Preprint of Paper ThA17.2 to be presented at the 2014 American Control Conference, Portland, OR, June 4 - 6, 2014.

A Hybrid Model Predictive Control Strategy for Optimizing a Smoking Cessation Intervention

Kevin P. Timms[1], Daniel E. Rivera[2], Megan E. Piper[3], and Linda M. Collins[4]

***Abstract*— The chronic, relapsing nature of tobacco use represents a major challenge in smoking cessation treatment. Recently, novel intervention paradigms have emerged that seek to adjust treatments over time in order to meet a patient's changing needs. This article demonstrates that Hybrid Model Predictive Control (HMPC) offers an appealing framework for designing these optimized, time-varying smoking cessation interventions. HMPC is a particularly appropriate approach as it recognizes that intervention doses must be assigned in predetermined, discrete units while retaining receding-horizon, constraint-handling, and combined feedback and feedforward capabilities. Specifically, an intervention algorithm is developed here in which counseling and two pharmacotherapies are manipulated to reduce daily smoking and craving levels. The potential usefulness of such an intervention is illustrated through simulated treatment of a quit attempt in a hypothetical patient, which highlights that prioritizing reduction in craving over total daily smoking levels significantly reduces craving levels, suppresses relapse, and successfully rejects time-varying disturbances such as stress, all while adhering to several practical operational constraints and resource use considerations.**

I. INTRODUCTION

After approximately 40 years of consistent improvement, decreasing smoking rates in the U.S. have stalled, where one in five adults smoke [1]. This is largely a consequence of the chronic, relapsing nature of tobacco use, which also leads to a quit attempt failure rate around 90% [2].

Seeking to support successful cessation attempts, many smoking interventions that are behavioral or pharmacological in nature have been designed and implemented. For example, 10 minute phone counseling sessions are shown to be effective, especially when centered around problem solving skills training and general encouragement [2]. Sustained-release bupropion (Zyban®, GlaxoSmithKline) has been shown to reduce nicotine withdrawal symptoms and is also commonly used as a cessation treatment; patients are currently prescribed a dose of 150 mg/day for 3 days starting 1-2 weeks prior to target quit day (TQD) before increasing to 300 mg/day for up to 12 weeks post-TQD [2], [3]. Similarly, the nicotine lozenge is a form of nicotine replacement therapy that patients take as needed starting on TQD (up to 20 lozenges/day) and delivers low doses of nicotine to the bloodstream through tissues in the mouth [2]. However, even the most effective combinations of medications have six month success rates below 35% [2], [4]. The limited success of current intervention protocols is of particular concern given the fact that tobacco use remains the greatest cause of preventable death in the U.S. [1], [2].

Novel paradigms in behavioral medicine seek to address the chronic, relapsing nature of smoking [2]. Contrasting current widespread treatment practices in which patients are assigned a single dose of a therapeutic component for the duration of the quit attempt, *adaptive smoking cessation interventions* seek to adjust treatment components over time based on a patient's changing needs. Specifically, adaptive interventions involve *decision policies* that define how an intervention component dosage (e.g., counseling frequency) should be adjusted over time based on personalized *tailoring variables* (e.g., patient-reported withdrawal symptoms) [5]. The feedback nature of this approach means adaptive interventions consist of closed-loop systems, which highlights the opportunity for introduction of control systems engineering principles into design of behavioral health interventions [6], [7]. Furthermore, employment of control theory to design novel smoking cessation treatment strategies may directly address the lack of a well-accepted algorithmic framework for optimal assignment of intervention components to individual patients, which is explicitly cited by the U.S. Department of Health and Human Services as a factor limiting the efficacy of existing treatment protocols [2].

This article proposes an adaptive smoking cessation intervention in the form of an algorithm featuring a Hybrid Model Predictive Control (HMPC) formulation that is intended to lay the initial conceptual and computational groundwork for future experimental trials and clinical utility. HMPC is employed in this problem setting as it facilitates treatment optimization in terms of efficacy and resource use, while adhering to patient care constraints. Specifically, an HMPC-based intervention is presented here that seeks to optimally adjust manipulated variables with categorical units—i.e., treatment component dosages—in an attempt to reduce cigarettes smoked per day (*CPD*) and average daily craving (*Craving*) levels over time. Furthermore, predictive control facilitates "just-in-time" interventions. In behavioral medicine, just-in-time interventions seek to avoid challenges associated with real-time measurement and adaptation (e.g.,

*This work was supported by the Office of Behavioral and Social Sciences Research (OBSSR) of the National Institutes of Health (NIH) and the National Institute on Drug Abuse (NIDA) through grants K25 DA021173, R21 DA024266, P50 DA10075, and F31 DA035035.

