

UKAEA-CCFE-CP(19)56

Stephanie Hall, Phillippe Jacquet, Graham Naylor

The MAST-U Real-Time Protection System

This document is intended for publication in the open literature. It is made available on the understanding that it may not be further circulated and extracts or references may not be published prior to publication of the original when applicable, or without the consent of the UKAEA Publications Officer, Culham Science Centre, Building K1/0/83, Abingdon, Oxfordshire, OX14 3DB, UK.

Enquiries about copyright and reproduction should in the first instance be addressed to the UKAEA Publications Officer, Culham Science Centre, Building K1/0/83 Abingdon, Oxfordshire, OX14 3DB, UK. The United Kingdom Atomic Energy Authority is the copyright holder.

The contents of this document and all other UKAEA Preprints, Reports and Conference Papers are available to view online free at <https://scientific-publications.ukaea.uk/>

The MAST-U Real-Time Protection System

Stephanie Hall, Phillippe Jacquet, Graham Naylor

The Real-Time Protection System for MAST-U

Stephanie Hall^a, Phillippe Jacquet^a, Graham Naylor^a

^a United Kingdom Atomic Energy Authority, Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, Oxon OX143DB UK

The Mega Amp Spherical Tokamak (MAST) is currently being extensively upgraded to provide increased coil system and power supply capability to the extent certain combinations of currents could cause catastrophic damage. To ensure operation of MAST-U within safe engineering limits, the machine protection system was also updated. The protection system is needed to protect against coil faults (such as flashovers) and monitor and calculate, in real-time, quantities such as coil currents, axial coil forces and the stored energy within the coils. The Real-Time Protection System was developed and uses Field Programmable Gate Arrays (FPGAs) which monitors up to 89 different signals collected from 7 different locations around MAST-U at a rate of 0.5MS/s. It filters and calibrates the signals and performs calculations to check if these signals correspond to valid operating ranges for currents, stored energy and coil forces. If thresholds are exceeded, then several actions can be carried out including stops to specific power supplies or a more gradual stop from the plasma control system. Initial testing of the system has shown it meets this challenging specification and the system is able to react within 200 μ s to breaches in thresholds as well as collecting 400MB/s of calibrated and calculated data.

Keywords: MAST-U, FPGA, protection, real-time

1. Introduction

The Mega Amp Spherical Tokamak (MAST) (see Fig.1) is currently being extensively upgraded to provide a system that will be able to add to the knowledge base for ITER as well as testing innovative reactor systems. MAST-U will have [1]:

- Increased toroidal field capabilities – improving confinement
- A new solenoid - increasing plasma current and pulse duration
- 19 new poloidal field coils - improving shaping
- A super-X divertor [2] – testing innovative divertor concepts and improve power handling
- Off axis NBI – modifying q and fast ion pressure profiles

To ensure operation of MAST-U within safe engineering limits, the machine protection system needs to be updated. The protection system needs to protect against coil faults and dangerous coil current combinations. To do this it monitors and calculates, in real-time, quantities such as coil currents, axial coil forces, coil heating and the stored energy within the coils. The aim is to be able to react within hundreds of microseconds if there are any inconsistencies that breach pre-defined thresholds.

The following protection criteria need to be monitored:

- Currents on toroidal, poloidal, and divertor coils are below pre-defined limits
- Currents on the coils and currents from the power supplies are consistent within a pre-defined margin
- Toroidal Field voltage measurement across sections of the coil are self-consistent within a

pre-defined margin (in order to identify any shorted-turn path fault conditions)

- Current measurements on the central solenoid are self-consistent within a pre-defined margin
- The energy stored in the coils and coil heating (I²t calculation) is kept below set pre-defined thresholds.
- The force exerted on all coils by other coils and the plasma is below pre-defined limits.
- Langmuir probes are used to ensure that plasma does not approach a set point in the divertor region (in order to protect fragile windows).

