, however, some system architectures of the future may incorporate a general-purpose host processor as a module within the preprocessor box. In these cases, it is imperative that software of the host module be general-purpose, applications-oriented, and capable of independence from the preprocessor for diagnostics and for general-purpose processing, and also that the processor module be accessible through bus architecture for expansion by the designer or user.

Processing power in the display workstation will continue to increase, with more and better choices of display methods and formats. The system of the near future will distribute a current value data table to all workstations and maintain the freshness of that table, so that each analyst can select, format, process, and display data as desired. The EMR System 90 architecture has this capability already, using Ethernet as the local-area distribution network.

Telemetry users continue to search for an optimum method for high-speed large-volume data storage. Some of the stored data may be recalled quickly for quick-look analysis, and some or all may be archived onto a transfer medium for distribution to other processing and analysis locations. The system of the near future is likely to incorporate a cluster of Winchester-technology disks with the capability for data transfer at rates of more than ten megabytes per second, plus a lower-rate archival transfer device such as an optical disk or high-density computer tape.

On this subject, disks with high transfer rates are available. The problem comes as we fill one cylinder with data, and must cut off data flow until heads are moved to a new cylinder.

The solution of the near future is a cluster of disks and controllers, such that realtime data can be sent to disk A cylinder 1 until it is filled, then can be transferred immediately to disk B cylinder 1. While that cylinder is being written, the heads on disk A are moved to cylinder 2, and it is ready when disk B cylinder 1 has been filled, and so on. In some cases, several disks may be in a cluster to take care of extremely high rates. On data playback, the software must be able to cycle through the cluster in the same sequence, of course.

Data for transfer to other systems and analysts can be reformatted off-line to an industry-standard 9-track tape. Hopefully, optical disk formatting standards will develop soon, and we will be able to use this low-cost medium as an archival and transfer device.

DATA DISTRIBUTION TECHNIQUES FOR THE NEAR FUTURE

Referring again to Figure 4 for consideration of data distribution, it can be recognized that data volume expands at the input to the preprocessor unit. Here, possibly more than one data stream is input onto a common bus. Each data word from each stream is tagged to identify the measurement point represented. Time is merged with data and tagged. So the first data distribution challenge to the designer occurs at the merge point. In the EMR 8715 used already as an example, the internal data path is via an extremely high speed bus, a proprietary bus called the “priority/command/data bus”, or PCD bus, with a rate of 10 million transfers per second. Special bus drive logic and architecture were chosen here to support the preprocessor’s capability for data rates to 2.4 megawords (32 bits) in and 2.4 megawords (32 bits) out per second, with ample growth capability.

Key to the speed of this unique bus, a significant step in improvement of telemetry data distribution rates, is the arbitration scheme. The PCD bus is synchronous, operating on a 100 nanosecond clock. During each clock period, while up to 32 bits of data and 32 bits of command are being transferred from a driver module to a receiver module, arbitration is being performed on the 16 priority arbitration lines to decide which driver module is to be the bus master during the next clock period. Should a given driver module require control of the bus for more than one clock period in order to make multi-word transfers, it sets a “bus busy” line and retains that control. Unlike many so-called hi-speed buses which in fact require several clock cycles for arbitration, handshake, and information transfer, this bus can transfer information from random sources to random destinations at a sustained rate of 10 million 64-bit transfers per second. This is a significant step forward in data distribution within a high-speed component in the modern telemetry system, the preprocessor.

High-speed data transfer from the host computer to disks presents a continuing problem to us in the world of telemetry. A superminicomputer often uses a proprietary high-speed bus

system to service disks, and we anticipate that in the near future even higher-speed buses, possibly fiber optic, will be offered. Still the throughput from input port through memory onto a disk is unacceptable with present-day telemetry rates, and the problem will get worse as rates increase. Considering the fact that data for disk storage does not get processed in the computer, it seems logical that system designers in the near future will use ways to bypass the host computer entirely with disk data, possibly using some popular distribution bus in order to give options on which type of storage device to install on a given system. One such distribution bus is Motorola's VME bus, and another is the "Small Computer System Interface" (SCSI) bus, each of which has attracted a large number of peripheral storage device manufacturers. The VME bus has a theoretical transfer rate of 10 megawords (32 bits) per second, the SCSI-1 is rated at 5 megabytes per second, and the upcoming SCSI-2 will be rated at 20 megabytes per second.

