2008 American Control Conference
Westin Seattle Hotel, Seattle, Washington, USA
June 11-13, 2008

Optimization of Switched-Mode Systems with Switching Costs

X.C. Ding, Y. Wardi, D. Taylor, and M. Egerstedt

***Abstract*— This paper concerns the problem of optimizing switched-mode hybrid dynamical systems, where it is required to balance the tracking of a reference signal with an attempt to reduce a state-dependent cost-penalty associated with the switchings among the modes. We propose an algorithmic approach to the problem and we illustrate it on a specific example. The investigation has been motivated by design issues in power electronics, where it is desirable to regulate the load current in a switching circuit by controlling the state of a switch. On one hand tighter regulation requires switching at a higher frequency, and on the other hand each switching requires a certain amount of energy. Our algorithmic approach brings together various techniques that we recently have developed for optimizing hybrid systems, and its demonstrated application on a specific example suggests its potential viability in a broader array of problems in power electronics.**

I. Introduction

Optimal mode-scheduling in switched-mode hybrid dynamical systems has been extensively investigated in the past few years, and several algorithmic techniques have been developed [1], [2], [3], [4], [7], [8]. In such problems, the system's dynamics typically have the following form,

$$\dot{x} \ \in \ \{f_\alpha(x) : \alpha \in \mathcal{A}\}, \tag{1}$$

$t \in [0, T]$, where $x \in R^n$ is the state variable, the initial state $x_0 := x(0)$ and the final time $T > 0$ are given, and $f_\alpha : R^n \to R^n$ are continuously differentiable functions indexed by α in a given finite set $\mathcal{A}$. Each one of the functions f_α corresponds to a mode of the system and hence it is labelled a modal function. The modes' schedule consists of the sequence of deployed modal functions in the interval $[0, T]$ and the timing of switching between consecutive modes. The sequence of modal functions, henceforth called the modal sequence, is characterized by a finite or countable sequence $\{\alpha_1, \alpha_2, \ldots,\}$ of elements in $\mathcal{A}$, and the switching times are denoted by $\tau_1, \tau_2, \ldots$, in increasing order. If the modal sequence is finite, and consisting of $N+1$ modes for some integer N, then there are N switching times among them, $\tau_1, \ldots, \tau_N$, and in this case we use the vector notation $\bar{\tau} := (\tau_1, \ldots, \tau_N)^T \in R^N$ to denote them. Furthermore, we define $\tau_0 := 0$ and $\tau_{N+1} := T$, and observe that the switching-time vector satisfies the following inequalities,

$$0 = \tau_0 \leq \tau_1 \ldots, \leq \tau_N \leq \tau_{N+1} = T. \tag{2}$$

The ith modal function is active during the interval $[\tau_{i-1}, \tau_i)$, and thus, defining the function $F(x, t)$ by

$$F(x, t) \ := \ f_{\alpha_i}(x) \ \text{ for all } t \in [\tau_{i-1}, \tau_i), \ \ i = 1, \ldots, N+1, \tag{3}$$

The authors are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, Georgia, 30332, USA {ding,ywardi,dtaylor,magnus}@ece.gatech.edu

the system's dynamics are defined by the equation

$$\dot{x} \ = \ F(x, t). \tag{4}$$

Let $L : R^n \to R$ be a continuously differentiable function, and consider the cost functional J, defined by

$$J \ = \ \int_0^T L(x)dt. \tag{5}$$

This cost functional is a function of the modes' schedule, and the problem of minimizing it has been studied extensively in the past few years. Most of the published works, such as [3], [7], [8], considered a fixed mode-sequence and devised gradient-descent algorithms for optimizing J with respect to the switching times (and possibly a control input u which is not considered here). A systematic treatment of the discrete, mode-sequencing parameter was carried out in [2], where an inherently NP-hard problem was addressed by variational techniques yielding local minima in a suitable sense. The algorithms proposed there can result in increasing numbers of switching points, and therefore an iteration sequence can approach a sliding-mode control.

The optimal scheduling problem described above arises in a number of application domains; see, e.g., [3] for a survey. This paper is motivated by applications in power electronics, where switching devices must be controlled in order to balance the tracking by a certain current of a given reference signal, with energy-related costs associated with the switching actions. Switching at high frequency may yield a good tracking of the reference signal, but it may result in excessive heating of the switch or in the depletion of an energy source like a battery. We formulate such balancing problems in an abstract setting, where it is desirable to minimize the sum of a tracking cost function in the form of (5) with the sum of energy-related costs associated with the switchings.

