Experimental Real-Time Optimization of a Solid Oxide Fuel Cell Stack via Constraint Adaptation

Gene A. Bunin [a]***, Zacharie Wuillemin*** [b]***, Grégory François*** [a]***, Arata Nakajo*** [b]***, Leonidas Tsikonis*** [b]***, and Dominique Bonvin*** [a]

[a] *Ecole Polytechnique Fédérale de Lausanne (Laboratoire d'Automatique), Lausanne, Switzerland*
[b] *Ecole Polytechnique Fédérale de Lausanne (Laboratoire d'Energétique Industrielle), Lausanne, Switzerland*

Abstract: The experimental validation of a real-time optimization (RTO) strategy for the optimal operation of a solid oxide fuel cell (SOFC) stack is reported in this paper. Unlike many existing studies, the RTO approach presented here utilizes the constraint-adaptation methodology, which assumes that the optimal operating point lies on a set of constraints and then seeks to satisfy those constraints in practice via bias update terms. These biases correspond to the difference between predicted and measured outputs and are updated at each steady-state iteration, allowing the RTO to successfully meet the optimal operating conditions of a 6-cell SOFC stack, despite significant plant-model mismatch. The effects of the bias update filter values and of the RTO frequency on the power tracking and constraint handling are also investigated.

Keywords: Applied Fuel Cell Modeling, Constraint Adaptation, Optimal Fuel Cell Performance, Real-Time Optimization, SOFC Load Tracking, SOFC Operation.

1. Introduction

In the recent decade, fuel cells have received growing attention as viable energy alternatives, advocated as a cleaner and more efficient energy source. There remains, however, a number of open problems with fuel cell technology that must be resolved before it can be put into widespread use and become a practical, capable substitute for current methods. One of these issues lies in the life of a cell, which can be shortened significantly if the system does not successfully adhere to certain safe operating regions, qualified by constraints on certain input and output variables. In addition to simply being safe, the cell must also perform optimally and be able to operate at the highest efficiency for any immediate power demand. For these reasons, the domains of control and optimization have been increasingly called upon for improved fuel cell performance. Unfortunately, despite a large number of theoretical contributions - the majority of which have focused on control (see, for example, [5] or [13]) and few on optimization ([6] and [14]) - there still remains a large gap between simulation studies and reported experimental results. To the best of the authors' knowledge, all experimental studies so far have been limited to proton exchange membrane (PEM) cells ([1, 12, 9]), and many have sought efficiency based on specific criteria that were already known in advance, rather than formally treating the cell as a multiple-input, constrained optimization problem with changing optimal conditions.

In this paper, a previously developed and simulated constraint-adaptation methodology ([7]) is validated experimentally for a 6-cell SOFC stack. Unlike the method in [6], which seeks to track an optimality criterion via model-predictive control, or the approach in [14], which aims to achieve optimality by tracking the flow-dependent maximum power, the constraint-adaptation methodology discussed here works on the very simple - yet often true - premise that the optimum of the problem lies somewhere on the constraints. Therefore, if the proper set of constraints can be met in practice, then the optimality of the process is guaranteed as well. Because of uncertainty, the values of the constrained quantities given by the model will rarely match those provided by the real system, and so an adaptation - carried out by adding a bias term to the modeled constraints - is used to ensure that the constraints used by the optimization match those of the real system. In doing so, the RTO iteratively drives the system to the true constraints, with the speed of convergence dictated

Corresponding author: Dominique Bonvin, Email: dominique.bonvin@epfl.ch

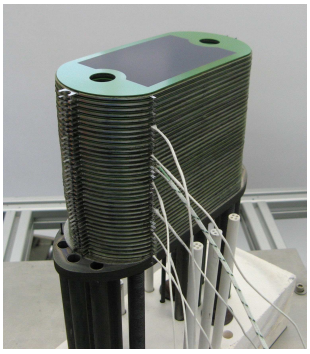

Figure 1: The setup of a typical stack.

by the way the bias update is filtered. As a demonstration of its robustness to uncertainty, the ability of this method to reject long-term system degradation has been shown via simulation in [2].

This paper will be structured as follows. In Section 2, the description of the experimental apparatus and a summary of the model used for the SOFC stack will be given. In Section 3, the constraint-adaptation methodology will be discussed in detail. Its application to the real system will then be outlined and the tested scenarios explained. Section 4 will focus on the results - looking not only at the general performance of the experimental stack, but also at the effects of altering the bias filters and optimization frequency, and Section 5 will conclude the paper.

