MATHEMATICAL BIOSCIENCES
AND ENGINEERING
Volume **12**, Number **5**, October **2015**

doi:10.3934/mbe.2015.12.1127

pp. **1127–1139**

COMBINING ROBUST STATE ESTIMATION WITH NONLINEAR MODEL PREDICTIVE CONTROL TO REGULATE THE ACUTE INFLAMMATORY RESPONSE TO PATHOGEN

Gregory Zitelli

Department of Mathematics
University of California, Irvine
340 Rowland Hall, Bldg #400
Irvine, CA 92697-3875, USA

Seddik M. Djouadi

Electrical Engineering and Computer Science Department
Masdar Institute of Science and Technology
Masdar City, Abu Dhabi, UAE
On leave from
Department of Electrical Engineering and Computer Science
University of Tennessee
1520 Middle Dr
Min H. Kao Bldg Ste 401
Knoxville, TN, 37996-2250, USA

Judy D. Day

Department of Mathematics
University of Tennessee
1403 Circle Dr
Ayres Hall 227
Knoxville, TN, 37996-2250, USA

(Communicated by Yang Kuang)

Abstract. The inflammatory response aims to restore homeostasis by means of removing a biological stress, such as an invading bacterial pathogen. In cases of acute systemic inflammation, the possibility of collateral tissue damage arises, which leads to a necessary down-regulation of the response. A reduced ordinary differential equations (ODE) model of acute inflammation was presented and investigated in [10]. That system contains multiple positive and negative feedback loops and is a highly coupled and nonlinear ODE. The implementation of nonlinear model predictive control (NMPC) as a methodology for determining proper therapeutic intervention for *in silico* patients displaying complex inflammatory states was initially explored in [5]. Since direct measurements of the bacterial population and the magnitude of tissue damage/dysfunction are not readily available or biologically feasible, the need for robust state estimation was evident. In this present work, we present results

2010 *Mathematics Subject Classification.* Primary: 92C50, 93C15; Secondary: 93C83.

Key words and phrases. Inflammation, particle filter, immuno-modulation, nonlinear controllability, nonlinear observability.

This work has been supported by the National Science Foundation under Contract NSF-DMS 1122462.

on the nonlinear reachability of the underlying model, and then focus our attention on improving the predictability of the underlying model by coupling the NMPC with a particle filter. The results, though comparable to the initial exploratory study, show that robust state estimation of this highly nonlinear model can provide an alternative to prior updating strategies used when only partial access to the unmeasurable states of the system are available.

1. **Introduction.** The acute inflammatory response is a biological process of host immunity which can be initiated by the presence of a bacterial pathogen. The response can lead to harmful and potentially life threatening effects to the host as much as to the pathogen for which it is intended. The phenomenon of unresolved inflammation due to infection is observed in intensive care units (ICU) and typically referred to as sepsis or the systemic inflammatory response syndrome (SIRS). Commonly, it leads to multi organ-failure resulting in death, making it one of the leading causes of mortality in the ICU and a cost intensive disease to treat [1]. Natural regulatory mechanisms exist in the host response process for inhibiting inflammatory mediators in order to prevent such excessive inflammation and resulting damage from occurring. In particular, these anti-inflammatory regulators act by damping the functions of inflammatory mediators and their responses to other cells and inflammatory by-products produced during the response [6]. When the response becomes dysfunctional and the regulatory mechanisms are insufficient to effectively resolve the response, intervention is required. However, the complex nature of the response makes it difficult to determine effective intervention strategies.

An ordinary differential equations (ODE) model for the acute inflammatory response to a generic (gram-negative) pathogen was developed in [10]. This four dimensional model considers a pathogen population P, activated phagocytes (immune cells) N^*, a marker for tissue damage/dysfunction D, and anti-inflammatory mediators C_A. (See Appendix A.) For a set of biologically admissible parameter values, the model exhibits tri-stabilty (in the positive octant). Each of the three stable equilibria are biologically consistent with clinically observed states: a healthy equilibrium where $P = N^* = D = 0$ and C_A is at a background level, a septic death equilibrium where all mediators, N^*, C_A, and D, as well as pathogen, P, are significantly elevated, and an aseptic death equilibrium where $P = 0$ but N^*, C_A, and D are elevated far above background.

In [5], this reduced ODE model was used in the exploration of nonlinear model predictive control (NMPC) as a methodology to systematically control the inflammatory response and direct response trajectories toward the healthy equilibrium. Both pro- and anti-inflammatory therapies were considered. In other words, two inputs (controls) were used to contribute positive impulse doses to the state equations for N^* and C_A. Although a thorough dynamical systems analysis of the model exists in [10], in Section 2 we present new results toward an understanding of controllability in this model. Specifically, we look at the Kalman-type rank condition for nonlinear systems, and show that this system demonstrates local accessibility at all points within the positive octant. In Section 3, we describe the particle filter used to provide robust state estimation on the model for use within the NMPC algorithm. The results are then compared to results in which no estimates of the unmeasureable states are made as well as to the initial study results in [5], where an *ad hoc* routine was necessarily employed to address this lack of information about the system states.

