Fusion Engineering and Design 85 (2010) 447–450

Contents lists available at ScienceDirect

Fusion Engineering and Design

journal homepage: www.elsevier.com/locate/fusengdes

Plasma control system upgrade and increased plasma stability in NSTX

D. Mastrovito*, D. Gates, S. Gerhard, J. Lawson, C. Ludescher-Furth, R. Marsala

Princeton Plasma Physics Laboratory, P.O. Box 451 Princeton, NJ 08543, USA

ARTICLE INFO

Article history:
Available online 18 February 2010

Keywords:
Real-time control
NSTX
PCS

ABSTRACT

Plasma control on the National Spherical Torus Experiment (NSTX) was previously accomplished using eight 333 MHz G4 processors built by Sky computers. Several planned improvements and additional control algorithms required significant upgrades to our real-time control computers and real-time data acquisition infrastructure. Several in-house modules have been designed and implemented including: the digital time stamp module (DITS) and for digital/analog front panel data port (FPDP) output, the FPDP output module digital/analog (FOMD/A). Standard Linux based Intel computers perform the real-time control tasks and InfiniBand as been employed for communication between a user-accessible "host" server and the real-time computer. In addition to several independent real-time processes the General Atomics developed PCS (Bell (2006) [1]) system infrastructure continues to be used on NSTX. While maintaining previous functionality, improvements in the control system software include: an RWM feedback algorithm, beta feedback NBI control, more comprehensive error logging and trapping, more user-friendly interface, more complete archiving and restoring functionality, and better status reporting and diagnostic tools. Once completed, we succeeded in increasing overall plasma stability and decreasing control system latency by several times.

1. Introduction

National Spherical Torus Experiment (NSTX) began operating in 1999 and performs with typical parameters Ip $\leq$ 1.5 MA, BT $\leq$ 0.55 T, Neutral Beam Injection (NBI) heating $\leq$7 MW and a 1.5 s maximum pulse length [1,2]. Plasma control on NSTX was until recently accomplished using a SKY computer consisting of eight 333 MHz G4 processors [3]. Scheduled programmatic goals including improved MHD mode stabilization and increased plasma stability for longer pulses motivated the addition of an RWM feedback algorithm and a beta feedback NBI control algorithm requiring the addition of 64 new real-time data acquisition diagnostic channels and 6 new output control commands. The increased demands on processing power and timing constraints given the additional *I*/*O* latency, required a major system upgrade. To accommodate these new demands we devised a plan to replace this system with more modern technology utilizing readily available commodity computers, which are more easily maintainable. Physically distributed real-time diagnostic inputs to the control system were previously acquired heterogeneously with some acquired over fiber optic FPDP serial link and others over VME bus [3]. The system upgrade included a plan to combine and thus streamline all real-time *I*/*O* over Front Panel Data Port (FPDP). Other upgrade goals included developing a more accurate timing system, increasing our ability to debug and troubleshoot, decreasing the overall latency in the system, and increasing the friendliness of the user interface. In this paper we will report on the current status of this upgrade.

2. Hardware improvements

As a replacement for our SKY computers for real-time computations, we purchased a Sun Fire V40z with 4 dual core AMD Opteron 880 2.4 GHz and 8 GB of memory, giving us the same number of processors at 7 times the previous speed. For "host" processes which provide for user interaction and reporting of real-time status and error information, we purchased an HP DL385 Proliant Server with one 2.4 GHz AMD Opteron 250. Communication between the host and real-time computers occurs over Infiniband using Mellanox Infiniband PCI cards.

The previously used SKY computers operated on a VME backplane which also functioned as the mechanism for collecting real-time data from several locations into the real-time computer via a VME-based FPDP digitizer. We eliminated the need for the VME backplane in our control system by replacing our previous FPDP Merlin digitizers with two new VMETRO FPDP PCI DPIO2 modules, one each for input and output. The boards are bi-directional and can be programmed at run-time via software to function as either an input or output board. Each has a 128 kB first in first out buffer (FIFO) and can transfer data at a maximum rate of 510 MB/s. Previously our gas injection system control/acquisition

---

* Corresponding author.
E-mail address: dmastrovito@pppl.gov (D. Mastrovito).