2. System Description

2.1. Experimental Setup

2.1.1. Design and Interface

The study presented in this article concerns an SOFC short-stack developed at EPFL for HTceramix-SOFCpower ([11, 4]). This stack consists of planar anode-supported cells with an active area of 50 cm^2, pressed between gas-diffusion layers (SOFConnex^{TM}) and metallic interconnector plates. The anodes are made of standard nickel/yttrium-stabilized-zirconia (Ni-YSZ) cermet, while the thin electrolyte consists of dense YSZ. The cathodes are made of screen-printed $(La, Sr)(Co, Fe)O_3$, allowing standard operation temperatures between 650°C and 850°C. A detailed description of its construction can be found in [11], and a photo of a typical assembly is given in Fig. 1. A stack of 6 cells was used for this study.

The stack was placed in a high-temperature furnace at 775°C and connected to a testing station providing controlled flow rates of air and preheated fuel (~770°C). Additionally, an active load was used to control the delivered current. The control of the testing station was ensured by a LabVIEW interface in which the RTO algorithm was implemented via a MATLAB Script function.

2.1.2. System Constraints

Two key constraints limit the efficiency in an SOFC. While the cell may deliver a given electrical power at several different operating conditions (different fuel flows and currents), the maximum electrical efficiency is usually found close to the highest achievable fuel utilization (70-90%) - defined as the percentage of the fed fuel that reacts. However, to prevent damages to the stack by local fuel starvation and reoxidation of the anode [11], a conservative maximum fuel utilization of 75% is set. In addition, it is known that significant internal losses are detrimental to SOFC stacks. These losses appear as differences between the ideal and measured cell potential (overpotentials). Therefore, a minimum cell voltage of 0.75V is set to protect the stack from accelerated degradation, resulting in the second major constraint for the system.

In addition, the air excess ratio (or just the "air ratio"), defined as the stoichiometric ratio between the oxygen fed to the system and the oxygen needed to react with the fuel, must be kept within certain bounds so as to avoid steep thermal gradients. For this setup, the ratio is kept between 4 and 7. A lower bound of 3.14 $ml/(min \cdot cm^2)$[1] is also placed on the fuel feed rate so as to avoid local (or widespread) fuel starvation, and an upper bound of 30A is placed on the current to avoid excessive heating ([7]).

To avoid damaging the stack, limits on the rate of input changes were defined as 0.54 $ml/(min^2 \cdot cm^2)$, 1.37 $ml/(min^2 \cdot cm^2)$, 2.0 A/min for the hydrogen flow, oxygen flow, and current, respectively. Conditional laws were written into the LabVIEW code so that any of these rates could be set to 0 in the case of a fuel utilization or air ratio violation.

2.2. Steady-State SOFC Model

The steady-state model used to optimize the stack is largely similar to that which has been previously reported in [7]. As such, only the most fundamental elements of the model, or any deviations from prior reported work, are given here. For a full and detailed treatment, the interested reader is referred to [7].

[1] All flux values, given in $ml/(min \cdot cm^2)$, are calculated under normal conditions.

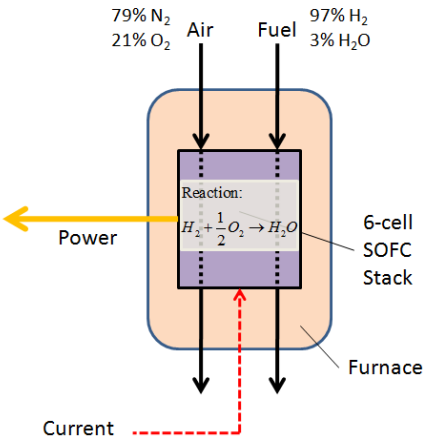

Figure 2: The diagram of the SOFC system. A fuel feed that is 97% hydrogen, 3% water is sent through the anode, while air is fed through the cathode. The stack is polarized with an exogenous current.

As mentioned in the previous section, the SOFC is a system fed with O_2 (air stream) and H_2 (fuel stream), which react electrochemically to produce electrical power and heat. The fuel cells are usually assembled in stacks in order to reach the desired voltage, and require an outside current source to operate (Fig. 2).