Once these protection criteria have been compared to the pre-defined thresholds then any breaches are mapped to actions. The actions that can be carried out are:

- controlled stops to the plasma control system which can bring the pulse down slowly
- stops directly to powers supplies
- shifts in the position of the plasma by using the vertical control system.

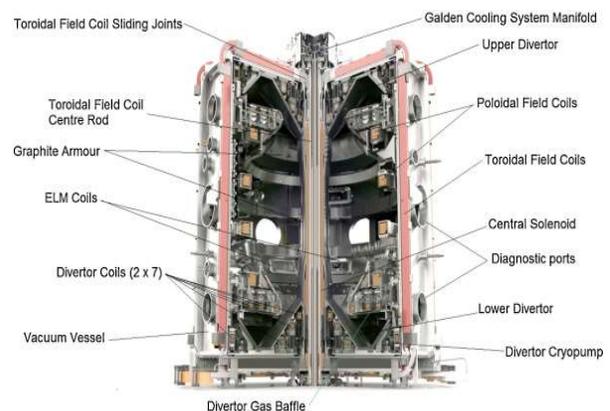


Fig. 1. Annotated cross section of the MAST-U machine [3]

2. Detailed Requirements

The Real Time Protection (RTP) System must be able to react to the protection criteria within 200 μ s (mainly determined by the signal filtering latency) of thresholds being exceeded. To achieve this, it needs to measure up to 89 signals from 7 different areas of plant, at a rate of 500kHz for each signal, meaning that the system would have to handle a throughput of just under 200MB/s. The system will need to filter, calibrate and perform calculations on signals in real-time. The system also needs to be highly configurable to allow changes to filter, calibration and calculation coefficients, calculation thresholds and mappings of events to plant. Changes to the configuration should be possible on a per pulse basis.

To achieve these challenging requirements, it was decided to use a network of interconnected FPGAs, controlled with a central FPGA and ARM processor control unit. FPGAs were chosen, as determinism and low latency are inherent to this technology. CCFE has built up a significant experience and expertise in FPGA development [3]. The FPGA can handle the real-time aspects of the system while the ARM processor is able to deal with the communication with other plant systems on a slower time-scale.

The system architecture is shown in Fig.2. The FPGA based acquisition units: ADCs, integrator and protection Langmuir probe units, located around MAST-U, collect signals and send them to a central FPGA on digital optical links using the aurora protocol [5]. This central FPGA performs the real-time processing on the data and sends requests to the plasma position controller or the FPGA based digital output modules which send optical outputs

to the Plasma Control System (PCS) and to power supplies. Data is also streamed to an archive via 2 FPGA boards in a data collection PC. The ARM processor handles the communications to:

- The Machine Control System which sequences RTP states throughout a MAST-U pulse
- The parameters database which stores all the thresholds and coefficients required by RTP
- The data acquisition management system which records any actions that occur during a pulse.

Note that a hardwired system also protects some critical components of MAST-U (TF coils, Solenoid), enhancing the overall protection system reliability.

3. Protection Design

The ADC units can each collect up to 32 channels and consist of ADC PCBs produced by CCFE and connected via a backplane to an FPGA board: the Simple PCIe FMC Carrier (SPEC) board [4] (see Fig.3) The SPEC board sends control signals to the ADCs, decodes the ADC data and multiplexes this data over a fiber link to the central FPGA unit using the Xilinx aurora protocol [5] through an SFP port. There are 6 of these ADC units and an ADC coupled with integrator units at the front end. The data from the ADCs and integrators is decoded in the central FPGA processing unit.

The digital input unit includes optical (ON/OFF) receivers used to trigger and synchronize the RTP system, input the Langmuir probes faults status in the RTP system, and input the status of various parts of the plants involved in MAST-U protection (hardwired protection tripped status for example).

