ECS Transactions, 25 (2) 1175-1184 (2009)
10.1149/1.3205646 ©The Electrochemical Society

Physically Based Model-Predictive Control for SOFC Stacks and Systems

Tyrone L. Vincent, Borhan Sanandaji
Andrew M. Colclasure, Huayang Zhu, and Robert J. Kee

Engineering Division, Colorado School of Mines, Golden, CO 80401, USA

This paper discusses model-predictive controllers (MPC) that can incorporate physical knowledge of fuel-cell behavior into real-time multiple-input–multiple-output (MIMO) process-control strategies. The controller development begins with a high-fidelity, transient, physical model that represents the physical and chemical processes responsible for fuel-cell function. However, because such large nonlinear models cannot be solved in real time as part of the controller logic, linear reduced-order state-space models are required. The model reduction is accomplished via a process called system identification. The controller is designed to interpret sensors in the context of the reduced-order model and determine optimal actuation sequences that cause the system to follow a desired output trajectory. The process is demonstrated for a tubular SOFC stack that could be used for auxiliary-power unit (APU) applications.

Introduction

Especially for small portable-power applications, solid-oxide fuel cells (SOFC) must deliver power profiles that meet the demands of transient loads. Consider, for example, an auxiliary power unit (APU) that is designed to satisfy the hotel loads for the sleeper cab on a long-haul truck. Depending upon the activities and appliances in the cab, the power demanded from the APU can vary considerably over time. Although such systems are usually designed with battery storage and power electronics that seek to limit the transients demanded from the SOFC, the fuel cell must still respond to transient load demands.

Any control system depends upon *actuation* and *sensing*. SOFC actuation is usually accomplished through a combination of operating voltage, fuel flow rate, and air flow rate. Sensing is usually accomplished with thermocouples that measure temperature and possibly gas analyzers that can measure chemical composition for flow streams. In addition, there are usually practical *constraints* on actuation and on system performance. For example, the operating voltage (actuation) may be constrained to stay above 0.6 V. The cell temperature (an observable output as measured by a sensor) may be constrained to stay within a specified range.

Model-predictive control (MPC) provides a means to incorporate quantitative physical understanding into real-time process-control decisions. An SOFC is a complex nonlinear system, with a widely disparate range of characteristic time scales. Moreover, there are multiple actuation possibilities and a variety of sensors. It is a challenging task to design and implement a control strategy that achieves optimal performance through the

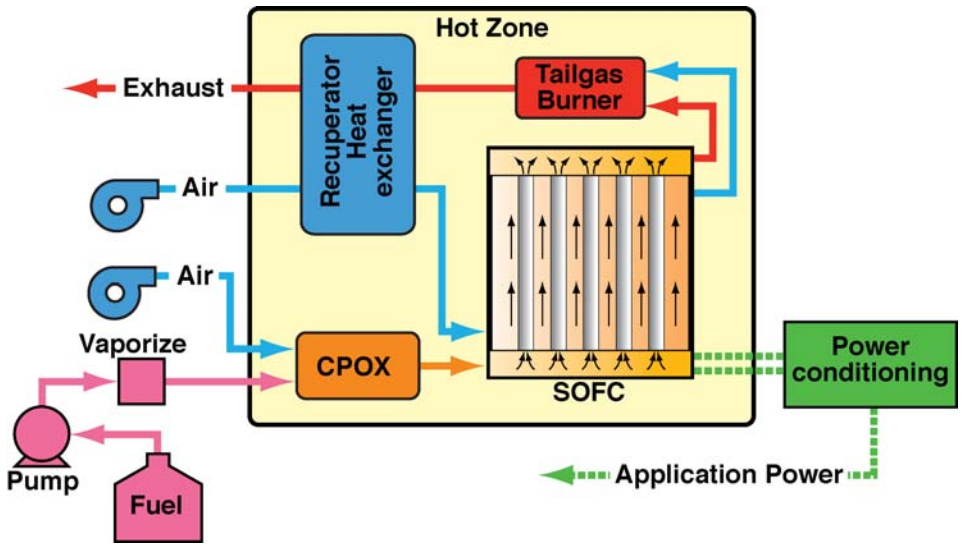

Figure 1. Schematic of an SOFC system that may be used for an APU application.

coordination of multiple sensors and actuators. The approach reported here builds upon a combination of physical and chemical models and advanced process-control theory.