Three inputs are used to control the operation of the stack: the hydrogen flux $\dot{n}_{H_2}$ (in ml/(min $\cdot$ cm^2)), the oxygen flux $\dot{n}_{O_2}$ (in ml/(min $\cdot$ cm^2)), and the current I (in A):

$$\mathbf{u} = \begin{bmatrix} \dot{n}_{H_2} \\ \dot{n}_{O_2} \\ I \end{bmatrix}. \tag{1}$$

As will be shown in the following section, the three outputs of interest are the power density p_{el} (in W/cm^2), the cell potential U_{cell} (in V), and the electrical efficiency η. The potential is modeled based on the equivalent circuit approach ([8]) :

$$\begin{aligned} U_{cell} = U_N - U_{act,c} - U_{i,e} - U_{dis,c} \\ -U_{dif,a} - U_{dif,c} - U_{MIC}. \end{aligned} \tag{2}$$

Here, U_N denotes the reversible cell voltage, while $U_{act,c}$, $U_{dis,c}$, $U_{dif,a}$ denote the non-Ohmic overpotential losses (cathode activation, cathode oxygen dissociation, and diffusion at the anode, respectively). $U_{i,e}$, $U_{dif,c}$, and U_{MIC} denote the Ohmic losses (ionic conductivity, diffusion through the cathode, and metallic interconnect, respectively). The latter, not mentioned in [7], is defined as:

$$U_{MIC} = (R_{MIC,1} + R_{MIC,2})\, I, \tag{3}$$

with $R_{MIC,1}$ and $R_{MIC,2}$ denoting the resistances of the two interconnects.

The values of p_{el} and η follow as functions of U_{cell}:

$$\begin{aligned} p_{el} &= \frac{U_{cell}\, N_{cells}\, I}{A_c} \\ \eta &= \frac{p_{el}\, A_c}{\dot{n}_{H_2}\, Q_L}, \end{aligned} \tag{4}$$

where N_{cells} is the number of cells in the stack, A_c is the active area of the cell, and Q_L is the lower heating value of the fuel. Unlike in the work in [7], the parasitic power demand of the air blower is not included in the definition of η.

Because U_{cell}, and thus also p_{el} and η, depend substantially on the temperature of the stack ([7]), an energy balance is also required, and may be expressed as:

$$mc_P \frac{dT}{dt} = -\Delta\dot{H}_{gas} - p_{el}A_c - \dot{Q}_{loss}, \tag{5}$$

with m, c_P, and T used to express the mass, specific heat capacity, and temperature of the stack, respectively. $\Delta\dot{H}_{gas}$ denotes the enthalpy change for the gases, while $\dot{Q}_{loss}$ denotes the radiative heat loss calculated as:

$$\dot{Q}_{loss} = A\alpha\sigma_{SB}(T^4 - T_{furn}^4), \tag{6}$$

where A is the area of the stack, α is a transfer factor, σ_{SB} is the Stefan-Boltzmann factor, and T_{furn} is the temperature of the furnace.

To calculate the steady-state values of U_{cell}, p_{el}, and η, it is first necessary to integrate Eq. 5, and to use the resulting steady-state temperature to obtain the values of the potential, power, and efficiency. As a result, while the response of these quantities to changes in the inputs is practically instantaneous, the true steady state of the system is governed by the temperature as it gradually reaches its new value. While some SOFC systems may have additional dynamics depending on their setup ([10]), it is assumed that, for the system at hand, these two time scales - one instantaneous and one on the magnitude of approximately 30 minutes - are the only significant ones.

In testing this model against the real SOFC stack, one can see a divergence between the predicted potential and the actual value when the current is increased (Fig. 3). This is particularly crucial for the current range 18 to 25A, which is used throughout many of the experiments (Section 4.). However,

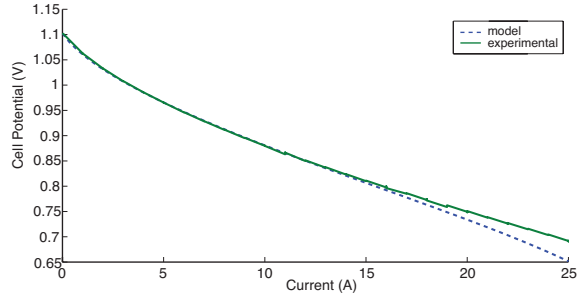

Figure 3: Current-potential (IV) curve for $\dot{n}_{H_2} = 5$ ml/(min · cm^2), $\lambda_{air} = 4$.

the constraint-adaptation methodology, introduced next, is an excellent tool for dealing with this deviation.