0920-3796/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.fusengdes.2010.01.005

was handled separately from the rest of our FPDP data, its *I*/*O* being accomplished purely with VME digitizers and D/A converters. Therefore, it was necessary to replace this VME-based hardware. To this end we developed two new boards for interfacing with the now consolidated FPDP *I*/*O* data stream, the FPDP output module analog (FOMA) and digital (FOMD) [4]. Both were developed using complex programmable logic devices (CPLD's). The FOMD provides four banks of 16-bit digital outputs, while the FOMA has eight analog output channels having 14-bit resolution.

In our previous system, time was computed relative to a hardware start trigger at a known time; subsequent input data time stamps were calculated based on the known input data frequency and the real-time CPU clock. In order to achieve an absolute time stamp for each data sample relative to the NSTX start of pulse (SOP), we developed a new digital input and time stamp module (DITS) [4] also using a CPLD. The DITS provides a 48-bit timestamp in microseconds for each input data sample, a 32-bit data-block counter, 56 bits of digital inputs and eight bits of digital outputs and a 16 MHz clock, which however, is limited by the speed of the FPDP and thus can operate with a maximum clock rate of about 1.5 MHz. One of the digital inputs is used as a signal for timestamp reset. This is done at 60 s before (T-60) the start of the shot (T-0) which signals the software that a new shot cycle has begun.

3. Software improvements

Standard Red Hat Enterprise Linux (RHEL) currently supports only 2 processors, therefore, in order to accommodate our 4 dual

