www.ietdl.org

Published in IET Control Theory and Applications
Received on 7th August 2011
Revised on 19th January 2012
doi: 10.1049/iet-cta.2011.0475

ISSN 1751-8644

Real-time control of human coagulation

J.G. Makin[1] *S. Narayanan*[2]

[1]*EECS Department, University of California, Berkeley, USA*
[2]*International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704, USA*
E-mail: makin@icsi.berkeley.edu

Abstract: In previous work, the authors showed that the dynamics of human blood clotting could be fruitfully modelled and simulated as a hybrid system (HS), that is, one with interacting continuous and discrete parts. Here, the authors show that, although a complete analysis of the HS is (computationally) infeasible, analysis and control techniques can indeed be applied to a large, critical subsystem a set of about 100 ordinary differential equations. The theory is outlined behind the control techniques and then demonstrate in a series of simulations their application to control of pathological blood clotting, both hypercoagulatory (factor-V Leiden) and hypocoagulatory (hæmophilia A). In particular, steering is simulated during a clotting event of the crucial blood–protein thrombin, via the controlled injection of (recombinant) factor VIII (for hæmophilia) or the anti-coagulant heparin (for FV Leiden). It remains to remedy the shortcomings of this control technique, and to extend it to the remainder of the HS of the previous work; methods for these are proposed, and addressed in a subsequent article.

1 Introduction

The complexity of the coagulation process renders treatment of its disorders difficult. Indeed, current treatments make no attempt to manipulate coagulation in real time; rather, periodic (of the order of 1 day to 1 week) measurements and interventions are made with the aim of keeping certain known risk factors within safe ranges. This limitation is a consequence of two facts: first, the state of current sensor technology, which precludes real-time measurements of blood proteins (and hence real-time feedback control); and second, the lack of theoretical techniques for such control. This paper addresses the second of these limitations by applying techniques from computer science and from mathematical control theory and demonstrating their validity – and limitations – both mathematically and in a series of simulations. By providing part of the solution to the theoretical problem of controlling blood clotting, we hope to provide an impetus for the development of the relevant sensor technologies.

The simulations of our previous work [1] demonstrate the utility of the hybrid-system (HS) model described therein for certain purposes, namely prediction, theoretical investigation and sensitivity analysis. However, ultimately we should like to use it as a basis for the real-time control of blood-clotting. Now, the standard technique for control of HSs is numerical solution of a set of partial differential equations – the Hamilton–Jacobi equations. For systems of dimension greater than about five, however, solution of the equations is computationally infeasible (a consequence of the curse of dimensionality) [2]. Unfortunately, our system has upwards of 100 state variables; so this technique is patently unworkable.

The approach taken here is to consider only a subsystem of the model, a set of non-linear ordinary differential equations (ODEs) that were originally drawn largely from [3, 4]. In fact, the model can be decomposed into this (purely continuous) subsystem and two HSs, the three interacting only through two state variables, as we shall see below (Section 3.1), so that control of clot formation in the complete model can be approached as a problem of continuous (non-linear) control. (A more elaborate justification for this decomposition can be found in Chapter 4 of [5].) Thus in the present study, we content ourselves with deriving and simulating the control techniques for the system of ODEs alone – which ODEs moreover are, in consequence of the mass-action kinetics they model, merely second-order polynomial functions of the state. Specifically, the control task is to steer the concentration of thrombin (activated blood factor II) along a desired trajectory during a clotting event, by controlled rate of injection of an exogenous pharmaceutical (e.g. heparin). We demonstrate two different techniques to effect this control: the first, following [6], using non-linear feedback and a change of variables to partially linearise the ODEs, and then controlling the linearised, single-input/single-output (SISO) system using standard methods; and the second based on step-input control. In the simulations, the input is either the anti-coagulant heparin or (recombinant) factor VIII, and the output is thrombin, the most important enzyme product of the coagulation cascade (see below).

We present the theory of the control techniques first, along with some preparatory results on alternative approaches, before presenting a series of simulations. We then simulate the clotting process in a patient with the pro-coagulatory disorder factor-V Leiden and in a patient with moderate

hæmophilia A, and then repeat the simulation under application of the control techniques. Finally, we discuss both the relevance of this control task to the overall task of controlling blood clotting, as well as implementation issues, and then propose an alternative technique to remedy some of the defects of the present approach.

2 Theory

The model to be controlled is a system of about 100 coupled, non-linear ODEs, of the most general form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, u) \tag{1}$$

where u is the (single) control variable (say, an anti- or pro-coagulant). The control task is to force one of the state variables (thrombin) to track a desired trajectory. In fact, the ODEs can be expressed, as we shall see, in a less general form; and the approach described in this paper will be to exploit some of the peculiarities of the system that distinguish it from the most general case.

Now, control of linear systems is comparatively easy, so a standard approach is to design the controller around a linear approximation to the true system, found by considering only the first term of the Taylor expansion, the Jacobian $[\partial \mathbf{f}/\partial \mathbf{x}](\mathbf{x}_0)$, near an equilibrium point $\mathbf{x}_0$. Unfortunately, the Jacobian in our model is singular, so (by the Hartman–Grobman theorem) the linearisation is not guaranteed to approximate the true system.

Alternatively, the system may be exactly linearised (as opposed to merely approximated by a linear system) by choosing the appropriate non-linear feedback $u = \psi(h(\mathbf{x}))$ and looking at the system through a change of variables $\boldsymbol{\xi} = \Phi(\mathbf{x})$

$$\begin{aligned} \dot{\boldsymbol{\xi}} &= A\boldsymbol{\xi} + \mathbf{b}v \\ y &= \xi_1 \end{aligned} \tag{2}$$

where in fact $(A, \mathbf{b})$ are in controllable canonical form, so the system is completely controllable. Here, $h(\mathbf{x})$ is an output function that reflects our observations of the state variables, and v is a synthetic input related to the true input u by a known function. (This discussion of feedback linearisation follows [6], but see also [7].) From this perspective, we can ask whether or not, given the dynamics in (1), there exists an output function rich enough to support the linearisation. The answer is particularly straightforward if, as in our case, the ODEs are affine in the control, that is can be written as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})u \tag{3}$$

in which case necessary and sufficient conditions for the linearisation can be given in the form of conditions on the vector fields $\mathbf{f}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$ and iterated Lie brackets thereof. Specifically, the matrix of vector fields

$$[\mathbf{g}(\mathbf{x}), ad_{\mathbf{f}}\mathbf{g}(\mathbf{x}), \ldots, ad_{\mathbf{f}}^{n-2}\mathbf{g}(\mathbf{x}), ad_{\mathbf{f}}^{n-1}\mathbf{g}(\mathbf{x})] \tag{4}$$

often referred to as the strong-accessibility distribution, must have full rank (n) in the region of interest; and the set of vector fields

$$\mathbf{g}(\mathbf{x}), ad_{\mathbf{f}}\mathbf{g}(\mathbf{x}), \ldots, ad_{\mathbf{f}}^{n-3}\mathbf{g}(\mathbf{x}), ad_{\mathbf{f}}^{n-2}\mathbf{g}(\mathbf{x}) \tag{5}$$

must be involutive in the region. (A set of vector fields is involutive if the Lie bracket of any two of those vector fields is within the span of the original set, where the Lie bracket $[\mathbf{f}, \mathbf{g}]$ of two vector fields is defined as $[\partial \mathbf{g}/\partial \mathbf{x}]\mathbf{f} - [\partial \mathbf{f}/\partial \mathbf{x}]\mathbf{g}$.) Here, by recursive definition

$$\begin{aligned} ad_{\mathbf{f}}^k\mathbf{g} &:= [\mathbf{f}, ad_{\mathbf{f}}^{k-1}\mathbf{g}], \quad k > 0 \\ ad_{\mathbf{f}}^k\mathbf{g} &:= \mathbf{g}, \quad k = 0 \end{aligned}$$

If these conditions are not met and full-state linearisation is not possible, one may attempt to linearise the system partially. Here, the linearisation is carried out with respect to some given output function, which will admit the formulation of some number q of new state variables with linear dynamics. Of course, if full-state linearisation is not possible, then q, called the relative degree of the affine-control system, is strictly less than n, the dimension of the state space.

Consider again the differential equations of our system as the affine-control system of (3), where this time $h(\mathbf{x}) = y$, the variable we wish to control. Differentiating the output with respect to time yields

$$\begin{aligned} \dot{y} &= (\nabla_{\mathbf{x}}h)\dot{\mathbf{x}} \\ &= (\nabla_{\mathbf{x}}h)\mathbf{f}(\mathbf{x}) + (\nabla_{\mathbf{x}}h)\mathbf{g}(\mathbf{x})u \\ &= L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})u \end{aligned} \tag{6}$$

where $L_{\mathbf{f}}h(\mathbf{x}) := (\nabla_{\mathbf{x}}h)\mathbf{f}(\mathbf{x})$, the Lie derivative of the function h along the vector field given by $\mathbf{f}$. Now if $L_{\mathbf{g}}h(\mathbf{x})$ is non-zero for all $\mathbf{x}$, then the system is said to have a strict relative degree of one, and we can force it to track a desired trajectory y_{d} by choosing

$$u_{\mathrm{d}} = \frac{1}{L_{\mathbf{g}}h(\mathbf{x})}(\dot{y}_{\mathrm{d}} - L_{\mathbf{f}}h(\mathbf{x})) \tag{7}$$

and making sure that the initial conditions match $(y(0) = y_{\mathrm{d}}(0), \dot{y}(0) = \dot{y}_{\mathrm{d}}(0))$. Changing variables according to $\xi = \Phi(\mathbf{x}) := h(\mathbf{x})$ and defining for the nonce $\dot{y}_{\mathrm{d}} =: v$, a synthetic input, we see a one-dimensional (1D) linear system $\dot{\xi} = 0 \cdot \xi + v$ and an $(n-1)$-dimensional non-linear system $\dot{\boldsymbol{\eta}} = \lambda(\xi, \boldsymbol{\eta})$. (The variable v is called a synthetic input in virtue of its role as the input to this linear system.) Note that the linear system is decoupled from the non-linear one, in the sense that $\dot{\xi}$ is not a function of $\boldsymbol{\eta}$.

If, on the other hand, $L_{\mathbf{g}}h(\mathbf{x})$ is zero for all $\mathbf{x}$, then we differentiate y a second time to obtain

$$\ddot{y} = L_{\mathbf{f}}^2h(\mathbf{x}) + L_{\mathbf{g}}L_{\mathbf{f}}h(\mathbf{x})u \tag{8}$$

and, if for all $\mathbf{x}$, $L_{\mathbf{g}}L_{\mathbf{f}}h(\mathbf{x}) \neq 0$ (i.e. the system has strict relative degree two), choose

$$u_{\mathrm{d}} = \frac{1}{L_{\mathbf{g}}L_{\mathbf{f}}h(\mathbf{x})}(\ddot{y}_{\mathrm{d}} - L_{\mathbf{f}}^2h(\mathbf{x})) \tag{9}$$

this time making sure that $\ddot{y}(0) = \ddot{y}_{\mathrm{d}}(0)$ as well. Choosing $\xi = \Phi(\mathbf{x}) := [h(\mathbf{x}), L_f h(\mathbf{x})]^{\mathrm{T}}$ again yields a linear system, this time 2D, which is decoupled from the remaining $(n-2)$-dimensional non-linear system.