3. RTO via Constraint Adaptation

3.1. Methodological Overview

Process optimization typically involves the minimization of a cost (or the maximization of a profit) that is subject to certain equality and inequality constraints. This results in a nonlinear programming (NLP) problem which, for the system described in Section 2., may be written as:

$$\begin{aligned} \max_{\mathbf{u}} \quad & \eta(\mathbf{u},\boldsymbol{\theta}) \\ \text{s.t.}: \quad & p_{el}(\mathbf{u},\boldsymbol{\theta}) = p_{el}^S \\ & U_{cell}(\mathbf{u},\boldsymbol{\theta}) \geq 0.75 \text{ V} \\ & \nu(\mathbf{u}) \leq 0.75 \\ & 4 \leq \lambda_{air}(\mathbf{u}) \leq 7 \\ & u_1 \geq 3.14 \text{ ml/(min cm}^2) \\ & u_3 \leq 30 \text{ A}, \end{aligned} \tag{7}$$

with the electrical efficiency η acting as the profit function to be maximized. Here, the superscript S denotes the setpoint for the power demand, effectively giving the optimization a second role as a load-following controller. The fuel utilization, ν, and the air ratio, λ_{air}, may be expressed in terms of the inputs as:

$$\nu = \frac{N_{cells} I}{2\dot{n}_{H_2} F} = \frac{N_{cells}}{2F}\frac{u_3}{u_1} \tag{8}$$

$$\lambda_{air} = 2\frac{\dot{n}_{O_2}}{\dot{n}_{H_2}} = 2\frac{u_2}{u_1}, \tag{9}$$

where F is the Faraday constant.

The vector of uncertain parameters $\boldsymbol{\theta}$ in (7) is used to represent the model parameters that do not match those of the real system.

With the steady-state model described in Subsection 2.2., the NLP problem (7) may be solved to obtain an optimal set of inputs that theoretically maximizes the cell efficiency while satisfying all the constraints. However, due to plant-model mismatch and process disturbances, this nominal solution is unlikely to be optimal for the actual SOFC system. For the system presented in this article, the optimization is very intuitive and follows the following general rules:

- At lower power demands, maximize ν to maximize efficiency, i.e. ν is the active constraint.
- For higher power demands, U_{cell} becomes the active constraint, and pushing ν to its boundary is no longer optimal.

Because the fuel utilization ν is independent of $\boldsymbol{\theta}$, it is known with certitude and thus can be satisfied exactly in (7). This is not true for the cell potential, however, which cannot be modeled perfectly and is often susceptible to uncertainty (as demonstrated in Fig. 3). Therefore, it is possible for the nominal problem to either underestimate or overestimate this value, resulting in "optimal" input values that will, in practice, either violate the constraint or reach an early limit by assuming it is active when it is not. This problem has been resolved in simulation in [7] with the use of "modifiers", or bias terms which are added to the uncertain constraint quantities to correct the model estimation. Using the modifiers, ε, for the power demand and the cell potential (but not for the efficiency, as adding a constant term would not affect the solution) results in a modified NLP problem that reads:

$$\begin{aligned} \max_{\mathbf{u}} \quad & \eta(\mathbf{u},\boldsymbol{\theta}) \\ \text{s.t.}: \quad & p_{el}(\mathbf{u},\boldsymbol{\theta}) + \varepsilon^{p_{el}} = p_{el}^S \\ & U_{cell}(\mathbf{u},\boldsymbol{\theta}) + \varepsilon^{U_{cell}} \geq 0.75 \text{ V} \\ & \nu(\mathbf{u}) \leq 0.75 \\ & 4 \leq \lambda_{air}(\mathbf{u}) \leq 7 \\ & u_1 \geq 3.14 \text{ ml/(min cm}^2) \\ & u_3 \leq 30 \text{ A}. \end{aligned} \tag{10}$$

As the modifiers are generally unable to converge to the optimal values in a single iteration, convergence is sought over a few iterations, using a low-pass filter with the filter constants $\mathbf{K}$ as suggested in [3]. At the k^{th} iteration, the optimization problem (10) is solved for $\mathbf{u}_k$ using the modifiers $\varepsilon_{k-1}^{p_{el}}$ and $\varepsilon_{k-1}^{U_{cell}}$ from the previous iteration. Then, the modifiers are

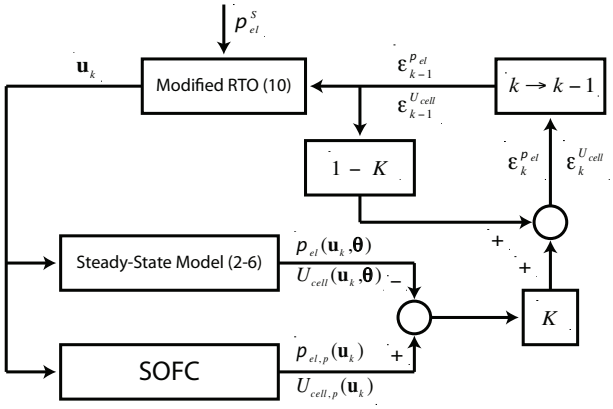

Figure 4: Constraint-adaptation RTO scheme.

