

Symbolic Computation of Differential Equivalences

Luca Cardelli

Microsoft Research and University of Oxford, UK
luca@microsoft.com

Mirco Tribastone Max Tschaikowski

Andrea Vandin

IMT Lucca, Italy
{name.surname}@imtlucca.it



Abstract

Ordinary differential equations (ODEs) are widespread in many natural sciences including chemistry, ecology, and systems biology, and in disciplines such as control theory and electrical engineering. Building on the celebrated molecules-as-processes paradigm, they have become increasingly popular in computer science, with high-level languages and formal methods such as Petri nets, process algebra, and rule-based systems that are interpreted as ODEs.

We consider the problem of comparing and minimizing ODEs automatically. Influenced by traditional approaches in the theory of programming, we propose *differential equivalence relations*. We study them for a basic intermediate language, for which we have decidability results, that can be targeted by a class of high-level specifications. An ODE implicitly represents an uncountable state space, hence reasoning techniques cannot be borrowed from established domains such as probabilistic programs with finite-state Markov chain semantics. We provide novel *symbolic* procedures to check an equivalence and compute the largest one via partition refinement algorithms that use satisfiability modulo theories.

We illustrate the generality of our framework by showing that differential equivalences include (i) well-known notions for the minimization of continuous-time Markov chains (lumpability), (ii) bisimulations for chemical reaction networks recently proposed by Cardelli *et al.*, and (iii) behavioral relations for process algebra with ODE semantics. With a prototype implementation we are able to detect equivalences in biochemical models from the literature that cannot be reduced using competing automatic techniques.

Categories and Subject Descriptors F.1.1 [Models of Computation]: Relations between models

General Terms Theory, Algorithms

Keywords Quantitative Equivalence Relations, Satisfiability Modulo Theory, Ordinary Differential Equations, Partition Refinement

1. Introduction

Ordinary differential equations (ODEs) are a widespread mathematical model to describe the time-course evolution of systems that can be characterized by continuously varying quantities. Classical

examples are concentrations of species in chemical reactions and in biological processes, pressure and temperature in a plant, and voltage and current in an electrical circuit. Much more recently there has been an increasing attention to quantitative models of computation based on ODEs, for example to use formal languages to describe biochemical models [11, 14, 17, 21, 26, 56, 60, 66] or as a deterministic approximation for languages with stochastic semantics [21, 42, 76].

In this paper we consider the fundamental problem of automatically comparing and minimizing programs with ODE semantics. From a mathematical viewpoint, the models of our interest are systems of coupled equations in n variables, $x = (x_1, \dots, x_n)$, where

$$\dot{x}_i = f_i(x), \quad i = 1, \dots, n,$$

and f_i is the *drift*, a real valued function giving the derivative with respect to time of variable x_i when the system's state is x .

IDOL. We study a basic formalism called IDOL—Intermediate Drift Oriented Language. It essentially gives a syntax for the drifts, covering a class of nonlinear ODEs for which the reasoning is decidable, hence amenable to automatic treatment. Although not every ODE system can be written with IDOL (e.g., ODEs with trigonometric functions or exponentials), it covers a wide range of ODE models including:

- **Linear ODE systems.** This is a very important class of models in many disciplines including control theory and electrical engineering. Here we remark that a continuous-time Markov chain (CTMC), a very popular stochastic semantics for higher-level quantitative languages (see [10] and references therein) can also be directly seen as a linear ODE system through its Kolmogorov equations (also called the *master equation*). These equations give the probability of being in each state of the chain at any point in time [62].
- **Chemical reaction networks.** Chemical reaction networks (CRNs) express interactions between chemical species or molecular compounds. IDOL allows to specify relevant (nonlinear) kinetics such as the well-known *law of mass action*, where the reaction rate is proportional to the product of the concentrations of the reagents; and the Hill kinetics, which involves rational expressions of polynomials in the species variables [80].
- **Quantitative models of computing systems.** Some formal methods with ODE semantics such as Petri nets [29] and process algebra [42, 76] have nonlinear laws of interaction based on threshold-like functions to model resource contention. For instance, these are used to model the firing rate of transition in a Petri net as being proportional to the *minimum* among the number of tokens at its incoming places [29].

