General Electric

Transient Efficiency, Flexibility, and Reliability Optimization of Coal-Fired Power Plants

DE-FE0031767

Final Report

Principal Investigator
Dr. Aditya Kumar
Sr. Principal Engineer
Controls & Optimization
GE Research
1 Research Circle, K1-5C29A
Niskayuna, NY 12309
Email: kumara@ge.com
Phone: 518-387-6716

GE Contracts Manager
Ahsan Zaidi
Contract Administrator
GE Research
1 Research Circle
Niskayuna, NY 12309
Email: ahsan.zaidi@ge.com

Feb 2023

This material is based upon work supported by the Department of Energy under Award Number DE-FE0031767.

Disclaimer: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

1 ABSTRACT

This program developed an advanced model-based monitoring and model-predictive control algorithms for a coal fired power plant (CFPP), and deployed these algorithms in a real-time platform to demonstrate performance benefits for transient flexibility and plant operation efficiency. More specifically, the objectives were successfully achieved through a combination of (i) developing a high-fidelity transient plant model in Apros, which was used as a high-fidelity plant simulation between $100 - 50\% TMCR$ where TMCR denotes the turbine maximum continuous rating, i.e., baseload, (ii) developing a very fast physics-based reduced-order model (ROM) of the plant, which ran more than $100\times$ faster than real-time, enabling its use as real-time embedded model for model-based estimation (MBE) and model predictive control (MPC) (iii) implementing a real-time MBE based on ROM using a robust unscented Kalman filter (UKF) to continuously tune the ROM to match the measurements from high-fidelity Apros plant model despite significant plant-model mismatch, and thus, obtain a Digital Twin of the plant (iv) designing and implementing a real-time MPC with dual objectives of transient plant load tracking with high ramp rates and minimizing coal consumption, i.e., improving plant efficiency in the baseload-partload operation range of $100 - 50\% TMCR$. Each key element above was developed and tested individually, and has been reported in corresponding Topical Reports. Finally, all the individual elements were integrated in an overall closed-loop system, that was successfully tested in desktop Simulink test harness simulations with ROM or high-fidelity model as the plant. Thereafter, the Simulink implementation was used to auto-generate C-code and deploy as real-time Docker microservice containers in Linux, and validate that they can run in real-time in the hardware-in-the loop (HIL) setup and produce the same results as in Simulink. The results of the integrated simulation tests in Simulink as well as the real-time HIL deployment are documented in this final report, showing good load tracking for load ramps at $3 - 4\%/min$ ramp rates, and achieving up to 5.5% reduction in coal relative to baseline operation at $50\% TMCR$ load. The desktop and HIL simulations show successful performance of the overall model based estimation and control solution and achieve the key objectives of the program for flexible, efficient and reliable operation of subcritical coal fired power plants.

2 Project Goals and Achievements

Transient Efficiency, Flexibility, and Reliability Optimization of Coal-Fired Power Plants program will benefit coal-fired power plants (CFPPs) by improving the flexibility, efficiency, and reliability of the CFPPs to help them meet the emerging operation demands. More specifically, as older CFPP plants are retired towards meeting decarbonization goals, the remaining active CFPPs have to operate even better, i.e., (i) be more flexible in terms of cyclic operation between baseload and partload in response to growing renewable generation that is inherently variable, and (ii) be more efficient.

Figure 1 shows the overview of the program and key goals. In this program, GE Research, working closely with GE Steam Power, seeks to develop advanced model-based controls technology to enhance the transient flexibility and efficiency of CFPPs. The two key metrics are (i) achieve 30% faster transients for load variation between $100 - 50\%$ TMCR compared to current baseline operation at $2 - 3\%/min$ load ramps, and (ii) achieve up to 5% relative improvement in plant efficiency across this load operation. Note that for a typical CFPP with a baseload efficiency of around 37%, a 5% relative improvement implies an improvement in overall efficiency by 1.85%, i.e., improved efficiency of 38.85%. Such an improvement in efficiency would lead to a corresponding reduction in the cost of fuel as well as CO_2 emissions on a per MW basis. The technology is developed and tested both in desktop simulations as well as in real-time HIL simulations, in the presence of significant plant-model mismatch for demonstrated robustness, to achieve target maturity of TRL 5.

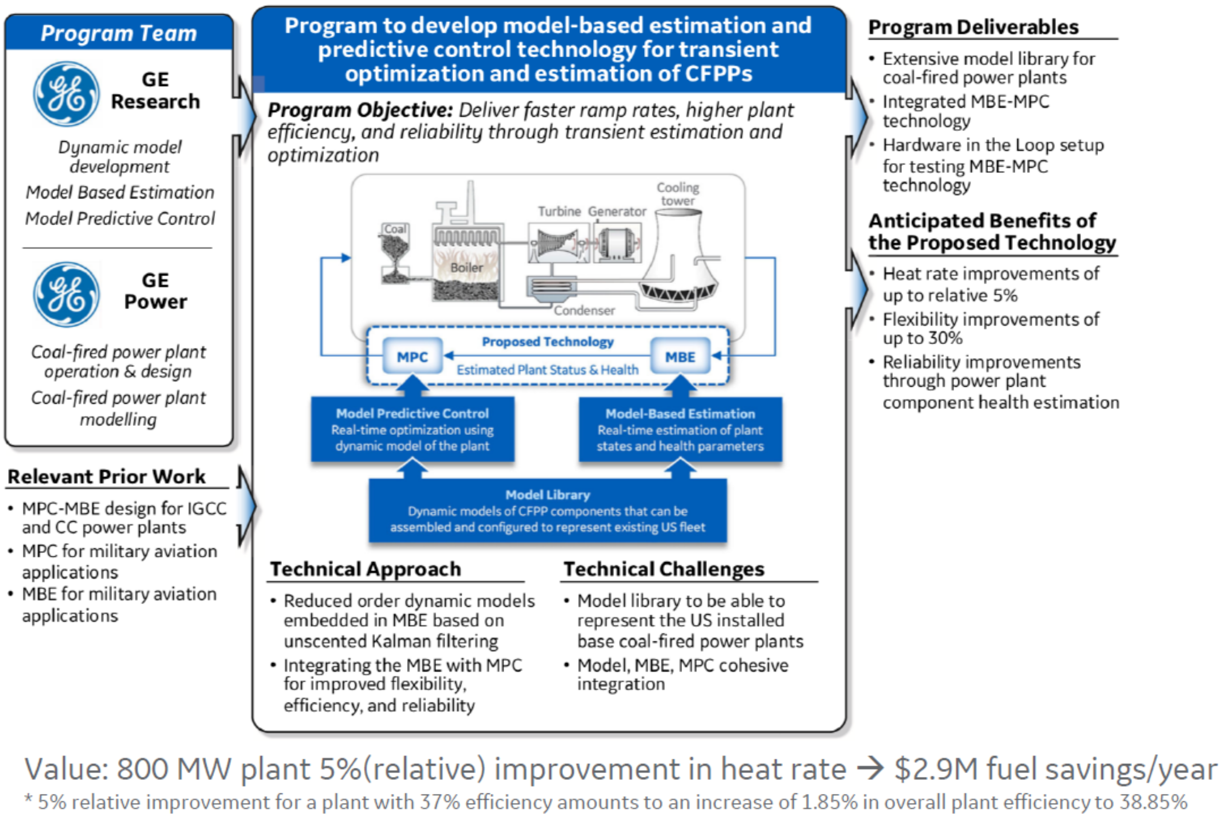

Figure 1: Overall program goals - improved transient flexibility, efficiency and reliability of coal fired power plants.

To achieve these goals, a model-based optimal control solution is proposed as shown in Figure 2 that will demonstrate the technology in both desktop closed-loop simulations as well as in a real-time hardware-in-the-loop (HIL) deployment. The overall architecture consists of four key elements:

1. *High-fidelity transient model of CFPP* - A high-fidelity transient model is implemented in Apros based on design specifications of a large 820 MW plant. The model includes all key components of the plant as well as the current baseline controller to allow load transients between $100 - 50\%$ TMCR. This high-fidelity plant model serves as the real-time simulation of a real plant where specific measurements are obtained for MBE, and MPC computed optimal control actions can be applied. The high-fidelity plant is accurate and runs faster than real-time - it just needs to be slightly faster than real-time to enable closed-loop real-time simulations.

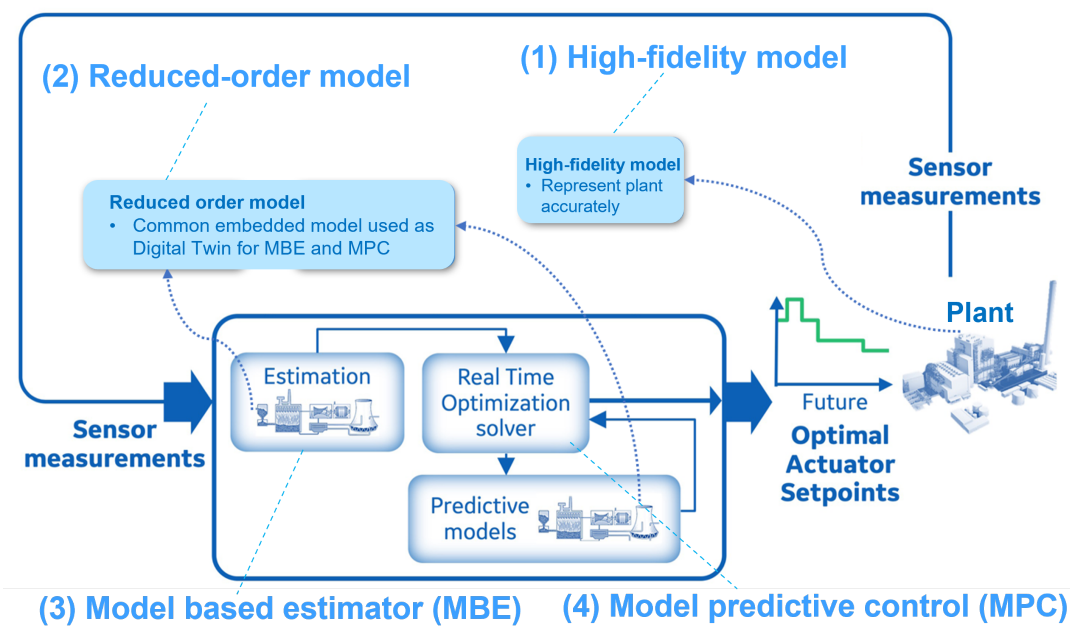

Figure 2: Model-Based Estimator (MBE) and Model Predictive Controller (MPC) test architecture for optimizing a coal-fired power plant. MBE and MPC have embedded reduced-order CFPP models whereas the actual CFPP is represented by a high-fidelity model.

2. *Reduced order model (ROM) of CFPP* - A reduced-order physics-based transient model was developed in a modular fashion for each key module in the plant, along with simple low level control loops (P or PI controllers) to enable very fast (more than 100× faster than real-time) transient simulations. The model is also parameterized with adjustable parameters in key modules to allow using an MBE to continuously tune the model with measurements from the high-fidelity plant simulation, thereby having a Digital Twin of the plant. The reduced order model needs to be robust and very fast to be used as embedded model in MBE as well as MPC - both need to make many calls to the ROM at each sample time.

3. *Model-based estimation (MBE)* - A model-based estimation algorithm was developed using unscented Kalman filter (UKF) to perform a joint estimation of the transient model states as well as the tunable parameters in ROM to keep the ROM matched to measurements from plant simulation. The MBE computes real-time estimates of the states and parameters to obtain the matched ROM as a Digital Twin of the plant, that can be used in MPC to perform real-time transient optimization.

4. *Model predictive control (MPC)* - The ROM was embedded in MPC to perform transient optimization to achieve the two goals of transient load setpoint tracking during fast load ramps, as well as simultaneously minimizing the coal consumption, thereby improving plant efficiency. To achieve a robust, and computationally efficient solution, a hybrid MPC solution was implemented that leverages the full nonlinear model prediction, coupled with online linearization and quadratic programming (QP) solution.

While the high-fidelity model was implemented in Apros, the rest of the modeling and algorithms were implemented in Matlab & Simulink. The closed-loop integration between Simulink and Apros was achieved using real-time two-way communication between the two using Apros API and real-time Data Distribution service (DDS). The high-fidelity model serves as a real plant to demonstrate robustness in the proposed solution even in the presence of significant plant-model mismatch. The high-fidelity plant model is structurally different since it captures the key components of the CFPP and baseline controls in great detail. In contrast, the ROM is deliberately simpler to achieve high simulation speeds and numerical robustness for use as embedded model. Demonstrating such robustness to plant-model mismatch even in simulation studies is very important to ensure that when the model-based technology is applied to a real plant, the solution can work robustly - all models are approximations of the real system so there is always a plant-model mismatch.

The above four elements were developed individually and tested, and then finally integrated into the overall closed-loop simulation setup to test its performance against the objectives. Initially, the closed-loop simulation was tested using the ROM as the plant itself, i.e., no structural plant-model mismatch. These simulations were used to systematically demonstrate a fast load ramp at $4\% TMCR/min$ both going down from 100% to 50% without active minimization of coal. Thereafter, the coal minimization objective was also included to demonstrate that the proposed solution could simultaneously achieve good tracking of load ramps at $4\%/min$ as well as achieve up to 5.5% reduction in coal use at 50% load in relative terms compared to the MPC simulation without active coal minimization. These runs were then replicated in the real-time HIL deployment. To this end, the developed MBE+MPC+ROM implementation in Simulink were processed through Mathwork's auto C-code generation and thereafter encapsulated into Docker container microservices running in Ubuntu Linux. The HIL simulation validated (i) same results as the desktop Simulink simulations, and (ii) ensured it can easily run in real-time. Finally, the same transient simulations were performed with the high-fidelity Apros model as the plant, i.e., test robustness of ROM+MBE+MPC solution in the presence of a realistic plant-model mismatch. Again, initial set of simulations were performed in desktop simulations in Simulink to test its performance, and then replicated in the HIL deployment in the same manner to ensure that the algorithms could work robustly in real-time. In contrast to the ROM-based plant simulations where the ROM plant simulation was also embedded in the microservice container, with Apros high-fidelity model based plant simulations, the two-way communication between the deployed microservice container and high-fidelity model in Apros software had to be verified across Linux and Windows. These simulation results are discussed in detail in Section 4, and they demonstrate the maturity of the developed technology to TRL 5.

3 Technology Development

This section summarizes the development of the four individual technology elements mentioned in the previous section.

3.1 High-fidelity Transient Model

GE Steam Power team developed and implemented a high-fidelity physics-based transient model of an 820 MW CFPP in Apros. The model contains detailed implementations of key sub-components including boiler, drum, superheaters, steam turbine, and air and water preheaters (see Figure 3 for representative key components modeled in Apros). The implementation is based on design specifications of an actual 820 MW power plant, and includes baseline lower level and boiler control to achieve both steady-state and transient operation between 100% and 50% TMCR load. This model is described in detail in a Topical report on this topic submitted to DOE[1].