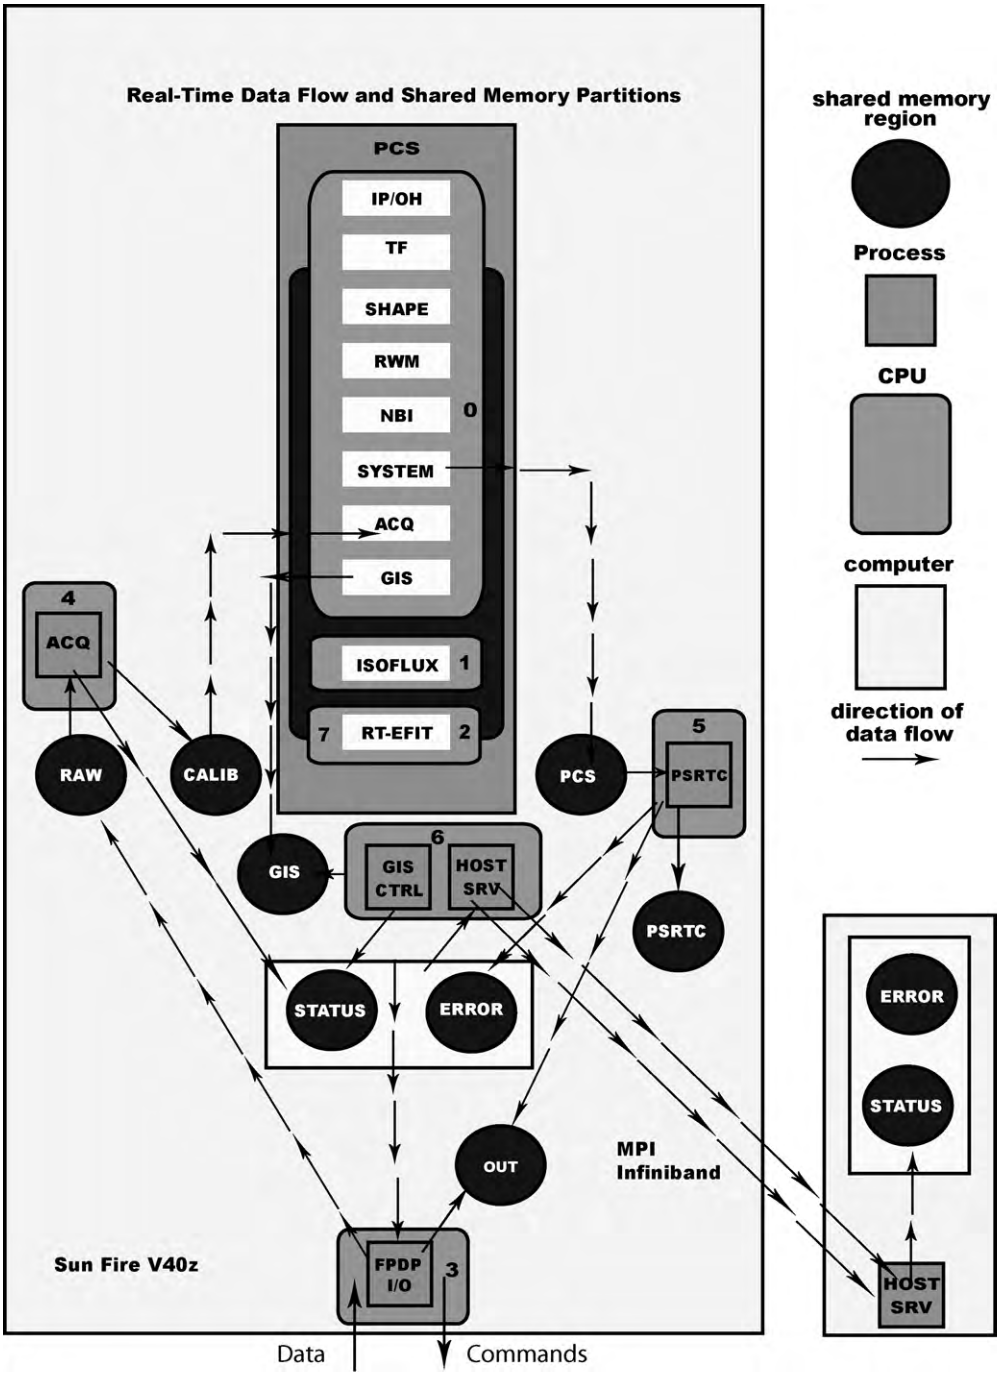

Fig. 1. Diagram of NSTX plasma control system shared memory regions, processes and data flow.

core Sun Fire, we purchased a license for RHEL Advanced Server (RHEL AS). RHEL AS is not a real-time operating system, however, we have customized the operating system using kernel modifications developed at GA and modified for use with linux 2.6 at PPPL which allow operating system interrupts to be shut off and turned back on just before and after the shot thereby allowing unhindered access to the processors for real-time computation for the duration of the shot [5].

Real-time control on NSTX is now accomplished using five different processes: front panel data port input/output (FPDP_IO), acquisition (ACQ), General Atomics developed plasma control system (PCS) [6], gas injection system (GIS), and power supply real-time control (PSRTC). Control processes are kept modular in case of any software problem resulting in the failure of intermediate portions of the system. This allows more reliable access to raw data coming into/out of the system, including reliable timestamp information to aid in the trouble-shooting process. FPDP_IO is the first process in the chain, responsible for retrieving data from the FPDP input board direct memory access (DMA) region and placing it in shared memory. FPDP_IO also functions as the last process in the chain, retrieving final output voltage requests from shared memory and placing them into a second DMA region for the output FPDP board. ACQ reads the raw data from shared memory, converts it to physical units, applies calibration factors, monitors the time coming from the DITS, calculates and subtracts raw data baselines, and places this calibrated data into shared memory. During plasma shots, the PCS, discussed in detail below, reads these values and uses them to generate voltage requests. The PSRTC [7] is responsible for retrieving those voltage requests and after applying engineering constraints, converts them to appropriate output units, pairing them with power supply addresses, and places them into shared memory for access by the FPDP_IO process. The GIS is responsible for generating commands to the injection valves, opening and closing them in real-time as needed based on commands from the PCS. Both the GIS and the PSRTC are also able to operate independently of the PCS and, in combination with ACQ, continue to collect data at a slower rate (1 Hz) between shots to allow for monitoring of coil temperatures, etc. Currently real-time data is acquired during shots at a rate of 5 kHz.

