SLAC-PUB-2891
MARCH 1982
(N)

3081 EXPERIENCE AT SLAC *

T.Y. Johnston

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

March 8, 1982

It has been over a year since our 3081 was installed. I would like to cover in this presentation:

- o Installation and the First Months
- o Current Configuration
- o Hardware Transitions
- o Current Status
- o Architectural Improvements
- o Reliability
- o Summary

Installation and the First Months

The 3081 arrived early on the morning of Saturday January 31, 1981. Four hours later it was in the machine room and the installation started. Twenty four hours later the 3081 was up and IBM started connecting channel cables. Monday morning we started bringing up the software. By Monday afternoon we had solved two problems (one ours and one IBM's) and we were ready to start software testing. Tuesday morning we started testing VM and by early afternoon we had cleaned out a few glitches and were ready to try running production. Wednesday morning we brought up VM production on the 3081. We have never even been tempted to move back.

Over the next few months we encountered several microcode problems. These have all been resolved. We have never had a TCM fail.

---

* Work supported by the Department of Energy, contract DE-AC03-76F00515

(Presented at the Annual SHARE Conference 1982, Los Angeles, CA. March 15-20, 1982)

Configuration

Our 3081 is a minimum machine - 16 Megs and 16 channels. Channel 0 is a byte multiplexor and all others are blocked multiplexors. Our use of the channels is:

- o 1 Byte multiplexor - 2 3705s
- o 2 Block multiplexors - tapes
- o 1 Block multiplexor - Solid state drums for paging (3 2305-2 equivalent)
- o 4 Block Multiplexors - 3350s
- o 1 Block Multiplexor - 3330s and operators consoles
- o 4 Block Multiplexors - 3380s
- o 1 Block Multiplexor - 3800, CTC, 1403s, and strange devices.
- o 2 Block Multiplexors - 327x terminals and a Series 1 running the Yale IUP.

We have defined the system as an AP. We don't have enough channels to make the machine a 370 MP. Clearly, when XA support is available we will be a true MP. The 3330s 3350s and some of the tapes are shared with two 168s. We have set up the 3350s with string switching so that the loss of a 3880 will take a 168 out of service but not the 3081.

The 4 channels of 3380s consist of 4 strings on 2 3880s. We did not install dynamic path selection. We now have trouble justifying to ourselves the need for dynamic path selection for the 3380s. There appears to be more than adequate channel and control unit capacity, such that we don't need it for performance and we have never had a 3880-3 failure. We have also never had a 3380 media failure. The 3380s are the primary DASD for the 3081, and they account for over 80 percent of the VM space. The drums and the 3380s provide all paging and spooling space. 3380 are also used for large or high activity user minidisks and highly used system disks such as the Y and S disks.

Hardware Transitions

Since installing the 3081 we have gone through 2 major hardware upgrades.

In May we received our 4 strings of 3380s. Within a week one of our physics groups was using them, and by the end of the month we had migrated CP paging and spooling to the 3380s and had give a series of maxi disks on the 3380s to the physicists. These have allowed the groups to change their computing style from OS.

On January 31, 1982 our 3081 was upgraded to the level of normal customer ship. This included several significant microcode improvements including VMA.

Architectural Improvements

Even before going into XA mode there are several changes in the 3081 I/O architecture versus previous machines. The three most significant are increased I/O queueing, channel timeouts, and increased I/O configuration control. In addition channel timings are much closer to OEM specification.

The 3081 channels have much greater queueing capability than previous channels. The queueing of STARTIO instructions helps to reduce supervisor overhead, since the I/O is redriven from the channel rather than the CPU. The queueing of interrupts from control units is also to allow control units to perform more work since they can pass the interrupt to the channel and are now free to continue. The channel queueing has several interesting side effects.

On earlier systems it was recommended that channels containing solid state drums be defined to CP as selector channels so that CP would not try to do more than 1 operation at a time on the channel. On a 3033 a second I/O request to the channel while the solid state drum is working would be rejected and it would be necessary for CP to queue the I/O and redrive when the channel is free. In the case of the 3081 we have defined our drum channel as a block multiplexor. The start I/O timeout is 10 milliseconds, and it takes less than 3 milliseconds to transfer a page. With a demand paging system like CP and the very small latency of the solid state device the probability is very low that the control unit would be busy the 10 milliseconds. Thus the queueing is handled in the channel not in the CPU. This has even more interesting results on 3380s with their high data transfer rate. Here a control unit must be busy for about 6 4K blocks before the channel will present a busy condition back to the CPU. The probability of this is quite low.

The increased channel interrupt queueing necessitated some changes in CP. The channel now will accept an interrupt from a control unit even when contingent connection is setup for obtaining sense information from another device on that control unit. On control units such as tapes where the control unit has a shared subchannel this will cause that subchannel to appear busy and the system must enable and accept the interrupt to clear the subchannel.

On most 370s if a program reads down a tape to where the tape was blank the read would terminate when the tape ran off the end of the reel. This both caused extra work for the operators and tied

up the channel possibly for several minutes. The 3081 channel uses a timer to recognize that a read has been outstanding for more than 10 seconds. The channel sends a selective reset to the control unit stopping the tape. Well and good. It now signals to the CPU Interface Control Check and turns on a bit in the channel logout area indicating timeout. This is an excellent hardware implementation. The software, however, did not give any different indication back to the user. On most Interface Control Checks retrying the operation is a reasonable idea. In this case that is not a reasonable idea. We have modified both CP and CMS so as to pass a different return code back to CMS in the case of such errors and CMS in this case will ABEND the command and force the tape marking of any output tapes.

Recent engineering changes on the 3081 have significantly increased the ability to monitoring, validate, change, and understand the I/O configuration. Displays are available while the system is running showing I/O status from the cpu level, thru the channel and control unit to the device. This includes the ability to verify the I/O configuration and determine the last status associated with any channel or device.

Channel timings on the 3081 are significantly closer to OEM specifications than earlier systems. On a 168 for instance System Reset sends out a signal the duration of which can be measured in hundreds of milliseconds. On the 3081 System Reset is an eleven microsecond signal. We had initial problems with both IBM and Brand X control units not recognizing these short signals.

Reliability

The 3081 is a very reliable system. Even on our very early shipped machine we have had very few cpu problems. We have never had a machine check. We have had afew channel hangs that also caused system outages. Most of the problems have been corrected with new microcode, engineering changes, and improved CE procedures. We have had one long outage (20 hours). This was caused by a net in a mother board going out. Because of Murphy's law this was of course in a board that caused the machine to go into an IML loop that failed and the CEs were unable to run diagnostics. This was clearly a manufacturing defect. I would expect that such an occurrence would be extremely rare.

Summary

The 3081 has proven to be a very reliable system. The new architecture of the 3081 promises much better capabilities than we have seen in previous systems. We were very fortunate to have obtained a 3081 when we did and although we were pioneers, we never really got any arrows in our backs.