In general, for a system of strict relative degree q, the control law is

$$u_{\mathrm{d}} = \frac{1}{L_{\mathbf{g}}L_{\mathbf{f}}^{q-1}h(\mathbf{x})}(v - L_{\mathbf{f}}^q h(\mathbf{x})) \tag{10}$$

where we choose the synthetic input $v = y_{\mathrm{d}}^{(q)}$ and set all the initial conditions $y_{\mathrm{d}}(0), \dot{y}_{\mathrm{d}}(0), \ldots, y_{\mathrm{d}}^{(q)}(0)$ appropriately. Note

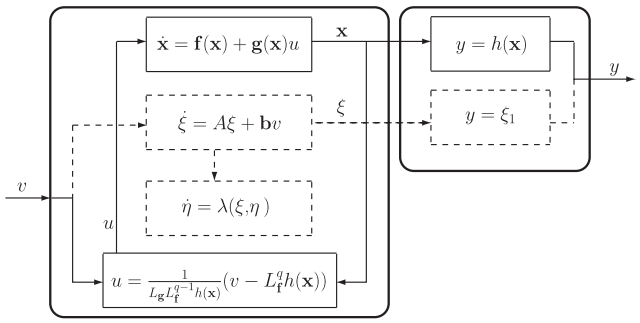

Fig. 1 *A conceptual rendering of feedback linearisation*
The arrangement can be viewed either as a controller in a feedback loop with the original control-affine system (solid lines), plus an output function; or as an equivalent output function along with coupled linear and non-linear systems (dashed lines), with the former evolving independently of the latter, and the output depending solely on the former

that, if $q = n$, the dimension of the state space, then

$$\Phi(\mathbf{x}) := [h(\mathbf{x}), L_\mathbf{f} h(\mathbf{x}), \ldots, L_\mathbf{f}^{n-1} h(\mathbf{x})]^\mathrm{T} \tag{11}$$

is a valid change of coordinates (locally diffeomorphic), which in fact puts the system into controllable canonical form, that is, (3) is fully linearisable by state feedback. The procedure is illustrated in Fig. 1.

Finally, we have so far ignored the case where $L_\mathbf{g} L_\mathbf{f}^{q-1} h(\mathbf{x})$ is zero for some values of $\mathbf{x}$, and non-zero for others. In this case, the relative degree of the system is not well defined, and complications ensue. (We shall see in the sequel that this is indeed the case in the present system.)

3 Application to the model

3.1 Decomposition

The set of ODEs in the present model leaves out two important sections of the coagulation cascade: (i) the intiation of the intrinsic pathway; and (ii) the final portion of the cascade, in which thrombin activates factors I and XIII and a clot forms. We claim, however, that control of the cascade that is indeed represented in the ODEs of this paper suffices, under certain general conditions, to ensure proper clotting.

We address the beginning of the intrinsic (contact-activation) pathway first. Although it is theoretically possible for defects in the intrinsic pathway to affect clotting, this has not been observed clinically. Indeed, deficiencies of the clotting factors of this early portion of the intrinsic pathway do not have hemorrhagic consequences [8].

Now, this consideration might make us wonder what, in fact, the point of the contact-activation system is. The answer appears to be that it offers a method for initiating clotting after blood-vessel trauma even without the action of tissue factor (viz., through the action of negatively charged surfaces on factor XII) – a redundancy of sorts. (One can imagine circumstances under which TF would fail to be released into the blood stream, or alternatively under which too little of it will have been been synthesised by the endothelial cells.) The control techniques of this paper, then, assume the presence of sufficient tissue factor.

The situation with the final portion of the cascade is slightly more complicated. Now, obviously, a control mechanism that does not include a model of the generation of factors XIIIa and Ia cannot guarantee healthy clotting if the production of these factors or the initial values of their unactivated forms are pathological; so use of the controllers derived in this paper assumes at least the the healthy functioning of this final portion of the cascade. However, there is the additional wrinkle that thrombin complexes with both factors I and XIII, which binding is not modelled by the system of ODEs, and which will perforce decrease the amount of free thrombin available for the modelled interactions. We assume, however, following [9], that very little of either complex forms.

Armed with these considerations, we restrict our attention from here on to a set of ODEs that describe the entire coagulation cascade except the initiation of the intrinsic pathway and the final portion of the cascade (factors I and XIII). Now, to prove that control of these ODEs suffices to induce healthy clotting, we should at least have to show that the amount of factor XIa produced by the former is negligible compared with the amount produced by reciprocal activation by factors Xa and IIa. In fact, we might be interested in showing for what range of initial conditions in this 'upstream' subsystem this result obtains. Similarly, a proof would require us to show that the thrombin trajectories we control are sufficient to produce a clot, and within a reasonable amount of time. This question, too, might be posed in terms of ranges: what range of thrombin is sufficient to effect this end?

Both of these questions are questions of reachability; they are addressed in detail in [5]. Presently, however, we make do without any proof. The idea, then, is to steer model thrombin concentrations, with the hope that it can be proven that these trajectories suffice to induce proper clotting. Lacking such a proof (and not wanting to rely on hope), we will throughout aim at the tightest control possible: that is, we will attempt to force thrombin concentrations in diseased-patient models to match the trajectory of thrombin during a healthy clotting event. This may be overkill in the sense that looser criteria might be adduced, via the reachability analysis lately outlined, to guarantee proper clotting, but it is certainly sufficient (at least insofar as the model is correct).

Finally, if our controller is going to neglect these two portions of the cascade, then certainly the input must not interact with them. Fortunately, the major drugs for the major clotting problems do not. In the following demonstrations, we focus on the two most common bleeding disorders, factor-V Leiden and hæmophilia A. The latter is treated with (recombinant) factor VIII, which lies in the subsystem modelled by the ODEs under consideration in this paper; and the former is treated with heparin, which – as we shall see – increases the efficacy of antithrombin in inhibiting thrombin and factor Xa, all of which proteins are again modelled by the ordinary differential equations appearing below.

3.2 Feedback linearisation

Full-state linearisation as lately described faces a significant issue when applied to the present model. Although verifying the two conditions associated with (4) and (5) is mathematically straightforward, it is computationally intensive. Computing the vector fields requires the computation of two Jacobians for every Lie derivative, although one of them, $\partial \mathbf{f}/\partial \mathbf{x}$, need only be computed once. However, each of the approximately 100 vector fields of (4) needs to be computed, and computation of each of the associated Jacobian matrices entails 100^2 derivatives of polynomials. Even neglecting the multiplication and addition operations, this

doi: 10.1049/iet-cta.2011.0475

brings the total to a million computations – each of which is a (symbolic) derivative of a polynomial, increasing the total number of computations yet more. And most unfortunately of all, each derivative (at least potentially) generates more variables, via the chain rule of differentiation, so the polynomials have increasingly more terms in later vector fields.

Apart from computational difficulties, it remains to construct the output function, $h(\mathbf{x})$. No mechanical procedure for this construction exists.

In fact, we can circumvent the computational obstacles. As lately noted, the rank of the strong-accessibility distribution, (4), must be full for the system to be fully linearisable. However, as [10] have pointed out, in the case of reaction systems like the present one, this rank is bounded from above by the rank of a matrix D (the 'accessibility matrix') that is independent of x, and is computed very simply. The existence of such a matrix depends once again on being able to express the governing ODEs in an even more specific form. In particular, (3) can be written

$$\dot{\mathbf{x}} = C\mathbf{r}(\mathbf{x}) + \mathbf{b}u \tag{12}$$

where each element in the vector $\mathbf{r}(\mathbf{x})$ is a monomial corresponding to one of the chemical reactions of the system (i.e. one of the arrows in the chemical formulae); and $\mathbf{b}$ is a constant vector, that is, does not depend on the state. The present case is a special case even of the one discussed in [10], since our model includes neither inflow nor outflow of reactants. For this system, the accessibility matrix is the augmented matrix $D := [\mathbf{b}, C]$, which we now show.

Consider a vector field $\mathbf{s}(\mathbf{x})$ in the span of D, that is

$$\mathbf{s}(\mathbf{x}) = \sum_i \mathbf{d}_i \phi_i(\mathbf{x}) \tag{13}$$

for arbitrary functions of $\mathbf{x}$, ϕ_i, and where $\mathbf{d}_i$ are the columns of D as defined above. Then the Lie bracket of the drift vector field $\mathbf{f}$ with $\mathbf{s}$ is

$$\begin{aligned}[\mathbf{f}, \mathbf{s}] &= \frac{\partial \mathbf{s}}{\partial \mathbf{x}}\mathbf{f}(\mathbf{x}) - \frac{\partial \mathbf{f}}{\partial \mathbf{x}}\mathbf{s}(\mathbf{x}) \\ &= \sum_i \mathbf{d}_i (\nabla_{\mathbf{x}} \phi_i)\mathbf{f}(\mathbf{x}) - C\frac{\partial \mathbf{r}}{\partial \mathbf{x}}\mathbf{s}(\mathbf{x})\end{aligned} \tag{14}$$

Now, since $(\nabla_{\mathbf{x}} \phi_i)\mathbf{f}(\mathbf{x})$ is just another arbitrary function of $\mathbf{x}$, the first term is within the span of D (compare this term with (13)). Similarly, the second term lies within the span of C and thus *a fortiori* lies within the span of D. Thus, we see that the Lie bracket of $\mathbf{f}$, the drift term of our ODE, with any vector field in the span of D, yields a vector field that is itself in the span of D. Since the control vector field $\mathbf{g}(\mathbf{x})$ is here just the constant vector $\mathbf{b}$, it too is in Sp$\{D\}$; and hence $ad^1_{\mathbf{f}}\mathbf{g} = [\mathbf{f}, \mathbf{g}]$ is as well, by the fact just noted. Moreover, now since $ad^1_{\mathbf{f}}\mathbf{g} \in$ Sp$\{D\}$, then so is $ad^2_{\mathbf{f}}\mathbf{g} = [\mathbf{f}, ad^1_{\mathbf{f}}\mathbf{g}]$ – and so on, extending the claim to all the vector fields in the strong-accessibility distribution. Thus we can say

$$\forall \mathbf{x}, \quad \mathrm{rk}(D) \geq \mathrm{rk}([\mathbf{g}(\mathbf{x}), ad_{\mathbf{f}}\mathbf{g}(\mathbf{x}), \ldots, ad^{n-2}_{\mathbf{f}}\mathbf{g}(\mathbf{x}), ad^{n-1}_{\mathbf{f}}\mathbf{g}(\mathbf{x})]) \tag{15}$$

Crucially, the rank of D is easy to compute, being but a constant matrix, and we avoid grinding through the iterated Lie derivatives of the right-hand side.

