

Research Article

Massive Exploration of Perturbed Conditions of the Blood Coagulation Cascade through GPU Parallelization

Paolo Cazzaniga,^{1,2,3} Marco S. Nobile,^{2,4} Daniela Besozzi,^{2,3,5}
Matteo Bellini,⁴ and Giancarlo Mauri^{2,4}

¹ Dipartimento di Scienze Umane e Sociali, Università degli Studi di Bergamo, 24129 Bergamo, Italy

² SYSBIO Centre for Systems Biology, 20126 Milano, Italy

³ Istituto di Analisi dei Sistemi ed Informatica “Antonio Ruberti”, Consiglio Nazionale delle Ricerche, 00185 Roma, Italy

⁴ Università degli Studi di Milano-Bicocca, Dipartimento di Informatica, Sistemistica e Comunicazione, 20126 Milano, Italy

⁵ Università degli Studi di Milano, Dipartimento di Informatica, 20135 Milano, Italy

Correspondence should be addressed to Paolo Cazzaniga; paolo.cazzaniga@unibg.it

Received 21 February 2014; Revised 7 May 2014; Accepted 18 May 2014; Published 16 June 2014

Academic Editor: Horacio Pérez-Sánchez

Copyright © 2014 Paolo Cazzaniga et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The introduction of general-purpose Graphics Processing Units (GPUs) is boosting scientific applications in Bioinformatics, Systems Biology, and Computational Biology. In these fields, the use of high-performance computing solutions is motivated by the need of performing large numbers of *in silico* analysis to study the behavior of biological systems in different conditions, which necessitate a computing power that usually overtakes the capability of standard desktop computers. In this work we present coagSODA, a CUDA-powered computational tool that was purposely developed for the analysis of a large mechanistic model of the blood coagulation cascade (BCC), defined according to both mass-action kinetics and Hill functions. coagSODA allows the execution of parallel simulations of the dynamics of the BCC by automatically deriving the system of ordinary differential equations and then exploiting the numerical integration algorithm LSODA. We present the biological results achieved with a massive exploration of perturbed conditions of the BCC, carried out with one-dimensional and bi-dimensional parameter sweep analysis, and show that GPU-accelerated parallel simulations of this model can increase the computational performances up to a 181× speedup compared to the corresponding sequential simulations.

1. Introduction

Mathematical modeling and computational analysis of biological systems nowadays represent an essential methodology, complementary to conventional experimental biology, to achieve an in-depth comprehension of the functioning of these complex systems [1, 2]. Given a model that describes the physical or logical interactions between the components of a biological system, different algorithms can be exploited to make predictions on the way this system behaves in both physiological and perturbed conditions. For instance, starting from distinct parameterizations of the model, simulation algorithms can be used to devise the different emergent behaviors that the system can present; the massive exploration of high-dimensional parameter spaces allows us to

better understand the system functioning across a wide spectrum of natural conditions, as well as to derive statistically meaningful properties. Indeed, standard investigations of biological systems usually rely on computational methods that require the execution of a large number of simulations, such as parameter sweep analysis [3], sensitivity analysis [4], structure and parameter identifiability [5], parameter estimation [6–8], and reverse engineering of model topologies [9–13].

In this context, the use of general-purpose Graphics Processing Units (GPUs) has recently boosted many applications in scientific computing, where CPUs have traditionally been the standard workhorses. As a matter of fact, when several batches of simulations need to be executed, the necessary computing power can rapidly overtake the capabilities

of standard desktop computers, therefore requiring high-performance computing solutions. After the introduction of general-purpose GPUs and of Compute Unified Device Architecture (CUDA, Nvidia's GPU programming language), the adoption of these graphics engines largely increased in research fields as Bioinformatics, Systems Biology, and Computational Biology (see an overview in [14–16]). Anyway, despite the remarkable advantages concerning the computational speedup, computing with GPUs usually requires the development and the implementation of *ad hoc* algorithms, since GPU-based programming substantially differs from CPU-based computing; as a consequence, scientific applications of GPUs might undergo the risk of remaining a niche for few specialists. To avoid such limitations, several packages and software tools have recently been released (see, e.g., [16–18]), so that also users with no knowledge of GPUs hardware and programming can access the high-performance computing power of graphics engines.

To investigate the dynamics of biological systems, either deterministic or stochastic approaches can be exploited [19], which are based on numerical integration (e.g., Euler's or Runge-Kutta methods [20]) or on Markov processes (e.g., Gillespie's algorithm [21]), respectively. To date, the most efficient algorithms to integrate a system of ordinary differential equations (ODEs), or to perform stochastic simulations of reaction-based models, are LSODA [22] and tau-leaping [23], respectively. In [24–26] we previously presented cupSODA and cuTauLeaping, the GPU-powered implementation of LSODA and tauleaping, respectively. cupSODA allows to run parallel deterministic simulations of a given mass-action based system of biochemical reactions, using the LSODA algorithm; cuTauLeaping represents a novel restructuring of the tau-leaping workflow that fits the GPU architecture and avoids any inefficiency drawback for coarse-grain massive parallel stochastic simulations.

In this work we introduce coagSODA, an extension of cupSODA that was specifically designed for the analysis of a model of the *blood coagulation cascade* (BCC). Blood is the subject of an intense scientific research, thanks to its key role in making diagnosis of numerous diseases [27]. Humans have evolved a complex hemostatic system that is able to maintain blood in a fluid state and allow the circulation through an intricate network of vessels; in particular, the presence of several fine-tuned feedback mechanisms in the BCC allows keeping all blood components within appropriate concentration ranges. The BCC consists in a complex network of cellular reactions which, under physiological conditions *in vivo*, are inhibited by the presence of intact endothelium [28]. Anyway, in response to any vascular injury, the hemostatic system is able to stop the blood leakage by rapidly sealing the defects in the vessels' wall [29].

In order to investigate the variations in blood coagulation components among individuals and to understand the corresponding response of the system to perturbed conditions, we consider here a computational perspective to study the BCC; we analyze the alterations (prolongation or reduction) of the time required to form the clot (i.e., the *clotting time*) by exploiting a reduced version of the mathematical model defined in [30]. This model describes

the intrinsic, extrinsic, and common pathways of the BCC and, more importantly, it accounts for platelets activation, as well as the presence of several inhibitors (e.g., the Tissue Factor Pathway Inhibitor, antithrombin III, and C1-inhibitor) [31].

Numerous mathematical models of blood coagulation were developed in the last years, as they represent a useful tool for systematic studies of the intricate network of the coagulation cascade and allow obtaining a suitable reconstruction of empirical observations (see, e.g., [32–34]). The earliest models considered only simple steps of the whole BCC, such as the conversion of the clot fibrin by thrombin [35]; lately, Hockin et al. [36] developed a comprehensive ODE-based model of the extrinsic blood coagulation system. This model was then considered as reference by several research groups to investigate the thrombotic risk in healthy and ill populations [37, 38], or to understand other complex biochemical processes: for instance, in [39] the roles of protein C, protein S, and phospholipid surface actions were considered, while in [40] the influence of trace amounts of key coagulation proteases on thrombin generation was investigated. Recently, other works modeling the blood clotting process in a comprehensive manner have been published; besides the already mentioned model developed by Chatterjee et al. [30], we mention the model defined by Wajima et al. [32], which simulates the intrinsic, extrinsic, and common pathways, the vitamin K cycle, the therapy with the anticoagulant drugs warfarin and heparin, and the laboratory tests PT and aPTT, as well as the Taipan snake bite, which causes coagulopathy.

coagSODA, the GPU-accelerated simulator that we present and exploit in this work for the analysis of the BCC model, is a user-friendly and efficient tool that circumvents the need of manually defining the system of ODEs that describe the blood coagulation network. More precisely, coagSODA is able to automatically derive the system of (mass-action and Hill function-based) ODEs and then perform their numerical integration starting from the given set of 96 biochemical reactions, which fully describe the molecular interactions between all the species involved in the BCC *in vivo*. We show that coagSODA allows us to efficiently execute a large number of parallel deterministic simulations of the BCC at a considerable reduced computational cost with respect to CPUs. In particular, we exploit coagSODA to carry out one-dimensional and bi-dimensional parameter sweep analysis of the BCC, to the purpose of investigating the prolongation and the reduction of the clotting time in response to perturbed values of some reaction constants and of the initial concentration of some molecular species, chosen according to their meaning within the whole pathway.

The paper is structured as follows. In Section 2 we fully describe the mechanistic model of the BCC used in this work and present the simulation method at the basis of cupSODA tool to introduce the coagSODA simulator. In Section 3 we present the results obtained from the parameter sweep analysis of the BCC model, as well as a comparison of the computational performance of coagSODA with respect to a CPU-based implementation of LSODA. Finally, in Section 4 we conclude the paper with a discussion of the presented work and future research directions.

2. Materials and Methods

2.1. A Mechanistic Model of the Blood Coagulation Cascade. Blood is an essential component in human life, whose primary functions are to feed cells by delivering a multitude of nutrients, such as oxygen, and to carry away the cellular wastes, such as carbon dioxide. Specialized cells and fluids in blood perform many physiological functions and can be isolated and analyzed through specific laboratory tests, giving the opportunity to settle a person's health condition. All blood components are kept within appropriate concentration ranges by means of fine-tuned regulatory mechanisms, ruled by several feedback controls; the constancy of blood composition is maintained thanks to the circulation through an intricate network of vessels. In particular, humans evolved a complex hemostatic system that, under physiological conditions, maintains blood in a fluid state; however, in response to any vascular injury, this system is able to rapidly react and seal the defects in the vessels' wall in order to stop the blood leakage [29]. Indeed, the circulatory system is self-sealing; otherwise, a continuous blood flow from even the smallest wound would become a threat for the individual's life.

To allow blood coagulation, in humans there exist 13 blood clotting proteins, called *coagulation factors*, which are usually designated by Roman numerals I through XIII. As a consequence of a vascular injury, platelets become active and the Tissue Factor (TF, also called factor III) is exposed in the subendothelial tissue, starting the *blood coagulation cascade* (BCC). The ultimate goal of the BCC is to convert prothrombin (factor II) into thrombin (factor IIa—i.e., the active factor II), the enzyme that catalyzes the formation of a clot. Traditionally, the BCC is divided into the *extrinsic* and *intrinsic pathways*, both leading to the activation of factor X [41]. The last part of the cascade, downstream of this factor, is called the *common pathway* and leads to the formation of fibrin monomers, whose polymers finally constitute the backbone of the clot.

Excluding thrombin, all the enzymes involved in blood clotting are characterized by a low activity, which increases upon binding to a specific protein cofactor (e.g., factors V and VIII) or to appropriate phospholipid surfaces (e.g., the plasma membranes of active platelets) [41]. Even calcium ions have a central role in coagulation, since they are essential to start and enhance numerous reactions; without calcium ions the blood coagulation cannot occur [42]. In the BCC pathways, the activity of the various active proteases is limited by several inhibitory factors, which allow regulating the whole cascade. When the hemostatic system is unregulated, thrombosis (i.e., the formation of a blood clot obstructing the blood flow in vessels) may occur due to impairment in the inhibitory pathway, or because the functioning of the natural anticoagulant processes is overwhelmed by the strength of the hemostatic stimulus [29].

