Using 200 Megabyte HIPPI Connection for Animation

Nancy Rowe, Applications Division, Visualization, Cray Research, A Silicon Graphics Company, Eagan, Minnesota USA

ABSTRACT: *High Performance Parallel Interface (HIPPI) is a high speed channel specification. Traditionally, 32-bit words are used to transfer data across a HIPPI connection. Cray Research has the unique ability among major computer vendors to transfer data across HIPPI connections using 64-bit words. By coupling a Cray Research supercomputer to a PsiTech frame buffer with a HIPPI connection and using 64-bit words to transfer data across the HIPPI connection we can display animations at film speed of images with a resolution of 2K x 1K.*

Introduction

Traditionally the only way to view film resolution animation was to record frames and play back the recording. This limitation was due to several factors. Computer monitors were not large enough to display film resolution animation. Computer monitors could not be refreshed fast enough to display film resolution animation. At a resolution of 2048 x 1024 each frame of film was more than 6 million bytes of data. Because the data size was so large, film data could not be read from disk and transferred to a display device at film speed. A large capacity storage device was required to store the film data. This amount of disk space needed was beyond the capacity of the average workstation disk.

Goal

Our goal was to provide a simple solution to display high resolution animation from disk. We stress the word simple. There are other alternatives which could deliver similar performance. Systems could be configured to perform at higher speed. Synchronization software could be used to synchronize multiple network channels. But none of these solutions offered the plug and play solution we were looking for.

Requirements

Although there are various definitions of film resolution, for the context of this paper we define film resolution as 2048 x 1107 pixels per frame of animation. The color information of each pixel is our data. For our data each pixel composed of 3 color components, red, green, and blue. We store each component as 8 bits or 1 byte of data. So each pixel is 24 bits or 3 Bytes of data. At film resolution this would require approximately 54 Mbits or 6.8 MBytes of data per frame (see table 1).

Table 1. Frame resolution requirements for film

The following calculations show the data storage requirements and the data transfer requirements for images of 2048 x 1107 pixels = 274536 pixels. Calculations are done in units of bits as well as units of Bytes

Pixel composition
3 components / pixel - red, green, blue (rgb)
3 components / pixel x 8 bits / component
= 24 bits / pixel = 3 Bytes / pixel

Frame data size
274536 pixels x 24 bits / pixel
= 54, 411, 264 bits / frame
274536 pixels x 3 Bytes / pixel
= 6,801,408 Bytes / frame

Data transfer rate required
24 frames / second x 54, 411, 264 bits / frame
= 1.305870 x 10^9 *bits / second*
24 frames / second x 6,801,408 Bytes / frame
= 1.632337 x 10^8
≈ 165 MBytes / second

Disk space requirement
1791 frames x 6,291,456 bytes / frame
= 11.26799 x 10^{10} *Bytes*
≈ 11.27 GigaBytes
approximately 74 seconds of animation

Film speed is 24 frames per second. In order to display a film resolution animation at film speed we need to transfer 24

6.8 MB frames per second to the display device. This means we need to transfer just under 165 MB of data per second to the display device.

Each frame of the animation is precalculated, so we need to store the data prior to transferring the image to the frame buffer. One second of film is 24 frames. At 6.8 MB per frame, one second of film requires just under 165 MB of storage capacity. One minute of film requires 9.8 GB of disk storage.

In order to view a film resolution animation we need a monitor capable of displaying images of 2048 x 1107 pixels. Typical monitor size for a workstation monitor is 1280 x 1024. Most workstation monitors have a display resolution of under 1600 x 1280.

By connecting a CRAY T90 with a Psitech frame buffer and a CRAY ND-40 RAID5 disk array and using a high resolution Sony monitor we are able to display high resolution animations at film speed.

Data

The frames were created in ppm format. Although the aim was to use film resolution images (2048 x 1107) for the initial tests we used a slightly lower resolution frame size of 2048 x 1024. This frame size required 50Mb or 6.3 MB of data per frame. Our animation consisted of 1791 frames. This was just over 74 seconds of animation. First we removed the header information from each ppm file. Then we concatenated the frames together to create a single file. This file was approximately 11.27 GB. The file was then stored on the CRAY ND-40 disk array.