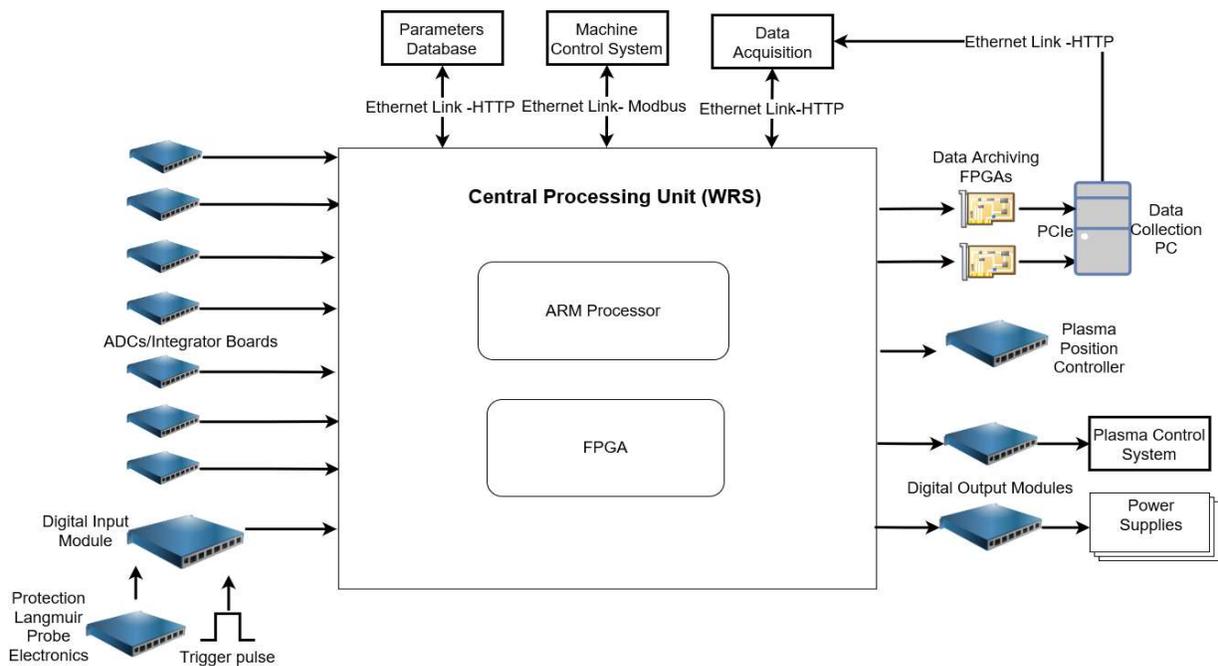


Fig. 2. The real-time protection system layout.



Fig. 3. SPEC board and White Rabbit Switch Hardware

The central processing unit hardware is the White Rabbit Switch (WRS [4]) (see Fig.3), which consists of a Virtex-6 FPGA connected to an ARM processor. The WRS was originally designed as an ethernet switch, however in this application, the FPGA is completely re-programmed to perform the MAST-U RTP specific function. The real-time aspects of the system are streamlined on the FPGA and the updating of remotely configurable coefficients, thresholds and mappings are handled by the processor as they can occur on a much slower timescale. The filtering and calibration occur on the decoded ADC and integrator data. The output of this is streamed to the calculation modules for additional real-time processing. The calculation modules perform the protection criteria mentioned previously.

The actions (mentioned previously), that can occur when the calculations breach pre-defined limits, are directed to outputs of the central FPGA. The communication to PCS and power supplies are done by connecting to digital output modules using the Xilinx aurora protocol. The digital output modules contain a SPEC board like the ADCs and integrators but are used to decode the aurora stream and send out digital outputs to on/off optical transmitters. Some of these optical transmitters connect to different power supplies and others encode different stop types to PCS. In addition to these actions the central FPGA can also communicate to the plasma position controller directly using the Xilinx aurora protocol and can request a kick of the plasma for a set period.

The central FPGA (the WRS hardware) embeds the real-time logic of the system, whereas the ARM processor is there to facilitate communications to the machine control system to enable the system and perform data acquisition to collect events and actions after the pulse. The FPGA also requires access to many coefficients for filters, calibration, offsets and calculations as well as thresholds and event/action mappings which are stored in a database. This is also handled by the ARM processor.