The BCC model we consider in this work is a slightly reduced version of the “Platelet-Plasma” deterministic model defined in [30], built upon a previous model [36], which describes all parts of blood coagulation: the platelets activation and aggregation; the extrinsic, intrinsic and common pathways (with the exception of factor XIII);

the action of several inhibitory molecules (Tissue Factor Pathway Inhibitor, antithrombin III, C1-inhibitor, α 1-antitrypsin, and α 2-antiplasmin). In addition, to simulate the coagulation process *in vitro*, Chatterjee et al. also modeled the action of corn trypsin inhibitor (CTI), which inhibits the activation of the so-called contact system, as well as the action of the fluorogenic substrate Boc-VPR-MCA widely used in laboratories for thrombin titration [30]. The role of calcium was not explicitly included in the model but considered as a nonlimiting factor, as in living organisms this ion very rarely drops to such low levels able to alter the kinetics of the formation of the fibrin clot [43].

Since the aim of this work is the investigation of blood coagulation *in vivo*, we exclude a small set of reactions given in [30] that have no effect on the clotting time. To be more precise, we do not consider the reactions occurring *in vitro* (namely, entries 28 and 35 in Table 1 in [30]), as well as the reactions downstream the fibrinogen conversion, that is, the interactions between the fibrin polymers, thrombin, and antithrombin III (namely, entries 55, 56, and 57 in Table 1 in [30]).

For the sake of completeness, we describe in Table 1 the BCC model considered in this work, overall consisting in 96 reactions among 71 molecular species. A graphical sketch of the main molecular interactions among the BCC components is given in Figure 1.

The model can be partitioned into four functional modules.

- (1) The first module corresponds to the *extrinsic pathway*, which consists in
 - (i) the formation of a complex between Tissue Factor and factor VII (modeled by reactions r_1, \dots, r_4);
 - (ii) the activation of factor VII by the complex between Tissue Factor and factor VIIa (reaction r_5);
 - (iii) the activation of factor IX by the complex between Tissue Factor and factor VIIa (reactions r_{13}, r_{14}, r_{15}) and by factor VIIa (reactions r_{78}, r_{79}, r_{80});
 - (iv) the activation of factor X by the complex between Tissue Factor and factor VIIa (r_8, \dots, r_{12}) and by factor VIIa (r_{81}, r_{82}, r_{83}).
- (2) The second module corresponds to the *intrinsic pathway*, which consists in
 - (i) the formation of a complex between factor VIIIa and factor IXa (modeled by reactions r_{18} and r_{19});
 - (ii) the activation of factor X by the complex between factor VIIIa and factor IXa (reactions r_{20}, r_{21}, r_{22}) and by factor IXa (reactions r_{72}, r_{73}, r_{74});
 - (iii) the activation of factor XII by factor XII itself (reaction r_{44}), by factor XIIa (reactions r_{45}, r_{46}, r_{47}), and by kallikrein (reactions r_{51}, r_{52}, r_{53});

TABLE 1: Reaction-based model of the blood coagulation cascade (reduced version of the “Platelet-Plasma” model described in [30]). The model consists in 96 reactions among 71 molecular species. With the exception of reaction r_{44} (see Section 3.1), the values of all reaction constants were taken from [30].

n	Reactants	Products	Constant (k)
r_1	TF + fVII	TF-fVII	$3.20 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$
r_2	TF-fVII	TF + fVII	$3.10 \cdot 10^{-2} \text{ s}^{-1}$
r_3	TF + fVIIa	TF-fVIIa	$2.30 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_4	TF-fVIIa	TF + fVIIa	$3.10 \cdot 10^{-5} \text{ s}^{-1}$
r_5	TF-fVIIa + fVII	TF-fVIIa + fVIIa	$4.40 \cdot 10^5 \text{ M}^{-1} \text{ s}^{-1}$
r_6	fXa + fVII	fXa + fVIIa	$1.30 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_7	fIIa + fVII	fIIa + fVIIa	$2.30 \cdot 10^4 \text{ M}^{-1} \text{ s}^{-1}$
r_8	TF-fVIIa + fX	TF-fVIIa-fX	$2.50 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_9	TF-fVIIa-fX	TF-fVIIa + fX	$1.05 \cdot 10^{-2} \text{ s}^{-1}$
r_{10}	TF-fVIIa-fX	TF-fVIIa-fXa	6.00 s^{-1}
r_{11}	TF-fVIIa-fXa	TF-fVIIa + fXa	19.00 s^{-1}
r_{12}	TF-fVIIa + fXa	TF-fVIIa-fXa	$2.20 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{13}	TF-fVIIa + fIX	TF-fVIIa-fIX	$1.00 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{14}	TF-fVIIa-fIX	TF-fVIIa + fIX	2.40 s^{-1}
r_{15}	TF-fVIIa-fIX	TF-fVIIa + fIXa	1.80 s^{-1}
r_{16}	fII + fXa	fIIa + fXa	$7.50 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{17}	fIIa + fVIII	fIIa + fVIIIa	$2.00 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{18}	fVIIIa + fIXa	fIXa-fVIIIa	$1.00 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{19}	fIXa-fVIIIa	fVIIIa + fIXa	$1.00 \cdot 10^{-4} \text{ s}^{-1}$
r_{20}	fIXa-fVIIIa + fX	fIXa-fVIIIa-fX	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{21}	fIXa-fVIIIa-fX	fIXa-fVIIIa + fX	$1.00 \cdot 10^{-5} \text{ s}^{-1}$
r_{22}	fIXa-fVIIIa-fX	fIXa-fVIIIa + fXa	8.20 s^{-1}
r_{23}	fVIIIa	fVIIIa ₁ -L + fVIIIa ₂	$6.00 \cdot 10^{-5} \text{ s}^{-1}$
r_{24}	fVIIIa ₁ -L + fVIIIa ₂	fVIIIa	$2.20 \cdot 10^4 \text{ M}^{-1} \text{ s}^{-1}$
r_{25}	fIXa-fVIIIa-fX	fVIIIa ₁ -L + fVIIIa ₂ + fX + fIXa	$1.00 \cdot 10^{-3} \text{ s}^{-1}$
r_{26}	fIXa-fVIIIa	fVIIIa ₁ -L + fVIIIa ₂ + fIXa	$1.00 \cdot 10^{-3} \text{ s}^{-1}$
r_{27}	fIIa + fV	fIIa + fVa	$2.00 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{28}	fXa + fVa	fXa-fVa	$4.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{29}	fXa-fVa	fXa + fVa	0.2 s^{-1}
r_{30}	fXa-fVa + fII	fXa-fVa-fII	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{31}	fXa-fVa-fII	fXa-fVa + fII	103.00 s^{-1}
r_{32}	fXa-fVa-fII	fXa-fVa + fmIIa	63.50 s^{-1}
r_{33}	fXa-fVa + fmIIa	fXa-fVa + fIIa	$1.50 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{34}	fXa + TFPI	fXa-TFPI	$9.00 \cdot 10^5 \text{ M}^{-1} \text{ s}^{-1}$
r_{35}	fXa-TFPI	fXa + TFPI	$3.60 \cdot 10^{-4} \text{ s}^{-1}$
r_{36}	TF-fVIIa-fXa + TFPI	TF-fVIIa-fXa-TFPI	$3.20 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{37}	TF-fVIIa-fXa-TFPI	TF-fVIIa-fXa + TFPI	$1.10 \cdot 10^{-2} \text{ s}^{-1}$
r_{38}	TF-fVIIa + fXa-TFPI	TF-fVIIa-fXa-TFPI	$5.00 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$
r_{39}	fXa + ATIII	fXa-ATIII	$1.50 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{40}	fmIIa + ATIII	fmIIa-ATIII	$7.10 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{41}	fIXa + ATIII	fIXa-ATIII	$4.90 \cdot 10^2 \text{ M}^{-1} \text{ s}^{-1}$
r_{42}	fIIa + ATIII	fIIa-ATIII	$7.10 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{43}	TF-fVIIa + ATIII	TF-fVIIa-ATIII	$2.30 \cdot 10^2 \text{ M}^{-1} \text{ s}^{-1}$
r_{44}	fXII	fXIIa	$5.00 \cdot 10^{-4} \text{ s}^{-1}$
r_{45}	fXIIa + fXII	fXIIa-fXII	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{46}	fXIIa-fXII	fXIIa + fXII	750.00 s^{-1}
r_{47}	fXIIa-fXII	fXIIa + fXIIa	$3.30 \cdot 10^{-2} \text{ s}^{-1}$
r_{48}	fXIIa + PKal	fXIIa-PKal	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{49}	fXIIa-PKal	fXIIa + PKal	$3.60 \cdot 10^3 \text{ s}^{-1}$

TABLE I: Continued.