2. **Controllability.** Previously in [2], properties of nonlinear system *observability* were explored for the model system in [10], which determine if system states are distinguishable with respect to the inputs. The results showed that the system was observable and this characteristic provides support for the use of nonlinear state estimation algorithms for the system, several of which were subsequently investigated in that work. However, the *controllability* properties for this system have not yet been discussed and in order to combine the nonlinear state estimation algorithms with a control scheme (in particular, model predictive control) as we wish to do here, it is important to understand these properties for the system in question.

Consider a general nonlinear affine system of the form

$$\dot{x} = F(x) + \sum_{k=1}^{m} g_k(x)u_k, \qquad x(0) = x_0 \in \mathbb{R}^n \tag{1}$$

$$y = h(x) \tag{2}$$

where $F : \mathbb{R}^n \to \mathbb{R}^n$, $h : \mathbb{R}^n \to \mathbb{R}^l$, and $g_k : \mathbb{R}^n \to \mathbb{R}^n$, for each $k = 1, \ldots, m$. The inputs $u_k : [0, \infty) \to \mathbb{R}$ for each $k = 1, \ldots, m$ are taken to be bounded, Borel-measurable, time dependent strategies. If the vector fields $F, g_1, \ldots, g_m$ are additionally taken to be continuous and there exists a constant $c > 0$ such that

$$|X(x_1)| \leq c\,(|x_1| + 1) \quad \text{and}$$
$$|X(x_1) - X(x_2)| \leq c|x_1 - x_2|, \qquad x_1, x_2 \in \mathbb{R}^n$$

for $X \in \{F, g_1, \ldots, g_m\}$, then it is shown in [13] that for each appropriate control u there exists a unique solution $x^{x_0,u}(t)$ of equation (1). The functions we consider herein have the above properties.

Given $x_0, x_1 \in \mathbb{R}^n$, some x_1 is said to be *accessible*[1] from x_0 with respect to the system if there exists some appropriate control u and finite time $T \geq 0$ such that $x^{x_0,u}(T) = x_1$. For linear control systems of the form $\dot{x} = Ax + Bu$, where $x \in \mathcal{R}^n, u \in \mathbb{R}^m$, accessibility between all pairs of points in the state space can be verified algebraically by what is known as the Kalman rank condition, specifically that

$$\operatorname{rank}\left(B|AB|A^2B|\ldots|A^{n-1}B\right) = n,$$

(where n is the dimension of x) and the system is then said to be controllable [9]. Unfortunately, for nonlinear systems, there is no analogous test to determine controllability of the system. However, the generalizations of these concepts do lead to a Kalman-like rank condition for nonlinear systems that provide some assurance that control may be possible. Although it appears the verification of whether or not a given nonlinear system can or cannot be controlled is usually left to practice, for completeness we explore this condition for the nonlinear system considered here.

Thus, we continue with the following framework. If $\mathcal{O}$ is an open set containing x_0, a point x_1 is said to be *accessible from* x_0 *through* $\mathcal{O}$ if, in addition, $x^{x_0,u}(t) \in \mathcal{O}$ for all $t \in [0, T]$. In keeping with one of a few standard conventions, we let $\mathbb{R}^{\mathcal{O}}(x_0, T)$ denote the set of all points accessible from x_0 in some time $T > 0$ through $\mathcal{O}$. The expression $\mathbb{R}^{\mathcal{O}}_T(x_0)$ indicates the set of all points accessible from x_0 through $\mathcal{O}$ in some time $\tau \in [0, T]$, so that $\mathbb{R}^{\mathcal{O}}_T(x_0) = \bigcup_{0 \leq \tau \leq T} \mathbb{R}^{\mathcal{O}}(x_0, \tau)$. If for every $T > 0$ and nonempty $\mathcal{O}$ containing x_0 the sets $\mathbb{R}^{\mathcal{O}}_T(x_0)$ have nonempty interior, then the system is said to be *locally accessible* at x_0. If the sets $\mathbb{R}^{\mathcal{O}}(x_0, T)$ have nonempty

[1] Alternatively, the terminology *reachable* is also used.

interior for sufficiently small values of T, then the system is said to be *strongly locally accessible* at x_0.

The properties of local accessibility and strong local accessibility at a point are linked directly to a Kalman-type rank condition on the distributions[2] associated with the vector fields governing the system (see [9], [13], [3], [4]). In the case of an affine system, these distributions are particularly straightforward to describe using the Lie algebra generated by $\{F, g_1, \ldots, g_m\}$.

Let $\mathcal{C}$ denote the Lie algebra generated by $\{F, g_1, \ldots, g_m\}$. Let $\mathcal{C}_0$ be the smallest Lie algebra containing $\{g_1, \ldots, g_m\}$ such that $[F, X] \in \mathcal{C}_0$ for any $X \in \mathcal{C}_0$, where we define the vector field $[F, X] : \mathbb{R}^n \to \mathbb{R}^n$ as the *Lie bracket* of (smooth) vector fields F and X, with $[F, X] = J_X F - J_F X$, where J_X denotes the Jacobian matrix of the vector field X. We write $C(x_0)$ and $C_0(x_0)$ for the linear span of $X(x_0)$, with X in $\mathcal{C}$ and $\mathcal{C}_0$, respectively.