For a hæmophiliac patient, we use factor VIII as the single input, giving $rk(D) = 83 < n = 98$ – so the system is not fully linearisable by state feedback; that is, there is no single-variable output function $h(\mathbf{x})$ that can be constructed to support such feedback. Now, factor-V Leiden implies the absence of certain reactions (and corresponding state variables), viz. those pertaining to the inactivation of factor Va. (The modelling of factor-V Leiden is explained at length below.) Treatment with heparin in turn requires the addition of several reactions and state variables. However, none of these changes, nor the different control vector field $\mathbf{b}$ – with its single non-zero entry in a different row – suffice to eliminate the deficit between the rank of D and the dimension of the state: $\mathrm{rk}(D) = 73 < n = 89$. The meaning of this deficit we shall discuss later.

So the system is not completely controllable. However, there is reason to believe that clotting can be controlled just by controlling the concentration of thrombin (activated factor II): clotting occurs downstream in the coagulation cascade of the ODEs of this study, and interacts with them only via thrombin (as we saw in the previous section). Thus we attempt to partially linearise the system of (3) with thrombin the output $y = h(\mathbf{x})$. Indeed, using heparin as the control variable and augmenting the system with the appropriate terms from the heparin chemical reactions, the strict relative degree of the system is two. (Lacking complete controllability, is there something the system has gained? That is, is there a trade-off in the area? An affirmative answer can be given if one considers the computational complexity: in order to compute the feedback law, we need only calculate [two] twice-iterated Lie brackets; whereas if the system were completely controllable, the two vector fields would require the calculation of n-iterated Lie brackets. In a system as complicated as ours, this would be a serious problem, even given a simple output function. Additionally, we need only specify two, rather than n, initial conditions. This question is even more relevant when an output function that will provide strict relative degree n is available.) Alternatively, using factor VIII as the input (for hæmophilia) yields a strict relative degree of three.

However, the relative degree for neither of these systems is well-defined; in particular, the relevant Lie derivative is zero at the initial condition, x_0. We circumvent this difficulty by allowing the plant to run uncontrolled until the concentration of thrombin is some distance δ from the origin, that is, until the state has drifted away from the singularity at which the relative degree changes, before turning the controller on. The parameter δ is tuned manually for ideal tracking. (The non-regularity of the system and its consequences are discussed in more detail in Section 6.2 below.)

3.2.1 Error correction: Now, in practice, small numerical discrepancies arise – for example, from a mismatch between the plant and the model, or from numerical approximations – so that the trajectory produced by the control law of (10) deviates from the desired trajectory. Compensation is made in the form of additional elements in the feedback loop, usually some variant on the PID controller. A proportional and an integral term are generally used [11], but here we content ourselves with proportional and derivative terms (as many as the strict relative degree of the system) since (i) the integral term did not seem to improve tracking in our simulations; and (ii) it can be shown [6] that, if the system of (3) is 'globally exponentially minimum-phase', then this control law guarantees bounded trajectory tracking, that is, the tracking error and its derivatives tend asymptotically to

zero. That is, we replace the synthetic input v in (10) with

$$v(t) = y_{\mathrm{d}}^{(q)}(t) + [K_p \;\; K_d^1 \;\; \cdots \;\; K_d^{q-1}][\mathbf{y}_d - \boldsymbol{\xi}] \qquad (16)$$

where $\mathbf{y}_{\mathrm{d}}$ is the vector of the desired output and its derivatives:

$$\mathbf{y}_{\mathrm{d}} = [y_{\mathrm{d}} \;\; \dot{y}_{\mathrm{d}} \;\; \cdots \;\; y^{(q-1)}]^{\mathrm{T}}$$

and ξ is the linearised state, which in our system corresponds to the actual output y and its derivatives.

The question of whether the system is indeed globally exponentially minimum-phase turns out to be complicated; it is discussed more fully in in a companion article, but here we summarise those conclusions. The usual test is to make sure that the 'zero-dynamics' – the evolution of the uncontrolled, non-linear subsystem $\dot{\boldsymbol{\eta}} = \lambda(\xi, \boldsymbol{\eta})$ (see Fig. 1), constrained to the manifold where the linear-system variables ξ are zero – is itself exponentially stable at the relevant equilibrium point. Unfortunately, the change of coordinates $(\boldsymbol{\xi}, \eta) = \Phi(\mathbf{x})$ breaks down (loses rank) precisely here owing to the non-regularity of the system – due in particular to the zero concentration of thrombin. More sophisticated approaches [12, 13] require computing the relative degree of the system at the singularity; but being much higher than 2 in our model, this turns out to incur similar, prohibitive, computational penalties in iterated Lie brackets as the ones lately discussed in generating the full accessibility distribution.

Failure to apply the usual tests does not, of course, show that the system is *not* exponentially minimum-phase, and moreover there are two additional reasons for suspecting the non-linear dynamics to be stable. First, the coagulation cascade as modelled here – even in the case of factor-V Leiden – ultimately eventuates in the binding of all active proteases into inactive complexes. (Whether this happens exponentially rather than merely asymptotically is not as obvious.) The introduction of the anti-coagulatory heparin into the system is unlikely to upset this stability. Second, the control schemes do appear to achieve bounded tracking in continuous-time simulation below.

The gains were calculated by treating the system as a linear-quadratic regulator, with – unless otherwise noted – state-penalty matrix set to the identity matrix and the input penalty set to unity. Solution of the optimisation problem with MATLAB's `lqr` function gives $K_p = 1, K_d = 1.7321$ for control under factor-V Leiden. However, control of the non-linear system was observed to be fairly insensitive to the choice of cost parameters. Increasing or decreasing the ratio of the state penalty to the input penalty by as much as an order of magnitude yielded qualitatively similar thrombin trajectories (but see the discussion in Section 6.2 below).

3.2.2 Discretisation: Finally, practical application will also require discretisation of the controller. In the simulations below, this was achieved simply by allowing it to update only every 0.5 s. The model was then simulated in continuous time using a zero-order hold. Now, the sampling rate of the controller affects the controllability of the system: under a sufficiently low rate the theory lately outlined will fail to achieve exact output tracking. Two hertz was therefore chosen simply by determining, via trial and error, the lowest rate that would not severely degrade the results.

3.3 Step-input control

A much simpler control technique was also applied; the rationale for it is discussed below. This strategy is predicated on the assumption that a single-step input might suffice to force the system to track the desired thrombin trajectory 'reasonably well'. More precisely, the assumption is that the step input that matches both the desired peak concentration of thrombin and the occurrence of this peak will result in a thrombin trajectory that deviates very little over its entirety from its desired counterpart. Therefore this technique was implemented by performing a parameter search over repeated trials for that step input that would minimise the peak-concentration and peak-time discrepancies. Thus the input u was modified on successive trials according to:

$$u \leftarrow u + \alpha \left(\max_t\{y(t)\} - \max_t\{y_{\mathrm{d}}(t)\} \right) + \beta \left(\mathrm{argmax}_t\{y(t)\} - \mathrm{argmax}_t\{y_{\mathrm{d}}(t)\} \right) \qquad (17)$$

where α and β-positive for anti-coagulant inputs and negative for pro-coagulants – scale the relative contribution of each term, and were adjusted by hand. The search was performed as follows. First, the time and location of the desired thrombin peak were computed from the desired thrombin trajectory (for its computation, see Section 4 below). Then, for fixed α and β and an input u initialised at 0, the cascade was simulated, from which the time and location of the thrombin peak was calculated. The input was then modified according to (17) and the process repeated until convergence. Since the relative weighting of peak-concentration and peak-time discrepancies yielded better and worse matches of the *overall* trajectory, it was adjusted (via α and β) by hand, another search performed, and so on; until visual inspection showed a good match. Useful values of α and β are reported in the results below (in general, a good ratio of the former to the latter was about five).

4 Methods

We draw our chemical reactions from three main sources, the first two of which collect their own data from a number of primary sources (the last is itself a primary source): [3, 4, 14]. A description of each follows. The entire system is summarised in Fig. 2.

Forty-eight chemical reactions for normal blood clotting, with no exogenous intervention, were taken directly from [3], from which also the numerical values of the rate constants were supplied. They are reproduced here in Table 1. The blood factors are referred to by their Roman-numeral designations, a lowercase 'a' denoting the activated form. Other abbreviations include: mIIa for meizothrombin, LBS for the concentration of lipid binding sites, PS for protein S, PC for protein C, APC for its activated form, TFPI for tissue-factor-pathway inhibitor, Tm for thrombomodulin, AT for antithrombin and TF for tissue factor. A subscripted 'i' indicates the inactivated form of an enzyme, and a subscripted 'L' indicates the lipid-bound form.

The simple mechanism of factor Va inactivation of the [3] model consisted of a single chemical equation. We substituted for this equation the model of factor-Va inactivation of [4], which was designed specifically to match data for factor-V Leiden. (For completeness: in fact, we used all of the chemical equations from [3] except two: the deactivation of $\mathrm{FVa_L}$ by APC; and the unbinding of $\mathrm{FVai_L}$, the lipid-bound, inactivated form of factor Va, from its phospholipid surface – since, after all, $\mathrm{FVai_L}$ never shows up in our model, being replaced by various cleaved forms of factor Va.) These 34 chemical equations are listed in Table 2. Factor-V Leiden is a genetic disorder in which the amino

doi: 10.1049/iet-cta.2011.0475

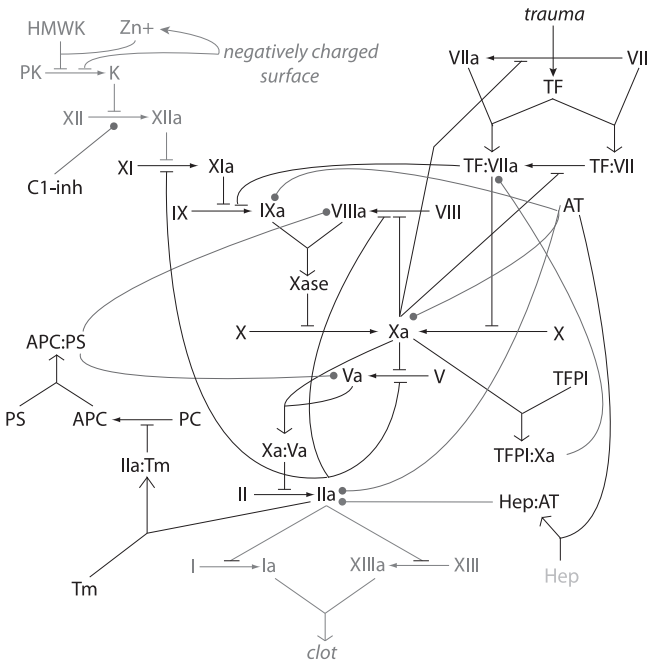

Fig. 2 *Coagulation cascade*
Triangle-headed arrows indicate transformations to activated forms; flat-headed lines connect the catalysts to the reaction they catalyse; line-arrowheaded arrows indicate binding reactions; and ball-ended lines denote inhibitory reactions. Reactions that are not treated in this model (see Section 3.1) have been greyed out. The input for treatment of factor-V Leiden, heparin, is in pale grey. The remaining arrows (true-arrowheaded) mark the initiation points of the cascade. Some reactions have been omitted to reduce clutter: the lipid bindings, those involving meizothrombin (mIIa; since its role is identical to thrombin), the sub-reactions that compose factor-V deactivation (Table 2), and the sub-reactions of heparin-mediated thrombin inactivation (Table 3)

acid Arg^{505} of the heavy chain of factor V is replaced by Gln^{505}, which effectively disables APC-mediated cleavage of factor V at this point. The disorder was discovered only as recently as 1994, but it is in fact quite common, affecting perhaps 5% of the white population, among whom it is the most prevalent of the inherited hypercoagulatory disorders. Now, [4] were able to match data for inactivation of FVa^{LEIDEN} by assuming that the form of factor V cleaved at Arg^{505}, Va_5, has no effect on clotting, that is, it is 'inactivated'. (This is so even despite its reported 60% efficacy in complex with factor Xa as prothrombinase. Since it has no overall effect, and for the sake of simplicity, we treated $Va5_L$ as inert.)

