Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Fusion Engineering and Design 81 (2006) 1975–1979
Fusion Engineering and Design
www.elsevier.com/locate/fusengdes

Storage and analysis techniques for fast 2D camera data on NSTX

W.M. Davis [a,*], D.M. Mastrovito [a], C.E. Bush [b], D.A. Gates [a], R.J. Maqueda [c], J.E. Menard [a], N. Nishino [d], P.G. Roney [a], A.L. Roquemore [a], S. Sabbagh [e], B.C. Stratton [a], S.J. Zweben [a]

[a] *Princeton Plasma Physics Laboratory, P.O. Box 451, Princeton, NJ 08543, USA*
[b] *Oak Ridge National Laboratory, Building 9201-2, P.O. Box 2009, Oak Ridge, TN 37831, USA*
[c] *Nova Photonics, P.O. Box 451, Princeton, NJ 08543, USA*
[d] *Department of Engineering, Hiroshima University, 1-4-1 Kagamiyama, Higash-Hiroshima 739-8527, Japan*
[e] *Department of Applied Physics, 200 S.W. Mudd Building, MC 4701, Columbia University, New York, NY 10027, USA*

Available online 9 June 2006

Abstract

Animations from fast 2D camera data are facilitating the investigation of spatially distributed phenomena in high-temperature plasmas. The National Spherical Torus Experiment (NSTX) now has six fast camera systems, and more are expected to be added. Image capture rates vary between 1000 and 500,000 frames per second. Archiving and retrieving this data is a challenge for data repositories and networks. For example, if all camera data had been archived during the 2004 run, the total amount of data from NSTX (300MB per pulse) would have doubled, and this year, one new camera alone can acquire 2GB per pulse. The paper will describe the storage strategies, and compare some data compression techniques used for NSTX. Tools which animate camera data, synchronized with displays from other plasma diagnostics and simulations, allow scientists to gain insights and observe correlations that would be difficult with conventional tools. A labor-saving technique is described for archiving fast camera images from a vendor's system into MDSplus and to AVI files. Examples of specific analysis and display techniques are presented. Future challenges are also discussed.

Keywords: NSTX; MDSplus; Fast cameras; Data compression; Plasma images

1. Introduction

The National Spherical Torus Experiment (NSTX) began fusion experiments at the Princeton Plasma Physics Laboratory (PPPL) in 1999 [1]. During an experimental cycle, or "shot," a plasma is produced for about 1 s and about 400MB (uncompressed) of raw data

* Corresponding author. Tel.: +1 609 243 2546; fax: +1 609 243 3086.
E-mail address: bdavis@pppl.gov (W.M. Davis).

0920-3796/$ – see front matter © 2006 Published by Elsevier B.V.
doi:10.1016/j.fusengdes.2006.04.024

is acquired from instruments in dozens of subsystems, hosted by Unix, VMS and Windows computers. This raw data is, for the most part, transferred to and stored on centralized data servers and is available within a few minutes to anyone "subscribed" to particular signals. Some of this raw data is also input to analysis programs which store their results in similar structures, or in databases, both to provide immediate feedback to the machine operators and diagnostic physicists so they may make adjustments for the next shot, and to provide a repository of information for later off-line analysis. A typical NSTX run day produces about 40 shots. NSTX runs for between 60 and 80 days a year. 2.7TB of raw and analyzed NSTX data (compressed) currently reside on disk. 2D camera data is accounting for an increasing proportion of the total data on NSTX. Two-dimensional camera data is a valuable tool for studying turbulence, gas injection, ion transport, pellet injection, wall interaction, and other phenomena, which affect energy confinement and heating efficiency in fusion energy production.

Most raw and analyzed data from NSTX, including raw camera images, is stored in MDSplus [2,3] but animation files are stored elsewhere currently. Better compression techniques are being investigated since our high-storage, high-availability disk space currently costs over $8 (US) per gigabyte, and networks and tape backups add to the cost.

2. 2D camera data storage requirements

Table 1 shows examples of four of the six fast camera diagnostics in use on NSTX. Some cameras are used for a visual inspection of large-scale features, such as gas injection and surface heating. For such purposes 30–1000 frames per second (fps) can be adequate. Longer exposure times, such as those for the large image in Fig. 1, blur small-scale, quickly changing phenomena, such as the filaments shown in the 9 μs exposure from the Phantom 7 camera in Fig. 2.