n	Reactants	Products	Constant (k)
r_{50}	fXIIa-PKal	fXIIa + Kal	40.00 s^{-1}
r_{51}	fXII + Kal	fXII-Kal	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{52}	fXII-Kal	fXII + Kal	45.30 s^{-1}
r_{53}	fXII-Kal	fXIIa + Kal	5.70 s^{-1}
r_{54}	PKal + Kal	Kal + Kal	$2.70 \cdot 10^4 \text{ M}^{-1} \text{ s}^{-1}$
r_{55}	Kal	Kal _i	$1.10 \cdot 10^{-2} \text{ s}^{-1}$
r_{56}	fXIIa + Clinh	fXIIa-Clinh	$3.60 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{57}	fXIIa + ATIII	fXIIa-ATIII	$21.60 \text{ M}^{-1} \text{ s}^{-1}$
r_{58}	fXI + fIIa	fXI-fIIa	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{59}	fXI-fIIa	fXI + fIIa	5.00 s^{-1}
r_{60}	fXI-fIIa	fXIa + fIIa	$1.30 \cdot 10^{-4} \text{ s}^{-1}$
r_{61}	fXIIa + fXI	fXIIa-fXI	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{62}	fXIIa-fXI	fXIIa + fXI	200 s^{-1}
r_{63}	fXIIa-fXI	fXIIa + fXIa	$5.70 \cdot 10^{-4} \text{ s}^{-1}$
r_{64}	fXIa + fXI	fXIa + fXIa	$3.19 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$
r_{65}	fXIa + ATIII	fXIa-ATIII	$3.20 \cdot 10^2 \text{ M}^{-1} \text{ s}^{-1}$
r_{66}	fXIa + Clinh	fXIa-Clinh	$1.80 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{67}	fXIa + A1AT	fXIa-A1AT	$1.00 \cdot 10^2 \text{ M}^{-1} \text{ s}^{-1}$
r_{68}	fXIa + A2AP	fXIa-A2AP	$4.3 \cdot 10^3 \text{ M}^{-1} \text{ s}^{-1}$
r_{69}	fXIa + fIX	fXIa-fIX	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{70}	fXIa-fIX	fXIa + fIX	41.00 s^{-1}
r_{71}	fXIa-fIX	fXIa + fIXa	7.70 s^{-1}
r_{72}	fIXa + fX	fIXa-fX	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{73}	fIXa-fX	fIXa + fX	0.64 s^{-1}
r_{74}	fIXa-fX	fIXa + fXa	$7.00 \cdot 10^{-4} \text{ s}^{-1}$
r_{75}	fXa + fVIII	fXa-fVIII	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{76}	fXa-fVIII	fXa + fVIII	2.10 s^{-1}
r_{77}	fXa-fVIII	fXa + fVIIIa	0.023 s^{-1}
r_{78}	fVIIa + fIX	fVIIa-fIX	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{79}	fVIIa-fIX	fVIIa + fIX	0.90 s^{-1}
r_{80}	fVIIa-fIX	fVIIa + fIXa	$3.60 \cdot 10^{-5} \text{ s}^{-1}$
r_{81}	fVIIa + fX	fVIIa-fX	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{82}	fVIIa-fX	fVIIa + fX	210.00 s^{-1}
r_{83}	fVIIa-fX	fVIIa + fXa	$1.60 \cdot 10^{-6} \text{ s}^{-1}$
r_{84}	Fbg + fIIa	Fbg-fIIa	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{85}	Fbg-fIIa	Fbg + fIIa	636.00 s^{-1}
r_{86}	Fbg-fIIa	Fbn1 + fIIa + FPA	84.00 s^{-1}
r_{87}	Fbn1 + fIIa	Fbn1-fIIa	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{88}	Fbn1-fIIa	Fbn1 + fIIa	742.60 s^{-1}
r_{89}	Fbn1-fIIa	Fbn2 + fIIa + FPB	7.40 s^{-1}
r_{90}	Fbn1 + Fbn1	(Fbn1) ₂	$1.00 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$
r_{91}	(Fbn1) ₂	2Fbn1	$6.40 \cdot 10^{-2} \text{ s}^{-1}$
r_{92}	(Fbn1) ₂ + fIIa	(Fbn1) ₂ -fIIa	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{93}	(Fbn1) ₂ -fIIa	(Fbn1) ₂ + fIIa	701.00 s^{-1}
r_{94}	(Fbn1) ₂ -fIIa	(Fbn2) ₂ + fIIa + FPB	49.00 s^{-1}
r_{95}	Fbn2 + fIIa	Fbn2-fIIa	$1.00 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$
r_{96}	Fbn2-fIIa	Fbn2 + fIIa	$1.00 \cdot 10^3 \text{ s}^{-1}$

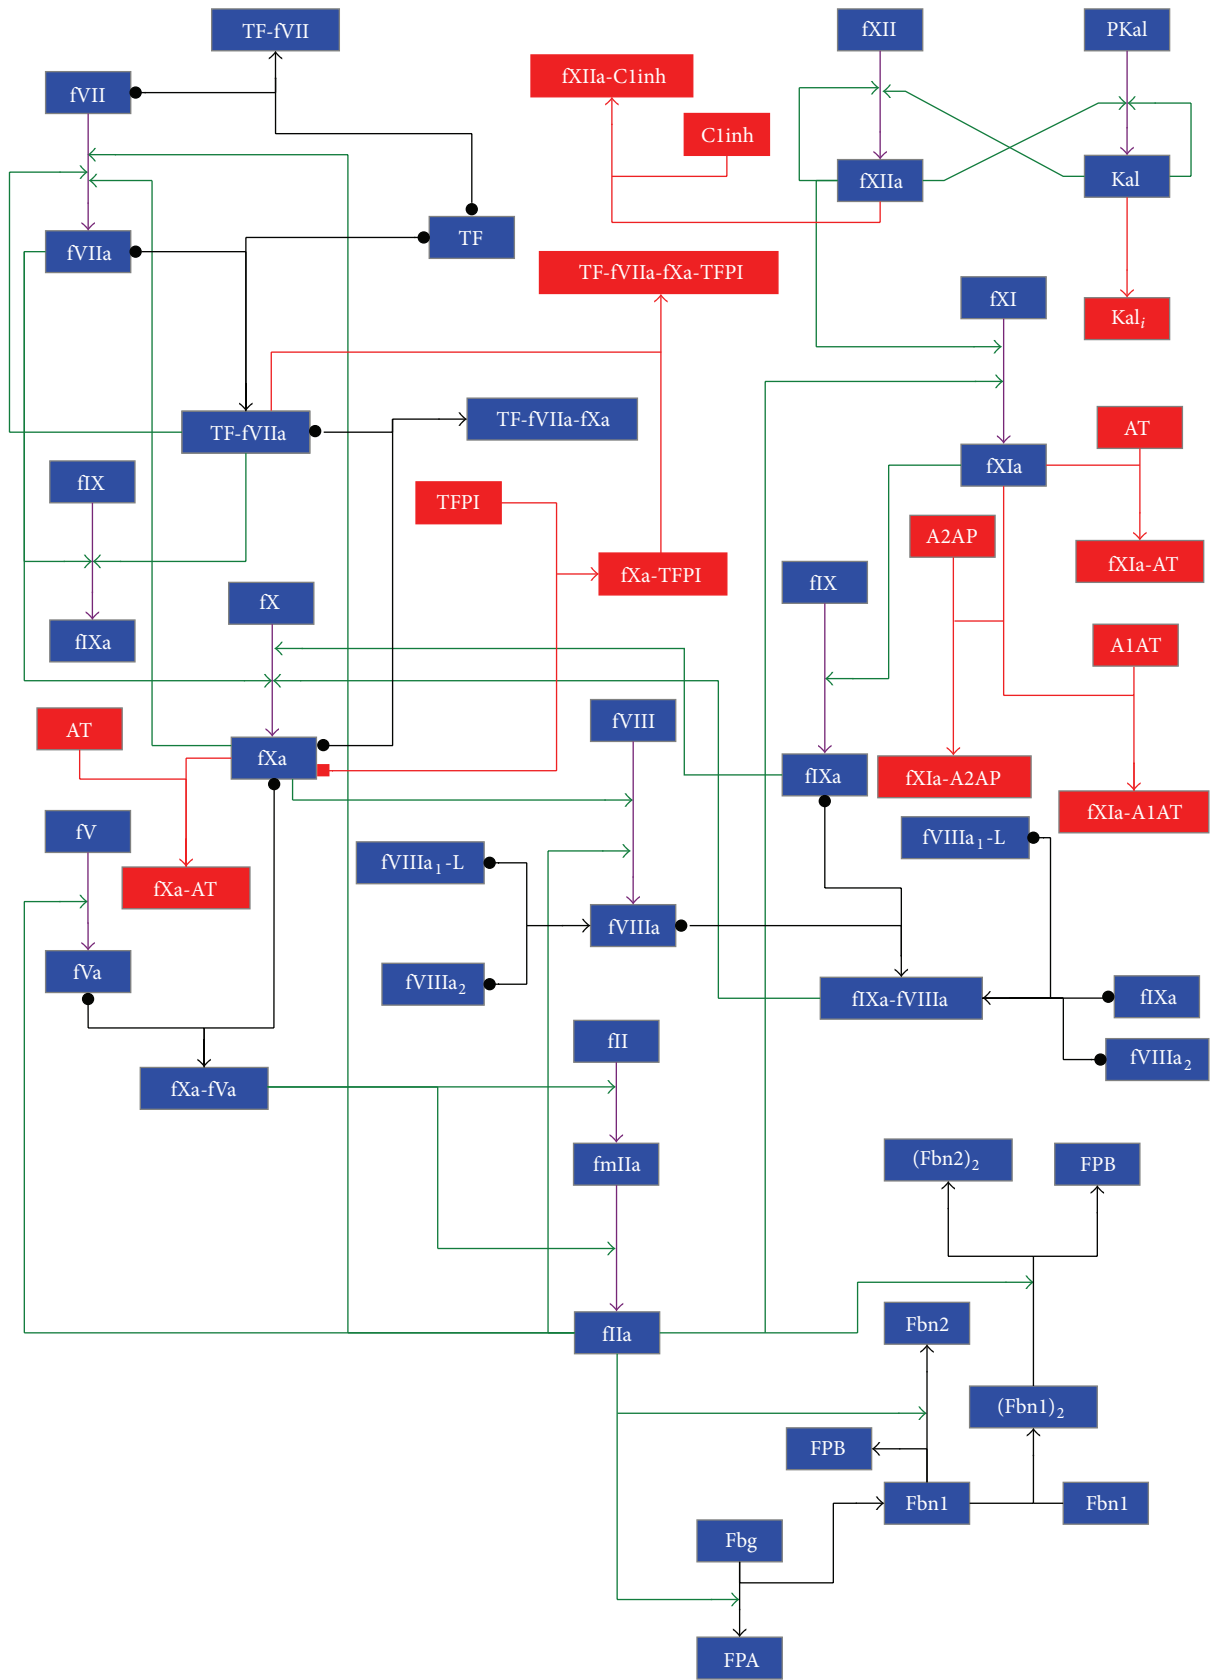


FIGURE 1: Graphical representation of the blood coagulation cascade model considered in this work. *Legend.* Blue box: coagulation factor; red box: inhibitor and related complexes. Black arrow: complex formation; green arrow: catalytic activation; violet arrow: activation; red arrow: inhibition. The reaction is reversible if the arrow tail consists in a dot.

- (iv) the activation of prekallikrein by factor XIIa (reactions r_{48} , r_{49} , r_{50}) and by kallikrein (reaction r_{54});
 - (v) the activation of factor XI by factor XIIa (reactions r_{61} , r_{62} , r_{63}) and by factor XIa (reaction r_{64});
 - (vi) the activation of factor IX by factor XIa (reactions r_{69} , r_{70} , r_{71});
 - (vii) the dissociation of free factor VIIIa (reactions r_{23} and r_{24});
 - (viii) the dissociation of factor VIIIa in complex with other factors (reactions r_{25} and r_{26}).
- (3) The third module corresponds to the *common pathway*, which consists in
- (i) the activation of factor II by factor Xa (modeled by reaction r_{16}) and by the complex between factor Xa and factor Va, through the formation of the intermediate meizothrombin (reactions r_{30} , \dots , r_{33});
 - (ii) the activation of factor VII by factor Xa (reaction r_6) and by factor IIa (reaction r_7);
 - (iii) the activation of factor VIII by factor Xa (reactions r_{75} , r_{76} , r_{77}) and by factor IIa (reaction r_{17});
 - (iv) the formation of a complex between factor Xa and factor Va (reactions r_{28} and r_{29});
 - (v) the activation of fibrinogen by factor IIa (reactions r_{84} , \dots , r_{96});
 - (vi) the activation of factor V by factor IIa (reaction r_{27});
 - (vii) the activation of factor XI by factor IIa (reactions r_{58} , r_{59} , r_{60}).
- (4) The fourth module describes the *inhibition of coagulation*, carried out through the main inhibitors of the BCC (antithrombin, Tissue Factor Pathway Inhibitor, CI-inhibitor, α 1-antitrypsin, and α 2-antiplasmin). These are tight binding inhibitors belonging to the serpin superfamily, which form irreversible complexes. This module consists in
- (i) the inhibition of factor Xa by Tissue Factor Pathway Inhibitor (modeled by reactions r_{34} and r_{35});
 - (ii) the inhibition of the complex between Tissue Factor, factor VIIa, and factor Xa by Tissue Factor Pathway Inhibitor (reactions r_{36} , r_{37} , r_{38});
 - (iii) the inhibition of factor IIa by antithrombin (reactions r_{40} and r_{42});
 - (iv) the inhibition of factor Xa by antithrombin (reaction r_{39});
 - (v) the inhibition of factor IXa by antithrombin (reaction r_{41});
 - (vi) the inhibition of factor XIa by antithrombin (reaction r_{65}) and by CI-inhibitor (reaction r_{66});

TABLE 2: Initial concentrations of molecular species in the blood coagulation cascade model (values taken from [30]).