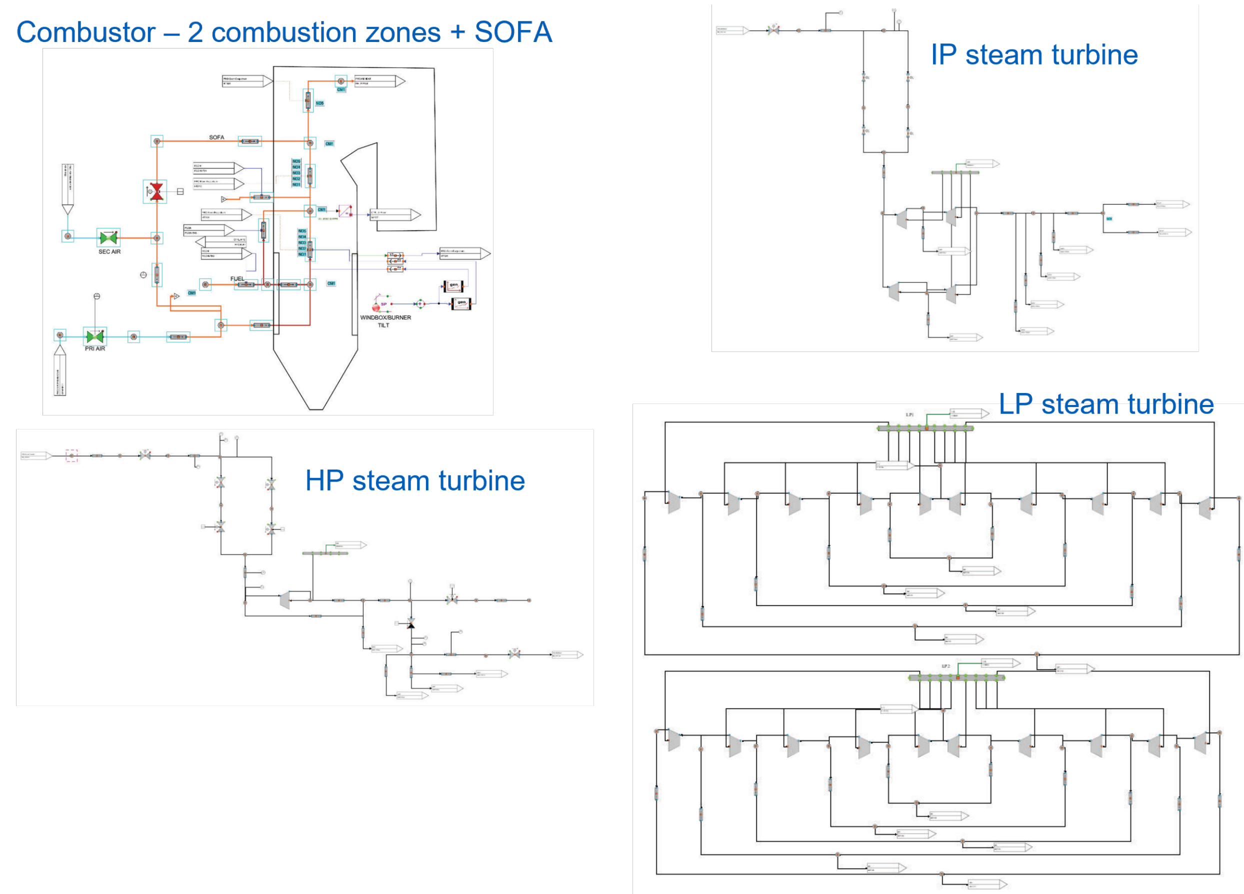

Figure 3: Representative components of coal fired power plant modeled with high fidelity in Apros.

3.2 Reduced Order Transient Model

The Reduced order Model (ROM) was implemented in Matlab/Simulink. In particular, individual modules including coal combustor with multiple zones, boiler drum, superheater heat exchangers for high and intermediate pressure sections, steam turbine for high/intermediate/low pressure sections, air preheaters and water preheaters & economizers (see Figure 4 for some key sub-components and modules in Simulink). Each components was implemented and tested individually in respective operation range, and then integrated in the desired configuration for the reference plant to obtain the overall CFPP transient model that ran robustly across the 100 – 50% TMCR range. Each module also includes specific adjustable parameters, e.g., multipliers on heat transfer coefficients, turbine efficiency, so that the ROM could be transiently tuned to match measurements from the high-fidelity plant model. The ROM was implemented with a careful consideration for the right trade-off between fidelity and computational efficiency to achieve robust transient simulations that are more than 100× faster than real time. This was critical to ensure it could be used as a real-time embedded model in MBE and MPC, which call the ROM multiple times in a single time sample. Also, each module was implemented to ensure it could meet the requirements for Mathworks' auto C-code generation for deployment. Details on the ROM are included in a Topical report submitted to DOE[2]

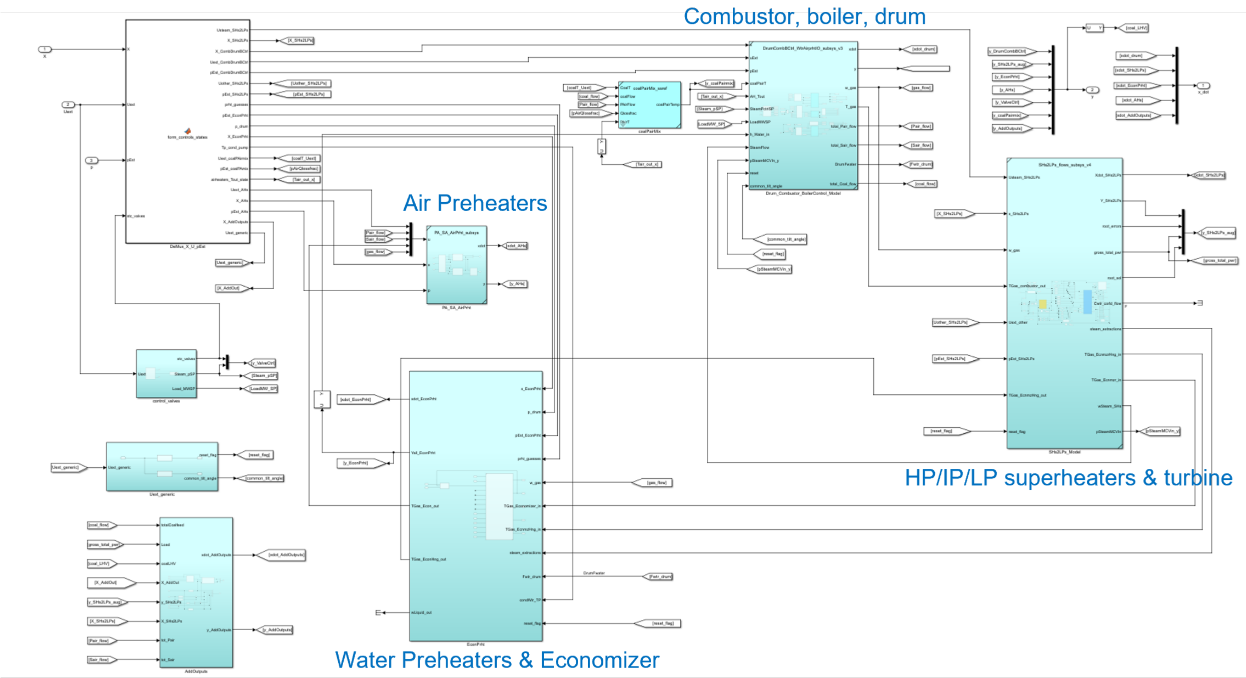

Figure 4: Modular implementation of key sub-components in reduced-order model (ROM) in Simulink.

3.3 Model Based Estimation

An unscented Kalman filter (UKF) was implemented in Simulink in a standardized library that can be integrated with the ROM as the embedded model. Figure 5 shows the use of ROM as embedded model used with UKF to tune the ROM with real-time input & output measurements from the plant, to obtain the resulting Digital Twin. The Simulink implementation allowed rapid prototyping and testing both at the MBE modular level as well as the final closed-loop simulation, and deployment via Mathworks' auto C-code generation and encapsulation in microservice container for HIL deployment. The MBE performs a joint state and parameter estimation, for all ROM transient states and selected subset of adjustable model parameters. It is easily customizable and tunable for each state and parameter being estimated via an Excel file. Details on the MBE implementation, tuning and testing are included in the Topical report submitted to DOE[3]

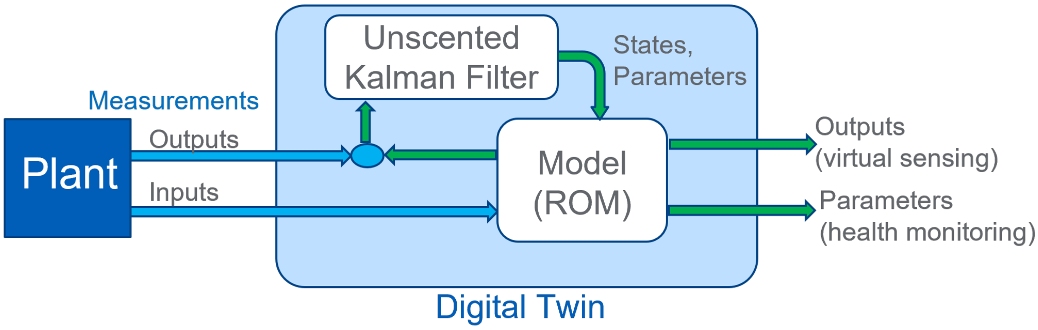

Figure 5: Model-based estimation (MBE) using Unscented Kalman filter with ROM and plant measurements to obtain a Digital Twin.

3.4 Model Predictive Control

A model predictive controller (MPC) was designed and implemented in a standardized library in Simulink and integrated with the ROM as the embedded model. Figure 6 shows the MPC integrating the tuned ROM-based predictions and online constrained optimization. The developed MPC uses a hybrid combination of the full nonlinear prediction with the ROM using the latest state and parameter estimates from MBE, and state-space linearization of the ROM for formulating and solving a quadratic programming (QP) optimization problem. This hybrid approach allows a good balance between a full nonlinear MPC that can be unduly sensitive and computationally expensive for real-time of such complex plants, and a linear MPC that completely ignores the plant nonlinearities. The hybrid approach achieves close to entitlement performance while guaranteeing a robust optimization convergence for a convex QP problem. Moreover, the MPC formulation was implemented with muti-objective optimization, in addition to all applicable input and output constraints. Specifically, it included a transient tracking objective to follow fast load transients, as well a performance optimization objective to maximize plant efficiency, or equivalently minimize coal use. The latter objective is achieved as a secondary goal while ensuring the primary goal of transient load tracking is not sacrificed. This multi-objective formulation is included as the following objective function:

$$J_k(U_k') = \frac{1}{2}\left(\Gamma_k - Y_k\right)^T Q_Y \left(\Gamma_k - Y_k\right) + L_{Y,perf} Y_k + \frac{1}{2}\Delta U_k^T R_U \Delta U_k. \tag{1}$$

where the first term denotes the quadratic penalty on output tracking error with a weight Q_Y, and the last term denotes the quadratic penalty on the change ΔU_k in control action U_k' over successive time samples k and $k-1$ with a quadratic weight R_U. The second term includes any performance optimization on performance outputs (e.g., plant efficiency) with a linear weight $L_{Y,perf}$. The linear weight can be positive or negative depending on whether the performance output is being minimized or maximized, respectively. Details on the MPC development, formulation and simulation studies are included in the Topical report submitted to DOE[4]

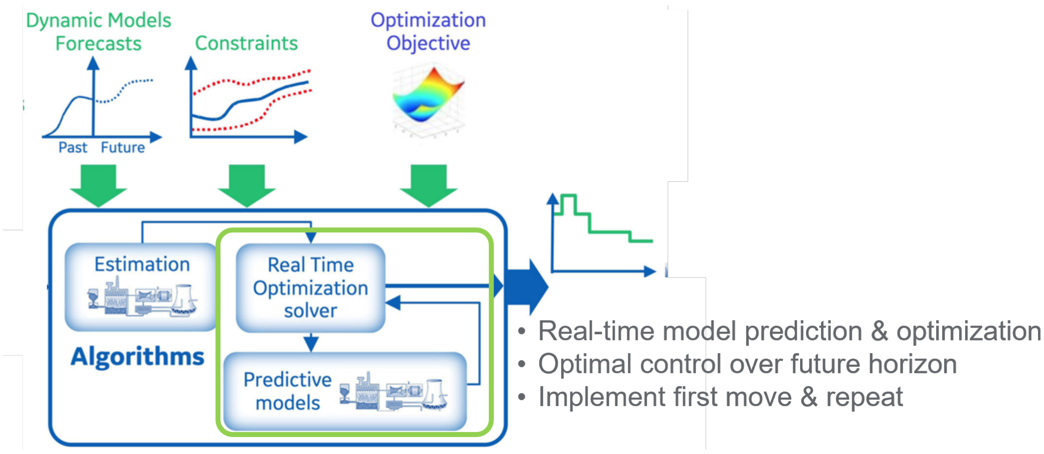

Figure 6: Model predictive control (MPC) using tuned ROM for a receding horizon prediction and constrained optimization.

4 Closed-Loop Simulation Studies

This section describes in detail the final simulation studies for the overall integrated closed-loop system.

4.1 Simulink Test Harness for Desktop Simulation Studies

Initially, these simulations were performed in Simulink with ROM as the plant model or Apros as the plant model - this allowed debugging, and fine tuning MBE and MPC as well as addressing issues identified in the closed-loop simulation, mainly on improving the embedded ROM robustness. The Simulink test harness for these studies is shown in Figure 7. It includes the control panel (e.g., to command a load transient and generate forecast setpoint and disturbances for MPC horizon), the monitoring block where key signals are plotted on scopes and logged for saving the results, and the core blocks for (i) PlantConnectivityBridge - allows using either the ROM as plant model, or provides a two-way communication bridge to the high-fidelity model in Apros on another Windows PC, (ii) UKFEstimation - the block for MBE using ROM as the embedded model and doing state and parameter estimation from the plant measurements available from the PlantConectivityBridge, (iii) MPCController - the block for MPC using the same ROM embedded model, using real-time state and parameter estimates from MBE, and computing the optimal control actions sent to the PlantConnectivityBridge to close the loop, and (iv) PersistentStore - used to initialize states of ROM, MBE, MPC and also store any intermediate results. This test harness leverages the individual ROM, MBE and MPC libraries in modular fashion using subsystem references, and is highly configurable, e.g., to enable or disable MPC (i.e. run MBE alone in open-loop, or run MBE+MPC in closed-loop), use ROM or two-way bridge with high-fidelity model in Apros as the plant model. Also, the MBE and MPC could be configured and tuned using Excel files to configure all I/Os and corresponding tuning of MBE and MPC, as reported in the Topical reports. The overall closed-loop simulation was executed with a 20s sample time, consistent with the time constants of the overall plant and the supervisory plant-level control via MPC.

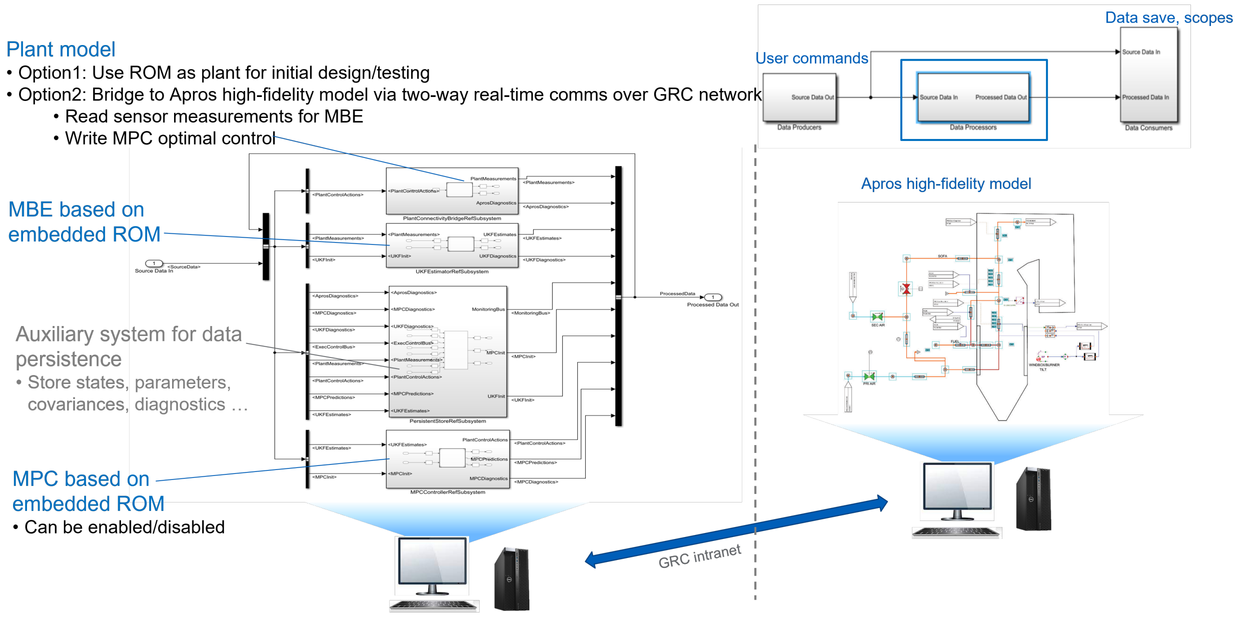

Figure 7: Overall test harness in Simulink for closed-loop simulation testing of MBE+MPC+Plant.

4.2 HIL Deployment for Real-time Simulation Testing

Once the closed-loop simulation performance was verified in the Simulink test harness, it was deployed in the HIL platform for real-time simulation testing. This deployment is achieved using a combination of Mathworks' auto C-code generation from the Simulink blocks in the harness, and encapsulation of the C-code into a Docker microservice container running in Ubuntu Linux. Initially, the aim was to deploy it on an Intel NUC computer shown in Figure 8 natively running Ubuntu 12 Linux OS with the Docker container management and middleware communication layers underneath. This is an affordable and very capable embedded software deployment platform adequate for the purposes of the project.

Figure 8: Target deployment platform: Intel NUC computer.