Table 1
2D camera capabilities of selected NSTX diagnostics

Diagnostic	Maximum fps	No. of frames	Typical resolution	No. of bits	Total MB
RF antenna	5000	300	512 × 240	14	36
Soft X-ray	500000	300	64 × 64	14	3.6
Divertor	40000	8192	64 × 64	8	32
Phantom-7	150000	175000	64 × 64	12	2000

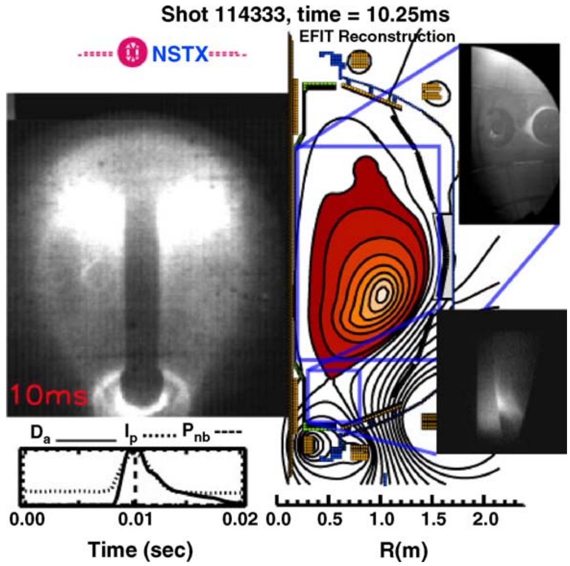

Fig. 1. Data from three cameras combined with an EFIT reconstruction [10] and plots of other plasma conditions vs. time. As time progresses in the composed animation, the vertical dashed line on the plasma-conditions vs. time plot moves to the right, and individual frames from each camera and the flux contours from EFIT, are displayed at the appropriate time.

The faster, higher-resolution cameras place a large burden on the data archiving system. The Phantom-7 camera [4] used for, among other things, the Gas Puff Imaging Diagnostic, can produce 2GB of data per shot, more than all the other NSTX diagnostics combined. Currently, a diagnostician monitors the data between shots and manually saves sequences of interest in the central repository. Any other data of potential interest is saved directly to DVDs. Similarly, the Divertor Camera (the small images in Fig. 1) [5] save most data only on the local PC.

3. Archiving 2D camera data into MDSplus and AVI files

NSTX 2D camera diagnostics use several ways to store their images in a central repository. Most write the frames into MDSplus, either as 2D individual

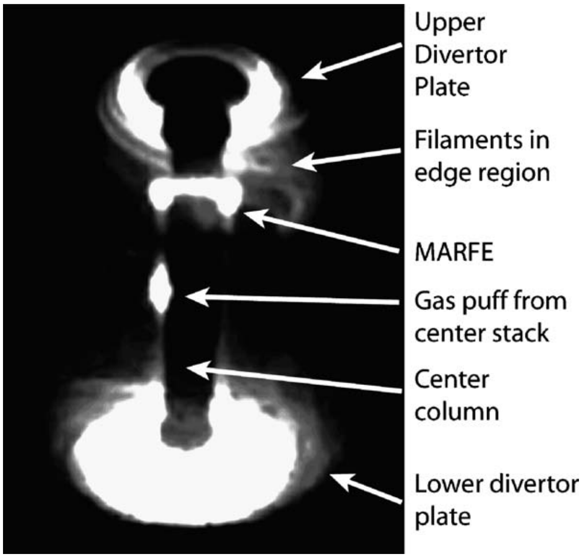

Fig. 2. Nine microsecond exposure from a Phantom 7 camera showing a MARFE on the center stack and fast filamentation on the edge (128 × 128 pixels, contrast enhanced).

frames, or as 3D arrays. Most camera diagnostics arrive at NSTX as "turn-key" systems, with control code written in LabView, Java, or C. When the source code is available, "hooks" may be added to write directly into MDSplus. However, in most cases it is easier and more maintainable to use a system that is independent of a vendor's software.