Species	Symbol	Concentration (M)
α 1-Antitrypsin	A1AT	$4.50 \cdot 10^{-5}$
α 2-Antiplasmin	A2AP	$1.00 \cdot 10^{-6}$
Antithrombin III	ATIII	$3.40 \cdot 10^{-6}$
CI-inhibitor	Clinh	$2.50 \cdot 10^{-6}$
Fibrinogen	Fbg	$9.00 \cdot 10^{-6}$
Factor II	fII	$1.40 \cdot 10^{-6}$
Factor V	fV	$2.00 \cdot 10^{-8}$
Factor VII	fVII	$1.00 \cdot 10^{-8}$
Active factor VII	fVIIa	$1.00 \cdot 10^{-10}$
Factor VIII	fVIII	$7.00 \cdot 10^{-10}$
Factor IX	fIX	$9.00 \cdot 10^{-8}$
Factor X	fX	$1.6 \cdot 10^{-7}$
Factor XI	fXI	$3.10 \cdot 10^{-8}$
Factor XII	fXII	$3.40 \cdot 10^{-7}$
Prekallikrein	Pkal	$4.50 \cdot 10^{-7}$
Tissue Factor	TF	$5.00 \cdot 10^{-12}$
Tissue Factor Pathway Inhibitor	TFPI	$2.50 \cdot 10^{-9}$

- (vii) the inhibition of factor XIIa by CI-inhibitor (reaction r_{56}), by antithrombin (reaction r_{57}), by α 1-antitrypsin (reaction r_{67}), and by α 2-antiplasmin (reaction r_{68});
- (viii) the inhibition of the complex between Tissue Factor and factor VIIa by antithrombin (reaction r_{43});
- (ix) the inhibition of kallikrein (reaction r_{55}).

The values of the initial concentrations of the molecular species occurring in the BCC model are given in Table 2. According to [30], the concentrations of complexes and active factors were set to 0, except for the active factor VII, which is physiologically present in the blood circulation, even in the absence of damage, in a concentration that is approximately equal to 1% of the corresponding inactive factor [42].

The system of ordinary differential equations (ODEs), needed to carry out the simulations and the parameter sweep analysis presented in Section 3, was derived from the reactions given in Table 1 according to the *mass-action law*, with the exception of 14 reactions belonging to the set $S_\epsilon = \{r_{19}, r_{21}, r_{23}, r_{29}, r_{31}, r_{35}, r_{46}, r_{49}, r_{52}, r_{62}, r_{70}, r_{73}, r_{76}, r_{82}\}$. The mass-action law is the fundamental empirical law that governs biochemical reaction rates which states that, in a dilute solution, the rate of an elementary reaction is proportional to the product of the concentration of its reactants raised to the power of the corresponding stoichiometric coefficient [44]. The reactions in S_ϵ are instead influenced by a specifically defined variable ϵ , depending on a Hill function, fit against experimental data which quantify the platelet activation status, which is used to model the physiological levels of thrombin concentration as a function of platelet activation, as thoroughly described in [30]. More precisely, the value of ϵ influences the formation of some complexes




FIGURE 2: Dynamics of thrombin (factor IIa) in physiological condition.

occurring on the platelets' surface, by modifying the activity of reactions of the form



Namely, ε intervenes with the dissociation constant of these reactions, so that the corresponding standard ODEs are changed to yield new equations of the form

$$\frac{d[AB]}{dt} = k_1 \cdot [A] \cdot [B] - \frac{k_2 \cdot [A] \cdot [B]}{\varepsilon}. \quad (2)$$

The value of ε in (2) depends on the following Hill function \mathcal{H} , which quantifies the state of platelets activation according to the thrombin concentration (here denoted as $[fIIa]$), that is, the factor catalyzing the formation of the fibrin clot:

$$\mathcal{H}([fIIa^*(t)]) = \frac{[fIIa^*(t)]^{1.6123}}{[fIIa^*(t)]^{1.6123} + (2.4279 \cdot 10^{-9})^{1.6123}}, \quad (3)$$

where $[fIIa^*(t)] = \max_{t' \in [0, t]} \{[fIIa(t')]\}$, for a chosen time interval $[0, t]$ of simulation.

The value $[fIIa^*(t)]$ represents the maximum transient thrombin concentration and is needed to simulate the fact that, in physiological conditions, the thrombin concentration starts decreasing after rising to a peak (Figure 2). So doing, $[fIIa^*(t)]$ never decreases once it reaches its maximum magnitude; $[fIIa^*(t)]$ is equivalent to $[fIIa(t)]$ until the concentration of factor IIa reaches the peak, while thereafter it remains constant at that value, which is the maximum in the considered time interval $[0, t]$. Function \mathcal{H} allows simulating the physiological condition, whereby platelets remain active also when the thrombin concentration decreases.

For a given concentration of factor IIa, the maximum platelets activation state ε_{\max} is defined as

$$\varepsilon_{\max} = \varepsilon_{\max_0} + (1 - \varepsilon_{\max_0}) \cdot \mathcal{H}([fIIa^*(t)]), \quad (4)$$

where ε_{\max_0} defines the basal activation state of the platelets at simulation time $t = 0$. The value ε_{\max_0} is initially set to 0.01, assuming a basal 1% binding strength of coagulation factors to the resting platelets' surface. When the full activation of platelets is reached, ε_{\max} is equal to 1 and the complex dissociation constants are minimized (see [30] for more details).

2.2. CUDA Architecture and the coagSODA Simulator. Introduced by Nvidia in 2006, Compute Unified Device Architecture (CUDA) is a parallel computing platform and programming model that provides programmers with a framework to exploit GPUs in general-purpose computational tasks (GPGPU computing). GPGPU computing is a low-cost and energy-wise alternative to the traditional high-performance computing infrastructures (e.g., clusters of machines), which gives access to the tera-scale computing on common workstations of mid-range price. However, due to the innovative architecture and the intrinsic limitations of GPUs, a direct porting of sequential code on the GPU is most of the times unfeasible, therefore making it challenging to fully exploit its computational power and massive parallelism [45].

CUDA combines the single instruction multiple data (SIMD) architecture with multithreading, which automatically handles the conditional divergence between threads. The drawback of such flexibility is that any divergence of the execution flow between threads causes a serialization of the execution, which affects the overall performances. Under CUDA's naming conventions, the programmer implements the *kernel*, that is, a C/C++ function, which is loaded from the host (the CPU) to the devices (one or more GPUs) and replicated in many copies named *threads*. Threads can be organized in three-dimensional structures named *blocks* which, in turn, are contained in three-dimensional *grids* (a schematic description is given in Figure 3, left side). Whenever the host runs a kernel, the GPU creates the corresponding grid and automatically schedules each block on one free streaming multiprocessor available on the GPU, allowing a transparent scaling of performances on different

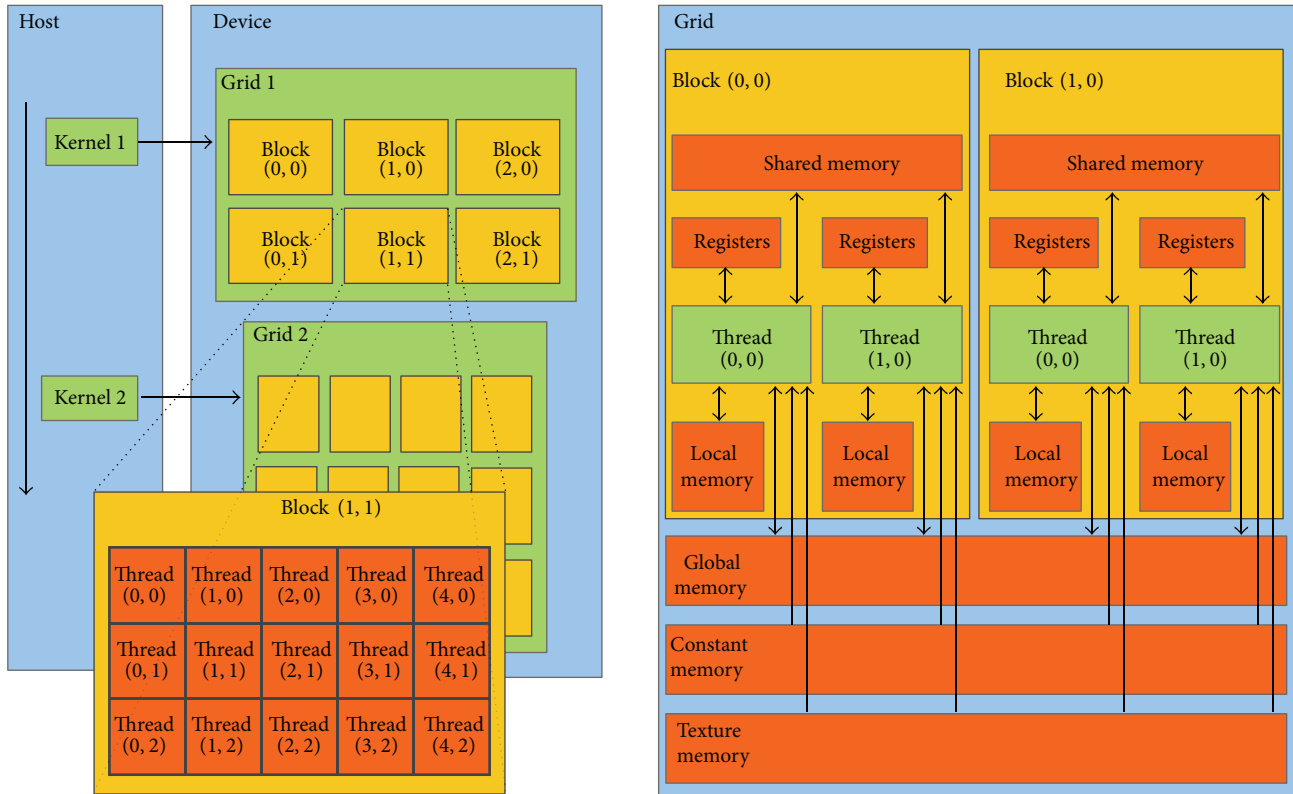


FIGURE 3: Schematic description of CUDA's architecture, in terms of threads and memory hierarchy. *Left Side*. Threads organization: a single kernel is launched from the host (the CPU) and is executed in multiple threads on the device (the GPU). Threads can be organized in three-dimensional structures named blocks which can be, in turn, organized in three-dimensional grids. The dimensions of blocks and grids are explicitly defined by the programmer. *Right Side*. Memory hierarchy: threads can access data from many different memories with different scopes. Registers and local memories are private for each thread. Shared memory lets threads belonging to the same block communicate and has low access latency. All threads can access the global memory, which suffers high latencies, but it is cached since the introduction of the Fermi architecture. Texture and constant memory can be read from any thread and are equipped with a cache as well; in this work we exploit the constant memory. Figures are taken from Nvidia's CUDA programming guide [60].

devices. Threads within a block are executed in groups of 32 threads named *warps*.

