Engineering Control Approaches for the Design and Analysis of Adaptive, Time-Varying Interventions

Daniel E. Rivera and Michael D. Pew
Control Systems Engineering Laboratory
Department of Chemical and Materials Engineering
Arizona State University
Tempe, Arizona 85287-6006

Linda M. Collins
The Methodology Center and
Department of Human Development and Family Studies
Penn State University
State College, PA

Susan A. Murphy
Institute for Social Research and
Department of Statistics
University of Michigan
Ann Arbor, MI

May 30, 2005

Technical Report 05-73
The Methodology Center, Pennsylvania State University

Abstract

Control engineering is the field that examines how to transform dynamical system behavior from undesirable conditions to desirable ones. Cruise control in automobiles, the home thermostat, and the insulin pump are all examples of control engineering at work. In the last few decades, significant improvements in computing and information technology, increasing access to information, and novel methods for sensing and actuation have enabled the extensive application of control engineering concepts to physical systems. While engineering control principles are meaningful as well to problems in the behavioral sciences, the application of this topic to this field remains largely unexplored.

This technical report examines how engineering control principles can play a role in the behavioral sciences through the analysis and design of adaptive, time-varying interventions in the prevention field. The basic conceptual framework for this work draws from the paper by Collins *et al.* (2004). In the initial portion of the report, a general overview of control engineering principles is presented and illustrated with a simple physical example. From this description, a qualitative description of adaptive time-varying interventions as a form of classical feedback control is developed, which is depicted in the form of a block diagram (i.e., a control-oriented signal and systems representation).

A simulation study of a hypothetical adaptive, time-varying intervention based on the Fast Track program is then presented. The results of a rule-based decision policy (similar to that proposed in Collins *et al.* (2004)) are compared to a Proportional-Integral-Derivative (PID)-type controller designed on the basis of model-based engineering control principles. The simulations are conducted under conditions of varying disturbance magnitudes, stochastic measurement error, and nonlinearity in the model parameters. While the rule-based approach is adequate under conditions of low disturbances, it faces problems of offset for large disturbance magnitudes and nonlinearity. The PID controller tuned using the principles of Internal Model Control does not exhibit these difficulties; however, this controller has an adjustable parameter λ that must be judiciously selected to achieve the proper tradeoff between a desired response in the tailoring variables and high variability and "aggressiveness" in the assigned intervention dosages.

In light of this analysis and the simulation study, a series of systems technologies that may be meaningful in future research activities on this problem are presented; these are dynamical modeling via system identification, Model Predictive Control, Model-on-Demand estimation, and the integration of Model-on-Demand estimation with Model Predictive Control. It is anticipated that the development of control approaches for time-varying, adaptive interventions will establish the basis for novel forms of control technnology involving the effective integration of data-centric estimation, hybrid decision-making, and constrained control.

Contents

List of Figures

List of Tables

1 Introduction

Control engineering, in a broad sense, refers to the discipline that examines how to transform dynamic system behavior from undesirable conditions to desirable ones. Cruise control in automobiles, the home thermostat, and the insulin pump are all examples of control engineering at work. In the last few decades, significant improvements in computing technology, increasing access and availability to information, and novel methods for sensing and actuation have enabled the extensive application of control engineering concepts to physical systems. Engineering control principles are applicable as well to problems in the behavioral sciences, yet this potential remains untapped. One area where engineering control principles can be applied productively is the analysis and design of time-varying, adaptive interventions (Collins *et al.*, 2004). We argue that this represents an important, significant and meaningful problem that can open new avenues in the application of control engineering to the behavioral sciences.

Adaptive interventions represent a promising approach to prevention and treatment. They are especially useful for prevention programs with numerous components aimed at different aspects of risk, and for treatment of chronic, relapsing disorders such as alcoholism, cigarette smoking, and other types of substance abuse. Contingency management, individualized treatments, stepped care programs, and case management all represent frameworks that enable the implementation of adaptive interventions. Adaptive interventions individualize therapy by the use of decision rules for how the therapy level and type should vary according to measures of adherence, treatment burden, and response collected during past treatment (MC-DATS, 2004). Table 1 contains a list of useful definitions relating to this problem. Adaptive interventions differ from conventional fixed interventions in significant ways. In fixed interventions, the same dosage is applied to all program participants without taking into account any of their individual characteristics. In an adaptive intervention, different dosages of prevention or treatment components are assigned to different individuals across time, with dosage varying in response to the needs of the individual. For example, the composition of a drug abuse prevention program might be varied somewhat depending on the ethnic make-up of each school in which it is delivered. Adaptive interventions are time-varying when the adaptation is repeated throughout the intervention. For example, a smoking cessation program may periodically assess each participant's progress along the stages of the Transtheoretical Model (Velicer and Prochaska, 1999), and accordingly adjust how key components of the intervention are presented. While adaptive interventions may on the surface appear similar to sensible clinical practice, in order to be successful, they must be much more tightly managed than typical clinical procedures. Interest in adaptive techniques is significant not only in the treatment of substance abuse (Sobell and Sobell, 1999; Velicer and Prochaska, 1999; Brooner and Kidorf, 2002; Murphy and McKay, Winter 2003/Spring 2004) but also in the treatment of hypertension (Glasgow *et al.*, 1989), depression (Rush *et al.*, 2004), Alzheimer's disease (Schneider *et al.*, 2001) and infectious diseases (Rosenberg *et al.*, 2000). As adaptive strategies play an increasingly prominent role as a methodological framework for many important prevention problems, it is evident that much research is needed on analysis, design, and effective implementation of these interventions.

Consider a hypothetical adaptive intervention, inspired by the Fast Track program (CPPRG, (1992; 1999*a*; 1999*b*)), which will serve as the basis for a number of example problems described in this report. The long-range purpose of the intervention is to prevent the development of conduct disorders in children. The intervention is family counseling. There are several possible levels of intensity, or doses, of family counseling. The idea is to vary the doses of family counseling depending

Table 1: Definitions adapted from material prepared for the Methodological Challenges in Developing Adaptive Treatement Strategies network (MC-DATS, 2004).

Phrase	Definition
Adaptive treatment strategies	These strategies individualize therapy by the use of decision rules for how the therapy level and type should vary according to measures of adherence, treatment burden and response collected during past treatment. Adaptive treatment strategies do *not* involve randomization; other names are stepped care models, dynamic treatment regimes, structured treatment interruptions, treatment algorithms and adaptive interventions.
Design of a Treatment	A treatment that might be implemented in clinical setting, e.g., provide 3 hours of counseling per week for 1 month. An adaptive treatment strategy is one type of treatment design. Does not involve randomization.
Design of an Experimental Trial	Statistical design of an experiment, usually involving randomization. The treatments to which a subject may be randomized may be constrained by preference, response to past therapy, side effects suffered from past therapy, etc. An experimental design is used to develop and/or evaluate different treatment designs.
Index	These might be risk indices or response indices. An index is a summary of the available information that is strongly predictive of future behavior or disease. For example, a risk index that includes results of biological tests and psychosocial factors may be highly predictive of future substance use.
Tailoring Variable	A tailoring variable is also a summary of the available information. However as opposed to an index, the primary purpose of a tailoring variable is to discriminate between different timings of treatment alterations or different intensities of treatment or different treatments. Since the purposes of a tailoring variable and an index differ (discrimination versus prediction), they may summarize similar information but they might not. Of course sometimes a good risk index is also a good tailoring variable but this need not be the case.

on the need of the family, in order to avoid both providing an insufficient amount of counseling for very troubled families and wasting counseling resources on families that do not need it. The decision about which dose of counseling to offer each family is based on two factors. One is the family's level of functioning, assessed by a family functioning questionnaire completed by one of the parents. The score on the family functioning questionnaire is called a tailoring variable, because it is used to tailor the treatment to the individual family. The other factor is the judgment of a clinician familiar with the family's case. Based on the questionnaire and the clinician's assessment, family functioning is determined to be very poor, poor, near threshold, or at/above threshold. The decision rule is as follows: families with very poor functioning are given weekly counseling; families with poor functioning are given biweekly counseling; families with below threshold functioning are given monthly counseling; and families at or above threshold are given no counseling. Family functioning is reassessed every three months, at which time the intervention dosage may change. This goes on for three years, with twelve opportunities for a dose of family counseling to be assigned. The final outcome of interest is a measure of conduct disorder in the target child, assessed one year after the end of the intervention period.

During the past two years, the authors have held some intriguing discussions on the topic of the report. In these discussions we have identified conceptual linkages between adaptive, time-varying interventions in the prevention field and engineering process control. Process control systems are widely used in the chemical industries to adjust flows to maintain level and product compositions at desired values (Ogunnaike and Ray, 1994). Effective time-varying adaptive interventions have goals similar to those of well-designed process control systems, in that both seek to 1) reduce negative effects, 2) increase intervention potency and 3) reduce waste. We believe that control engineering represents an exciting and very promising new framework for development of methodology for adaptive, time-varying interventions. The objective of this technical report is to examine the links between these fields, with the ultimate goal of developing novel and effective intervention strategies based on engineering control principles.

The report is organized as follows: Section 2 describes some control engineering fundamentals and introductory background in this field. Section 3 considers the link between adaptive, time-varying interventions and engineering control by examining the hypothetical Fast Track intervention described previously. An extensive simulation study is presented that contrasts a rule-based decision policy with a model-based approach based on engineering control principles. Section 4 describes some systems and control technologies that we believe would be applicable in future research on this problem. Section 5 summarizes the findings and observations of this report, as well as some of the fundamental challenges associated with this research.

2 Control Engineering Fundamentals

Control engineering is a broadly applicable field that has become a significant part of every day life. Many items that form part of our modern existence (such as automobiles, airplanes, and consumer appliances) rely on well-designed control systems to enable safe, profitable, and environmentally friendly operation. Control engineering is a subject that is included in the program of study of most engineering fields (aerospace, biomedical, chemical, electrical, industrial, and mechanical, to name a few). In much of the remainder of this report, many terms new to behavioral scientists will be introduced. Table 2 presents a list of definitions that will be useful in understanding this topic.

Feedback control represents one of the most useful and commonplace forms of control strategies applied in industrial practice. We can illustrate the concept of a feedback control system with a simple example from everyday life: taking a shower. Think of your ideal shower, with your preferred temperature and water flow. When you are taking a shower you *control* these features (see Figure 1) by using the hot and cold taps. In control theory language, you are the *controller*, and temperature and water flow are the *outputs*. In this example the outputs are *controlled variables*, and the particular temperature and water flow desired are the *setpoint values*. The purpose of the feedback control system is to keep the controlled variables as close as possible to the setpoints, or in this case, to keep the temperature and water flow as close as possible to your ideal settings. Manual feedback control of a shower as described here mimics conventional clinical practice in that a clinician decides on dosages and treatment based on his/her judgement of the current state of the patient, and the history of treatment.