Figure 1 is a schematic that illustrates essential aspects of a tubular SOFC system that may be used for an APU application (1). Vaporized fuel and air are converted within a catalytic partial oxidation (CPOX) reactor to produce a syngas mixture (i.e., a mixture of H_2, H_2O, CO, CO_2, and N_2). The syngas enters the anode side of the SOFC stack. Pre-heated air enters the cathode side of the SOFC stack. The SOFC stack is based upon a group of anode-supported tubes, with fuel flowing inside the tubes and air circulating outside the tubes. Air and unspent fuel leaving the stack are mixed and burned in a catalytic tail-gas combustor. Hot gases leaving the tail-gas burner are used in a recuperating heat exchanger to preheat the air entering the SOFC stack. Raw power from the SOFC is processed through power electronics before being delivered to the application load.

Fuel-cell developers may find the control theory and supporting mathematics to be unfamiliar, and possibly confusing. Nevertheless, most of the control approaches developed here are well established in the control community. At the same time, process-control developers may be unfamiliar with the detailed physics and chemistry of fuel cells. The effort reported here brings together these two bodies of knowledge in ways that facilitate the practical development of optimal control strategies.

Model Predictive Control Development

As itemized below, developing a model-predictive-control system follows an established set of steps:

Develop a physical model. The first step is to develop, and validate, a model that represents the essential physical and chemical processes that are to be controlled. The model must be transient in the sense that it can accept time-varying inputs (representing actuation) and can predict accurately time-varying responses (representing system observables).

Model reduction and system identification. Physical models are usually too complex and computationally expensive to be run as part of real-time controller software. Therefore, reduced-order models must be derived that capture the essential relationships between system actuation and sensor responses. As discussed subsequently, reduced-order models are locally linear models.

Develop the control algorithm. Based upon the reduced model, write a multiple-input–multiple-output (MIMO) control algorithm. Identifying actuation and sensing opportunities is an important element of the controller design. It is also important to consider operating constraints, both for actuation and system performance. The mathematics of the controller design are discussed in subsequent sections.

Test the controller in simulation. The controller is designed on the basis of the reduced-order model. However, its purpose is to control the nonlinear system. Once the controller is written, it can be used to drive the full nonlinear model. In this way, the controller performance can be evaluated for a range of potential operating trajectories.

Test and refine the controller with the actual hardware. The final objective is to use the controller to operate a real system. A model-predictive controller embeds a reduced-order model that is derived from a physical model. However, all models are based upon assumptions and simplifications. Therefore, it must be anticipated that iteration and refinement is needed to optimize the controller design and implementation.

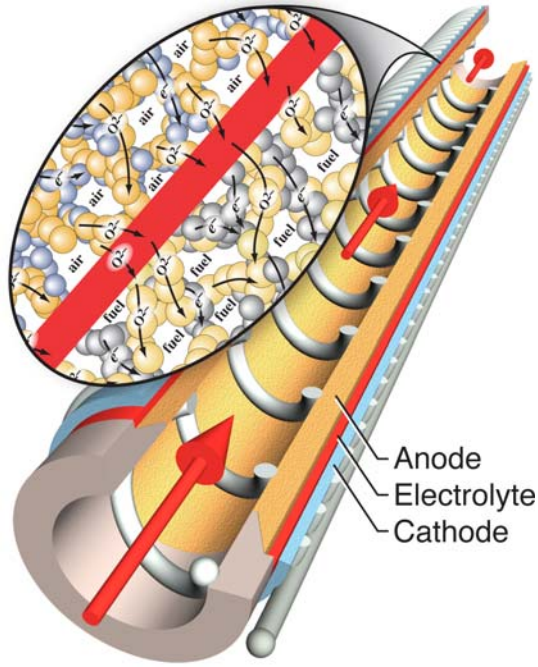

Figure 2. Illustration of a single SOFC tube. The balloon illustrates chemical and transport processes at the micron scale that are responsible for electrochemical charge transfer.

