

Real-Time Signal Communication between Diagnostic and Control in ASDEX Upgrade

Wolfgang Treutterer, Gregor Neu, Gerhard Raupp, Thomas Zehetbauer, Dieter Zasche, Klaus Lüddecke†, Richard Cole† and ASDEX Upgrade Team

Max-Planck Institut für Plasmaphysik, Garching, Germany, EURATOM Association

† Unlimited Computer Systems, Iffeldorf, Germany

Abstract

The ASDEX Upgrade tokamak experiment is equipped with a versatile discharge monitoring and control system. It allows to develop and use advanced control algorithms to investigate plasma physics under well-defined conditions with the objective of optimising plasma performance. The achievable quality depends on the accuracy with which the plasma state can be reconstructed from measurements under real-time conditions.

Today's advanced algorithms need physics quantities – scalar entities as well as profiles. These are obtained processing huge numbers of raw measurements with complex diagnostic algorithms. Adequate network communication for the resulting signals is crucial to satisfy real-time requirements, especially when several diagnostic systems cooperate in a feedback control loop. Support for the technology of choice, however, is not easily available for all of the diverse, highly specialised diagnostic systems.

We give an overview about the methods that have been explored at ASDEX Upgrade for real-time signal transfer. In particular, we investigated reflective shared memory and Ethernet technologies. Our solution strives to combine their strengths. For fast communication on dedicated computing nodes, reflective shared memory is used. For the majority of diagnostic systems producing large data blocks at moderate rates, Ethernet connections with UDP protocol are employed. Following ASDEX Upgrade's framework concept, a software layer hides the networks used from both diagnostic and control applications.

1. Introduction

Big efforts are being undertaken to reveal the physics background of magnetohydrodynamic plasma behaviour in thermonuclear fusion. For the upcoming joint experiment ITER as well as a later fusion reactor it is essential to exploit this knowledge for designing advanced operation regimes. Such regimes depend on many factors and can be established and sustained only with feedback control systems. An ever increasing multitude of diagnostic measurements and associated evaluation algorithms is required to compile the complex information about these factors into a plasma state reconstruction and to generate adequate control commands [1]. All big fusion plasma devices therefore have started to enable diagnostic systems to feed reconstructed physics quantities to their plasma controllers in real-time [2]. They have in common that massive parallel data acquisition requires the use of multiple computer hosts and real-time capable networks to stream data to evaluation and control nodes.

Moreover, on the RFX tokamak experiment combinations of Ethernet protocols with VxWorks and real-time Linux Systems have been studied [3]. ASDEX Upgrade has been using reflective shared memory for years. Driven by the design of a feedback loop for MHD mode stabilisation [4] [5] (see figure 1) we started to investigate whether also Ethernet technologies are suitable for this purpose. Real-time diagnostic systems collect raw measurements with sampling periods up to 2 MHz and condense these to a number of physics quantities with a rate of 100 Hz. Higher level evaluation processes combine inputs from several diagnostic systems to values appropriate for the control task. These are then shared with the control system, which reacts with an overall latency not worse than 10 ms to the raw measurements. The real-time data network has to provide bidirectional streaming data exchange between diagnostic systems, evaluation processes and control system. Due to limited resources, work concentrated on a qualitative analysis of applicability. So far, no detailed statistical analysis of speed and latency has been performed. Nevertheless, the experience gathered was sufficient for making decisions on the real-time network technology at ASDEX Upgrade.