We made a further simplification in the incorporation of this model: [4] model the inactivation of factor Va by APC, rather than the more efficacious APC:PS complex, as in [3]. Now in fact, uncomplexed APC competes with APC:PS for binding with factor Va, but [3] leaves out this competition, that is, only APC:PS is modelled as binding with Va – presumably because, under normal conditions, very little free APC circulates: the vast majority of it is immediately bound up by protein S. We follow this expedient in appropriating the [4] model, replacing APC with APC:PS in all the relevant chemical equations. Does not this substitution require an alteration of the rate constants? After all, APC:PS is supposed to be more efficacious. Yes: and here we follow the results of [15], who found that the increased efficacy of the complex results from a factor of three increase in the cleavage rate of Va at Arg^{306}. (Previously it had been reported that the factor was twenty rather than three, but as [15] notes, this would abolish the thrombophilic effect of factor-V Leiden: in the presence of protein C, the cleavage at Arg^{306} would dominate the cleavage at Arg^{505}, so that the disabling of the latter in FVL would have no overall effect on clotting.) In the interest of perspicuity, the factor of three and the original rate constant from [4] are listed explicitly in Table 2. Of course, if we wanted to model (e.g.) protein-S deficiency, we should need to put the equations for (uncomplexed) APC inactivation of FVa into the model.

Does this complicated replacement for the single equation of [3] yield similar results under normal (healthy) conditions? In fact, there is a significant change: under the more complicated model, thrombin peaks during normal clotting around 90 nM, whereas under [3]'s model, it peaked at almost 40 nM (although, curiously, the time of the peak is the same in both models). Now [3] verifies their model by comparing it to clinical observations of thrombin curves under various physiological conditions, in particular by comparison with Fig. 5 of [16]. However, oddly enough, these conditions assume no protein C or protein S, so [3]'s simulations do not actually verify this part of their model. In point of fact, [3]'s 'verifications' are anyway qualitative: the thrombin curves from [16], while peaking at similar concentrations, are somewhat delayed compared with those of [3]. Resolution of this issue, too, we defer to the discussion section below. For now it should be said that the more detailed model of [4] is probably more accurate than the single equation of [3], so we proceed with some confidence.

The heretofore unexplained abbreviations appearing in the table are as follows. The numbers 3, 5, and 6 denote the site or sites at which the species of factor Va has been cleaved, respectively, Arg^{306}, Arg^{505} and Arg^{662}. Thus, Va53 is factor Va cleaved at both Arg^{306} and Arg^{505}; and VaA3 is factor V's A2 peptide chain, cleaved at Arg^{306}. LC denotes the light peptide chain of factor Va and HC its heavy chain; and VaLC is the light chain and A1 peptide chain. The subscripted L again denotes the lipid-bound species.

A word on the lipid-binding: the reader may notice that not all of the proteins subscripted with an L have equations for dissociating from their lipid substrates. In this, we are for the most part following [3], in which only the simple proteins, but not the compounds, are allowed to dissociate. For the remainder, we have followed [4] in not allowing the uncomplexed proteins LC_L, $Va5_L$, $Va3_L$, $Va53_L$, $Va36_L$, $Va56_L$, $Va536_L$ and $VaLC_L$ to dissociate from their substrates. Now, in fact, the experiments in [4] assume a saturating amount of lipids; whereas in our model, the phospholipid concentration is something less than saturating. Technically, then, these proteins should be allowed to dissociate. However, assuming on- and off-rates equal to that for (uncleaved) factor Va, these dissociations have an entirely negligible effect even on the concentrations of the FVa-variants themselves, let alone on thrombin levels – which is to be expected, since phospholipids never fall below about 1500 nM throughout the simulations, at which level-binding forces vastly predominate over unbinding forces (cf. the rate constants of chemical reaction 4). Note, also, that the heavy-chain variants of factor Va are not lipid bound: the binding site for lipids lies in the light chain of factor V [4].

Finally, on this head, a correction was required: [3] claim that thrombin is released from the lipid surface when it is activated. However, conservation of mass requires these lipid-binding sites to be returned to the pool of uncomplexed LBS; thus in chemical reaction 33 we have added

Table 1 Chemical reactions and rate constants drawn from [3]; they govern all the reactions except factor Va inactivation and the interaction of heparin with the system

No.	Reaction	k_{on}, $nM^{-1}\ s^{-1}$	k_{off}, s^{-1}	k_{cat}, s^{-1}
1	$II + LBS \rightleftharpoons II_L$	0.0043	1	–
2	$mIIa + LBS \rightleftharpoons mIIa_L$	0.05	0.475	–
3	$V + LBS \rightleftharpoons V_L$	0.05	0.145	–
4	$Va + LBS \rightleftharpoons Va_L$	0.057	0.17	–
5	$VII + LBS \rightleftharpoons VII_L$	0.05	0.66	–
6	$VIIa + LBS \rightleftharpoons VIIa_L$	0.05	0.227	–
7	$VIII + LBS \rightleftharpoons VIII_L$	0.05	0.1	–
8	$VIIIa + LBS \rightleftharpoons VIIIa_L$	0.05	0.335	–
9	$IX + LBS \rightleftharpoons IX_L$	0.05	0.115	–
10	$IXa + LBS \rightleftharpoons IXa_L$	0.05	0.115	–
11	$X + LBS \rightleftharpoons X_L$	0.01	1.9	–
12	$Xa + LBS \rightleftharpoons Xa_L$	0.029	3.3	–
13	$APC + LBS \rightleftharpoons APC_L$	0.05	3.5	–
14	$PS + LBS \rightleftharpoons PS_L$	0.05	0.2	–
15	$VIIIa_i + LBS \rightleftharpoons VIIIa_{i,L}$	0.05	0.335	–
16	$PC + LBS \rightleftharpoons PC_L$	0.05	11.5	–
17	$TF_L + VIIa_L \rightleftharpoons TF{:}VIIa_L$	0.5	0.005	–
18	$TF_L + VII_L \rightleftharpoons TF{:}VII_L$	0.005	0.005	–
19	$TF{:}VIIa_L + IX_L \rightleftharpoons TF{:}VIIa{:}IX_L \rightarrow TF{:}VIIa_L + IXa_L$	0.01	2.09	0.34
20	$TF{:}VIIa_L + X_L \rightleftharpoons TF{:}VIIa{:}X_L \rightarrow TF{:}VIIa{:}Xa_L$	0.1	32.5	1.5
21	$TF{:}VIIa{:}Xa_L \rightarrow TF{:}VIIa_L + Xa_L$	–	1	–
22	$TF{:}VII_L + Xa_L \rightleftharpoons TF{:}VII{:}Xa_L \rightarrow TF{:}VIIa_L + Xa_L$	0.05	44.8	15.2
23	$IXa_L + VIIIa_L \rightleftharpoons IXa{:}VIIIa_L$	0.1	0.2	–
24	$Xa_L + Va_L \rightleftharpoons Xa{:}Va_L$	1	1	–
25	$IXa{:}VIIIa_L + X_L \rightleftharpoons IXa{:}VIIIa{:}X_L \rightarrow IXa{:}VIIIa_L + Xa_L$	0.1	10.7	8.3
26	$V_L + Xa_L \rightleftharpoons V{:}Xa_L \rightarrow Va_L + Xa_L$	0.1	1	0.043
27	$VIII_L + Xa_L \rightleftharpoons VIII{:}Xa_L \rightarrow VIIIa_L + Xa_L$	0.1	2.1	0.023
28	$V_L + IIa \rightleftharpoons V{:}IIa_L \rightarrow Va_L + IIa$	0.1	6.94	0.23
29	$VIII_L + IIa \rightleftharpoons VIII{:}IIa_L \rightarrow VIIIa_L + IIa$	0.1	13.8	0.9
30	$Xa{:}Va_L + II_L \rightleftharpoons Xa{:}Va{:}II_L$	0.1	100	–
31	$Xa{:}Va_L + mIIa_L \rightleftharpoons Xa{:}Va{:}mIIa_L$	0.1	66	–
32	$Xa{:}Va{:}II_L \rightarrow Xa{:}Va{:}mIIa_L$	–	–	13
33	$Xa{:}Va{:}mIIa_L \rightarrow Xa{:}Va_L + IIa + LBS$	–	–	15
34	$VII_L + Xa_L \rightleftharpoons VII{:}Xa_L \rightarrow VIIa_L + Xa_L$	0.05	44.8	15.2
35	$XI + IIa \rightleftharpoons XI{:}IIa \rightarrow XIa + IIa$	0.1	10	1.43
36	$APC{:}PS_L + VIIIa_L \rightleftharpoons APC{:}PS{:}VIIIa_L \rightarrow APC{:}PS_L + VIIIa_{i,L}$	0.1	1.6	0.4
37	$TFPI + Xa \rightleftharpoons TFPI{:}Xa$	0.016	0.000333	–
38	$TFPI{:}Xa + TF{:}VIIa_L \rightleftharpoons TFPI{:}Xa{:}TF{:}VIIa_L$	0.01	0.0011	–
39	$IXa + AT \rightarrow IXa{:}AT$	4.9×10^{-7}	–	–
40	$Xa + AT \rightarrow Xa{:}AT$	2.3×10^{-6}	–	–
41	$IIa + AT \rightarrow IIa{:}AT$	6.83×10^{-5}	–	–
42	$V_L + mIIa_L \rightleftharpoons V{:}mIIa_L \rightarrow Va_L + mIIa_L$	0.1	6.94	1.035
43	$VIII_L + mIIa_L \rightleftharpoons VIII{:}mIIa_L \rightarrow VIIIa_L + mIIa_L$	0.1	13.8	0.9
44	$IIa + Tm_L \rightleftharpoons IIa{:}Tm_L$	1	0.5	–
45	$IIa{:}Tm_L + PC_L \rightleftharpoons IIa{:}Tm{:}PC_L \rightarrow IIa{:}Tm_L + APC_L$	0.1	6.4	3.6
46	$mIIa + AT \rightarrow mIIa{:}AT$	6.83×10^{-6}	–	–
47	$APC_L + PS_L \rightleftharpoons APC{:}PS_L$	0.1	0.5	–
48	$XIa + IX_L \rightleftharpoons XIa{:}IX_L \rightarrow XIa + IXa_L$	0.01	1.417	0.183

an LBS term that was missing from [3]. Of course, as we have just finished saying, the concentration of binding sites is in excess over the other proteins, so this addition makes no discernible difference to the simulations.

