IMPLEMENTATION OF A FAST ORBIT FEEDBACK SYSTEM AT THE ALS*

E. Williams, A. Biocca, C. Steier, T. Scarvie, M. Bennett,
Lawrence Berkeley National Laboratory, Berkeley, CA 94720, U.S.A.

Abstract

A fast global orbit feedback system has been in user operation at the ALS for over 5 years. This system was constructed using custom control software running in the EPICS environment and off-the-shelf computer and Ethernet network hardware to provide improved beam stability. An overview of the feedback system design and implementation will be presented, as well as unique issues encountered during commissioning and proposals for future improvements.

PREFACE

The requirements and operational performance of the ALS fast global orbit feedback system have been discussed in previous papers[1][2][3]. This paper will concern itself with design and implementation issues.

THEORY AND DESIGN

The ALS Fast Feedback (FFB) system was premised on certain fundamental characteristics of 100BASE-T Ethernet networks: That Ethernet switches reproduce broadcast packets received on one port and re-send them out all other ports; that, if such broadcasts arrive at multiple ports at the same time, the Ethernet switch will resolve the overlap and efficiently sequence all the broadcasts to every port; and that full-duplex Ethernet links can handle bursts of packet traffic without collisions and their associated delays or loss of data.

The ALS FFB Network operates by distributing measurement, calculation, and adjustment among 12 Compact PCI crates around the storage ring of the ALS, one crate per sector. Each crate can be configured to manage as many as 16 Beam Position Monitors (BPMs) and 8 Corrector Magnets (CMs), though rarely is more than half that capacity used. On every cycle of the FFB, each crate's computer would read its BPMs, send these readings to the other crates via UDP multicast packets, receive the readings from the other 11 crates, apply a matrix multiplication between all these readings and the inverse of the S-matrix that represents the effect of each CM on the beam at each BPM position, and apply the resulting correction to the appropriate CM.

The process begins with a Master Crate sending out a multicast UDP packet bearing the SYNC opcode. The SYNC packet also contains a sequence number which is incremented on each cycle, and a FFB state field which turns the FFB network on, off, or places it in open-loop mode (described later). Upon reception of this SYNC packet, all the crates in the FFB network sample their BPMs and compare them to their corresponding EPICS setpoint Process Variables (PVs) to determine the beam position error for each BPM. These BPM position errors are then multicast to the other sectors in a UDP packet bearing the SECTOR opcode along with the sequence number from the SYNC packet which triggered them.

As the multicasts from the other sectors arrive, each crate verifies they have the correct sequence number, and checks them off, until the BPM errors from all sectors are accounted for. Once this happens, the matrix of errors is multiplied by the portion of the inverse-S-matrix that applies to that crate's sector, resulting in correction values for the crate's CM power supply settings. These corrections are filtered through a proportional–integral–derivative (PID) filter, but are not applied to the corrector magnets until the next SYNC packet is received. If the SYNC packet is received before all sectors' data has been received, a network error is indicated and the data for that cycle is discarded.

As well as BPM setpoint PVs, each FFB crate has a PV for PID filter parameters, for CM error trip points (which, if exceeded, cause the FFB system to go off-line), and status PVs such as state and network error rate. The master crate also has PVs for turning the FFB network on or off, and for setting the operating frequency of the feedback loop, i.e. the frequency at which SYNC packets are broadcast.

IMPLEMENTATION ISSUES

The FFB network is set up on an isolated IP subnet to keep the massive network traffic from affecting other accelerator controls, and vice versa. The original configuration consisted of a single Cisco C3500 switch connecting all 12 sector crates together. However, initial testing revealed increasing transport level error counts at the switch end of some of the connections, leading to the conclusion that the interconnecting cables were too long, even though they were within the 100M 100BASE-T maximum length. It was decided to split the subnet between two switches with the switches linked to each other via gigabit optical fibre cable in order to reduce the maximum cable length to 75M. This configuration eliminated this particular error, and our theory is that some CPU boards have less robust Ethernet transmitters which are unable to reliably drive the full-length cables. While this weakness would not be apparent to most IP protocols which incorporated retries, the real-time nature of the FFB system left it vulnerable to these marginal links. Later, the C3500 switches were upgraded to C2970s to support future expansion to gigabit.

Processing at each crate is handled with a single Motorola MCP750 single-board computer. The 350MHz

*Work supported by the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

PowerPC CPU on the MCP750 proved to be entirely adequate for the needs of the FFB system, though the single 10/100 Ethernet port is somewhat of a bottleneck.

When the FFB network was initially set up, several network issues arose. Initial performance was well below what was calculated, and eventually the cause was traced to several of the crates auto-negotiating link parameters with the Cisco switches to half-duplex. There being no easy way to fix the VxWorks Ethernet driver code, we elected to lock the ports in the Cisco switch to full-duplex, and network performance then increased to within expected levels.

As testing proceeded and the FFB system operating frequency was gradually increased to 1kHz, we began to receive frantic phone calls from our network group telling us that the CPU utilization in our network routers was hitting 98% as the router examined each multicast to decide how to route it. Access Control List (ACL) rules were added to the router to ignore the multicasts from the FFB subnet, but the high CPU load persisted. Finally it was discovered that the VxWorks default TTL of 1 was forcing the CPU to examine packets despite the ACLs. Changing the TTL to 2 reduced router load to a more normal 12%.