Whereas a cost-performance like (5) can be handled by variational techniques, the energy cost has a discrete variable, since it is related to the number of switchings. Such a variable has not been treated in the aforementioned references, and it appears to provide a new challenge to the general problem of optimizing switched-mode systems.

The power-electronics problem we have in mind is described in Section II, and it is then simplified by heuristic considerations. Our algorithmic approach is described in detail in Section III, and numerical examples are provided. Finally, Section IV concludes the paper and suggests directions for future research.

978-1-4244-2079-7/08/$25.00 ©2008 AACC.

II. Example: Switching Circuit

Consider the current-regulation circuit shown in Figure 1, whose purpose is to deliver a specified current level, denoted by $I_{\rm ref}$, to the load inductor. The voltage source has a given constant value $V > RI_{\rm ref}$, and therefore the actual current delivered to the load is regulated by opening and closing the transistor switch. The state of the switch is controlled by the base drive circuit. When the switch is turned on (closed) the diode is reverse biased, and most of the current $i(t)$ flows through the transistor. On the other hand, when the switch is turned off (opened), the diode is forward biased, and most of the current $i(t)$ flows through it.

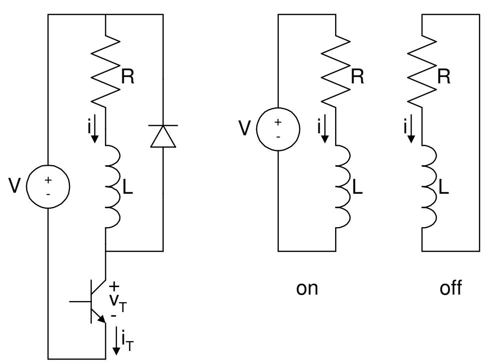

Fig. 1. Switching circuit for current regulation.

In order to study the load current $i(t)$, it is reasonable to assume that the switch is opened and closed instantaneously; in the on state the switch is a short circuit and the current through it is the load current $i(t)$, while the diode is an open circuit. On the other hand, in the off state the switch is an open circuit and the voltage across it is V, while the diode is a short circuit. The two corresponding circuit diagrams, labeled "on" and "off", are shown in Figure 1.

To study the power loss at the switch, we cannot assume that its change-of-state is instantaneous, but rather that it takes a certain finite amount of time due to parasitic junction capacitances in the transistor. We assume that it takes a constant amount of time, t_s seconds, to open or close the switch. During that period the voltage and current at the switch change rapidly, and will be simultaneously large over an interval of time, giving rise to power loss. Moreover, the load current $i(t)$ is continuous in t during a change of state, and hence if t_s is much smaller than the time constant of the circuits, then $i(t)$ can be assumed to have a constant value during a closing or opening of the switch. It must be mentioned that power loss also occurs during conduction, but we chose to neglect it in the present study in order to focus on the salient features of our algorithms when applied to optimal control problems having discrete, discontinuous switching costs.

Consider an opening of the switch which starts at time τ, when the switch is closed, and ends at time $\tau+t_s$, when the switch becomes fully open. Making the common assumption of linear voltage and current during the opening of the switch (see, e.g., [5], Chapter 2, Section 4), we note that first the voltage across the switch, $v_T(t)$, changes in a linear fashion from 0 to V volts, and then the current through the switch, $i_T(t)$, changes in a linear fashion from the load current $i(\tau)$ to 0. This is shown in Figure 2, where the time it takes the voltage to rise is t_m seconds, and the time it takes the current to decline is t_s-t_m seconds. Accordingly, the power dissipating at the switch during its opening, $p(t)$, is given by

$$p(t) \ := \ v_T(t)i_T(t) \ = \\ \begin{cases} \frac{V}{t_m}(t-\tau)i(\tau), & t\in[\tau,\tau+t_m), \\ \big(i(\tau)-\frac{i(\tau)}{t_s-t_m}(t-\tau-t_m)\big)V, & t\in[\tau+t_m,\tau+t_s). \end{cases} \tag{6}$$

Integrating $p(t)$ over the interval $[\tau,\tau+t_s]$ we obtain the energy lost due to the opening of the switch. Denoting it by $q(\tau)$, it follows from (6) after some algebra that

$$q(\tau) \ = \ \frac{1}{2}Vt_si(\tau). \tag{7}$$

This expression for the energy loss is also valid when the switch is being closed, since we assume that both closing and opening the switch takes the same amount of time, t_s seconds. As for the diode, its energy loss can be neglected.