Theorem 2.1. [9].*Consider the system in (1) where* $F : \mathbb{R}^n \to \mathbb{R}^n$.

- *If* $\dim C(x_0) = n$, *then the system is locally accessible at* x_0.
- *If* $\dim C_0(x_0) = n$, *then the system is strongly locally accessible at* x_0.

Now we consider a specific system referred to earlier that represents a model of the acute inflammatory response given below as a function $F : \mathbb{R}^4 \to \mathbb{R}^4$ and modified with control input terms, with $x = (P, N^*, D, C_A)$ (in keeping with the original notation of the model described in [10]) and with $g_1(y) = [0\ 1\ 0\ 0]^T$ and $g_2(y) = [0\ 0\ 0\ 1]^T$. As such, this system does indeed exhibit affine control.

$$F(P, N^*, D, C_A) + g_1(y)u_1 + g_2(y)u_2 = \begin{bmatrix} F_1(P, N^*, D, C_A) \\ F_2(P, N^*, D, C_A) + u_1(t) \\ F_3(P, N^*, D, C_A) \\ F_4(P, N^*, D, C_A) + u_2(t) \end{bmatrix}$$

$$= \begin{bmatrix} \frac{dP}{dt} \\ \frac{dN^*}{dt} + u_1(t) \\ \frac{dD}{dt} \\ \frac{dC_A}{dt} + u_2(t) \end{bmatrix},$$

where $\frac{dP}{dt}$, $\frac{dN^*}{dt}$, $\frac{dD}{dt}$, and $\frac{dC_A}{dt}$ are the differential equations from [10] reproduced for convenience in Appendix A. Since the controls g_1, g_2 are independent of the states, x, we can ignore elements of $\mathcal{C}_0$ which involve Lie brackets of the form $[g_i, g_j]$, since these terms would all produce zero values. We may therefore construct a strong accessibility matrix by considering a smaller collection of entries, such as

$$S = \begin{bmatrix} g_1 & g_2 & [F, g_1] & [F, g_2] & [F, [F, g_1]] & \ldots \end{bmatrix}$$

[2]Using the definition from [9], a distribution D on a manifold M is a map which assigns to each $p \in M$ a linear subspace $D(p)$ of the tangent space T_pM.

We can truncate the above matrix after the first four columns, since we are only concerned with the matrix having a rank equal to the dimension of $\mathbb{R}^4$. Furthermore, under standard assumptions on the parameters as suggested in [10], algebraic criteria can be used to show that the resulting truncated matrix given as $S_{(4)}$ below is nonsingular for $(P, N^*, D, C_A) > (0, 0, 0, 0)$. (See Appendix B.). The result is as follows:

$$S_{(4)} = \begin{bmatrix} 0 & 0 & -\frac{\partial F_1}{\partial N^*} & -\frac{\partial F_1}{\partial C_A} \\ 1 & 0 & -\frac{\partial F_2}{\partial N^*} & -\frac{\partial F_2}{\partial C_A} \\ 0 & 0 & -\frac{\partial F_3}{\partial N^*} & -\frac{\partial F_3}{\partial C_A} \\ 0 & 1 & -\frac{\partial F_4}{\partial N^*} & -\frac{\partial F_4}{\partial C_A} \end{bmatrix}.$$

Thus, this implies that the system maintains strong local accessibility throughout the first octant from the above theorem. If it were possible to further guarantee that the interiors of sets $\mathbb{R}^{\mathcal{O}}_T(x_0)$ contained x_0, then the system would be said to be *locally controllable* at x_0. Local controllability at all points in a compact region of $\mathbb{R}^n$ implies that control is possible between any two points in the region, and additionally that there exists some finite upper limit on the time that *any* such transport will take. Unlike the property of strong local accessibility for nonlinear systems, the topological demands of local controllability cannot be demonstrated using algebraic criteria on the Lie algebras $C(x_0)$ and $C_0(x_0)$ in the absence of additional symmetries to exploit (i.e. symmetric controllers as in [13]). As a result, the local controllability of nonlinear systems is usually left to practice to verify. The possibility of successful control is encouraged by the necessary condition of strong local accessibility (see [4]) and was, in fact, verified in practice in [5] when enough information about the states was acquired. The positive observability conditions demonstrated in [2] further suggest that the use of a robust state estimator will allow good control with access to less state information than that used in [5].