Relating IDOL programs. We cast the problem of relating IDOL programs into the traditional context of equivalences for more classical models of computation based on labeled transition systems

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive version was published in the following publication:

POPL'16, January 20–22, 2016, St. Petersburg, FL, USA
© 2016 ACM. 978-1-4503-3549-2/16/01...
<http://dx.doi.org/10.1145/2837614.2837649>

(LTS). We put forward the analogy between states of an LTS and IDOL variables. Thus, our equivalences are between variables, (exactly) preserving their ODE solutions in some appropriate sense. We propose two variants of *differential equivalence*.

The first variant is *forward differential equivalence* (FDE). This is such that an ODE system can be written for the variables that represent the equivalence classes. To be more concrete, let us consider the following trivial, yet illustrative, ODE example:

$$\dot{x}_1 = -x_1, \quad \dot{x}_2 = k_1 \cdot x_1 - x_2, \quad \dot{x}_3 = k_2 \cdot x_1 - x_3, \quad (1)$$

where k_1 and k_2 are constants. Then, it turns out that there is an FDE relating x_2 and x_3 . Indeed, we have

$$\dot{x}_1 = -x_1, \quad (x_2 + x_3) = \dot{x}_2 + \dot{x}_3 = (k_1 + k_2) \cdot x_1 - (x_2 + x_3).$$

By the change of variable $y = x_2 + x_3$, this is equivalent to writing

$$\dot{x}_1 = -x_1 \quad \dot{y} = (k_1 + k_2) \cdot x_1 - y.$$

This *quotient ODE model* recovers the sum of the solutions of the variables in each equivalence class. That is, we have that setting the *initial condition* $y(0) = x_2(0) + x_3(0)$ yields that $y(t) = x_2(t) + x_3(t)$ at all time points t .

Our second variant is *backward differential equivalence* (BDE). It equates variables that have the same solutions if they start from the same initial conditions. In (1), it can be shown that x_2 and x_3 are related also by BDE when $k_1 = k_2$. In this case, we obtain a quotient ODE by removing either equation, say x_3 , and rewriting every occurrence of x_3 into x_2 :

$$\dot{x}_1 = -x_1 \quad \dot{x}_2 = k_1 x_1 - x_2.$$

If one starts with $x_2(0) = x_3(0)$ in (1) then the solution of the quotient ODE gives that $x_2(t) = x_3(t)$ at all time points t .

Since in BDE every variable in the same equivalence class has the same solution, the original model can be fully recovered. On the other hand, from the quotient FDE model one cannot recover the original solutions, but FDE poses no restriction on the initial conditions.

Checking and computing IDOL equivalences. An IDOL variable corresponds to a real function, thus it represents a *continuous state space*: proving two IDOL variables equivalent concerns relating two real-valued functions for all possible assignments—which involves reasoning over an uncountable state space. A major consequence is that established techniques for checking and computing equivalence relations over models based on LTSs with discrete state spaces (e.g., [6, 45, 48, 51, 64]) do not carry over.

We tackle this problem by proposing a *symbolic* approach based on satisfiability modulo theories (SMT) [7]. We encode differential equivalences into satisfiability problems of quantifier-free first-order logic formulae containing IDOL terms. Checking candidate relations amounts to establishing their validity, as usual through the unsatisfiability of their negation. The SMT solver that we use, the well-known Z3 [30], is a *decision procedure* for such formulae, so it can answer whether or not they are valid.

More importantly, we provide an automatic technique to compute the *largest* differential equivalence for an IDOL model, which is very relevant for minimization because it yields the smallest quotient ODE system. We do this by developing a partition refinement algorithm (cf. [64]) to which we introduce two novelties.