The form of the heparin equations comes from [14] and appear in Table 3. The abbreviated proteins are heparin, thrombin (activated factor II), and antithrombin; (IIa-AT) is a stable complex, which can no longer dissociate into inhibitor and protease.

Such are the forms of the heparin reactions; however, to our knowledge, the exact values of the on- and off-rates under physiological conditions have never been determined. In [14], however, are found estimates of the ratios of off- to on-rates; the following expedient was therefore adopted:

doi: 10.1049/iet-cta.2011.0475

Table 2 Chemical reactions and rate constants of the [4] model of factor-Va deactivation; factor-V Leiden is modelled as knocking out the reactions that cleave factor V at the Arg^{505} site; and proteins cleaved at this site are marked by a 5 (see the text for abbreviations)

No.	Reaction	k_{on}, nM/s	k_{off}, s^{-1}	k_{cat}, s^{-1}
49	$LC_L + HC \rightleftharpoons Va_L$	2.63×10^{-6}	1.72×10^{-5}	–
50	$LC_L + HC5 \rightleftharpoons Va5_L$	2.63×10^{-6}	1.72×10^{-5}	–
51	$LC_L + HC3 \rightleftharpoons Va3_L$	2.63×10^{-6}	1.72×10^{-5}	–
52	$LC_L + HC53 \rightleftharpoons Va53_L$	2.63×10^{-6}	1.72×10^{-5}	–
53	$LC_L + HC36 \rightleftharpoons Va36_L$	2.63×10^{-6}	1.72×10^{-5}	–
54	$LC_L + HC56 \rightleftharpoons Va56_L$	2.63×10^{-6}	1.72×10^{-5}	–
55	$LC_L + HC536 \rightleftharpoons Va536_L$	2.63×10^{-6}	1.72×10^{-5}	–
56	$LC_L + APC{:}PS_L \rightleftharpoons LC{:}APC{:}PS_L$	2.63×10^{-6}	1.72×10^{-5}	–
57	$Va_L + APC{:}PS_L \rightleftharpoons Va{:}APC{:}PS_L$	0.1	0.7	–
58	$Va5_L + APC{:}PS_L \rightleftharpoons Va5{:}APC{:}PS_L$	0.1	0.7	–
59	$Va3_L + APC{:}PS_L \rightleftharpoons Va3{:}APC{:}PS_L$	0.1	0.7	–
60	$Va53_L + APC{:}PS_L \rightleftharpoons Va53{:}APC{:}PS_L$	0.1	0.7	–
61	$Va36_L + APC{:}PS_L \rightleftharpoons Va36{:}APC{:}PS_L$	0.1	0.7	–
62	$Va56_L + APC{:}PS_L \rightleftharpoons Va56{:}APC{:}PS_L$	0.1	0.7	–
63	$Va536_L + APC{:}PS_L \rightleftharpoons Va536{:}APC{:}PS_L$	0.1	0.7	–
64	$VaLC_L + APC{:}PS_L \rightleftharpoons VaLC{:}APC{:}PS_L$	0.1	0.7	–
65	$VaLC_L + VaA3 \rightleftharpoons Va3_L$	2.57×10^{-6}	0.028	–
66	$VaLC_L + VaA53 \rightleftharpoons Va53_L$	2.57×10^{-6}	0.028	–
67	$VaLC_L + VaA36 \rightleftharpoons Va36_L$	2.57×10^{-6}	0.028	–
68	$VaLC_L + VaA356 \rightleftharpoons Va536_L$	2.57×10^{-6}	0.028	–
69	$VaLC{:}APC{:}PS_L + VaA3 \rightleftharpoons Va3{:}APC{:}PS_L$	2.57×10^{-6}	0.028	–
70	$VaLC{:}APC{:}PS_L + VaA53 \rightleftharpoons Va53{:}APC{:}PS_L$	2.57×10^{-6}	0.028	–
71	$VaLC{:}APC{:}PS_L + VaA36 \rightleftharpoons Va36{:}APC{:}PS_L$	2.57×10^{-6}	0.028	–
72	$VaLC{:}APC{:}PS_L + VaA356 \rightleftharpoons Va536{:}APC{:}PS_L$	2.57×10^{-6}	0.028	–
73	$Va{:}APC{:}PS_L \rightharpoonup Va3{:}APC{:}PS_L$	–	–	3*0.064
74	$Va5{:}APC{:}PS_L \rightharpoonup Va53{:}APC{:}PS_L$	–	–	3*0.064
75	$Va5{:}APC{:}PS_L \rightharpoonup Va56{:}APC{:}PS_L$	–	–	5.2×10^{-4}
76	$Va3{:}APC{:}PS_L \rightharpoonup Va36{:}APC{:}PS_L$	–	–	5.2×10^{-4}
77	$Va53{:}APC{:}PS_L \rightharpoonup Va536{:}APC{:}PS_L$	–	–	5.2×10^{-4}
78	$Va56{:}APC{:}PS_L \rightharpoonup Va536{:}APC{:}PS_L$	–	–	3*0.064
79	$Va{:}APC{:}PS_L \rightharpoonup Va5{:}APC{:}PS_L$	–	–	1
80	$Va3{:}APC{:}PS_L \rightharpoonup Va53{:}APC{:}PS_L$	–	–	1
81	$Va36{:}APC{:}PS_L \rightharpoonup Va536{:}APC{:}PS_L$	–	–	1

Table 3 Chemical equations of the heparin reactions; the equations were drawn from [14], but the rate constants were estimated as in the text

No.	Reaction	k_{on}, nM/s	k_{off}, s^{-1}	$k_{cleavage}$, s^{-1}
82	$Hep + IIa \rightleftharpoons Hep{:}IIa$	0.1	$0.1 * 3 \times 10^4$	–
83	$Hep + AT \rightleftharpoons Hep{:}AT$	0.1	$0.1 * 2 \times 10^2$	–
84	$Hep{:}IIa + AT \rightleftharpoons Hep{:}AT{:}IIa$	0.1	$0.1 * 20$	–
85	$Hep{:}AT + IIa \rightleftharpoons Hep{:}AT{:}IIa$	0.1	$0.1 * 3 \times 10^3$	–
86	$IIa{:}AT + AT \rightleftharpoons Hep{:}AT{:}IIa$	0.1	$0.1 * 0.6$	–
87	$Hep{:}AT{:}IIa \rightarrow IIa\text{-}AT$	–	–	5

On-rates were set to an intermediate value within the biologically normal range for enzyme-substrate reactions (see for example Table 4.4 in [17], but also [3]), viz. 0.1 s/nM. The corresponding off-rates were then computed by multiplying the aforementioned ratios by the on-rates, that is, 0.1. Again, this is represented explicitly in the table. This approximation does not vitiate the theoretical apparatus, but it does have numerical consequences; we take up this question again in the discussion below.

The final rate constant is a catalytic rate, and as such is taken directly from [14]. However, we also deviate in one other particular from that paper: whereas [3] model AT binding to IIa as irreversible, with a very slow on-rate, [14] considers the reaction to be reversible, with a very high off-to-on ratio. These mechanisms are certainly not equivalent, and we use the more recent mechanism from [3]. However, in any case, this reaction will be dominated by the much higher affinity reaction between AT and heparin, the primary path to the Hep:AT:IIa complex, so the choice has little practical consequence.

Finally: heparin is also thought to facilitate the inactivation of factor Xa, but once again for simplicity these interactions have been neglected.

Table 4 lists all the non-zero initial concentrations. The concentration of lipid binding sites and the initial activating amount of tissue factor were chosen to match [16], as in [3]. All other concentrations are normal physiological values.

Table 4 Non-zero initial conditions; all other proteins are initialised at zero

Species	Conc., nM	Species	Conc., nM	Species	Conc., nM
tissue factor	0.005	factor VIII	0.7	TFPI	2.5
factor II	1400	factor IX	90	antithrombin	3400
factor V	20	factor XI	30	protein C	60
factor VII	10	factor X	170	protein S	300
factor VIIa	0.1	thrombomodulin	1	LBS	3396

The chemical reactions of Tables 1–3 are equivalent to a set of coupled, non-linear, ODEs. For completeness we illustrate the transformation for the first chemical equation, $\mathrm{II} + \mathrm{LBS} \rightleftharpoons \mathrm{II_L}$. This is equivalent to

$$\frac{d[\mathrm{II}]}{\mathrm{d}t} = -\mathrm{k_{on}}[\mathrm{II}][\mathrm{LBS}] + \mathrm{k_{off}}[\mathrm{II}_L]$$
$$\frac{d[\mathrm{LBS}]}{\mathrm{d}t} = -\mathrm{k_{on}}[\mathrm{II}][\mathrm{LBS}] + \mathrm{k_{off}}[\mathrm{II}_L]$$
$$\frac{d[\mathrm{II}_L]}{\mathrm{d}t} = \mathrm{k_{on}}[\mathrm{II}][\mathrm{LBS}] - \mathrm{k_{off}}[\mathrm{II}_L]$$

where square brackets denote (time-varying) concentration. More succinctly, as in (12), we may write

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix}[\mathrm{II}]\\ [\mathrm{LBS}]\\ [\mathrm{II}_L]\end{bmatrix} = \begin{bmatrix}-1 & 1\\ -1 & 1\\ 1 & -1\end{bmatrix}\begin{bmatrix}k_{\mathrm{on}}[\mathrm{II}][\mathrm{LBS}]\\ k_{\mathrm{off}}[\mathrm{II}_L]\end{bmatrix} \quad (18)$$

Incorporating the other chemical reactions amounts to adding monomials to each ODE as necessary, and of course adding an ODE for each new chemical introduced; or, alternatively, to adding new columns to the numerical matrix, C, for each new rate constant, and adding a new row for each new chemical. This simple structure is a consequence of the mass-action kinetics of the coagulation reactions, and in turn makes possible the algebraic proof exploited above (Section 3.2).

All simulations were performed in MATLAB [18] using the stiff ODE solver `ode15s`. Computation of the strict relative degree of the system was also performed in MATLAB, using the symbolic toolbox. For all simulations, the 'desired trajectory' (time course) of thrombin concentration was precomputed by first simulating healthy coagulation, that is, numerically solving the unmodified system of ODEs. Simulation of the control techniques then took place in models that had been modified appropriately to account for the disease of interest, as above.

5 Results

5.1 Factor-V Leiden: unconstrained, continuous, feedback-linearising control

We first simulate coagulation in a patient with the hypercoagulatory disorder factor-V Leiden, with and without intervention by the anti-coagulant heparin, as well as normal (non-pathological) clotting. Initiation of the clotting event is assumed to take place via the intrinsic pathway and is therefore modelled by initialising tissue factor at a concentration of 5 pM (following [3]).

Fig. 3 shows the thrombin profile during normal clotting (dark line), clotting in a 'patient' with factor-V Leiden (light line), and that clotting in the same patient but with the feedback-linearising controller (medium line). Exact output tracking has been achieved; however, several defects are immediately obvious. First, the controller operates continuously, that is, updating as often as the numerical simulation of the ODEs does, whereas any practical controller must be digital. The input, appearing in at right in Fig. 3, is obviously not feasible for any imaginable drug-delivery mechanism. In fact, to avoid numerical errors, the input was clamped at 100 nM/s. (We explain this issue in the discussion section below.) Finally, and relatedly, a realistic controller will have a maximum and minimum input rate; certainly, the rate cannot be negative, since this implies withdrawal of heparin from the site of the injury. The present controller was allowed us to do just that.