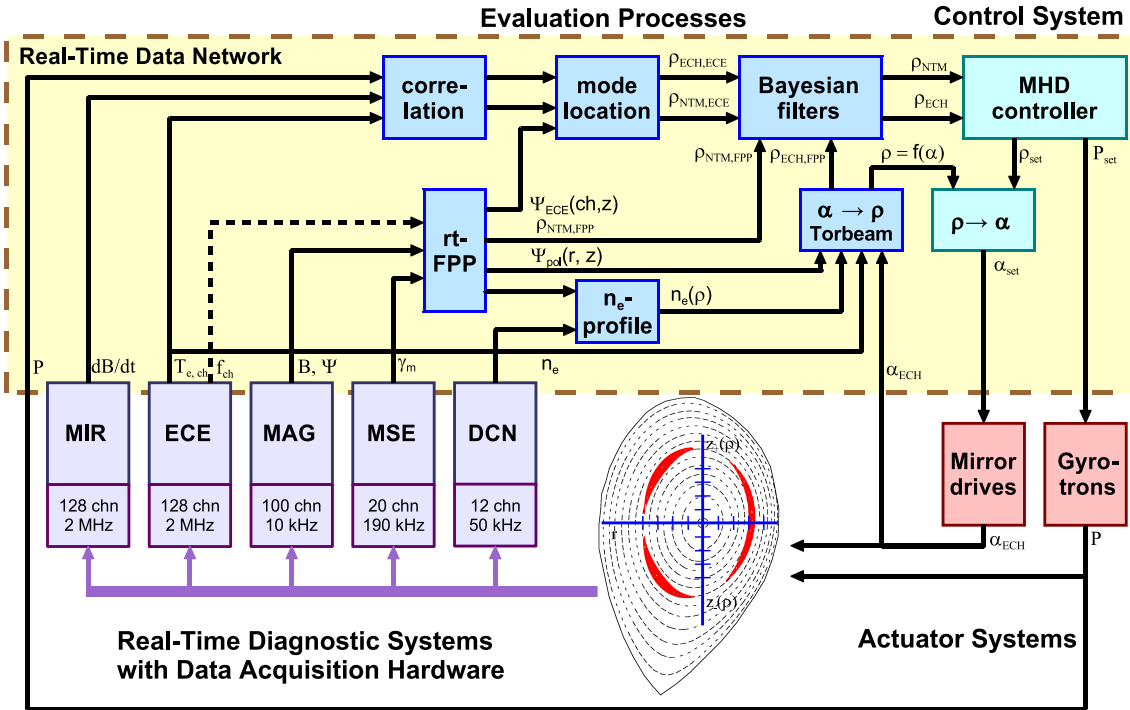


Fig. 1: Diagnostic and control integration for MHD feedback control

Facing the requirement of fast and reliable data communication, Ethernet based technology seems to be an obvious solution. Gigabit Ethernet promises satisfactory speed, commodity hardware is cheap and easily available for almost any computer platform and TCP/IP or UDP are proven transmission protocols with built-in support in most operating systems. Section 2 deals with the experiences and capabilities of this approach. The outstanding performance, however, deteriorates significantly for small but frequent data packets with low communication latency, which are common requirements for real-time plasma control. Reflective shared memory technology is a resort which is examined in section 3, but it has deficiencies in the support on different hardware platforms and operating environments.

Our solution strives to combine the strengths, selecting the most adequate technology for each use case, but hiding the networking details from diagnostic and control applications using a generic software layer. The concept of this layer is described in section 4.

Finally, section 5 will summarize the experience gathered and give a preview of future development.

2. Ethernet based real-time communication

Ethernet technology connects computer nodes in a standardised, fast and efficient way. It is a mature,

well-proven technology. The software layer is often a core part of the operating system (UNIX). A big variety of hardware is available, and due to the large quantities produced, costs are low. Ethernet technology is installed almost everywhere, forming the backbone of local and world-wide networks.

For this reason Ethernet is a good candidate real-time network for ASDEX Upgrade's Discharge Control System (DCS).

Several network protocols are built on top of Ethernet hardware. Among those TCP (Transmission Control Protocol) is the most popular one. It features easy-to-use and reliable data transfer with balanced network load [6]. However, several self-adjusting methods, like packet retransmission and congestion control, particularly Nagle's algorithm [7], lead to intrinsic non-deterministic communication latencies. Although this is contrary to basic real-time requirements, there was the expectation that in a switched 100 MBit/s network, the data buffering performed by the switches would avoid collisions and retransmissions, thus reducing latency variations to acceptable values below 5% of the DCS control cycle (50 μ s).

Disabling Nagle's algorithm would force to send all sample data sets, also the small ones, as individual entities. Still, in our installation, sporadic outages of 20 ms and longer occurred in the VxWorks network stack, which is not tolerable for discharge control. As the underlying reasons could not be revealed, using TCP protocol for real-time data communication was finally abandoned.