3. **Particle filter.** Particle filters have become a popular tool for performing robust state estimation on nonlinear systems, where Kalman-type filters often prove inadequate. The utility of a particle filter in a model predictive control study of a deterministic ODE model of the inflammatory response to endotoxin was nicely displayed in [8]; however, the model used in that study has significantly less complex dynamics than the inflammatory model used herein due to the use of an exponentially decaying stimulus (endotoxin) compared to a dynamically changing pathogen population that is coupled to other system states. This difference poses a particular challenge in that the goal to minimize damage must now also be combined with the goal to minimize pathogen load; yet, these goals are at odds with one another since the remedy to accomplish one will not remedy the other and might further exacerbate the situation. For example, increasing inflammation to minimize pathogen will also cause more damage; whereas, increasing anti-inflammatory levels to minimize damage might make conditions more favorable for pathogen growth.

It became clear through the work in [5] that having a good estimation of the system states, especially the pathogen state, was critical to meet this challenge effectively. In that initial work, robust state estimation was not employed and instead

a procedure was added to the algorithm that indirectly updated the pathogen state in the model to better match that of the virtual patient's pathogen state. Although the procedure was biologically justified and resulted in more effective control outcomes over a diverse virtual patient population, the question remained as to whether such an *ad hoc* procedure could be replaced by a more systematic technique that did not rely on *faux* measurements or physical observation. Thus, while extending this work, we explored various filtering techniques for estimating the states of this nonlinear system (in the absence of the control problem) and determined that the particle filter proved the most promising choice [2].

3.1. **Setup.** The particle filter is a sequential Monte-Carlo scheme, sometimes referred to as a condensation algorithm. In the current implementation, a batch of particles are generated and have artificial control applied to them at each discrete time step (one hour). The endpoint of their simulated trajectories are compared to the observed patient states, with particles whose values match the data more closely propagating to further time steps more successfully. The design and implementation of our particle filter is based on [11] and [12].

Our model is characterized as a nonlinear ODE system with affine control. We take $x = (P, N^*, D, C_A)$ and define the function $F : \mathbb{R}^4 \to \mathbb{R}^4$ to be the system described in the previous section. (See also Appendix A.) Measurements of the observable variables N^* and C_A introduce noise, and so, we have rewritten the original ODE as a stochastic differential equation (SDE) for the purpose of filtering.

$$dX_t = \left(F(X_t) + \sum_{k=1}^{2} g_k(X_t)u_k \right) dt$$
$$Y_t = \int_0^t h(X_s)ds + W_t$$

where W_t represents a 2-dimensional Brownian motion, used to indicate measurement noise, and the expression

$$h(X_s) = \begin{bmatrix} F_2(X_s) + g_1(X_s)u_1 \\ F_4(X_s) + g_2(X_s)u_2 \end{bmatrix}$$

isolates the observable second and fourth components of the state. The σ-algebras $\mathcal{G}_t = \sigma(Y_0, Y_1, \ldots, Y_t)$ are understood to represent the data collected from measurements Y_t for discrete time steps $t \in \mathbb{N}$. We let π_t be the regular conditional probability distribution of X_t given $\mathcal{G}_t$, which is such that for any Borel $B \in \mathbb{R}^4$, $\pi_t(A) = \mathbb{P}(X_t \in B|\mathcal{G}_t)$. Consequently, integration of the identity function over $\mathbb{R}^4$ with respect to the measure π_t will provide the best estimate for X_t. As in [12], the particle filter aims to approximate π_t by Dirac-masses at the locations of each particle,

$$\pi_t \approx \frac{1}{N} \sum_{i=1}^{N} \delta_{X_t^i}$$

which indicates that our estimate $\widehat{X}_t$ for X_t is recovered as the mean in each component of the collection of particles X_t^i.

3.2. **Design.** The particle filter is initiated with $N = 500$ particles for all of our simulations. We create the particles with values of N^* and C_A matching the initial measurement Y_0 (with noise), while values of P and D are initialized to 0.01. The

choice of particle initialization values for P and D is discussed in section 3.3. Implementation. In keeping with the original application of NMPC in [5], we run simulations in time steps of one hour increments for one week so that $t = 1, 2, \ldots, 168$. At each hour t, the following algorithm refines the particles X_{t-1}^i into the new particles X_t^i.

- For each $i \in \{1, \ldots, N\}$, we generate temporary particles $X_t^{-,i}$ according to a simulation of the former particles X_{t-1}^i. This is done using the trajectory $x_i(t)$ with initial point $x_i(t-1) = X_{t-1}^i$ and desired plant control u:

$$X_t^{-,i} = X_{t-1}^i + \int_{t-1}^{t} \left(F(x_i(s)) + \sum_{k=1}^{2} g_k(x_i(s)) u_k \right) ds$$

- Calculate weights q_i for particles $X_t^{-,i}$, under the assumption that measurements Y_t are taken with Gaussian noise having covariance $R = 10^{-2}$:

$$q_i = \frac{1}{2\pi\sqrt{R}} \exp\left(- \left(Y_t^{-,i} - Y_t \right)^T \left(Y_t^{-,i} - Y_t \right) / 2R \right)$$

  where $Y_t^{-,i}$ is simply an array of the second and fourth coordinates of $X_t^{-,i}$, corresponding to those coordinates in Y_t, the measurable states.
- Normalize weights:

$$\tilde{q_i} = \frac{q_i}{\sum_{j=1}^{N} q_j}$$

- Resample the temporary particles $X_t^{-,i}$ according to weights $\tilde{q_i}$. In other words, the value $X_t^{-,j}$ is assigned to new particle X_t^i with probability $\tilde{q_j}$.
- Generate the X_t estimate as the mean $\widehat{X}_t = \frac{1}{N} \sum_{i=1}^{N} X_t^i$. The value of $\widehat{X}_t$ is used in the predictive model of the NMPC to generate the control u for the future time period $[t, t+1)$.