Figure 9 shows the original plan for the target HIL setup, which includes (i) Windows desktop PC that runs Simulink test harness modules for data generation (e.g., load setpoint variation), and data viewing/saving via live scopes, (ii) another Windows desktop PC that runs the APROS high-fidelity model and the APROS communication bridge client for live

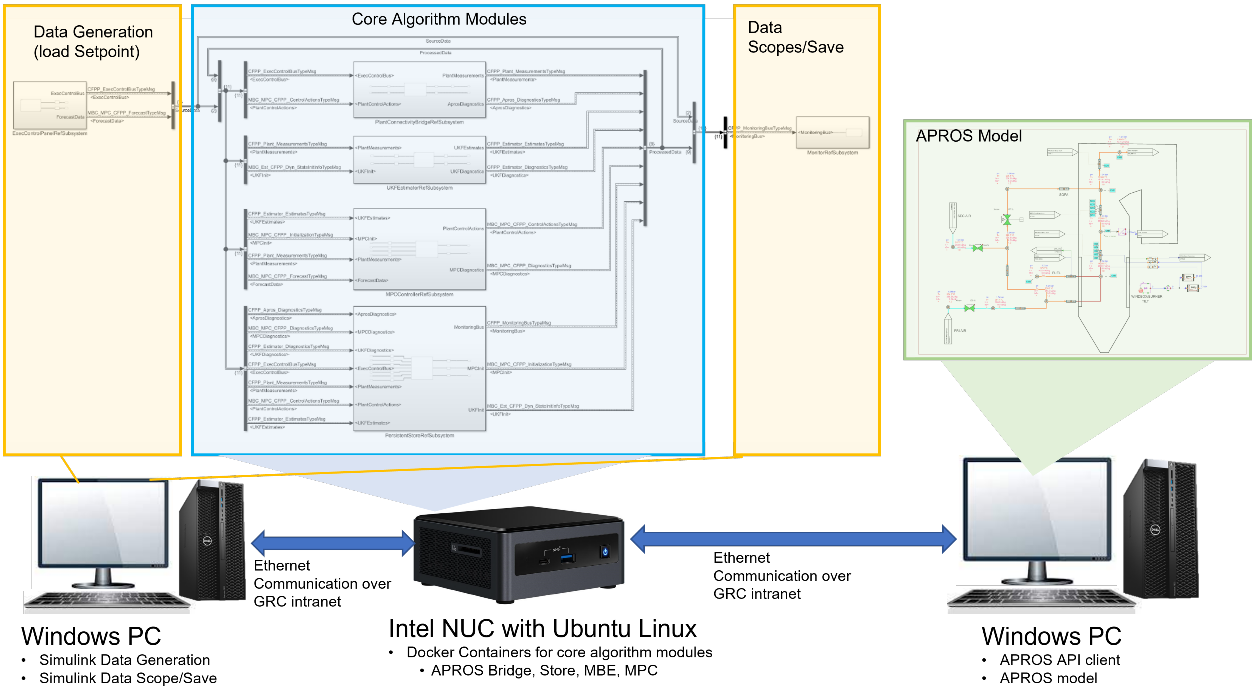

Figure 9: Target deployment Architecture with separate Windows Desktop Computers, Intel NUC and real-time two way communication over GE Research intranet

two-way read/write communication (e.g., writing MPC commands, and reading MPC and UKF measurements), and (iii) an Intel NUC PC with Ubuntu Linux where containerized implementations of the core algorithms (data store, APROS bridge, MBE and MPC) execute in real time. One challenge is the enabling of appropriate network communication access on the Intel NUC over the GE Research intranet. Since this is not a standard image computer, there are limitations in getting that network access and communication working over the GE Research intranet. In light of this, a modified deployment architecture was pursued as shown in Figure 10. This modified HIL architecture proves the same real-time deployment capabilities, while avoiding the pitfalls of network access with non-standard Linux PC.

In this architecture, the Windows desktop PC runs the Apros high-fidelity model natively in Windows OS, and also has an Ubuntu Linux VM (virtual machine), where the deployment code for ROM+MBE+MPC are executed real-time in a Docker microservice container, along with a real-time two-way communication via the bridge between the Apros model running in Windows OS and the microservice container running in the Linux VM. Also, the ability to run Docker containers for the core modules on a Linux VM in real-time will prove that they can be deployed real-time on the NUC platform with native Ubuntu Linux OS as well - if anything deployment in native Ubuntu Linux will run faster than in a Linux VM on a Windows PC.

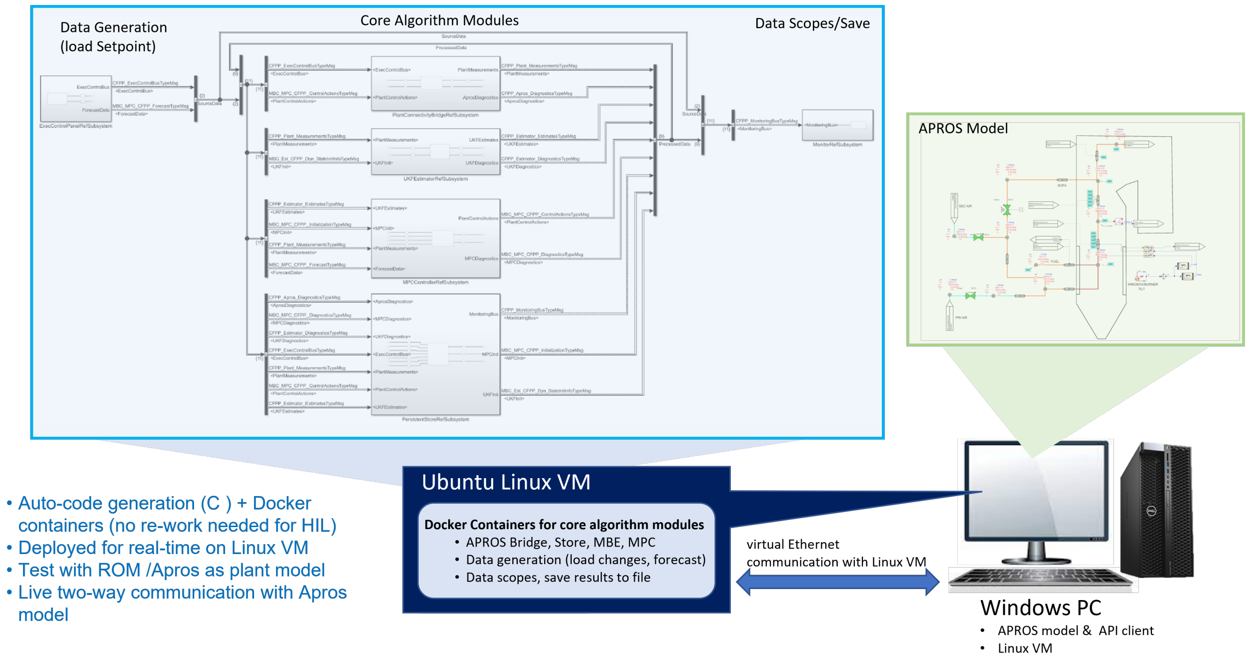

Figure 10: Target deployment Architecture with separate Windows Desktop Computers, Linux VM and real-time two way communication over GE Research intranet

4.3 Closed-Loop Simulations with ROM as Plant Model

This section summarizes the simulation results for the closed-loop test harness desktop simulations as well as real-time simulations with the HIL deployment using container microservices. For these simulations, MBE and MPC were enabled and the ROM was used as the plant simulation in closed-loop - i.e., there is no structural mismatch between the plant model and the embedded ROM used in MBE and MPC. This is to establish performance entitlement in terms of transient tracking for fast load ramp rates, and plant efficiency improvements. For these simulations, three control inputs are enabled (i) steam pressure setpoint for the existing boiler controller that calculates coal and air flows, (ii) master control valve (MCV) opening for the high-pressure steam flow to the steam turbine, and (iii) common tilt angle for the coal and air feeds in the combustor - the tilt angle is the same for all three combustor zones as is currently implemented in plant controls (in future there could be further improvements by allowing independent positive or negative tilt in each combustor zone). The MPC objective includes a quadratic term for load tracking error, with a weight Q_Y, a linear objective for coal minimization with a weight $L_{Y,perf}$, and a quadratic term for the changes in control actions with a weight R_U (see Eq 1). The weight $L_{Y,perf}$ is used to enable or disable coal minimization objective to see improvement relative to the baseline case with $L_{Y,perf} = 0$, i.e., only focused on load tracking.

4.3.1 Desktop simulations with Simulink test harness

The initial set of desktop simulations were performed with the Simulink test harness shown in Figure 7. The first set of runs were performed with $L_{Y,perf} = 0$, i.e., no emphasis to minimize coal use, and focus only on tracking fast load ramps. In particular, a fast load ramp from $100-50\%$ TMCR and then back from $50-100\%$ TMCR was studied with a ramp rate of $4\%TMCR/min$. Figure 11 shows the output profiles. Clearly, the plant load (MW) follows the fast ramps in load setpoint (MWref) very well throughout the transient run with no overshoots or oscillations, both during the ramp down and ramp up. Figure 12 shows the corresponding profiles for the three MPC inputs. The pressure setpoint changes the most as the load setpoint is changed between 100% and 50%, while MCV changes comparatively less, closing down to about 75% at 50% load, and tilt doesn't change much at all. Note, there are three control inputs for a single tracking output, so there are excess degrees of freedom available. This can be exploited to perform the secondary objective of minimizing coal use, i.e. maximizing plant efficiency, while maintaining the primary objective of load setpoint tracking. This is achieved by using a non-zero weight $L_{Y,perf} > 0$.

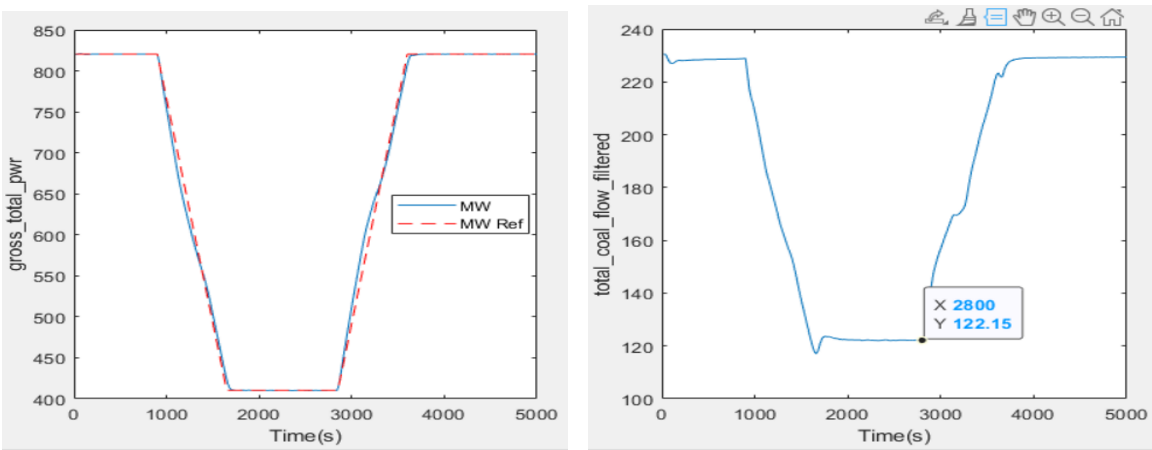

Figure 11: Load and coal feed profiles for a fast load ramp at $4\%/min$ between $100-50\%$ TMCR with no coal minimization objective.

The second set of simulation runs in Figure 13 and Figure 14 show the corresponding output and input profiles, for the same primary tracking objective, along with the secondary objective for coal minimization. First, the primary tracking objective is met as well as in the previous run, clearly illustrating that the secondary objective is achieved without sacrificing the primary objective. Second, indeed the coal use is reduced from $122.5lb/s$ in previous run to $117.5lb/s$ in this run, i.e., achieving a 3.8% relative reduction in coal use. To achieve this coal use reduction, the common tilt angle clearly changed to a negative tilt and saturated at the minimum allowed tilt of -15 deg. If the tilt limit were to be further reduced (e.g., to -30 deg), it is possible to achieve even further reduction in coal, or equivalently higher improvement in plant efficiency.

The key benefit by reducing the tilt to achieve the reduction in coal use, is to change the temperature profile in the combustor. More specifically, as the tilt is reduced to negative angle, the hot spot in the combustor is moved to lower zones. This, in turn does two things (i) it generates more steam with higher heat transfer to the water wall in the lower zones, and (ii) reduces the temperature of the flue gas from the top zone into the superheaters. Thus, it allows increasing steam flow and reducing superheating. The latter is especially important when there is already more than necessary superheating requiring water attemporation

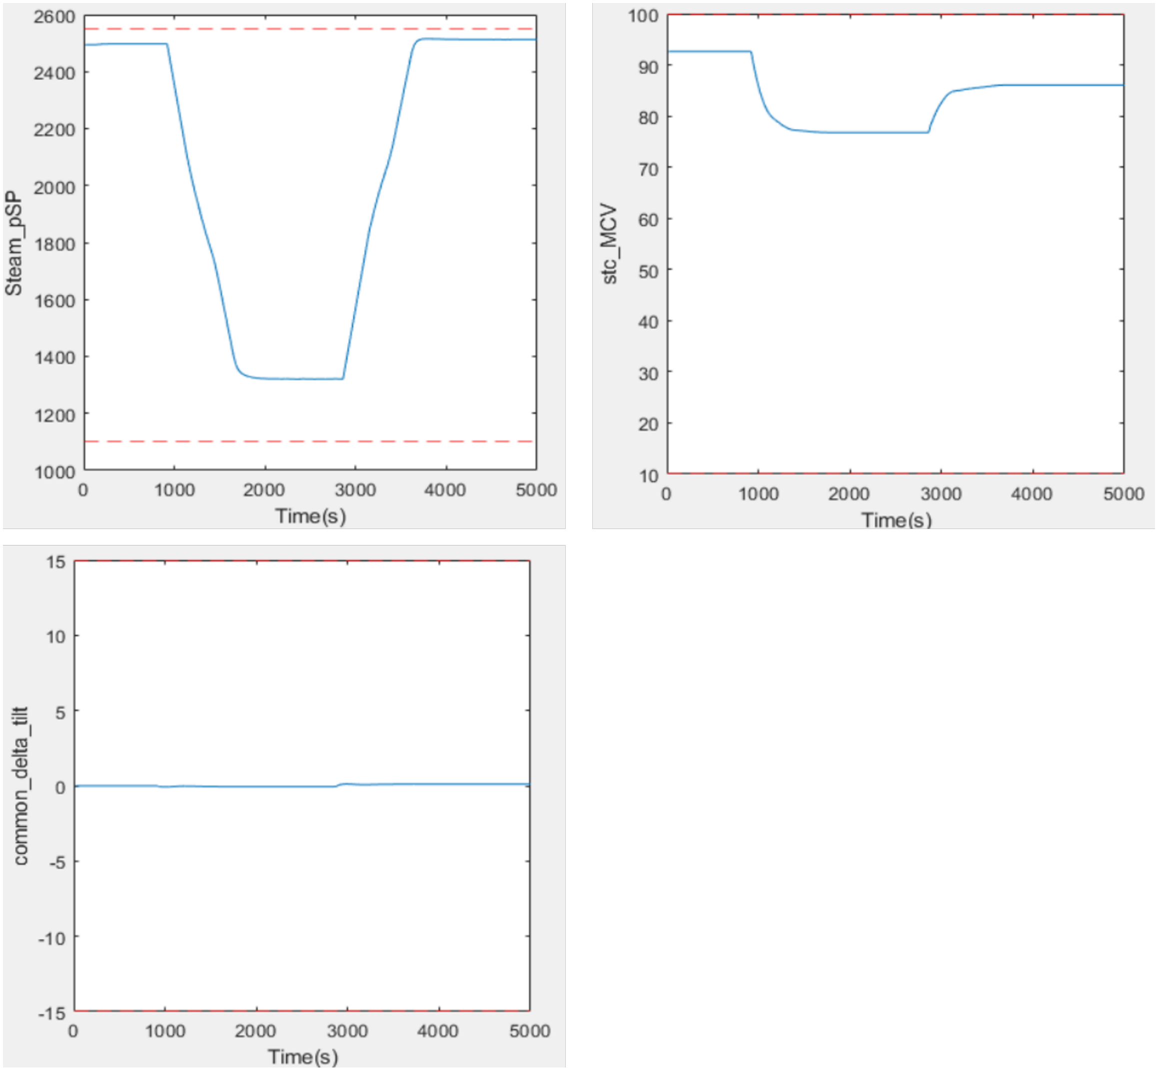

Figure 12: MPC input profiles for a fast load ramp at $4\%/min$ between $100-50\%$ TMCR with no coal minimization objective.