Each of the one to eight output ports of the EMR 8715 has the ability to handle up to two megawords (32 bits each) per second from the PCD bus. Where a system has several data destinations, the preprocessor will distribute as programmed in order to service the desired devices and hold the transfer rate for each device to an acceptable level.

Often the input port to a host computer is a telemetry product. Computer manufacturers do not need high-speed realtime input ports for many applications, and have hesitated to develop them. Telemetry system designers, then, must provide such ports in their systems. Our company, for example, has developed unique interface devices for Digital's Unibus, Q-bus, and BI bus, and for Concurrent's and Gould's high-speed memory buses. Typically a device operates in an externally-specified mode (where a tag from the preprocessor with each word tells the port where the word is to be placed in memory) or in a ping-pong buffer-building mode (where buffers are built in computer memory to be transferred to disk or tape). In either mode, the unique interface device does memory addressing and thus provides high-speed data transfer. Telemetry system architecture of the near future will see continuing development and improvement in these unique telemetry-related devices for the more popular realtime computers.

One thing which makes the input port simpler than it was a few years ago is the evolving change in function. Earlier input ports were required to merge time with data, provide a setup path for front-end equipment, and control the high-speed data input. Now, setup is generally handled by a separate port and time is often merged at the preprocessor input rather than at the computer input. Further, preprocessed data is usually tagged, thus the start-stop capability of an input port is not as critical as with data which is not tagged. For these reasons we would expect a new generation of simpler, faster input ports within the next few years, with sustained rates up to 6 or more megabytes per second.

Data for display, assembled into a current-value table in the host computer and refreshed several times each second, must be distributed to display workstations over a high-speed bus to support the needs of data analysts at those stations. With present current-value table sizes and workstation processor capabilities, the Ethernet Local Area Network (LAN) is marginally suitable. We expect it to become inadequate for many large systems within a few years as table sizes increase and workstations are upgraded to receive and process larger data packages. This necessitates the search for a different LAN having data rates to ten times that of Ethernet. One possible candidate for such display data distribution is the Fiber Distributed Data Interface (FDDI), a fiber optic LAN operating at 100 megabits per second (a ten-fold increase over Ethernet). Up to 500 stations can be put on the FDDI network. A controller for this LAN will generally include a microprocessor, random-access memory, programmable read-only memory, first-in first-out buffers, built-in network test capability, and other operating conveniences which make it valuable to the telemetry system designer and user.

SUMMARY

The telemetry system designer and user have reached the point where system throughput rates are not simply a function of the hardware and software, but may be restricted by data distribution techniques. This necessitates improvements in several areas of system architecture, as defined in this paper. We see such improvements as feasible and imminent; in fact, some are under way already, and others will certainly be developed within the next few years.