3.3. **Implementation.** In [5], a mechanism to update the predictive model with respect to the pathogen level was implemented to deal with instances of gross mismatch between pathogen levels in the predictive model and those in the patient model. They implemented this with the following routine: every four hours the pathogen level in the virtual patient is directly compared to that in the predictive model. If the pathogen level is significantly lower in the predictive model than in the patient model, then the predictive model's pathogen level is set to $P = 0.5$ on the next time step. Alternatively, if the predictive model's pathogen level is significantly higher than in the patient model, then the predictive model's pathogen level is set to $P = 0$. Such a process was justified by the clinical manifestations of infectious markers such as high fever or blood cultures and does not use the actual measurement value directly. However, we wished to improve on this implementation so as not to rely on this form of updating for one of the non-measurable states.

The removal of the pathogen update routine described above is a critical change to the current implementation, since without it the controller performance is extremely poor when considering a particularly aggressive pathogen. (See section 4. Results & Conclusions.) Thus, in this current work, the pathogen update routine is eliminated and instead, the state estimates found by the particle filter are used. The particle filter estimates for N^* and C_A are updated to reflect the measurements of the patient system (with measurement noise); however, we stress that the

estimates for both P and D (the unmeasurable states) are now solely reliant on the performance of the particle filter.

With the exception of this major change in how state estimation is handled and the introduction of noise to the measurements taken, the original Matlab implementation from [5] is maintained, where the algorithm used is a modified version of that developed by Florian et al., [7]. All simulations were conducted using MatLab®. The ODE system is numerically integrated using MatLab's Simulink default solver routine, *ode45*. The optimization routine uses the *fmincon* built-in MatLab® function. Equation (3) gives the objective function, J, we seek to minimize which uses the standard weighted sum of squares form. It contains terms to minimize pathogen levels (P), damage levels (D), and the total amount of control inputs, $u_1(t)$ and $u_2(t)$ given over the prediction horizon. All have reference trajectories of zero and the weights are $\Gamma_D = \Gamma_P = \Gamma_{u_1} = \Gamma_{u_2} = 1$.

$$J = \min_{\substack{u_1(t) \\ u_2(t)}} ||\Gamma_D D||_2^2 + ||\Gamma_P P||_2^2 + ||\Gamma_{u_1} u_1(t)||_2^2 + ||\Gamma_{u_2} u_2(t)||_2^2 \tag{3}$$

The same cohort of 1,000 virtual patients that were generated in the previous study are used here as well. Each virtual patient is represented by a copy of the four equation ODE model that is used as the predictive model but with the values of six parameters and two nonzero initial conditions (P and C_A) randomly selected from feasibly bounded intervals. (See Appendix C.) Thus, a selection of parameters have mismatched values among the patient models and the predictive model, with the remaining parameters set to their baseline values determined in [10]. The remaining initial conditions for each virtual patient for the pro-inflammatory mediator (N^*) and tissue damage/dyfunction (D) variable are set to zero. From these patient-specific initial conditions, virtual patients are simulated without any form of control or state estimation. If a virtual patient's level of N^* rises above 0.05 (a somewhat elevated inflammatory level for this model), then they are selected for treatment. The time at which this occurs represents the zero hour of the NMPC process, which then continues for one week (168 hours). Of the initial 1,000 virtual patients, 620 were selected to receive treatment, while the remaining did not display inflammatory responses large enough to warrant treatment.

At this "time zero," the particle filter is initialized for each virtual patient simulation in the following way: the initial measurements of the pro-inflammatory mediator (N^*) and the anti-inflammatory mediator (C_A) of that virtual patient (including measurement noise) are used to initialize the respective particles for N^* and C_A and the particles for the remaining variables (P and D) are randomized in feasible intervals. As was also done in [5], the predictive model is first initialized with initial condition $(P_0, N_0^*, D_0, C_A^0) = (0.5, 0, 0, 0.125)$ and is then integrated forward in time until $N_0^* > 0.05$, the same threshold used for the selection of virtual patients to receive treatment. The values of the system variables at the time at which the N_0^* threshold is crossed is then used as the initial condition of the predictive model at the zero hour of the NMPC process as described above. Thus, the first control move is based on the predictions of the model starting at this initial condition to be consistent with [5].