The GPU is equipped with different kinds of memory. In this work, we exploit the *global memory* (accessible from all threads), the *shared memory* (accessible from threads of the same block), the *local memory* (registers and arrays, accessible from owner thread), and the *constant memory* (cached and not modifiable). A schematic representation of this memory hierarchy is shown in Figure 3, right side. To achieve the best performances, the shared memory should be exploited as much as possible; however, it is very limited (i.e., 49152 bytes for each multiprocessor, since the introduction of the Fermi architecture) and it introduces restrictions on the blocks' size. On the other hand, the global memory is very large (thousands of MBs) but suffers from high latencies. A solution to this problem was implemented on the Fermi architecture, where the global memory is equipped with a L2 cache. This architecture also introduced the possibility to balance 64 KB of fast on-chip memory between the shared memory and L1 cache using two possible configurations, 48 KB for the shared memory and 16 KB for L1 cache, or 16 KB for the shared memory and 48 for L1 cache. The Kepler architecture, used in this paper, allows a third and perfectly

balanced configuration, where shared memory and L1 cache obtain the same amount of memory (32 KB). See Figure 4 for a schematization of the memory architecture.

The systematic analysis of models of biological systems often consists in the execution of large batches of simulations. One of the standard analyses that can be executed on such kind of models regards an intensive search within the parameters space, which requires large numbers of independent simulations. In order to reduce the computational burden, we previously implemented on the CUDA architecture one of the most efficient numerical integration algorithms for ODEs, LSODA, that is able to automatically recognize stiff and nonstiff systems and to dynamically select between the most appropriate integration procedure (i.e., Adams method in the absence of stiffness and the Backward Differentiation Formulae otherwise) [22]. This GPU-powered tool is called cupSODA [24, 25] and exploits CUDA's massive parallelism to execute different and independent simulations in each thread, thus reducing the computational time required by a standard CPU counterpart of LSODA.

Besides being very efficient for the simulation of many independent simulations, cupSODA is also user-friendly. LSODA was originally designed to solve ODEs systems




FIGURE 4: Schematic description of memory hierarchies in Fermi and Kepler architectures. GPUs relying on these architectures are equipped with a two-level data cache and a read-only data cache. Shared memory and L1 cache share the same on-chip 64 KB memory banks; the amount of memory can be reconfigured by the user, according to the specific needs of the application. Figure taken from Nvidia's Kepler GK110 whitepaper [57].

written in the canonical form, but the user is supposed to specify the system of ODEs by implementing a custom C function that is passed to the algorithm; moreover, in order to speed up the computation when dealing with stiff systems, in LSODA the Jacobian matrix associated with the ODEs system must be implemented as a custom C function as well. On the contrary, cupSODA was conceived as a *black-box* simulator that can be easily used without any programming skills. cupSODA consists in a tool to automatically convert a generic mechanistic reaction-based model of a biological system into the corresponding set of ODEs, to comply with the mass-action kinetics [46], and to directly encode the obtained system, along with the corresponding Jacobian matrix, as C arrays.

This fully automatic methodology can be exploited for the simulation of models whose chemical kinetics is based on the mass-action law. The BCC model described in Section 2.1, though, includes a set S_ε of reactions that do not follow the mass-action kinetics [30]; the platelets activity is not explicitly modeled by biochemical reactions, but it is realized by modulating the rate of the dissociation of the complexes formed on a platelet's surface by means of the variable ε , which is calculated with a special equation during the integration steps. For this reason, cupSODA was used as a starting point for the development of a new tool able to compute at *run-time* the specific kinetics of this set of reactions.

This new CUDA-powered simulation tool, named coagSODA, is specifically tailored for the simulation of the BCC

model developed in [30]. In particular, coagSODA realizes the run-time calculation of the ε value required to correctly simulate the activity of platelets, which is determined according to the equations described in Section 2.1. The platelets activation state ε is calculated at each time instant by solving the following differential equation:

$$\frac{d\varepsilon}{dt} = k(\varepsilon_{\max} - \varepsilon), \quad (5)$$

where the constant k is inversely proportional to the time scale of platelets activation and is set to 0.005. This is consistent with the fact that platelets do not instantly achieve their maximum attainable activation state (ε_{\max}), but they reach it on a physiologically relevant timescale [30].

Dealing with the 14 reactions in the set S_ε that are influenced by the Hill function \mathcal{H} (see Section 2.1), the value of $[fIIa^*(t)]$ must be stored on the GPU because, during each integration step, coagSODA recalculates (4) which exploits the value $[fIIa^*(t)]$ to determine a new ε value. Since CUDA's architecture does not offer static variables, the information for each thread has to be memorized in the global memory. The accesses to the global memory and the computational costs due to these additional calculations slow down the integration process, with respect to the integration of ODEs performed according to strictly mass-action based kinetics (as in the cupSODA simulator). Nevertheless, in Section 3.3 we show that our parallel implementation largely outperforms a sequential counterpart of the LSODA algorithm.




FIGURE 5: Dynamics of fibrinogen in physiological condition. The clotting time is defined as the time necessary to convert the 70% of the fibrinogen into fibrin.

As in the case of cupSODA [24, 25], coagSODA exploits the shared memory to improve performances by storing the current state and time of the simulations in these low-latency memory banks. Despite the improvement of performances ensured by this solution, it strongly affects the occupancy of GPU's multiprocessors and therefore it represents the limiting factor for the number of blocks that can be executed simultaneously; as a matter of fact, coagSODA is limited to 2 blocks per streaming multiprocessor on GPUs based on the Fermi architecture, with a reduced exploitation of the GPU with respect to the theoretical 8 blocks allowed by this architecture.

3. Results and Discussion

In this section we discuss the results of the parameter sweep analysis (PSA) carried out on the BCC model, to the aim of investigating either the prolongation or the reduction of the clotting time in response to perturbed values of some reaction constants and of the initial concentration of some molecular species, chosen according to their meaning within the whole pathway. PSA was performed by generating a set of different initial conditions, corresponding to different parameterizations of the model, and then automatically executing the deterministic simulations with coagSODA. The use of GPU technology is fundamental in this type of analysis, especially for large biological systems as the BCC is, because it drastically reduces the computational time.

The sweep analysis for single parameters (PSA-1D) was performed considering a logarithmic sampling of numerical values of each parameter under investigation (reaction constant or initial molecular concentration) within a specified range with respect to its physiological reference value. The sweep analysis over pairs of parameters (PSA-2D) was performed by simultaneously varying the values of two parameters within a specified range, considering a logarithmic sampling on the resulting lattice. The logarithmic sampling allows uniformly spanning different orders of magnitude of the parameters value using a reduced and fine-grained set

of samples, therefore efficiently analyzing the response of the system in a broad range of conditions.

To determine the response of the BCC to perturbed conditions, we chose the *clotting time* (CT) as output of the PSA. The CT is defined as the time necessary to convert the 70% of the fibrinogen into fibrin (see Figure 5), conventionally assumed to correspond to the time required to form the clot [47], and it is generally used in laboratory tests for monitoring the therapy with anticoagulant drugs. According to the model defined in [30], the reference value of CT is around 300 seconds in physiological conditions. We investigate here the response of the BCC by evaluating the CT in various conditions, corresponding to different values of the reaction constants, varying over six orders of magnitude with respect to their physiological values (i.e., three below and three above the reference values, if not otherwise specified), as well as to different values of the initial molecular concentrations, and varying over twelve orders of magnitude with respect to their physiological values (i.e., six below and six above the reference values, if not otherwise specified).

The total number of parallel simulations executed to carry out these analyses was 100 for PSA-1D over reaction constants, 200 for PSA-1D over initial concentrations, and 1600 for PSA-2D.

Finally, we present the comparison of the performance of the CPU and GPU to run an increasing number of simulations of the BCC model, to prove the efficiency of coagSODA.

3.1. PSA-1D of the BCC Model

3.1.1. Reaction r_{44} . The first PSA was performed to determine the value of the kinetic constant for the autoactivation of factor XII (reaction r_{44} in Table 1), which corresponds to an upstream process in the intrinsic pathway. This analysis was motivated by two considerations. Firstly, by using the full parameterization given in [30], the action of the intrinsic pathway turns out to be fundamental for the BCC *in vivo*, which is in contrast to experimental observations which




FIGURE 6: Clotting time according to PSA-ID over constant k_{44} of reaction r_{44} : XII \rightarrow XIIa. The reference value of k_{44} used in [30] is $5 \cdot 10^{-4} \text{ s}^{-1}$; the sweep range is $[5 \cdot 10^{-10}, 5 \cdot 10^{-4}]$.




FIGURE 7: Clotting time according to PSA-ID over constant k_{27} of reaction r_{27} : IIa + V \rightarrow IIa + Va. The reference value is $2 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$; the sweep range is $[2 \cdot 10^4, 2 \cdot 10^{10}]$.

indicate that the extrinsic pathway is the main responsible of clot formation. Secondly, in [30] all constants values have a reference to experimental measurements, except for the constant of this reaction (which corresponds to entry 29 in Table 1 in [30]).

Figure 6 shows the results of this PSA-ID, where k_{44} was varied over six orders of magnitude, considering the value given in [30] as the upper limit of the sweep interval. We can observe that, by decreasing the value of k_{44} with respect to the value considered in [30], the CT increases with respect to its reference value; however, for values of this constant lower than $1.00 \cdot 10^{-6} \text{ s}^{-1}$ the CT remains unaltered, and this can be intuitively explained by the fact that in this situation the fibrinogen is mainly activated by the extrinsic pathway.

Consequently, we assigned the value $7.00 \cdot 10^{-6} \text{ s}^{-1}$ to k_{44} , achieving a CT that is comparable to the experimental observations of the BCC *in vivo*. This new value was used in all PSA discussed in what follows.

3.1.2. Reactions r_{27} and r_{58} . In the next PSA we investigated the effect of the perturbation of the kinetics of two pivotal

reactions of the BCC model. Reaction r_{27} , which describes the catalytic activation of factor V by factor IIa, was chosen because it represents the main positive feedback within the common pathway; reaction r_{58} , which is involved in the activation of factor XI by binding to factor IIa, was chosen because it represents the main positive feedback in the intrinsic pathway (Figure 1). Moreover, preliminary PSA over all reaction constants of the BCC model given in Table 1 evidenced that these two reactions are among the most relevant steps of the coagulation network.