The PCS is responsible for calculating power supply request voltages based on PPPL-defined physics models as well as a real-time EFIT approximation. The PCS has been used for plasma control on NSTX previously, however, a new version of the PCS has been implemented for use on multiprocessor computers. This version allows for communication via shared memory between various PCS algorithms in real-time and consolidates all user interaction processes on the "host" machine which communicate with the real-time computer between shots via socket connection with a message server. The PCS utilizes the first 3 processors on the real-time computer (0–2) and a thread has recently been added on CPU 7 for a separable portion of rt-EFIT, which contributes beta calculations to a new NBI control algorithm.

The NBI algorithm operates using a proportional gain scheme and is used to adjust the injected beam power to achieve a desired value of β below instability thresholds.

$$\Delta P_{\text{inj}} = G \cdot C \cdot \text{Error}$$

$$\text{Error} = \beta_{\text{rt-EFIT}} - \beta_{\text{req}}$$

where ΔP_{inj} is the required increase or decrease in injected neutral beam power, G is a gain coefficient, and C is a function of plasma current, volume, magnetic field, typical energy confinement time and minor radius of a typical discharge. $\beta_{\text{rt-EFIT}}$ is the value of β calculated in rt-EFIT, β_{req} the requested β. β feedback can be selected at run-time for β normal β_N or β torroidal β_T.

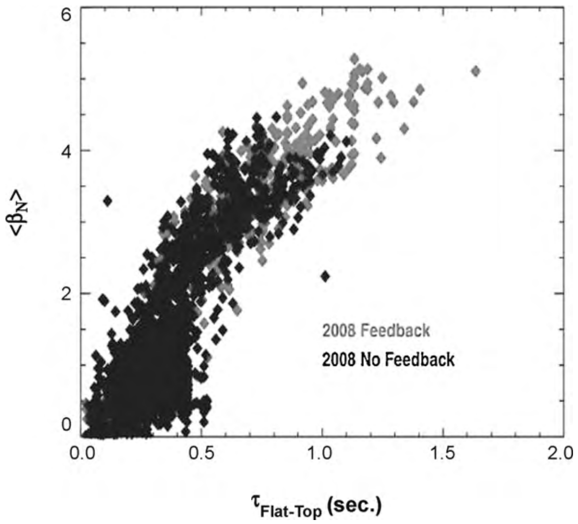

Fig. 2. Comparison of β_N averaged over the plasma current flat-top for shots with (grey) and without (black) error field.

In addition to the real-time processes, there is an independent process called the hostsrv which does not actively participate in control. It runs partly on the real-time computer and partly on the "host" computer acting as an overall system monitor responsible for reporting real-time status and error information to the operator. This information is stored in a shared memory region on the host machine where it is retrieved and displayed in a new user interface. Prior to this system upgrade, there was no user interface or status information reported to the NSTX operators. Other new functionality provided by this interface is the ability to change, save, and restore PSRTC settings. Communication between the two halves of the hostsrv process is accomplished using Message Passing Interface (MPI) over Infiniband. Fig. 1 shows the current layout of processes, shared memory regions, and data flow of the control system.

3.1. Improved plasma stability

Non-axisymmetric field control gained through use of the RWM algorithm, in addition to other factors such as lithium condition-

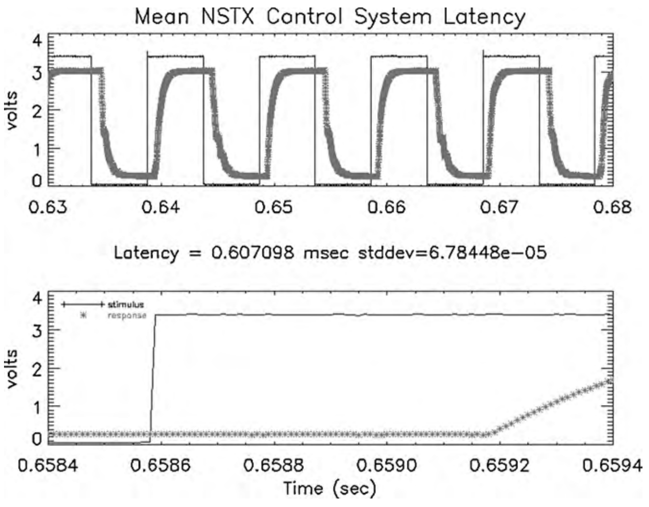

Fig. 3. System latency determined by the difference in time between stimulus signal and control response (indicated by *) measured to be .6 ms.