The initial predictions of the particle filter for the pathogen variable are quite important since P is not in a positive feedback loop with the measured variables N^* and C_A in the same manner that D is. This means that the information, especially received from the measurement of N^*, is not as informative for determining P as it is

for determining D. The initial guess for the P state, based on the fact that N^* values have passed the threshold value of 0.05, is that P is a relatively high value. This assumption will most likely cause the first determined control dose implemented in the first hour that treatment is initiated to be pro-inflammatory therapy with no anti-inflammatory therapy. However, after this dose is implemented, it is not safe to overestimate where P is. Thus, P is specifically initialized around a small value of 0.01 in the particle filter, which will be used to determine all subsequent state estimates for P and D. Later, in section 4. Results & Conclusions, we discuss this issue further.

In [5], it was noticed that the pathogen growth rate parameter (k_{pg}) affects the performance of the NMPC, and therefore, simulation sets were run for several k_{pg} values: 0.52, 0.6, and 0.8. Here we do the same for comparison. At each time step, the predictive model is used by the NMPC algorithm to determine a proper control (therapy) input for the next time step. This is implemented in the virtual patient model and in the predictive model. New (noisy) measurements are then taken to update the predictive model as well as the particle filter. The particle filter subsequently provides the predictive model with new state estimates for P and D and the next iteration begins.

4. **Results & Conclusions.** Results for three NMPC simulation groups are presented in Table 1. Each simulation group considers measurement noise for the measurable states, N^* and C_A, and also gives results for three k_{pg} values: 0.52, 0.6, and 0.8, as discussed above. Simulation Group (i) considers simulations like that of [5] but without using any pathogen updating routine nor robust state estimation to update the unmeasurable states, in particular, pathogen (P). Simulation Group (ii) gives results of simulations as performed in [5]. We note that there are slight differences between the percentages seen in Table 1 and those seen in [5] which are due to the now present addition of noise in the measurements, which was not considered in the prior work. Lastly, Group (iii) presents the simulation results in which a particle filter was solely used for state estimation; in particular, for pathogen (P) and damage (D) without any further updating. Included among these initial results is a placebo (no therapy) case, in which the NMPC is effectively turned off.

The results for the $k_{pg} = 0.52$ and $k_{pg} = 0.6$ cases on the whole appear quite comparable across all three simulation groups. However, in the case of an aggressively growing pathogen ($k_{pg} = 0.8$), there is a dramatic increase in the number of patients rescued in Groups (ii) and (iii) compared to Group (i) which does not update the pathogen state. See the first row of entries in Table 1 for the $k_{pg} = 0.8$ columns. Although Groups (ii) and (iii) show comparable results in terms of number of patients rescued in the $k_{pg} = 0.8$ category, there is a dramatic decrease in the number of patients harmed in Group (iii) which uses the particle filter (2 harmed) versus Group (ii) which uses the *ad hoc* pathogen update routine (16 harmed). In other words, 14 less virtual patients who would have been healthy without treatment (i.e. had a healthy placebo outcome) were pushed to an unhealthy state by the administration of therapy when the *ad hoc* updating routine was used to update the pathogen state instead of the particle filter. See the second to last row of Table 1 for the $k_{pg} = 0.8$ columns.

Hence, the ability to estimate the crucial but difficult-to-measure pathogen state is essential to the successful performance of the NMPC, especially when the pathogen state plays a strong role in driving the system dynamics like in the case when

Therapy Type:	**Placebo**	**(i) Results with no updating of the pathogen state**			**(ii) Results with ad hoc pathogen update mechanism (as in Day, 2010)**			**(iii) Results with particle filter state estimation**		
kpg:	***n/a***	***0.52***	***0.6***	***0.8***	***0.52***	***0.6***	***0.8***	***0.52***	***0.6***	***0.8***
Healthy:	251	366	496	50	369	510	513	368	456	459
Aseptic:	228	123	62	570	120	49	107	121	77	74
Septic:	141	131	62	0	131	61	0	131	87	87
Harmed (out of 251):	n/a	0	1	243	0	2	16	0	2	2
Rescued (out of 369):	n/a	115	246	42	118	261	278	117	207	210

FIGURE 1. Patient outcome results (number of patients out of 620) from three different NMPC simulation groups, Column (i) - (iii), and for the case in which no therapy is given, Placebo. All simulations included measurement noise for the two measurable states, N^* and C_A. Using the convention in [5], each simulation grouping shows results over three different k_{pg} values for the underlying predictive model: 0.52, 0.6, and 0.8. Column (i): Results using no updating strategy for the unmeasurable states, pathogen (P) and damage (D). Column (ii): Results based on the *ad hoc* pathogen updating mechanism from [5]. Column (iii): Results using a particle filter to perform state estimation for all the system states, in particular D and P, in place of the mechanism used in the simulations for Column (ii).

$k_{pg} = 0.8$. The particle filter strategy provides the means to gain access to this state from the measurements of other states. Through additional simulation studies in which no pathogen updating/estimation strategy was used, it was seen that, for this particular system, an apparently effective dosing strategy is a strong initial pro-inflammatory boost (achieved by setting the initial value of P in the predictive model to a higher value than used in the simulations described above). However, this can come at the cost of harming more patients compared to the results in which the initial pro-inflammatory input was less, but where some form of updating/estimation strategy is used (results not shown).