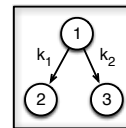
For FDE, we are able to establish a key technical result. We start from its classic definition in terms of a linear transformation of the ODE variables that preserves the aggregated dynamics, e.g., [2, 63, 74]. This definition requires to check “higher-order” properties, i.e., it involves (partition) blocks of variables, instead of individual variables. Thus, if a block does not satisfy the FDE condition, no information can be derived about how to split the current candidate partition. Instead, we equivalently verify FDE in

terms of checks that involve only two IDOL variables at a time; that is, we characterize FDE in a form that does enable partition refinement.

For BDE, the partition refinement is counter-example guided. We fully exploit the ability of an SMT solver to produce an assignment of the variables that falsifies the assertion that a candidate partition is a BDE. The algorithm splits partition blocks according to such assignment; the iterative procedure terminates with the coarsest BDE partition when a distinguishing assignment is not found.

Applications. Many apparently unrelated formalisms and languages can benefit from the common framework provided by IDOL. In order to support our claim we consider three applications: CTMCs, CRNs, and process algebras.

Continuous-time Markov chains. The properties captured by FDE and BDE are analogous to ordinary and exact lumpability for CTMCs [12], respectively, and many behavioral equivalences for higher-level languages based on these notions (e.g., [10, 13, 27, 31, 37, 44, 46, 57, 71]). Indeed, we use the terms “forward” and “backward” to align with the terminology used in some of this literature (e.g., [20, 37, 71]). Actually, we show that FDE and BDE correspond to their respective variants of lumpability when the IDOL program is a linear ODE system representing a CTMC. For instance, the ODEs (1) are the Kolmogorov equations of the simple CTMC with state-transition diagram in the inset below, where x_i is the probability of finding the process in state i . However our differential equivalences are more general: they do not require k_1 and k_2 to be nonnegative, and can establish relations also for nonlinear ODEs. For instance, using the nonlinear ODE $\dot{x}_1 = -x_1^2$ in (1) we would equate x_2 and x_3 in the forward as well as in backward sense.



Chemical reaction networks. CRNs have received increased attention in computer science due to the powerful analogy between computational processes and biological systems [19, 38, 68]. In addition to being a relevant model *per se*, CRNs are also closely related to many other languages. Cardelli establishes a correspondence between his Chemical Ground Form and CRNs [17]; rule-based languages such as κ [26] and BioNetGen [11] provide compact descriptions of biomolecular systems that can be “compiled down” to CRNs; Petri nets with an appropriate mass-action semantics on the transitions correspond to CRNs (e.g., [43]).

The idea of formally relating the dynamics of CRNs has recently emerged. In [18] Cardelli presents the notion of *emulation* between two CRNs as a property that exactly relates the ODE trajectories of a source CRN to those of a target CRN. Syntactic conditions are given to establish an emulation under the assumption of mass-action kinetics. Here we show that BDE is more general than emulation. As an application, we find that the emulations found for a class of biological processes in [18] are preserved even when an alternative dynamics based on the Hill kinetics is considered. This reinforces the findings in [18] that networks with different biological functionality are indeed related structurally, in a way that is insensitive to the underlying kinetics assumed.

In [20] Cardelli *et al.* present *forward* and *backward* bisimulations. The intent is analogous to ours, but those are equivalences that can be detected syntactically by inspecting the set of reactions. However, they only apply to a class of CRNs with mass-action semantics (namely, reactions with at most two reagents). Under these restrictions we show that forward bisimulation is only a sufficient condition for FDE; instead, backward bisimulation corresponds to BDE. Furthermore, in [20] a polynomial time partition refinement algorithm to compute the largest bisimulations is provided. We use benchmarks including those from [20] to present two findings of experimental nature: (i) FDE minimizes CRNs of biological interest that cannot be reduced by forward bisimulation; and (ii) the

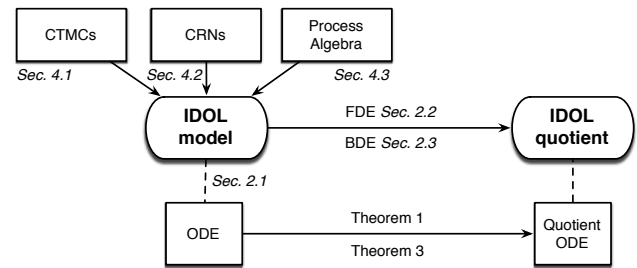


Figure 1. Paper overview.

symbolic implementation of BDE may empirically provide faster answers than the specialized partition-refinement algorithm.