updated as follows:

$$\begin{aligned} \varepsilon_k^{p_{el}} &= (1 - K_{p_{el}})\varepsilon_{k-1}^{p_{el}} + \\ &\quad K_{p_{el}}[p_{el,p,k} - p_{el}(\mathbf{u}_k, \boldsymbol{\theta})] \\ \varepsilon_k^{U_{cell}} &= (1 - K_{U_{cell}})\varepsilon_{k-1}^{U_{cell}} + \\ &\quad K_{U_{cell}}[U_{cell,p,k} - U_{cell}(\mathbf{u}_k, \boldsymbol{\theta})], \end{aligned} \tag{11}$$

with the subscript k indicating the iteration number and the subscript p used to denote a plant value. At the optimum (for $k \to \infty$), the modifiers will have converged and will simply be the difference (or bias[2]) between the actual and estimated values:

$$\begin{aligned} \varepsilon_\infty^{p_{el}} &= p_{el,p}(\mathbf{u}_\infty) - p_{el}(\mathbf{u}_\infty, \boldsymbol{\theta}) \\ \varepsilon_\infty^{U_{cell}} &= U_{cell,p}(\mathbf{u}_\infty) - U_{cell}(\mathbf{u}_\infty, \boldsymbol{\theta}) \end{aligned} \tag{12}$$

With the addition of the modifiers, the solution given by the optimization is guaranteed, upon convergence, to satisfy the constraints of the plant. The general algorithm proceeds as follows:

1. Set $k = 1$ and choose initial values for the modifiers $\varepsilon_0^{p_{el}}$ and $\varepsilon_0^{U_{cell}}$.
2. Solve the modified optimization problem (10) to obtain new input values $\mathbf{u}_k$.
3. Apply these input values and let the system converge to a new steady state.
4. Update the modifiers according to (11). If $\|\mathbf{u}_k - \mathbf{u}_{k-1}\| \leq \delta$ (where δ is a user-specified criterion), assume convergence. If not, set $k := k + 1$ and return to Step 2.

The algorithm is presented schematically in Fig. 4.

[2]To clarify, the complete differences between the model and plant values are referred to as "bias", while the partial, filtered differences used in the optimization are called "modifiers".

3.2. Application to the Real Stack

To test the effectiveness of the methodology presented above, a preset power demand profile,

$$p_{el}^S(t) = \begin{cases} 0.30 \frac{\mathrm{W}}{\mathrm{cm}^2} & t \leq 90\,\mathrm{min} \\ 0.38 \frac{\mathrm{W}}{\mathrm{cm}^2} & 90\,\mathrm{min} < t \leq 180\,\mathrm{min} \\ 0.30 \frac{\mathrm{W}}{\mathrm{cm}^2} & t > 180\,\mathrm{min} \end{cases} \tag{13}$$

was taken to demonstrate how the change in active constraints (from ν to U_{cell}) may occur. Note that the power demand profile of Eq. 13 acts as a disturbance at the RTO layer: in other words, it is not known *a priori* when and how the power demand may change. An RTO iteration frequency of 30 minutes was used, as this was generally the time it took for the actual system to reach steady state. The initial (sub-optimal) steady-state inputs were 5 ml/(min · cm^2), 12.77 ml/(min · cm^2), and 20 A for the fuel flux, oxygen flux, and current, respectively.

4. Results and Discussion

4.1. Optimal Power Tracking with Different Filters

The scenario described in Subsection 3.2. was tested in the SOFC system for different values of the filter constants $\mathbf{K}$. For this set of scenarios, $K_{U_{cell}}$ was set equal to $K_{p_{el}}$, and the two were varied together. Three different filter values of 0.4, 0.7, and 1.0 (the latter corresponding to full adaptation) were investigated. Complete results, including the input and efficiency graphs, are presented in Fig. 5 for a filter value of 0.4. The subsequent sets (Fig. 6-7) are given in their cut versions, and focus only on the power tracking and constraints (the inputs and efficiencies in these latter cases follow trends that are very similar to those in the former).

It can be observed that the optimizer immediately seeks to maximize the air ratio for all cases. This is because, without any parasitic losses to the air blower in the objective function, there is no reason for the optimizer to keep it at low levels. The fuel utilization also seems extremely sensitive to small disturbances in the hydrogen flux, which leads to occasional fluctuations and violations in this constraint. Finally, there is a "dip" in the fuel utilization during power changes, which is due to the fact that the system must keep the air ratio below its upper limit during the transient and, for this reason, does not decrease the hydrogen flux quickly enough to