As an alternative, the User Datagram Protocol (UDP) was tested. UDP is preferred for streaming data applications. It works with low administrative overhead and thus offers shorter and more deterministic latencies. On the other hand, it uses neither handshake mechanisms nor retransmission in case of packet loss. Hence there is no guarantee that every packet is delivered to the recipient [8]. In the context of a switched Ethernet environment, this deficiency seems tolerable, though, since the network switches can buffer packets and thus avoid packet loss. Anyway, control applications have to define a strategy for the case when a network connection breaks unexpectedly and the input data stream is interrupted. Another issue with UDP can be the handling of datagrams larger than the Maximum Transmission Unit (MTU) length for Ethernet (1476 bytes). These are broken down into a series of smaller packets. Unlike TCP, UDP does not provide an intrinsic mechanism for reassembling the pieces at the recipient side. At ASDEX Upgrade, however, this was no real problem, since typical real-time signal samples need just 16 bytes for value, timestamp and state tag and can be handled in a simple way [9].

Our experience with UDP has been very satisfactory so far. We connected 5 real-time diagnostic systems at sampling periods between 2 ms and 40 ms, and their number is expected to grow. Typical latencies for network transfer amount to 100 μ s, including the overhead from the software framework layer described in section 4, but can hit peaks of 200 μ s. Loss of packets was not observed. We ascribe this to the use of a switched network. Some effects, however, deserve special attention. First, diagnostic systems might send data at an even higher rate, eventually causing an overloading of the receiving systems, if they cannot keep pace in reading, internally distributing and processing these samples. Thus the control system's real-time capability can get lost. When many diagnostic systems at high data rates are connected to a receiver, the available network bandwidth can be exceeded. The consequences in such a case depend on the setup and the capabilities of the networking hardware like network switches and network cards.

Other options, which have not yet been investigated, are the use of alternative dedicated Ethernet protocols like Ethernet Powerlink [10], RTnet [11], communication middleware software like DDS [12] or TIPC [13] or Infiniband technology [14]. Instead, we focussed the evaluation on reflective shared memory products.

3. Reflective shared memory based real-time communication

Reflective shared memory is a technology where several cooperating computer nodes communicate via mirrored copies of a memory segment. Whenever the data content of this segment is modified on one node, the change is propagated to the other nodes. This operation is implemented in hardware: each node is equipped with an adapter card, and these cards are interconnected with fibre optical links. Seen from the user application side, communications are as simple as memory read and write operations. Network transfer rates are in the range of 40 MByte/s and better, depending on the product and computer bus adaptor. Two common network topologies exist. In star architectures all nodes are connected to a central hub or switch which broadcasts input data from one node simultaneously to the other nodes. In a ring architecture the nodes are connected directly to their neighbouring nodes, and data are shifted around from node to node. Thus, data are not available on all nodes at the same time. But in our setup with just a few nodes this property had no impact on performance. Important for synchronised real-time control applications is the ability to notify recipients when data cells have changed. Reflective shared memory technology is tailor-made for real-time applications. The hardware ensures fast and reliable data transfer

with guaranteed latency. Small and big data segments can be efficiently handled..

ASDEX Upgrade's discharge control system is operating a reflective shared memory as its communication backbone between six nodes since 2003. The SBS Broadcast Memory [15] features a star network topology and memory mapped interrupt notification. Our experience has been very good, but this product has unfortunately been discontinued.

For the connection to upcoming real-time diagnostic systems we therefore envisaged a current product. The GE Fanuc 5565 board [16] is already in use in the JT-60 fusion experiment [17]. It comes with a ring network topology. Notification of recipients is possible by remote interrupt, meaning that the data modifier node must inform all data user nodes. The system has been tested so far with results comparable to the SBS Broadcast Memory. The main application of GE Fanuc 5565 will be the real-time transfer of large arrays of plasma and diagnostic quantities. An example is a 39x69 magnetic flux matrix, which is exchanged between higher level plasma diagnostic systems. We obtained a latency of 70 μ s for a matrix sample and 5 μ s for scalar samples, including software layer overhead.

Another interesting product of this family, which could be examined in the future, is the Dolphin DXH510 switch based reflective memory [18]. It uses cached system memory instead of local on-board memory buffers, thus reducing the number of copying operations. Network transfer rates of 1.3 Gbyte/s are advertised.

Up to now reflective shared memory is, however, a niche technology with a comparatively small number of users, especially in the research and military areas. Product diversity, innovation cycle and platform availability are not comparable to the Ethernet product line. Often, platform and use case specific custom drivers have to be procured by the user.

