2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

Targeting a fixed percentage of granulocyte differentiation using experiments designed via nonlinear model predictive control

Sarah L. Noble and Ann E. Rundell

Abstract—**Previous efforts to control cellular differentiation have been largely experimental. Although some mathematical models for this process exist, rarely has a quantitative approach been employed to design experiments that predictably direct the cell fate. As an initial step towards this aim, a control strategy for sustaining a desired constant level of differentiated human promyelocytic leukemia (HL60) cells is proposed. Nonlinear model predictive control is applied using a model of HL60 cell differentiation which reflects the dominant observable cell states along the known granulocyte/monocyte differentiation pathway. After implementing changes identified by a preliminary experimental trial, the experimental realization of the controller successfully achieved the target differentiation level and demonstrated the need for a quantitative approach to experiment design.**

I. Introduction

Controlling cellular processes involves unique challenges not encountered in the control of traditionally engineered systems. Real-time continuous feedback is generally not available and realizable control actions are limited. Mathematical models of the cellular processes are uncertain abstractions of complex biochemical and gene regulatory networks. As a result, there have been minimal efforts to apply control theory at the cellular level. Recent work in this area has included circadian phase entrainment [1], bacterial chemotaxis [2], calcium oscillations [2], and cellular differentiation [3], with none of the proposed control strategies experimentally evaluated. The application of control theory to these types of problems will help quantitatively design experiments to direct cellular processes and minimize expensive and exhaustive experimental efforts. There is no doubt that developing a quantitative method for predictably controlling stem cell differentiation into specific lineages would be a significant contribution to stem cell-based technologies, such as tissue engineering [4].

In this work, a nonlinear model predictive control (MPC) strategy is used to control the differentiation of human promyelocytic leukemia (HL60) cells into mature granulocytes and monocytes. MPC is employed as it is known to be robust to model uncertainties, measurement noise, and output disturbances [5], and it naturally accommodates the discrete-time feedback necessary for the slow-moving dynamics of a differentiating cell population. Control is realized through periodic additions, or boluses, of the differentiation-inducing chemical. The sequence of boluses is found by the optimal selection of the controller parameter such that a desired level of granulocytes is sustained.

The supporting mathematical model for this work is developed in Section III based upon known HL60 cell differentiation behaviors briefly described in Section II. The parameters for the model are identified from the results of a set of preliminary experiments as described in Section IV. Section V presents the nonlinear model predictive controller design approach with the experimental control strategy implementation described in Section VI. An initial experimental trial evaluated the first proposed experimental protocol as described in Section VII and resulted in modifications to the model and process. The final experimental results are presented in Section VIII, and conclusions and future work are described in Section IX.

II. Background Information

The human promyelocytic leukemia HL60 cell line was established from a human patient in 1977 and is commonly used for studying cell differentiation and leukemia treatment. HL60 cells can be induced to differentiate into granulocytes, monocytes, macrophage-like cells, and eosinophils [6]. While HL60 cells will naturally differentiate into monocytes and granulocytes in very small numbers, full-scale differentiation must be induced using chemical agents. Of interest herein, exposing cells to dimethyl sulfoxide (DMSO) initiates differentiation into granulocytes [7]. The dynamics of HL60 cell differentiation are largely dependent on the concentration of DMSO present during the period of incubation.

III. Model Development

A system of nonlinear ordinary differential equations describes how the population of cells progresses through discrete maturation stages over time. A graphical representation of the system is shown in Fig. 1. The maturation stages represent benchmarks which are experimentally distinguishable using flow cytometry.

All cells begin in the first stage as undifferentiated HL60 cells. The first indication of differentiation is the expression of the cell-surface-localized cluster of differentiation (CD) CD11b. A maturing granulocyte will then express the marker CD16, and a maturing monocyte will express the

Manuscript received March 13, 2009.

S.L. Noble is with the Electrical and Computer Engineering Department, Purdue University, West Lafayette, IN 47906 (phone: 765-496-7953; e-mail: slnoble@purdue.edu).

