Improved Data Acquisition Methods for Uninterrupted Signal Monitoring and Ultra-Fast Plasma Diagnostics in LHD[*]

Hideya NAKANISHI, Setsuo IMAZU, Masaki OHSUNA, Mamoru KOJIMA, Miki NONOMURA,
Mamoru SHOJI, Masahiko EMOTO, Masanobu YOSHIDA, Chie IWATA, Hitoshi MIYAKE,
Yoshio NAGAYAMA and Kazuo KAWAHATA

National Institute for Fusion Science, Toki 509-5292, Japan

To deal with endless data streams acquired in LHD steady-state experiments, the LHD data acquisition system was designed with a simple concept that divides a long pulse into a consecutive series of 10-s "subshots". Latest digitizers applying high-speed PCI-Express technology, however, output nonstop gigabyte per second data streams whose subshot intervals would be extremely long if 10-s rule was applied. These digitizers need shorter subshot intervals, less than 10-s long. In contrast, steady-state fusion plants need uninterrupted monitoring of the environment and device soundness. They adopt longer subshot lengths of either 10 min or 1 day. To cope with both uninterrupted monitoring and ultra-fast diagnostics, the ability to vary the subshot length according to the type of operation is required. In this study, a design modification that enables variable subshot lengths was implemented and its practical effectiveness in LHD was verified.

1. Introduction

To deal with endless data streams, steady-state fusion experiment in the Large Helical Device (LHD) used a new concept that divides a long-pulse shot into a consecutive series of 10-s "subshot" chunks [1, 2]. The concept was implemented in the LHD data acquisition and storage system, i.e., the LABCOM system. Hence, long-pulse plasma diagnostic data have been acquired according to this simple concept since the beginning of LHD's steady-state experiments. Currently, the LABCOM system comprises many 100 MB/s high-speed real-time (RT) data acquisitions (DAQs) in addition to non-RT batch processing ones using conventional CAMAC digitizers. However, the latest digitizers applying high-speed PCI-Express (PCIe) technology output a nonstop gigabyte per socend data stream from each DAQ frontend. The chunk size becomes too big for a "subshot" if the 10-s rule is applied.

Furthermore, for steady-state fusion plants, uninterrupted monitoring of the environment and state of device soundness are required. For the past 10 years, the acquisition of device soundness information and environmental radiological dosage in LHD has been done by dedicated computer systems. These older computers have recently been replaced by the LABCOM system as a part of the standard DAQ framework of LHD.

To cope with both uninterrupted monitoring and ultra-

fast diagnostics, a fixed subshot length of 10 s was modified to enable acquisition at variable lengths according to the need. In this study, the design modification for variable subshot lengths was implemented in LHD and its effectiveness verified.

2. Data Trend in LHD

The amount of data in fusion plasma diagnostics continues to grow at a rate of approximately 50% per year, which corresponds to an increase of 10 times (1000%) in 5 years and 100 times in 10 years. This tendency can be observed in LHD very clearly (Fig. 1), and similar trends are also found in many fusion experiments worldwide. This growth rate shows a very good fit to the well-known *Moore's law*, doubling every 18 months.

In the last experimental campaign of LHD, an ultra-fast DAQ of approximately 600 MB/s and two uninterrupted 24/7 continuous monitoring DAQs contributed considerably to data growth. Even though the production rate of the latter is only about 1.1 MB/s, the daily amount adds up to 97 GB/day, which corresponds to 6 shots of plasma diagnostic data, 16.6 GB/shot, and 6% of the weekly amount. The former fast DAQ produces 1–2 GB/shot, which also corresponds to 6%–12% of the shot data.