5.2 Factor-V Leiden: constrained, continuous, feedback-linearising control

We address the final defect first. Note, however, that the theory outlined above does not guarantee the success of any these remedies. As lately noted, the input cannot be a negative number, nor can it be much higher than a single-digit nanomolarity per second. (Continuous intravenous infusions over a 24 h period total between 20 000 and 40 000 International units, which comes to about 10 pM/s. Since our treatment lasts only over the course of a few minutes, we can presumably use higher doses over this short interval without overmuch hæmophilic risk downstream.) We impose an upper bound (somewhat arbitrarily) at 20 nM/s and a lower bound of zero by 'squashing' inputs outside this range, that is, setting any u calculated from (10) that exceeds one of these bounds to the value of the bound. The result appears in Fig. 4.

There is indeed some degradation from the unconstrained controller; in particular, the controller overshoots the peak trajectory somewhat, and is unable to compensate for the final undershoot. These defects are respectively the consequences of the maximum and minimum constraints; but the latter at least might be remedied by some form of anticipation; that is, if the controller could predict the final undershoot, it might ease off the input earlier. This possibility is explored in the discussion below. As for the overshoot, it can be eliminated by raising the maximum input to 50 nM/s (graph not shown). However, in the absence of an established maximum, it might be safer to try to address this issue with an anticipatory control as well.

5.3 Factor-V Leiden: constrained, discrete, feedback-linearising control

The corresponding input at right in Fig. 4 still exhibits some chatter, naturally, since it is allowed to vary continuously. That shortcoming is rectified by discretising the controller, which results in the output and input of Fig. 5 (solid lines). Here the output again overshoots the desired

doi: 10.1049/iet-cta.2011.0475

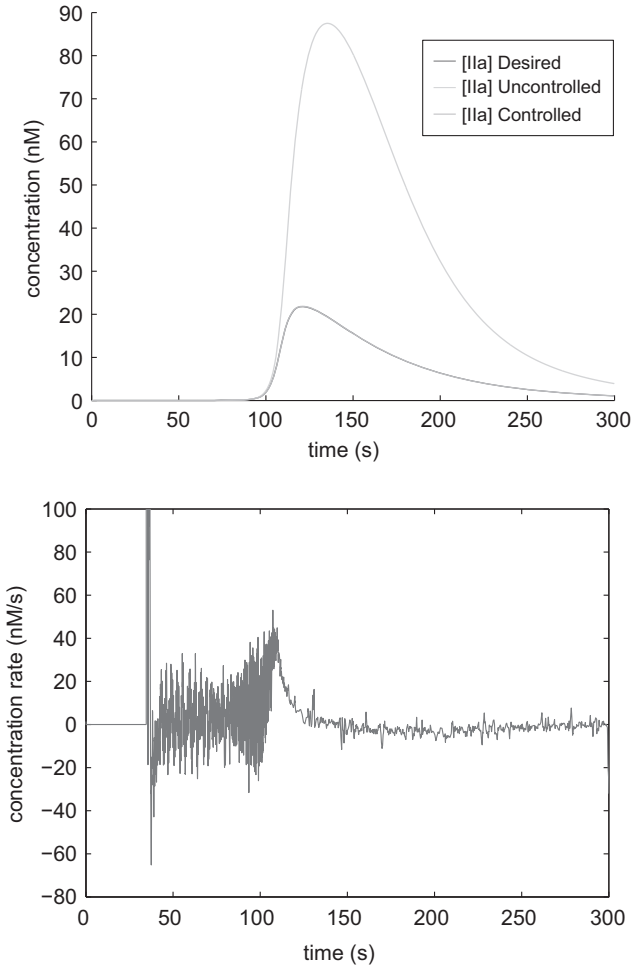

Fig. 3 *Graphs showing thrombin profile during normal clotting*
Top: simulated control of thrombin concentration during a clotting event in a patient with factor-V Leiden. The controller was unconstrained, and was updated continuously, allowing arbitrarily close tracking of the desired trajectory. Bottom: the heparin input in nM/s that generated the output above. Although this input yields perfect tracking, it is obviously unacceptable for any realistic drug-delivery mechanism, operating as it does on a continuous time scale and using negative inputs (see the text)

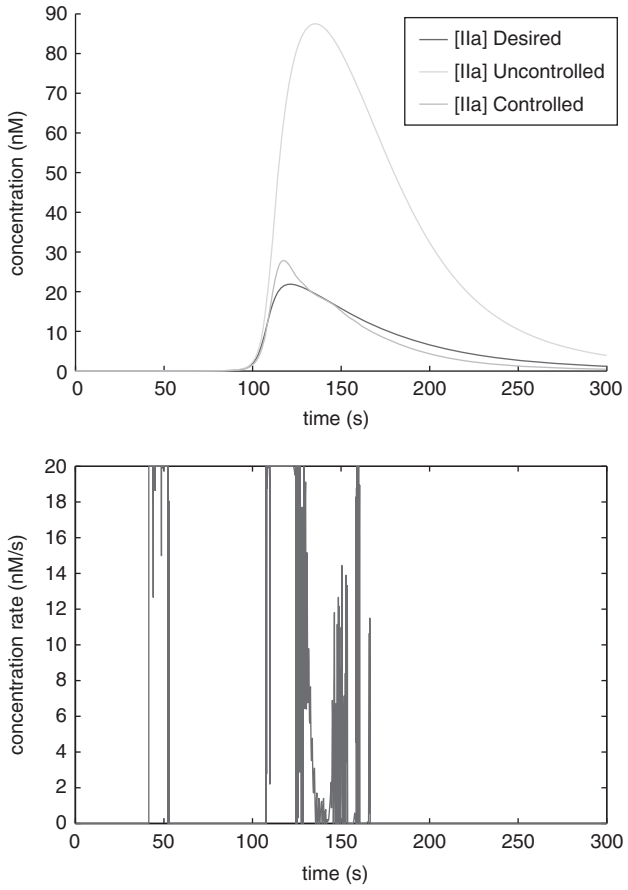

Fig. 4 *Graphs showing*
Top: simulated control of thrombin concentration during a clotting event in a patient with factor-V Leiden, again continuously sampling. Bottom: the input generating the output above, it was constrained to the range 0–20 nM/s

thrombin trajectory, but also undershoots the pre-peak trajectory, evidently because it was unable to recover from the initial input; and, for similar reasons, undershoots the post-peak thrombin curve slightly more than its counterpart continuous controller (Fig. 4).

A smaller controller step yields, as expected, superior results (not shown); and raising the input maximum again eliminates the overshoot. In the limit, of course, we should be able to reproduce the results of the continuous controller. However, and for that very reason, the smaller the step, the more unrealistic the controller. Then, in any case, we shall presumably not be able to eliminate the undershoot that the continuous controller exhibits, since that results from the input minimum; nor, if we take seriously the input maximum of 20 nM/s, shall we be able to eliminate the overshoot.

The 'clipping' that the upper constraint imposes on the input can be eliminated by increasing the input penalty (relative to the state penalty) in the linear-quadratic regulator (producing feedback gains $K_p = 0.1414$, $K_d = 0.5503$), but at the cost of some degradation in output tracking. Nor could adjustment of the relative penalties eliminate clipping at the lower input bound ($u = 0$). In Fig. 5 (dashed lines), the input penalty has been increased from unity to 50 (while the state penalty remains the identity matrix).

Whether the controlled trajectories of Fig. 5 present a hypercoagulatory risk is something of an open question: Rapid product formation 'downstream' in the cascade, that is, the formation of fibrin from fibrinogen and activation of factor XIII – requires concentrations of thrombin less than 2 nM [19], so it is not clear how important the exact trajectory of the thrombin curve is. Nevertheless, it may be important, and so we propose a more sophisticated approach below.

5.4 Hæmophilia A: 'impulse' controller

We turn to 'moderate' hæmophilia A, at which native levels of the zymogen factor VIII are about 2.5% of normal levels. Now, if indeed we have a controller that can be switched on precisely at the onset of a clotting event (perhaps by the release of tissue factor), and can apply the control locally – assumptions under which we have been operating so far – then feedback linearisation is overkill: we can simply dump in the healthy initial concentration of factor VIII in the first controller time step, and then turn the controller off. Using again a control sampling rate of 2 Hz, this technique yields the near-perfect results of Fig. 6. Here the non-zero input at the first time step (not shown) is $x(0)/T$, where $x(0)$ is the desired initial concentration of factor VIII, 0.7 nM, and

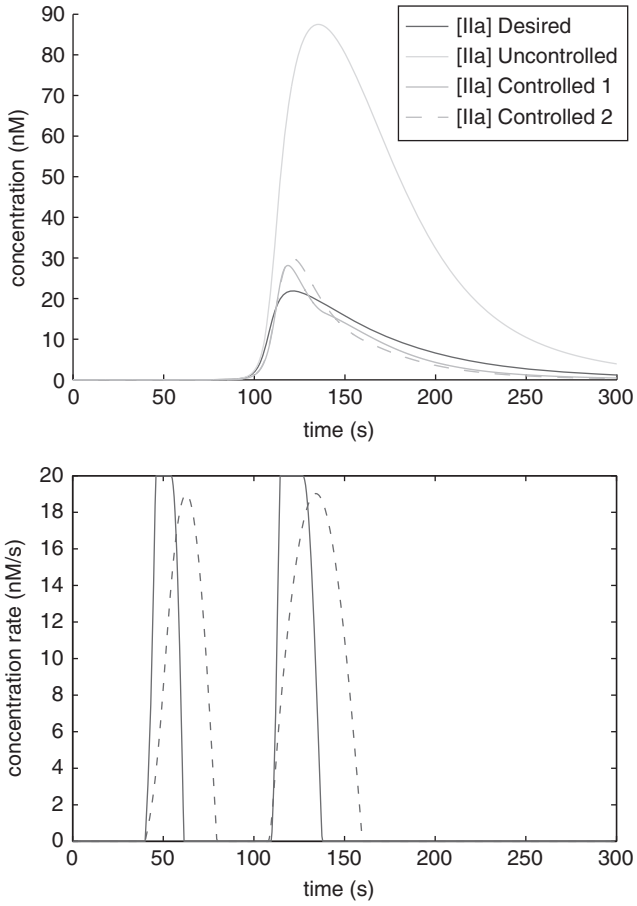

Fig. 5 *Graphs showing*
Top: simulated control of thrombin concentration during a clotting event in a patient with factor-V Leiden, this time using a discrete controller sampling every 0.5 s. The dashed line shows a controller with a high input penalty (relative to the state penalty). Bottom: the discrete input, constrained to lie between 0 and 20 nM/s, that generated the output above. Increasing the relative input penalty avoids 'clipping' at the upper input bound, but degrades tracking

T is the controller time step, 0.5 s. An added boon of this technique is that it requires no sampling at all: the controller would not require sensors.