[1]Biological Design Program, Arizona State University, Tempe, AZ ktimms@asu.edu

[2]Control Systems Engineering Laboratory, School for Engineering of Matter, Transport and Energy, Arizona State University, Tempe, AZ daniel.rivera@asu.edu

[3]Department of Medicine, Center for Tobacco Research and Intervention, University of Wisconsin, Madison, WI MEP@ctri.medicine.wisc.edu

[4]The Methodology Center and Department of Human Development and Family Studies, Pennsylvania State University, State College, PA lmcollins@psu.edu

minute-to-minute measurement and dosage reassignment); instead, adaptation is said to be done in a timeframe that will be effective, but minutes to days after a patient's needs actually change [8], [9]. The three treatment components included in this novel intervention design are weekly telephone counseling, bupropion, and nicotine replacement lozenges. Furthermore, the control-based intervention design problem addressed here features multiple cessation objectives, a combination of behavioral and pharmacological treatment components, time delays in treatment effects, manipulated variable switching limits, and anticipation of both set point and disturbance changes.

The paper is organized as follows: Section II outlines the intervention design problem; Section III formulates the controller, which is then examined through simulation of a quit attempt of a hypothetical patient predisposed to quit attempt failure in Section IV; Section V describes possible directions for future work.

II. Intervention Structure, Components, & Constraints

A. Smoking Cessation Intervention Structure

An optimized smoking cessation intervention proposed in this work employs the general structure in Fig. 1. The controller-design problem presented in subsequent sections is based in the following considerations:

- An adaptive smoking intervention should personalize treatment over time in order to support abstinence directly as well as manage factors that contribute to relapse. Previous work, [7], has shown that there is a self-regulatory relationship between *CPD* and *Craving* during a quit attempt, suggesting *Craving* reduction can support and sustain cessation success. Consequently, both *CPD* and *Craving* are controlled variables, where the overall treatment goals are $CPD = 0$ and *Craving* reduction.
- The intervention should entail just-in-time treatment adaptation. The predictive and receding-horizon nature of Model Predictive Control (MPC) offer appealing capabilities for just-in-time intervention designs, as these features can determine optimal control action that addresses deficiencies in current outcomes as well as those predicted for future time points.
- Treatment adjustment based on changing patient needs is facilitated by feedback of the measured *CPD* and *Craving* outcomes each day.
- Three manipulated variables consisting of a combination of behavioral and pharmacological treatments are considered: 10 minute phone counseling sessions, sustained-release bupropion, and nicotine replacement lozenges. These treatments are available in predetermined, discrete dosages [2]. Consequently, the appealing aspects of MPC must be implemented in the context of a hybrid linear dynamical system [10].
- Inclusion of feedforward signals can diminish the harmful effects of exogenous influences on the quit attempt. Feedforward components are a unique aspect of the intervention presented here. Notably, the patient's transition from not trying to quit smoking to trying to quit (*Quit*, represented by as a step from 0 to 1 on TQD) is treated as a measured and anticipated disturbance, which actually contributes to the goal of $CPD = 0$ for a period of time. *Stress* is considered to be a second measured and anticipated disturbance.

B. Open-Loop Models for Smoking Cessation Treatment

The open-loop models used in this work represent an extension of modeling efforts documented in prior work; [11] developed dynamical systems models that describe a self-regulatory relationship between *CPD* and *Craving* during a quit attempt, as captured in data collected via self-reports through a mobile technology as part of a 400-participant smoking cessation clinical trial [3], [7]. The models from these efforts, as well as additional clinical trial data, inform the dose-response and behavioral open-loop models that act as the basis for controller-design.

The overall open-loop dynamics of the cessation problem considered here are represented mathematically in (1).