Even though the legacy CAMAC digitizers have low RT acquisition capability, the WE7000 and CompactPCI (cPCI) systems can achieve RT acquisition with uplink throughputs of 2.2 MB/s and approximately 600 MB/s, respectively. As Fig. 1 shows, the number of cPCI DAQs is
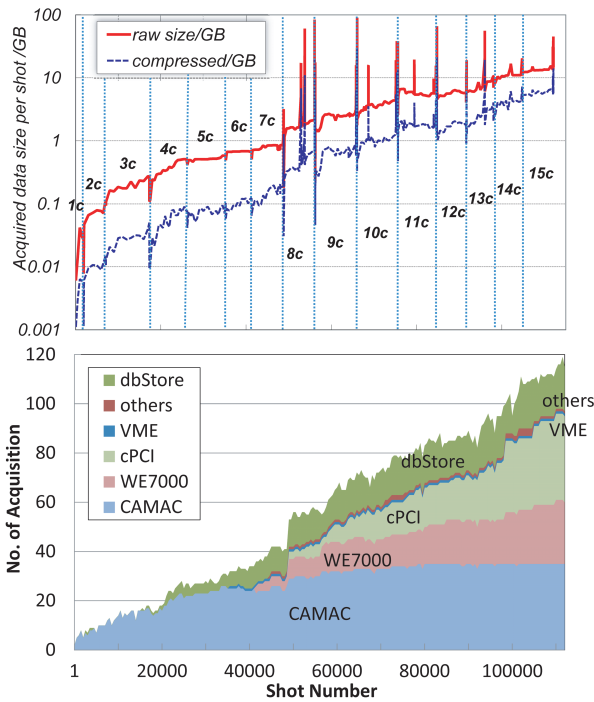
Fig. 1 Trend of data growth and DAQ increase: Acquired data per shot shows steady growth over the past 10 years and more. Compressed size shows the usage of storage capacity. Spikes indicate long-pulse experiments (top). The number of DAQs has increased since the beginning of the LHD experiment (bottom). "dbStore" indicates manual registrations of user data.

increasing at a high rate which is the main cause of the recent data growth.

3. LABCOM Real-Time Data Acquisition and Storage

The LABCOM DAQ system primarily consists of a number of distributed DAQ nodes, storage servers, and an index database that manages all distributed data locations. LHD usually has approximately 180 shots per day, each of which now outputs 16.6 GB/shot every 3 min. The key objectives of the system are as follows:

1. RT DAQ at the same rate of burst sampling, typically 1 MS/s/ch
2. Performance scalability and topological extendibility.

In general, massively parallel processing architecture is essential for providing scalable I/O performance for large-scale data archives. Currently, LHD uses 100 DAQ nodes in parallel. As modern plasma diagnostics often uses several tens of signal channels for 2-D or 3-D profile measurements, each DAQ is typically required to handle data streams exceeding 100 MB/s. We have been developing high-speed RT DAQ since 2001 [3] in which we have adopted the cPCI standard for the digitizer frontend platform.
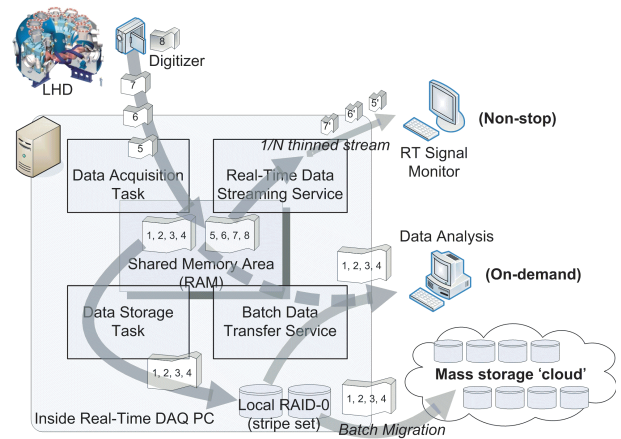

Fig. 2 Internal diagram of real-time DAQ: Consecutive incoming blocks from the digitizer are stored in the shared volatile memory (RAM) until a full subshot is accumulated, and then the blocks are written onto a persistent subshot file on a local HDD. Another independent task migrates the subshot file into the outer mass storage area. On-demand data clients may retrieve the last subshot data from the RAM before its local storage is completed [4,5]. Sequential numbers "1, 2, 3, ..." indicates the order of consecutive data stream slices.