One such system is illustrated in Fig. 3. A hardware start trigger (usually from a CAMAC module set through MDSplus) is received by vendor-provided camera and PC system (1). The vendor software writes individual TIFF files to a directory on the PC whose name is set to the NSTX shot number initially, and incremented for every shot. A PPPL-written program installed on the PC, PCDA_plus, waits for an MDSplus event known to occur after all the frames have been written, "ACQ_DONE" in this example (2), and initiates a secure copy (scp), using a public-domain program called pscp.exe [6] (3). Previously generated Kerberos [7] tickets allow this transfer to occur securely without passwords needing to be entered each time or being sent in clear text. When the transfer is completed, the MDSplus event "MOVED" is declared (4). Another program running on the Linux server (or anywhere with an MDSplus connection) waits for the "MOVED" event and writes all the frames into MDSplus as one 3D signal (5), and declares the event "LOADED" when finished (6). Another program on the AVI PC waits for this event, reads the signal from MDSplus (7) and creates an animation file in a web-accessible directory (8). This program uses public-domain IDL code and DLLs [8] to create AVI (Audio-Video Interleave) files

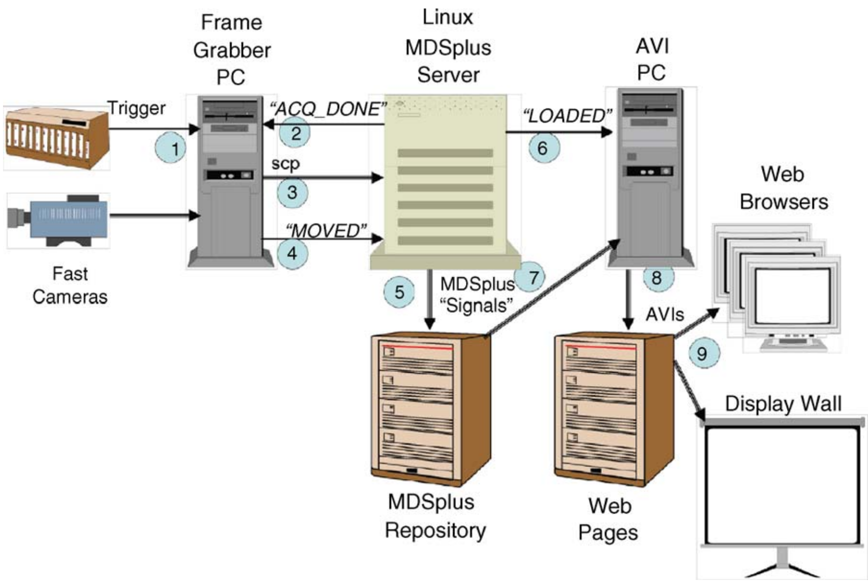

Fig. 3. Event-driven fast camera data archiving.

with a choice of compression–decompression methods (codecs). Anyone can then play these animations in a browser, download them to their favorite movie player, or send them to the NSTX display wall (9). The codec used for compression must also be on the play back computer.

4. Compression techniques

MDSplus has an option for writing data in compressed format, but there is only one compression method for all data types. For our eight-bit integer image data, we typically see a compression ratio of 1.3–2.0 using MDSplus compression. Standard image compression methods vary substantially depending on the data (see Table 2), but usually do much better. PNG (Portable Network Graphics) gives the best compression ratio without loss of data quality.

Many images of plasmas are non-distinct and can tolerate the lower-quality JPEG compressions and still show the relevant features. Since video compression methods can compress in time, as well as within a frame, they can achieve even higher compression ratios. Table 3 shows some results of 24-bit animations from NSTX for various codecs. Clearly, for optimal storage, the results of various compression methods should be examined for each type of image.

Table 2
Compression ratios for two types of 2D image files[a] using different methods

Method	Plasma only	Complex graphs[b]
GIF	3.5	4.0
JPEG Qual = 100	4.6	2.5
PNG	6.5	6.7
JPEG Qual = 75	28.2	7.7
JPEG Qual = 25	39.1	14.2

[a] Eight-bit images.
[b] The complex graphs contain lines and text as well as plasma images, as shown in Fig. 1.

Table 3
Compression ratios for two types of animation files[a] using different compression methods (codecs)

Codec used	Plasma only	Complex graphs
Cinepak by Radius	11	32
µSoft Video 1	22	14
MPEG-1	41	89
Indeo Video 5.10	64	104
XviDMPEG-4	184	170
VP31 compressor	384	70

[a] 24-bit images.