The PSA-ID over k_{27} (Figure 7) shows that the CT is very sensitive to the perturbation of the rate of this reaction, when the reference value of its constant (i.e., $k_{27} = 2 \cdot 10^7 \text{ M}^{-1} \text{ s}^{-1}$) is either increased or decreased; in particular, when k_{27} is very low, a plateau in CT is reached since the strength of the positive feedback exerted by factor IIa is largely reduced, a condition where the contribution of the amplification of the hemostatic stimulus (due by the common pathway) to the formation of the clot is basically not effective. On the other side, the PSA-ID over k_{58} (Figure 8) shows that, while a decrease of the constant with respect to its reference value




FIGURE 8: Clotting time according to PSA-1D over constant k_{58} of reaction r_{58} : $\text{IIa} + \text{XI} \rightarrow \text{IIa-XI}$. The reference value is $1 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$; the sweep range is $[1 \cdot 10^5, 2 \cdot 10^{11}]$.




FIGURE 9: Clotting time at different initial concentrations of factor VIII. The reference value is $7.00 \cdot 10^{-10} \text{ M}$ (dashed black line); the sweep range is $[7.00 \cdot 10^{-16}, 7.00 \cdot 10^{-4}]$. The blue area indicates a condition in which factor VIII concentration is less than 1% of its physiological concentration, corresponding to a situation of severe haemophilia; the red area indicates a condition in which factor VIII concentration is between 1% and 30% of its physiological concentration, corresponding to a situation of mild haemophilia; the green area indicates a condition in which factor VIII concentration is greater than 130% of its physiological concentration, corresponding to a situation of hypercoagulability.

(i.e., $k_{58} = 1 \cdot 10^8 \text{ M}^{-1} \text{ s}^{-1}$) does not have any substantial effect, an increase of its value leads to a progressive increase in the CT. This increase is due to the fact that factor IIa is sequestered in the formation of a complex with factor XI, and hence it is no longer available as a free component in blood to participate in other reactions, especially those reactions of the extrinsic pathway which principally lead to the clot formation *in vivo*. This behavior highlights that the intrinsic pathway has a secondary role in blood coagulation *in vivo*, compared with the extrinsic pathway, as also evidenced by various experimental observations [48].

3.1.3. *Factors VIII, IX, and II.* The next set of PSA-1D was realized by varying the initial concentrations of factors VIII, IX, and II. These factors were selected since both an excess

and a deficiency of their concentrations lead to diseases related to blood clotting.

The PSA-1D over factor VIII (Figure 9) shows that increasing the initial concentration of this factor results in decreasing the CT, suggesting the possible presence of hypercoagulable states in these perturbed conditions. As a matter of fact, high levels of factor VIII cause an increased risk of deep vein thrombosis and pulmonary embolism [49]. On the other hand, individuals with less than 1% of the average concentration of factor VIII show a severe haemophilia A, characterized by higher CT (Figure 9), and require infusions of plasma containing the deficient factor; otherwise, frequent spontaneous bleeding would occur [50]. When the concentration of this factor is between 5% and 30% of the average concentration, individuals still risk bleeding in case of trauma [50].




FIGURE 10: Clotting time at different initial concentrations of factor IX. The reference value is $9.00 \cdot 10^{-8}$ M; the sweep range is $[9.00 \cdot 10^{-14}, 9.00 \cdot 10^{-2}]$.




FIGURE 11: Clotting time at different initial concentrations of thrombin. The reference value is $1.40 \cdot 10^{-6}$ M; the sweep range is $[1.40 \cdot 10^{-12}, 1.40 \cdot 10^0]$.

The PSA-ID over factor IX (Figure 10) shows only a slight decrease of the CT as the initial concentration of fIX increases; this is in contrast to recent studies that demonstrated how the excess of factor IX leads to an increased risk of deep vein thrombosis [51]. This result can be explained by considering that (i) a high concentration of factor IX is not sufficient to bring about coagulation problems, though when the concentration of other factors is above the average value (yet not at pathological levels), prothrombotic states can be observed; (ii) in this model we consider average values as initial concentrations of factors; however, individuals are characterized by different (balanced) combinations of procoagulant and anticoagulant factor levels that altogether contribute to define a unique coagulation phenotype that reflects the developmental, environmental, genetic, nutritional, and pharmacological influences of each individual [52]. On the contrary, the lack of factor IX causes

haemophilia B, characterized by higher CT with respect to the reference value (Figure 10).

Furthermore, by comparing the PSA-ID of factors VIII and IX (Figures 9 and 10) it is clear that haemophilia A is more serious than haemophilia B, since the CT achieved in conditions of factor VIII deficiency is higher than the CT obtained in the case of factor IX deficiency.

In both PSA-ID over factors VIII and IX we observed, after the initial decrease of the CT, an unexpected increase of the CT as the factor concentration increases. This counter-intuitive behavior arises at very high concentrations of these factors (with respect to the average physiological levels) and, to the best of our knowledge, it was never observed *in vivo*. Nonetheless, it would be interesting to verify, by means of *ad hoc* laboratory experiments, if the model correctly describes the behavior of the BCC even in these conditions or, on the contrary, it is not predictive in these extreme situations.

The PSA-1D over factor II (Figure 11) shows a dramatic decrease of the CT as the initial concentration of this factor increases (with respect to the average physiological level). This behavior resembles the effects of hypercoagulability (or thrombophilia), a disease caused by mutation G20210A in the prothrombin gene [53] that causes an increase of the prothrombin level (factor II) in the blood flow, resulting in an excessive formation of the active form of this factor, thus heightening venous thrombosis risks [54]. Hypercoagulability is usually treated with warfarin therapy, or with other anticoagulants with a similar effect. These drugs decrease the capacity of coagulation factors to become active, preventing the formation of unwanted thrombi.

On the other hand, when the initial concentration of factor II is low, we achieved the effects of prothrombin deficiency, a rare autosomal recessive disease that causes a tendency to severe bleeding [29, 55]. As shown in Figure 11, a concentration equal to 10% of the physiological value of factor II (i.e., $1.4 \cdot 10^{-6}$ M) leads to clotting effects similar to severe haemophilia A.

3.2. PSA-2D of the BCC Model. We present here the results of the PSA-2D on the BCC model, where couples of parameters were varied to analyze the possible effects arising from the combined perturbation of their values.

3.2.1. Reactions r_{27} and r_{58} . Figure 12 shows the effect of the simultaneous variation of constants k_{27} and k_{58} (over the same sweep ranges considered in the two PSA-1D, Section 3.1). This result remarks that reaction r_{27} , involved in the common pathway, has a stronger influence on the BCC, and that there is a synergic interplay between these two reactions. In particular, when the value of k_{27} is low and the value of k_{58} is high, the CT is higher than the values achieved when only a single constant is changed, because in this condition both the intrinsic and the common pathways are simultaneously inhibited.

3.2.2. Factor VIII, Factor IX, and Tissue Factor. In the last two PSA-2D we varied the initial concentration of factor VIII and Tissue Factor and the initial concentration of factor IX and Tissue Factor, respectively.

The initial concentrations of factors VIII and IX were varied over four orders of magnitude, using their physiological values as upper limit for the sweep ranges; the concentration of Tissue Factor was varied over four orders of magnitude, two above and two below its reference value (see Table 2). The rationale behind this choice is to observe how the BCC model, in conditions corresponding to different states of haemophilia (obtained by decreasing the concentrations of factors VIII and IX), behaves with different initial concentrations of the Tissue Factor, which is the upstream factor of the extrinsic pathway, that is, the most important element of the BCC.

The results of these PSA-2D show that, with respect to the condition of haemophilia B, in the case of haemophilia A the amount of Tissue Factor (below its reference value) has a negligible influence on the CT, as indicated by the

presence of a plateau in Figure 13; on the contrary, concerning haemophilia B, a deficiency of Tissue Factor leads to an increase of the CT, especially when factor IX is present in low concentrations (Figure 14).

The different results achieved in the two PSA-2D are due to the presence, in the BCC model, of a direct interaction between Tissue Factor and factor IX by means of the TF-fVIIa complex (see reactions r_{13}, \dots, r_{15} in Table 1). Indeed, the lack of Tissue Factor directly affects the concentration of active factor IX, which results in a strong alteration of the CT with respect to physiological conditions.

3.3. CPU versus GPU Performance Comparison. In order to show the relevant speedup achieved by coagSODA, we present here the comparison of the computational effort required by GPU and CPU for the simulation of the BCC model. The performances of our GPU simulator were compared with those obtained using the LSODA algorithm implemented in the software COPASI [56], executing on the CPU the same set of simulations that were run on the GPU. COPASI is single-threaded and does not exploit the physical and logical cores of the CPU; therefore, it represents a good benchmark as a single-node CPU-bound simulator of biological systems.

In all simulations, executed on both GPU and CPU, we stored 100 samples, uniformly distributed in the time interval considered for each simulation, that is, $[0, 700]$ seconds, of the dynamics of all chemical species involved in the BCC model. The settings for the LSODA algorithm were the following: relative error $1 \cdot 10^{-7}$, absolute error $1 \cdot 10^{-13}$, and maximum number of internal steps set to 20000.

3.3.1. Benchmark Details. The GPU used for the simulations is a Nvidia Tesla K20c, equipped with 2496 cores organized in 13 streaming multiprocessors, GPU clock 706 MHz, and 5 GB of DDR5 RAM. In all tests, coagSODA was compiled and executed with version 5.5 of CUDA libraries. Even though this GPU has compute capability 3.5 and is based on the GK110 Kepler architecture, currently coagSODA does not exploit any of the new features (e.g., Dynamic Parallelism, Hyper-Q, Remote DMA) with the exception of the new ISA encoding, which allows threads to exploit an increased number of registers (255 instead of 63), reducing register spilling into global memory and increasing performances [57]. Moreover, as described in Section 2.2, the Kepler architecture offers the possibility to reconfigure the 64KB on-chip cache, balancing between L1 cache and shared memory. Since coagSODA exploits the shared memory to reduce the latencies during the access to the concentration values of the BCC model, we opted for the following configuration: 48 KB for the shared memory, 16 KB for the L1 cache.

The performance of the GPU was compared against a personal computer equipped with a dual-core CPU Intel Core i5, frequency 2.3 GHz, 4 GB of DDR3 RAM, running the operating system Mac OS X Lion 10.7.5. The software used as a coagSODA-equivalent single-threaded CPU implementation is COPASI version 4.8 (build 35).




FIGURE 12: Clotting time according to PSA-2D over reaction constants k_{27} and k_{58} . The sweep ranges are $[2 \cdot 10^4, 2 \cdot 10^{10}]$ and $[1 \cdot 10^5, 2 \cdot 10^{11}]$, respectively.




FIGURE 13: Clotting time at different initial concentrations of Tissue Factor and factor VIII. The sweep ranges are $[5.00 \cdot 10^{-14}, 5.00 \cdot 10^{-10}]$ and $[7.00 \cdot 10^{-14}, 5.00 \cdot 10^{-10}]$, respectively.