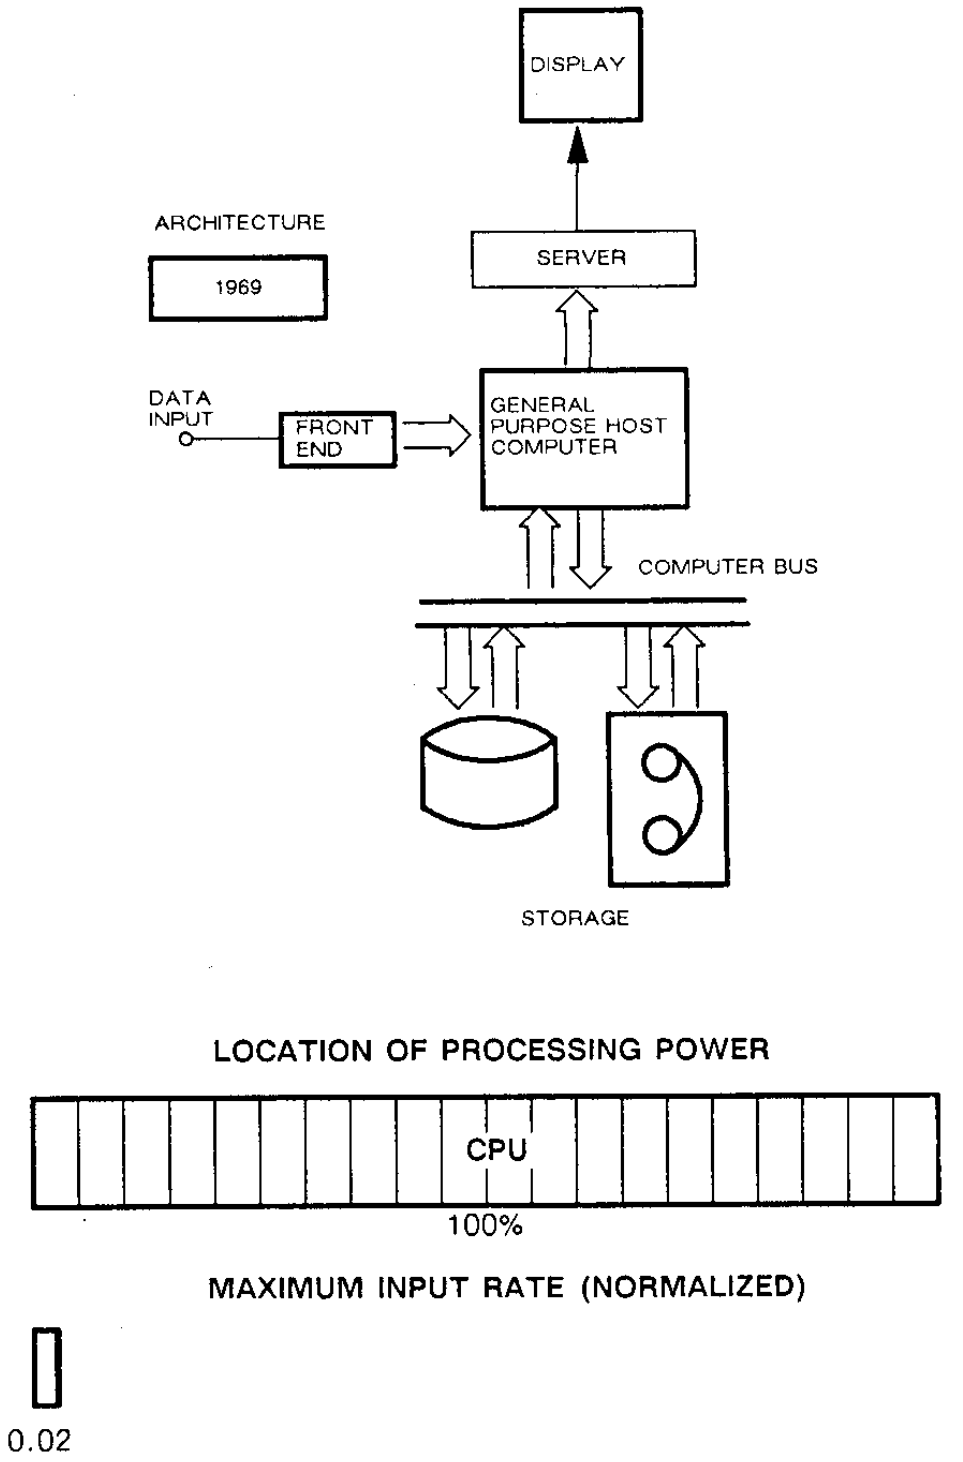

Figure 1. Typical Architecture, 1969