Physically Based Models

Because of limited space and the fact that they are well documented, the present paper provides only a brief summary of the physical models. Moreover, the present paper concentrates on model-predictive control for the fuel-cell stack, not the entire system. High fidelity models have been developed to represent the physics and chemistry of anode-supported tubular cells (2–7), such as the one illustrated in Fig. 2. These models consider fluid flow within the tube, porous-media gas transport within the pore spaces of the composite electrodes, catalytic reforming chemistry within the anode structure, and ion and electron transport throughout the membrane-electrode assembly. The models predict details of the flow fields, chemical composition, temperature, and current density throughout the tube. They also predict overall performance, including cell efficiency and fuel utilization.

Model Reduction and System Identification

For system-identification and control purposes, the SOFC stack is considered to have three inputs and four outputs (illustrated in Fig. 3). Although the stack has many tubes, the model here considers that all tubes behave exactly alike; thus only one representative tube is modeled. The reduced-order model is defined in terms of three vectors: the state vector x, the actuation vector u (cell voltage, fuel flow rate, and air flow rate), and the observable vector y (cell current, exhaust composition, MEA temperature, and air-exhaust temperature). The state vector is an abstraction that has no specific physical meaning. The linear-time-invariant reduced-order model is written in state-space form as

$$x_{k+1} = \Phi x_k + \Gamma \delta u_k,$$
$$\delta y_k = C x_k + D \delta u_k. \qquad [1]$$

In this equation, the index *k* represents time, with equal time increments. For this work, a time interval of 0.1 seconds is utilized. The first equation is essentially a linear ordinary differential equation that represents the evolution of the state vector, which depends upon the previous state and the actuation. The second equation represents the observable vector as a linear combination of the state x and the actuation u. The equations are written in terms of small perturbations from operating points u and y. That is, $u = \bar{u} + \delta u$ and $y = \bar{y} + \delta y$. The four matrices Φ, Γ, C, and D are found in such a way that solution of Eq. 1 provides a good representation of the physical-model response.

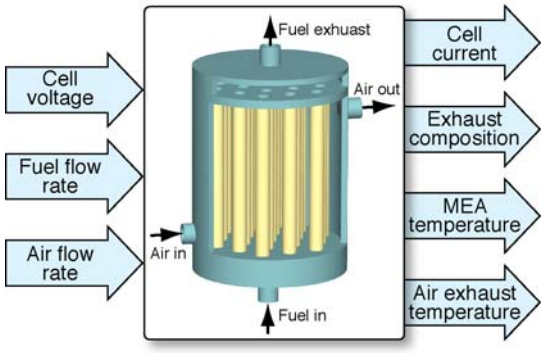

Figure 3. SOFC block diagram, showing input and output variables.

Subspace identification (8) is used to establish the matrices in Eq. 1. Assuming a nominal operating condition $(\bar{u}, \bar{y})$, a small-signal input sequence δu is designed with a frequency content that represents the expected system bandwidth. This sequence is used as input (i.e., $u = \bar{u} + \delta u$) to the physical model and the response sequence (i.e., $\delta y = y - \bar{y}$) is recorded. The small-signal input δu is a pseudo-random binary series (PRBS) of small step changes for the actuation. The identification process is implemented using toolboxes in MATLAB/SIMULINK (9–10). Usually the reduced-order model requires only a few states (i.e., the x vector has only 6-8 components).

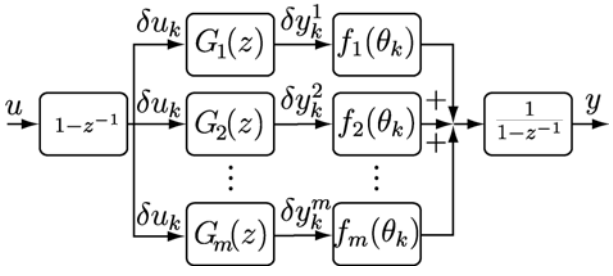

Figure 4. Schematic of the linear-parameter-varying (LPV) model structure for low-order representation.