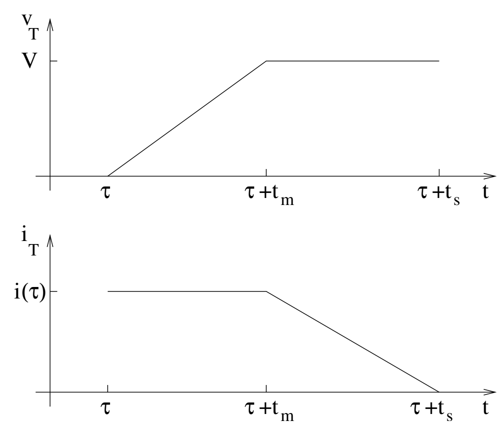

Fig. 2. Voltage and current during switch opening.

The particular circuit we shall consider in the next section has the following parameter values: $R = 5.1\ \Omega$, $L = 10$ mH, $V = 150$ V, and $t_s = 0.5\ \mu$s. By Equation (7), the energy term $q(\tau)$ is given by

$$q(\tau) \ = \ \frac{1}{2}Vt_si(\tau) \ = \ 0.375\times10^{-4}i(\tau). \tag{8}$$

We also point out that the time constant of the circuit is roughly 2 ms, which is two orders of magnitude larger than t_s.

III. Case Study: Optimal Current Regulation

This section considers the current regulation problem for the circuit shown in Figure 1. It can be seen that, when the switch is closed, the steady-state load current is $I_s := V/R$, and when the switch is open, the steady-state load current is 0. Let $I_{\rm ref}\in(0,I_s)$ be a given reference value; the objective of current regulation is to control the switch in order to have the current $i(t)$ track the reference value $I_{\rm ref}$. Perfect tracking, where $i(t) = I_{\rm ref}$ for every t in an open interval, requires a sliding-mode control, and hence an infinite switching cost. Thus, the optimal regulation problem that we consider aims at balancing the tracking performance measure with the cumulative cost associated with the switchings.

This current regulation problem can be transcribed into an optimal control problem with a discontinuous cost functional. The underlying dynamical system consists of the circuit, and its state variable $x(t)$ is the current, namely $x(t) = i(t)$. According to the model set forth in Section II, the state equations have the following forms,

$$\dot{x} \;=\; -\frac{R}{L}x + \frac{V}{L} \;=\; -510x + 15000 \qquad (9)$$

when the switch is closed, and

$$\dot{x} \;=\; -\frac{R}{L}x \;=\; -510x \qquad (10)$$

when the switch is open. It can be seen that $I_s = \frac{V}{R} = 29.4$ A, and we assume, somewhat arbitrarily, that at the initial time $t = 0$ the switch is closed and $x_0 = I_s$ is the initial condition of the state variable. We will track half this value, namely $I_{\rm ref} = 14.7$ A.

Following the notation established in Section I, given a final time $T > 0$, let τ_i, $i = 1, \ldots, N$, be the switching times between the modes associated with the states of the switch, and let $\bar{\tau} := (\tau_1, \ldots, \tau_N)^T$ denote the vector of switching times. Define the cost function $L(x)$ as $L(x) = (x - I_{\rm ref})^2$, and following (5), define the cost functional J by

$$J \;=\; \int_0^T (x - I_{\rm ref})^2 dt; \qquad (11)$$

we chose $T = 0.01$ seconds and, as mentioned earlier, $I_{\rm ref} = 14.7$ A. Obviously $J := J(\bar{\tau})$ is a function of $\bar{\tau}$ by the state equations (9) – (10). Furthermore, define $Q(\bar{\tau})$ by

$$Q(\bar{\tau}) \;=\; \sum_{i=1}^{N} q(\tau_i), \qquad (12)$$

where $q(\tau_i)$ is given by (8). Given a constant $K > 0$, we define the combined cost functional, $W(\bar{\tau})$, by

$$W(\bar{\tau}) \;=\; J(\bar{\tau}) + KQ(\bar{\tau}). \qquad (13)$$

The objective of the optimal regulation problem is to minimize $W(\bar{\tau})$ over all switching schedules. This section describes an algorithm for minimizing W. Its presentation proceeds in the following four steps, corresponding to increasingly more complicated problems.