There are two general types of *inputs* in feedback control systems. One is *manipulated variables*, which can be adjusted by the controller to achieve the desired effect. In our example these are the hot and cold taps. The other is *disturbance* (or *exogenous*) *variables*, that induce changes in the

Table 2: Fundamental Control Engineering Terminology

Phrase	Definition
Adaptive Control	A body of control engineering that considers mechanisms for updating model and/or controller parameters as these change with operating conditions. To be distinguished from an adaptive intervention, where the term adaptive implies the presence of feedback control.
Block Diagram	A graphical representation of the signals and systems that comprise a closed-loop control system.
Closed-loop	Refers to system behavior once a controller/decision policy is implemented.
Controller	A mathematical set of relationships that translate error (i.e., deviation from a goal or setpoint) into settings for a manipulated variable (which defines an intervention dosage). Also referred to as a decision policy or decision rule in the context of this report.
Control Engineering	The science that considers how to manipulate system variables in order to transform dynamic behavior to desirable from undesirable.
Control Error ($e = r - y$)	The difference between the controlled variable and the setpoint; the ultimate goal of a control system is to take this variable to zero.
Control Loop	Refers to a closed-loop system.
Control Structure	Refers primarily to whether feedback or feedforward strategies (or their combination) are applied in a closed-loop system.
Controlled Variables (y)	System variables that we wish to keep at a reference value or setpoint (r).
Disturbance Rejection	Refers to the ability of the control system to manipulate system variables such that the controlled variable is kept as close as possible to the setpoint, in spite of significant changes in the disturbance variables.
Disturbance Variable (d)	A system variable that influences the controlled variable response, but cannot be manipulated by the controller; disturbance changes occur external to the system (hence sometimes referred to as exogeneous variables).
Feedback control	A control strategy in which the a controlled variable (y) is examined and compared to a reference value or setpoint (r). The controller issues actions (decisions on the values of a manipulated variable (u)) on the basis of the discrepancy between y and r.
Feedforward control	A control strategy in which changes in a disturbance variable (d) are monitored and the manipulated variable (u) is chosen to counteract anticipated changes in y as a result of d.
Manipulated Variable (u):	System variable whose adjustment influences the response of the controlled variable y; the magnitude of u is determined by the controller.
Open-loop	Refers to dynamical system behavior without a controller or decision policy.
Offset	A sustained discrepancy between the controlled variable response and the setpoint in a closed-loop system.
Process	The dynamical system under study, for which a closed-loop controller or decision rule will be applied.
Setpoint Tracking	refers to the ability of the control system to manipulate system variables such that the controlled variable follows a reference (setpoint) trajectory as closely as possible.

controlled variables that keep them from attaining the setpoint values. There are many possible sources of disturbances in this example, among them ambient temperature changes, fluctuations in the operation of the home water heater, abrupt changes in the water flow sources to the shower (for instance, yard sprinklers going off or a nearby toilet being flushed).

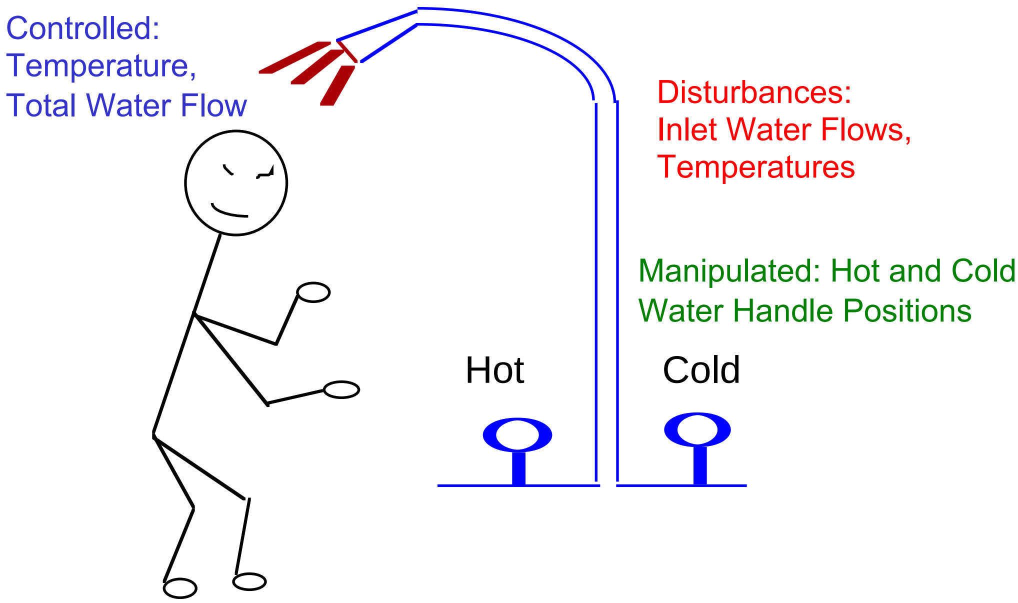

Figure 1: Person in the shower: an everyday control problem.

It is customary in control engineering to depict feedback control problems in a block diagram like the one in Figure 2 related to the shower example. The quantities discussed above are represented in the block diagram as follows:

- The vector of output variables, y (temperature and water flow);
- The vector of setpoint values, r (preferred values for temperature and water flow);
- The vector of error signals, e, which are the differences between output variables and setpoint values (differences between water temperature and preferred temperature, and between water flow and preferred water flow);
- Manipulated input variables, u (hot and cold taps);
- Disturbance input variables, d (fluctuations in water heater etc.).
- C represents the controller; this could either be the person taking a shower (akin to clinical judgement) or in an automatic closed-loop system, a set of decision rules or similar algorithm.
- P represents the effect of the manipulated variables u on the output variables y in the absence of any disturbance variables (the effect on water temperature and flow of manipulating the taps if there were no fluctuations in the water heater, etc.).
- P_d represents the effect of the disturbance variables d on the output variables y in the absence of any manipulated variables (the effect on water temperature and flow of fluctuations in the water heater, etc. if no one was manipulating the taps).

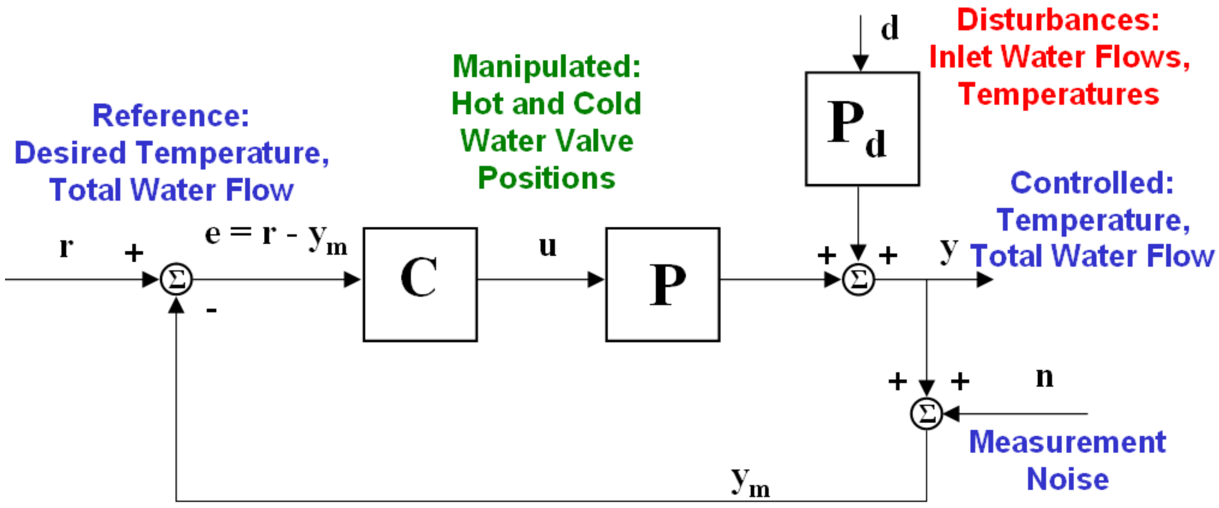

(a) Block Diagram

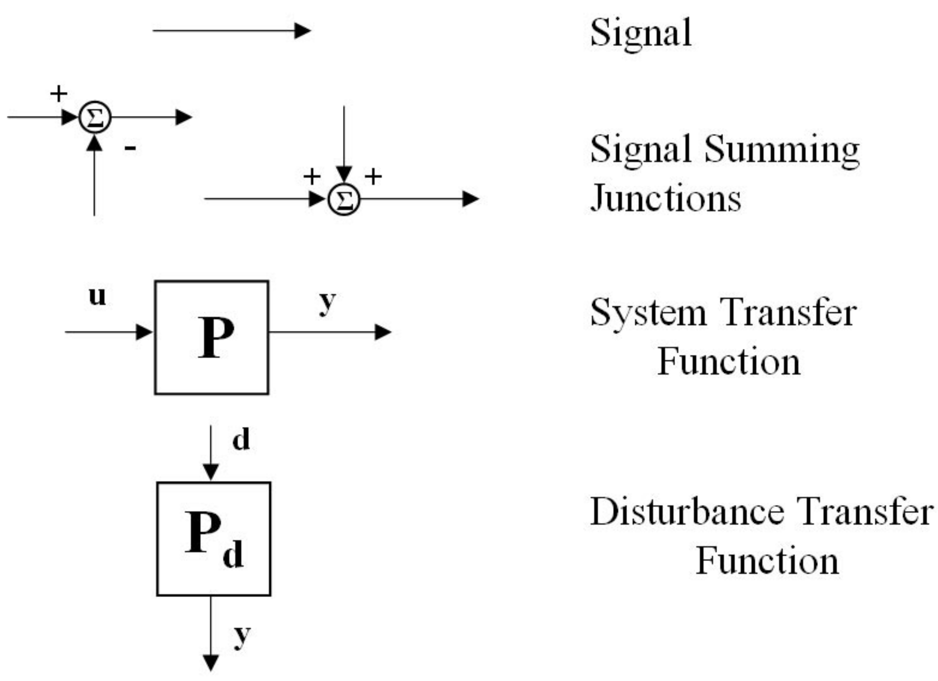

(b) Legend

Figure 2: Block diagram representation of the closed-loop system entailed in shower operation.

- The symbol (arrow) represents a signal (e.g., a measure of water temperature and flow).
- The symbol (signal summing junction) represents the addition of two or more signals.

Let us work through the process depicted in Figure 2 from left to right and back again, following the arrows. The designated setpoint vector r (representing desired settings for temperature and water flow) are specified externally . The signal which contains the *measured* values for temperature and water flow y is compared to the setpoints to obtain the error signal e. The controller (or decision rule) C uses the measured and error signal to decide on values for the manipulated variables u, that is, the settings for the hot and cold taps. P, the "open-loop" model, expresses the effect this has on water temperature and flow. The actual water temperature and flow are also altered by the disturbance P_d; as a result the actual values of the controlled variables y include the effects of both manipulated variable and disturbance changes. The action of the control system is conducted continuously in time; in a well-designed control system, the controller/decision rule continues to suggest changes in u until the control error e is minimized.

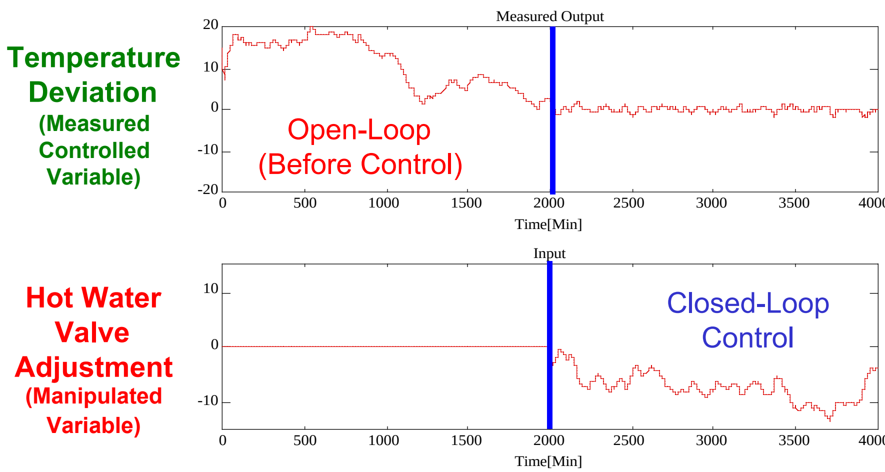

Figure 3: Time series representing the effects of control action in a stochastic disturbance setting. From time = 0 to 2000, the system is in open-loop (manual) operation, with no changes made to the manipulated variable. Closed-loop (automatic) feedback control is engaged at time = 2000. The transfer of variance from the controlled variable to the manipulated variable as a result of feedback control action is clearly illustrated.