to manage the steam temperature in the superheaters for high-pressure and intermediate pressure sections. Such water injection in superheaters is detrimental to plant efficiency. Figure 15 shows the profiles for high and intermediate pressure steam temperatures at the inlet of MCV and ICV, and the corresponding attemporator water flows. Note that there is significant attemporator water flow, indicating excessive superheating. For a new plant design at 100% load, often it is designed to not need attemporation (to achieve high plant efficiency) at TMCR. However, often with equipment degradation over time, the same plant can reach operation conditions where attemporation is needed, with a corresponding reduction in efficiency. The high-fidelity model in Apros was tuned to such a degraded operation condition. The ROM was also tuned to match the same baseload condition. For the same transient run with explicit inclusion of the secondary objective for coal use minimization, the steam temperature and attemporator water flow profiles are shown in Figure 16, starting

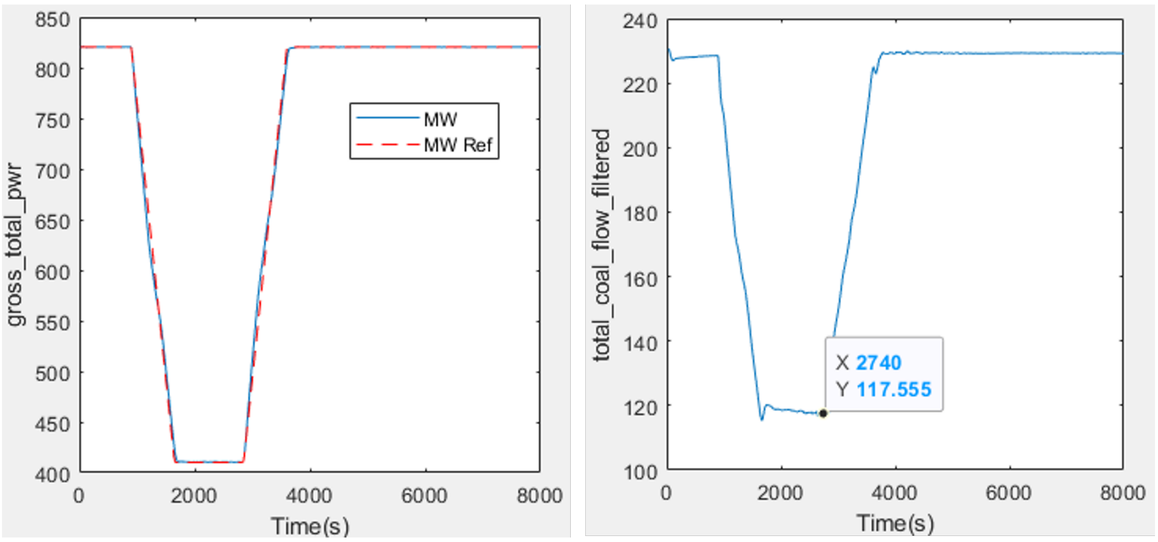

Figure 13: Load and coal feed profiles for a fast load ramp at $4\%/min$ between $100 - 50\%$ TMCR, including a secondary objective for coal minimization

from the same initial condition at $t = 0$. Note that, owing to the negative feeder tilt in the combustor, the steam temperature and consequently, the attemporator water flows are reduced at the 50% load as well as the final return to 100% load.

These simulations clearly show that the developed ROM+MBE+MPC solution successfully achieves the dual objective of fast load transients as well as plant efficiency improvements, when there is no significant plant-model mismatch.

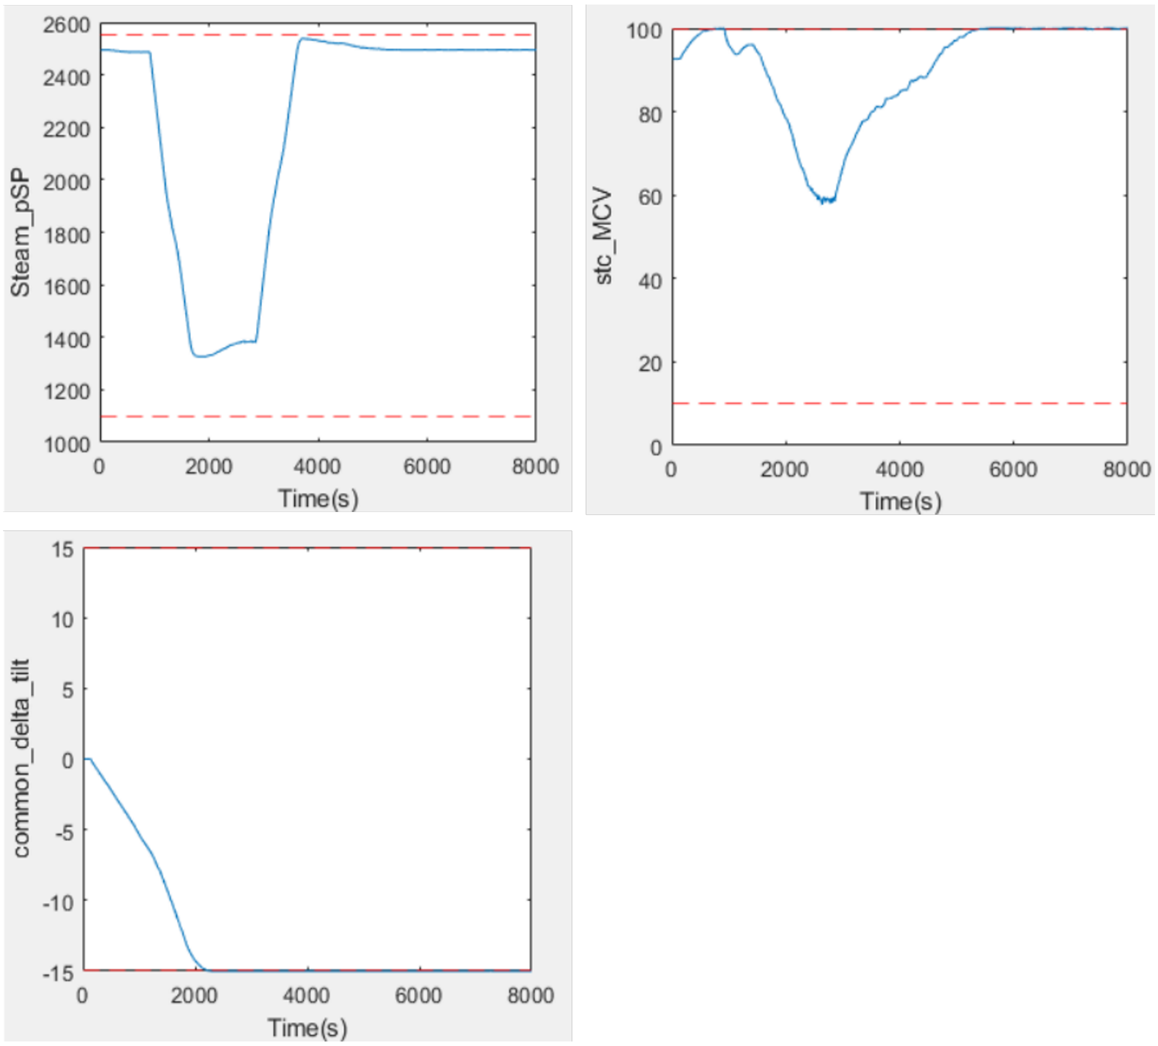

Figure 14: MPC input profiles for a fast load ramp at $4\%/min$ between $100 - 50\%$ TMCR, including a secondary objective for coal minimization.

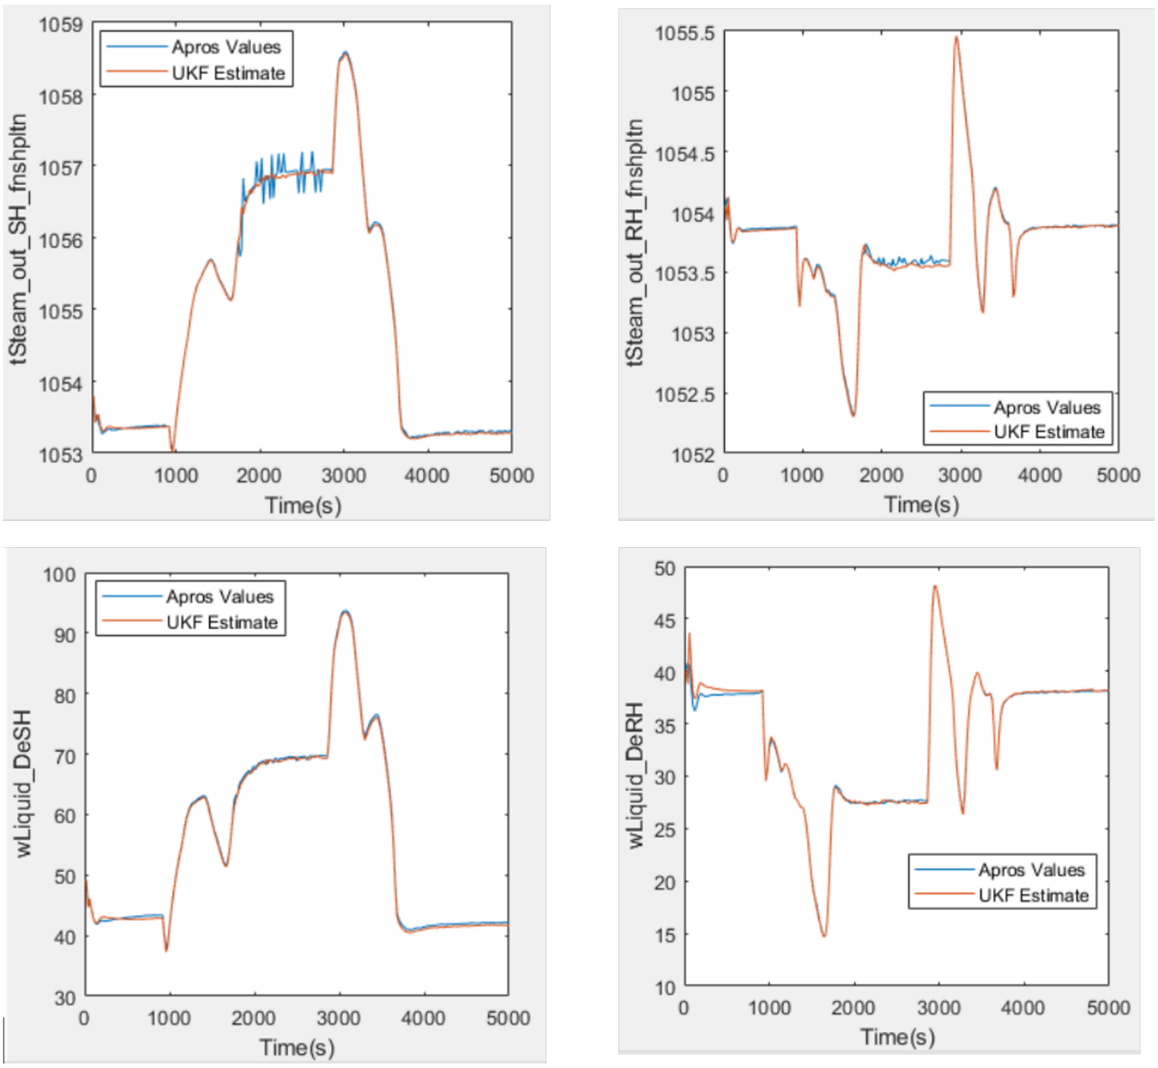

Figure 15: High pressure and Intermediate pressure steam temperatures and attemporator water flows for 100 − 50% TMCR load transient.

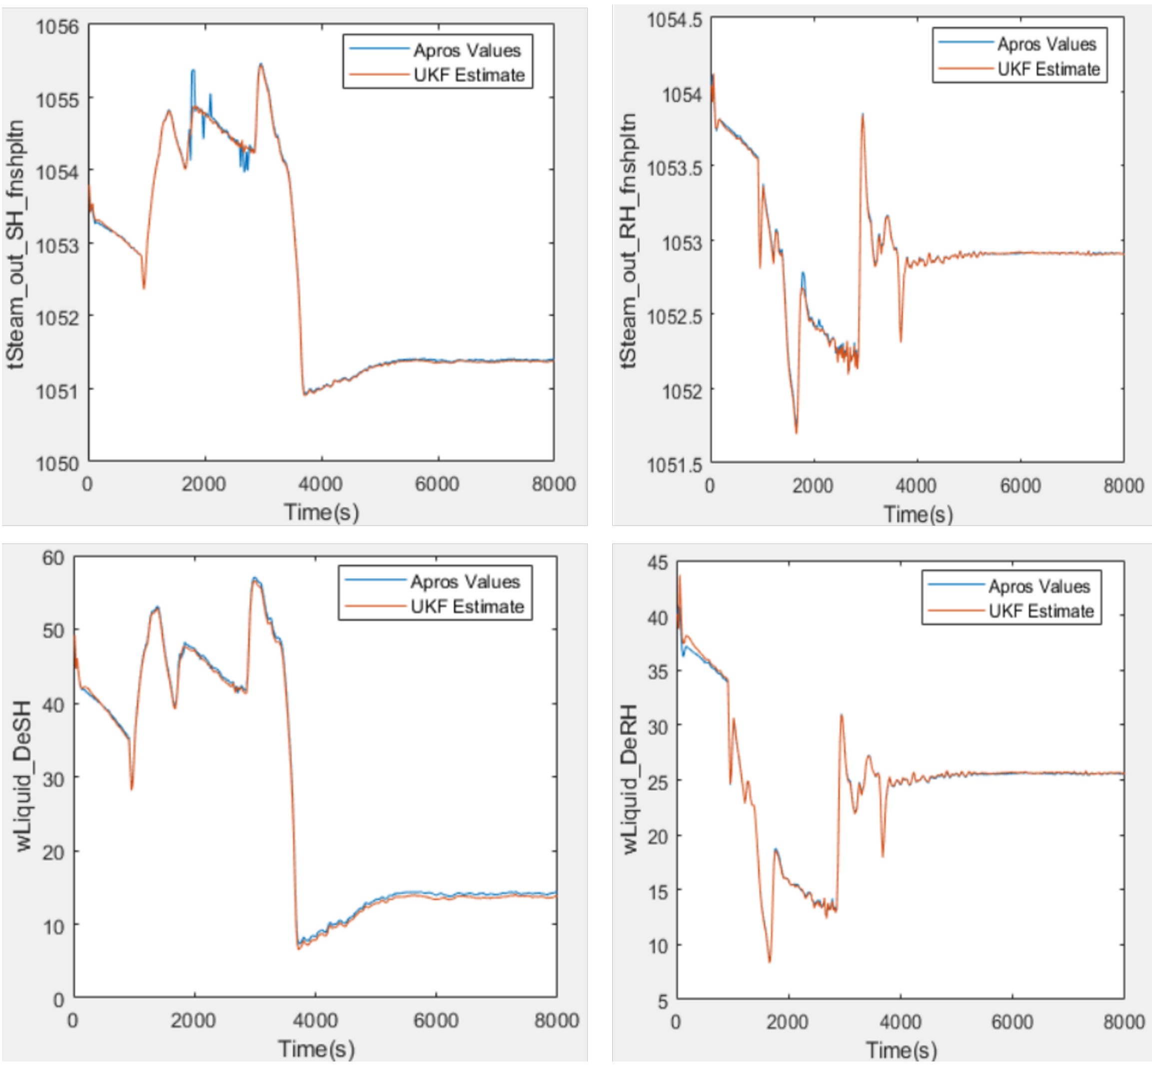

Figure 16: High pressure and Intermediate pressure steam temperatures and attemporator water flows for 100 – 50% TMCR load transient, including secondary objective for coal use minimization.

4.3.2 Real-time HIL simulations with container microservices

The Simulink implementation of ROM for plant model, MBE and MPC was then used to generate C-code using Mathworks' Matlab and Simulink coders, and then encapsulated in Docker microservice container to be deployed for real-time HIL simulation in an Ubuntu Linux VM. Figure 17 shows the transient profiles for both MPC outputs (load and coal use) and the three MPC inputs for the same load transient with a $4\%/min$ ramp rate between $100-50\%$ load, with explicit inclusion of the secondary objective for coal minimization. The real-time HIL deployment achieved the same exact results as with the desktop Simulink simulations in Figures 13-14. The entire simulation ran more than $3\times$ faster than real time despite running in a Linux VM on a Windows PC. A deployment on a native Linux OS would achieve even faster execution, well within the real-time requirements corresponding to the $20s$ sample time for these runs.