A single linear system (i.e., Eq. 1) cannot adequately represent the nonlinear physical model over wide operating ranges. Therefore, a model structure is needed that smoothly combines linear systems at different

operating points. Figure 4 illustrates such a model. The forward shift operator is written as z and the linear system $G_i(z)$ is meant to accurately represent small-signal behavior at m different operating points (represented by the m outputs δy_k^i). The variable θ is a scheduling parameter that is a function of the (measured) system output, and $f_i(\theta)$ represents gains that are nonlinear functions of θ. For the fuel-cell stack, the scheduling parameters were chosen to be current and air-exhaust temperature. Because the linear models capture only small-signal behavior, the operating point needs to be removed from the input u, and added back to the output y. This is accomplished by applying a first difference $(1-z^{-1})$ to the input, and integrating the outputs $1/(1-z^{-1})$. In other words, the constant nominal actuation u is removed by differentiation, and the actual response y is recovered from the small-signal output δy by integration. The following paragraphs discuss the procedure to estimate the linear models $(G_i(z))$ and the parameter-varying $(f_i(\theta_k))$ portions of the model.

Identification of Scheduling Functions

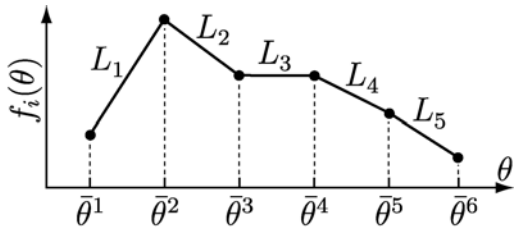

Figure 5. Illustration of a scheduling function $f_i(\theta)$.

As a result of the system-identification process discussed above, small-signal models $G_i(z)$ for m different operating points are available (the notation $G_i(z)$ represents a model in the form of Eq. 1. Thus, each system $G_i(z)$ represents the small-signal behavior for one value of the scheduling parameter θ. Let ϕ^i be the value of θ that corresponds to the operating point for system $G_i(z)$. The scheduling functions $f_i(\theta)$ are determined such that the nonlinear physical system response is predicted over a large operating range.

The scheduling parameters θ are quantized into a fixed number of "bins" with centers $\bar{\theta}^j$, $j=1,\ldots,l$, with the objective being to estimate the values for $f_i(\bar{\theta}^j)$ for all i,j (illustrated in Fig. 5). In other words, the set of values $f_i(\bar{\theta}^j)$ for all i,j become parameters to be identified. Note that the l can be much larger than m. Once the values of $f_i(\theta^j)$ are found, the functions $f_i(\theta)$ can be represented by an interpolation of these points.

With all $G_i(z)$ having been determined, the small-signal responses are calculated as

$$\delta y_k^i = G_i(z)(1-z^{-1})u_k \qquad [2]$$

where u_k is input to the physical model. The output y_k and quantized scheduling sequence $\bar{\theta}_k$ are also known. The objective is to select the values of $f_i(\bar{\theta}^j)$ such that $\delta y_k \approx \delta\hat{y}_k$, where δy_k is the first difference of the measured output from the physical model. That is,

$$\delta y_k = (1-z^{-1})y_k, \text{ and } \delta\hat{y}_k = \sum_{i=1}^{m} f_i(\bar{\theta}_k)\delta y_k^i. \qquad [3]$$

Without placing restrictions on $f_i(\bar{\theta}^j)$, this problem is under-determined. One set of constraints on the scheduling functions comes from their definition: they should range from 0 to 1, and the sum of the scheduling functions over *i* for any $\bar{\theta}^j$ should always sum to unity. In addition, the functions are required to be sufficiently smooth. Hsu et al. introduced a dispersion function, which is a smoothness measure for functions that are described point-wise (11). Given a point-wise description of a function in terms of points $(\bar{\theta}^j, f(\bar{\theta}^j))$, let

$$\Psi(f(\bar{\theta}^j)) = \sum_{j=1}^{l-1} L_j^2 \qquad [4]$$

be the dispersion of *f*, where L_j are the lengths of a linear interpolant of $(\bar{\theta}^j, f(\bar{\theta}^j))$, as shown in Fig. 4. Note that the dispersion is a quadratic function of the values of $f_i(\bar{\theta}^j)$.