During the course of commissioning the FFB system, it became useful to run the system without actually adjusting the CMs, allowing parameters to be tuned without triggering the trip points which would bring the FFB offline, or perturbing the storage ring beam enough to cause it to be lost. A third mode of operation, "open loop", was developed to operate the FFB system with these sensitive elements disabled.

In addition, a log buffer was added to store a record of each crate’s BPM and CM errors, so that the system's performance, impulse response, and response to parameter changes could be evaluated. A PV was added to the master crate to trigger a capture of several (currently 1000) cycles’ worth of data from all crates simultaneously, along with the sequence numbers of the cycles which allowed the different sectors’ data to be aligned. These are then dumped to an NFS-mounted file for analysis. This feature has proved useful in normal ALS operations, as the data can be used to pinpoint many sources of problems such as noisy power supplies, faulty BPMs, etc.

The ALS FFB system went on-line for beamline operations for the first time on April 22, 2004. While effectiveness of the system has been good, the reported network error rate has never reached the insignificant levels that were expected. Instead, the errors tended to periodically oscillate between .3 and 1.4%, for reasons that were never clear. While this was annoying, FFB network parameters could be tuned to where the errors did not affect the reliability of the system while still providing useful noise suppression. However, with the upgrade of the ALS for top-off mode, the errors began to occasionally cause beam dumps when a peak in network errors coincided with a fresh injection of electrons into the storage ring, prompting deeper investigation.

A network analyser was employed to capture packets on the FFB net, and software was written to analyse the packets. What we found is summarized in Figure 1. The graph shows time in microseconds on both scales, the horizontal scale shows approximately 1 minute of the operation of the FFB network. The vertical scale shows the time delay between when a SYNC packet arrives from the master crate to when its corresponding SECTOR packet is transmitted. Most of the responses are well within the 1000µs response window required for a 1kHz FFB system. But every 5 seconds there is a cluster of delayed responses – as many as 5 of them – indicating that the CPU is periodically becoming unavailable for as long as 5ms. The cause of this unavailability is as yet unknown, though we do know it is not caused by EPICS since a debugging system without EPICS showed the same characteristics.

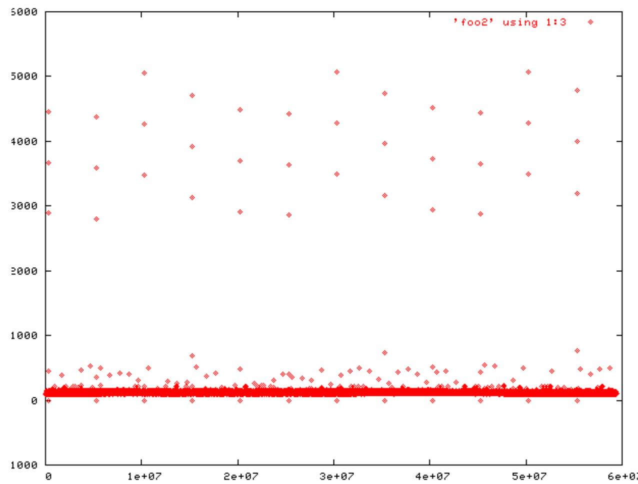

Figure 1: FFB packet response time relative to SYNC for a single sector’s crate.

While the ALS has been able to continue to operate with this source of error, it is at the expense of gain in the feedback loop which limits the system's ability to suppress noise. We will continue to investigate this phenomenon and are confident that a solution will be found.

FUTURE IMPROVEMENTS

The next planned improvement in the ALS FFB system is to upgrade the network to gigabit. When the network was originally built, gigabit was a very expensive option, whereas today it is almost impossible to find a compact PCI computer that does not have it.

Likewise, dual-port Ethernet has become widely available, and a second port will allow us to move the EPICS controls traffic to a different subnet than the one used for the more time-critical BPM multicasts, giving us the flexibility to add more diagnostic data gathering in the form of waveform records and other controls.

CPU power as well has both improved and become cheaper, and even the least expensive cPCI computer boards with Intel Atom processors offer several times the computing power of our current MCP750.

Once the issue of CPU availability has been resolved and the other upgrades completed, the frequency of operation of the FFB system can be increased to its target of 2kHz. At this frequency, the system will be operating close to the bandwidth limit of our power supplies, giving the ALS the maximum orbit stability possible.

CONCLUSIONS

At the time it was first proposed, some questioned the advisability of an Ethernet-based real-time orbital feedback system, but our experience has proven that it is a viable alternative to more exotic hardware. Since its initial design the reliability and performance of standard network hardware has continued to increase, creating a clear upgrade path for such systems without major re-design and re-build. The installation and maintenance costs of such a network can more easily be calculated and justified, as the important factors are well understood by management and the information technology industry.

But perhaps more important is that the problems encountered in installing and maintaining an Ethernet-based feedback system are identical, or closely related to, the kinds of problems all computer networks experience. It is not just the network hardware that is off-the-shelf, but also the debugging tools, diagnostics, and trained personnel who are necessary to bring up and keep such a system running that are readily available in the commercial IT market. Most large research laboratories already have the equipment and organizations in place that can support such a feedback system with little or no augmentation.

REFERENCES

[1] C. Steier, A. Biocca, et al., "Design of a Fast Global Orbit Feedback System for the Advanced Light Source," PAC'01.
[2] C. Steier, A. Biocca, et al., "Orbit Feedback Development at the ALS," EPAC'02.
[3] C. Steier, E. Domning, et al., "Operational Experience Integrating Slow and Fast Orbit Feedbacks at the ALS," EPAC'04.