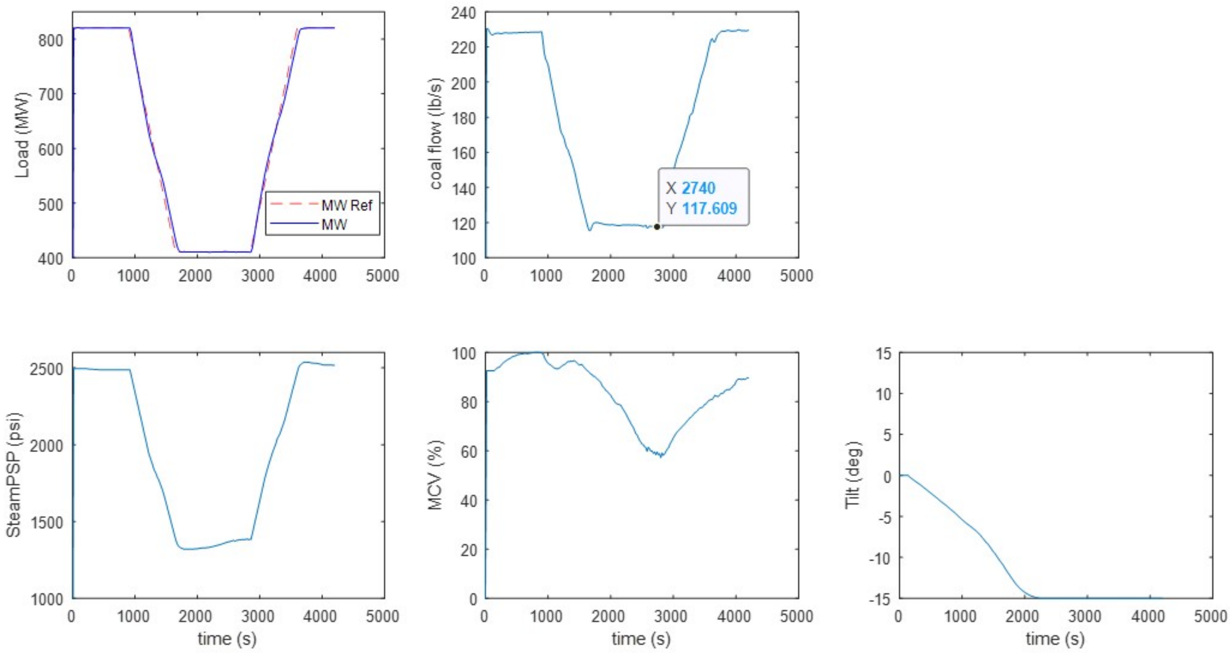

Figure 17: Transient profiles for MPC outputs and inputs for $100-50\%$ TMCR load transient in a real-time HIL deployment, including secondary objective for coal use minimization.

4.4 Closed-Loop Simulations with High-Fidelity Apros Model as Plant

This section summarizes the simulation results for the closed-loop test harness desktop simulations as well as real-time simulations with the HIL deployment using container microservices. For these simulations, MBE and MPC were enabled and the high-fidelity Apros model was used as the plant simulation in closed-loop - i.e., there is significant structural mismatch between the plant model and the embedded ROM used in MBE and MPC. This is to establish robustness of the MBE+MPC performance despite the presence of significant plant-model mismatch that is very likely when applying such model-based technologies in a real plant. For these simulations, three control inputs are enabled (i) steam pressure setpoint for the existing boiler controller that calculates coal and air flows, (ii) master control valve (MCV) for the high-pressure steam flow to the steam turbine, and (iii) common tilt angle for the coal and air feeds in the combustor - the tilt angle is the same for all three combustor zones as is currently implemented in plant controls (in future there could be further improvements by allowing independent positive or negative tilt in each combustor zone). The MPC objective includes a quadratic term for load tracking error, with a weight Q_Y, and a linear objective for coal minimization with a weight $L_{Y,perf}$. The weight $L_{Y,perf}$ is used to enable or disable coal minimization objective to see relative improvement.

4.4.1 Desktop simulations with Simulink test harness

The initial set of desktop simulations were performed with the Simulink test harness shown in Figure 7. The first set of runs were performed with $L_{Y,perf} = 0$, i.e., no explicit objective to minimize coal use, and focus only on tracking fast load ramps. In particular, a fast load ramp from $100 - 50\%$ TMCR and then back from $50 - 100\%$ TMCR was studied with a ramp rate of $3\% TMCR/min$. This is in contrast to the baseline controller and schedules for a nominal ramp rate of $2.5\%/min$. Also, in the first set of runs, the tilt angle control input was disabled by setting a very high weight R for that MPC input. Figure 18 shows the output profiles. The plant load (MW) follows the fast ramp in load setpoint (MWref) very well during the ramp down, and is slower than the $3\%/min$ setpoint ramp up. Also, there is small load transient at $1200s$ when the ramp down ends and setpoint changes to a constant 50% load. This is mainly due to a sudden jump in the coal flow, as can be seen in Figure 18. This jump in coal flow happens due to the existing baseline control logic for the boiler controller in the Apros model, designed to prevent a load overshoot/undershoot after a transient load ramp. A similar sudden drop in coal flow happens at $3300s$ when the ramp up finishes and the load setpoint stabilizes at 100% load, for the same logic avoid a load overshoot. However, this detailed control logic is not known to MPC which is a multivariable supervisory plant-level controller coordinating all control inputs to achieve a good transient tracking without undershoot or overshoot - this sudden unknown jump in coal feed acts like a disturbance. Despite this significant disturbance, MPC is able to recover from it once the measurements are used in MBE to provide updated state estimates. This transient performance can be improved by changing the baseline boiler controller logic to prevent such abrupt changes in coal feed and letting MPC perform the transient control to avoid overshoot/undershoot, since it is anticipating such load setpoint changes over the prediction horizon and will naturally change the pressure setpoint, MCV (and tilt when enabled) to achieve good tracking (see previous section on MPC results with ROM as the plant model).

Figure 19 shows the corresponding profiles for the three MPC inputs. The pressure setpoint changes the most as the load setpoint is changed between 100% and 50%, almost reaching the minimum limit, while MCV changes relatively less, closing down to about 35% at 50% load, and tilt doesn't change since it is disabled with the high R in this run.

The above run with two enabled MPC inputs, steam pressure setpoint, and MCV, to track setpoint changes in the main output, plant gross electric output, were replicated in the HIL deployment to demonstrate that it can be deployed in real-time on the HIL platform and produce the same results, as seen in Section 4.4.2. Thereafter, additional simulation runs with all three MPC inputs (including common feed tilt angle) enabled and varying

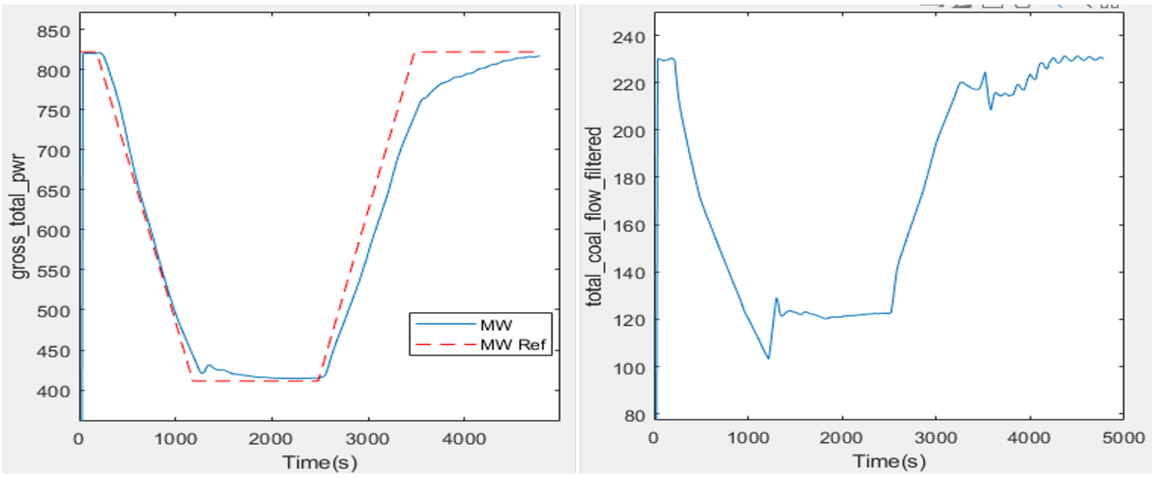

Figure 18: Load and coal feed profiles for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR with no coal minimization objective.

the weight $L_{Y,perf}$ for the additional performance output, i.e., minimizing coal use. The results from these three simulation runs with $L_{Y,perf} = 0$ (i.e, no performance output), and $L_{Y,perf} = 0.8$ (moderate weight on performance output), and $L_{Y,perf} = 3.0$ (high weight on performance output) are shown below.

Figures 20-21 show the two controlled outputs and three control inputs for the first run with $L_{Y,perf} = 0$, i.e. no emphasis on performance output, and only focused on load tracking. As seen from the load output plot, the overall load tracking is good for ramp down and up at $3\%/min$ ramp rates, except for the transients at the end of the ramp down, and the slower load increase at the end of ramp up - again these are caused by the same jumps in coal feed in the baseline controller in Apros designed to prevent overshoots at the end of the ramp down/up. The corresponding control input profiles are also shown. Note that in this run with zero emphasis on performance output, there are multiple combinations of these three inputs that achieve the single tracking objective. In particular, the tilt angle increases to a positive value by the end of the run. For reference, the coal use at 50% load steady-state point is about $122.95 lb/s$.

Figures 22-23 show the two controlled outputs and three control inputs for the second run with $L_{Y,perf} = 0.8$, i.e. with a moderate emphasis on performance output, in addition to load tracking. As seen from the load output plot, the overall load tracking is good for ramp down and up at $3\%/min$ ramp rates, except for the transients at the end of the ramp down, and the slower load increase at the end of ramp up. Now that the performance objective for coal feed reduction is also included, the sudden jumps in coal use due to the baseline controller in Apros, especially the sudden increase at the end of ramp down causes MPC to respond to this unknown disturbance. With the dual objective on minimizing coal (performance) and load tracking, there is an increased transient in the load tracking due to this disturbance. The corresponding control input profiles are also shown. With the inclusion of a moderate emphasis on coal feed minimization, the tilt angle increases to a small positive value at the end of the run compared to the first run described above. Also, in comparison to the first baseline run, the coal use at 50% load steady-state point is about $120.8 lb/s$, which is about 1.7% relative reduction in coal use, or equivalently in plant efficiency, compared to the first

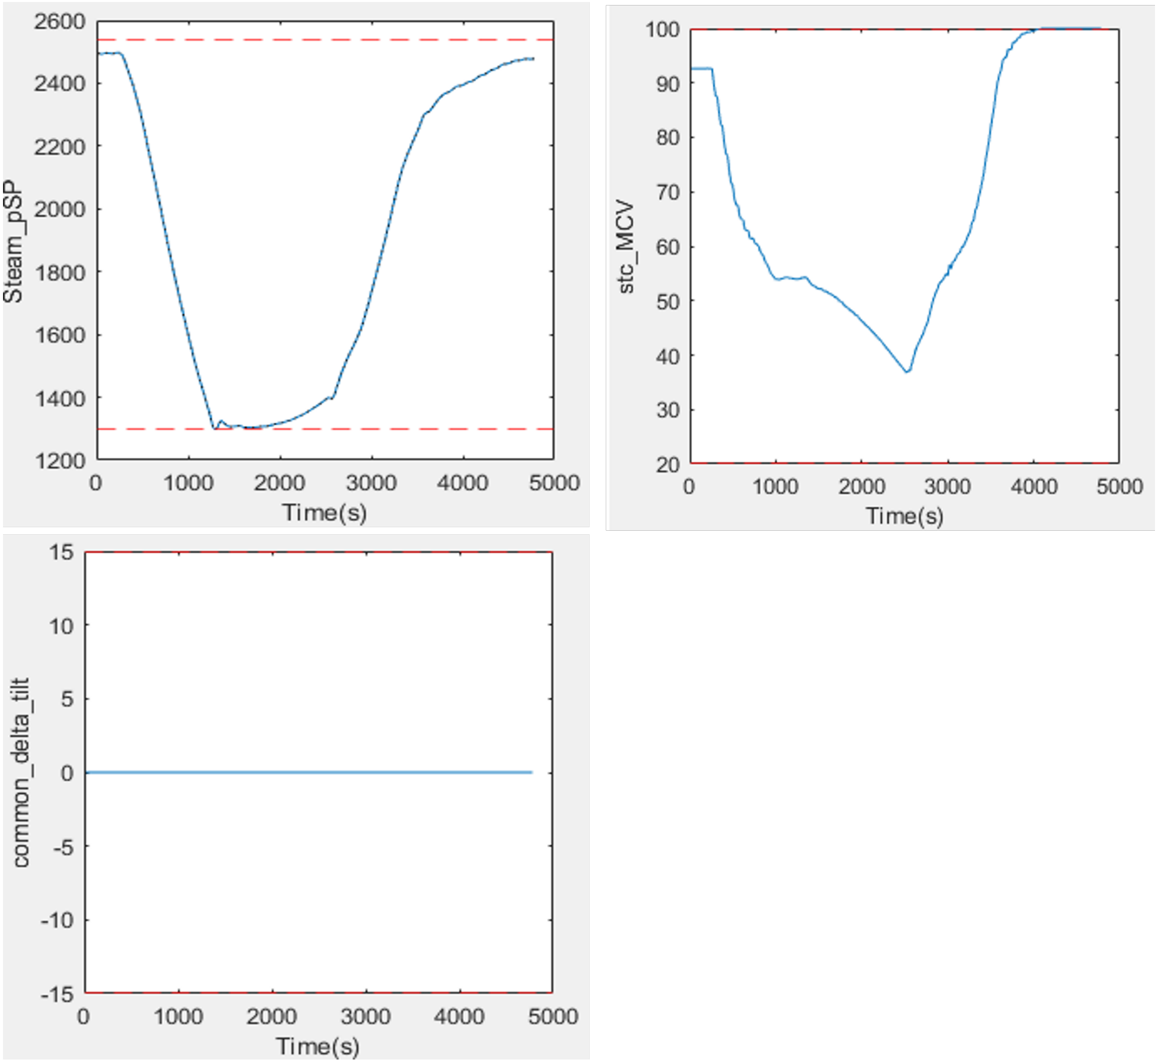

Figure 19: MPC input profiles for a fast load ramp at $3\%/min$ between $100-50\%$ TMCR with no coal minimization objective.

MPC run.

Figures 24-25 show the two controlled outputs and three control inputs for the last run with $L_{Y,perf} = 3.0$, i.e. with a higher emphasis on performance output, in addition to load tracking. Again, the overall load tracking is good for ramp down and up at $3\%/min$ ramp rates, except for the transients at the end of the ramp down, and the slower load increase at the end of ramp up. Similar to the previous case, the undesired transient in load at the end of the ramp is worsened due to the unknown disturbance in the coal feed jump from the baseline controller in Apros - with the increased emphasis on minimizing coal (performance), there is a larger transient in the load tracking due to this disturbance. The corresponding control input profiles are also shown. With the higher emphasis on coal feed minimization, the tilt angle actually now decreases to a negative value - this is expected as with the ROM based plant simulations. The negative tilt angle reduces the flue gas temperature from the top combustion zone, which in turn reduces the superheating of the steam and reduces the

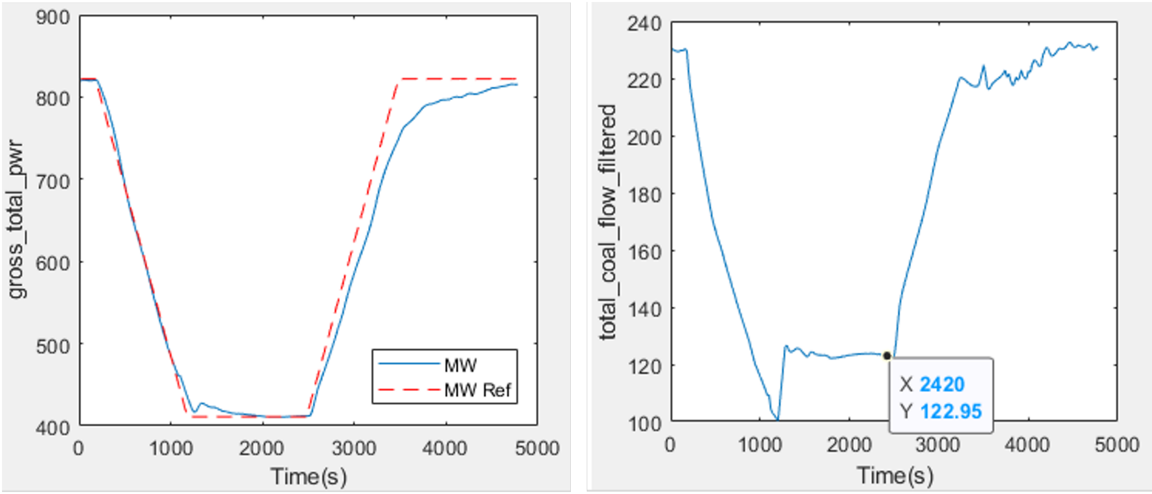

Figure 20: Load and coal feed profiles for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR with no coal minimization objective ($L_{Y,perf} = 0$), using all 3 control inputs.

attemporation water spray, as discussed in Section 4.3.1. Also, in comparison to the first baseline run, the coal use at 50% load steady-state point is about $116.61 lb/s$, which is about 5.0% relative reduction in coal use, but due to the increased transient in load at partload operation, this is not an accurate measure and likely an over-estimate.