HIPPI

HIPPI is High-Performance Parallel Interface. HIPPI is an ANSI standard developed by the ANSI X3T9.3 committee. HIPPI is a point to point interface that can transmit data between devices at either 100 MB/s or 200 MB/s. Using 2 cables for bi-directional transfers, HIPPI can transfer 32 bits of data in parallel for a theoretical maximum transfer rate of 98.8 MB/s (rounded to 100 MB/s). Using 4 cables, HIPPI can transfer 64 bits of data in parallel for a theoretical maximum transfer rate of just under 200 MB/s. HIPPI uses 50 shielded twisted pair of copper cable to transmit data up to 25 meters. HIPPI switches can be cascaded to extend a HIPPI network to hundreds of meters. Serial HIPPI which defines a fiber-optic connection can be used to extend a HIPPI network 10 kilometers. Although, HIPPI was designed as a “machine room” network, HIPPI networks can be extended over long distances by working in conjunction with ATM or Sonet.

We chose HIPPI as a networking interface for several reasons. HIPPI is a networking standard. The HIPPI standard specifies a transfer rate of 200 MB/s. This met our transfer rate requirement of 165 MB/s. The HIPPI standard has been in existence for nearly a decade and HIPPI is supported by a large number of vendors, so all the equipment we needed was readily available. HIPPI disks, HIPPI switches, computers with HIPPI interfaces, and HIPPI frame buffers have been on the market for a number of years and have become reliable products.

Among major computer manufacturers, Cray Research is the only one to offer a 200 MB/s HIPPI channel option. But, even Cray Research does not currently support 200 MB/s HIPPI on all configurations. The CRAY J90 series currently only supports the 32 bit 100 MB/s HIPPI interface. The CRAY Y-MP, C90 and T90 all support 200 MB/s HIPPI. With the availability of CRAY’s Gigaring the CRAY J90SE and Cray T3E will also be able to approach 200 MB/s HIPPI speed.

Equipment Configuration

For our initial tests we used an NSC PS32 HIPPI switch to connect a CRAY J916, a PsiTech HFB360A frame buffer, and a CRAY ND-40 RAID5 disk array. We used a Sony DDM2802C monitor with a display resolution of 2k x 2k as a display device. The CRAY J916 is only able to run 32 bit HIPPI, so as work progressed we replaced the CRAY J916 system with a CRAY T94. On the Cray Research systems we ran the UNICOS operating system version 9.0. PsiTech provided a software library that ran on the Cray Research systems and issued commands to the PsiTech frame buffer.

The CRAY J916 is a vector architecture machine. The CRAY J916 has a clock speed of 10 nanoseconds. The CRAY J916 was configured with 16 processors and 256 MW of memory.

The CRAY T94 is a vector architecture machine. The T94 has a clock speed of 2.2 nanoseconds. The CRAY T94 was configured with 4 processors and 128 MW of memory.

The PsiTech HFB360A frame buffer is a high-speed frame buffer. The HFB360A uses a HIPPI interface to connect to the host computer, the Cray Research system. The HFB360A can be configured to accept data transfers of 200MB/s. The HFB360A is capable of displaying images of 2k x 2k x 24 bits deep. We connected the HFB360A to a Sony DDM2802C monitor. The Sony DDM2802C monitor has a display area of 2K x 2K.

The CRAY ND-40 Network Disk Array is a single chassis RAID5 storage device. The CRAY ND-40 has 4 slots for HIPPI connection. The ND-40 can sustain data transfer rates of up to 160 MB/s. The ND-40 can store up to 376 GB of data in a RAID5 configuration.

Initial Work

Although the ultimate goal was to read film resolution images, 2048 x 1107, off disk and display them to a monitor at film speed, 24 frames per second, we began testing simpler scenarios. We built up step by step to our ultimate goal. We started with a smaller frame size and with a system with a theoretical maximum transfer rate of only 100 MB/s. First we worked on a CRAY J916 configured with a single wide HIPPI channel, 32 bits. The single wide HIPPI channel was limited to 100 MB/s theoretical maximum speed. We used a frame size of 2048 x 1024. We began by reading 2 images into memory and on the CRAY J916 system, transferring the images over the

HIPPI channel through a HIPPI switch to a PsiTech HFB360A frame buffer to a high resolution Sony monitor to determine possible display speed. Using this configuration we achieved transfer rates of 89 MB/s. Results of all tests are displayed in table 2.