4. Integrated real-time network

Finally, it was decided that the ASDEX Upgrade discharge control system DCS should support both technologies, Ethernet with UDP as well as reflective shared memory, thus combining their strengths. Control applications which are latency critical, operate at sampling periods shorter than 2 ms, or exchange large data blocks, communicate via reflective shared memory. The others work with UDP protocol over Ethernet. Additionally, data must be relayed between those networks. The modular framework structure

of DCS makes the various networks and relays entirely transparent to user applications [19]. It presents a uniform interface to user algorithms encapsulating the diversity of network technologies, hardware platforms and operating systems. This facilitates future upgrades like the migration to a more powerful network product without modifications to the application code.

In control systems real-time data streams convey signal samples from the output of an algorithmic block to the inputs of connected blocks. The framework will arrange for all routing, conveying and notification services across all networks for input and output signals with a uniform interface to the applications. In order to do this, a registration service collects publication and subscription requests containing signal names, data types and sampling frequencies. It checks for consistency and computes the necessary network setup for the subsequent real-time phase.

Figure 2 shows the implementation. On each node a gate process controls access to the attached network. The exchanged signal samples are retained in an area of local memory, the signal buffer. The application interface to the local store, however, is network independent. The framework's publish and subscribe service, in turn, computes routing tables and configures the gate processes with network specific parameters. All actions are coordinated by a state machine (Run Control). A real-time signal server completes the picture. Attached to several networks, it performs the routing between those networks via the corresponding gate and hosts evaluation algorithms combining input from the various diagnostic sources.

Owing to its versatile framework concept, DCS comprises a SBS Broadcast Memory network with 6 control application nodes, a GE Fanuc 5565 Reflective Shared Memory with two diagnostic application nodes, an UDP via Ethernet network with 8 diagnostic system participants and one real-time signal server. The nodes are based on Intel x86 and Sun Sparc hardware and running in VxWorks, Sun Solaris, rtLinux and rtLabView operating environments.

5. Summary

Real-time capable networks form the glue between data acquisition, data processing and control nodes in ASDEX Upgrade's discharge control system. The evaluation of several network technologies has finally led to a combined solution. Reflective shared memory satisfies all real-time communication requirements,

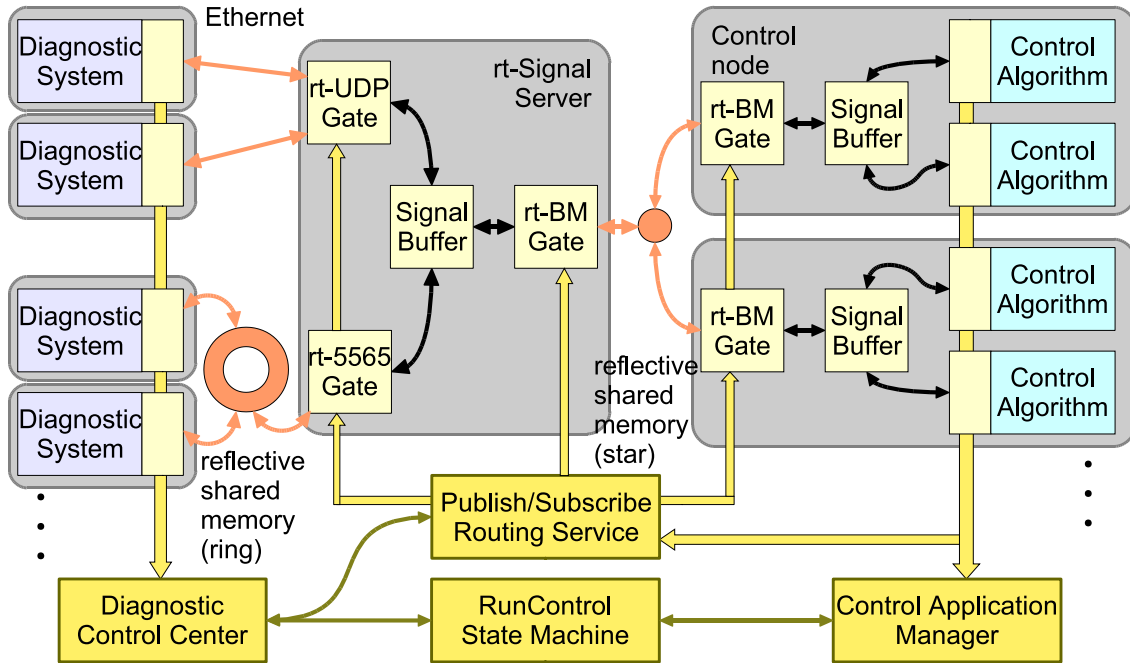


Fig. 2: Real-time communication scheme with multiple networks

but is expensive to port to a big variety of computer platforms and operating systems. Ethernet technology with UDP is an adequate alternative for many diagnostic applications with relaxed time scales of 2 ms and above. TCP for real-time usage, on the other hand, failed due to intolerable non-deterministic latency variations.