$$\begin{bmatrix} CPD \\ Craving \end{bmatrix} = \begin{bmatrix} P_{cpd_c} & P_{cpd_b} & P_{cpd_l} \\ P_{crav_c} & P_{crav_b} & P_{crav_l} \end{bmatrix} \begin{bmatrix} u_c \\ u_b \\ u_l \end{bmatrix} + \begin{bmatrix} P_{cpd_Q} & P_{cpd_S} \\ P_{crav_Q} & P_{crav_S} \end{bmatrix} \begin{bmatrix} Quit \\ Stress \end{bmatrix} \quad (1)$$

where u_c, u_b, and u_l are the counseling, bupropion, and lozenge dosage inputs; *Quit* and *Stress* are measured disturbances; P_{cpd_c}, P_{cpd_b}, P_{cpd_l}, P_{cpd_Q}, and P_{cpd_S} are the models describing the response of *CPD* to changes in u_c, u_b, u_l, *Quit*, and *Stress*, respectively, for a hypothetical patient predisposed to quit attempt failure; P_{crav_c}, P_{crav_b}, P_{crav_l}, P_{crav_Q}, and P_{crav_S} are representations of the open-loop dynamics corresponding to the *Craving* outcome for the hypothetical patient. Informed by prior work (see [7]), clinical trial data when possible (see [3]), and additional knowledge of dynamic relationships between smoking variables, the components of (1) are represented as continuous-time models below.

The dose-response models that describe the effects of one 10 minute phone counseling session on both *CPD* and *Craving* employ a single transfer function structure:

$$P_{y_c}(s) = K_{p_c} \frac{\tau_{a_c} s + 1}{\tau_c^2 s^2 + 2\tau_c \zeta_c s + 1} \quad (2)$$

where y is either *CPD* or *Craving* and the parameters for the $P_{cpd_c}(s)$ and $P_{crav_c}(s)$ expressions are found in Table I. Drawing from clinical trial data [3], both $P_{cpd_b}(s)$ and $P_{crav_b}(s)$ employ the structure in (3).

$$P_{y_b}(s) = K_{p_b} \frac{\tau_{a_b} s + 1}{\tau_b^2 s^2 + 2\tau_b \zeta_b s + 1} e^{-\theta_b s} \quad (3)$$

Table I contains the parameter values corresponding to the effect on *CPD* and *Craving* per-150 mg dose of bupropion.

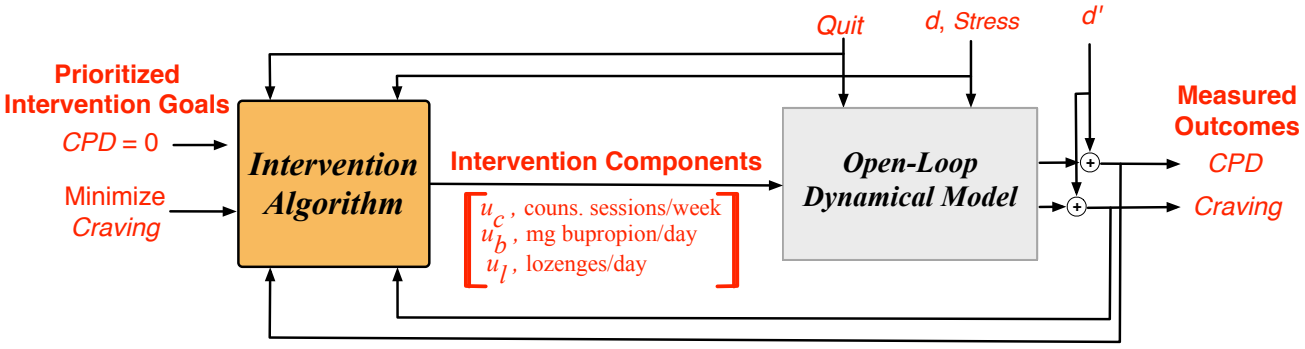

Fig. 1. General structure of the proposed control-oriented approach to designing adaptive smoking cessation interventions.

TABLE I

PARAMETERS OF THE TRANSFER FUNCTION MODELS USED IN DEVELOPMENT OF AN HMPC-BASED SMOKING CESSATION INTERVENTION.

Model	Parameters	Parameter Values
P_{cpd_c}	$K_{p_c}, \tau_{a_c}, \tau_c, \zeta_c$	-0.26, 18.00, 4.24, 1.30
P_{crav_c}	$K_{p_c}, \tau_{a_c}, \tau_c, \zeta_c$	-0.77, 13.50, 3.67, 1.43
P_{cpd_b}	$K_{p_b}, \tau_{a_b}, \tau_b, \zeta_b, \theta_b$	-4.34, 2.50, 1.50, 1.00, 3.00
P_{crav_b}	$K_{p_b}, \tau_{a_b}, \tau_b, \zeta_b, \theta_b$	-7.69, 2.00 , 1.10 1.00, 3.00
P_{cpd_l}	$K_{p_l}, \tau_{a_l}, \tau_l$	-0.12, 0.44, 0.22
P_{crav_l}	$K_{p_l}, \tau_{a_l}, \tau_l$	-0.25, 0.44, 0.22
P_{cpd_Q}	$K_{p_Q}, \tau_{a1_Q}, \tau_{a2_Q}, \tau_Q, \zeta_Q$	-0.24, 3.45, 90.53, 5.10, 0.59
P_{crav_Q}	$K_{p_Q}, \tau_{a1_Q}, \tau_{a2_Q}, \tau_Q, \zeta_Q$	0.60, 3.45, 90.53, 5.10, 0.59
P_{cpd_S}	$K_{p_S}, \tau_{a_S}, \tau_S, \zeta_S$	0.50, 0.75, 0.55, 0.49
P_{crav_S}	$K_{p_S}, \tau_{a_S}, \tau_S, \zeta_S$	0.70, 0.75, 0.55, 0.49