Despite very significant plant-model mismatch (deliberate choice to emulate a real-world condition) between the high-fidelity Apros model used for plant simulation, and a simpler ROM used in MBE and MPC, these simulations show great results in transient tracking of load even for fast ramp rates up and down at $3\%/min$, while also reducing the coal feed (or equivalently increasing plant efficiency) at 50% TMCR load. The results, in particular, load tracking, can be further improved if (i) the baseline controller in Apros were to be modified to avoid the sudden jumps in coal feed (designed to avoid load overshoots), and (ii) MBE could be further tuned for improved transient tracking. In particular, in light of these results, the MBE tuning was updated to track key measured outputs (e.g., coal and air flows, steam flows, temperature and pressure, MW output) more accurately. This allows improved tracking of the transient variation from the Apros high-fidelity plant model, specifically the transient jumps in coal flow at the end of a transient load ramp as seen in the next set of simulations.

Figures 26-27 show the transient profiles for the load and coal flow and the three control inputs for load ramps between $100\% - 50\%$ at ramp rates of $3\%/min$. With the improved MBE tuning, note that the transient variation in controlled load is significantly reduced (compare with Figure 20). Also, the coal flow at 50% load steady-state is reduced to $116.3 lb/s$ which is about 5.4% relative reduction in coal flow, and a corresponding improvement in efficiency.

For the same simulation run, Figure 28 shows the comparison of key measured outputs from Apros model and matched outputs from ROM. Despite significant structural differences between the high-fidelity Apros model used for plant, and the simple ROM, MBE is able to track transient variations in these key outputs (e.g., steam pressures and temperatures, and flows of coal, air and steam and the gross MW output) very well. The corresponding

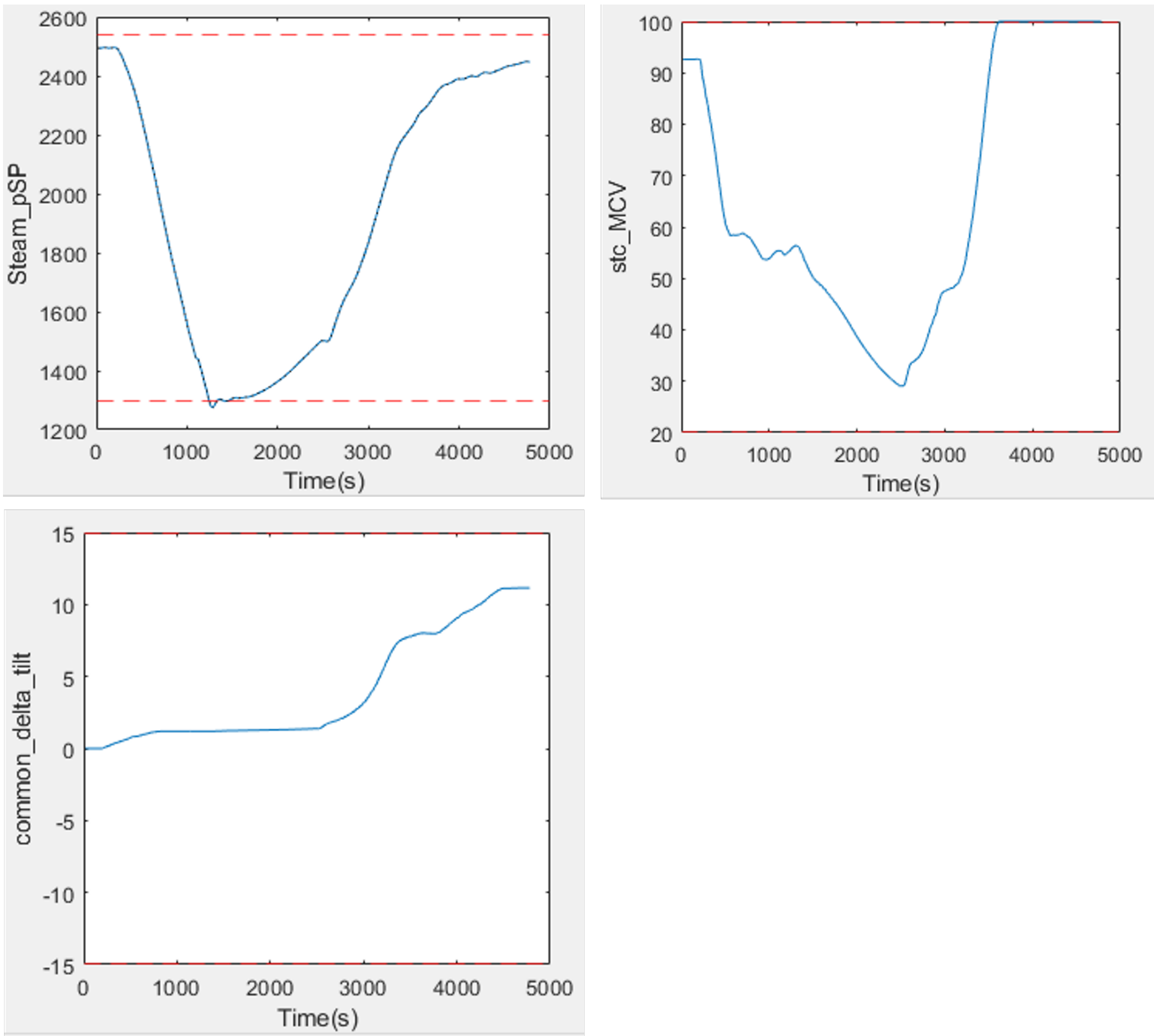

Figure 21: MPC input profiles for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR with no coal minimization objective ($L_{Y,perf} = 0$), using all 3 control inputs.

profiles for a subset estimated parameters in key components in ROM that are used to match the ROM to Apros model are shown in Figure 29. These estimated health parameters for key sections (e.g., heat transfer coefficients in combustor, superheater and reheater) can be trended over time to monitor the corresponding equipment health (e.g., fouling of the heat exchangers) and enable condition-based maintenance. While many of the estimated parameters are essentially constant, some of them, specifically parameters for the IP steam section vary systematically with the operating load - indicating a systematic structural mismatch between the ROM and Apros across load, that is compensated for by the variation in these adapted parameters.

The final set of simulations studied the performance for the developed solution for a faster load ramp at $4\%/min$ between $100-50\%$ TMCR. As shown in Figures 30-33, again it is able to follow the fast load ramp, especially ramp down, very well and yields a similar 5.7% relative reduction in coal flow at 50% load steady-state compared to the baseline MPC run

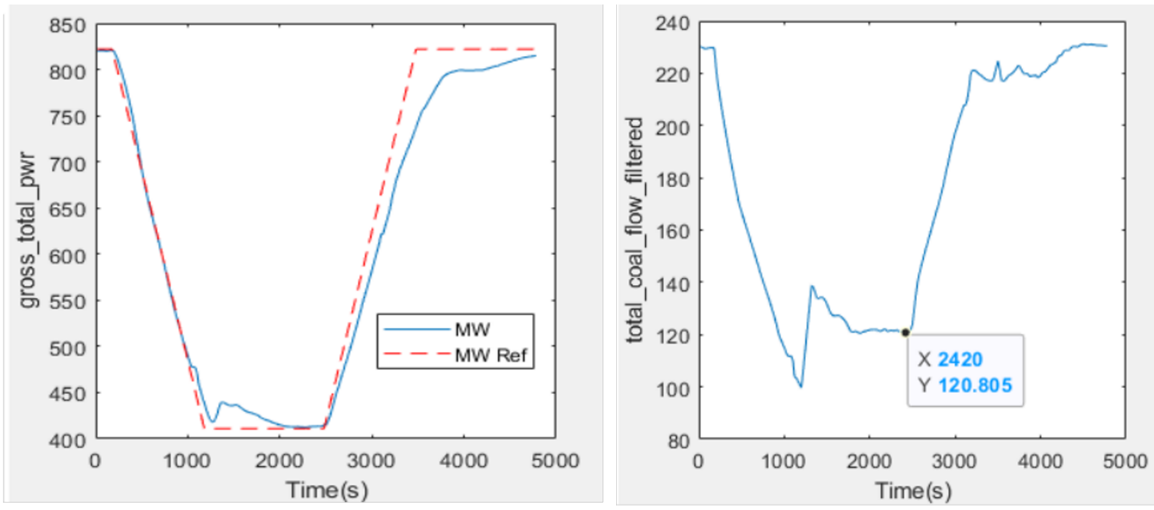

Figure 22: Load and coal feed profiles for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR with moderate emphasis on coal feed minimization objective ($L_{Y,perf} = 0.8$), using all 3 control inputs.

with no emphasis on coal flow reduction (Figure 20). The one aspect that could be further improved is the transient load tracking during the ramp up - by further tuning the MBE and, more importantly, updating the baseline boiler controller in Apros to avoid transient jumps in coal flow at the end of the ramp down/up.

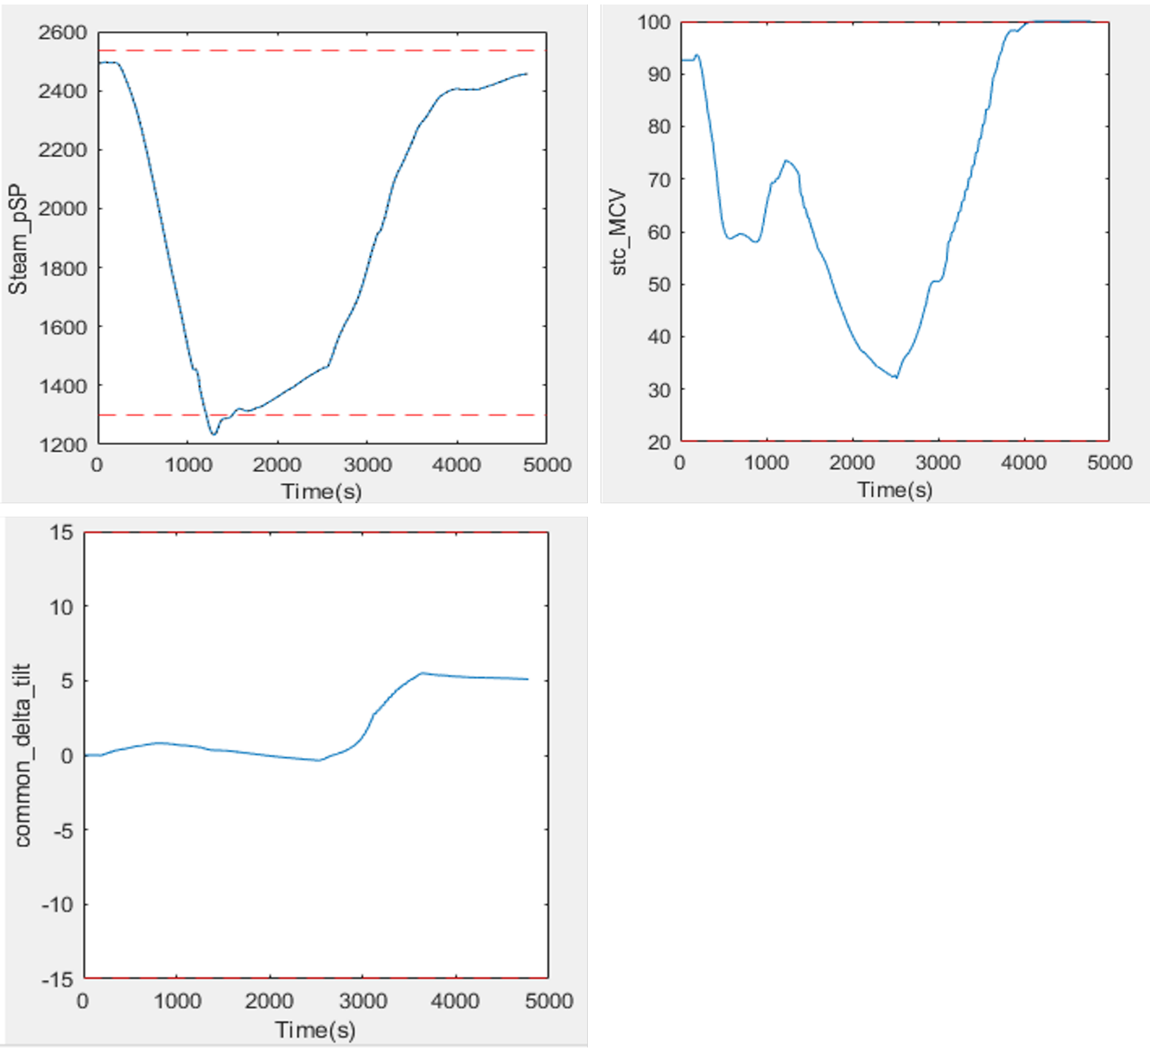

Figure 23: MPC input profiles for a fast load ramp at $3\%/min$ between $100-50\%$ TMCR with moderate emphasis on coal feed minimization objective ($L_{Y,perf} = 0.8$), using all 3 control inputs.

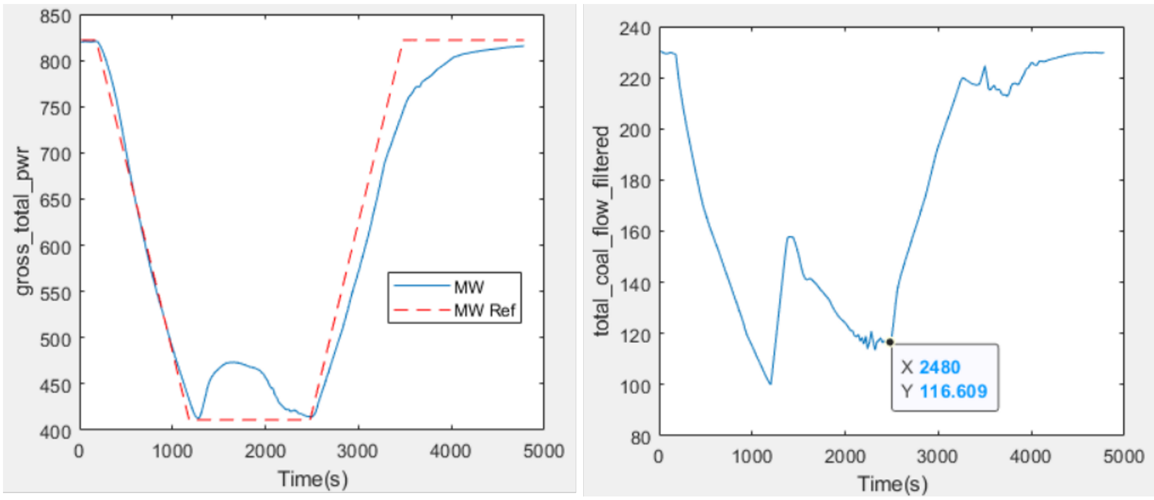

Figure 24: Load and coal feed profiles for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR with high emphasis on coal feed minimization objective ($L_{Y,perf} = 3.0$), using all 3 control inputs.