A.E. Rundell is with the Biomedical Engineering Department, Purdue University, West Lafayette, IN 47906 (e-mail: rundell@purdue.edu)

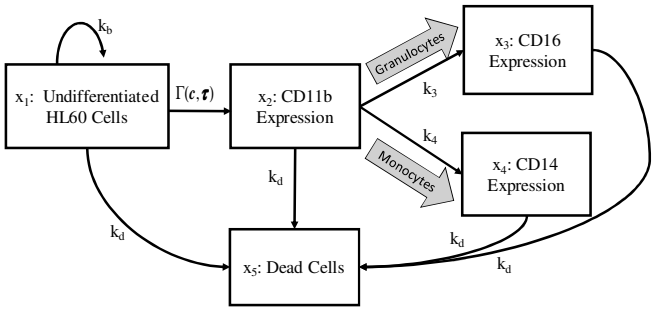

Fig. 1. States of the HL60 differentiation model. Cells begin to differentiate according to a concentration- and duration-dependent transition rate, $\Gamma(\mathbf{c},\boldsymbol{\tau})$. Because only undifferentiated HL60 cells can proliferate, a constant birth rate, k_b, is associated with the first state. Cells can die in any phase according to a constant, phase-independent death rate, k_d.

marker CD14 [9].

A constant birth rate, k_b, is present in the first state as only undifferentiated cells can proliferate. Cells transition into the second state according to a concentration- and duration-dependent transition rate, $\Gamma(\mathbf{c},\boldsymbol{\tau})$. The number of cells transitioning depends on the concentration of each previous DMSO bolus, c_i, and the time since the bolus administration, τ_i. It is assumed that the concentration of DMSO affects only the initiation of differentiation, so the transition rates into the third and fourth maturation states are constant. Cells die according to a constant, state-independent death rate, k_d.

The ODE model is constructed from the dynamics of Fig. 1 as:

$$\begin{aligned}
\dot{x}_1 &= k_b x_1 - \Gamma(\mathbf{c},\boldsymbol{\tau}) x_1 - k_d x_1 \\
\dot{x}_2 &= \Gamma(\mathbf{c},\boldsymbol{\tau}) x_1 - k_3 x_2 - k_4 x_2 - k_d x_2 \\
\dot{x}_3 &= k_3 x_2 - k_d x_3 \\
\dot{x}_4 &= k_4 x_2 - k_d x_4 \\
\dot{x}_5 &= k_d (x_1 + x_2 + x_3 + x_4)
\end{aligned} \tag{1}$$

where x_i represents the number of cells in the i^{th} maturation state. The transition rate, $\Gamma(\mathbf{c},\boldsymbol{\tau})$, describes the cells beginning to differentiate:

$$\Gamma(\mathbf{c},\tau) = \mathbf{V}(\mathbf{c})^T \mathbf{G}(\tau)$$

where

$$\mathbf{V}(\mathbf{c}) = k_m \mathbf{c} + k_i$$

describes the maximum transition rate as a linear function of differentiation-inducing agent concentration and

$$G_j(\tau_j) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\tau_j - \mu)^2}{2\sigma^2} \right)$$

describes the likelihood of transitioning after the j^{th} bolus administration as a Gaussian distribution function of the duration since that bolus, τ_j. The Gaussian distribution with mean, μ, and standard deviation, σ, represents the period of time for which the undifferentiated cells are responsive to the bolus. At any given time, the effects of each Gaussian distribution (resulting from each administered bolus) are additive.

The concentration vector, $\mathbf{c}$, stores the concentration of each administered bolus, and the duration vector, $\boldsymbol{\tau}$, tracks how much time has elapsed since each bolus was given. The j^{th} bolus, corresponding to concentration c_j, is administered at τ_{0_j}. These duration dynamics are modeled as Heaviside functions:

$$\dot{\tau}_j = 1^+(t - \tau_{0_j}).$$

Most cells begin as undifferentiated HL60 cells, so the initial condition is assumed to be:

$$\mathbf{x}(0) = [N_0 \quad 0 \quad 0 \quad 0 \quad 0]^T,$$

where N_0 is the initial number of cells. The duration vector, $\boldsymbol{\tau}$, is initialized to zero.