By solving the following optimization problem, the values of the scheduling functions $f_i(\bar{\theta}^j)$ for all i,j can be found that minimize a weighted sum of the fitting error and the dispersion, subject to constraints:

$$\begin{aligned} &\min_{f_{i(\bar{\theta}^j)}} \ \sum_{k=1}^{N} \left\| \delta y_k - \delta \hat{y}_k \right\|^2 + \beta \sum_{i=1}^{m} \Psi(f_i(\bar{\theta}^j)) \\ &\text{subject to:} \quad 0 \le f_i(\theta^j) \le 1 \\ &\qquad\qquad \sum_{i=1}^{m} f_i(\bar{\theta}^j) \\ &\qquad\qquad f_i(\bar{\theta}^i) = 1 \end{aligned} \qquad [5]$$

The parameter β is used to set the relative weights between fitting error and dispersion. Because the dispersion is a quadratic function of $f_i(\bar{\theta}^j)$, this is a convex optimization problem that can be solved using modern optimization packages.

By augmenting the linear state-space representation with the past input and output (u_{k-1} and y_{k-1}), the final result is a parameter-varying state-space model of the form

$$\begin{bmatrix} u_k \\ x_{k+1} \\ y_k \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ -\Gamma & \Phi & 0 \\ -D(\theta) & C(\theta) & I \end{bmatrix} \begin{bmatrix} u_{k-1} \\ x_k \\ y_{k-1} \end{bmatrix} + \begin{bmatrix} I \\ \Gamma \\ D(\theta) \end{bmatrix} u_k \qquad [6]$$

where

$$C(\theta) = \begin{bmatrix} f_1(\theta) & 0 & \cdots \\ 0 & f_2(\theta) & \\ \vdots & & \ddots \end{bmatrix} C, \qquad D(\theta) = \begin{bmatrix} f_1(\theta) & 0 & \cdots \\ 0 & f_2(\theta) & \\ \vdots & & \ddots \end{bmatrix} D \qquad [7]$$

PRBS Results

Figure 6 illustrates the transient response of a physical SOFC model to PRBS input signals for cell voltage, fuel flow rate, and air flow rate. The most apparent response is the relationship between cell voltage and output current. Upon a sudden voltage decrease, the current increases rapidly, and vise versa. The sudden change in current is followed by a slower decay, which is caused by fuel composition variations associated with electrochemical oxidation rates. The characteristic time for composition changes is on the order