The dose-response models describing the effect of u_l on *CPD* and *Craving* both employ the following equation structure:

$$P_{y_l}(s) = K_{p_l}\frac{\tau_{a_l}s+1}{\tau_l s+1} \tag{4}$$

Nicotine lozenges deliver nicotine through the bloodstream—a slower delivery method compared to cigarette smoking—which contributes to patient statements that a lozenge is less satisfying than a cigarette. However, nicotine delivery in this form is relatively rapid compared to timeframe on which this article focuses, with a single lozenge lasting less than one hour [4]. These effect speed considerations are reflected in Table I, as is the relatively modest effect of a unit dose of lozenge.

Quit is considered to be the primary measured disturbance in this controller design problem. Section IV simulates an intervention for a hypothetical smoker who is predisposed to quit attempt failure; that said, the response of *CPD* and *Craving* to changes in *Quit* for this hypothetical patient are similar to those observed for one participant of a previously examined clinical trial (see [3]) and modeled in [7]. As depicted in Fig. 2 (dotted grey), the failed quit attempt is characterized by initial quit success, $CPD = 0$ on TQD, and a simultaneous increase in *Craving*; this is followed by full resumption of smoking, with *CPD* and *Craving* returning to approximately pre-TQD levels. The models describing the response of *CPD* and *Craving* to a quit attempt both employ the equation structure in (5), with the parameters corresponding to each set of dynamics found in Table I.

$$P_{y_Q}(s) = K_{p_Q}\frac{(\tau_{a1_Q}s+1)(\tau_{a2_Q}s+1)}{\tau_Q^2 s^2+2\tau_Q\zeta_Q s+1} \tag{5}$$

Our approach also considers *Stress* as an independent, exogenous influence on the cessation process. These effects are represented by the following transfer function structure and corresponding parameters in Table I:

$$P_{y_S}(s) = K_{p_S}\frac{\tau_{a_S}s+1}{\tau_S^2 s^2+2\tau_S\zeta_S s+1} \tag{6}$$

C. Controller Constraints

The use of counseling, bupropion, and nicotine replacement lozenges, as well as the fact that the system to be intervened upon is human health, implies that certain constraints must be imposed on the actions of the controller. The constraints to be imposed are described in general terms in the following with equivalent mathematical representations defined in Section III.

Medication toxicity limits cap possible dosing levels: maximum bupropion and lozenge doses are 2 doses/day and 20 lozenges/day, respectively, with lower bounds of 0 doses. The medically-necessitated ramp-up to one 150 mg/day bupropion dose before increasing to 300 mg/day translates to a 1-dose move size constraint for increases in u_b, although there are no medically-necessitated limitations on decreasing u_b dose. A weekly switching time for u_b is also imposed on the controller to avoid unreasonably frequent bupropion dose adjustment. A 5 lozenge move size limit is also considered.

As the combination of lozenge and bupropion is one of the most effective cessation pharmaco-therapies [4], the controller is further constrained such that nonzero lozenge doses can only be assigned when bupropion dose is also nonzero. However, no constraints to lozenge dosing frequency are imposed as long as the bupropion dose condition is met.

In terms of constraints on u_c, resource use and logistical concerns associated with delivering counseling sessions suggests the controller should allow only one dose of counseling be delivered on a given day, and at a maximum frequency

of weekly counseling. Consequently, there are time-varying limits and move size constraints associated with u_c.

Upper bounds on the measured outcomes are also imposed so as to ensure that dosage decisions that would lead to excessive smoking or craving are penalized. Lower outcome boundaries that correspond to physically-realizable *CPD* and *Craving* values are not defined as constraints, but are imposed in simulation (e.g., no fewer than 0 cigarettes can be smoked in a day); this reflects physically-realizable measurements, but avoids any hesitation in decision-making as the controlled variables near 0.