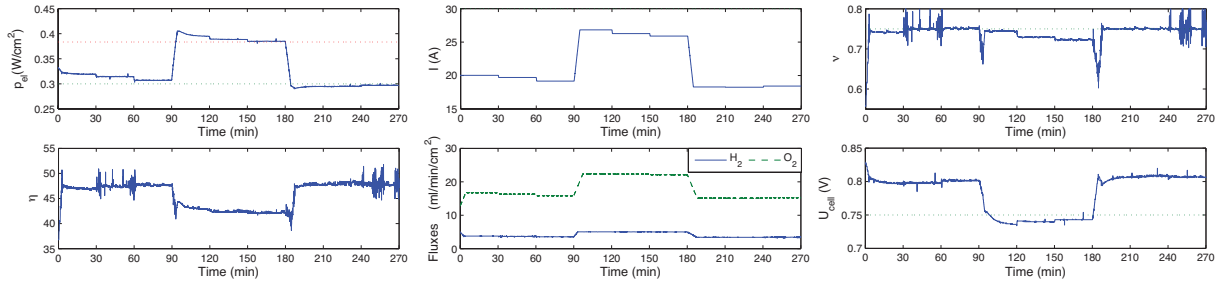

Figure 5: RTO performance with $K_{p_{el}} = K_{U_{cell}} = 0.4$.

match a decrease in current. What results is a temporarily low fuel utilization.

Otherwise, as expected, the filter constants affect the speed of convergence to the optimum. With a low filter, as in Fig. 5, convergence is very slow and damped. For the medium-sized filter in Fig. 6, it is quicker but still damped. For the full adaptation case in Fig. 7, convergence is fast but oscillatory.

Of additional interest is the way the algorithm handles the constraints. For fuel utilization, there are practically no issues (except for the noisy performance of the PI fuel flux controller), as there is no uncertainty. More interesting is the electric potential, whose constraint is initially violated when the algorithm tries to use the modifiers obtained for a low power demand to compute the optimum for a higher one. With steady-state RTO alone, there seems to be no means to solve this problem, as the converged modifiers from the first power demand always lead to this sort of violation in the second.

4.2. Optimal Filter Design

To improve power tracking and constraint satisfaction, it is possible to assign different values to the modifiers. From the bias values in Table 1 (taken from the experiment depicted in Fig. 7), it is clear that the bias for the electric potential constraint changes little during operation at the given conditions. Therefore, it is of interest to get to this value as quickly as possible so as to obtain and maintain an accurate prediction. One way is to use the highest possible value (i.e. 1) for $K_{U_{cell}}$ (it is readily seen that using lower values, as in Fig. 5, only leads to greater violations and slower convergence).

For the power demand, it is more difficult to estimate an optimal filter value. A value of 0.7 results in a second-iteration step that is too small, while a value of 1.0 leads to a step that is too big (compare the instances at $t = 120, 210$ between Fig. 6 and 7). Assuming that the optimal value lies somewhere in between, one may choose $K_{p_{el}} = 0.85$. Implementing this value does indeed result in much better power tracking, as shown in Fig. 8. The potential constraint is violated, but returns back to its bound quickly due to the large filter value.

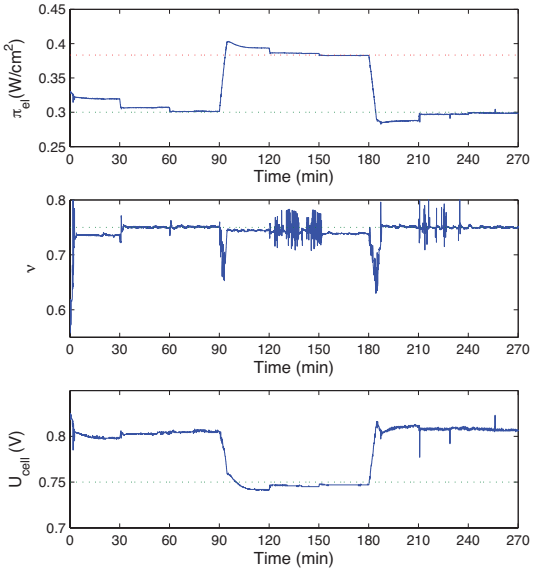

Figure 6: RTO performance with $K_{p_{el}} = K_{U_{cell}} = 0.7$.