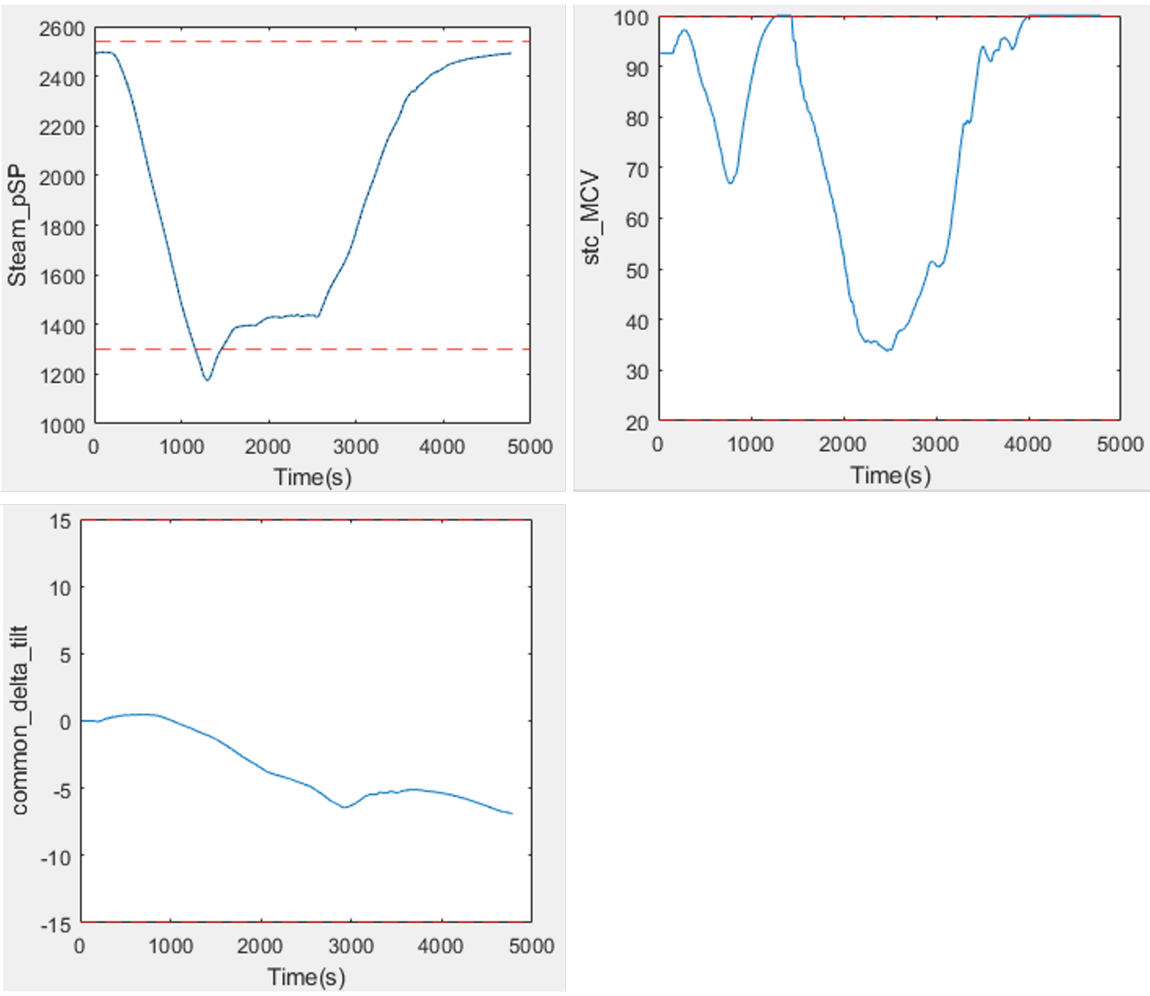

Figure 25: MPC input profiles for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR with high emphasis on coal feed minimization objective ($L_{Y,perf} = 3.0$), using all 3 control inputs.

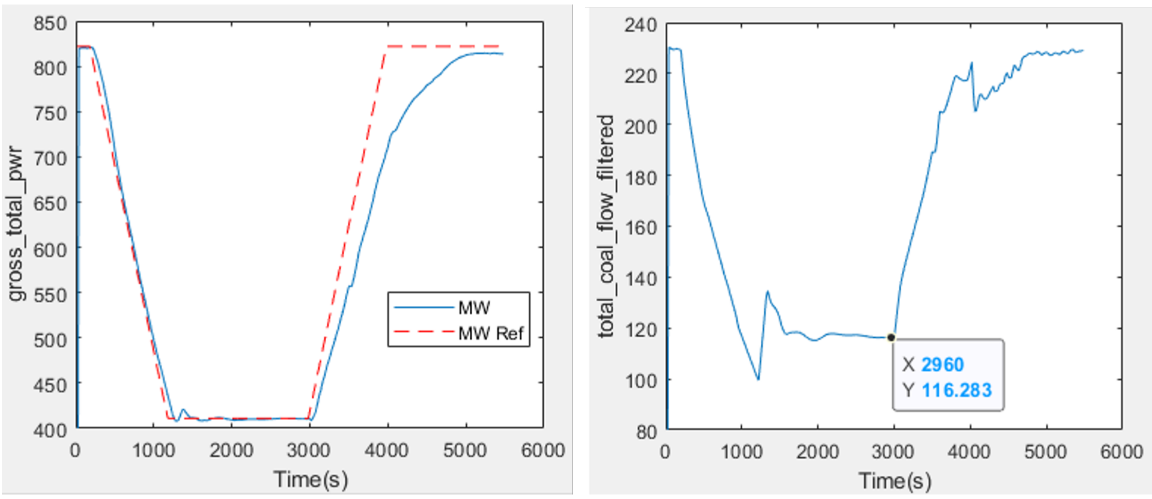

Figure 26: Load and coal feed profiles for a fast load ramp at $3\%/min$ between $100-50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

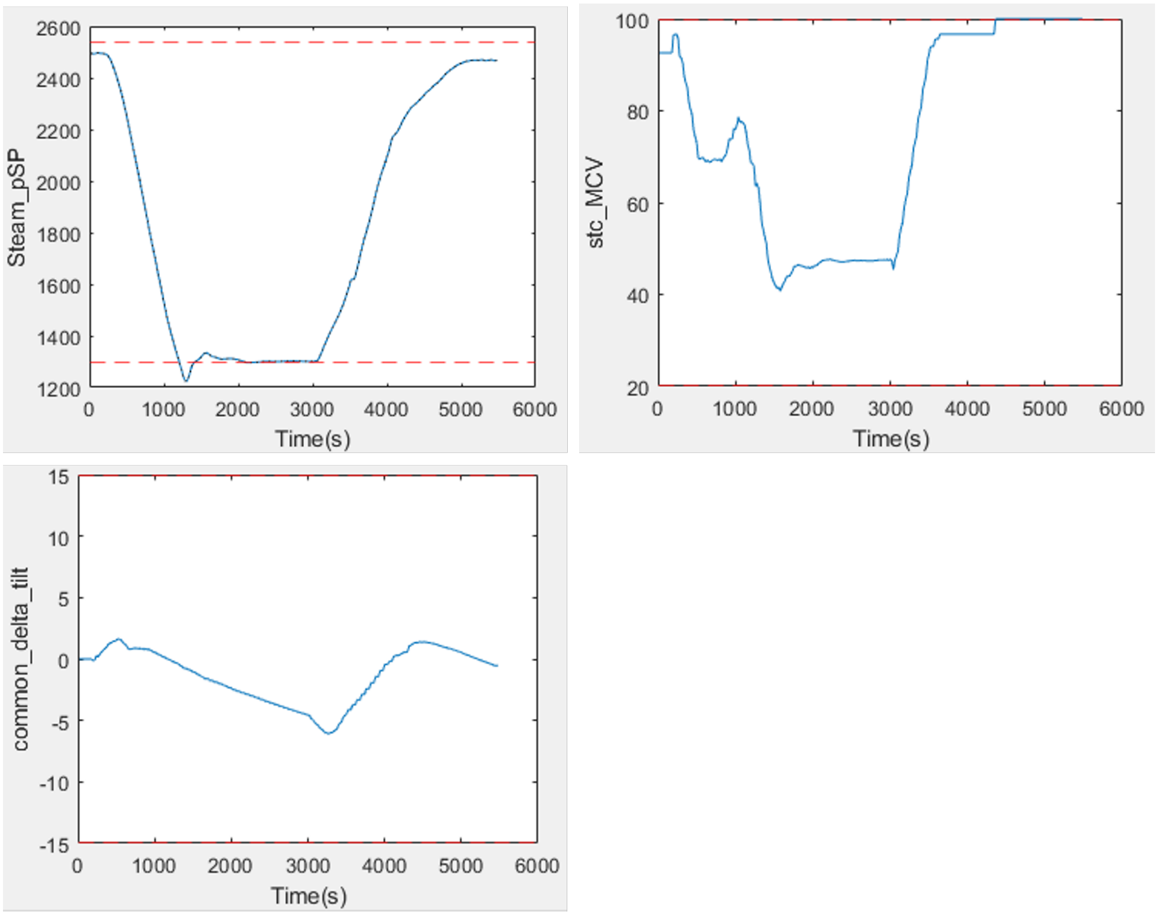

Figure 27: MPC input profiles for a fast load ramp at $3\%/min$ between $100-50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

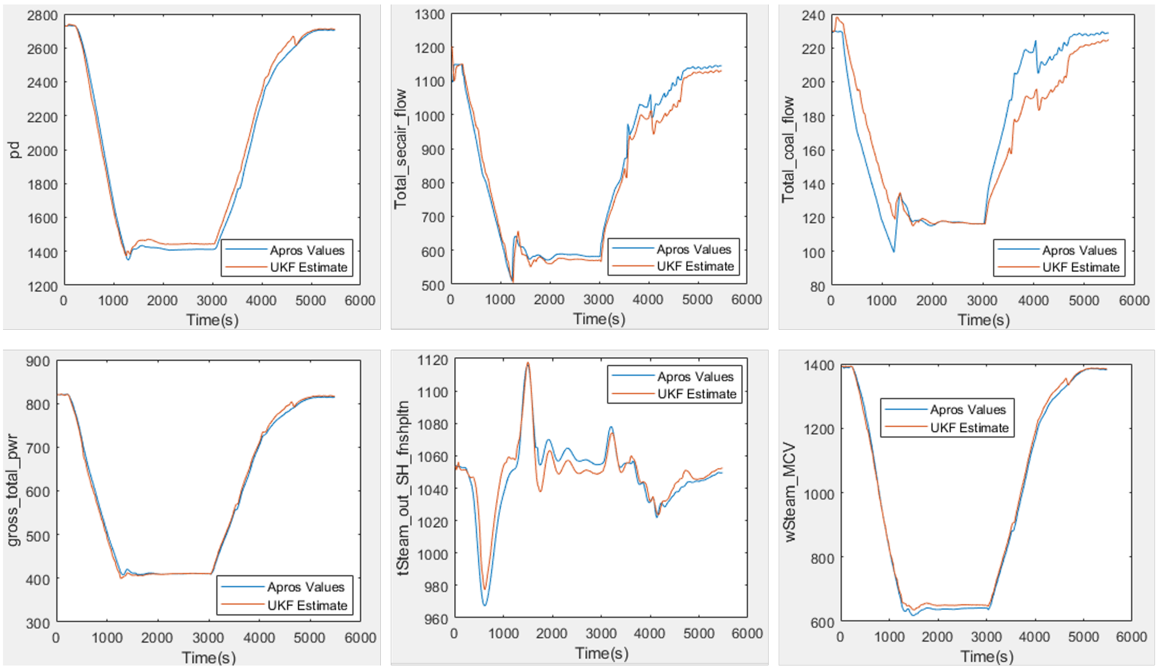

Figure 28: Tracking of key measured outputs by MBE for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

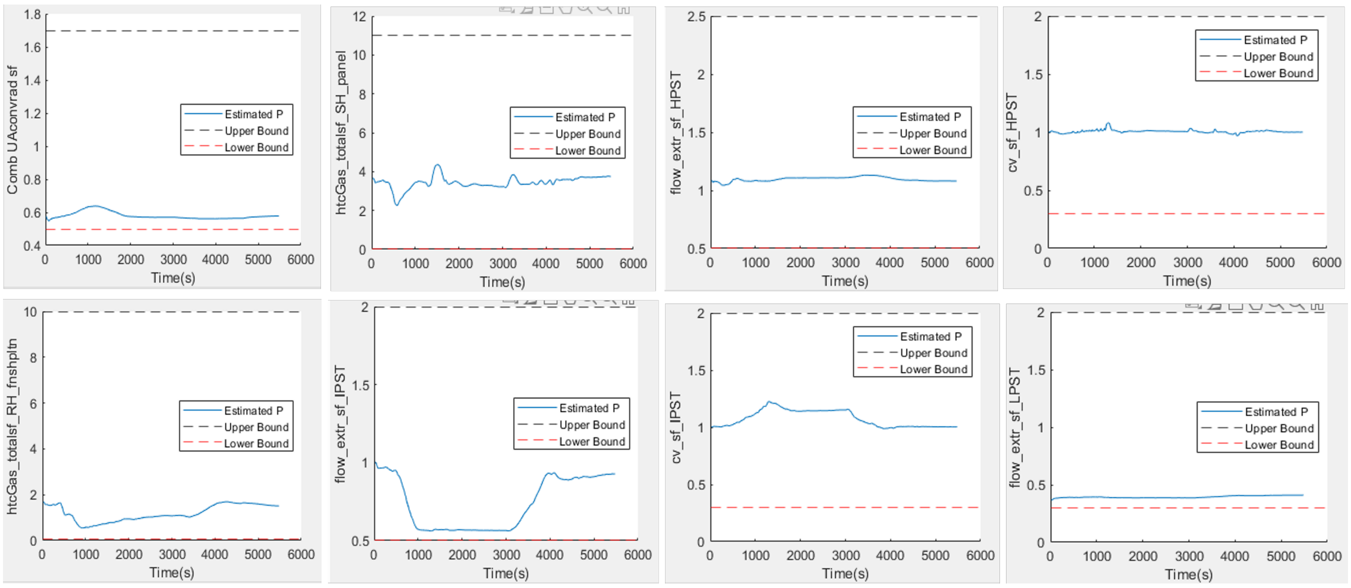

Figure 29: Profiles for MBE parameter estimates for key sections for a fast load ramp at $3\%/min$ between $100 - 50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

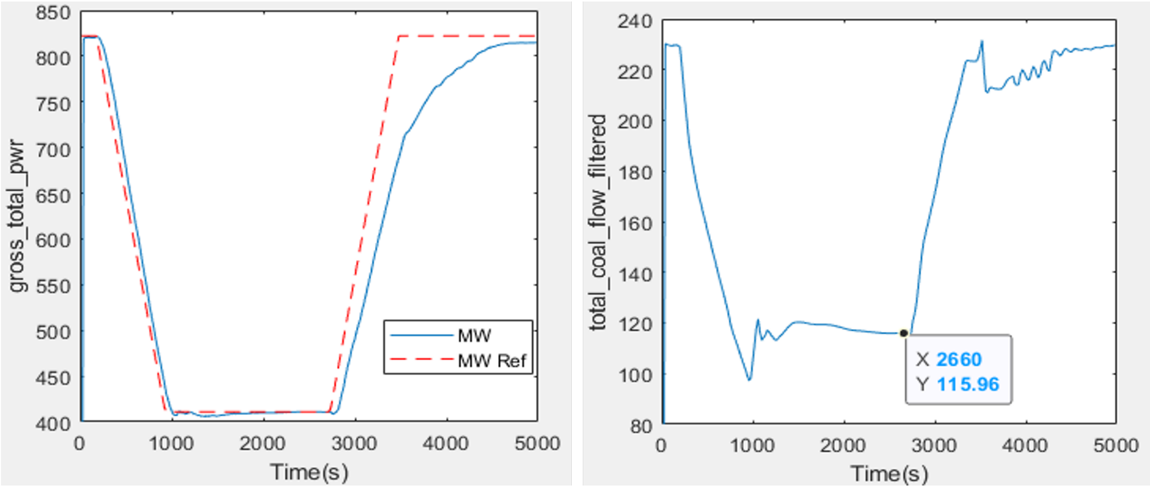

Figure 30: Load and coal feed profiles for a fast load ramp at $4\%/min$ between $100 - 50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