IV. Experimental Methods for Model Parameter Identification

A. Cell Culturing

HL60 cells were grown in suspension at 37ºC and 5% CO_2 in 20mL of an 80/20 mixture of RPMI-1640 (ATCC) and fetal bovine serum (Harlan), supplemented with 1% P/S. Balanced growth conditions were maintained by passing the cells to a concentration of 10^5 – 10^6 cells/mL every two days. One day before an experiment, cells were passed down to a concentration of $6\text{x}10^5$ cells/mL. Cells were not passed during the experiment.

B. Staining Procedure

Cells were induced to differentiate using five concentrations (v/v) of DMSO: 1.2%, 1.0%, 0.8%, 0.6%, and 0.4%. One sample was taken from each population every day for seven days. For each sample, a 1mL aliquot was aspirated and washed in 1mL stain buffer (Pharmingen). The cell pellet was exposed to 10µL of each antibody (BD Bioscience) for 20 minutes in a covered ice bucket. PE, PE-Cy5, and FITC conjugated fluorescent antibodies were used to label CD11b, CD16, and CD14, respectively. Cells were washed twice in 1mL stain buffer.

In preparation for flow cytometry, cells were resuspended in 100µL Cytofix/Cytoperm (BD Bioscience), a 4% paraformaldehyde fixation buffer, for 30 minutes in a covered ice bucket. Cells were washed and resuspended in 1mL stain buffer.

C. Flow Cytometry Analysis

Cytometry analysis was performed within one week of sample collection. Samples were analyzed with a BD-Elite flow cytometer with 20,000 cells analyzed per sample. Emissions were measured at 530, 575, and 683nm. The percentage of cells expressing CD11b, CD16, and CD14 were found by gating the flow cytometry data in WinMDI 2.8.

D. Model Parameter Identification

A model parameter set was identified using Matlab's genetic algorithm to fit (1) to the experimental data. To ensure the simulated model dynamics captured the response to different DMSO concentrations, the model was fit to experimental data taken at five different DMSO

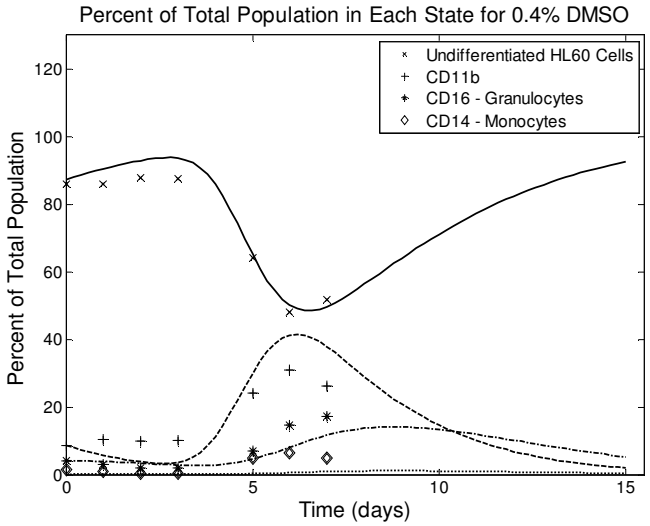

Fig. 2. Model simulation (lines) fit to experimental data (symbols) for cells exposed to 0.4% DMSO. The model also predicts future system behavior. Model parameters: k_b = 0.321/day, k_d = 0.005/day, k_3 = 0.126/day, k_4 = 0.010/day, k_m = 2.393, k_i = 0.126, μ = 5.137 days, σ = 0.957 days.

concentrations. Fig. 2 shows an example of the model fit to experimental data taken at 0.4% DMSO.