Table 2. Test results

Transfer rate from memory of J90

image size	transfer rate	frames/s
1024 x 1024	*85 MB/s*	*27*
2048 x 1024	*89 MB/s*	*14.1*

Transfer rate from ND-40 connected to J90

image size	transfer rate	frames/s
1024 x 1024	*59 MB/s*	*18.8*
2048 x 1024	*66 MB/s*	*10.5*

Transfer rate from memory of T90 (single wide, 32 bit HIPPI)

image size	transfer rate	frames/s
1024 x 1024	*87 MB/s*	*27.7*
2048 x 1024	*90 MB/s*	*14.3*

(double wide, 64 bit HIPPI)

image size	transfer rate	frames/s
1024 x 1024	*139 MB/s*	*44.2*
2048 x 1024	*152 MB/s*	*24.2*

Next we connected an CRAY ND-40 disk array through the HIPPI switch to the CRAY J916 to provide a storage location for the multiple frames of the animation. This added the extra factor of time required to read the information off disk.

In our third test we replaced the CRAY J916 with a CRAY T94. First we configured the T90 with a 32 bit HIPPI channel capable of 100 MB/s theoretical maximum transfer rate. The frames were read into memory and then displayed to the frame buffer. We achieved slightly better transfer rates than for a 100 MB/s HIPPI channel on the Cray J916. We then configured the CRAY T94 with a 64 bit parallel HIPPI connection capable of a theoretical maximum transfer rate of 200 MB/s. First we tested this configuration with the frames already in memory. Then we read the frames from the ND-40 disk array to more closer approximate our ultimate configuration goal. With the frames in memory we achieved transfer rates of 152 MB/s. We were unable to test reading the frames from disk on the T90 because of our T90 system was not configured with 2 200 MB/s HIPPI channels. We need a channel for the frame buffer configured raw and a separate channel for the ND-40 configured to run IPI3 protocol.

What we learned

200 MB/s HIPPI is a viable option for displaying high resolution animation. HIPPI is a relatively mature specification. Configuring 200 MB/s is easy. Software synchronization is not required. HIPPI is supported by many vendors. The HIPPI committee continues to progress the specification.

But there are some limitations. Although the distance limitations of HIPPI can be overcome, doing so can be complicated and expensive. Our solution is best suited to a computer room type setting. For the simplest solution distances should be kept to under 50 meters. We also found that working with HIPPI connectors can be frustrating. The connectors are large and have 100 pins, which can be easily bent or broken. In our environment we needed to run the HIPPI cables through a channel under the floor. Because of the large connector heads, fitting multiple cables into a small space of a channel was problematic. HIPPI has been in existence for a number of years and is supported by many vendors, but, the 200 MB/s HIPPI channel is not generally a system default. Frequently only the 100 MB/s HIPPI option is available from the vendors. Although computer systems and peripheral equipment can be configured with a HIPPI interface, frequently it is not. So in order to use the HIPPI solution an additional board may need to be added to an existing system.

Future plans

While our initial results are positive, we would like to run additional tests. We hope to improve our test results by utilizing asynchronous reads from disk. We would like to reconfigure our system with 2 200 MB/s HIPPI channels. We intend to increase the data size to full film resolution of 2048 x 1107. We intend to substitute striped disks for the ND-40. This should be a more cost-effective solution. We intend to investigate other networking options such as Fibre Channel. As we continue our work we will be monitoring the progress of the HIPPI specification. HIPPI 6400 will offer at least 800 MB/s throughput. With the introduction of Cray's Gigaring, the 200 MB/s HIPPI solution will be available on the CRAY J916. When Gigaring becomes available on the CRAY J916, we will test this configuration. We anticipate results from the J916 with Gigaring, similar to the results we have seen on the CRAY T90. We may also investigate 3rd party data transfers to reduce the transfer time required. Third party data transfers would transfer the data directly from the ND-40 disk to the frame buffer instead of sending the data from the disk through the Cray system to the frame buffer.

Conclusion

At the present time, 200 MB/s HIPPI provides a simple solution to viewing high resolution animation at film speed. By understanding the potential and the limitations of HIPPI, we can create a simple animation playback system from a Cray supercomputer. More work is needed to realize the full potential of the solution. Since 200 MB/s HIPPI is a relatively mature standard, 200 MB/s equipment is available. We expect that our animation system will benefit by the work being done on HIPPI 6400 and Fibre Channel. For a simple 200 MB/s solution to viewing high resolution animation, HIPPI is a good choice.

References:

1. Renwick, John, Cray's Perspective on the High Performance Parallel Interface Standard
2. Tolmie, Don and Flanagan, Don, HIPPI: It's Not Just for Supercomputer Anymore, Data Communications on the Web May 1995 http://www.cic-5.lanl.gov/lanp/HIPPI_Data_Comm.html
3. Tolmie, Don and Tenbrink, Steve, A Short History of Gigabit Research at Los Alamos http://juggler.lanl.gov/lanp/hiphis.html
4. HIPPI in Detail, EANTE 1994