Control Engineering Practice 68 (2017) 15–22

Contents lists available at ScienceDirect

Control Engineering Practice

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

GPU-based optimal control for RWM feedback in tokamaks

Mitchell Clement [a,b,*], Jeremy Hanson [b], Jim Bialek [b], Gerald Navratil [b]

[a] Mechanical and Aerospace Engineering, University of California, San Diego, La Jolla, CA 92093, USA
[b] Department of Applied Physics and Math, Columbia University, New York, NY 10027, USA

ARTICLE INFO

Keywords:
Optimal control
LQG control
GPU computing
Tokamaks
Nuclear fusion
Plasma physics
Magnetohydrodynamics
Plasma control

ABSTRACT

The design and implementation of a Graphics Processing Unit (GPU) based Resistive Wall Mode (RWM) controller to perform feedback control on the RWM using Linear Quadratic Gaussian (LQG) control is reported herein. The control algorithm is based on a simplified DIII-D VALEN model. By using NVIDIA's GPUDirect RDMA framework, the digitizer and output module are able to write and read directly to and from GPU memory, eliminating memory transfers between host and GPU. The system and algorithm was able to reduce plasma response excited by externally applied fields by 32% during development experiments.

1. Introduction

Tokamaks can excite kink modes which can lock or nearly lock to the vacuum vessel wall, whose rotation frequencies and growth rate vary in time but are generally inversely proportional to the magnetic flux diffusion time of the wall (Chu et al., 2010). The growth times of such modes are determined in part by the magnetic flux diffusion time of the resistive wall and they are therefore called resistive wall modes (RWMs). The need to maintain adequate margin to the onset of such unstable modes can limit the maximum achievable β for tokamak plasma. β is defined as the ratio or the plasma's hydrodynamic pressure to the confining magnetic field pressure(%). RWMs can also lead to catastrophic events in tokamaks called disruptions. Disruptions lead to a rapid termination of the magnetically confined plasma and deposit energetic particles and heat on plasma facing components. In future devices where the plasma's energy content is expected to be significant, physical damage to the machine itself may occur (Bateman, 1978). Advanced Tokamak (AT) designs are predicated to operate at high plasma β in order to maximize fusion power gain (Najmabadi et al., 2006) and will have to operate with very few, if any, disruptions over their lifetimes.

There are two generally accepted approaches to stabilizing the RWM. The first being plasma rotation in the toroidal direction, in which the plasma moving relative to the wall generates eddy currents and in turn magnetic fields which have a stabilizing effect (Chu et al., 2010). The second method is by applying stabilizing magnetic fields using electromagnetic control coils external to the plasma (Chu et al., 2010; Strait, 2015). Simulations show that the RWM may or may not be fully suppressed by rotation alone in future tokamak devices operating with high performance plasma conditions (Berkery et al., 2010; Chapman et al., 2012; Liu, 2006, 2009). Therefore advanced techniques for active feedback control of this magnetohydrodynamic (MHD) instability are further investigated.

The DIII-D tokamak has two sets of magnetic field control coils, 12 interior coils (I-coils) and six non-axisymmetric external coils (C-coils), which are used for error field correction and other 3D field control. The I-coils are typically used for RWM feedback due to their proximity to the plasma edge and because they are not encumbered by the relatively lengthy magnetic flux diffusion time through the vacuum vessel. The present algorithm for RWM feedback uses Proportional Integral Derivative (PID) control, using only some of the available poloidal field sensors, to determine voltages to be applied to the coils (Strait et al., 2004). Future ATs built to demonstrate burning plasma conditions will likely not have control coils internal to the vacuum vessel. Simulations and modeling predict that classical control techniques such as PID control, used with external coils on the proposed ITER tokamak may be ineffective at stabilizing the RWM or may use excessive current in doing so (Katsuro-Hopkins et al., 2007). Thus more advanced control techniques will need to be investigated for stabilizing the RWM using external coils relevant to future machines.

For active feedback, a variety of control schemes have been investigated to counteract the RWM, including classical control and state-space methods (Strait, 2015). A Linear Quadratic optimal controller for RWM feedback on ITER using the in-vessel control coils has also been proposed (Ariola et al., 2014). Other fusion experiments have had

---

[*] Corresponding author at: Mechanical and Aerospace Engineering, University of California, San Diego, La Jolla, CA 92093, USA.
E-mail address: mdclemen@eng.ucsd.edu (M. Clement).

http://dx.doi.org/10.1016/j.conengprac.2017.08.002
Received 4 January 2017; Received in revised form 22 June 2017; Accepted 8 August 2017

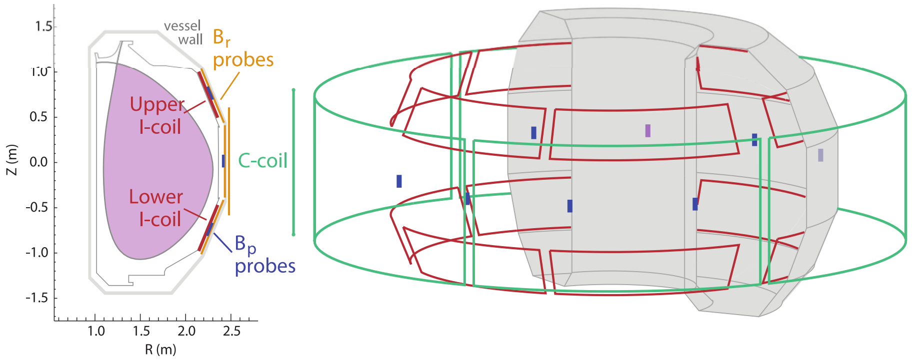

Fig. 1. Simple cross sectional view and schematic of DIII-D sensors and 3D control coils.