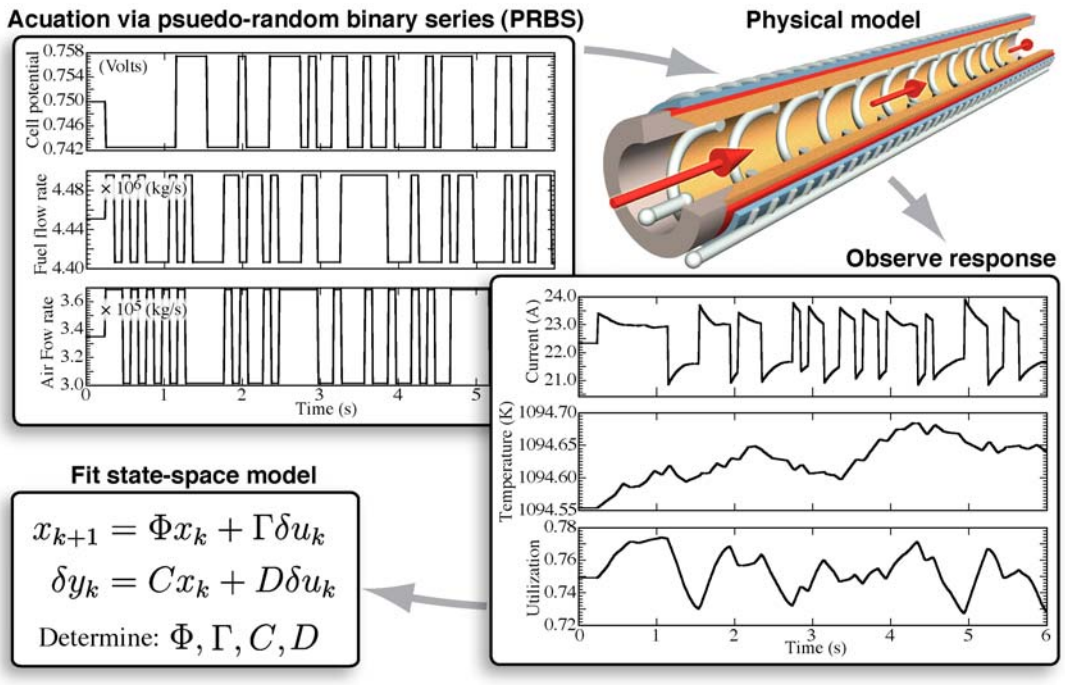

Figure 6. Illustration of the state-estimation process using a PRBS actuation of the physical model and the observation of observable responses.

of a second and is related to the residence time of the fuel channel flow and radial porous-media transport within the anode. The cell temperatures are also affected by the inputs. However, because of relatively high thermal capacity, the changes are small on the time scale of seconds illustrated in Fig. 6. A reduced-order model with only 6 to 8 states is found to be sufficient to capture the behavior of the physical model at a nominal operating condition.

Model Predictive Control

The model-predictive controller design is accomplished in two distinct parts: 1) estimate the current system state x from observations of past inputs u and outputs y, and 2) establish future actuation trajectories u to take the system from a current state through a desired output trajectory.

State Estimation

The first task is to use past sensor measurements y to determine the current state x. The state estimator is split into two parts–estimating the state of the small-signal model and estimating the state of the operating point (the integrator in Fig. 4). Assume that the process started at time $k = 0$, and the current time is $k = k_0$. Past information $(y_k, u_k, k = 0, \ldots, k_0)$ is available. Let $\delta y_k = y_k - y_{k-1}$ and $\delta u_k = u_k - u_{k-1}$. Choose weighting matrices R^e, Q^e and define the following optimization problem:

$$\min_{x_k} \sum_{k=0}^{k_0} \|e_k\|^2 + \|w_k\|^2, \text{ where} \qquad [8]$$

$$e_k = R^{\mathrm{e}}\left[\delta y_k - \left(C(\theta)x_k - D(\theta)\delta u_k\right)\right] \quad [9]$$

$$w_k = Q^{\mathrm{e}}\left[x_{k+1} - \left(\Phi x_k + \Gamma \delta u_k\right)\right] \quad [10]$$

This can be interpreted as follows: find the state sequence x_k such that the error e_k between the (small signal) measurements and the model output and the adjustment w_k needed to match the dynamic update for x_{k+1} is minimized. The relative weights R^{e} and Q^{e} can be chosen based upon engineering judgment, or they can be chosen based on a stochastic model for measurement errors and disturbances. In any case, this is a quadratic optimization problem. If the system is observable, the minimizer is unique. As stated, this optimization problem grows in size as k_0 gets larger. However the Kalman Filter can implement this estimation process recursively (8). That is, a gain sequence K_k can be calculated such that the output of the system

$$\hat{x}_{k+1} = \Phi\hat{x}_k + \Gamma u_k + K_k\left[y_k - \left(C(\theta)x_k + D(\theta)u_k\right)\right] \quad [11]$$

produces the same sequence of estimates for $\hat{x}_k$.

Control

The role of the control is to choose future actuation to guide the system according to a desired trajectory. Assume that a desired output trajectory can be specified as y_k^{d} for $k = k_0, \ldots, k_0 + p$. The control is chosen to best match the desired output trajectory over a horizon of p sampling intervals, while meeting fixed constraints on input and output excursions. At every time k_0 the current state $\hat{x}_{k_0}$ is estimated according to the foregoing optimization problem (i.e., Eq. 8) and the system output at time $k_0 - 1$ is measured. After choosing weighting matrices R^{c} and Q^{c} and constraints b_{ui} and b_{yi}, the following optimization problem is solved

$$\begin{aligned} \min_{u_k} \quad & \sum_{k_0}^{k_0+p} \left\| R^{\mathrm{c}}\left(\mathrm{y}_k^{\mathrm{d}} - y_k\right)\right\|_1 + \left\| Q^{\mathrm{c}} u_k \right\|_1 \\ \text{subject to:} \quad & x_{\mathrm{k+1}} = \left(\Phi x_k + \Gamma u_k\right) \\ & y_k = C(\theta)x_k - D(\theta)\delta u_k + y_{k-1} \\ & b_{u1} < u_k < b_{u2} \\ & b_{y1} < y_k < b_{y2} \end{aligned} \quad [12]$$

and the first element u_{k_0} of the optimal sequence for u is applied as the control. In the example that follows, p is chosen to be 12, for a horizon of 1.2 seconds.