To facilitate post processing of events and actions it is necessary to collect both the calibrated and calculated data that is acted on during the pulse. The throughput of calibrated and calculated data is 400MB/s and will need to be collected over a 7s period. This data is not able to be stored locally as there is not enough space on the central FPGA unit, therefore this data is streamed over two

optical links to two AC701 FPGAs [6] using the aurora protocol. Once the data is on the AC701 FPGAs it can be streamed over PCIe to a data collection PC using the Xillybus FPGA IP which allows direct memory access to the PC using a linux driver [7].

3. Testing

To ensure the high level of reliability needed for the real-time protection system, a through testing strategy was developed. This testing strategy also needed to be able to be as streamlined as possible to reduce delays in functional testing and commissioning. The strategy developed involves application of virtual signals for testing that enables unit tests to be run on the FPGA but controlled from the RTP data collection PC. By running a script on this PC, testing instructions are sent across the duplex link and data is streamed back to the PC and compared to previously validated data. In addition, manual testing will also be carried out using an of the shelf 32 channel DAC unit that can be controlled remotely. This will produce predefined waveforms as inputs to mimic waveforms seen during a normal pulse. In implementing the virtual signal tests, external signal tests and automating them both, testing can be completed quickly. This facilitate any changes needed during testing as well as ensuring a high level of reliability, as the output can be compared to predefined waveforms on multiple test runs with ease.

4. Results

The setup in Fig.4 was used to test the latency given by a loopback of a signal. A signal from a generator (ON/OFF) is input into the optical receiver modules and looped back via the WRS (with no filtering or calculations), to the optical transmitter modules, the latency is measured on an oscilloscope. The transfer and decoding of the data using the aurora protocol without any processing is $2\mu\text{s}$ with this setup.

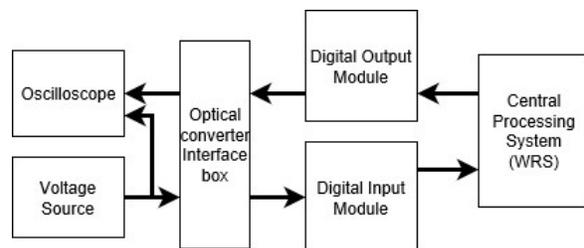


Fig. 4. loopback experiment setup

Filtering (see Fig.5) has also been tested by applying the filter with coefficients shown in Fig.5. This was calculated by convoluting a low pass filter and a notch filter at 125kHz to attenuate a systematic noise contribution from the ADCs. The noise was reduced from 150 to 7 bits as shown from the comparison of graphs in Fig.5. The 64-coefficient symmetric filter setup in the WRS will add an additional $60\mu\text{s}$ to the $2\mu\text{s}$ measured on the loopback test. The calibrations, calculations and mapping add a further $\sim 2\mu\text{s}$. These latencies match the overall latency requirements needed for MAST-U.

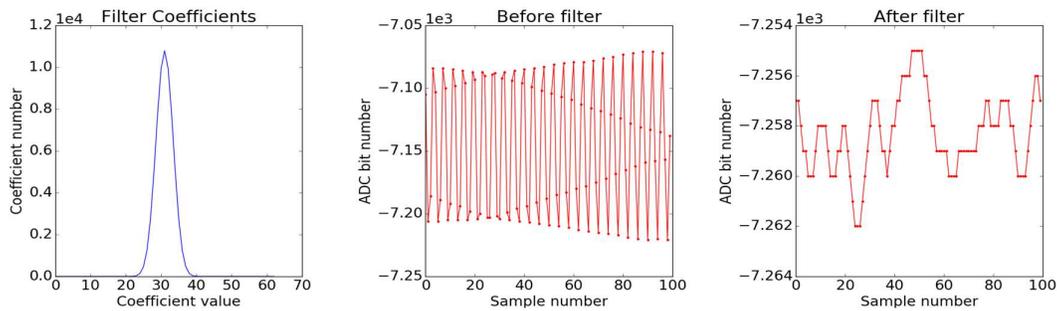


Fig. 5. The filter coefficients applied in real-time to a 125kHz modulated signal which attenuated the signal