2.1 Representing Control Problems using Systems of Equations

In this section we show how to represent control problems by means of systems of equations. In a manual control setting, the individual taking the shower senses what he or she believes is an adequate water flow and temperature, and makes adjustments accordingly. Manual control is somewhat challenging for this system because of the transportation delay created by the long pipe between the hot and cold water valves and the shower head. For simplicity we consider a single-variable problem with one output signal $y(t)$ (shower temperature), one manipulated variable $u(t)$ (hot water valve position), and one disturbance variable $d(t)$ (ambient temperature changes). This is an *open-loop* system. Unlike the shower example, an open-loop system does not have a controller. A dynamical model that can describe the system dynamics as a result of arbitrary changes in these inputs is written as a first-order continuous-time differential equation as follows:

$$\tau \frac{dy}{dt} + y(t) = K_p \, u(t-\theta) + K_d \, d(t) \tag{1}$$

Equation (1) can be derived on the basis of physical conservation of mass and energy in the system. The system dynamics are characterized by τ, the time constant, K_p and K_d the process and disturbance gains, respectively, and θ, the time delay. Details of the derivation are not presented here, but it can be shown that the values for τ, K_p, K_d, and θ are determined by the length and width of the pipe, the volume in the shower head, the heat capacities of the materials and other physical attributes of the shower. An equivalent representation of the model per Equation (1) in the Laplace domain (s) leads to the transfer functions $P(s)$ and $P_d(s)$:

$$y(s) = P(s)u(s) + P_d(s)d(s) \tag{2}$$

$$P(s) = \frac{K_p \, e^{-\theta s}}{\tau s + 1} \tag{3}$$

$$P_d(s) = \frac{K_d}{\tau s + 1} \tag{4}$$

To meet the operational goals of *setpoint tracking* and *disturbance rejection*, it is necessary to "close the loop", that is, to introduce some form of automatic control. In setpoint tracking, the controller determines the position of the manipulated variables to take the controlled variables from an initial setting to a final desired goal. For instance, the process of starting the shower from an initial closed (or shutdown) state to one in which the desired final temperature setting is reached is one example of setpoint tracking. The second goal, disturbance rejection, describes the adjustments made by the control system to address the effect of disturbances (such as ambient temperature changes, unexpected flow variations , and so forth) on the total water flow and temperature in the shower.

A closed-loop system consists of a controller as well as the required *sensors* and *actuators*. Sensors represent devices or mechanisms that measure the controlled variables; these measurements are potentially prone to errors. Controllers are mathematical relationships that act on the basis of measurements of the controlled variables or other signals in the system to determine the settings for the manipulated variables. Actuators translate the settings received from the controller into final positions of the manipulated variable. Sensors and actuators are both "hardware" components that need to be properly designed to obtain a well-performing control system.

It should be noted that feedback control is not the only control strategy available for this problem. In *feedforward control*, changes in the disturbance variables (d) are measured and the manipulated variable values (u) are chosen to counteract anticipated changes in y as a result of d. In industrial practice, the majority of control systems are feedback-only, but combined feedforward-feedback control systems are nonetheless commonplace. Regardless of the control strategy used, a main objective of engineering control is the transfer of variability from an expensive resource (the controlled variable) to a cheaper one (the manipulated variable). In the shower example, if we do not attempt to control water temperature by adjusting the hot and cold taps, in other words, if essentially we have an uncontrolled open-loop system, the hot water valve remains constant while water temperature fluctuates. Once control is introduced via adjustment of the taps, i.e., we move to a closed-loop system, water temperature variability is greatly reduced while the hot water valve position varies. This transfer of variance is critical to the important role that engineering-based control systems play in industrial practice; Figure 3 illustrates this.

3 Extending Control Engineering Principles to the Behavioral Sciences

3.1 A Qualitative Description Using Block Diagrams

Our first step is establishing how engineering control principles can be used to conceptually describe a time-varying, adaptive intervention is to develop a corresponding block diagram. We draw on the discussion of adaptive, time-varying interventions per Collins *et al.* (2004). Prevention and treatment data are nearly always longitudinal, which involves multiple variables interacting in potentially complicated ways. The problem of how to capture the mathematical relationship between intervention and outcome(s) in a manner meaningful for control design represents a significant challenge. We noted previously that engineering control systems are traditionally applied to physical systems where the principles of conservation of mass, energy, charge, and momentum are used as the basis for describing phenomena. Problems in the prevention field are not so easily characterized, and furthermore, these systems display behavior that can be highly nonlinear, uncertain, and stochastic. Nonetheless, we capture this relationship in an abstract setting with a model block as shown in Figure 4.

We illustrate the nature and benefits of a process control understanding of adaptive, time-varying interventions through a description of its application to the hypothetical Fast Track intervention described in the Introduction. A block-diagram representation of the parental function-home visit time-varying, adaptive intervention as an engineering feedback control system is shown in Figure 4, and a table relating the intervention parameters to engineering control variables is presented in Table 3. In this representation, the tailoring variable (a measured level of parental function, denoted by the variable $PF_{meas}(t)$, where the time variable t is expressed as an integer value of a review interval T) represents a controlled variable signal that is used by a feedback controller (i.e., decision rules) to determine the dosage level of the intervention (in this case, the frequency of home visits, denoted by the signal $I(t)$) based on a measured level of parental function. The adaptive nature of the time-varying intervention implies that feedback control is taking place; hence the term adaptive in the prevention field equates feedback (not to be confused with the term adaptive control, which is a distinct area within the control engineering field). Disturbance signals in this problem represent individual, time-varying characteristics that can deplete parental function, and for purpose of this report will be lumped together into a depletion signal $D(t)$. In addition, we need to note the possibility of unreliability in the measurements; such "noise" in the parental function measurement is captured via the signal $N(t)$.

While the control-oriented representation depicted in Figure 4 is conceptual in nature, nonetheless it is extremely useful for articulating some fundamental questions on the design and implementation of adaptive, time-varying interventions; these are topics that merit additional investigation.

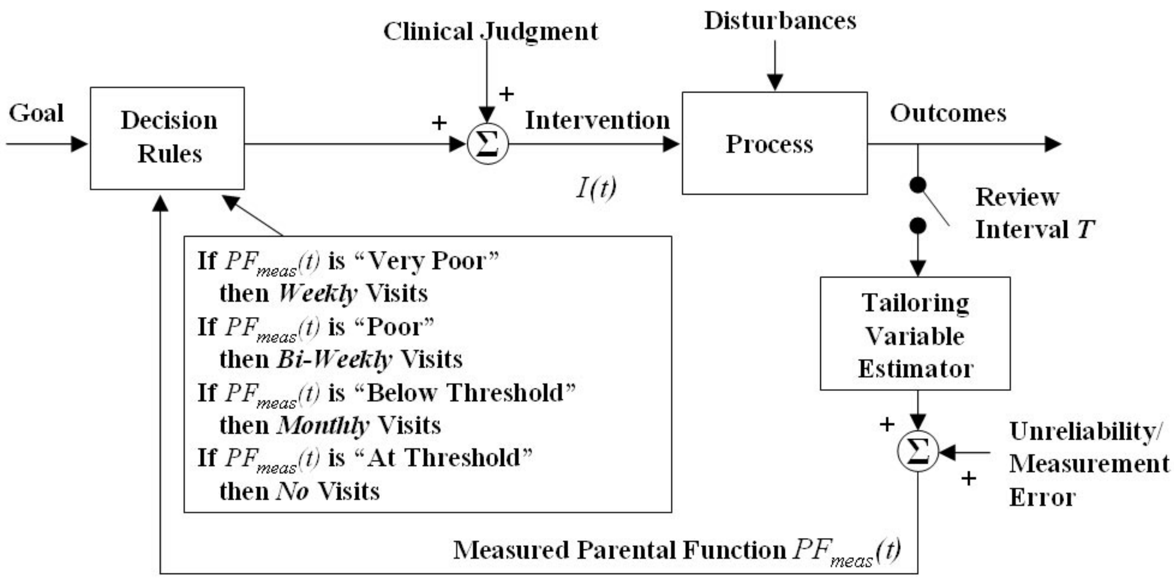

(a) Block Diagram

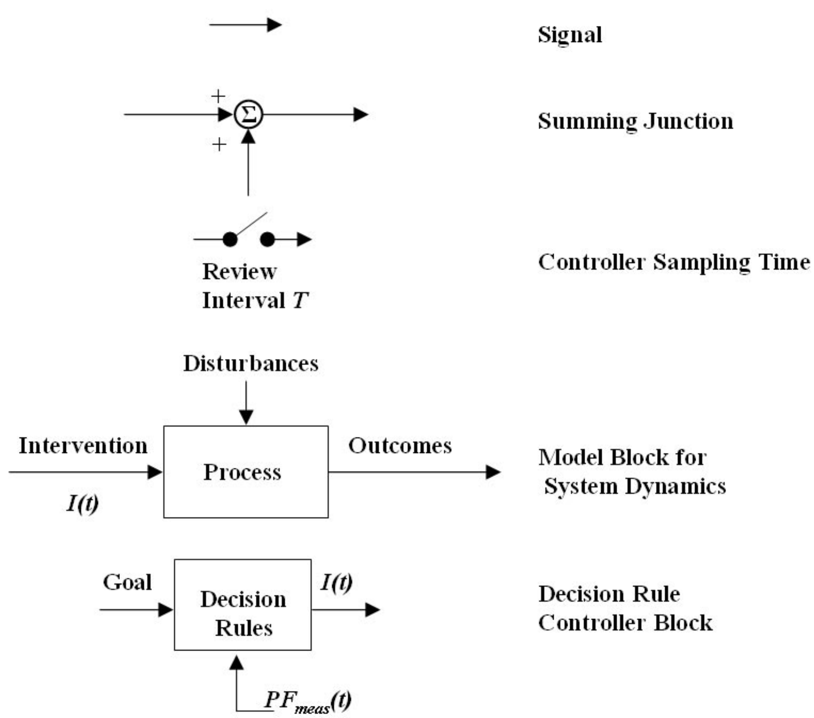

(b) Legend

Figure 4: Block-diagram feedback control representation of a hypothetical home visits adaptive component for the FAST TRACK program

Among these questions are:

1. *How can model-based engineering control design techniques enhance the development of decision rules for time-varying, adaptive interventions?* Having developed models for prevention phenomena that are relevant to control, the issue of how to systematically arrive at appropriate decision policies for these problems must be considered. In general, the sophistication of the control laws will be dictated by the complexity of the model. In prevention problems, these control laws will likely be hybrid in nature, incorporating both continuous and discrete-event decisions.

2. *How reliably and frequently must the tailoring variable(s) be estimated, when the purpose of the measurement is a time-varying adaptive intervention?* Previous work by Collins and Graham (2002) has demonstrated the importance of judicious selection of sampling time when drawing inference in longitudinal studies of prevention phenomena. The requirements for both measurement and modeling accuracy in a time-varying adaptive intervention will be influenced by the performance expected from the control system. Furthermore, there are additional issues related to how to translate potentially multiple outcomes into an estimate for one or more tailoring variables.

3. *What role do disturbances (i.e., individual time-varying characteristics) play in adaptive, time-varying interventions, and to what extent can they be effectively managed by the actions of the decision policy?* Individual time-varying characteristics represent external (exogenous) conditions that will influence how an individual (or a group of individuals) responds to an intervention. These "disturbances" (which can be either measured or unmeasured) form part of the control system and need to be effectively managed by the actions of a well-designed adaptive time-varying intervention. Specifically, the control system must suppress the deleterious effects of the detrimental disturbances, and take advantage of those that will take the system to goal. If these disturbances are measurable, it is possible to enhance the decision algorithm by incorporating them in a feedforward strategy in the control system, as noted previously.

4. *How can effective adaptive interventions be implemented in the presence of external clinical actions?* As noted in Figure 4 the clinician will ultimately make the final decision on dosage levels in an intervention. An effective adaptive intervention should create a synergism between decision rules and clinical judgement; however, being able to accomplish this will depend in large part on the judicious design of the intervention.