V. Model Predictive Control Problem Formulation

The objective of this study was to determine a sequence of DMSO boluses to reach and sustain a cell population containing a fixed percentage of mature granulocytes over a three-week period. Nonlinear MPC uses the mathematical model of the underlying process to predict the future behavior of the system over a finite prediction horizon. The controller calculates an appropriate input sequence by solving a constrained optimization problem. In this case, the optimization problem minimized the deviation of the system trajectory from the reference trajectory, as well as the magnitude of the control input:

$$\arg\min_{u}\left(Q\sum_{i=1}^{H_p}\left[\hat{y}(k+i|k)-s\right]^2 + R\sum_{i=1}^{H_u}\left[u(k+i|k)\right]^2 \right), \quad (2)$$

where $\hat{y}(k+i|k)$ is the predicted output trajectory, s is the target percentage of granulocytes, and $u(\cdot)$ is the control input, DMSO dose. Because the controlled output was the percent of total cells that are mature granulocytes (state x_3), the predicted output trajectory was:

$$\hat{y}(k+i|k) = \frac{\begin{bmatrix}0 & 0 & 1 & 0 & 0\end{bmatrix}\hat{x}(t_{k+i})}{\begin{bmatrix}1 & 1 & 1 & 1 & 0\end{bmatrix}\hat{x}(t_{k+i})}\times 100$$

where $\hat{x}(t_{k+i})$ represents the model-predicted state vector at each sampling time. The parameters used in the optimization are shown in Table I.

Matlab's constrained optimization solver, `fmincon`, was used to minimize the objective function subject to the following control input constraint:

$$u(k+i|k) \in [0,\ 1.2].$$

The rationale for the higher limit comes from experimental results which indicate that a DMSO concentration higher than ~1.2% inhibits differentiation by preventing cell replication [8].

VI. Experimental Control Strategy Implementation

To experimentally implement the MPC-derived control strategies, feedback data came from a population of differentiating HL60 cells, which served as the plant. Samples of the differentiating cell population were analyzed using flow cytometry to find the percent of the population in each maturation stage. Cells were stained in the same manner outlined in Section IV, except samples were immediately analyzed using flow cytometry rather than being fixed for later analysis. At each sampling time, the experimental results served as the initial conditions from which the model of (1) predicted the expected results for the prediction horizon (H_p). The control parameters were determined by solving the optimization problem of (2) to find the control sequence (of length H_u). The first element of the control sequence was implemented as a bolus of DMSO in the flask of differentiating HL60 cells. At the next sampling time, another sample of the cell population was taken and analyzed via flow cytometry to find the state of the plant. While equally-spaced sampling times are usually used for an MPC experiment, an exception was made in this case as the flow cytometry facilities were unavailable on weekends. For this reason, sampling times alternated between 4 days and 3 days.

The long duration of the experiment necessitated changes to the original cell culturing procedure outlined in Section IV. Undifferentiated HL60 cells continually proliferate and, by the end of a three-week experiment, cell growth would have halted due to a lack of nutrients. Even with the addition of fresh media without the removal of any cells, the final cell population would be prohibitively large. Hence, a portion of the volume (containing cells) in each flask was removed daily and replaced with fresh media to maintain a constant concentration of $6\text{x}10^5$ cells/mL. This daily removal of cells was reflected in the model implementation used for the MPC by assuming a fixed fraction of the cell population in each maturation stage was removed.

VII. Initial Experimental Trial Suggests Modifications

The MPC-design process from Section V was experimentally implemented as outlined in Section VI. The

TABLE I

Parameter	Description	Value
H_p	Prediction horizon	21 days
H_u	Control horizon	7 days
s	Target trajectory	15% of population granulocytes
Q	Weight for matching trajectory	10
R	Weight for minimizing control input	1

Parameters for MPC constrained optimization problem of (2).