A direct comparison of the performances between these two different architectures and implementations is not an easy task, since CPUs are optimized for single-thread execution and exploit a large number of optimizations (e.g., higher clock frequency, instruction caching, pipelining, out of order execution, and branch prediction), whereas GPUs are optimized for graphics processing and parallel execution of identical code, relying only on multilevel data caching. For this reason, in this paper we propose an empirical comparison of the performances based on the running times of identical batches of simulations of the BCC model.

3.3.2. Benchmark Results. In Figure 15 we show the comparison of the running times required to run several batches of simulations, executed to carry out the PSA-ID over the reaction constant k_{27} . The choice of comparing the performances of the GPU and CPU implementations by executing batches of simulations that are related to a PSA, instead of running n independent but identical simulations (i.e., all characterized by the same parameterization of the model), is due to the fact that these results represent a more realistic scenario in the computational analysis of biological systems, whereby it is useful to investigate large search spaces of parameters, corresponding to different perturbed conditions of the model [26]. Moreover, for the evaluation of the running

time, the execution of a batch of n identical deterministic simulations would be futile. The figure clearly shows that coagSODA always performs better than the CPU counterpart. In particular, while the CPU performance increases linearly with the number of simulations, the running times are strongly reduced on the GPU; in this case, the overall running time roughly corresponds to the time required for the execution of the slowest simulations. This is due to the fact that different parameterizations may require different time steps for LSODA to converge, and the execution of a block terminates as long as all the threads it contains terminate. In turn, the execution of a kernel terminates when all blocks terminate; for this reason, a single simulation, whose dynamics requires more steps than the others, may affect the overall running time.

In Table 3 we report the running times of all batches of simulations, along with the speedup achieved on the GPU. In particular, these results highlight that the advantage of exploiting the GPUs for the simulation of the BCC model becomes more evident as the number of simulations increases, with a 181 \times speedup when a PSA with 10000 different parameterizations is executed. Therefore, the GPU accelerated analysis of the BCC model with coagSODA represents a novel, relevant computational means to investigate the behavior of this complex biological system under nonphysiological conditions and could be exploited




FIGURE 14: Clotting time at different initial concentrations of Tissue Factor and factor IX. The sweep ranges are $[5.00 \cdot 10^{-14}, 5.00 \cdot 10^{-10}]$ and $[9.00 \cdot 10^{-12}, 9.00 \cdot 10^{-8}]$, respectively.




FIGURE 15: Comparison of the computational time required to execute an increasing number of simulations of the BCC model on CPU (Intel Core i5, 2.3 GHz) and GPU (Nvidia Tesla K20c, GPU clock 706 MHz). The computational time is expressed in seconds. The time values related to 1000, 5000, and 10000 simulations on the CPU were estimated by regression (see also Table 3).

TABLE 3: CPU versus GPU performance comparison.

Number of simulations	CPU time (sec)	GPU time (sec)	Speedup
5	380.8	127.8	2.98×
10	1060.4	232.9	4.56×
50	5170.4	805.8	6.42×
100	11652.8	1097.7	10.61×
500	53605.8	1739.4	30.81×
1000	107358.2*	1998.2	53.73×
5000	536350.7*	3096.0	173.2×
10000	1072497.5*	5895.1	181.9×

* Values estimated by regression.

to efficiently determine the response of the BCC to different therapeutical drugs.

4. Conclusions

Thanks to their high-performance computing capabilities and the very low costs, GPUs nowadays represent a suitable

technology for the development and the application of parallel computational methods for *in silico* analysis of complex biological systems. However, the implementation of efficient computational tools able to fully exploit the large potentiality of GPUs is still challenging, since good programming skills are required to implement GPU-based algorithmic methods, and to handle specific features as an optimal usage of memory or the communication bandwidth between GPU and CPU. Moreover, algorithms cannot be directly ported on the GPU because of the limited programming capabilities allowed by GPU kernels; as a matter of fact, they need to be redesigned to meet the requirements of the underlying SIMD architecture.

In this work we presented coagSODA, a GPU-powered simulator specifically developed to carry out fast parallel simulations of the BCC model. coagSODA was designed to offer a black-box solution usable by any final user in an easy way. It relies on cupSODA [24], a numerical integrator for ODEs that we previously implemented for the GPU architecture, based on the LSODA algorithm and capable of automatically translating a reaction-based model into a set of coupled ODEs. In addition to mass-action kinetics, coagSODA implements specific functions to compute

the kinetics of Hill function based reactions, such as those involved in the platelets activity of the BCC model. coagSODA exploits the massive parallel capabilities of modern GPUs, and our results demonstrated that it can achieve a relevant reduction of the computational time required to execute many concurrent and independent simulations of the BCC dynamics. The mutual independence of the simulations allows fully exploiting the underlying SIMD architecture; moreover, coagSODA benefits from an additional speedup, thanks to our choice of storing the state of the system into the low-latency shared memory (a solution that was already implemented in cupSODA). Since the BCC model is large (96 reactions, 71 chemical species), a large amount of shared memory was assigned to each thread, strongly reducing the theoretical occupancy of the GPU, that is, the ratio of active warps with respect to the maximum number of warps supported by each streaming multiprocessor of the GPU. However, the results of the analysis presented in this work, performed on the BCC model, show that coagSODA achieves a relevant boost with respect to a reference CPU implementation. For instance, in the case of 10000 simulations, we achieved a noticeable 181× speedup. Interestingly, the performances of coagSODA are better than COPASI even for small batches of simulations. These results indicate that, for biological models consisting in many reactions and many species, our GPU implementation of LSODA becomes more profitable than the CPU counterpart as the number of concurrent simulations increases, making it suitable especially when performing demanding computational analysis such as, for example, parameter sweep, parameter estimation, or sensitivity analysis.

As a test case, in this paper we presented several parameter sweep analyses over reaction constants and initial concentrations of factors involved in the BCC model. Other computational analyses on mathematical models of this pathway were previously presented. For instance, in [58], a parameter estimation of the reaction constants of different models involving the activation of factor X, factor V, prothrombin, and the inactivation of factors was performed. The aim of this analysis was to discriminate between different models to unravel the mechanisms on the basis of the BCC. Similar analyses were executed on a model describing thrombin generation in plasma, since the reaction mechanism, the reaction constants, and initial concentrations were unknown [59]. In [34], a sensitivity analysis of a model consisting of 44 species over 34 chemical species was presented; reaction constants were varied between 10% and 1000% of their reference value, to the aim of identifying the most influential factors of the BCC.

All these analyses were performed by means of sequential simulations of the models under investigation; therefore, in general, only small batches of simulations could be run (for instance, in [34] only 836 simulations were executed to execute the sensitivity analysis of the model). On the other hand, in this paper we presented the results of different PSA, efficiently executed by means of coagSODA, overall providing useful information regarding unknown parameters and interesting insights into the functioning of the BCC. A thorough sensitivity analysis of the whole BCC model, consisting in

around $5 \cdot 10^5$ simulations, is currently in progress on GPUs by our group, exploiting the great computational efficiency of coagSODA.

The results of our computational analyses should be now validated by means of specifically designed laboratory experiments. In particular, we identified a plausible value for the constant of the reaction describing the autoactivation of factor XII (since no reference values can be found in literature); moreover, the results of the PSA over factor IX suggest that a deficiency of this factor is not enough to cause severe bleeding disorders as haemophilia B, but an alteration of other factors seems to be necessary for the occurrence of such condition (e.g., the lack of Tissue Factor, as suggested by the PSA-2D over Tissue Factor and factor IX). Finally, the PSA over factors VIII and IX showed, in a situation characterized by very high concentrations of these factors, a counterintuitive behavior in which the clotting time is increased with respect to the value obtained in physiological conditions. Indeed, it would be interesting to design *ad hoc* laboratory experiments to verify if the BCC model is actually predictive in such extreme situations.

The coagSODA software and the SBML version of the BCC model used in this work are available from the authors upon request.

Conflict of Interests

The authors declare that there is no conflict of interests regarding the publication of this paper.

Acknowledgment

This work is partially supported by SYSBIO Centre for Systems Biology.

References

- [1] B. B. Aldridge, J. M. Burke, D. A. Lauffenburger, and P. K. Sorger, "Physicochemical modelling of cell signalling pathways," *Nature Cell Biology*, vol. 8, no. 11, pp. 1195–1203, 2006.
- [2] I.-C. Chou and E. O. Voit, "Recent developments in parameter estimation and structure identification of biochemical and genomic systems," *Mathematical Biosciences*, vol. 219, no. 2, pp. 57–83, 2009.
- [3] D. Besozzi, P. Cazzaniga, D. Pescini, G. Mauri, S. Colombo, and E. Martegani, "The role of feedback control mechanisms on the establishment of oscillatory regimes in the Ras/cAMP/PKA pathway in *S. cerevisiae*," *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2012, no. 10, 2012.
- [4] A. Saltelli, M. Ratto, T. Andres et al., *Global Sensitivity Analysis: The Primer*, Wiley-Interscience, 2008.
- [5] A. Raue, C. Kreutz, T. Maiwald et al., "Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood," *Bioinformatics*, vol. 25, no. 15, pp. 1923–1929, 2009.
- [6] A. Dräger, M. Kronfeld, M. J. Ziller et al., "Modeling metabolic networks in *C. glutamicum*: a comparison of rate laws in combination with various parameter optimization strategies," *BMC Systems Biology*, vol. 3, article 5, 2009.