Process algebra. Our last application is a fragment of Hillston’s PEPA [46]. This is a stochastic process algebra which has been more recently equipped also with an ODE semantics with nonlinear minimum-based drifts that approximate the average evolution of underlying CTMCs with massively parallel computations [42, 47, 76]. We take PEPA as the representative of a family of languages for the quantitative evaluation of computing systems — indeed it is expressive enough to cover the semantics of a popular class of queuing networks [75] and Petri nets [39]. We show that recently proposed behavioral equivalences for PEPA [49, 77] are special, language-specific cases of our differential equivalences.

Summary. Figure 1 provides a pictorial representation of the structure and main results of this paper. All the CRN experiments herein reported are replicable using a prototypal implementation available at <http://sysma.imtlucca.it/erode>.

Further related work. We are not aware of general automated approaches to ODE equivalences as done in this paper, though there is a large literature of techniques in domain-specific situations. The combinatorial explosion of CRN biochemical models has spurred considerable research in this area, e.g., [15, 16, 20, 24, 25, 28, 35, 36]. The fragmentation approach for κ identifies a coarse-grained ODE system for models with mass-action semantics through sums of variables; this is weaker than an equivalence relation over species, because one variable may appear in more than one block (a *fragment*) [28, 36]. Using the terminology of [63], fragmentation is a form of *improper lumping*, as opposed to our differential equivalences where species belong to a single block. As such it can still be seen as an ODE aggregation obtained through a transformation of the variables by a linear matrix, for which the general theory is well established (see [2, 63, 74]) but no general algorithms for computing the largest equivalences are available.

SMT has become a cornerstone in the programming languages and in the verification community, with contributions to program synthesis [41], constraint programming [54], and symbolic optimization [58]. The combination of SMT and equivalence relations has been the subject of recent investigations. In [9] partition-refinement algorithms are proposed to compute equivalences between terms over arbitrary theories inferred from a set of axioms. Applied to our context, these partition-refinement algorithms could be used to check if a candidate partition is a differential equivalence, but not to compute the largest equivalence for an IDOL program. In [32] the authors present an SMT-based approach for the computation of the coarsest ordinary lumpable partition of a Markov chain, but for a fragment of the PRISM language [55].

Finally, links between ODEs and SMT are established in the formal verification community, especially for hybrid systems (e.g., [40, 59, 69]); however none of these works considers ODE comparisons and minimizations through equivalence relations. Still at the interface between control theory and computer science, the idea of

bisimulation for dynamical systems has been developed in a series of works by Pappas and coauthors [65] and van der Schaft [79]. These works are similar in spirit to ours, but the setting is different because the focus is on *control systems*, i.e., dynamical systems with internal states, external inputs, and output maps. In that context, bisimulation relates internal states mapped to the same output, i.e., they cannot be told apart by an external observer. The largest bisimulation is therefore related to the maximal unobservability subspace of a control system (e.g., [79, Corollary 6.4]) while our largest differential equivalences provide the coarsest partition of ODE variables that preserves the dynamics.

2. Intermediate Drift Oriented Language

We first introduce IDOL as a language to define a class of ODEs. More precisely, IDOL can describe nonlinear, first-order, autonomous, and explicit finite systems of coupled ODEs. Then, we present the notions of FDE and BDE as equivalence relations over IDOL variables, with their characterizations in terms of properties enjoyed by the underlying semantics.