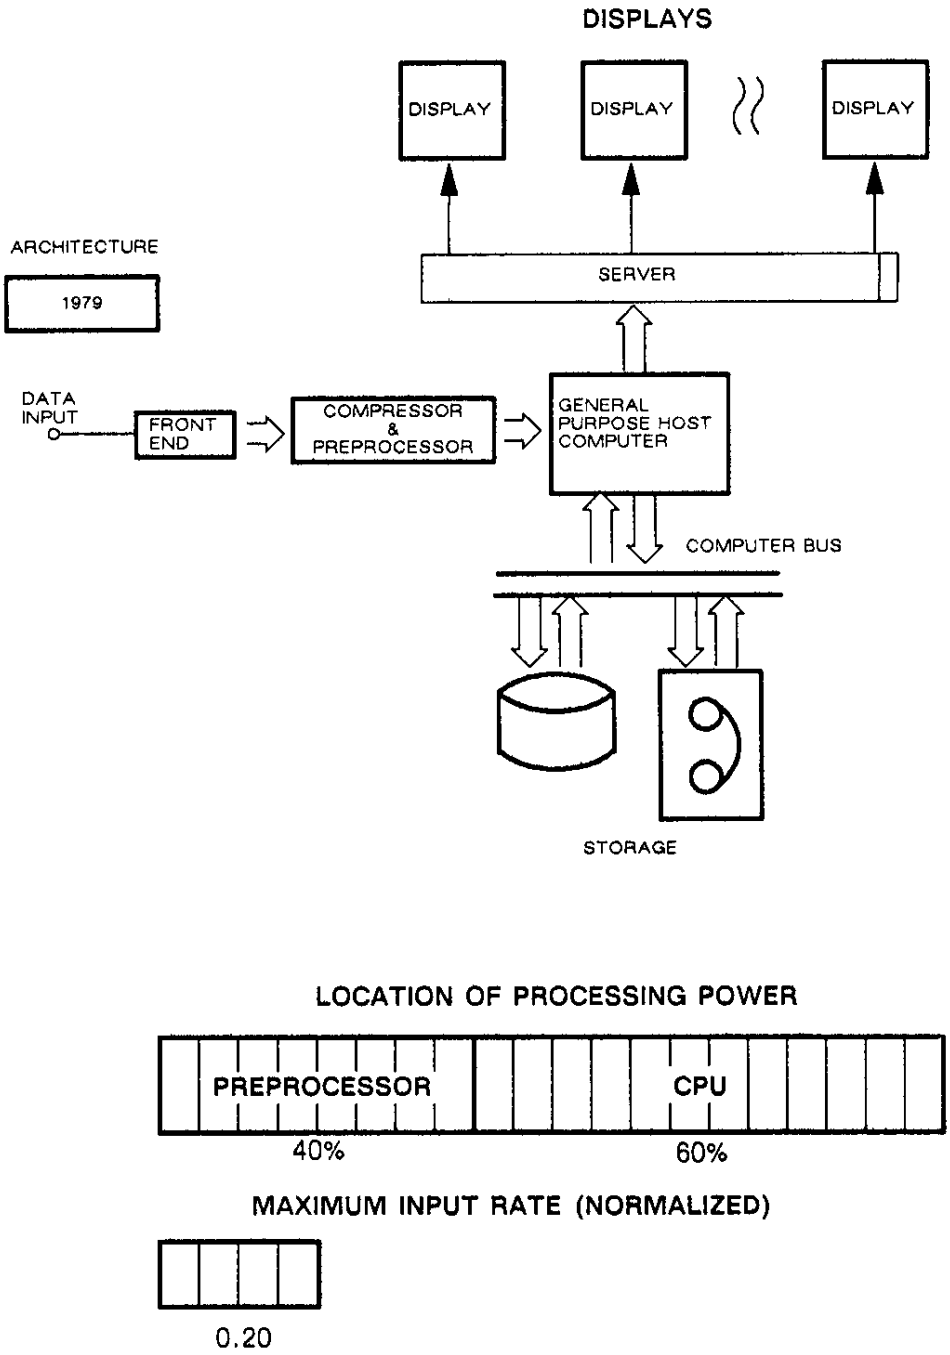

Figure 2. Typical Architecture, 1979