1) We first consider minimizing only J without regard to Q, and we assume a given fixed number of switching points.
2) We still consider only J, but now we allow the number of switching points to be part of the variable.
3) We consider only minimizing Q without regard to J.
4) We put everything together and minimize $W(\bar{\tau})$ as a function of the switching schedule $\bar{\tau}$.

In order to reduce the cost (whether J, Q, or W) with respect to the continuous timing variable we use a feasible gradient descent technique with Armijo step sizes; see [6] for its details, analysis, and convergence properties. This algorithm performs approximate line minimization along a suitable feasible descent direction, whose choice will be described later.

A. The tracking problem with fixed numbers of switchings

Consider the problem of minimizing J as defined in (11), where the number of times the state of the switch is changed (N) is fixed. Recall that the initial condition is $x_0 = I_s$, and hence the state trajectory depends on the switching-times vector $\bar{\tau} := (\tau_1, \ldots, \tau_N)^T \in R^N$. Since the order of the modes is given, $\bar{\tau}$ must lie in the constraint set defined by Equation (2). When two or more switchings co-occur, namely $\tau_{i-1} = \tau_i$ for some $i = 1, \ldots, N+1$, we consider them as two distinct switchings, since this allows them to diverge at a later time and keeps the problem's dimension (N) a constant.

We applied a feasible direction algorithm with Armijo step sizes, whose descent directions are defined as follows. Let $\bar{\tau}_i$ denote the ith switching schedule computed by the algorithm. Then, define $\tilde{h}_i$ to be the projection of -$\nabla J(\bar{\tau}_i)$ onto the feasible set, as defined by (2). Note that if $\bar{\tau}_i$ lies in the interior of the feasible set then $\tilde{h}_i = -\nabla J(\bar{\tau}_i)$, while if $\bar{\tau}_i$ lies on the boundary of the feasible set, then $\tilde{h}_i$ is similar to the descent direction used by the simplex method in linear programming, which is computationally attractive when the constraint set is a polygon. Now $\tilde{h}_i$ may overshoot the feasible set, and hence we scale it if necessary, and define the descent direction, denoted by h_i, via $h_i := \gamma_i \tilde{h}_i$, where the scaling factor γ_i is defined by

$$\gamma_i \;:=\; \max\{\gamma \in (0, 1] \;:\; \bar{\tau}_i + \gamma \tilde{h}_i \text{ is feasible}\}. \qquad (14)$$

The gradient $\nabla J(\bar{\tau})$ was computed by the following formula (see [3]). Recall the definition of the function $F(x,t)$ in Equation (3). Define the costate $p(t)$ by the equation

$$\dot{p} \;=\; -\Big(\frac{\partial F}{\partial x}(x,t)\Big)^T p - \Big(\frac{\partial L}{\partial x}(x)\Big)^T, \qquad (15)$$

with the boundary condition $p(T) = 0$. Then, for all $i = 1, \ldots, N$,

$$\frac{dJ}{d\tau_i} \;=\; p(\tau_i)^T \big(f_i(x(\tau_i)) - f_{i+1}(x(\tau_i))\big). \qquad (16)$$

We made several runs of the algorithm with various numbers of switchings. The results of a typical run, with $N = 23$, are shown in Figures 3 – 4. The final time is $T = 10$ ms, and the initial iteration consists of equally-spaced switching times. Figure 3 shows the graph of the cost J as a function of the iteration count, and we notice a significant decline in the cost until it flattens out. Figure 4 shows the trajectory of the current at the final iteration, where the horizontal line at $i = 14.7$ indicates $I_{\rm ref}$. These results are not surprising, since we expect the circuit to open immediately in order to let the current go down to $I_{\rm ref}$, and then bounce back and forth around that value at roughly equally-spaced switching times. Several additional runs with different initial schedules yielded similar results.

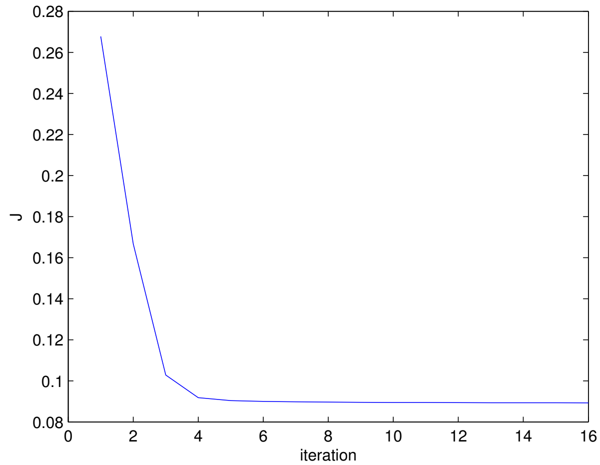

Fig. 3. Problem 1: J vs. iteration count.