5. Analysis techniques

Image analysis techniques of NSTX plasmas usually consist of visual examination and manual correlation with other phenomena, such as MHD activity, the onset of H-mode, structure in light emission signals, etc. [9]. Composing synchronized animations from camera images, EFIT reconstructions [10], and time histories, as shown in Fig. 1, facilitate making visual correlations. Some quantitative studies have measured the size and motion of regions of enhanced local line emission near the plasma edge, commonly referred to as "blobs," and compared them to other diagnostics and to theoretical models [11]. More of such work is needed in order to utilize the increasingly large amounts of camera data.

6. Future challenges

Fast 2D camera data will continue to make up an increasing portion of the data acquisition load on NSTX and other machines. Saving this data will be challenging, especially for long pulse and continuous-operation machines. We need to determine which compression techniques are best suited for various diagnostics and what the acceptable trade-offs are between detail losses and storage costs, if any.

It is convenient to access one- and two-dimensional signals stored in MDSplus with general tools that do not have to be custom made, but few generalized tools for MDSplus deal with animations. Some lossless compression occurs when images are stored in MDSplus (a factor of 1.3–2), but built-in methods for images would save much more space.

Although storage requirements can be met given sufficient expenditure, effectively analyzing huge amounts of data is still a problem. As the operator of the Gas Puff Imaging Diagnostic said of the data shown in Fig. 2 “almost every sequence of 300 frames within the 45,000 has something interesting . . .and different from the previous 300 frames sequence” [12]. Techniques are needed to categorize image sequences automatically, and to store and retrieve them in databases. Image databases exist [13], and descriptive keywords can be entered by hand, but automatic discovery techniques for plasma phenomena would be valuable.

7. Summary

Fast 2D camera diagnostics are important for understanding important phenomena in plasmas, and are becoming increasingly prevalent on NSTX. Storing and analyzing the data is challenging, both in terms of labor (programming and analyzing) and in terms of storage media. Compression methods help, but should be applied intelligently for the proper trade off between resolution loss, cost and accessibility. Automated analysis methods are only beginning to aid NSTX physicists.

Acknowledgement

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

References

[1] S. Kaye, M. Ono, Y.-K.M. Peng, D.B. Batchelor, M.D. Carter, W. Choe, et al., The Physics design of the National Spherical Torus Experiment, Fusion Technol. 36 (1999) 16, http://www.nstx.pppl.gov/.

[2] J.A. Stillerman, T.W. Fredian, K.A. Klare, G. Manduchi, MDSplus data acquisition system, Rev. Sci. Instrum. 68 (1) (1997) 939, http://www.mdsplus.org/.

[3] W. Davis, P. Roney, T. Carroll, T. Gibney, D. Mastrovito, The use of MDSplus on NSTX at PPPL, Fusion Eng. Des. 60 (2002) 247–251.

[4] Phantom 7 Fast Camera, http://www.scitech.com.au/cameras/highspeed/pdf/visionres/Phantomv7.0spec&info.pdf.

[5] N. Nishino, L. Roquemore, R. Kaita, S. Zweben, D. Johnson, H. Kugel, et al., Results with the NSTX fast divertor camera, J. Plasma Fusion Res. 78 (2002).

[6] Open ssh for Windows, http://www.openssh.com/windows.html.

[7] Kerberos, http://web.mit.edu/kerberos/www/.

[8] IDL to AVI code for PCs http://www.rsinc.com/codebank/search.asp?FID=139.

[9] R. Maingi, et al., H-mode turbulence, power threshold ELM and pedestal studies in NSTX, Nucl. Fusion, 2006, submitted for publication.

[10] S.A. Sabbagh, S.M. Kaye, J. Menard, F. Paoletti, M. Bell, R.E. Bell, et al., Equilibrium properties of spherical Torus plasmas in NSTX, Nucl. Fusion 41 (2001) 1601.

[11] S. Zweben, R. Maqueda, D.P. Stotler, A. Keesee, J. Boedo, C. Bush, et al., High-speed imaging of edge turbulence in NSTX, Nucl. Fusion 44 (2004) 134–153.

[12] e-mail correspondence from R. Maqueda to W. Davis June 30, 2005.

[13] Almagest OpenSource software, http://www.princeton.edu/~almagest/opensource/.