Although we demonstrated that the model exhibits a strong local accessibility property, this is not the same as exhibiting local controllability. Toy models in the literature can easily demonstrate that accessibility may very well allow control along certain directions and not along others. Furthermore, controllability is often framed using general controls meeting certain integrability conditions, such as L^1_{loc}, which are often unrealistic. Under the current setup, we assumed impulse-like controls to simulate instantaneous injections, each separated by whole hour increments. Additionally, biological demands also insisted on nonnegative control inputs and certain constraints to guard against excessive dosing, as described in [5].

Furthermore, the objective function formulated for the NMPC algorithm (equation (3)) considers both minimizing pathogen as well as damage as a part of the goal; yet, the chosen therapies have opposite effects on each of these, creating a tug-of-war effect: implementing the inflammatory therapy, while helping to minimize pathogen, works against the goal to minimize damage since it is a direct cause of damage; likewise, implementing the anti-inflammatory therapy, while helping to minimize damage, may work contrary to the goal of minimizing pathogen since it suppresses the response against it. This work again highlights the importance of obtaining accurate state estimation early on in the process to ensure more effective control and unfortunately speaks to the complexity of treating patients under many constraints. Consequently, it eludes to the fact that sometimes success might just not be possible under certain given circumstances.

However, the particle filter is an effective state estimator for this highly nonlinear system, since we see in our simulations that convergence occurs for every state, even for the unmeasurable ones. Time to convergence can vary, however, and this needs to be taken into consideration regarding the effects this may have on the performance of the NMPC to find successful therapeutic control strategies. Overall, though, we demonstrate that the particle filter is capable of estimating pathogen levels well enough to eliminate the need for the type of *ad hoc* pathogen update routine that was previously critical for course corrections related to mismatch in this variable in [5], as discussed in section 3.3. Implementation. This is a key insight, in that the ability to have access to this state through the available measurements of other variables defers the need to get a direct measurement. Thus, in general, a combination of dynamical modeling of a process, incorporation of data measurements, and state estimation techniques offers a possible approach for inferring information about other clinically relevant states that are difficult or impossible to measure but could provide critical feedback for therapeutic decision making.

REFERENCES

[1] D. C. Angus and T. van der Poll, Severe sepsis and septic shock, *New Eng J Med*, **369** (2013), 840–851.

[2] O. Bara, J. Day and S. Djouadi, Nonlinear state estimation for complex immune responses, *Proceedings of the* 52^{nd} *IEEE Conference on Decision and Control, Florence, Italy*, **December 10–13** (2013), 3373–3378.

[3] G. Conte, C. H. Moog and A. M. Perdon, *Nonlinear Control Systems: An Algebraic Setting*, Springer-Verlag London, Ltd., London, 1999.

[4] J. M. Coron, *Control and Nonlinearity*, American Mathematical Society, Providence, RI, 2007.

[5] J. Day, J. Rubin and G. Clermont, Using nonlinear model predictive control to find optimal therapeutic strategies to modulate inflammation, *Math Biosci Eng*, **7** (2010), 739–763.

[6] M. de Waal, J. Abrams, C. Bennett, B. Figdor and J. de Vries, Interleukin 10(il-10) inhibits cytokine synthesis by human monocytes: An autoregulatory role of il-10 produced by monocytes, *J Exp Med*, **174** (1991), 1209–1220.

[7] J. A. Florian Jr., J. L. Eiseman and R. S. Parker, Nonlinear model predictive control for dosing daily anticancer agents using a novel saturating-rate cell-cycle model, *Comput. Biol. Med.*, **38** (2008), 339–347.

[8] J. Hogg, G. Clermont and R. S. Parker, Acute inflammation treatment via particle filter state estimation and mpc, *9th International Symposium on Dynamics and Control of Process Systems*, **9** (2010), 272–277.

[9] H. Nijmeijer and A. van der Schaft, *Nonlinear Dynamical Control Systems*, Springer-Verlag, New York, 1990.

[10] A. Reynolds, J. Rubin, G. Clermont, J. Day, Y. Vodovotz and G. B. Ermentrout, A reduced mathematical model of the acute inflammatory response. i. derivation of model and analysis of anti-inflammation, *J Theor Bio*, **242** (2006), 220–236.
[11] D. Simon, *Optimal State Estimation: Kalman, H-infinity and Nonlinear Approaches*, Wiley-Interscience, Hoboken, NJ, 2006.
[12] J. Xiong, *An Introduction to Stochastic Filtering Theory*, Oxford University Press, Oxford, 2008.
[13] J. Zabczyk, *Mathematical Control Theory: An Introduction*, Birkhäuser Boston, Inc, Boston, MA, 1992.