3.2 A Simulation-Based Analysis

3.2.1 Open-Loop Model Definition

As a means to further understand the benefits of control systems engineering to adaptive interventions in drug abuse prevention, we present the results of a simulation-based approach that allows us to examine one possible means by which a process control perspective can be used to model prevention phenomena, and develop decision rules for an adaptive intervention. In engineering systems, conservation and accounting of extensive properties such as mass, energy, momentum, and charge

Engineering Control Variable	Adaptive, Time-varying Intervention Variable
setpoint $r(t)$	Goal or threshold on the Measured Parental Function $PF^{Goal}(t)$
output $y(t)$	Measured Parental Function $PF(t)$
disturbance input $d(t)$	Sources of Parental Function Depletion $D(t)$
manipulated input $u(t)$	Intervention Frequency of Home Visits $I(t)$

Table 3: Relationship with engineering control variables for the Fast Track adaptive time varying intervention example.

serve as the basis for developing models that describe dynamical system behavior (Ogunnaike and Ray, 1994). The general accounting principle is represented by the equation:

$$\text{Accumulation} = \text{Inflow - Outflow + Generation - Consumption}, \tag{5}$$

The general accounting principle per (5) represents one possible approach to describe the "open-loop" dynamics of phenomena occurring in prevention. For the Fast Track problem, we postulate a difference equation model to model the relationship between home visits and parental function; the model is written as follows:

$$PF(t+T) = PF(t) + K_I\, I(t-\theta) - D(t) \tag{6}$$

$$D(t) = \sum_{i=1}^{n_d} D_i(t) \tag{7}$$

$$PF_{meas}(t) = PF(t) + N(t) \tag{8}$$

$PF(t)$, $PF_{meas}(t)$, $I(t)$, $D(t)$, and $N(t)$ have been previously defined in this document; K_I is the intervention gain, while θ represents the time delay between the intervention and its actual effect on parental function. The equation per (6) states that the parental function at the end of a review period $PF(t+T)$ equals the parental function at the start of the review period $PF(t)$ plus the scaled effect of the intervention (implemented θ time units prior) less the depletion occuring during that time period ($D(t)$). Equation (7) considers that the disturbance signal is a collective effect of multiple (n_d) individual time-varying characteristics. Equation (8) indicates that the parental function measurement is potentially corrupted by error and unreliability, denoted by $N(t)$.

A *fluid analogy* for this model is presented in Figure 5, which corresponds to a draining tank being fed parental function (the "fluid") from a potentially long pipe. In the fluid analogy, parental function is accumulated as an inventory, which can either be depleted by disturbances, or replenished by an intervention. The length of the pipe is proportional to the delay between the intervention and tailoring variable. While particular disturbances could be beneficial in character (i.e., a new job, leaving a bad neighborhood, a positive religious experience, etc.), for purposes of this study we will treat the net sum of these disturbances (denoted by $D(t)$) as a depletion.

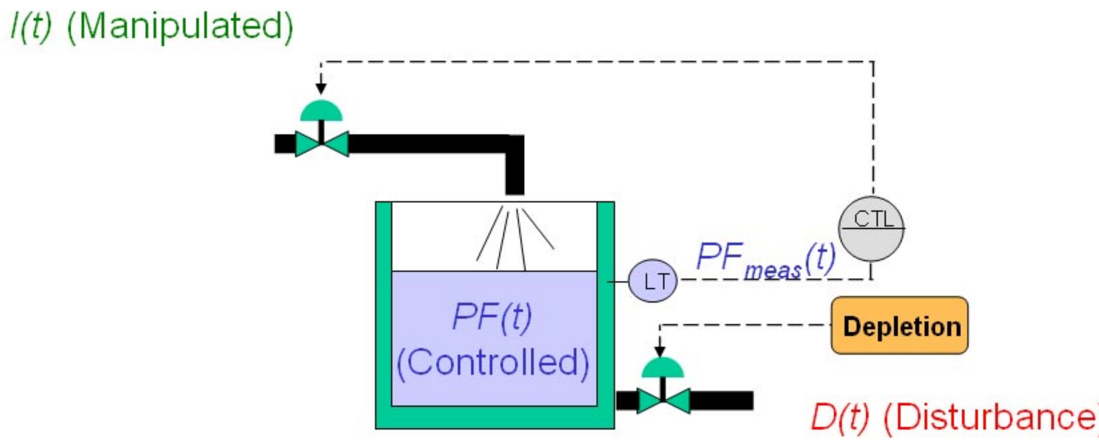

Figure 5: Fluid analogy for the model and control system describing the hypothetical Fast Track adaptive intervention. LT = level transmitter; CTL = feedback controller.

It is important to note that the model per (6) and (8) hints at some of the significant modeling challenges associated with dynamically modeling prevention phenomena. The mechanisms by which interventions translate into outcomes that define the tailoring variable will most likely be highly stochastic and nonlinear in nature, both in K_I and θ. Furthermore, the disturbance effects leading to the depletion rate $D(t)$ will have both deterministic and stochastic components; $N(t)$, on the other hand, will be principally a stochastic signal but may be subject to bias.

3.2.2 A Rule-Based Decision Policy

The decision rules described in the Introduction (henceforth referred to as the rule-based control policy or rule-based controller) can be mathematically represented as follows:

- If the measured parental function is "Very Poor" ($0 \leq PF_{meas}(t) \leq PF^{Very\ Poor}$) then the intervention dosage should correspond to weekly home visits ($I(t) = I^{weekly}$),
- If the measured parental function is "Poor" ($PF^{Very\ Poor} < PF_{meas}(t) \leq PF^{Poor}$) then the intervention dosage should correspond to bi-weekly home visits ($I(t) = I^{biweekly}$),
- If the measured parental function is "Below Threshold" ($PF^{Poor} < PF_{meas}(t) \leq PF^{Goal}$) then the intervention dosage should correspond to monthly home visits ($I(t) = I^{monthly}$),

Though not explicitly mentioned in Collins *et al.* (2004), we consider that if parental function is meets or exceeds the goal ($PF_{meas}(t) > PF^{Goal}$) then the intervention dosage should correspond to no home visits ($I(t) = 0$). It should be recognized that the performance of the adaptive, time-varying intervention will depend on the values selected for the threshold parameters $PF^{Very\ Poor}$, PF^{Poor} and PF^{Goal} as well as the intervention potency reflected in the values of I^{weekly}, $I^{biweekly}$, and $I^{monthly}$.

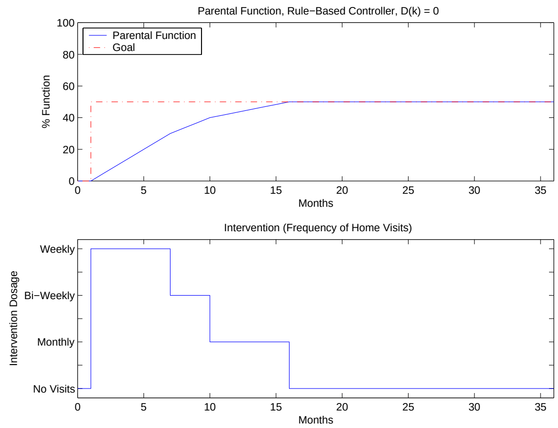

(a) $D(t) = 0$

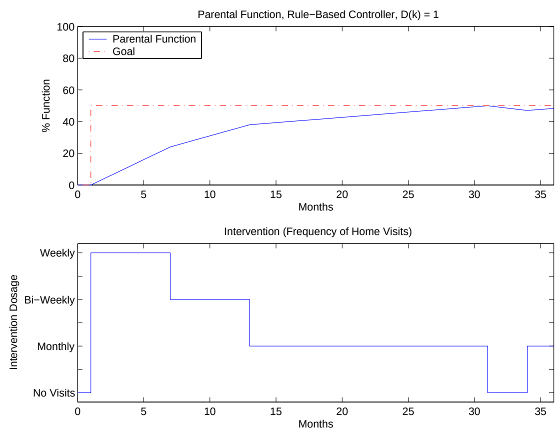

(b) $D(t) = 1$

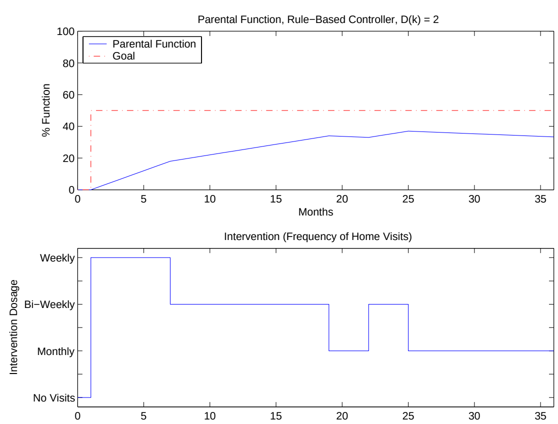

(c) $D(t) = 2$

Figure 6: Closed-loop response of the rule-based control system per Collins *et al.* with $K_I = 0.05$, $\theta = 0$, $T = 3$ months, no measurement error ($N(t) = 0$ for all t) and increasing magnitude of depletion.

Figure 6 demonstrates the response obtained by applying an implementation of the rule-based controller described in Collins *et al.* (2004) to the system per (6) under deterministic conditions with $K_I = 0.05$, $\theta = 0$, $T = 3$ months for the review interval, no measurement error ($N(t) = 0$ for all t) and no depletion ($D(t) = 0$ for all t). Both parental function $PF(t)$ and intervention dosage $I(t)$ are considered as normalized measurements with values ranging from 0 to 100%. Settings for the parental function thresholds are set to $PF^{Very\ Poor} = 16.7\%$, $PF^{Poor} = 33\%$, and $PF^{Goal} = 50\%$; the intervention potency is assumed to be linearly scaled and is defined according to $I^{weekly} = 100\%$, $I^{biweekly} = 66.7\%$, and $I^{monthly} = 33\%$.

Initially, the subject is considered to possess 0% parental function with the intervention dosage determined at $t = 1$ month and reviewed every 3 months thereafter for a 36 month total program. As seen in Figure 6a, under these circumstances, the rule-based controller works as expected; initially the rules dictate an intervention dosage of weekly home visits, with the frequency of visits decreasing as the parental function of the family improves. At 16 months the estimated parental function meets the goal, and remains there for the duration of the assigned time period for the intervention. Because estimated parental function has achieved the goal, the rule-based policy determines that there is no need for additional home visits and the intervention is concluded.

We consider the presence of nonzero depletion of parental function, and for the sake of simplicity assume that the rate of depletion occurs at a constant value. Consider the case where $D(t) = 1$ for all t, as shown in Figure 6b. Under these circumstances it takes longer for parental function to reach the goal (30 months) and the intervention dosages are on average higher than before and never go to zero (since there is a constant loss of parental function per month). When the rate of depletion is increased to $D(t) = 2$ for all t, the rule-based controller, while recommending higher dosages, fails to attain the goal throughout the 36 months of the intervention. This unattainment of the goal is referred to as "offset" in control engineering terminology, and is depicted graphically in Figure 6c. Despite an overall increase in parental function in the family at the conclusion of the intervention, the presence of offset is an undesirable phenomenon that results from the simplified nature of the decision rules applied to this intervention.

3.2.3 A Decision Policy Based on Engineering Control Principles

In this subsection we examine how to apply a model-based design procedure stemming from engineering control principles to arrive at a controller/decision policy for the hypothetical adaptive, time-varying intervention previously described. A corresponding block diagram for a closed-loop system using an engineering controller is presented in Figure 7. In general, the requirements for the engineering controller are defined by the character and sophistication of the open-loop model, the desired closed-loop performance characteristics, and the external signals (reference setpoints and disturbances) that will be faced by the control system. There are multiple ways in which engineering control policies can be developed for systems with dynamics per (6). One approach towards obtaining a model-based, engineering feedback control law in this case is to recognize that the model per (6) corresponds to a delayed, integrating plant. The use of a first-order Padé approximation