III. Controller Development

In behavioral health settings, treatment component dosages are often categorical in nature. Ensuring manipulated variable assignments adhere to predetermined discrete levels means the algorithmic structure of a control-oriented intervention must consider hybrid systems [10]. Specifically, the cessation system considered in this article is described in a mixed logical and dynamical (MLD) structure. An MLD representation of linear hybrid systems consists of real and integer states, inputs, and constraints [12]:

$$x(k+1) = Ax(k) + B_1u(k) + B_2\delta(k) + B_3z(k) + B_dd(k) \quad (7)$$

$$y(k) = Cx(k) + d'(k) + \nu(k) \quad (8)$$

$$E_2\delta(k) + E_3z(k) \leq E_5 + E_4y(k) + E_1u(k) - E_dd(k) \quad (9)$$

where x and u are system states and inputs (both continuous and discrete) and y are the outputs; d, d', and ν are measured disturbances, unmeasured disturbances, and measurement noise, respectively; δ and z are discrete and continuous auxiliary variables that facilitate conversion of logical/discrete decisions into the corresponding linear inequality constraints. Details of the controller formulation developed for systems represented in the MLD structure can be found in [10].

A cost function accounting for cessation success, resource use, and adherence to operational constraints is in quadratic form:

$$\begin{aligned} J \triangleq & \sum_{i=1}^{p} ||y(k+i) - y_r)||^2_{Q_y} + \sum_{i=0}^{m-1} ||(\Delta u(k+i) - u_r)||^2_{Q_{\Delta u}} \\ & + \sum_{i=0}^{m-1} ||(u(k+i) - u_r)||^2_{Q_u} + \sum_{i=0}^{p-1} ||(\delta(k+i) - \delta_r)||^2_{Q_d} \\ & + \sum_{i=0}^{p-1} ||(z(k+i) - z_r)||^2_{Q_z} \end{aligned} \quad (10)$$

where p is the prediction horizon, m is the control horizon, Q_i is the matrix weight of the vector 2-norm of variable i, and the r subscript indicates a reference. Optimization consists of finding the control action sequences for u, δ, and z that minimize J according to:

$$\min_{\{[u(k+i)]_{i=0}^{m-1}, [\delta(k+i)]_{i=0}^{p-1}, [z(k+i)]_{i=0}^{p-1}\}} J \quad (11)$$

subject to the mixed integer constraints in (9):

$$y_{min} \leq y(k+i) \leq y_{max}\ ,\ 1 \leq i \leq p \quad (12)$$

$$u_{min} \leq u(k+i) \leq u_{max}\ ,\ 0 \leq i \leq m-1 \quad (13)$$

$$\Delta u_{min} \leq \Delta u(k+i) \leq \Delta u_{max}\ ,\ 0 \leq i \leq m-1 \quad (14)$$

In this problem setting, constraints that result from logistical, resource availability, and medically-necessitated restrictions must be imposed (described in Section II). The manipulated variable bounds, move size limits, and move frequency limits (i.e., switching times) are represented mathematically as:

$$\begin{aligned} 0 \leq u_c(k) \leq 0\ ,\ & k \neq 1 + n\,T_{sw} \\ 0 \leq u_c(k) \leq 1\ ,\ & k = 1 + n\,T_{sw}\ ,\ n = 0,1,2,... \end{aligned} \quad (15)$$

$$0 \leq u_b(k) \leq 2 \quad (16)$$

$$0 \leq u_l(k) \leq 20 \quad (17)$$

$$-1 \leq \Delta u_c(k) \leq 1 \quad (18)$$

$$\begin{aligned} 0 \leq \Delta u_b(k) \leq 0\ ,\ & k \neq 1 + n\,T_{sw} \\ -2 \leq \Delta u_b(k) \leq 1\ ,\ & k = 1 + n\,T_{sw}\ ,\ n = 0,1,2,... \end{aligned} \quad (19)$$

$$-5 \leq \Delta u_l(k) \leq 5 \quad (20)$$

where T_{sw} is switching time. (Additional details on addressing switching restrictions in control-based behavioral health interventions can be found in [13].) To ensure dosage assignments that may move patient smoking and craving to overly excessive values, maximum bounds that are approximately 50% larger than baseline levels are also defined for the controlled variables:

$$-\infty \leq CPD(k) \leq 14.25 \quad (21)$$

$$-\infty \leq Craving(k) \leq 24.76 \quad (22)$$

An HMPC decision framework calculates a control sequence by assigning discrete dosages of each intervention component at each sample time, where all components are initially off. The discrete nature of the system's manipulated variables are represented as:

$$\delta_i(k) = 1 \Leftrightarrow z_i(k) = i; \quad i \in \{0,1\} \quad (23)$$

$$u_c = \sum_{i=0}^{1} z_i(k) \qquad \sum_{i=0}^{1} \delta_i(k) = 1 \quad (24)$$

$$\delta_j(k) = 1 \Leftrightarrow z_j(k) = j-2; \quad j \in \{2,3,4\} \quad (25)$$

$$u_b = \sum_{j=2}^{4} z_j(k) \qquad \sum_{j=2}^{4} \delta_j(k) = 1 \quad (26)$$

$$\delta_k(k) = 1 \Leftrightarrow z_k(k) = k-5; \quad k \in \{5,...,26\} \quad (27)$$

$$u_l = \sum_{k=5}^{26} z_k(k) \qquad \sum_{k=5}^{26} \delta_k(k) = 1 \quad (28)$$

Equations (23) and (24) state that one of two possible levels of counseling must be assigned for each sample—reflecting that either 0 or 1 counseling sessions are delivered on a given day. Equations (25) and (26) state that one of three levels (0, 1, or 2 doses) of bupropion must be assigned for each sample. Equations (27) and (28) state that u_l must assume one of 21 possible levels (0 to 20 lozenges) for each sample.

IV. Intervention Simulation

To demonstrate the potential utility of a control-oriented smoking cessation intervention, the controller formulation described in the previous sections is implemented in the context of a quit attempt made by a hypothetical patient. This patient smokes 9.25 cigarettes daily (the *CPD* scale is assumed to be continuous in this simulation) and has an average daily craving level of 16.40. Represented in the continuous-time model structures and parameters in (5) and Table I, this hypothetical patient is predisposed to quit attempt failure. This is clearly demonstrated in the intervention-free quit attempt simulation (dotted grey in Fig. 2, where the only nonzero input is *Quit*).

In the intervention simulation, no treatment components are active initially. As the goal of cessation is to eliminate smoking and ultimately the desire to smoke, $CPD = 0$ and $Craving = 0$ set points initialize on TQD. *Quit* and set point changes are all anticipated, where this measured disturbance and the predicted set point sequences are adjusted over the prediction horizon for each sample relative to TQD. For simplicity, *Stress* is not introduced until day 45 and is modeled as Gaussian noise with magnitudes between 0 and 2. Anticipation of this disturbance consists of assuming *Stress* will equal the most recent measured value for the duration of the prediction horizon.

The controller begins making dosage decisions beginning on day 1, 14 days prior to the *Quit* and set point changes. Switching limitations associated with counseling and bupropion, mathematically quantified in (15) and (19), are imposed such that the first day for which u_c and u_b control action can occur is day 4. The restriction of dose reassignments for counseling and bupropion means u_c and u_b switching is limited to days 4, 11, 18, and so on. This contrasts u_l, for which dosages can be adjusted on a daily basis once $u_b \neq 0$.

The parameters of the simulated HMPC-based intervention are as follows: $p = 30$, $m = 8$, $Q_{cpd} = 1$, $Q_{craving} = 100$, $Q_c = Q_b = 0$, and $Q_l = 1$. Despite the fact that this hypothetical patient is predisposed to quit attempt failure, the adaptive intervention is able to facilitate a successful quit attempt, as depicted with the dashed blue responses in Fig. 2.

With a 30 day prediction horizon, the very first control decision considers outcome predictions through TQD and the first weeks of the quit attempt. With an 8 day control horizon, controller predictions consider at least the very next time at which u_c and u_b adjustments can be made. Contrasting current treatment protocols that implement counseling and full bupropion dosing prior to TQD [2], counseling and the maximum bupropion dose are not implemented in this simulation until 3 days into the quit attempt. Initial hesitation of manipulated variable adjustments may suggest the controller seeks to avoid significant overshoot around TQD, as the *Quit* step induces an immediate 9.25 cigarette reduction of smoking and no lower limit on *CPD* is imposed on the controller (although it is included in simulation). Regardless, the initial dose of bupropion lead to a reduction in both *CPD* and *Craving* prior to the quit attempt. Furthermore, aggressive u_l dosing beginning around TQD and aggressive u_c and u_b dosing starting on day 18 promotes set point tracking in *CPD* for most of the simulation. These dosage assignments similarly promote reduction in *Craving* below baseline levels for even the first five days of the quit attempt, which is entirely suppressed within one week of TQD.