Example Result

To illustrate controller performance consider the following example. The inlet fuel is mixture of 37.9% H_2, 1% CH_4, 18.9% CO, 0.4% CO_2, 3% H_2O, and 38.1% N_2. The system is operating nominally at 800 °C and 0.75 V, with fuel flow rate of 4.45e-6 kg/s (0.44 m/s to each tube) and an air flow rate equivalent to four stoichs. Initially, the cell is operating at steady state and delivering 22.3 A with a fuel utilization of 75% (9% H_2 in exhaust). Because dynamics of the fuel-processing and air-delivery components are not modeled, first-order systems with unity DC gains and time constants of 1 second are

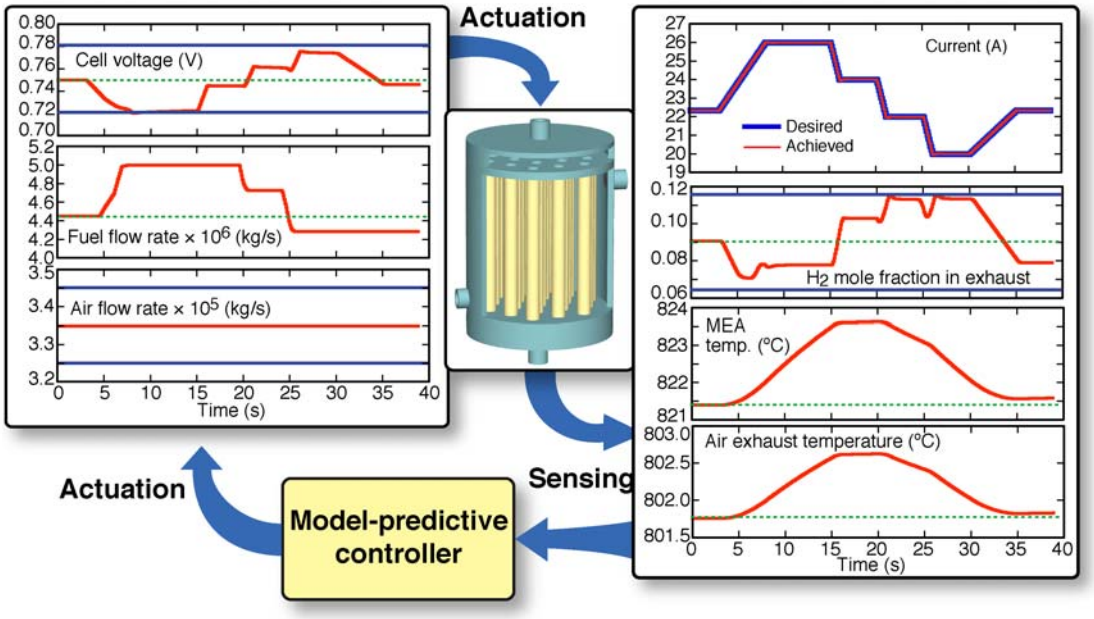

Figure 7. Example of control to a specified current using model predictive control. For current, the desired trajectory is the thick blue line, while the achieved trajectory is in orange. For other plots, the orange line is the trajectory, the dashed green line is the nominal value, and thin solid blue lines indicate desired upper and lower bounds.

added to the fuel- and air-input channels of both the physical and reduced-order models to simulate their effects.

The controller seeks to deliver a desired current profile that consists of several set-point changes between 20 and 26 A. This profile is illustrated by the thick blue line in the upper right of Fig. 7. The trajectories between setpoints are specified as ramps with slopes ranging between 0.5 and 2.0 A/s. Recall that the controller is designed with a trajectory preview of 1.2 s, corresponding to a scenario in which the fuel cell is integrated with power electronics and energy storage components (e.g., batteries or supercapacitors) that accommodate demand fluctuations of less than 1.0 s. An outer control loop specifies the desired current trajectory 1.2 seconds into the future. In addition to following the current-demand profile, the controller is written to respect the following constraints: cell voltage between 0.72 to 0.78 volts, fuel flow rate between 3.45e-6 to 5.45e-6 kg/s, and H_2 mole fraction in the anode exhaust between 6.6% and 11.6%.