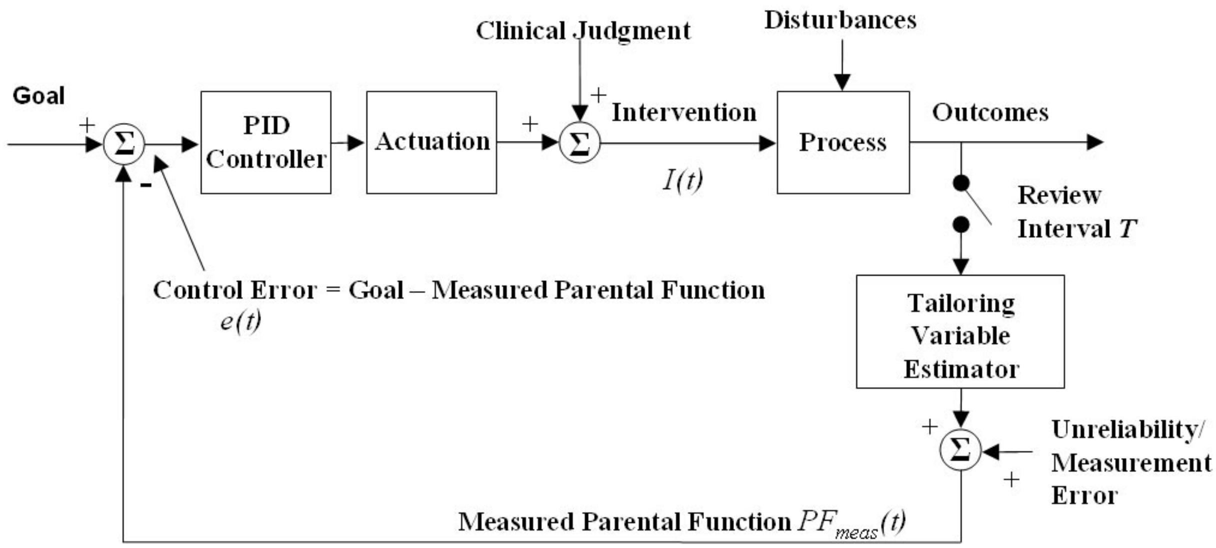

(a) Block Diagram

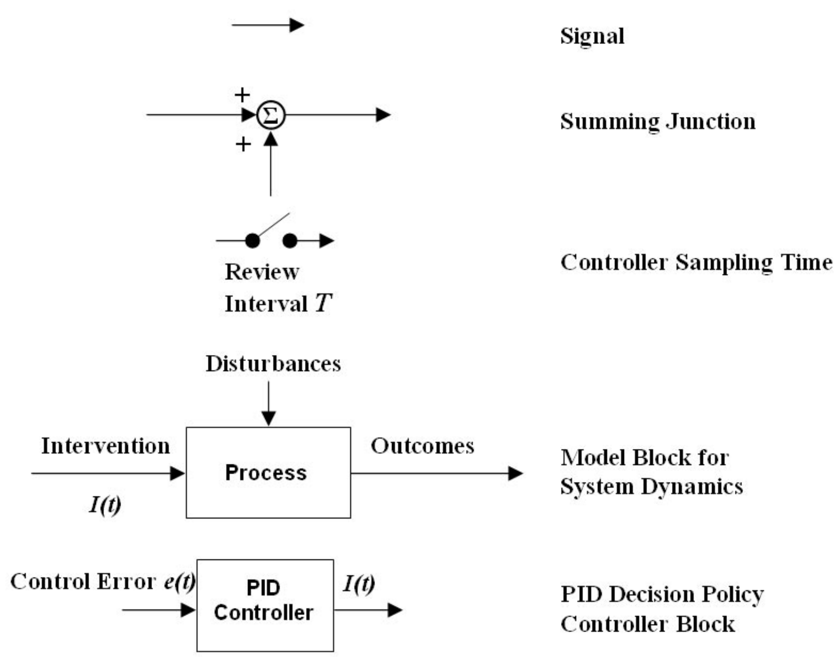

(b) Legend

Figure 7: Block-diagram for an engineering-based feedback control approach to the home visits adaptive component for the FAST TRACK program

on the delay leads to an integrating system with lag and a Right-Half Plane zero,

$$PF(s) = P(s)\,I(s) + P_d(s)\,D(s) \tag{9}$$

$$= \frac{e^{-\theta s}}{s}\,I(s) - \frac{1}{s}\,D(s) \tag{10}$$

$$\approx \frac{\frac{-\theta}{2}s+1}{s(\frac{\theta}{2}s+1)}\,I(s) - \frac{1}{s}\,D(s) \tag{11}$$

The resulting model per (11) is thus amenable to the Internal Model Control (IMC) tuning rules for Proportional-Integral-Derivative (PID) type controllers with filter

$$\frac{u(s)}{e(s)} = C(s) = K_c(1 + \frac{1}{\tau_I s} + \tau_D s)\frac{1}{(\tau_F s+1)} \tag{12}$$

as developed in Rivera *et al.* (1986) and Morari and Zafiriou (1989). The PID controller is one of the most commonplace control algorithms in industrial practice; K_c is the proportional gain, τ_I is the integral time constant, τ_D is the derivative time constant and τ_F the filter time constant for this control algorithm. $e(s) = r(s) - y_m(s)$ is the error signal that denotes the magnitude of the deviation between controlled variable and setpoint. An IMC-PID tuning rule corresponding to the structure per (12) which results in no offset for ramp-like disturbances is:

$$\beta = \tau = \frac{\theta}{2} \qquad K_c = \frac{2(\beta+\lambda)+\tau}{K_I(2\beta^2+4\beta\lambda+\lambda^2)} \qquad \tau_I = 2(\beta+\lambda)+\tau \tag{13}$$

$$\tau_D = \frac{2\tau(\beta+\lambda)}{2(\beta+\lambda)+\tau} \qquad \tau_F = \frac{\beta\lambda 2}{2\beta^2+4\beta\lambda+\lambda^2} \tag{14}$$

This tuning rule is model-based because the controller parameters K_c, τ_I, τ_D, and τ_F are all defined on the basis of the open-loop model parameters (K_I and θ) and an adjustable parameter λ which can be used to determine the speed-of-response and to influence the peformance-robustness tradeoff inherent to all feedback control systems. In general, λ is inversely proportional to the closed-loop speed-of-response, that is, the speed it takes the controlled variable to reach the goal. In general, increasing λ will result in less responsiveness of the manipulated variable to setpoint and disturbance changes (thus making the closed-loop system sluggish), while decreasing λ will promote more aggressive changes in the manipulated variable and lead to faster response in the controlled variable.

The PID with filter controller per Equation (12) is written in continuous form, but needs to be numerically discretized to be implemented in a sampled data environment (where measurements are available only during specified review intervals T). By using first-order backward-difference approximations on the integral and derivative modes, the continuous-time PID controller per (12) can be expressed as a difference equation, written as follows:

$$I(t) = I(t-T) + K_1 e(t) + K_2 e(t-T) + K_3 e(t-2\,T) + K_4(I(t-T) - I(t-2T)) \tag{15}$$

$$K_1 = \frac{TK_c}{\tau_F+T}\left(1+\frac{T}{\tau_I}+\frac{\tau_D}{T}\right) \quad K_2 = -\frac{TK_c}{\tau_F+T}\left(1+\frac{2\tau_D}{T}\right) \quad K_3 = \frac{K_c\tau_D}{\tau_F+T} \quad K_4 = \frac{\tau_F}{\tau_F+T}$$

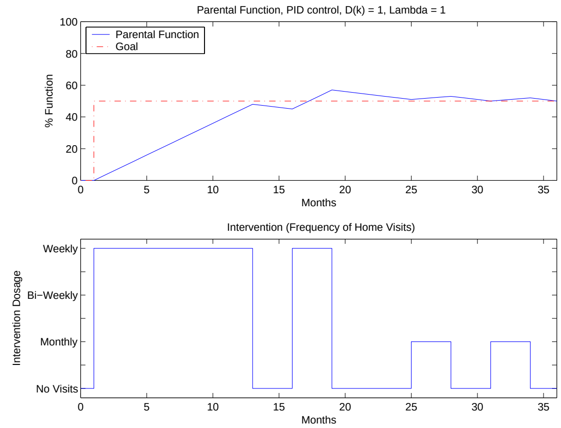

(a) $\lambda = 1$

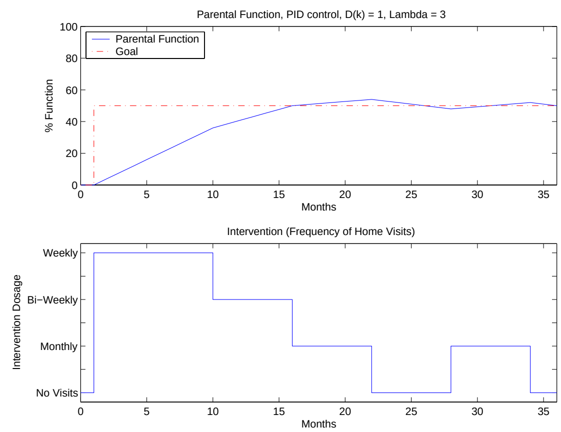

(b) $\lambda = 3$

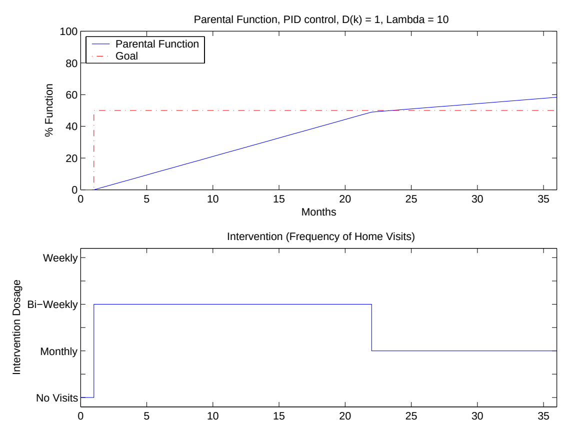

(c) $\lambda = 10$

Figure 8: Closed-loop response comparison of the discrete-time IMC-PID control law for various λ values in the case of constant one percent depletion per month ($D(t) = 1$)

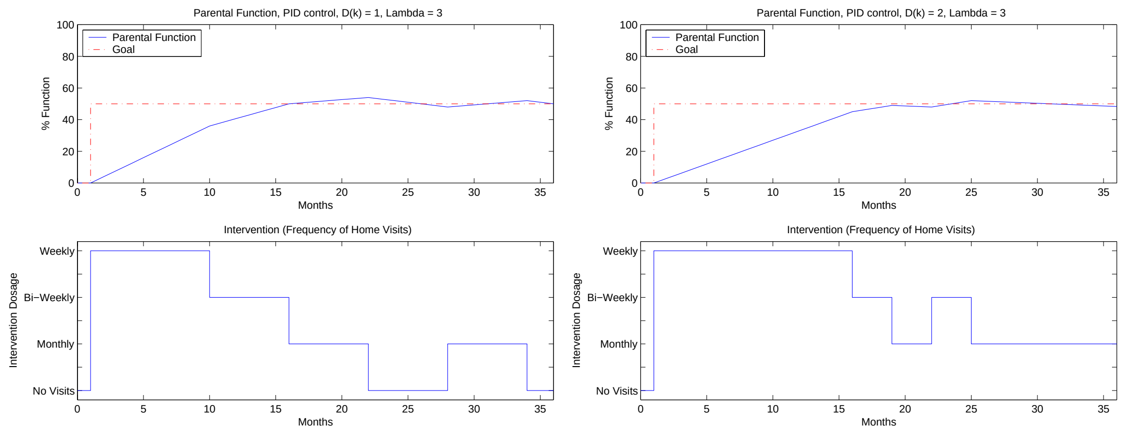

(a) (b)

Figure 9: Closed-loop responses of the discrete-time IMC-PID control law for the case of one percent depletion per month ($D(t) = 1$, left, (a)) and two percent depletion per month ($D(t) = 2$, right, (b)). $\lambda = 3$ in both cases.

Here $e(t) = R(t) - PF_{meas}(t)$ is the control error signal, which represents the discrepancy between the measured parental function and a reference setpoint $R(t)$ which is the desired goal of intervention. As a result, the current dosage of the intervention is determined as the previous dosage, plus some scaled corrections from the current and previous control errors (scaled according to the controller tuning parameters K_1, K_2, K_3, and K_4). It should be noted that the controller per (15) assumes that the intervention dosage is a continuous variable over a range (for instance, a valve position that can assume any value between 0 and 100%). In this evaluation, we quantize the intervention $I(t)$ to correspond to the closest of one of the four dosage levels (I^{weekly}, $I^{biweekly}$, $I^{monthly}$, and 0). While quantizing the action of the PID controller in this way is somewhat *ad hoc*, it does lead to acceptable results (as will be noted in the discussion of the simulation cases in the ensuing paragraphs) and motivates the need for developing controllers for this problem that include discrete-event decision-making (as in discussed in Section 4).