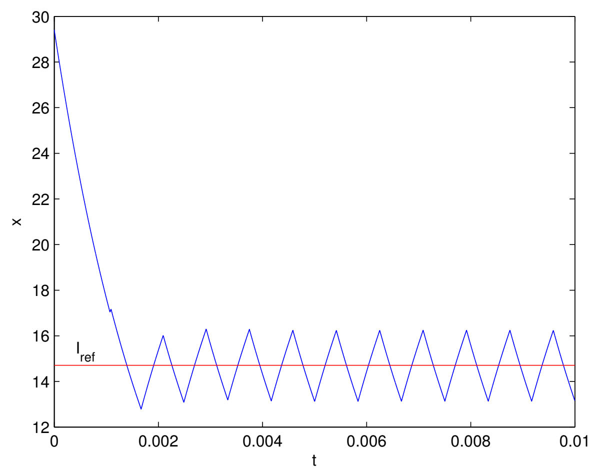

Fig. 4. Problem 1: Current trajectory vs. time.

B. The tracking problem with variable numbers of switchings

Consider the problem of minimizing J as a function of the entire schedule of modes, namely the number and timing of the switchings. Reference [2] has addressed this problem in a general theoretical setting, and devised an algorithm that computes locally-minimal schedules in a suitable sense. To explain that algorithm, consider for a moment a fixed value of N (the number of switchings), and suppose that the algorithm described in the previous subsection has computed a solution point to the problem with the given N. Now consider a time $t \in [0,T]$ that is not a switching time, and denote the modal function at that time by $f(x)$; this modal function is given by the right-hand side of (9) or (10), depending on whether the switch is closed or open at time t. Let f_{-1} denote the other modal function corresponding to the complementary state of the switch. Now suppose that we insert the complementary mode for a brief interval of length δ seconds, centered at time t. As a result we introduce two additional switchings at the times $t-\frac{1}{2}\delta$ and $t+\frac{1}{2}\delta$. Consider the dependence of the cost functional J on δ, and denote it by $\tilde{J}(\delta)$. Now denote by $D_f(t)$ the one-sided derivative of this functional at $\delta=0$, which was shown in [3] to have the following form:

$$D_f(t) \;:=\; \frac{d\tilde{J}}{d\delta^+}(0) \;=\; p(t)^T\big(f_{-1}(x(t))-f(x(t))\big), \quad (17)$$

where $p(t)$ is the costate defined by (15). Observe that if $D_f(t)>0$ then such an insertion would result in a higher cost, while if $D_f(t)<0$ then the insertion would result in a lower cost. In this case the insertion of the two switchings can be done at the same time t, and if an Armijo step follows the insertion, it would separate the two switchings while reducing the cost. The following algorithm makes this point clear.

Algorithm 3.1: Step 0. Fix $N_0>0$, and choose a switching schedule $\bar{\tau}_0$ having N_0 switching times. Set $k=0$.
Step 1. Minimize $J(\bar{\tau})$ under the constraint that the number of switching points must be N_k, by using the algorithm described in the last subsection. Denote the final point by $\bar{\tau}_k$.
Step 2. Fix a fine grid of the time-interval $[0,T]$, denoted by G, and for each point t on the grid, compute $D_f(t)$. Define $G^- := \{t \in G : D_f(t)<0\}$.
Step 3. Make an insertion of the complementary mode at each point $t \in G^-$, and call the resulting schedule $\bar{\tau}_{k+1}$. Let N_{k+1} be the resulting number of switchings. Set $k=k+1$, and goto Step 1.

The results of a typical run are shown in Figures 5 – 6. Figure 5 shows the cost performance J as a function of the iteration count, where by "iteration" we mean a descent according to the Armijo stepsize. It is readily seen that during a run of the algorithm in Subsection A for a fixed number of switching times, the cost goes down until it flattens out. Then an insertion is made, which results in a dramatic decline in the cost. The number of switching times was initially 3, it was 11 at the knee of the curve, and eventually it grew to 150. This number is shown at selective iteration counts in the bottom of the figure. We note that most of the decline in the cost occurred by the time we had 11 switching times, and from this point on the algorithm declined slowly. Figure 6 shows the trajectory of the current at the final iteration point, and it is not surprising to see it bouncing around I_{ref}, in light of the fact that the optimal switching schedule is known to be a sliding mode at that value.