4.3. Fast RTO

Although SOFCs with the capability to track constant power demand profiles may be of use industrially, many applications involve power demand changes that occur much more frequently than on the scale proposed in the preceding experiments. For this reason, the use of "fast" RTO was investigated. Instead of waiting for the system to reach true steady state with constant temperature (~30 min) before implementing the RTO, it was assumed that

Table 1: Values of the bias/modifiers for the experiment in Fig. 7.

t (min)	0	30	60	90	120	150	180	210	240	270
$\varepsilon^{p_{el}}$ (W/cm^2)	11.9	19.9	17.8	18.2	26.7	24.8	24.7	17.3	18.5	18.3
$\varepsilon^{U_{cell}}$ (V)	0.101	0.164	0.163	0.162	0.166	0.163	0.161	0.161	0.163	0.165

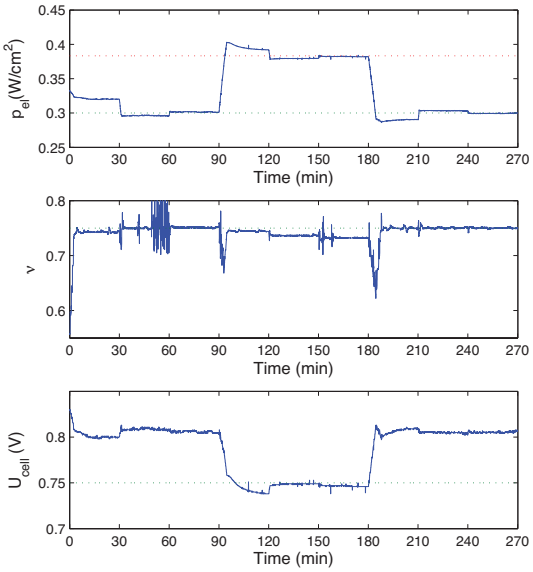

Figure 7: RTO performance with $K_{p_{el}} = K_{U_{cell}} = 1.0$.

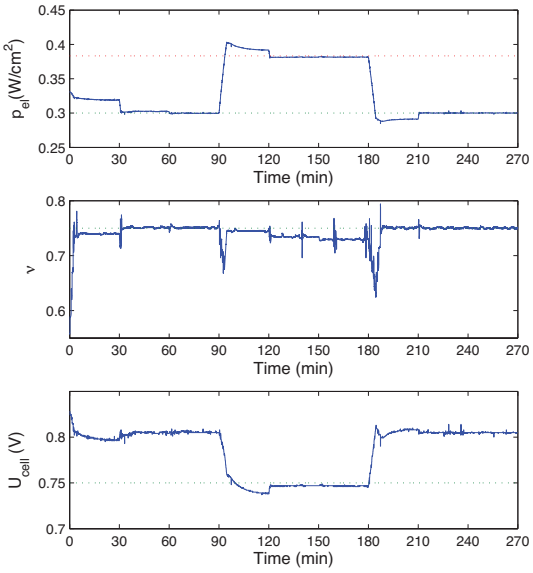

Figure 8: RTO performance with $K_{p_{el}} = 0.85$ and $K_{U_{cell}} = 1.0$.

the majority of the output response had already occurred at the electrochemical scale (< 1 sec). With this assumption, the temperature dynamics were ignored and treated like a slow-scale parametric drift, and the RTO frequency was increased to an action every 10 seconds. The optimal filter constants found in the previous section were retained. A power-demand profile spanning one hour of operation was generated, with a new random power demand between 0.30 and 0.38 W/cm^2 being given every 5 minutes. Towards the end of the experiment, a 15 minute stretch was used to manually test the ability of this algorithm to meet the maximum power without violating the constraint. A converged plant at $p_{el}^S = 0.30$ W/cm^2 was used as a starting point. The results are presented in Fig. 9.

The outcome is very promising. Owing to the fact that there is a very large difference between the two time scales, the optimizer does not suffer from the lack of true steady-state bias, and is able to also act as a very effective *controller* - quickly tracking the appropriate power demand without needing any extensive tuning. Via its role as an optimizer, it maintains the efficiency at near-optimal levels throughout the course of operation. Finally, unlike in the previous cases where the slow updates allowed violations, the potential constraint is approached and met, rather than violated, with this method. One does notice, however, that new power demands cannot be met if the change in the demand is too large (the last iteration in Fig. 9), but this is a limitation of the physical system, rather than of the algorithm.

5. Conclusions

An RTO with constraint adaptation was investigated for an experimental SOFC stack. It was shown that, despite uncertainty and plant-model mismatch, the adaptive optimization algorithm was able to successfully drive the system to its true optimum, converging to the specified power demand and to the proper active constraint. With additional studies, it was shown that tuning the filter values could result in even better performance and faster convergence. A high-frequency RTO was also attempted, and it was demonstrated that ignoring the transient effects of the temperature did not harm the performance of