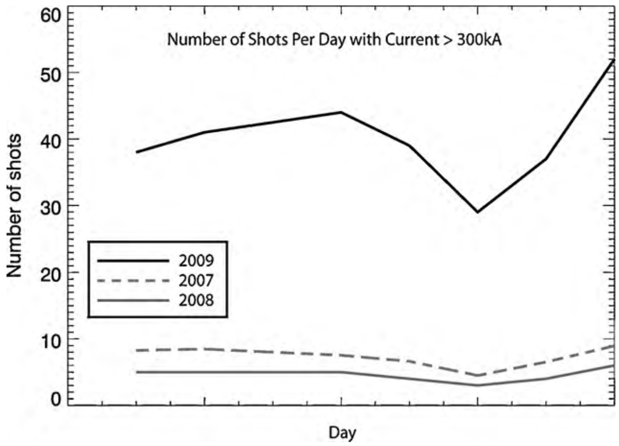

Fig. 4. Increased reliability of the control system as shown by the number of successful shots (with plasma current >300 kA) per day per year.

ing, has been a significantly contributing factor in achieving shots with high $\beta_N = 7.2$ [8] and record breaking shots lasting 1.8 s having plasma flat-top of 1.6 s. Fig. 2 shows average β_N (averaged over the plasma current flat-top) plotted versus the length of the flat-top, spanning the entire NSTX database for 2008 [9], with black points representing discharges that did not have error field control, and red points representing ones plasmas that did.

4. Status of the system

Latency tests are conducted by replacing the PF3 current signal with one from a function generator producing output of 0–10 V at 100 Hz acting as a square wave stimulus to the control system. Then a comparison can be made between those stimulus signals and the measured corresponding response from the control system. Fig. 3 shows 5 cycles of this square wave stimulus and the control system response during a latency test (top). The latency was measured to be about .6 ms or an improvement of 5 times over the latency in our previous system. Fig. 4 demonstrates the increased reliability of the overall control system following our system upgrade. On average we have 4 times the successful shots having plasma current >300 kA per day as previously achieved.

Acknowledgement

This work was supported by DOE Contract DE-AC02-76CH03073.

References

[1] M.G. Bell, et al., New capabilities and results for the National Spherical Torus Experiment, Nucl. Fusion 46 (2006) S565–S572.
[2] M. Ono, et al., Exploration of spherical torus physics in the NSTX device, Nucl. Fusion. 40 (2000) 557–561.
[3] D.A. Gates, J.R. Ferron, M. Bell, T. Gibney, R. Johnson, R.J. Marsala, et al., Status of the control system on the National Spherical Torus Experiment, Fusion Eng. Des. 81 (2006) 1911.
[4] M. Isaacs, R.J. Marsala, J.E. Lawson, P. Sichta, Flexible FPDP I/O and time stamping for the NSTX plasma control system, in: Sixth IAEA technical meeting on control, data acquisition and remote participation for fusion research, Inuyama, Japan, 2007.
[5] B.G. Penaflor, J.R. Ferron, D.A. Piglowsky, R.D. Johnson, M.L. Walker, Real-time data acquisition and feedback control using Linux Intel computers, Fusion Eng. Des. 81 (2006) 1923–1926.
[6] B.G. Penaflor, D.A. Piglowski, J.R. Ferron, M.L. Walker, Current, status of DIII-D real-time digital plasma control, IEEE Trans. Nucl. Sci. 47 (2000) 201.
[7] D.A. Gates, J.R. Ferron, M. Bell, T. Gibney, R. Johnson, R.J. Marsala, et al., Plasma shape control on the National Spherical, Torus Experiment (NSTX) using real-time equilibrium reconstruction, Nucl. Fusion 46 (2006) 17.
[8] S.A. Sabbagh, J.M. Bialek, R.E. Bell, et al., The resistive wall mode and feedback control physics design in NSTX, Nucl Fusion 44 (2004) 560–570.
[9] D.A. Gates, et al., Overview of results from the National Spherical Torus Experiment (NSTX), Nucl Fusion 49 (2009).