TABLE II

Time Point (days)	Control Input
0	0.42%
4	0.03%
7	0.34%
11	0.13%
14	0.26%

Each control input gives the application dose of DMSO (v/v). MPC-derived control strategy for experimental study of Fig. 3.

administered control sequence is stated in Table II.

It was noticed after day 7 of the experiment that the level of differentiation was nearly zero (see Fig. 3). This was inconsistent with the previous results obtained for the parameter identification with a similar DMSO dosage (see Fig. 2). We reexamined the originally-perceived minimal change in the protocol (namely, the daily passage of cells) and discovered that the new protocol lead to a higher proliferation rate. Retrospectively, this was not surprising due to the daily addition of fresh media. But more importantly, there was a dilution of the DMSO concentration in the media during the attempt to keep the number of cells in the flask constant.

To investigate the effects of this dilution, on day 11 the experimental flask was divided into two flasks. The first flask was cultured with the same DMSO dilution as before, but the second flask was cultured such that DMSO was replaced after each daily passing to simulate a constant concentration. The time- course experimental results for the diluted-DMSO flask are shown in Fig. 3 as open bars, and the time-course experimental results for the constant-DMSO flask are shown as solid bars. The revision of the experimental procedure on day 11 to keep the DMSO concentration constant between sampling points did cause a small but visible increase in the level of differentiated cells. These results motivated changes to the experimental protocol and MPC parameters.

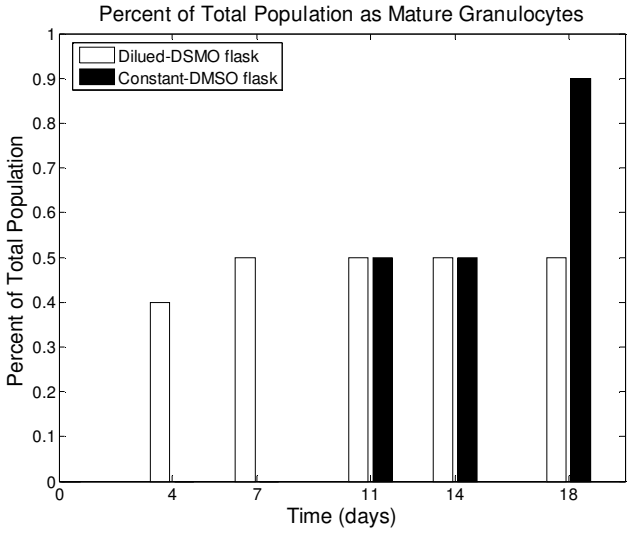

Fig. 3. Experimental time course for the differentiation of HL60 cells to granulocytes. Open bars indicate diluted-DMSO flask data and solid bars (starting at day 11) show the result with the modified experimental procedure keeping the DMSO level constant at the controller input. The trajectories do not achieve the desired level of 15% granulocytes. Model parameters were given in Fig. 2.

The results of the initial experimental trial indicated that it was necessary to maintain a constant DMSO concentration. However, adding a portion of fresh DMSO and media daily is not consistent with the procedure used to gather data for model identification, resulting in model-plant mismatch. In the model-identification experiment, DMSO and media were added to the flask once and allowed to naturally degrade, be taken up and consumed by cells over the course of several days. When a portion of the DMSO and media is replaced daily, this alters the anticipated differentiation dynamics. The work of [9] shows that exposing cells to 1.2% DMSO for 3 days results in a peak differentiation level of ~90%, which represents full-scale differentiation for this concentration [8]. Thus, it is expected that 3 days of exposure should provide sufficient time for any DMSO bolus to achieve its maximal level. Hence the HL60 cell population will be reduced to $6x10^5$ cells/mL only at sampling times to maintain the constant DMSO concentration for the necessary period of time while providing fresh media to maintain the nutrient level every 3 to 4 days. Passing the cells less frequently should not have an adverse effect on cell growth as it is common practice to culture cells by passing twice weekly. This will allow the controller-designed experiment to more closely mimic the model-identification experiment.