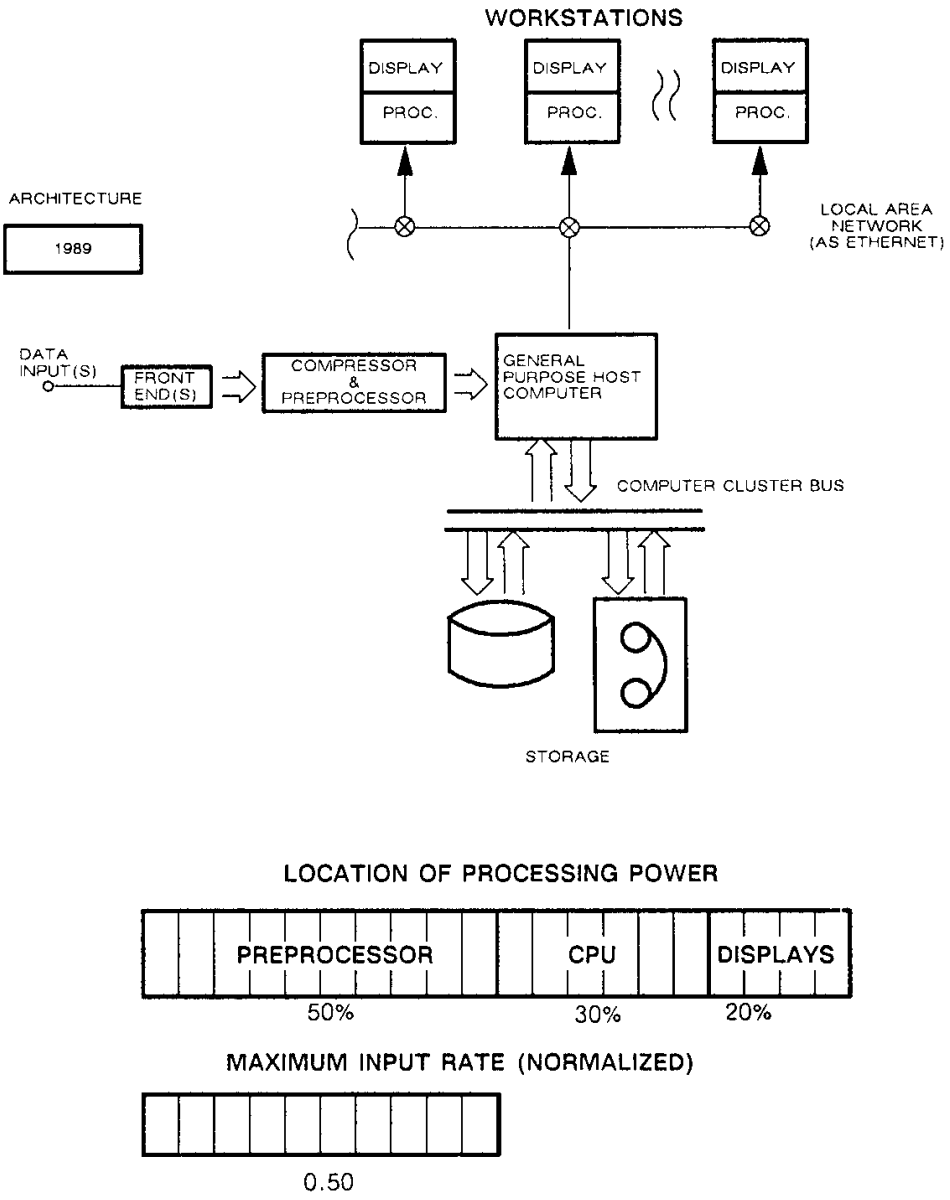

Figure 3. Typical Architecture, 1989