C. Optimizing the switching cost

To further test the descent algorithm with Armijo stepsizes, we apply it to minimizing the switching-cost function $Q(\bar{\tau})$ as defined in (12) and (8). Of course the minimum is obtained when there are no switchings at all, since in this case $Q(\bar{\tau})=0$. However, our objective is to test the gradient-descent algorithm with Armijo stepsizes on problems involving switching costs, and hence this subsection serves to illustrate its efficacy on $Q(\bar{\tau})$ in isolation. Observe that $q(\tau)$ is proportional to $x(\tau)$, and hence the algorithm is expected to drive the switchings to occur at times of low currents. When two switching times co-occur we remove them from

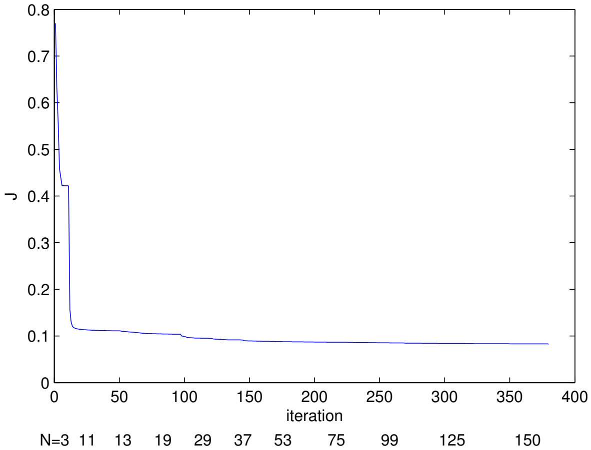

Fig. 5. Problem 2: J vs. iteration count.

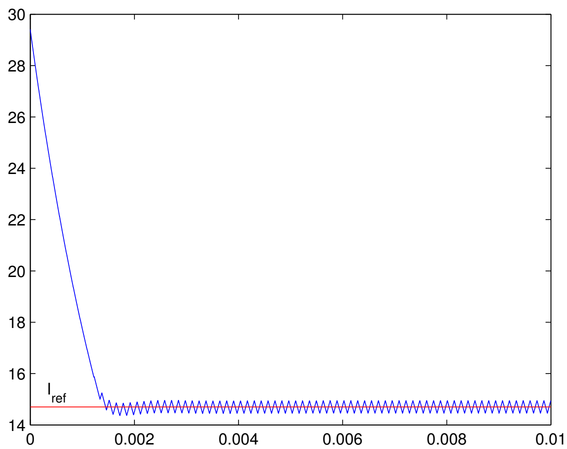

Fig. 6. Problem 2: Current trajectory vs. time.

the schedule and eliminate the mode between them; this is a reasonable step because it results in an instantaneous reduction in the cost Q. Since initially the switch is closed and the current is set to I_s, any small increase in the ith switching time will result in larger currents if i is odd, and in smaller currents if i is even. Consequently, $\frac{dQ}{d\tau_i} > 0$ if i is odd, and $\frac{dQ}{d\tau_i} < 0$ if i is even. If N is even, the algorithm is expected to push pairs of odd and even switching times towards each other until they merge and are eliminated, leaving no remaining switchings. A similar situation arises when N is odd except that τ_1 would move towards 0 and will be the time of the sole remaining switching.

To verify these predictions, we started the algorithm with 51 switchings equally spaced in the time-interval $[0, T]$, and let it run its course from there. The cost declined rapidly over 30 iterations, from 0.035 down to 0.001, while the number of switchings declined from 50 to 1.