To further minimize any potential model-plant mismatch, the model parameters were also refined prior to the next iteration. The revised parameters were close to the original parameters. Changes can be found by comparing the values reported in the captions of Figs. 2 and 4.

Adjustments were also made to several MPC parameters given in Table I. The prediction horizon, H_p, was deemed to be too large because the controller did not aggressively adjust the inputs mid-experiment when the controlled output failed to reach the target. A new prediction horizon of 7 days was chosen to address this problem. The target percentage of granulocytes, s, was also adjusted. Due to error inherent in the experimental process and flow cytometry analysis, we decided that a granulocyte objective higher than 15% was necessary to be confident that the results were safely outside the bounds of measurement error. For this reason, we increased our target level to 25% granulocytes. A revised MPC parameter set is given in Table III.

TABLE III

Parameter	Description	Value
H_p	Prediction horizon	7 days
H_u	Control horizon	7 days
s	Target trajectory	25% of population granulocytes
Q	Weight for matching trajectory	10
R	Weight for minimizing control input	1

Revised parameter set for MPC constrained optimization problem of (2).

VIII. Results

An experimental evaluation of the MPC-derived strategy was performed using the revisions outlined in the previous section. The experimental time-course results are shown in Fig. 4. The control strategy was successful in achieving and sustaining the target percentage of granulocytes. The implemented control strategy that achieved these results is provided in the second column of Table IV. The controller implemented a "bang-bang" strategy to drive the granulocyte population to the desired level before applying a bolus at an intermediate concentration to sustain the target level.

It is interesting to note that the controller opted for no DMSO input on several days during the experimental evaluation despite the fact that the controlled output had not reached the target level (see days 3 and 10). This occurred due to the percentage of cells in x_2 of the model (cells beginning to differentiate, CD11b). On days when the controller opted for no input, the population of CD11b cells was significant (~50%-60%): these cells transitioned slowly into granulocytes during the next 4-7 days. Thus, even with no DMSO added, the granulocyte population rose at the next sampling point. Controlling the cell population given this observed time lag may be non-obvious to an experimenter using intuition alone to determine DMSO inputs. The predictive capabilities of the model allow for future state projections to be quantitatively factored into the current input decision. In this way, the controller designed an easily-implementable control strategy while taking into account quantitative projections not obvious to the experimenter.

To investigate the implications of a mismatch between the model and the plant on the MPC-derived sequence of DMSO boluses, a simulation of the MPC implementation without any model-plant mismatch was conducted. It found that the granulocyte population would reach and sustain the target level with minimal error (±2%) (results not shown) using the control strategy shown in the third column of Table IV. Without any model-plant mismatch, the necessary sustaining dosages of DMSO seem to alternate between two low levels administered at each 3 and 4 day sampling point.

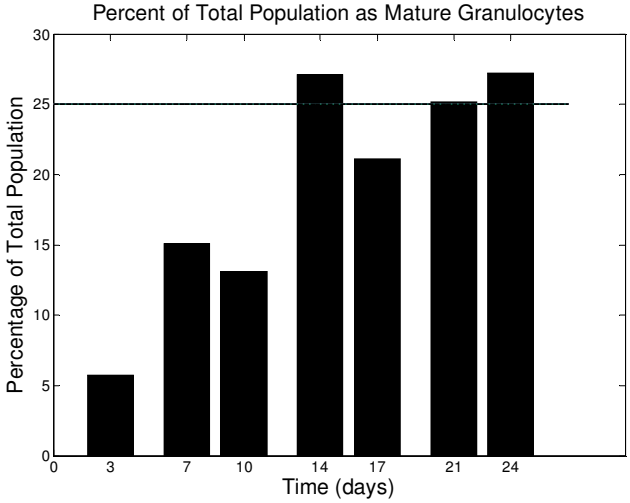

Fig. 4. Experimental time course for the differentiation of HL60 cells to granulocytes (bars) and target trajectory (line). The output achieves the desired level of 25% granulocytes within the first 14 days and sustains that level for the next 10 days. Revised model parameters: k_b = 0.302/day, k_d = 0.005/day, k_3 = 0.170/day, k_4 = 0.010/day, k_m = 1.378, k_i = 0.150, μ = 4.845 days, σ = 0.915 days.