3.1 Nonstop steaming DAQ

To achieve throughput greater than 100 MB/s in the DAQ frontend, we adopted the NI PXI and PXI-Express (PXIe) digitizers that are compatible with cPCI and CompactPCI Express standards, respectively. These digitizers can stream a maximum of 110 MB/s and approximately 600 MB/s consecutive data blocks, respectively.

As Fig. 2 shows, the digitizer's output blocks accumulate in the dedicated shared area of the DAQ PC's volatile memory (RAM). The data stream may be delivered from the RAM via different routes for storing, RT streaming, and batch transferring. For each route, a different block size must be chosen to obtain the best I/O performance through the respective device ports.

To store so-called endless data streams having indefinite time duration, we developed a new idea of a "subshot" that cuts the stream into 10-s time chunks [4, 6]. This enables us to access the data even while data writing continues until the end of the long pulse discharge. These subshot data can be retrieved by the key numbers of "shot" and "subshot" such as #110000.1, #110000.2, #110000.3, ... for the long pulse experiment.

The 10-s "subshot" rule is also implemented in the data storage and retrieval functions of the LABCOM system. However, this rule was designed primarily for plasma diagnostic DAQs by considering earlier PC performance. Hence, the fixed length of the subshot has gradually become unsuitable for other types of slow continuous monitoring and much faster sampling rates.

3.2 Ultra-fast vs. 24/7 slow monitoring

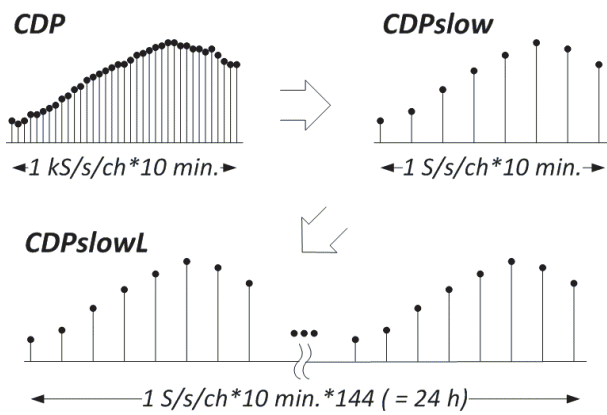A 100 MB/s DAQ produces 1 GB/shot data in which

Fig. 3 Resampling of raw CDP signal: By using the original 1 kS/s CDP signal, a 1 S/s resampled set is automatically generated to enable quick examination of the weekly/monthly signal fluctuations.

Table 1 DAQs for 24/7 slow monitoring: CDP and RMSAFE denote the raw data acquisitions, and CDPslow(L) denotes the acquisitions resampled from CDP. "ss" denotes "subshot".

	rate	sample/ss	period	ss/day
CDP	10^3 s^{-1}	600 000	10 min.	144
CDPslow	1 s^{-1}	600	10 min.	144
CDPslowL	1 s^{-1}	86 400	24 h.	1
RMSAFE	0.2 s^{-1}	120	10 min.	144

one channel has typically tens of megabytes of data. A new NI PXIe-5186 digitizer, however, can sample at 12.5 GS/s/ch, and from the PXIe frontend, we can sustain data grabbing at approximately 600 MB/s. Processing and analyzing 6 GB raw data is still challenging on modern PCs. Hence, shorter time chunks are preferable, e.g., 1 s.

In contrast, continuous monitoring of environmental signals or device soundness often adopt slow sampling at 1 S/s–1 kS/s, which produces numerous short chunks if the 10-s rule is applied. For easy assessment of daylong trends in graphs we need to adopt longer time chucks, such as 10 min or 1 day.