5.5 Hæmophilia A: step controller

Now, the simple, step controller discussed in Section 3.3 above operates on similar principles, although it makes no assumptions as to the controller operating frequency. Applying the simple learning algorithm of (17) with $\alpha = 1 \times 10^{-4}$ and $\beta = 1 \times 10^{-5}$ yields the trajectory in Fig. 7. The constant input (not shown) of factor VIII is the very low rate of 6.92 pM/s. The price we pay for the simplicity of this control scheme is the overshoot on the back half of the trajectory.

5.6 Factor-V Leiden: step controller

Finally, we ask how well we can manage factor-V Leiden with the step controller. In fact, we can do quite well. Choosing $\alpha = 0.05$ and $\beta = 0.01$ in (17), the input will settle on $u = 4.39$ nM/s, at which rate the peak concentration and its occurrence can be very nearly matched (Fig. 8). The only question that remains is how significant the undershoot is from a clinical perspective.

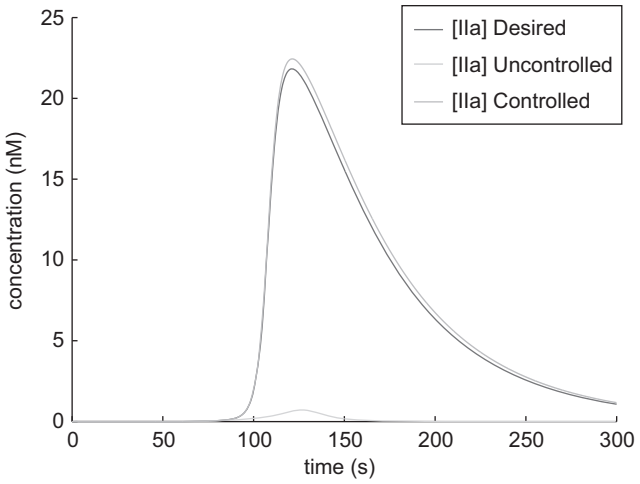

Fig. 6 *Simulated control of thrombin in a patient with moderate hæmophilia A (see text), where the controller inputs the proper initial concentration of factor VIII at the first time step, then shuts off*

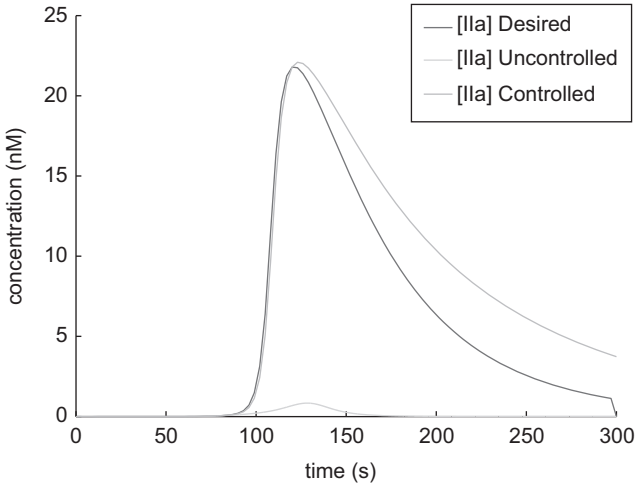

Fig. 7 *Control of thrombin concentration in a hæmophiliac using a single-step input at time $t = 0$*
The input is 6.92 pM/s

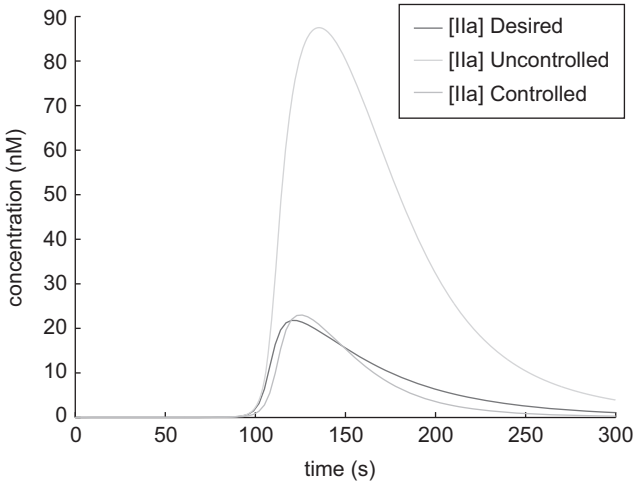

Fig. 8 *Control of thrombin concentration in a patient with factor-V Leiden by a single-step input at time $t = 0$*
The input is 4.39 nM/s

doi: 10.1049/iet-cta.2011.0475

6 Discussion

6.1 Controllability

We have shown via an algebraic criterion that our system of ODEs is not full-state linearisable. As a matter of fact, we can say more. It follows from the method of transforming chemical equations into ODEs that every reversible reaction will introduce into C two columns that are additive inverses of each other. Thus, while the rank of C is obviously upper-bounded by the number of uni-directional reactions, that is, the number of arrows in all our chemical equations, or again the number of rate constants – since this is the number of columns in C, it is also upper-bounded by the (generally smaller) number of uni- or bidirectional reactions, that is, counting each pair of arrows as just one reaction; henceforth simply called 'reactions'. If this number falls short of the dimension of the state space, then full-state linearisation requires the deficit to be covered by the control vector fields – which, in the case of a single input, can provide additional rank of at most one. So in chemical systems of mass-action kinetics with no outflow, a single output, and a single input, a necessary (though insufficient) condition for full-state linearisation is that there be at least as many reactions as state variables.

We can translate this result into a statement about the controllability of the system. The rank of the strong-accessibility distribution (at some point x_0) is the dimension of the locally accessible manifold (from x_0); or alternatively, the difference between this rank and the dimension of the state, n, is the number of uncontrollable modes of the system. So the dimension of the locally accessible manifold is upper-bounded by the total number of reactions (again, counting bidirectional reactions just once). More formally:

Theorem 1: Consider a chemical system of mass-action kinetics with m_u unidirectional reactions, m_b bidirectional reactions, and n states (reactants). Under control with p inputs, the rank q of the strong-accessibility distribution is everywhere less than or equal to the total number of reactions – each bidirectional reaction counting just once –, plus p. That is

$$q \leq m_u + m_b + p \leq n \tag{19}$$

The proof is immediately from the discussion above. That the result is global (rather than at some particular x_0) follows from the 'augmented matrix' D of Section 3.2 being independent of the state. Note also that, although p is obviously increased by the addition of an exogenous pharmaceutical, like heparin, this addition will likewise increase n, both because heparin (say) was not part of the original system, and because it will introduce as well additional complexes (e.g. Hep:AT, Hep:IIa, etc.). It follows that, for an exogenous pharmaceutical to increase the controllability of the system, it must introduce more reactions than complexes.

Considering the present system, we find that without the heparin reactions there are 96 reactions and 98 state variables. Factor-V Leiden knocks out some reactions and some state variables, leaving 77 and 84, respectively. Now, we are obviously interested in pharmaceutical interventions in addition to heparin (a whole host of pro-coagulants, for instance). However, it can now be claimed that in order even to stand a chance of fully linearising the system with one of these interventions as the (sole) input, or equivalently of fully controlling the system, the drug must interact with the system via at least six more reactions than the number of variables that these reactions introduce. (We say six because the control vector field may not be in the span of C and hence cover the remaining dimension; recall (15).) And indeed, heparin introduces only six reactions for its five new state variables, yielding 83 and 89, respectively (see Table 3). So even though there are more unidirectional reactions (rate constants, columns of C) than state variables (proteins, rows of C), we know that there are uncontrollable modes even without constructing C simply by noting that the number of reactions falls short of the number of state variables.

Finally on the topic of full-state linearisation, it can be shown that outflow of the blood factors – in particular, heterogeneous outflow – can increase the rank of the accessibility matrix, and hence of the strong-accessibility distribution [10]. For simplicity, the model as it stands neglects outflow, but in fact *in vivo* coagulation will perforce have some flow (both in and out) of clotting factors, and if it be not negligible over the time scale of interest, this may provide for greater controllability of the system.

6.2 Non-regularity

The coagulation system under application of heparin is, we have said, non-regular; that is, the relative degree is not constant across the state space. What consequences ensue for control?

It appears that the technique of allowing the system to drift away from the singularity has no adverse effects on control. We describe now in detail why this should be so. The feedback non-linearity that shows up in the denominator of the control law, (10), is in the present case, which has strict relative degree ($q = 2$)

$$L_{\mathbf{g}} L_{\mathbf{f}} h(\mathbf{x}) = -k_{\text{on}}[\text{IIa}](t) \tag{20}$$

inducing a singularity precisely when concentrations of thrombin (factor IIa) are zero. However, during a clotting event, thrombin concentrations will always be greater than zero, so the state is safely bounded away from the singularity. We could make this bound precise by choosing some δ in thrombin-concentration space below which we do not care to regulate the thrombin trajectory. And the simulations have demonstrated that in the concentration range of interest – viz., on the order of nanomolarity – the state is not near enough to the singularity to introduce numerical errors into the control.

Now, on the other hand, some (too large) choice of δ would violate the initial-condition criteria too grossly for the controller to be able to recover. That it does recover at all is a consequence (presumably) of the error correction terms added to the control law in (16). Again, the simulations have demonstrated that the controller can indeed recover as late as 60 seconds into the clotting event.

We saw in the Results section above that the unconstrained (continuous) controller led to numerical errors in the ODE simulation – which is why the data shown in Fig. 3 were produced with input clamped at 100 nM/s. These errors are indeed the consequence of the non-regularity of the system, but a peculiar consequence that could not arise in a realistic controller: since the heparin input rate was allowed us to assume negative values, the concentration of heparin itself as well as other state variables could be driven below zero. And of course pushing thrombin levels through zero means operating the controller in the region of the singularity. Lower-bounding the input at zero, as we must in

any realistic controller and as we do in the other simulations, eliminates this threat to the numerical stability of the controller.

Are there not more theoretically sound, less heuristical, bases for avoiding the singularity? There are, but these [12, 13] involve variations on the following theme: compute the relative degree of the system at the singularity, and use the corresponding control law (again (10), but with a higher strict relative degree, q) in the region of the singularity. The problem with this type of technique for our model is that the relative degree at the singularity is higher than the ability of our machine to compute iterated Lie derivatives; that is, we are back to the same, apparently insuperable, computational obstacle that prevented direct computation of the rank of the strong-accessibility distribution. Fortunately, the simulations presented here evidently demonstrate the superfluity of these more sophisticated methods.

6.3 Controller merits and demerits

How precise does our control need to be, after all? The answer is not known, but there is some suggestion that the required amount of thrombin for clotting is much less than the actual peak concentration: Brummel *et al.* [19] have demonstrated in an *in vitro* study that less than 2 nM of thrombin is required for rapid product formation downstream in the cascade. Certainly this suggests that step inputs suffice to steer hæmophilia-A and factor-V-Leiden patients through safe clotting, given the qualitative matches of Figs. 7 and 8. In case these matches are not sufficient, however, or more critically in case the upper bound on the input needs to be lowered, we propose in the following section a model-predictive approach to remedy the defects of the feedback-linearising controller.