TABLE IV

Time Point (days)	MPC-derived With Experimental Data* Control Input	MPC-derived With Simulated Data** Control Input
0	1.20%	1.2000%
3	0%	1.1350%
7	1.20%	0.0361%
10	0%	0.4591%
14	0%	0.5276%
17	0.40%	0.2767%
21	0%	0.5637%
24	0%	0.2686%

Each control input gives the application dose of DMSO (v/v).
*MPC-derived control strategy for the experimental study of Fig. 4.
**A simulation study assuming no plant-model mismatch.

IX. Summary and Future Work

This work presents a model for the differentiation of HL60 cells and implements a model predictive controller to achieve a desired level of granulocytes. After an initial experimental study that highlighted key changes to the MPC problem formulation and experimental protocol, the MPC-derived DMSO schedule and dosage was successful in achieving and maintaining the desired level of granulocytes in the laboratory. Additionally, the control sequence was non-intuitive and relied on the model's ability to predict responses in the calculation of each dosage. This demonstrates the need for a quantitative approach to design experiments, as relying on intuition alone may necessitate a series of expensive experiments.

This work represents a first step toward providing a more quantitative framework for systematically controlling cellular differentiation. Further improvements to the model to minimize model-plant mismatch are anticipated to improve the precision of the results. Additionally, extensions will be made to the single-input, single-output controller to design multivariable controllers capable of simultaneously directing HL60 differentiation into several specific cell lineages. The evolution of these techniques for quantitative experiment design will ultimately develop an approach to design bioreactors and experimental protocols for predictably manipulating stem cell differentiation.

References

[1] Bagheri, N., J. Stelling, and F.J. Doyle, "Circadian phase entrainment via nonlinear model predictive control." *Int. J. Robust Nonlinear Control*, 2007. 17(17): p. 1555-1571.

[2] Lebiedz, D., "Exploiting optimal control for target-oriented manipulation of (bio)chemical systems: A model-based approach to specific modification of self-organized dynamics." *International Journal of Modern Physics B*, 2005. 19(25): p. 3763-3798.

[3] Noble, S., et al., "Obtaining target levels of granulocyte differentiation through optimization techniques," presented at the American Control Conference, Seattle, WA, June 11–13, 2008, pp. 2571–2572.

[4] Viswanathan, S. and P.W. Zandstra, "Towards predictive models of stem cell fate." *Cytotechnology*, 2003. 41(2-3): p. 75-92.

[5] Fontes, F.A.C.C., M. L, and G. E, *Sampled-Data Model Predictive Control for Nonlinear Time-Varying Systems: Stability and Robustness in Assessment and Future Directions of Nonlinear Model Predictive Control*, R. Findeisen, F. Allgower, and L.T. Biegler, Editors. 2007, Springer Berlin p. 115-129.

[6] Birnie, G.D., "The HL60 cell line: a model system for studying human myeloid cell differentiation." *Br J Cancer Suppl*, 1988. 9: p. 41-5.

[7] Collins, S.J., "The HL-60 promyelocytic leukemia cell line: proliferation, differentiation, and cellular oncogene expression." *Blood*, 1987. 70(5): p. 1233-44.

[8] Collins, S.J. et al, "Terminal differentiation of human promyelocytic leukemia cells induced by dimethyl sulfoxide and other polar compounds." *Proc Natl Acad Sci*, 1978. 75(5): p. 2458-2462.

[9] Tarella, C. et al, "Induction of Differentiation of HL-60 Cells by Dimethyl Sulfoxide: Evidence for a Stochastic Model Not Linked to the Cell Division Cycle." *Cancer Research*, 1982. 42: p. 445-449.