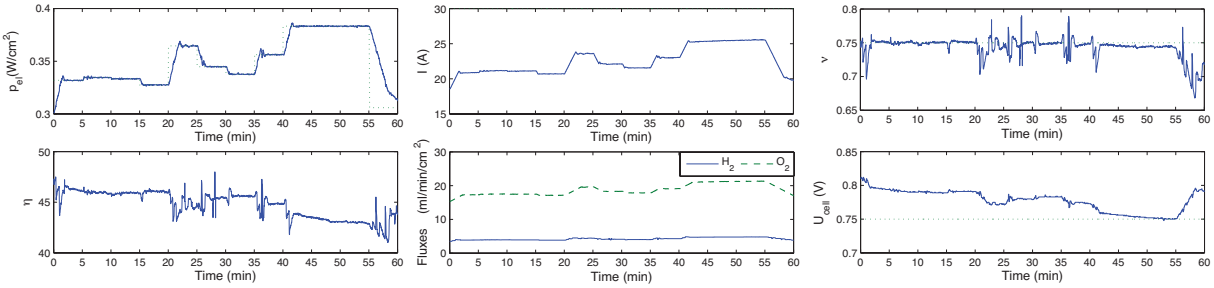

Figure 9: Performance of the fast RTO.

the algorithm. As a result, the RTO acted as both an optimizer and a controller in this case.

Though not addressed in this paper, this mechanism still has open issues that must be looked at. A more rigorous theoretical treatment of the filter tuning is still needed. The efficacy of the proposed method for more complicated SOFC problems, such as those involving steam reformers, cost criteria with parasitic losses, or heat demand-following with co-generation, is yet to be studied.

References

[1] A. Arce, A.J. del Real, C. Bordons, and D.R. Ramírez. Real-time implementation of a constrained MPC for efficient airflow control in a PEM fuel cell. *Accepted for publication in: IEEE Trans. on Industrial Electronics*, 57, 2010.

[2] G.A. Bunin, G. François, and D. Bonvin. Two-layered real-time optimization of a solid oxide fuel cell stack. In *DYCOPS, Leuven (accepted)*, 2010.

[3] B. Chachuat, A. Marchetti, and D. Bonvin. Process optimization via constraints adaptation. *J. Process Contr.*, 18:244–257, 2008.

[4] S. Diethelm, J. Van Herle, Z. Wuillemin, A. Nakajo, N. Autissier, and M. Molinelli. Impact of materials and design on solid oxide fuel cell stack operation. *J. Fuel Cell Sci. and Tech.*, 5(3):–, 2008.

[5] R. Gaynor, F. Mueller, F. Jabbari, and J. Brouwer. On control concepts to prevent fuel starvation in solid oxide fuel cells. *J. Power Sources*, 180:330–342, 2008.

[6] J. Golbert and D.R. Lewin. Model-based control of fuel cells: (2) optimal efficiency. *J. Power Sources*, 173:298–309, 2007.

[7] A. Marchetti. *Modifier-adaptation methodology for real-time optimization (No. 4449)*. PhD thesis, EPFL, 2009.

[8] A. Nakajo, Z. Wuillemin, J. Van Herle, and D. Favrat. Simulation of thermal stresses in anode-supported solid oxide fuel cell stacks. Part I: Probability of failure of the cells. *J. of Power Sources*, 193(1):203–215, 2009.

[9] P. Rodatz, G. Paganelli, A. Sciarretta, and L. Guzzella. Optimal power management of an experimental fuel cell/supercapacitor-powered hybrid vehicle. *Control Eng. Practice*, 13:41–53, 2005.

[10] C. Stiller, B. Thorud, O. Bolland, R. Kandepu, and L. Imsland. Control strategy for a solid oxide fuel cell and gas turbine hybrid system. *J. Power Sources*, 158:303–315, 2006.

[11] Z. Wuillemin. *Experimental and Modeling Investigations on Local Performance and Local Degradation in SOFC (No. 4525)*. PhD thesis, EPFL, 2009.

[12] Y.P. Yang, F.C. Wang, H.P. Chang, Y.W. Ma, and B.J. Weng. Low power proton exchange membrane fuel cell system identification and adaptive control. *J. Power Sources*, 164:761–771, 2007.

[13] X.W. Zhang, S.H. Chan, H.K. Ho, J. Li, G. Li, and Z. Feng. Nonlinear MPC based on the moving horizon state estimation for the solid oxide fuel cell. *Int. J. Hydrogen Energy*, 33:2355–2366, 2008.

[14] Z. D. Zhong, H. B. Huo, X. J. Zhu, G. Y. Cao, and Y. Ren. Adaptive maximum power point tracking control of fuel cell power plants. *J. Power Sources*, 176:259–269, 2008.

Acknowledgments: The authors would like to thank Stefan Diethelm and Luis Quina, from the Laboratoire d'Energétique Industrielle of EPFL, for their assembly of the stack, and for all subsequent laboratory assistance thereafter.