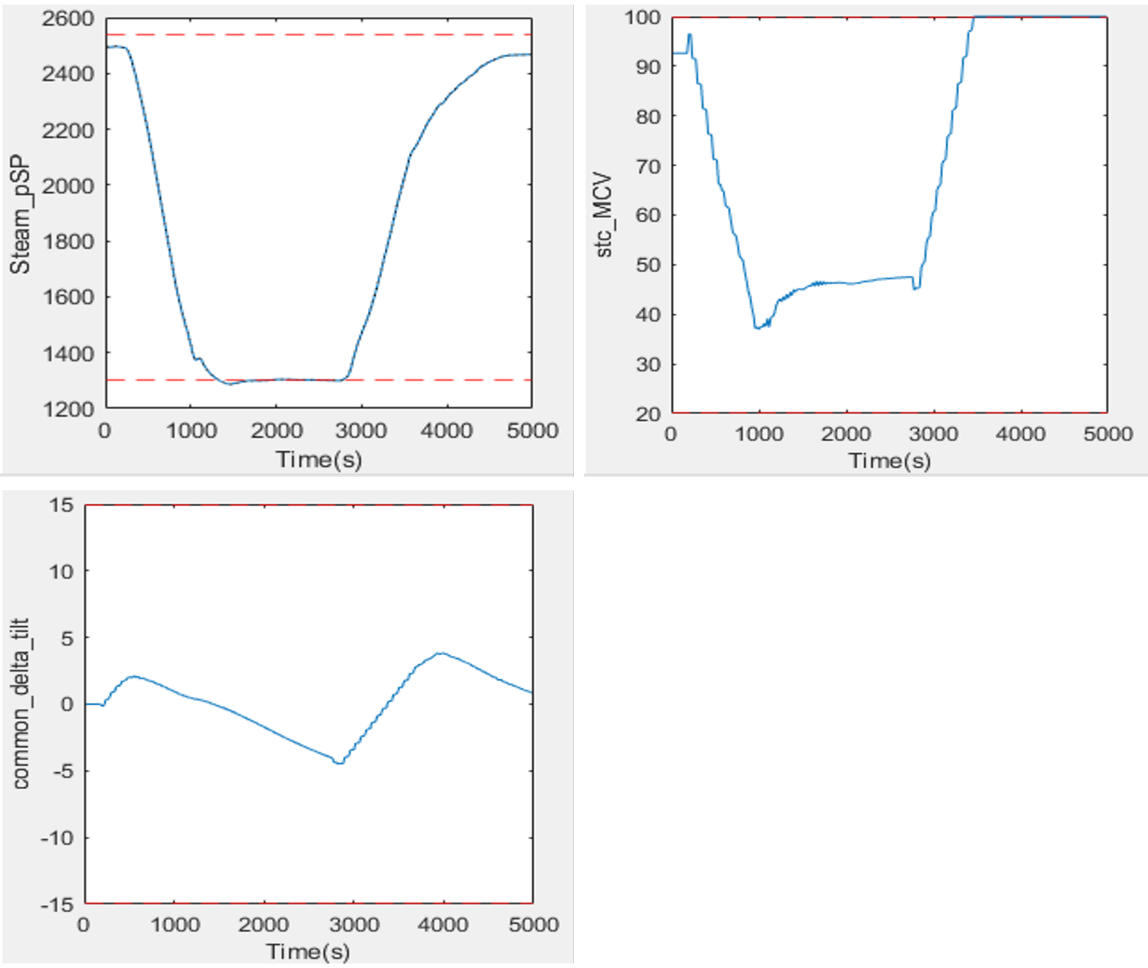

Figure 31: MPC input profiles for a fast load ramp at $4\%/min$ between $100 - 50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

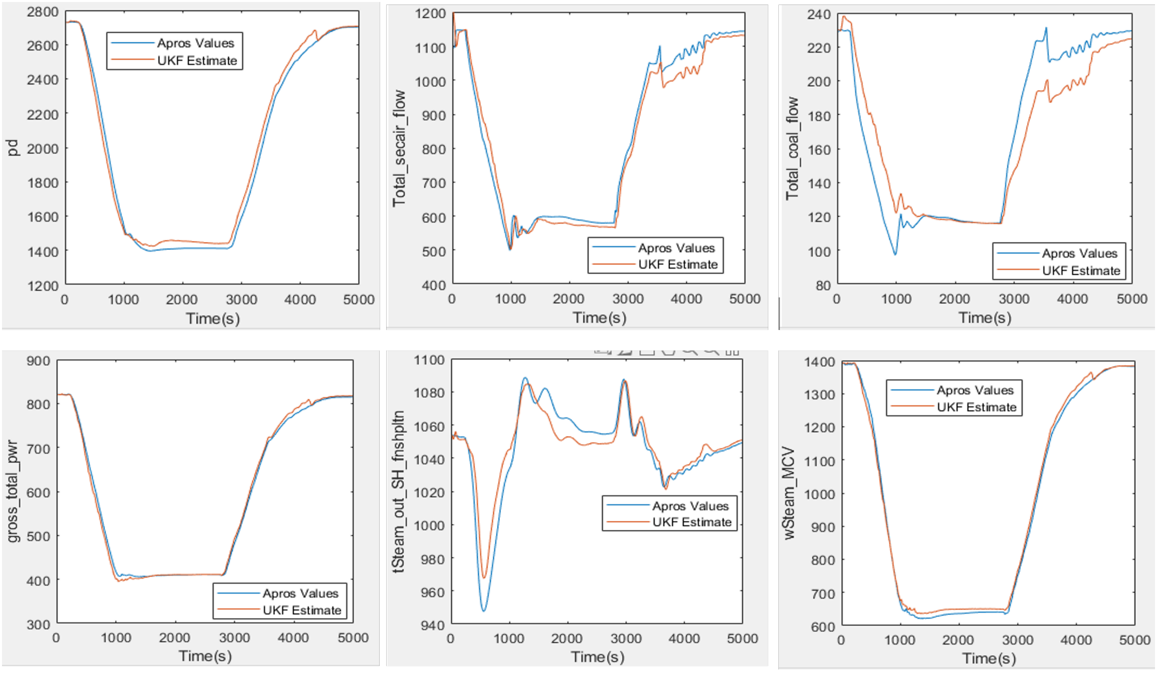

Figure 32: Tracking of key measured outputs by MBE for a fast load ramp at $4\%/min$ between $100 - 50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

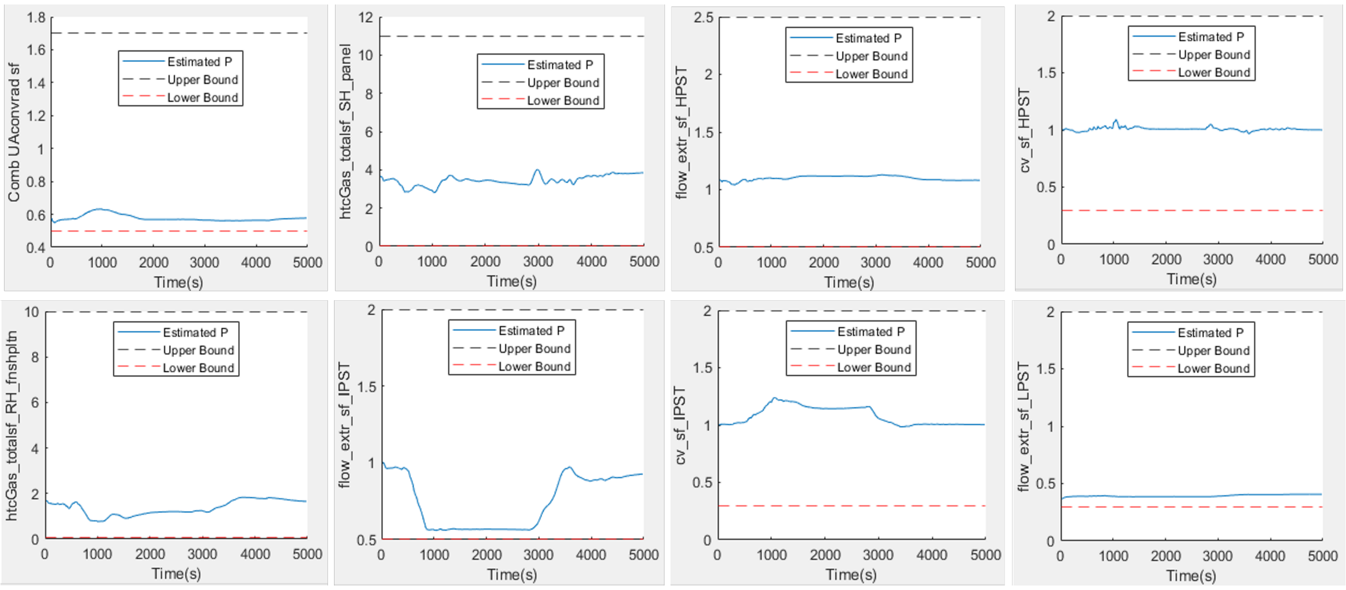

Figure 33: Profiles for MBE parameter estimates for key sections for a fast load ramp at $4\%/min$ between $100 - 50\%$ TMCR using all 3 control inputs, and improved MBE tuning.

4.4.2 Real-time HIL simulations with container microservices

The Simulink implementation MBE and MPC was then used to generate C-code using Mathworks' Matlab and Simulink coders, and then encapsulated in Docker microservice container to be deployed for real-time HIL simulation in an Ubuntu Linux VM. The deployed MBE+MPC algorithm in Linux VM used a real-time two-way communication with the high-fidelity Apros model running on the same PC in Windows. Despite both the deployed algorithms and the high-fidelity Apros model running on the same PC, the closed-loop HIL simulation ran easily in real-time, without any violation of real-time computation limits. Figure 34 shows the transient profiles for both MPC outputs (load and coal use) and the three MPC inputs for the same load transient with a $3\%/min$ ramp rate between $100-50\%$ load. The real-time HIL deployment achieved the same results as with the desktop Simulink simulations in Figures 18-19. The very small differences (e.g., in coal profile towards the end of the run) occur due to slight jitter in the real-time communication with Apros model with the Apros API, which is not really meant for a precise real-time communication.

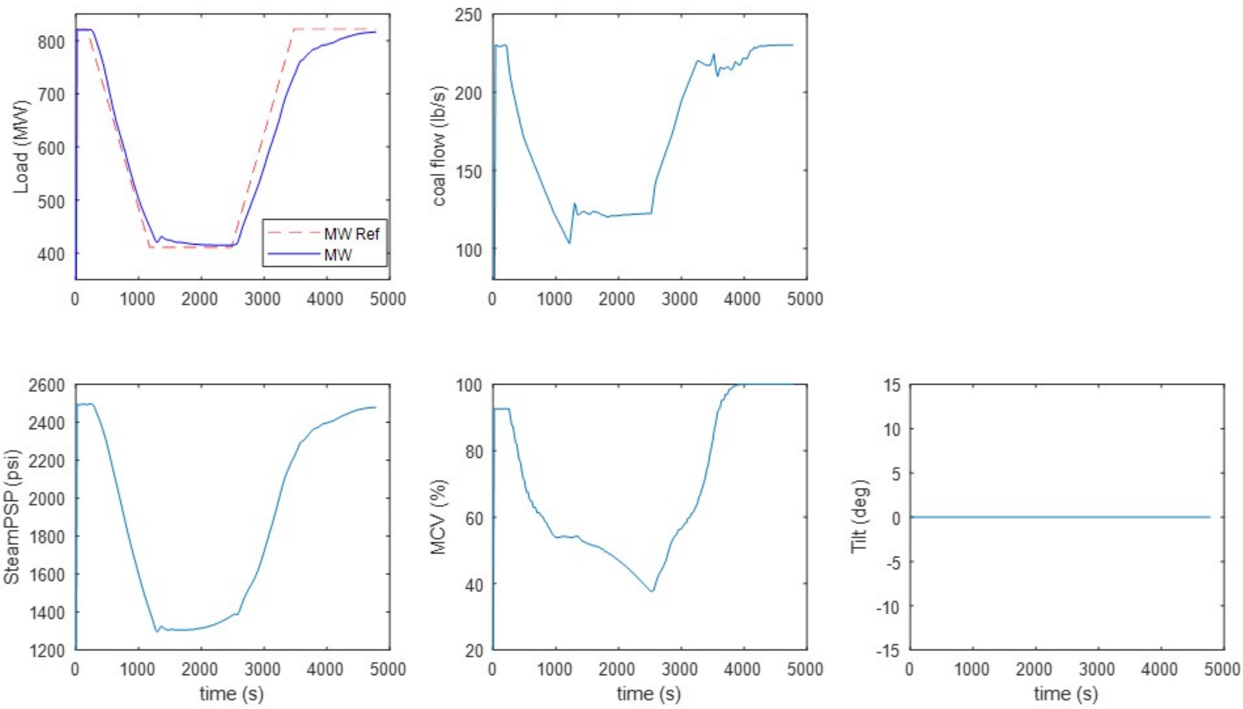

Figure 34: Transient profiles for MPC outputs and inputs for $100-50\%$ TMCR load transient in a real-time HIL deployment.

The HIL results are the same as the ones obtained by the desktop simulation in the test harness, establishing that the deployed ROM+MBE+MPC software can run in the target Linux OS in real-time (it actually runs about $3\times$ faster than real-time despite running both the APROS model and the control software on a Linux VM on the same Windows PC, and will likely run even faster on a real Linux hardware instead of a VM), and work well with two-way real-time communication with the Apros model used for high-fidelity plant simulation.

5 Conclusions

This program successfully achieved development and application of model-based estimation and control for a subcritical coal fired power plant that enables a faster load transient for improved flexibility. High-fidelity plant model-based simulations demonstrated tracking up to $4\%/min$ load transients in a $820MW$ plant. This is important as these plants cycle more frequently between baseload and partload, with increasing fluctuation in power generation on the grid due to increasing renewables. This faster load transient is achieved simultaneously with a reduction in coal use, or equivalently plant efficiency improvement, achieving a $3.8 - 5.5\%$ relative improvement at 50% load. The integrated closed-loop MBE+MPC simulations were performed both in Simulink desktop simulations as well as in a HIL real-time deployment using container microservices in a Linux VM. The performance of the developed solution was tested against a high-fidelity transient model of CFPP implemented in Apros, with two-way real-time communication between the MBE+MPC algorithms and the Apros model. Despite significant plant-model mismatch (between the high-fidelity plant model in Apros, and the reduced order model used as embedded model for MBE and MPC), the solution achieved good transient performance. Further improvements are possible by updating the design of the lower-level boiler controller to work synergistically with the optimized plant-level supervisory control from MPC. In the current studies, the lower level boiler controllers were kept as the baseline design with existing schedules and control logic, which worked well, but led to some abrupt transients in coal feed during load changes, that weren't anticipated by MPC and acted like a transient disturbance.

The developed ROM and MBE, MPC algorithms are implemented in Simulink in a manner conducive for auto C-code generation using Mathworks' Matlab and Simulink coder, and encapsulation of the resulting C-code in Docker microservice containers. While the desktop simulations achieved 2× faster than real time, largely owing to a very fast ROM used as embedded model that is capable of more than 100× faster than real-time simulation, it ran about 3× faster than real time in the final HIL, despite the fact that it was implemented in a Linux VM in a Windows PC. It would run even faster, e.g., 4× or more, than real time, if the eventual deployment is done on a dedicated PC with native Linux OS instead of a VM. Thus, this program successfully demonstrated the goals for the program. In particular, the final set of simulation runs (both in desktop harness and HIL) using high-fidelity Apros model as the plant simulation, showed the robustness of the overall ROM+MBE+MPC solution developed and implemented as shown in Figure 2 inspite of significant plant-model mismatch between the high-fidelity Apros model for plant and the much simpler (and faster) ROM used in MBE and MPC. This gives confidence that despite significant plant-model mismatch in a real plant applications (all models are only approximations of a real system), the proposed solution can work very well to achieve the objectives of fast load ramps and increased plant efficiency. The simulation results shown in Section 4.4 demonstrate this capability.

The results of this program demonstrate the viability of the proposed model-based estimation and control technology for subcritical coal-fired power plants towards achieving flexible and efficient operation of these plants across the load range as these plants will cycle more often. The extensive desktop and HIL simulations in the presence of significant plant-model mismatch demonstrate the robustness of the solution, matured to TRL5. It is desirable to pursue further maturation of these technologies to TRL7, wherein the developed technology is tested at a beta test site in a potential follow-on program. That maturation can incrementally pursue (i) MBE demonstration for real time process monitoring, (ii) MBE + MPC in advisory mode - to recommend to operators from small/slow transients, (iii) fully automated MBE+MPC (with appropriate safeguards) to achieve best transient flexibility and efficiency results.

In another extension, the developed closed-loop fast simulation models can be used to generate data for AI/ML algorithms, e.g., for fault diagnostics, cyber-physical security. Also, generalizations of modeling techniques combining complementary simplified physics-based modeling and data-based modeling with limited data, is another interesting research direction that can address practical challenges in modeling complex systems.

References

[1] "Transient Efficiency Flexibility and Reliability Optimization of Coal-Fired Power Plants - Report on High-Fidelity Dynamic Modeling of a Coal-Fired Steam Power Plant," October 2021.

[2] "Transient Efficiency Flexibility and Reliability Optimization of Coal-Fired Power Plants - Report on Reduced Order Model Development," October 2020.

[3] "Transient Efficiency Flexibility and Reliability Optimization of Coal-Fired Power Plants - Report on Model Based Estimation Library Development," January 2021.

[4] "Transient Efficiency Flexibility and Reliability Optimization of Coal-Fired Power Plants - Report on Model-Predictive Control Library Development," March 2022.