We first examine the effect of changing λ on the response of parental function and intervention dosages for the hypothetical Fast Track problem using the same gain (K_I) and delay (θ) parameters as before. Closed-loop responses for three values of λ ($\lambda = 1, 3,$ and 10) are presented in Figure 8. $\lambda = 1$ represents the most aggressive tuning settings (Figure 8a). While parental function is close to goal after 13 months, over the life of the intervention, the controller makes dosage assignments that jump from weekly visits to no visits (and back to weekly) in successive review intervals. Such drastic dosage changes may have detrimental effects and may be deemed unacceptable by clinical personnnel. On the other hand, setting $\lambda = 3$ (Figure 8b) leads to much more gradual dosage changes while resulting in parental function reaching the goal in a similar amount of time as when $\lambda = 1$ (around 16 months). For this value of λ, the effectiveness of the intervention is essentially the same as before, but without the erratic changes in intervention dosage. Increasing λ further to $\lambda = 10$ significantly reduces the potency of the intervention, since no dosages beyond bi-weekly visits are applied. The concomitant effect is a reduction in the speed-of-response and a longer time to reach goal (21 months versus 16 months for $\lambda = 3$), as seen in Figure 8c.

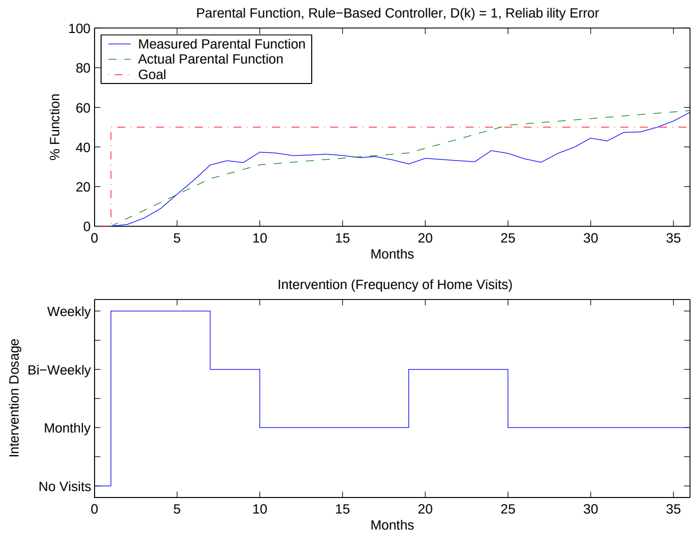

Figure 10: Closed-loop response comparison of the rule-based decision policy in the case of constant one percent depletion per month ($D(t) = 1$) under stochastic measurement error.

As an adjustable design parameter, the final value for λ will be determined by clinician preferences and any requirements imposed by the problem. From the responses presented in Figure 8, it would appear as $\lambda = 3$ represents an appropriate balance between the parental function response and the dosages applied in the intervention. It is interesting to compare the IMC-PID controller response to that of the rule-based controller for different values of the depletion rate $D(t)$. The IMC-PID response clearly outperforms the rule-based policy (with a much shorter settling time) in the case of $D(t) = 1$ (Figure 9a), and is able to eliminate the offset that plagued the rule-based control system in the case of $D(t) = 2$ (Figure 9b). An explanation for why the IMC-PID controller is able to eliminate the offset under higher rates of depletion can be attributed to two factors: 1) the IMC-PID controller acts on current and lagged values of the error signal (the difference between measurement and goal) as opposed to just the current measurement in the rule-based system, and 2) both the choice of controller structure and tuning parameters is model-based, so the character of the depletion effect is systematically recognized in the control algorithm.

In addition, we evaluate a scenario in which measurement noise $N(t)$ is introduced as a uniform zero-mean white noise signal with magnitude between 1 and -1. The response of the rule-based controller for $D(t) = 1$ is presented in Figure 10 and can be compared to the results of the IMC-PID controller under different settings of λ, as shown in Figure 11. The same realization of the measurement noise is applied in all cases. From these simulations, one is able to observe that a consequence of increasing λ is a reduction in the sensitivity of the intervention to noise. By comparing the intervention dosages calculated in Figures 8 and 11, one is able to see that the least discrepancy occurs when $\lambda = 10$. However, the price to be paid for this reduced sensitivity to the noise signal is a more sluggish closed-loop response and a less agressive intervention in general.

The final scenario we consider is to apply a nonlinear gain relationship in K_I. This is accomplished by expressing (6) as

$$PF(t+T) = PF(t) \; + \; K_I(PF)\, I(t-\theta) \; - \; D(t) \tag{16}$$

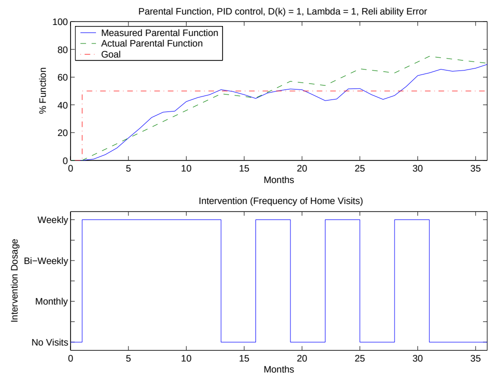

(a) $\lambda = 1$

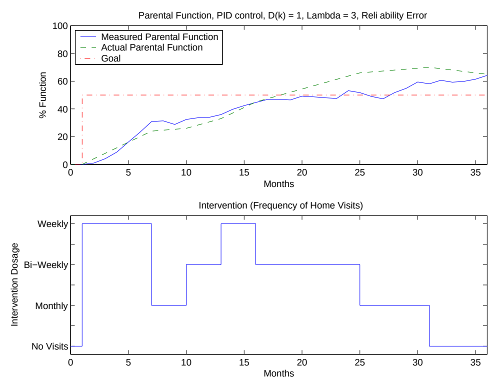

(b) $\lambda = 3$

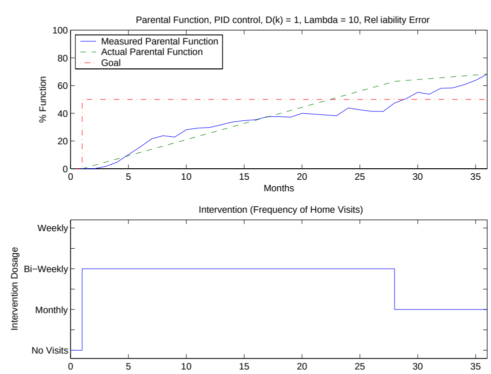

(c) $\lambda = 10$

Figure 11: Closed-loop response comparison of a discrete-time IMC-PID control law for various λ values in the case of constant one percent depletion per month ($D(t) = 1$) under stochastic measurement error.

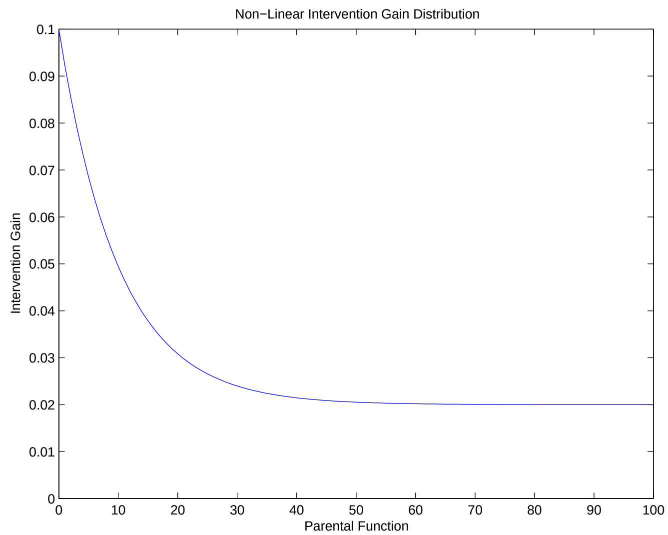

Figure 12: Graphical depiction of the nonlinear gain relationship per (17) for K_I.

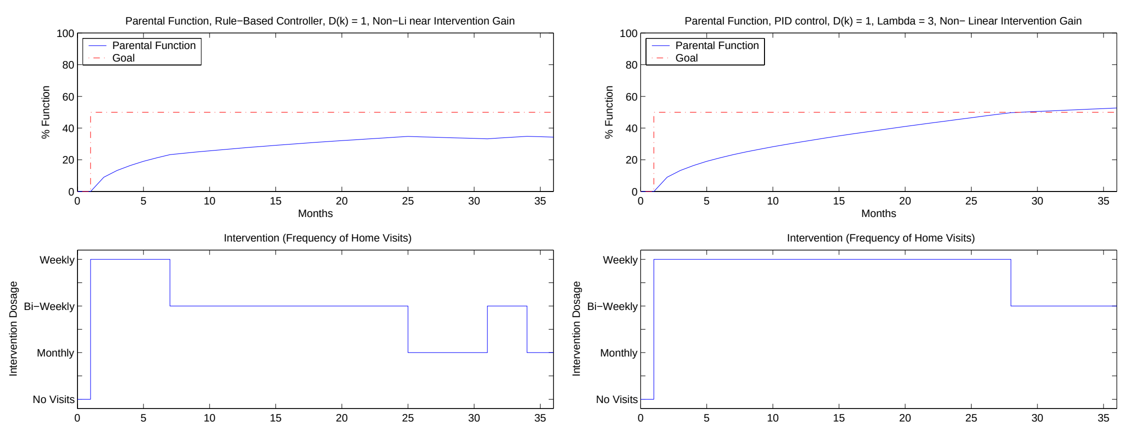

(a) (b)

Figure 13: Closed-loop response comparison of the rule-based control system per Collins *et al.* (left, (a)) and a discrete-time IMC-PID control law with $\lambda = 3$ (right, (b)) for the case of one percent depletion per month ($D(t) = 1$), with nonlinearity in K_I per (17).

where the magnitude of K_I is dependent on the value of the parental function. In this scenario we examine a nonlinear gain relationship described by the formulas

$$K_I = be^{aPF(t)} + c \tag{17}$$

$$c = K_{max} - b \qquad b = (K_{min} - K_{max})/(e^{-a} - 1)$$

A plot of (17) for $K_{max} = 0.1$, $K_{min} = 0.02$, and $a = 10$ are shown in Figure 12. Such a nonlinear relationship could describe, for instance, the possibility that higher functioning parents may receive lesser benefits from an intervention than lower functioning ones. The rule-based and IMC-PID controllers (tuned for $\lambda = 3$ and with a nominal setting of $K_I = 0.05$ as before) are contrasted in Figure 13. It is interesting to note that significant offset occurs in the rule-based case, while the engineering control policy is able to apply higher intervention dosages to take the system to goal.

4 Description of Applicable Systems Technologies

The discussion of the simulation examples in the previous section indicates that control engineering provides a systematic, quantitative framework for understanding adaptive, time-varying interventions. Control design techniques in particular can be used to develop novel decision policies for adaptive, time-varying interventions on the basis of a postulated model describing the relationship between interventions and tailoring variables. This section provides brief descriptions of some control engineering topics that promise to be useful in the pursuit of this work. Specifically we address the topics of system identification, Model Predictive Control, and Model-on-Demand estimation combined with Model Predictive Control.

4.1 System Identification

System identification refers to the field of study that is concerned with the modeling of dynamical systems from experimental data. A block diagram is shown in Figure 14. In the system identification problem, records of input and output data from the plant are used to obtain a dynamical model that best approximates the system under study. Black-box models obtained from system identification experiments are used in many fields, and are among the most common form of dynamical models used for advanced control purposes in the chemical industries.

System identification is traditionally broken down into four sub-steps: 1) experimental planning and execution, 2) data preprocessing and model structure selection, 3) parameter estimation and 4) model validation. As noted in Figure 15 (Lindskog, 1996), system identification in practice is an iterative procedure. The lack of *a priori* information regarding the plant model will require that initially each step be examined in a superficial manner. After each stage, the user must discern if the previous stages were properly accomplished; if this is not the case, the stage(s) must be redone until a satisfactory model is obtained. A satisfactory model is one that meets the requirements of the intended application (e.g., simulation, prediction, or control).

There are a number of very good texts in system identification (Ljung, 1999; Ljung and Glad, 1994) and good software for evaluating classical identification methods (e.g, the System

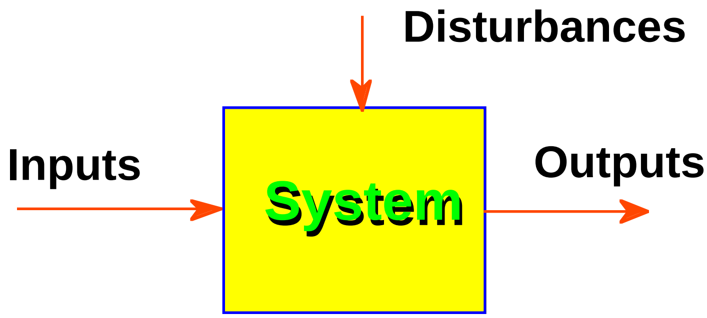

Figure 14: General block diagram for the system identification problem.