To satisfy these contradictory requirements, we modified the DAQ system to cope with variable chunk sizes. Currently, the system is already implemented for 560 channels of the device monitoring control data processing system (CDP) [7] and 16 channels of the environmental radiation monitoring Radiation Monitoring System Applicable to Fusion Experiments (RMSAFE) [8]. Both systems adopt 10 min chunks and have been running since the past two years.

Figure 3 shows a schematic view of sparse resampling of daylong data used in device soundness monitoring in the "CDP" case. The daylong "CDPslowL" is usually used for browsing long-term trends, while the original CDP is used for analyzing fine signals. Table 1 shows the details of slow monitoring in LHD.

4. Uninterrupted Data Storing
4.1 High-availability data store

In addition to scalable I/O performance, large data archives must provide data redundancy and high availability. LHD's primary storage previously consisted of a cluster of hard drive arrays on the 4-Gbps FibreChannel storage area network (FC-SAN). However, this system sometimes suffered from a complete deadlock due to the strong interaction between nodes because trouble at one node often affected all other nodes. The typical throughputs of the system did not exceed 150 MB/s.

To obtain higher availability and upgraded I/O performance, we adopted a new distributed "Key-Value Store (KVS)" data archiving system based on the popular "cloud" storage technology and 10-Gbps Ethernet network attached storage [1]. This system can provide many beneficial characteristics such as automatic file replication, dynamic load balancing, on-demand scalable I/O speed, and capacity reinforcement.

In addition, cloud storage is also equipped with the function of hot plug-on/off-and-play. In other words, it has high fault tolerance due to hardware redundancy and has no single point of failure (SPOF). Once continuous monitoring starts, it is not posssible to stop the data storage service for any maintenance or upgrade. Although the FC-SAN cluster storage has a hot plug-and-play ability, it requires a minimal switchover time when service is stopped for maintenance.

Cloud storage's ability of hot plugging, therefore, helps us maintain or modify the storage system during runtime. Indeed, in the 2011 experimental campaign of LHD, we added a new storage node at the peak of the experiment sequences without any momentary stop of service.

4.2 High-reliability common data platform

Since 2008, the LABCOM data system has been used as a common data store and remote access platform not only for the LHD experiment but also for the QUEST project of Kyushu University and the GAMMA10 project of the University of Tsukuba. LABCOM enables bidirectional remote acquisition and re-distribution through the fusion dedicated virtual private network *"SNET"* [9]. This project is named the *"Fusion Virtual Laboratory (FVL)"* [10] in Japan and currently covers the three abovementioned experiments.

The FVL framework is advantageous for remote experiment applications. However, the storage system must be operational during all three different campaign periods, and it has less time slots in which services can be interrupted for maintenance. For multiple site collaborations such as in FVL, the hot plug-and-play function of the cloud KVS is essential for enabling zero-downtime maintenance or upgrade without interrupting the system's operation.

High-performance computing systems often use so-called parallel filesystems or cluster filesystems [11]. LHD and FVL have verified the use of cloud storage in data
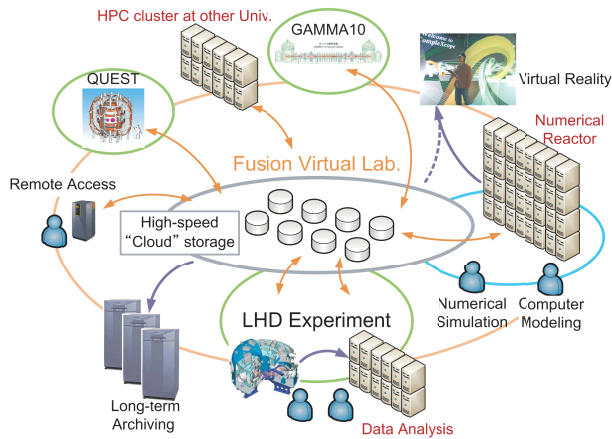
Fig. 4 New FVL concept used in the unified data platform for both numerical computation and experimental analysis

archiving. Furthermore, cloud storage can be used for long-term data storage of numerical modeling or simulations. In other words, it could be a common data platform for both experimental and computational research. If all computational data can be freely accessed in the same manner as the experiment data, it could usher a new era in fusion research.