However, the step controllers have another enormous advantage over the feedback-linearising control: they require no sampling, only initiation at the inception of a clotting event. There is not currently a method for measuring blood-protein concentrations *in vivo* in real time – certainly not at the required 2 Hz sampling rate. The term $L_{\mathbf{f}}^2 h(\mathbf{x})$ from the control law, (10), is a function of some 32 proteins. (The other term in the control law, $L_{\mathbf{g}} L_{\mathbf{f}} h(\mathbf{x})$, is a function only of thrombin. And in fact, thrombin can indeed be measured in real-time *ex vivo*; see [20].) We interpret our results on feedback linearisation, then, as showing that if indeed a greater degree of accuracy is required for thrombin tracking, then it would be fruitful to devise the appropriate sensors for real-time sampling, since they would make possible this suitable technique. Now, naïvely constraining the input vitiates the accuracy of this technique (Figs. 4 and 5), hence the need for the model-predictive controller.

On the other hand, two significant advantages of the feedback controllers must be stressed. First of all, they require no training. Patient-to-patient variation and model errors (e.g. in rate constants, which vary fairly widely from study to study) make pre-training on a model insufficient for the development of actual treatments, though it certainly provides a good starting point. So in the limit, the training of the step-controller approaches the actual (barbarous) state of the art, namely testing the patient once a day and modifying doses accordingly. Second, and perhaps more significantly, the feedback controllers operate fairly robustly over a wide range of delays in turning on. So, for example, in the case of constrained, discrete control (the most plausible scenario), turning on the controller a full minute after the release of tissue factor increases the overshoot by less than 5%, with the rest of the trajectory tracking essentially identical (data not shown). The issue here is how similar the state at (for example) $t = 60$ s is to the desired initial conditions – the closer, the more easily the controller can recover – and as is clear from the plots above, thrombin concentrations (presumably, *inter alia*) do not increase significantly until about 85 s.

One envisages, then, the following real-world scenario in which this type of controller would be particularly useful: an individual suffering from factor-V Leiden cuts herself, and – within the first minute – applies to the site of injury a device that injects heparin, in response to the local concentration of all (or some subset of) the other blood proteins.

6.4 Significance of model assumptions

Two model assumptions require elaboration. First, both the feedback and step controllers show some sensitivity to the choice of rate constants. An earlier version of the model, for instance, estimated both on- and off-rates for the heparin reactions, again setting the former to 0.1 nM/s but also fixing the latter at $10\,\mathrm{s}^{-1}$, and achieved essentially perfect tracking with the same input limits and sampling rate. As we have said repeatedly, reported rate constants in the blood clotting literature vary somewhat widely, so this consideration is quite relevant. However, even if we assume the affinities in [14] are correct, as in the present model, there is still the matter of our assumption of the on-rates (recall that the unknown heparin on- and off-rates were estimated from their ratio, fixing the on-rate in the middle of the reasonable physiological range at 0.1 nM/s and then computing the off-rate from their known ratio; see Section 4).

The effect on the controller turns out to be quite small. Similar tracking results (not shown) can be achieved by both controllers for rate-constant values that span three orders of magnitude, with k_{on} from 0.01–1 nM/s and k_{off} set to maintain the desired ratio. Values outside this range are unlikely, but there the step controller can match the thrombin peak's timing and height only on pain of greatly distorting the subsequent trajectory.

A related question is how sensitive the control is to *unknown* deviations in the rate constants. Given the heavy dependence of the control law on the state, it will not be surprising that this sensitivity is indeed acute. Simulations (not shown for reasons of space) reveal that even minor perturbations of the model rate constants – leaving the control-law constants at their nominal values – degrades the tracking of thrombin intolerably. Nor is the step controller immune from this problem: a suitable step input was found by optimisation over repeated runs of the model; no such luxury is available for the patient, so if model and patient ('plant') differ, the step controller will issue in poor control.

Nevertheless, the degeneration of tracking can be eliminated by increasing the gain on the proportional term of the feedback-linearisation controller. This is true for rate-constant deviations of up to 5%. (In the simulations, rate constants were randomly perturbed by multiplicative Gaussian noise with unity mean and standard deviation $1/3 \times 5\%$, keeping 99% of the rate constants within 5% of their original values.) The insensitivity of the model to high proportional-term gains is of itself interesting; we pursue it at length in a companion article.

The final model assumption adverted to at the beginning of this section is that of locality: the supply of unactivated (zymogen) blood factors was treated as limited. This choice was made to conform to the model of [3] as well as to the

doi: 10.1049/iet-cta.2011.0475

in vitro results of [16]. In reality, however, zymogens are replenished by circulating blood, and activated proteins are similarly removed. On the other hand, the limited amount of lipids also restricts the amount of each zymogen that can ever take part in the reaction, since nearly all of the clotting reactions take place on a phospholipid surface. Congruently, simulations (not shown) indicate that even if the zymogens are modelled as inexhaustible, the thrombin curve remains qualitatively the same. We proceeded then with the locality assumption, in order to avoid the complexities of an added circulation model.

7 Conclusions

Regulation of many clotting disorders requires daily visits to a doctor both for the administration of blood thickeners or thinners, and for assaying blood–protein concentrations to determine dosages. In exceptional circumstances, continuous intravenous administration is necessary. Neither of these alternatives is desirable, and both are in some sense a consequence of taking into account only the most primitive knowledge about human coagulation: certain drugs encourage clotting, whereas others discourage it.

This paper has proposed and demonstrated two control techniques that are much more powerful in virtue of exploiting a mathematical model of a large portion of the coagulation cascade and the effects of the relevant pharmaceuticals. Both controllers assume local application and the ability to detect a clotting event (release of tissue factor); under which assumptions both have been shown capable of regulating the time-varying concentrations of thrombin, the key blood protein which alone of the proteins in this model determines clot formation.

Each controller has its merits and demerits: the step controller is more critically dependent on the accuracy of the model but requires no sensors; whereas the feedback controller is robust to model changes but requires real-time (of the order of 2 Hz) sampling of a number of protein concentrations (viz., all those involved in the feedback terms). Feedback control is also (slightly) degraded by constraining the input rate to lie within a feasible range. We address both of these deficiencies of feedback linearisation in a companion article (although not, unfortunately, at the same time), demonstrating a feedback controller that requires observation of only one state variable, and a different (model-predictive) controller, which overcomes some of the constraint-induced tracking degradation.

The significance of both the controllers presented in this paper for clinical application would be greatly advanced by various improvements to the model. First, as lately noted, the rate constants for heparin interactions with the clotting process are not known; if a precise controller is to be constructed along the lines proposed here, these must be determined. Second, a circulation model would dispense with the somewhat dubious locality assumption, as well as possibly afford more powerful control, since inflow and outflow can increase the strict relative degree of the system (and hence the number of linearised variables that can be controlled).

Finally, this paper has made some useful general observations about chemical kinetics and controllability. In particular, we have shown that the dimension of the locally accessible manifold is upper-bounded by the number of reactions in a chemical system, where bi-directional reactions are counted just once. The authors have not seen this result elsewhere. We can say then, e.g. that since full-state linearisation requires that the dimension of the locally accessibly manifold be equal to the dimension of the state, in chemical systems it requires that the number of chemicals not exceed the number of reactions.

8 References

1 Makin, J.G., Narayanan, S.: 'A hybrid-system model of human blood clotting'. Technical Report, International Computer Science Institute, 1947 Center Street, Berkeley, CA 94703, 2008

2 Mitchell, I., Bayen, A.M, Tomlin, C.J.: 'Validating a Hamilton–Jacobi approximation to hybrid system reachable sets', in Benedetto M.D.D., Sangiovanni-Vincentelli A.L. (Eds.): Fourth Int. Workshop Hybrid Systems: Computation and Control, HSCC 2001, Rome, Italy, 28–30 March 2001, (*LNCS*, **2034**), 2001, pp. 418–432

3 Bungay, S., Gentry, P., Gentry, R.: 'A mathematical model of lipid mediated thrombin generation', *Math. Med. Biol.*, 2003, **20**, pp. 105–129

4 Hockin, M.F., Cawthern, K.M., Kalafatis, M., Mann, K.G.: 'A model describing the inactivation of factor Va by APC: bond cleavage, fragment dissociation, and product inhibition', *Biochemistry*, 1999, **38**, pp. 6918–6934

5 Makin, J.G.: 'A computational model of human blood clotting'. Technical Report, International Computer Science Institute, 1947 Center Street, Berkeley, CA, 94703, 2008

6 Sastry, S.S.: 'Nonlinear systems: analysis, stability, and control' (Springer–Verlag, New York, 1999)

7 Isidori, A.: 'Nonlinear control systems' (Springer, 1995, 3rd edn.)

8 Adcock, D.M., Jensen, R., Johns, C.S.: 'Coagulation handbook' (Esoterix Coagulation, 2002)

9 Khanin, M.A., Rakov1, D.V., Kogan, A.E.: 'Mathematical model for the blood coagulation prothrombin time test', *Thrombosis Res.*, 1998, **89**, pp. 227–232

10 Bastin, G, Lévine, J.: 'On state accessibility in reaction systems', *IEEE Trans. Autom. Control*, 1993, **38**, (5), pp. 733–742

11 Henson, M.A., Seborg, D.E.: 'Feedback linearizing control', in Seborg D.E., Henson M.A. (Eds.): Nonlinear Process Control, Prentice-Hall, Upper Saddle River, NJ, 1996, pp. 149–231

12 Tomlin, C.J., Sastry, S.S.: 'Switching through singularities'. Proc. 36th Conf. Decision and Control, 1997, pp. 1–6

13 Hauser, J., Sastry, S., Kokotovic, P.: 'Nonlinear control via approximate input–output linearization: the ball and beam example'. Proc. 28th Conf. Decision and Control, Tampa, Florida, 1989, pp. 392–398

14 Olson, S.T.: 'Transient kinetics of heparin-catalyzed protease inactivation by antithrombin III', *J. Biol. Chem.*, 1988, **263**, (4), pp. 1698–1708

15 Egan, J., Kalafatis, M., Mann, K.G.: 'The effect of $Arg^{306} \rightarrow$ Ala and $Arg^{506} \rightarrow$ Gln substitutions in the inactivation of recombinant human factor VA by activated protein C and protein S', *Protein Sci.*, 1997, **6**, pp. 2016–2027

16 Butenas, S., van't Veer, C., Mann, K.G.: ' "Normal" thrombin generation', *Blood*, 1999, **94**, pp. 2169–2178

17 Fersht, A.: 'Structure and mechanism in protein science' (W.H. Freeman and Company, New York, 1999)

18 MATLAB. version 7.1 (R14SP3). Natick, Massachusetts: the MathWorks Inc., 2005

19 Brummel, K., Paradis, S., Butenas, S., Mann, K.: 'Thrombin functions during tissue factor-induced blood coagulation', *Blood*, 2002, **100**, (1), pp. 148–152

20 Hemker, H.C., Giesen, P., AlDieri1, R., *et al.* 'The calibrated automated thrombogram(CAT): a universal routine test for hyper- and hypocoagulability', *Pathophysiol. Haemostasis Thrombosis*, 2002, **32**, pp. 249–253