- [7] D. Besozzi, P. Cazzaniga, G. Mauri, D. Pescini, and L. Vanneschi, "A comparison of genetic algorithms and particle swarm optimization for parameter estimation in stochastic biochemical systems," in *Proceedings of the 7th European Conference, Evolutionary Computation, Machine Learning and Data Mining in Bioinformatics (EvoBIO '09)*, C. Pizzuti, M. Ritchie, and M. Giacobini, Eds., vol. 5483 of *Lecture Notes in Computer Science*, pp. 116–127, Springer, 2009.
- [8] M. S. Nobile, D. Besozzi, P. Cazzaniga, G. Mauri, and D. Pescini, "cupSODA: a GPU-based multi-swarm PSO method for parameter estimation in stochastic biological systems exploiting discrete-time target series," in *Proceedings of the 10th European Conference, Evolutionary Computation, Machine Learning and Data Mining in Bioinformatics (EvoBIO '12)*, M. Giacobini, L. Vanneschi, and W. Bush, Eds., vol. 7246 of *Lecture Notes in Computer Science*, pp. 74–85, Springer.
- [9] M. S. Nobile, D. Besozzi, P. Cazzaniga, D. Pescini, and G. Mauri, "Reverse engineering of kinetic reaction networks by means of Cartesian Genetic Programming and Particle Swarm Optimization," in *Proceedings of the IEEE Congress on Evolutionary Computation (CEC '13)*, pp. 1594–1601, June 2013.
- [10] M. S. Nobile, "Evolutionary inference of biochemical interaction networks accelerated on Graphics Processing Units," in *Proceedings of the 11th International Conference on High Performance Computing & Simulation 2013 (HPCS '13)*, pp. 668–670, Helsinki, Finland, 2013.
- [11] K. Kentzoglanakis and M. Poole, "A swarm intelligence framework for reconstructing gene networks: searching for biologically plausible architectures," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 9, no. 2, pp. 358–371, 2012.
- [12] J. R. Koza, W. Mydlowec, G. Lanza, Y. Jessen, and M. A. Keane, "Automatic computational discovery of chemical reaction networks using genetic programming," in *Computational Discovery of Scientific Knowledge*, S. Džeroski and L. Todorovski, Eds., vol. 4660 of *Lecture Notes in Computer Science*, pp. 205–227, Springer, 2007.
- [13] S. Ando, E. Sakamoto, and H. Iba, "Evolutionary modeling and inference of gene network," *Information Sciences*, vol. 145, no. 3–4, pp. 237–259, 2002.
- [14] L. Dematté and D. Prandi, "GPU computing for systems biology," *Briefings in Bioinformatics*, vol. 11, no. 3, Article ID bbq006, pp. 323–333, 2010.
- [15] J. L. Payne, N. A. Sinnott-Armstrong, and J. H. Moore, "Exploiting graphics processing units for computational biology and bioinformatics," *Interdisciplinary Sciences: Computational Life Sciences*, vol. 2, no. 3, pp. 213–220, 2010.
- [16] M. J. Harvey and G. De Fabritiis, "A survey of computational molecular science using graphics processing units," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 2, no. 5, pp. 734–742, 2012.
- [17] Y. Zhou, J. Liepe, X. Sheng, M. P. H. Stumpf, and C. Barnes, "GPU accelerated biochemical network simulation," *Bioinformatics*, vol. 27, no. 6, Article ID btr015, pp. 874–876, 2011.
- [18] M. Vigelius, A. Lane, and B. Meyer, "Accelerating reaction-diffusion simulations with general-purpose graphics processing units," *Bioinformatics*, vol. 27, no. 2, Article ID btq622, pp. 288–290, 2011.
- [19] D. J. Wilkinson, "Stochastic modelling for quantitative description of heterogeneous biological systems," *Nature Reviews Genetics*, vol. 10, no. 2, pp. 122–133, 2009.
- [20] J. C. Butcher, "Trees and numerical methods for ordinary differential equations," *Numerical Algorithms*, vol. 53, no. 2–3, pp. 153–170, 2010.
- [21] D. T. Gillespie, "Stochastic simulation of chemical kinetics," *Annual Review of Physical Chemistry*, vol. 58, pp. 35–55, 2007.
- [22] L. Petzold, "Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations," *SIAM Journal of Scientific and Statistical Computing*, vol. 4, no. 1, pp. 136–148, 1983.
- [23] Y. Cao, D. T. Gillespie, and L. R. Petzold, "Efficient step size selection for the tau-leaping simulation method," *The Journal of Chemical Physics*, vol. 124, no. 4, 2006.
- [24] M. S. Nobile, D. Besozzi, P. Cazzaniga, G. Mauri, and D. Pescini, "cupSODA: a CUDA-powered simulator of mass-action kinetics," in *Proceedings of the 12th International Conference on Parallel Computing Technologies (PaCT '13)*, V. Malyskhin, Ed., vol. 7979 of *Lecture Notes in Computer Science*, pp. 344–357.
- [25] M. S. Nobile, P. Cazzaniga, D. Besozzi, and G. Mauri, "GPU-accelerated simulations of mass-action kinetics models with cupSODA," *The Journal of Supercomputing*, 2014.
- [26] M. S. Nobile, P. Cazzaniga, D. Besozzi, D. Pescini, and G. Mauri, "cuTauLeaping: a GPU-powered tau-leaping stochastic simulator for massive parallel analyses of biological systems," *PLoS ONE*, vol. 9, no. 3, 2014.
- [27] K. Rogers, *Blood: Physiology and Circulation*, The Human Body, Britannica Educational Publishing, 2010.
- [28] A. D. Michelson, *Platelets*, Elsevier, 2nd edition, 2007.
- [29] R. W. Colman, J. Hirsh, V. J. Marder, A. W. Clowes, and J. N. George, *Hemostasis and Thrombosis: Basic Principles and Clinical Practice*, Lippincott Williams & Wilkins, 2006.
- [30] M. S. Chatterjee, W. S. Denney, H. Jing, and S. L. Diamond, "Systems biology of coagulation initiation: Kinetics of thrombin generation in resting and activated human blood," *PLoS Computational Biology*, vol. 6, no. 9, Article ID e1000950, 2010.
- [31] M. Bellini, D. Besozzi, P. Cazzaniga, G. Mauri, and M. S. Nobile, "Simulation and analysis of the blood coagulation cascade accelerated on GPU," in *Proceedings of the 22nd Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, pp. 590–593, IEEE Computer Society Conference Publishing Services, 2014.
- [32] T. Wajima, G. K. Isbister, and S. B. Duffull, "A comprehensive model for the humoral coagulation network in humans," *Clinical Pharmacology and Therapeutics*, vol. 86, no. 3, pp. 290–298, 2009.
- [33] K. Lo, W. S. Denney, and S. L. Diamond, "Stochastic modeling of blood coagulation initiation," *Pathophysiology of Haemostasis and Thrombosis*, vol. 34, no. 2–3, pp. 80–90, 2006.
- [34] C. M. Danforth, T. Orfeo, K. G. Mann, K. E. Brummel-Ziedins, and S. J. Everse, "The impact of uncertainty in a blood coagulation model," *Mathematical Medicine and Biology*, vol. 26, no. 4, Article ID dqp011, pp. 323–336, 2009.
- [35] S. D. Lewis, P. P. Shields, and J. A. Shafer, "Characterization of the kinetic pathway for liberation of fibrinopeptides during assembly of fibrin," *The Journal of Biological Chemistry*, vol. 260, no. 18, pp. 10192–10199, 1985.
- [36] M. F. Hockin, K. C. Jones, S. J. Everse, and K. G. Mann, "A model for the stoichiometric regulation of blood coagulation," *The Journal of Biological Chemistry*, vol. 277, no. 21, pp. 18322–18333, 2002.
- [37] K. Brummel-Ziedins, C. Y. Vossen, F. R. Rosendaal, K. Umezaki, and K. G. Mann, "The plasma hemostatic proteome: thrombin

- generation in healthy individuals,” *Journal of Thrombosis and Haemostasis*, vol. 3, no. 7, pp. 1472–1481, 2005.
- [38] K. E. Brummel-Ziedins, C. Y. Vossen, S. Butenas, K. G. Mann, and F. R. Rosendaal, “Thrombin generation profiles in deep venous thrombosis,” *Journal of Thrombosis and Haemostasis*, vol. 3, no. 11, pp. 2497–2505, 2005.
- [39] S. D. Bungay, P. A. Gentry, and R. D. Gentry, “A mathematical model of lipid-mediated thrombin generation,” *Mathematical Medicine and Biology*, vol. 20, no. 1, pp. 105–129, 2003.
- [40] S. Butenas, T. Orfeo, M. T. Gissel, K. E. Brummel, and K. G. Mann, “The significance of circulating factor IXa in blood,” *The Journal of Biological Chemistry*, vol. 279, no. 22, pp. 22875–22882, 2004.
- [41] D. Voet and J. G. Voet, *Biochemistry*, Wiley, 2011.
- [42] A. C. Guyton and J. E. Hall, *Textbook of Medical Physiology*, Saunders, 2010.
- [43] M. F. M. James and A. M. Roche, “Dose-response relationship between plasma ionized calcium concentration and thrombelastography,” *Journal of Cardiothoracic and Vascular Anesthesia*, vol. 18, no. 5, pp. 581–586, 2004.
- [44] D. L. Nelson and M. M. Cox, *Lehninger Principles of Biochemistry*, W. H. Freeman Company, 2004.
- [45] R. M. Farber, “Topical perspective on massive threading and parallelism,” *Journal of Molecular Graphics and Modelling*, vol. 30, pp. 82–89, 2011.
- [46] O. Wolkenhauer, M. Ullah, W. Kolch, and K.-H. Cho, “Modeling and simulation of intracellular dynamics: choosing an appropriate framework,” *IEEE Transactions on Nanobioscience*, vol. 3, no. 3, pp. 200–207, 2004.
- [47] B. H. Estridge, A. P. Reynolds, and N. J. Walters, *Basic Medical Laboratory Techniques*, Cengage Learning, 4th edition, 2000.
- [48] T. Renné, A. H. Schmaier, K. F. Nickel, M. Blombäck, and C. Maas, “In vivo roles of factor XII,” *Blood*, vol. 120, no. 22, pp. 4296–4303, 2012.
- [49] P. V. Jenkins, O. Rawley, O. P. Smith, and J. S. O’Donnell, “Elevated factor VIII levels and risk of venous thrombosis,” *British Journal of Haematology*, vol. 157, no. 6, pp. 653–663, 2012.
- [50] C. E. Walsh and K. M. Batt, “Hemophilia clinical gene therapy: brief review,” *Translational Research*, vol. 161, no. 4, pp. 307–312, 2013.
- [51] A. Van Hylckama Vlieg, I. K. Van Der Linden, R. M. Bertina, and F. R. Rosendaal, “High levels of factor IX increase the risk of venous thrombosis,” *Blood*, vol. 95, no. 12, pp. 3678–3682, 2000.
- [52] K. Brummel-Ziedins, “Models for thrombin generation and risk of disease,” *Journal of Thrombosis and Haemostasis*, vol. 11, no. 1, pp. 212–223, 2013.
- [53] U. Seligsohn and A. Lubetsky, “Genetic susceptibility to venous thrombosis,” *The New England Journal of Medicine*, vol. 344, no. 16, pp. 1222–1231, 2001.
- [54] I. Martinelli, P. Bucciarelli, and P. M. Mannucci, “Thrombotic risk factors: basic pathophysiology,” *Critical Care Medicine*, vol. 38, no. 2, pp. S3–S9, 2010.
- [55] K. Brummel-Ziedins, A. Undas, T. Orfeo et al., “Thrombin generation in acute coronary syndrome and stable coronary artery disease: dependence on plasma factor composition,” *Journal of Thrombosis and Haemostasis*, vol. 6, no. 1, pp. 104–110, 2008.
- [56] S. Hoops, R. Gauges, C. Lee et al., “COPASI—a complex pathway simulator,” *Bioinformatics*, vol. 22, no. 24, pp. 3067–3074, 2006.
- [57] Nvidia, NVIDIA’s Next Generation CUDA Compute Architecture: Kepler GK110, 2012.
- [58] W. Stortelder, W. Hemker, and H. Hemker, “Mathematical modelling in blood coagulation: simulation and parameter estimation,” *Modelling, Analysis and Simulation*, vol. 20, pp. 1–11, 1997.
- [59] S. M. C. Orfao, G. Jank, and K. Mottaghy, “A new mathematical approach to modelling thrombin generation,” *International Journal of Artificial Organs*, vol. 29, no. 7, pp. 703–708, 2006.
- [60] Nvidia, Nvidia CUDA C Programming Guide v5.0, 2012.