The concept under the FVL framework, therefore, is to construct a composite of cluster computer and experiment sites, as shown in Fig. 4. A similar idea is also discussed for the ITER project.

5. Conclusions

In this study, variable length of "subshot" time chunks have been implemented and verified with different practical uses in LHD. We have shown that variable subshot lengths are essential and effective for ultra-fast sampling and slow endless data monitoring. The former uses shorter variable intervals – less than 10 s. Currently, we adopt a sampling rate of 1 s or less for the 12.5 GS/s digitizer. The latter adopts dual subshot lengths of 10 min or 1 day. For slow endless data monitoring, we use 1-Hz sparse data resampled from the original 1-kHz sampled data.

Uninterrupted data storage was demonstrated using new distributed KVS based on the so-called "cloud" technology. This successful prototype system demonstrates the large-scale use of distributed and diverse DAQs for ITER and next-generation fusion experiments.

[1] H. Nakanishi, M. Ohsuna, M. Kojima, S. Imazu, M. Nonomura, T. Yamamoto, M. Emoto, M. Yoshida, C. Iwata, M. Shoji, Y. Nagayama, K. Kawahata, M. Hasegawa, A. Higashijima, K. Nakamura, Y. Ono, M. Yoshikawa and S. Urushidani, Nucl. Fusion **51**, 113014 (2011).
[2] H. Nakanishi, M. Ohsuna, M. Kojima, S. Imazu, M. Nonomura, M. Hasegawa, K. Nakamura, A. Higashijima, M. Yoshikawa, M. Emoto, T. Yamamoto, Y. Nagayama, K. Kawahata and LHD Experiment Group, Fusion Sci. Technol. **58**, 445 (2010).
[3] H. Nakanishi, M. Kojima and LABCOM group, Fusion Eng. Des. **56**-**57**, 1011 (2001).
[4] H. Nakanishi, M. Ohsuna, M. Kojima, S. Imazu, M. Nonomura, M. Emoto, H. Okumura, Y. Nagayama, K. Kawahata and LHD exp. group, J. Plasma Fusion Res. **82**, 171 (2006) [in Japanese].
[5] M. Ohsuna, H. Nakanishi, S. Imazu, M. Kojima, M. Nonomura, M. Emoto, Y. Nagayama and H. Okumura, Fusion Eng. Des. **81**, 1753 (2006).
[6] H. Nakanishi, M. Ohsuna, M. Kojima, S. Imazu, M. Nonomura, H. Okumura, Y. Nagayama and K. Kawahata, Plasma Fusion Res. **2**, S1117 (2007).
[7] S. Yamaguchi, M. Shoji, M. Emoto, S. Sudo, J. Kariya, H. Okumura, Y. Teramachi, M. Tamura and LHD Data Processing Group, Fusion Eng. Des. **48**, 9 (2000).
[8] H. Yamanishi, H. Miyake, J. Kodaira, H. Obayashi, M. Isobe and K. Matsuoka, Nucl. Instrum. Methods A **501**, 589 (2003).
[9] K. Tsuda, Y. Nagayama, T. Yamamoto, R. Horiuchi, S. Ishiguro and S. Takami, Fusion Eng. Des. **83**, 471 (2008).
[10] T. Yamamoto, Y. Nagayama, H. Nakanishi, S. Ishiguro, S. Takami, K. Tsuda and S. Okamura, Fusion Eng. Des. **85**, 637 (2010).
[11] K computer, http://en.wikipedia.org/wiki/K_computer (2011).