D. Minimize the combined tracking/switching cost

At last, we consider the problem of minimizing $W(\bar{\tau})$ as defined by (13), whose solution point is expected to comprise a balance between the tracking performance measure and the cost of switching. The idea is to alternate between the gradient-descent algorithm with Armijo stepsizes, and the insertion of new modes. During the gradient-descent algorithm, it may happen that two switching times coalesce; in this case we cancel the two switchings and eliminate the mode between them. This is a reasonable strategy since it results in reductions in the switching costs while it does not change the tracking performance measure. Mode insertion is more problematic, since it results in an instantaneous rise in the switching cost. In this case, it is hoped that the subsequent gradient descent algorithm would offset this rise. Thus, whenever we make an insertion, we store the switching schedule prior to the insertion, $\bar{\tau}$, and its cost, $W(\bar{\tau})$. We then use the schedule resulting from the insertion as the starting point for the gradient descent procedure, and when that converges to a switching schedule $\bar{\tau}_+$, we compare the cost terms $W(\bar{\tau}_+)$ and $W(\bar{\tau})$. If the former is smaller, we take $\bar{\tau}_+$ as the next point, and otherwise, we return to $\bar{\tau}$ and increase the grid size for the next insertion.

The procedure for selecting insertion times is similar to what was used in subsection B, except for the following two differences. First, the insertion gradient $D_f(t)$ (see (17)) has to take into account not the cost $J(\bar{\tau})$, but the cost $W(\bar{\tau})$, and this results in a modification of (17). Second, instead of inserting a mode at every time-point t on the grid such that $D_f(t) < 0$, we make only a single insertion, at the grid-point t where the insertion gradient is the most negative (if none is negative, there are no insertions). This prevents the addition of too many new modes and a fast rise in the number of switching times, which could slow down the gradient-descent algorithm. We point out that this change from the procedure in subsection B amounts to an ad-hoc heuristic, and no comparative study between the two has been made.

All of this is put together in a formal way by the following algorithm.

Algorithm 3.2: Data. A sequence of progressively finer grids of the interval $[0, T]$, denoted by $\{G_j\}_{j=1}^{\infty}$.
Initialize. Pick an integer $N_0 > 0$, and let the initial schedule $\bar{\tau}^0$ consist of N_0 equally-spaced switchings in the interval $[0, T]$. Adopt the first grid G_1.
Step 0. Set $k = 0$, and define $W(\bar{\tau}_{-1}) := \infty$.
Step 1. Starting with $\bar{\tau}^k$, perform a gradient descent algorithm with Armijo stepsizes until convergence is noted. Whenever two switching times coalesce, eliminate them and the mode between them, thereby reducing N_k by 2. Denote the resulting schedule by $\bar{\tau}_k$.
Step 2. If $W(\bar{\tau}_k) - W(\bar{\tau}_{k-1}) > 0$, then set $\bar{\tau}_k = \bar{\tau}_{k-1}$, and take the next-finer grid in the grid-sequence.
Step 3. With the schedule $\bar{\tau}_k$, compute the insertion gradients $D_f(t)$ for every t on the grid. If $D_f(t) \geq 0$ for every t on the grid, then take the next-finer grid in the grid-sequence, and repeat Step 3.

Step 4. Perform a mode insertion, set $N_{k+1} = N_k + 2$, and denote the resulting schedule by $\bar{\tau}^{k+1}$. Set $k = k + 1$, and goto Step 1.

A few remarks are due.

I. In Step 1, convergence of the gradient-descent algorithm is noted whenever either the gradient $\nabla W(\bar{\tau})$, or the Armijo stepsize, is small enough. By "small enough" we mean that $||\nabla W(\bar{\tau})|| \leq 1.0$, or the Armijo stepsize is under 10^{-16}. The former bound may raise an eyebrow since 1.0 does not appear to be too small. However, due to the stiffness of the circuit and the associated optimization problem, at the initial stages of the gradient-descent algorithm the magnitude of the gradients typically is in the order of 500 – 800, and the Armijo stepsize is in the order of 10^{-5}.

II. The grids we chose are equally spaced in the interval $[0, T]$; the first grid has 10 points, and each successive grid has 50% more points than the previous one.

We tested the algorithm on $W(\bar{\tau}) := J(\bar{\tau}) + KQ(\bar{\tau})$, with $K = 30$, and $K = 5$. The results for $K = 30$ are shown in Figures 7 – 8 for $N_0 = 3$ and equally-spaced initial switchings; additional runs with $N_0 = 9$ and $N_0 = 15$ yielded similar final-cost results. As seen in Figure 7, the cost declined until it flattened, and after 11 iterations we performed an insertion, marked by a rise in the cost. The cost then declined, until another insertion was made at iteration 24. Finally, convergence was obtained with 7 switchings, and the resulting cost was 0.247. The graph of the state trajectory at the final point is shown in Figure 8, where we note that the first switching time is $t = 0$, and subsequent switchings occur at the times of sharp turns in the graph.