Intuitively, *CPD* is directly indicative of a smoking cessation intervention's efficacy. Interestingly, though, the hundredfold greater weight on the *Craving* set point tracking penalty compared to that for *CPD* during daily decision-making ($Q_{cpd} = 1$ vs. $Q_{craving} = 100$) leads to a favorable *CPD* response. Specifically, the maximum post-TQD smoking is less than one cigarettes ($CPD = 0.50$ on day 30).

In terms of a concern for resource use, the controller's aggressive lozenge assignment begins to diminish on day 21—after the major spike in *Craving* in the intervention-free scenario would have subsided—leading to maximum lozenge dose assignments for under one week. Setting $Q_l = 1$ helps facilitate the lozenge dose reduction overall as compared to an alternate simulation (not shown) in which $Q_l = 0$; u_l dose in the alternate simulation is approximately double that depicted in Fig. 2 after day 27, despite having very similar outcomes. As seen in days 45 through 70 in Fig. 2, a nonzero penalty on u_l does not significantly reduce the controller's ability to reject the measured *Stress* disturbance; the controller's reintroduction of counseling after three counseling-free weeks and small adjustments to u_l in response to *Stress* measurements actually make the *Stress* disturbance virtually undetectable in the outcomes.

V. Summary & Future Work

This article illustrates an HMPC structure that offers a potentially valuable framework for adaptive smoking cessation interventions. The potential of this approach is significant given smoking's continued status as a major public health issue and the fact that over 90% of attempts to quit smoking fail and [1], [2]. Specifically, an HMPC algorithm formulated to manage MLD systems is employed in order to manage the discrete nature of the manipulated variables, while preserving the receding-horizon, feedback, and feedforward features of MPC [10]. Simulation shows that this decision framework adjusts counseling, bupropion, and nicotine replacement lozenge doses over time in a manner that promotes successful cessation in a hypothetical patient otherwise predisposed to quit attempt failure; these dosage adjustments are optimized according to measured and predicted intervention efficacy and resource use criteria, while adhering to a variety of operational constraints.

Ultimately, validation of an HMPC-based smoking cessation intervention requires novel clinical trials—a significant undertaking. However, future work is also required to build upon the valuable proof-of-concept established here in a clinically-meaningful way. Notably, standard MPC draws from a representation of a plant in discrete-time state-space form—where there are no direct feedthrough terms in the measurement equation. However, day-to-day changes

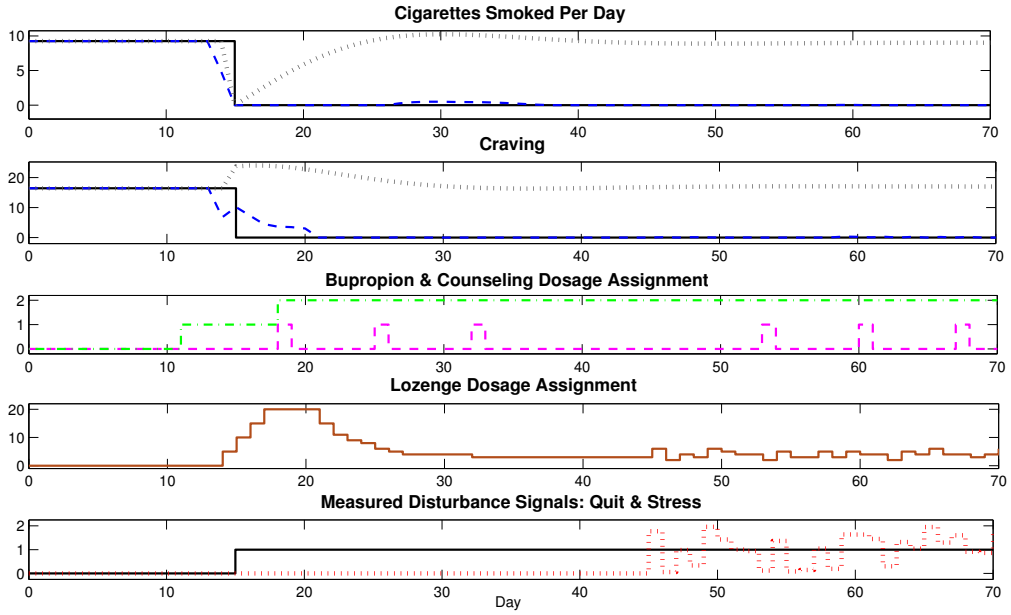

Fig. 2. Simulation of a quit attempt made by a hypothetical patient, where TQD is day 15: dotted grey (top two plots) depicts the simulated intervention-free response of *CPD* and *Craving* to a quit attempt (*Quit* [solid black, bottom plot] is the only input); dashed blue depicts the responses resulting from the HMPC-based intervention, which implements the counseling (dashed magenta), bupropion (dash-dot green), and lozenge (solid brown) dosing assignments in response to the set point changes (solid black, top two plots), *Quit* (solid black, bottom plot), and *Stress* (dotted red, bottom plot).