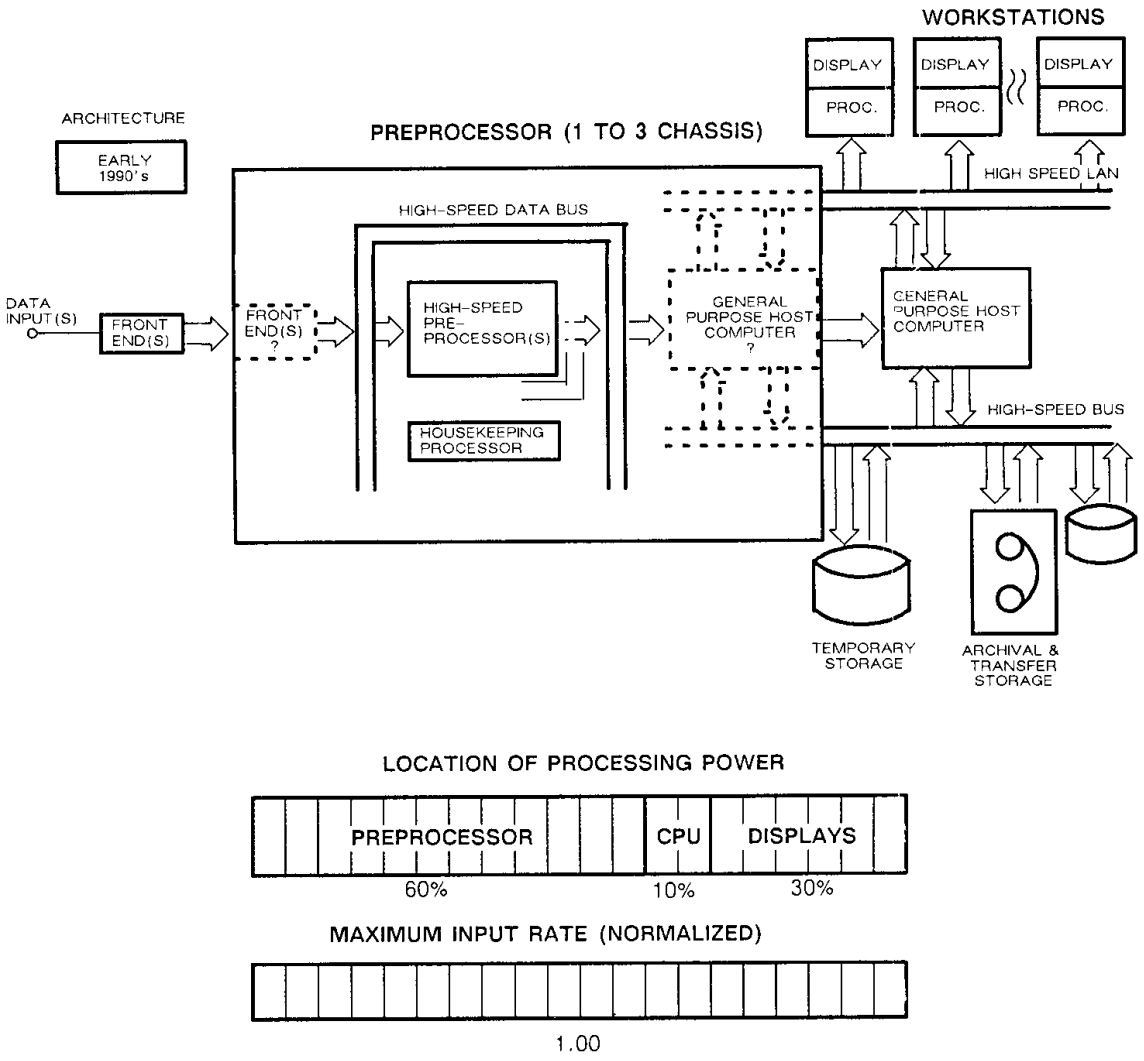

Figure 4. Typical Architecture, Near Future