A software framework layer encapsulating network operation from control and data acquisition algorithms is capable of integrating a variety of networks, computer platforms and operating environments in a uniform and transparent way. In the coming years many of ASDEX Upgrade's diagnostic systems will be connected to the real-time network. Thus the expertise of diagnostic algorithms can be made available to plasma control, allowing for enhanced plasma scenarios, even better stability and performance optimisation on the way to ITER.

References:

- [1] W. Treutterer, et al., Real-time diagnostic integration with the ASDEX upgrade control system, Fusion Eng. Des. (2008), doi:10.1016/j.fusengdes.2008.12.026
- [2] K. Kurihara, et. al, Plasma control systems relevant to ITER and fusion power plants, Fusion Eng. Des. (2008), doi:10.1016/j.fusengdes.2008.06.027
- [3] A. Lucchetta, et. al, Real-time communication for distributed plasma control systems, Fusion Eng. Des., Vol. 83 2-3 (2008), pp. 520-524, doi:10.1016/j.fusengdes.2007.12.026
- [4] A. Manini, et. al., Development of a feedback system to control MHD instabilities in ASDEX Upgrade, 24th SOFT (2006), Warsaw, to be published in Fus. Eng. and Design
- [5] W. Treutterer, et. al., Real-time diagnostics at ASDEX Upgrade—Integration with MHD feedback control, Fusion Eng. Des., Vol. 83 2-3 (2008), pp. 300-3-3, [doi:10.1016/j.fusengdes.2007.09.006](https://doi.org/10.1016/j.fusengdes.2007.09.006)
- [6] DARPA Internet Program, Transmission Control Protocol, Internet Engineering Task Force (1981), <http://tools.ietf.org/html/rfc793>
- [7] R. Braden, Requirements for Internet Hosts -- Communication Layers, Internet Engineering Task Force, Network Working Group, RFC 1122, section 4.3.2.4 (1989), <http://tools.ietf.org/html/rfc896>
- [8]] J. Postel, User Datagram Protocol, Internet Engineering Task Force (1980), <http://tools.ietf.org/html/rfc768>
- [9] W. Treutterer, et. al., Software structure and real-time signal exchange for the ASDEX Upgrade tokamak control system, Proceedings of the International Conference on Accelerator and Large Experimental Physics Control Systems Gyeon-Gju, South Korea (2003), pp. 548–550.
- [10] Ethernet POWERLINK Standardization Group (EPSPG), The real-time Ethernet protocol POWERLINK, <http://www.ethernet-powerlink.org>
- [11] RTnet homepage, Hard Real-Time Networking for Real-Time Linux, <http://www.rtnet.org>
- [12] Real Time Innovations, Distributed Data Service (DDS), <http://www.rti.com>
- [13] Wikipedia, Transparent Inter-process Communication (TIPC), <http://en.wikipedia.org/wiki/TIPC>

- [14] Wikipedia, InfiniBand, <http://en.wikipedia.org/wiki/Infiniband>
- [15] GE Fanuc, SBS Broadcast Memory, <http://www.sbs.com/products/family/112/>
- [16] GE Fanuc, Reflective Memory, <http://www.gefanuc.com/products/family/reflective-memory>
- [17] K. Kurihara, et. al., Status and prospect of JT-60 plasma control and diagnostic data processing systems for advanced operation scenarios, Fusion Eng. and Design, Vol. 81 (2006), 1729 – 1734, doi:10.1016/j.fusengdes.2006.04.058
- [18] Dolphin Interconnect Solutions, Dolphin Express Reflective Memory, <http://www.dolphinics.com/products/embedded-system-reflective-memory.html>
- [19] G. Raupp, et. al., Control processes and machine protection on ASDEX Upgrade, Fusion Engineering and Design, Vol. 82 5-14 (2007), pp. 1102-1110, doi:10.1016/j.fusengdes.2007.07.044