success in controlling the RWM using largely classical algorithms (Drake et al., 2005; Martin et al., 2009), but it is worth noting that these machines have very long flux diffusion times for their walls compared to DIII-D and are not able to reach plasma parameters relevant to advanced tokamaks. Two Reverse Field Pinch (RFP) experiments, which are similar in configuration to tokamaks, have reported using system identification techniques to develop state space models for RWM feedback (Olofsson et al., 2013). Other tokamaks have reported success in stabilizing the RWM via active control methods (Sabbagh et al., 2006). Simulation of controllers based on lower dimensional RWM models have yielded promising results for applications on DIII-D (Dalessio et al., 2008; In et al., 2006). Numerical simulations of feedback using Linear Quadratic Gaussian (LQG) control with external coils and the VALEN RWM model of the ITER tokamak have also yielded promising results (Katsuro-Hopkins et al., 2007). LQG and state space control in general allow for flexibility in the number of inputs and outputs as well as the ability to filter measurement and process noise. Pole-placement does not necessarily work for Multiple Input/Multiple Output systems to find a stabilizing control law, whereas LQG methods guarantee a stabilizing solution so long as the model is controllable and observable. Model-based control is expected to help optimize the penetration of control fields through the vacuum vessel wall.

The system described herein is largely based on the system installed at the HBT-EP tokamak (Sankar et al., 1993), as it demonstrated the concept of using a GPU for fast real-time control computation in the microsecond regime, a task that is non-traditional in the relatively young field of GPU computing (Rath et al., 2012). GPUs are programmed in conventional programming languages like C, and therefor can leverage all the benefits of digital control such as decision-making or logic flow and extremely flexible control programs (Franklin, Powell, & Workman, 1990). In this regard a GPU has advantages over a conventional CPU for real-time control. For instance, the GPU has no operating system to run or subject to processor interrupts and can therefore be used entirely for control purposes. The CPU can initialize the GPU and acquisition hardware, and then sit back and relax while the GPU handles the real-time control computations. Major differences between DIII-D's and HBT-EP's system are host computer architecture, 64bit(DIII-D) versus 32bit(HBT-EP) host, a superior GPU, and the use of an LQG algorithm for MHD mode control.

Section 2 describes the control system hardware, the VALEN RWM model and algorithm design. Section 3 provides an overview of initial system testing, performance and control development experiments on DIII-D.

2. Control system and algorithm design

2.1. Requirements

DIII-D can measure low frequency (<20 kHz) toroidal modes with toroidal mode number, n>0 with 34 poloidal can-type probe difference pairs arranged in five toroidal arrays, and radial magnetic field with 38 saddle loop difference pairs arranged in six toroidal arrays (King et al., 2014). Both $\mathbf{b}_p$ and $\mathbf{b}_r$ arrays span 360 degrees of toroidal angle. Due to space constraints, the control system temporarily only has access to 24 of the 72 total difference pairs. Sensors located at, above, and below the low field side midplane are used and their respective locations are shown in Fig. 1. The I and C coils are typically configured in $n = 1$ quartets or pairs respectively for feedback on the $n = 1$ component of the RWM, which is expected to be dominant. This translates to three control commands to amplifiers connected to the lead coil in each quartet or pair. Vacuum coupling between sensors and coils, i.e., poloidal field coils, needs to be accounted for and eliminated prior to using sensor measurements for feedback. System latency should be kept well below the lowest expected RWM growth time, 2.5 ms. A simple schematic of DIII-D with coil and sensor configurations is shown in Fig. 1. Not shown in Fig. 1 are the axisymmetric coils responsible for the equilibrium magnetic field which is much larger than the fields produced by the control coils.

2.2. Hardware

The DIII-D RWM control system integrates the following components into a low latency, high performance system:

- NVIDIA Tesla K20c GPU, 5 GB RAM GPU.
- D-TACQ Solutions ACQ196 96 channel, 16 bit digitizer.
- D-TACQ Solutions AO32 32 channel, 16 bit analog output module.
- National Instruments PXI-PCIe8362, MXI-Express, 2 Port PCIe host bus adapter.
- Supermicro X9DAI-0 Motherboard running 64 bit CentOS 6.5 with kernel 2.6.32.

The GPU resides in a PCIe-x16 slot on the motherboard in the same root complex as the National Instruments host bus adapters(HBAs). The HBAs each connect via cable to a Rear Transition Module (RTM-T) attached to either the ACQ196 or AO32. The ACQ196, AO32, and both RTM-Ts are housed in a 2U Compact Peripheral Component Interconnect (CPCI) chassis. When memory buffers on the GPU are allocated for input and output, their physical addresses on the system are found with NVIDIA's GPUDirect Remote Direct Memory Access(RDMA) framework (NVIDIA Corporation, 2017). These physical addresses are passed to the RTM-T device drivers in the OS' kernel for realtime streaming of data to the GPU's onboard memory. On every sample, new data is streamed from the digitizer, the algorithm operates on that data, and writes its command to the output buffer. The output buffer is read by the output module and does a zero order hold, i.e., the value in the output buffer is maintained throughout the sample interval. A diagram of the control system is shown in Fig. 2.