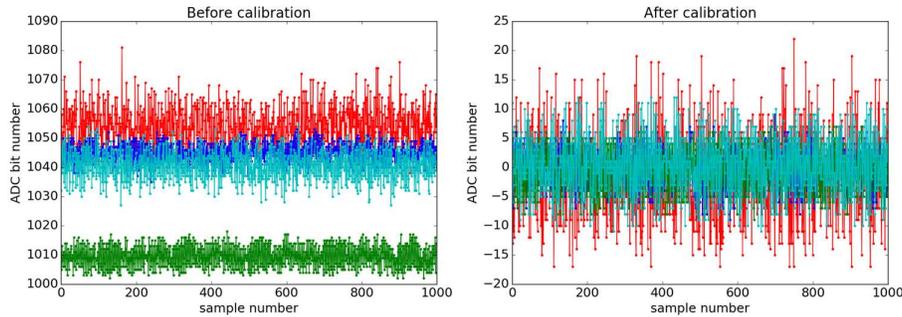


Fig. 6. The collected data calibrated in real-time to zero

In addition, the analog signal calibration carried-out in RTP allows precise correction for the ADC offsets as shown when comparing graphs in Fig.6. Initial testing of the calculation modules has also been carried out and shown that the functionality of these modules works as expected with minimal latency in comparison to the necessary latency added by the filter module. Furthermore, the system has shown reliable streaming of 400MB/s from the central FPGA to data storage PC without any loss of data and has routinely been used during testing. The WRS with all ADCs/integrators has also been tested and has been able to collect all signals simultaneously from all the different ADC/integrator modules for data processing in the FPGA.

The status of the system is nearing the end of the virtual signal testing stage, interface and testing involving external input of signals is expected to start in September 2018 ready for MAST-U power supplies off-coil commissioning later this year.

5. Future Plans

Long term goals for exploitation of this system would include integration of the system with wider control networks. There are already plans for development of a wider real-time network, which could be employed to perform divertor detachment control on fusion devices. This system would use the same technologies and firmware as the real-time protection system and would integrate with it. This would allow real-time control of the detachment front to ensure that divertor heat loads can be kept within reasonable operating limits

5. Conclusion

To ensure MAST-U operates within a safe operating range a protection system is needed that is able to react in

real-time: within hundreds of microseconds of the machine moving towards a non-safe operating range. The Real-Time Protection system was developed, which was able to meet the challenging criteria set. Current results have met the specification and the system is currently undergoing a period of testing to validate the system ready for off coil commissioning.

Acknowledgments

The authors would like to thank the MAST-U team who have been integral in the development and integration of the RTP system. This work has been funded by the RCUK Energy Programme [grant number EP/I501045]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ccfe.ac.uk

References

- [1] A. W. Morris, MAST: Results and Upgrade Activities, IEEE Transactions on Plasma Science, 40 (3) (March 2012)
- [2] E. Havlíčková et al, Investigation of conventional and Super-X divertor configurations of MAST Upgrade using scrape-off layer plasma simulation, Plasma Physics and Controlled Fusion, 56 (7) (2014)
- [3] J. Milnes, MAST Upgrade- Construction Status, Fusion Engineering and Design, 96-97 (2015) (42-47)
- [4] E van der Bij et al, Open Hardware for CERN's accelerator control system, Journal of Instrumentation, 7 (2012).
- [5] Xilinx Aurora Protocol (August 2018) <https://www.xilinx.com/products/intellectual-property/aurora8b10b.html>
- [6] Xilinx AC701 Evaluation Kit (August 2018) <https://www.xilinx.com/products/boards-and-kits/ek-a7-ac701-g.html>
- [7] Xillybus FPGA IP and PC driver (August 2018) <http://www.xillybus.com>