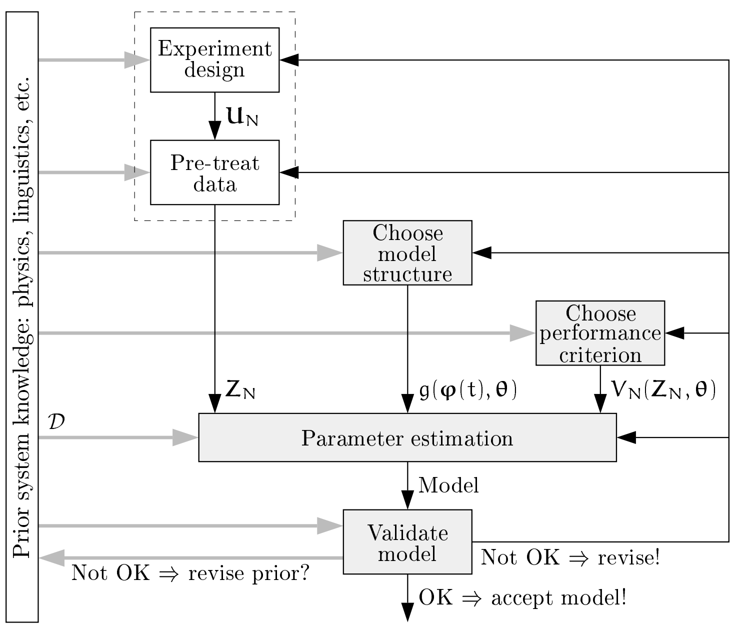

Figure 15: System Identification Loop, per Lindskog (1996) (reproduced with permission).

Identification toolbox in Matlab). While a system identification viewpoint is merited for this problem, the highly stochastic and nonlinear character of the phenomena associated with prevention problems, as well as the difficulties associated with measurement, pose significant challenges for traditional identification techniques. Extensions of identification techniques to address these difficulties merit further investigation.

System identification techniques have the potential to influence the development and dissemination of effective adaptive interventions in significant ways. Among these is the use of insights generated from system identification into the design of experimental trials geared for adaptive interventions, and the use of "control-relevant" identification to obtain simple yet useful models from data for the construction of decision rules. In experimental design for system identification, the concern is to decide on the appropriate sequence of inputs of the system so that the output measurements display sufficient information to properly estimate a model of the system dynamics. The system identification literature defines the concept of "persistence of excitation" (Ljung, 1999) to characeterize an input signal rich enough to generate an informative model estimate. In practice, however, time and operating constraints are considerations that must be factored in the design of a suitable input signal and subsequent experimental trial. To this end, the concept of "plant-friendliness" in system identification has been developed (Rivera *et al.*, 2003). A plant-friendly test will produce data leading to a suitable model within an acceptable time period, while keeping the variation in both input and output signals within user-defined constraints. How to contextualize the notion of plant-friendly identification testing to problems in prevention and treatment in behavioral health represents an open research problem. Of particular interest is the interplay between these ideas and the design of sequential multiple assignment randomized (SMAR) trials, which are currently being evaluated by Collins *et al.* (in press) and Murphy (in press) in the context of adaptive interventions.

4.2 Model Predictive Control

Model Predictive Control (MPC) stands for a family of methods that select control actions based on on-line optimization of an objective function. MPC has gained wide acceptance in the chemical and other process industries as the basis for advanced multivariable control schemes (Prett and García, 1988; Camacho and Bordons, 1999; Qin and Badgwell, 2003). In MPC, a system model and current and historical measurements of the process are used to predict the system behavior at future time instants. A control-relevant objective function is then optimized to calculate a sequence of future control moves that must satisfy system constraints. The first predicted control move is implemented and at the next sampling time the calculations are repeated using updated system states; this is referred to as a Moving or Receding Horizon strategy (as illustrated in Figure 16).

Model Predictive Control represents a general framework for control system implementation that accomplishes both feedback and feedforward control action on a dynamical system. The appeal of MPC over traditional approaches to feedback and feedforward control design include 1) the ability to handle large multivariable problems, 2) the explicit handling of constraints on system input and output variables, and 3) its relative ease-of-use. In this latter category, MPC represents a completely "time-domain" technology that avoids the cumbersome closed-form solutions associated with classical optimal control and some modern approaches (such as H_2 and H_∞ control).

The goal of the MPC decision policy is to seek a future profile for u, the manipulated variables,

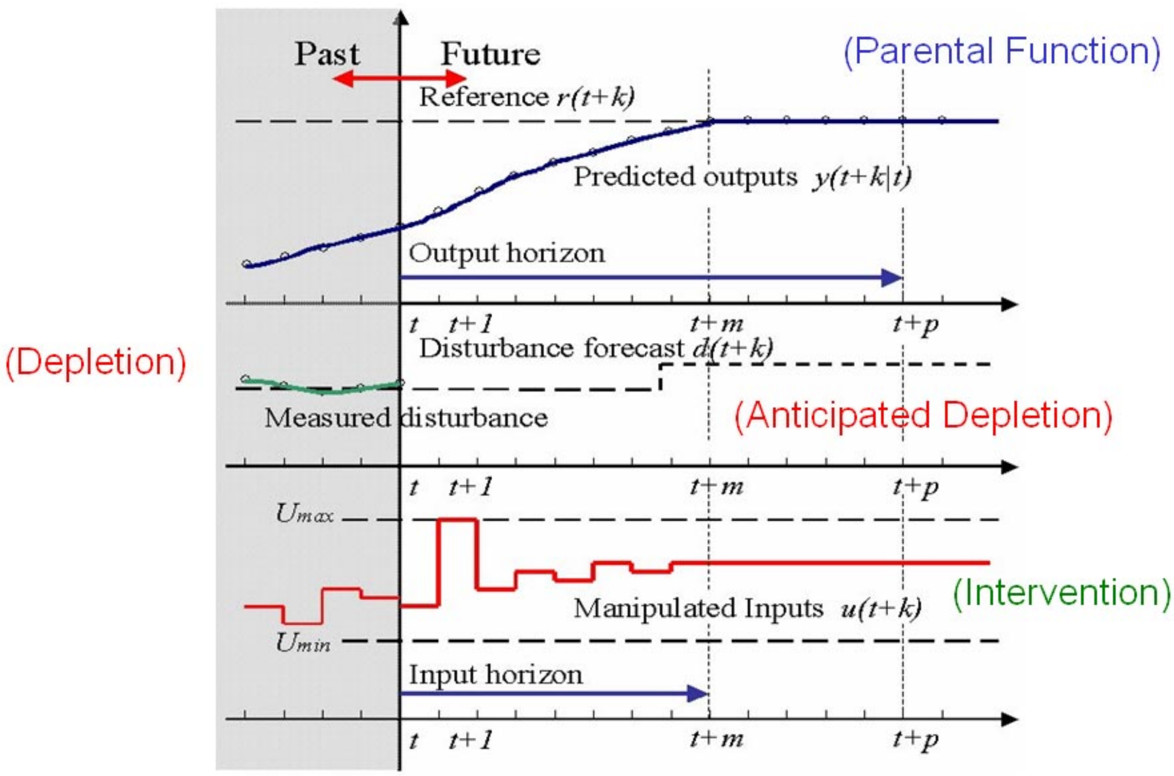

Figure 16: Moving horizon representation of Model Predictive Control.

that brings the system variables to some desired conditions per the minimization of an objective function. There is significant flexibility in the form of the objective function that can be used in MPC; a meaningful formulation for the adaptive, time-varying intervention problems considered in this technical report is as follows:

$$\min_{\Delta u(t|t) \ldots \Delta u(t+m-1|t)} J \tag{18}$$

where the individual terms of J correspond to:

$$J = \overbrace{\sum_{\ell=1}^{p} Q_e(\ell)(\hat{y}(t+\ell|t) - r(t+\ell))^2}^{\text{Keep Controlled Variables at Goal}} + \overbrace{\sum_{\ell=1}^{m} Q_{\Delta u}(\ell)(\Delta u(t+\ell-1|t))^2}^{\text{Penalize Changes in the Intervention Dosages}} \tag{19}$$

$$+ \overbrace{\sum_{\ell=1}^{m} Q_u(\ell)(u(t+\ell-1|t) - u_{target}(t+\ell-1|t))^2}^{\text{Maintain Intervention Dosages at Desired Targets}}$$

The objective function in MPC is a multi-term expression that can be used to address the main operational objectives in an adaptive, time-varying intervention:

1. A *setpoint tracking* term that is intended to maintain controlled variables (which could be either tailoring variables, outcomes, or their combination) at user-specified targets over time. These targets need not be constant and can change over the prediction horizon p.

2. A *move suppression* term that penalizes successive changes (also referred to as *moves*) in the intervention dosages. Move suppression also serves an important control-theoretic purpose as the primary means for achieving robustness in the controller in the face of uncertainty (Prett and García, 1988); the larger the weight on the move suppression, the greater the emphasis in the control system to not introduce any changes in the assigned dosages. Such control action may also be desirable to clinical personnel, who may not want to see drastic changes in dosage recommendations from the intervention between review intervals.

3. An *input target* term that is meant to maintain the intervention dosages close to desired pre-determined target values (such as those that may be associated with a fixed intervention).

The emphasis given to each one of the sub-objectives in (19) (or to specific system variables within these objective terms) is achieved through the choice of weights ($Q_e(\ell)$, $Q_{\Delta u}(\ell)$, and $Q_u(\ell)$); these can potentially vary over the move and prediction horizons (m and p, respectively). Constraints are an important feature of most real-life control problems, and the ability to address these explicitly in the controller formulation is part of the significant appeal of MPC. Constraints can be imposed on the magnitudes of the intervention dosages, the changes in these dosages, tailoring variables and outcomes, and similar system variables. For an MPC problem with an objective function per (18)-(19), relying on linear discrete-time state-space models to describe the dynamics, and subject to linear inequality constraints, a numerical solution is achieved via quadratic programming. However, depending on the nature of the objective function, model and constraint equations, other programming approaches (linear programming or nonlinear programming) may be involved (Morshedi *et al.*, 1985; Vargas-Villamil *et al.*, 2003). Efficient solvers for linear and quadratic programming problems are widely available, and can be integrated with readily available tools such as Excel (Frontline Systems, 2005).

It must be noted the nature of the decision making inherent in the adaptive intervention problem calls for MPC paradigms that consider both discrete-event and discrete-time decisions. Such decision policies would include choosing from among a set of different possible treatments, and assign the appropriate dosages. Recent developments in so-called "hybrid" predictive control (Bemporad and Morari, 1999) which enable the explicit inclusion of propositional logic into the decision-making algorithms associated with MPC, are particularly appropriate. The extension of Model Predictive Control with data-centric estimation is also appropriate and is discussed in the following section.

4.3 Model-on-Demand Estimation and Model Predictive Control

Model-on-Demand Predictive Control represents a data-centric decision framework that represents an effective integration of identification and Model Predictive Control. The MoD approach provides the potential for performance rivaling that of global methods (such as neural networks) while involving less complex *a priori* knowledge from the user and providing more reliable numerical computations. Furthermore, in comparison to fuzzy-modeling and similar local modeling techniques, the user is not forced to decide how many local models are required or how the controller will transition between them.