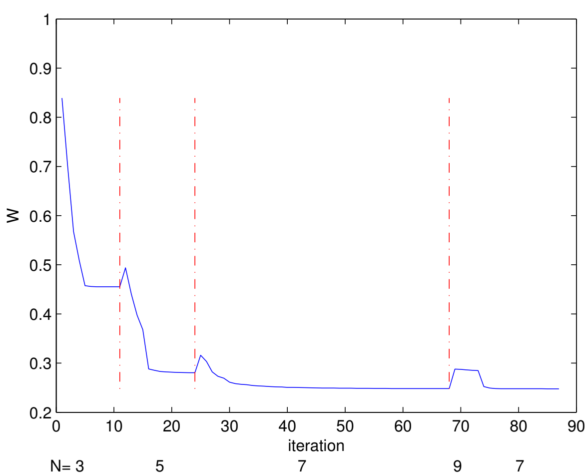

Fig. 7. Problem 4: W vs. iteration count, $N_0 = 3$.

IV. CONCLUSIONS

This paper addresses the problem of optimizing the switching schedule of a current regulator circuit so as to minimize a weighted sum of a tracking performance functional and a switching-related cost measure. It proposes an algorithmic framework consisting of variational principles for the tracking performance functional, and ad-hoc heuristics for handling the switching costs. We tested the algorithmic framework on a generic problem, and found it to be effective for computing locally optimal minima. Encouraged by the obtained numerical results, it is our plan to apply the proposed framework to realistic, concrete problems in power regulation, as well as to more general problems in power electronics.

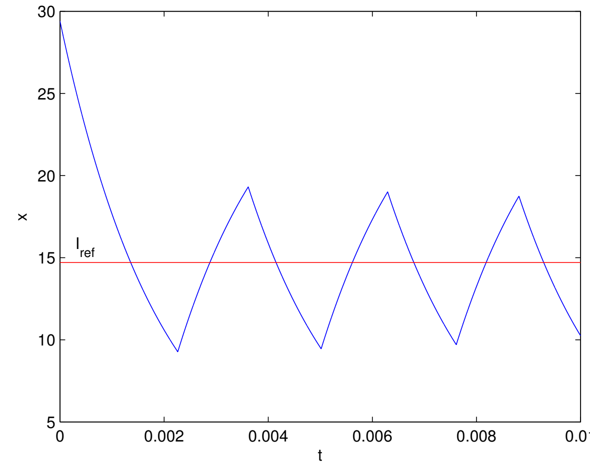

Fig. 8. Problem 4: Current trajectory at the final iteration.

V. ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation Under Grant Number 0509064.

REFERENCES

[1] S.A. Attia, M. Alamir, and C. Canudas de Wit. Sub Optimal Control of Switched Nonlinear Systems Under Location and Switching Constraints. *Proc. 16th IFAC World Congress*, Prague, the Czech Republic, July 3-8, 2005.

[2] H. Axelsson, Y. Wardi, M. Egerstedt, and E.I. Verriest. Gradient Descent Approach to Optimal Mode Scheduling in Hybrid Dynamical Systems. *Journal of Optimization Theory and Applications*, Vol. 136, pp. 167-186, 2008.

[3] M. Egerstedt, Y. Wardi, and H. Axelsson. Transition-Time Optimization for Switched Systems. *IEEE Transactions on Automatic Control*, Vol. AC-51, No. 1, pp. 110-115, 2006.

[4] S. Hedlund and A. Rantzer. Optimal Control of Hybrid Systems. *Proc. IEEE Conf. on Decision and Control*, pp. 3972-3977, 1999.

[5] N. Mohan, T.M. Undeland, and W.P. Robbins. *Power Electronics: Converters, Applications and Design*, 2nd Edition, John Wiley & Sons, New York, NY, 1995.

[6] E. Polak. *Optimization Algorithms and Consistent Approximations*, Springer-Verlag, New York, NY, 1997.

[7] M.S. Shaikh and P. Caines. On Trajectory Optimization for Hybrid Systems: Theory and Algorithms for Fixed Schedules. *Proc. IEEE Conf. on Decision and Control*, pp. 1997–1998, 2002.

[8] X. Xu and P.J. Antsaklis. Optimal Control of Switched Systems via Nonlinear Optimization Based on Direct Differentiations of Value Functions. *Int. J. of Control*, Vol. 75, pp. 1406-1426, 2002.