in *CPD* and *Craving* may be best represented with semi-proper transfer functions [7]. Consequently, the $P_{y_Q}(z^{-1})$ and $P_{y_l}(z^{-1})$ models used in the controller formulation developed here feature a significant delay compared to the models for responses to *Quit* and u_l, thereby introducing nontrivial plant-model mismatch and ultimately limiting the true optimality of dosing decisions. To address this challenge, future modifications to this work may consider decision-making on a time scale where the open-loop dynamics can adequately be represented by fully proper models (i.e., within-day dosing). Similarly, explicitly accounting for additional sources of variability may help limit plant-model mismatch due to variability in patient-to-patient responses to treatments and disturbances. Also, the controller may need to consider interaction effects. Intuitively, treatment effects may be a function of time or other variables (e.g., the observed effects of counseling that enhances stress management skills are likely to be greatest when a patient is facing stressful situations). Implementing linear time-varying open loop models or more complex dose effect logic in the hybrid system structure may help address this consideration. Furthermore, incorporating additional tuning capabilities into the HMPC framework, as outlined in [10], could facilitate more nuanced decision-making in a smoking intervention.

References

[1] Centers for Disease Control and Prevention. Current ciagette smoking among adults – united states, 2011. *Morbidity and Mortality Weekly Report (MMWR)*, 61(44):889–894, November 2012.

[2] Tobacco Use and Dependence Guideline Panel. *A Clinical Practice Guideline for Treating Tobacco Use and Dependence: 2008 Update*. U.S. Department of Health and Human Services, Rockville, MD, May 2008.

[3] D.E. McCarthy, T.M. Piasecki, D.L. Lawrence, D.E. Jorenby, S. Shiffman, M.C. Fiore, and T.B. Baker. A randomized controlled clinical trial of bupropion SR and individual smoking cessation counseling. *Nicotine and Tobacco Research*, 10(4):717–729, 2008a.

[4] M.E. Piper, S.S. Smith, T.R. Schlam, M.C. Fiore, D.E. Jorenby, D. Fraser, and T.B. Baker. A randomized placebo-controlled clinical trial of 5 smoking cessation pharmacotherapies. *Archives of General Psychiatry*, 66(11):1253–1262, 2009.

[5] L.M. Collins, S.A. Murphy, and K.L. Bierman. A conceptual framework for adaptive preventive interventions. *Prevention Science*, 5(3):185–196, 2004.

[6] D.E. Rivera, M.D. Pew, and L.M. Collins. Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction. *Drug and Alcohol Dependence*, 88(Supplement 2):S31–S40, 2007.

[7] K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper. Continuous-time system identification of a smoking cessation intervention. *International Journal of Control*, in press.

[8] W.T. Riley, D.E. Rivera, A.A. Atienza, W. Nilsen, S.M. Allison, and R. Mermelstein. Health behavior models in the age of mobile interventions: Are our theories up to the task? *Translational Behavioral Medicine*, 1(1):53–71, 2011.

[9] S. Kumar, W. Nilsen, M. Pavel, and M. Srivastava. Mobile health: Revolutionizing healthcare through trans-disciplinary research. *Computer*, 46(1):28–35, 2013.

[10] N.N. Nandola and D.E. Rivera. An improved formulation of Hybrid Model Predictive Control with application to production-inventory systems. *IEEE Transactions on Control Systems Technology*, 21(1):121–135, 2013.

[11] K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper. Control systems engineering for understanding and optimizing smoking cessation interventions. *Proceedings of the 2013 American Control Conference, Washington, DC*, pages 1967–1972, 2013.

[12] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35:407–427, 1999.

[13] Y. Dong, S. Deshpande, D.E. Rivera, D.S. Downs, and J.S. Savage. Hybrid Model Predictive Control for sequential decision policies in adaptive behavioral interventions. *Proceedings of the 2014 American Control Conference*, in press, 2014.