In recent years, data-centric estimation methods have received significant attention in the systems literature, and signficant applications have been reported (Kulhavý, 2003). One such novel concept for nonlinear identification and control is the Model-on-Demand (MoD) framework (Stenman, 1999), which is inspired by ideas from local modeling and database systems technology (Atkeson *et al.* (1997*a*; 1997*b*)). In MoD estimation all observations are stored on a database, and the models are built "on demand" as the actual need arises. The Model on Demand predictor relies on small portions of the data relevant to the region of interest to determine a model as needed. The variance/bias tradeoff inherent to all modeling is optimized locally by adapting the number of data and their relative weighting. From a practical standpoint, MoD estimation allows process engineers to naturally extend insight gained from linear modeling to nonlinear identification and control. Rather than spending time in a difficult structure selection and parameter initialization, the user can focus on developing informative datasets – a common requirement for all nonlinear black-box identification approaches. The user can fully examine the uniformity of coverage the excitation signals produced in the input and output spaces and better understand the impact it has on the nonlinear model validation and control problem. MoD can be formulated into a comprehensive methodology for nonlinear identification and predictive control (Braun, 2001; Braun *et al.*, submitted). A Matlab-based tool for MoD estimation and control, developed at the Control Systems Engineering Laboratory at ASU in collaboration with researchers from the Division of Automatic Control at Linköping University, is available in the public domain (Braun *et al.*, 2002).

We believe that the adaptive, time-varying intervention problem will call for novel forms of MoD-based predictive control systems that can address the challenges of a low signal-to-noise datasets and exhibit hybrid decision making (where both continuous-time and discrete-event decision-making may be necessary). Such a body of control engineering would be useful for a broad class of problems in behavioral health and could also impact problems in enterprise systems (such as supply chains).

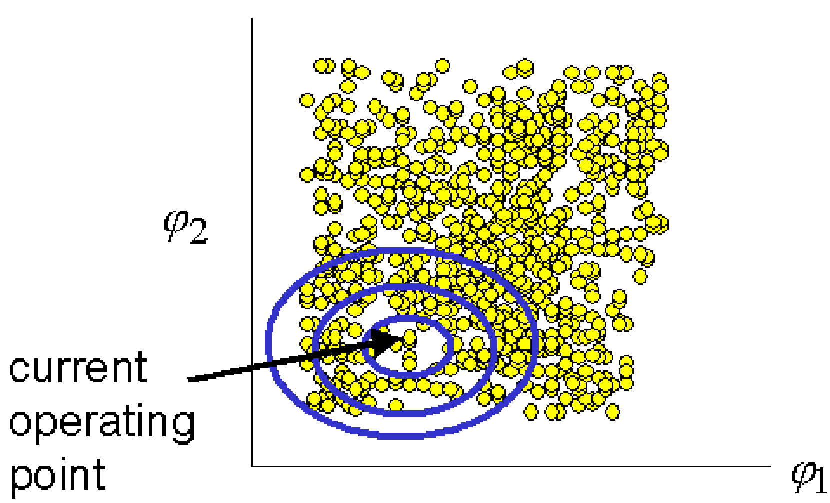

Figure 17: Depiction of local regressor selection for Model-on-Demand estimation

5 Summary and Conclusions

This report has established some conceptual linkages between the problem of adaptive, time-varying interventions in prevention and engineering process control. Specifically, we have shown that adaptive, time-varying interventions are feedback control systems, with the tailoring variable acting as the controlled variable, the intervention representing the manipulated variable, and decision rules serving the role of feedback control laws. A simulated study involving a hypothetical intervention inspired by the Fast Track program compared the performance of a control policy based on a series of IF-THEN decision rules with a model-based engineering PID controller tuned using Internal Model Control. While the IF-THEN controller performed adequately under mostly deterministic conditions with little or no disturbances (i.e., depletion), the IMC-PID controller was demonstrated to perform well under conditions of significant disturbances, measurement noise, and nonlinearity. In particular, the IMC-PID controller includes an adjustable parameter that can be used to influence the aggressiveness of the intervention and thus the sensitivity of dosage assignment to changes in the estimated tailoring variables. The ability to determine the "speed-of-response" of the intervention in this way enables the clinician making use of this adaptive intervention to influence the performance/robustness tradeoff associated with feedback control systems, and thus makes the use of engineering control principles a viable option in the control of highly nonlinear and stochastic phenomena, as that encountered in prevention and treatment problems in behavioral health.

From an engineering control perspective, there are fundamental modeling challenges associated with lack of a direct link to physical conservation relationships in the phenomena under study, the presence of low signal-to-noise ratios in the data, and confounding/bias among variables. Decision rule design challenges include the presence of both discrete-event and continuous decision variables, which leads to the potential need for hybrid data-centric modeling and control decision frameworks that currently do not exist in the control literature. To this end, we propose taking advantage of developments in the fields of system identification, Model Predictive Control, and Model-on-Demand estimation and control to develop a novel body of comprehensive engineering control technology geared towards adaptive, time-varying interventions and related problems in the behavioral sciences.

Note: The Excel spreadsheet used to obtain the simulation results in this technical report can be downloaded from the ASU Control Systems Engineering Laboratory website at:

`http://wwww.fulton.asu.edu/~csel/Software.html`.

6 Acknowledgements

Support for this work has been provided by the National Institute on Drug Abuse (NIDA) through a subcontract to Arizona State University (2778-ASU-DHHS-0075) from the Penn State University Methodology Center (prime award P50 DA010075).

References

Atkeson, C.G., A.W. Moore and S. Schaal (1997*a*). Locally weighted learning. *Artificial Intelligence Review* **11**(1-5), 11–73.

Atkeson, C.G., A.W. Moore and S. Schaal (1997*b*). Locally weighted learning for control. *Artificial Intelligence Review* **11**(1-5), 75–113.

Bemporad, A. and M. Morari (1999). Control of systems integrating logic, dynamics, and constraints. *Automatica* **35**(3), 407 – 427.

Braun, M. W. (2001). Model-on-Demand nonlinear estimation and Model Predictive Control: Novel methodologies for process control and supply chain management. PhD thesis. Dept. of Chemical and Materials Engineering, Arizona State University.

Braun, M.W., A. Stenman and D.E. Rivera (2002). Model-on-Demand Model Predictive Control Toolbox. `http://www.eas.asu.edu/~csel/Software.htm`.

Braun, M.W., D.E. Rivera and A. Stenman (submitted). Model-on-Demand Model Predictive Control for nonlinear process systems. *Computers and Chemical Engineering*.

Brooner, R.K. and M. Kidorf (2002). Using behavioral reinforcement to improve methadone treatment participation. *Science and Practice Perspectives* **1**, 38–48.

Camacho, E. F. and C. Bordons (1999). *Model Predictive Control*. Springer-Verlag. London.

Collins, L.M. and J.W. Graham (2002). The effect of the timing and spacing of observations in longitudinal studies of tobacco and other drug use: temporal design considerations. *Drug and Alcohol Dependence* **68**, S85 – S96.

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

Collins, L.M., S.A. Murphy, V.N. Nair and V.J. Strecher (in press). A strategy for optimizing and evaluating behavioral interventions. *Annals of Behavioral Medicine*.

Frontline Systems, Inc. (2005). Optimization with Excel solver - Frontline Systems. `http://www.solver.com`.

Glasgow, M.S., B.T. Engel and B.C. D'Lugoff (1989). A controlled study of a standardized behavioral stepped treatment for hypertension. *Psychosomatic Medicine* **51**, 10 – 26.

Kulhavý, Rudolf (2003). A developer's perspective of a decision support system: A look at the data-centric decision support system's technical origins and systems concepts. *IEEE Control Systems Magazine* **23**(6), 40–49.

Lindskog, P. (1996). Methods, Algorithms, and Tools for System Identification Based on Prior Knowledge. PhD thesis. Linköping University, Sweden. Dept. of Electrical Engineering.

Ljung, L. (1999). *System Identification: Theory for the User*. 2nd ed.. Prentice-Hall. New Jersey.

Ljung, L. and T. Glad (1994). *Modeling of Dynamic Systems*. Prentice-Hall. New Jersey.

Conduct Problems Prevention Research Group (1992). A developmental and clinical model for the prevention of conduct disorders: The Fast Track program. *Development and Psychopathology* **4**, 509 – 527.

Conduct Problems Prevention Research Group (1999*a*). Initial impact of the Fast Track prevention trial for conduct problems: I. The high-risk sample. *Journal of Consulting and Clinical Psychology* **67**, 631 – 647.

Conduct Problems Prevention Research Group (1999*b*). Initial impact of the Fast Track prevention trial for conduct problems: II. Classroom effects. *Journal of Consulting and Clinical Psychology* **67**, 648 – 657.

MC-DATS (2004). Methodological challenges in developing adaptive treatment strategies. http://neuromancer.eecs.umich.edu/dtr/twiki/bin/view.

Morari, M. and E. Zafiriou (1989). *Robust Process Control*. Prentice Hall. Englewood Cliffs, NJ.

Morshedi, A. M., C. R. Cutler and T. A. Skrovanek (1985). Optimal solution of dynamic matrix control with linear programming techniques (LDMC). In: *Proceedings of 1985 American Control Conference*. Boston, Massachusetts.

Murphy, S.A. (in press). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*.

Murphy, S.A. and J.R McKay (Winter 2003/Spring 2004). Adaptive treatment strategies: An emerging approach for improving treatment effectiveness. *Clinical Science* (Newsletter of the American Psychological Association Division 12, section III: The Society for the Science of Clinical Psychology).

Ogunnaike, B. A. and W. H. Ray (1994). *Process Dynamics, Modeling, and Control*. Oxford University Press. New York.

Prett, D.M. and C.E. García (1988). *Fundamental Process Control*. Butterworth. Stoneham, MA.

Qin, S. J. and T. A. Badgwell (2003). A survey of industrial model predictive control technology. *Control Engineering Practice* **11**, 733 – 764.

Rivera, D. E., M. Morari and S. Skogestad (1986). Internal Model Control. 4. PID controller design. *Ind. Eng. Chem. Process Des. Dev.* **25**, 252 – 265.

Rivera, D.E., H. Lee, M.W. Braun and H.D. Mittelmann (2003). Plant-friendly system identification: a challenge for the process industries. In: *13th IFAC Symposium on System Identification (SYSID 2003)*. Rotterdam, Netherlands. pp. 917–922.

Rosenberg, E.S., M. Altfield, S.H. Poon, M.N. Phillips, B.M. Wilkes, R.L. Eldridge, G.K. Robbins, R.T. D'Aquila, P.J.R. Goulder and B.D. Walker (2000). Immune control of HIV-1 after early treatment of acute infection. *Nature* **407**, 523–526.

Rush, A.J., M. Fava, S.R. Wisniewski, P.W. Lavori, M.H. Trivedi, H.A. Sackeim, M.E. Thase, A.A. Nierenberg, F.M. Quitkin, T.M. Kashner, D.J. Kupfer, J.F. Rosenbaum, J. Alpert, J. Stewart, P.J. McGrath, M.M. Biggs, K. Shores-Wilson, B.D. Lebowitz, L. Ritz and G. Niederehe (2004). Sequenced Treatment Alternatives to Relieve Depression (STAR*D): Rationale and design.. *Controlled Clinical Trials* **25**(1), 119 – 142.

Schneider, L.S., P.N. Tariot, C.G. Lyketsos, K.S. Dagerman KS, K.L. Davis, S. Davis, J.K. Hsiao, D.V. Jeste, I.R. Katz, J.T. Olin, B.G. Pollock, P.V. Rabins, G.W. Small, B. Lebowitz and J.A. Lieberman (2001). National Institute of Mental Health Clinical Antipsychotic Trials in Intervention Effectiveness (CATIE). *American Journal of Geriatric Psychiatry* **9**, 346 – 360.

Sobell, M.B. and L.C. Sobell (1999). Stepped care for alcohol problems: an efficient method for planning and delivering critical services. In: *Changing Addictive Behavior: Bridging Clinical and Public Health Strategies* (J.A. Tucker, D.M. Donovan and G.A. Marlatt, Eds.). pp. 331 – 343. New York:Guilford Press.

Stenman, Anders (1999). Model on Demand: Algorithms, Analysis and Applications. PhD thesis. Dept. of Electrical Engineering, Linköping University, Sweden.

Vargas-Villamil, F. D., D. E. Rivera and K. G. Kempf (2003). A hierarchical approach to production control of reentrant semiconductor manufacturing lines. *IEEE Transactions On Control Systems Technology* **11**(4), 578–587.

Velicer, W.F. and J.O. Prochaska (1999). An expert system intervention for smoking cessation. *Patient Education and Counseling* **36**, 119 – 129.