Appendix A - Acute inflammatory model. The underlying patient and predictive model, as presented in [10].

$$\frac{dP}{dt} = k_{pg}P\left(1 - \frac{P}{P_\infty}\right) - \frac{k_{pm}s_mP}{\mu_m + k_{mp}P} - k_{pn}f(N^*)P$$

$$\frac{dN^*}{dt} = \frac{s_{nr}R(P,N^*,D)}{\mu_{nr} + R(P,N^*,D)} - \mu_nN^*$$

$$\frac{dD}{dt} = k_{dn}\frac{f(N^*)^6}{x_{dn}^6 + f(N^*)^6} - \mu_dD$$

$$\frac{dC_A}{dt} = s_c + k_{cn}\frac{f(N^* + k_{cnd}D)}{1 + f(N^* + k_{cnd}D)} - \mu_cC_A$$

and where

$$f(x) = \frac{x}{1 + \left(\frac{C_A}{c_\infty}\right)^2}, \text{ and}$$

$$R(P,N^*,D) = f(k_{np}P + k_{nn}N^* + k_{nd}D).$$

Detailed explanations of each parameter, the ranges for their values, and the sources upon which the values were estimated are published elsewhere (see Table 1 in [10]).

Appendix B - Kalman-type rank condition. As a differential equation in $\mathbb{R}^4$ with affine control vectors $g_1(y) = [0\ 1\ 0\ 0]^T$ and $g_2(y) = [0\ 0\ 0\ 1]^T$, we recall the form of a general Lie bracket as

$$[f_1, f_2] = \frac{\partial f_2}{\partial x}\cdot f_1 - \frac{\partial f_1}{\partial x}\cdot f_2.$$

We are concerned, as in the Nijmeijer and van der Schaft Theorem [9] provided in the text, with showing that the Lie algebra $C_0(x_0) = \mathbb{R}^4$ for all x_0 in the positive octant. Since $C_0(x_0)$ is spanned by the control vectors and iterated Lie brackets with the vector field $F = \left[\frac{dP}{dt}\ \frac{dN^*}{dt}\ \frac{dD}{dt}\ \frac{dC_A}{dt}\right]^T$, we can translate the condition $C_0(x_0) = \mathbb{R}^4$ into a condition on the rank of sub-matrices of what we now refer to as the strong accessibility matrix.

$$S = \begin{bmatrix} g_1 & g_2 & [F,g_1] & [F,g_2] & [F[F,g_1]] & \dots \end{bmatrix}$$

Showing that for any point in the positive octant, we can find a 4×4 sub-matrix of S that achieves full rank when evaluated at the point will satisfy the conditions of the theorem. Truncating the matrix to the first four terms, we have

$$S_{(4)} = \begin{bmatrix} 0 & 0 & -\frac{\partial F_1}{\partial N^*} & -\frac{\partial F_1}{\partial C_A} \\ 1 & 0 & -\frac{\partial F_2}{\partial N^*} & -\frac{\partial F_2}{\partial C_A} \\ 0 & 0 & -\frac{\partial F_3}{\partial N^*} & -\frac{\partial F_3}{\partial C_A} \\ 0 & 1 & -\frac{\partial F_4}{\partial N^*} & -\frac{\partial F_4}{\partial C_A} \end{bmatrix}$$

$$= \begin{bmatrix} 0 & 0 & k_{pn} f(1) P & 0 \\ 1 & 0 & -\frac{\mu_{nr} s_{nr} k_{nn} f(1)}{(\mu_{nr} + R(P, N^*, D))^2} + \mu_n & -\frac{\mu_{nr} s_{nr} k_{nd} f(1)}{(\mu_{nr} + R(P, N^*, D))^2} \\ 0 & 0 & -\frac{6 x_{dn}^6 k_{dn} f(1) f(N^*)^5}{\left(x_{dn}^6 + f(N^*)^6\right)^2} & \mu_d \\ 0 & 1 & -\frac{k_{cn} f(1)}{(1 + f(N^* + k_{cnd} D))^2} & -\frac{k_{cn} k_{cnd} f(1)}{(1 + f(N^* + k_{cnd} D))^2} \end{bmatrix}$$

Using Maple$^{\text{TM}}$ (© Maplesoft), under standard parameters given by [10] this matrix was confirmed to be nonsingular for all $(P, N^*, D, C_A) > (0, 0, 0, 0)$.

Appendix C - Virtual patient parameters. The following table provides the parameters which have differing values between virtual patients and the predictive model, chosen within the ranges given.

Parameter	Patient Parameter Range
P_0	0.0 - 1.0
$C_{A,0}$	0.0938 - 0.1563
k_{pg}	0.3 - 0.6
k_{cn}	0.03 - 0.05
k_{nd}	0.015 - 0.025
k_{np}	0.075 - 0.125 (Co-varies with k_{nd})
k_{cnd}	36.0 - 60.0 (Co-varies with k_{nd})
k_{nn}	0.0075 - 0.0125 (Co-varies with k_{nd})

Received September 09, 2014; Accepted March 23, 2015.

E-mail address: judyday@utk.edu
E-mail address: djouadi@eecs.utk.edu
E-mail address: gzitelli@uci.edu