Figure 7 illustrates the input (left hand side) and output (right hand side) trajectories. The controller is able to achieve tight tracking of the current trajectory while respecting all specified constraints. Achieving this result requires the controller to modify both the cell voltage and fuel flow in a coordinated manner. For example, at 3 seconds, the specified current trajectory calls for an increase in current. Initially the current ramp can be achieved by decreasing the voltage. However, because the desired current trajectory goes above the level that can be attained by varying the voltage alone, at 5 seconds the controller commands an increase in the fuel flow rate. Between 8 and 15 seconds, the desired current of 26 A is attained while the voltage remains at the lower bound of 0.72 V. The changes in fuel flow rate at 20 and 25 seconds are needed to ensure that the fuel-utilization (H_2 mole fraction) peaks at 21.5 and 26.5 seconds fall below the specified upper bound. In this example, covering an interval of 49 s, the air flow rate does not change.

This is because the air flow rate primarily affects cell temperature, which responds on a much longer time scale.

The controller logic is not unique. The real-time optimization involves weighing factors for the state estimation (Eq. 5) and the controller actuation (Eq. 12). Also, choices can be made in establishing the reduced-order models and the gain scheduling. Thus, there is considerable flexibility in designing and implementing the control strategy.

Summary and Conclusions

Model-predictive control offers significant opportunities for controlling fuel-cell systems and stacks. Using system identification, large-scale physical models can be reduced to low-order models that are embedded into real-time controller logic. Process-control strategies can be written to incorporate a predictive knowledge of fundamental fuel-cell physics and chemistry into control decisions. The controller can use and coordinate multiple sensors and actuation possibilities. The MPC controller can also be designed to respect constraints on both actuation and output. Such controllers can facilitate the optimal performance of the SOFC while following transient demand requirements. The present paper focused on a tubular fuel cell alone. Ongoing efforts are focused on incorporating all the balance-of-plant components into the MPC paradigm, thus optimizing the system as a whole.

Acknowledgements

This work has been funded in part by the Department of Energy, Office of Energy Efficiency and Renewable Energy (DE-FG36-08GO88100) and by the Office of Naval Research via an RTC grant (N00014-05-1-03339).

References

1. R. J. Kee, H. Zhu, A. M. Sukeshini, and G. S. Jackson, *Combust. Sci. and Tech*., **180**, 1207-1244 (2008).
2. H. Zhu, R. J. Kee, V. M. Janardhanan, O. Deutschmann, and D. G. Goodwin, *J. Electrochem. Soc*., **152**, A2427-A2440 (2005).
3. H. Zhu, A. M. Colclasure, R. J. Kee, Y. Lin, and S. A. Barnett, *J. Power Sources*, **161**, 413-419 (2006).
4. G. K. Gupta, J. R. Marda, A. M. Dean, A. M. Colclasure, H. Zhu, and R. J. Kee, *J. Power Sources*, **162**, 553-562 (2006).
5. H. Zhu and R. J. Kee, *J. Electrochem. Soc*., **153**, A1765-A1772 (2006).
6. H. Zhu and R. J. Kee, *J. Power Sources*, **169**, 315-326 (2007).
7. H. Zhu and R. J. Kee, *J. Electrochem. Soc*., **155**, B715-B729 (2008).
8. L. Ljung, *System identification: Theory for the user*, Prentice-Hall, 1999.
9. J. Löfberg, YALMIP: A toolbox for modeling and optimization in MATLAB, Proceedings of the CACSD conference, Taipei, Taiwan, http://control.ee.ethz.ch/joloef/yalmip.php, (2004).
10. M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, http://stanford.edu/~boyd/cvx, (2009).
11. K. Hsu, K. Poolla, and T. L. Vincent. *IEEE Trans. Automatic Control*, **58**(11), 2497-2513 (2008).