2.1 Syntax and Semantics

Definition 1 (IDOL syntax). *The syntax of programs of the intermediate drift oriented language (IDOL) is given by*

$$p ::= \varepsilon \mid \dot{x}_i = f, \quad p$$

$$f ::= n \mid x_i \mid f + f \mid f \cdot f \mid f^{\frac{1}{m}},$$

where $x_i \in \mathbb{V}$ and $n, m \in \mathbb{Z}$ and $m \neq 0$.

The set \mathbb{V} represents ODE variables. A program is a list of elements $\dot{x}_i = f$ where each element gives the drift f for ODE of the variable x_i . Given an IDOL program p , we define $\mathcal{V}_p = \{x_1, \dots, x_n\}$ as the set of variables in p . We say that p is *well-formed* if for every $x_i \in \mathcal{V}_p$ there exists a unique term $\dot{x}_i = f$ in p . We denote its drift by f_i . From now on we assume to work with well-formed programs only. Finally, we remark that the restriction to integer parameters allows us to encode rationals, which is without loss of generality in practice (e.g., [4]).

Although some results presented below hold for richer classes of ODE systems, drifts expressible in IDOL form a class for which our differential equivalences are decidable. Despite the minimality of IDOL, it is possible to encode frequently used dynamics such as:

- the law of mass action, with drifts such as $x_1 \cdot x_2$;
- the Hill kinetics for CRNs, with drifts such as $x_1^2 / (1 + x_1^2)$;
- and the minimum function for threshold based drifts, where

$$\min(x_1, x_2) := \frac{1}{2}(x_1 + x_2 - |x_1 - x_2|), \quad \text{with } |x| := (x \cdot x)^{\frac{1}{2}}.$$

For the semantics of IDOL, we interpret each term f_i in a standard way, as a real function of real variables on an appropriate domain, $D(f_i) \subseteq \mathbb{R}^{\mathcal{V}_p}$, where the function is *well-defined*, i.e., with no division by zero or negative arguments in roots. We denote by σ an assignment of variables in p , thus $\sigma \in \mathbb{R}^{\mathcal{V}_p}$. The semantics of IDOL depends on a *context*: this is a pair $c = (T, \hat{\sigma})$ that contains a time horizon $T > 0$ and an *initial assignment* $\hat{\sigma}$. The semantics of a program p is a function that maps the variables \mathcal{V}_p to a continuous trajectory $\llbracket x \rrbracket_c^p : [0; T] \rightarrow \mathbb{R}^{\mathcal{V}_p}$ that describes the time course of every variable when starting from a given initial assignment $\hat{\sigma}$. In other words, $\llbracket x_i \rrbracket_c^p(t)$ is the value of the variable x_i at time t when $\llbracket x_i \rrbracket_c^p(0) = \hat{\sigma}(x_i)$.

Definition 2 (IDOL Semantics). *The semantics of an IDOL program p in a context $c = (T, \hat{\sigma})$ is the unique differentiable function*

$$\llbracket x \rrbracket_c^p := (\llbracket x_i \rrbracket_c^p)_{x_i \in \mathcal{V}_p}, \quad \llbracket x \rrbracket_c^p : [0; T] \rightarrow \mathbb{R}^{\mathcal{V}_p}$$

that satisfies

$$\llbracket x_i \rrbracket_c^p(t) = \hat{\sigma}(x_i) + \int_0^t \llbracket f_i \rrbracket_c^p(\llbracket x \rrbracket_c^p(s)) ds, \quad \text{for all } 0 \leq t \leq T,$$

where $\llbracket f_i \rrbracket_c^p : D(f_i) \rightarrow \mathbb{R}$ is recursively defined as follows:

$$\begin{aligned} \llbracket n \rrbracket_c^p : \mathbb{R}^{\mathcal{V}_p} &\rightarrow \mathbb{R}, & \llbracket n \rrbracket_c^p(\sigma) &= n \\ \llbracket x \rrbracket_c^p : \mathbb{R}^{\mathcal{V}_p} &\rightarrow \mathbb{R}, & \llbracket x \rrbracket_c^p(\sigma) &= \sigma(x) \\ \llbracket g + h \rrbracket_c^p : D(g) \cap D(h) &\rightarrow \mathbb{R}, & \llbracket g + h \rrbracket_c^p(\sigma) &= \llbracket g \rrbracket_c^p(\sigma) + \llbracket h \rrbracket_c^p(\sigma) \\ \llbracket g \cdot h \rrbracket_c^p : D(g) \cap D(h) &\rightarrow \mathbb{R}, & \llbracket g \cdot h \rrbracket_c^p(\sigma) &= \llbracket g \rrbracket_c^p(\sigma) \cdot \llbracket h \rrbracket_c^p(\sigma) \\ \llbracket g^{\frac{1}{m}} \rrbracket_c^p : D(g^{\frac{1}{m}}) &\rightarrow \mathbb{R}, & \llbracket g^{\frac{1}{m}} \rrbracket_c^p(\sigma) &= (\llbracket g \rrbracket_c^p(\sigma))^{\frac{1}{m}} \end{aligned}$$

with $D(g^{\frac{1}{m}}) = \{\sigma \in D(g) : (\llbracket g \rrbracket_c^p(\sigma))^{\frac{1}{m}} \text{ is defined}\}$. If no such unique function exists, we call p ill-posed.

As usual we call $\llbracket f \rrbracket_c^p = (\llbracket f_i \rrbracket_c^p)_{x_i \in \mathcal{V}_p}$ the *vector field* of program p in context c . Also, the semantics of an even root term is given by the nonnegative solution; e.g., in any context c , $(\llbracket 4 \rrbracket_c^p)^{1/2}$ is 2 and not -2 . We remark that, in general, no closed-form expressions for $\llbracket x_i \rrbracket_c^p$ exist. However these functions can be computed using standard numerical integration algorithms, cf. [3].

Assumptions. IDOL is permissive enough to define somewhat degenerate ODEs with no solutions like $\dot{x}_1 = x_1^{-1}$, with $\hat{\sigma}(x_1) = 0$, or multiple solutions as in $\dot{x}_1 = |x_1|^{1/2}$, with $\hat{\sigma}(x_1) = -1$. We exclude these cases making certain assumptions that are usual when dealing with ODEs (e.g., [69]). For this, we define a notion of invariance which considers a subset of the drifts' domain containing the trajectories of the IDOL variables starting from any initial condition within that set.

Definition 3. Given a program p and a time horizon $T > 0$, a set $E \subseteq \bigcap_{x_i \in \mathcal{V}_p} D(f_i)$ is invariant with respect to p and T if, for all $\hat{\sigma} \in E$ and $t \in [0, T]$, it holds that $\llbracket x \rrbracket_c^p(t) \in E$, where $c = (T, \hat{\sigma})$.

Now we make the following assumptions:

A1 For a given time horizon $T > 0$, an IDOL program p has an invariant set $E(p) \in \{\mathbb{R}_{>0}^{\mathcal{V}_p}, \mathbb{R}_{\geq 0}^{\mathcal{V}_p}, \mathbb{R}^{\mathcal{V}_p}\}$.

A2 For all $x_i \in \mathcal{V}_p$, the function $\llbracket f_i \rrbracket_c^p : D(f_i) \rightarrow \mathbb{R}$ is locally Lipschitz continuous at any point of $E(p)$.

A1 is a technical assumption that allows us to work with nice enough domains when reasoning about differential equivalences. Our theory can be developed for more general invariant sets, but at the expense of significantly more convoluted mathematical definitions which do not seem to add substantial value to our contribution. Instead, **A2** is a standard textbook condition to ensure the existence of a unique solution, hence to exclude ill-posedness.

In many applications, models are typically such that a) the solution will be positive if the initial condition is positive and b) the drift is well-defined on positive reals. Under such circumstances, local Lipschitz continuity is usually immediate. Indeed, all IDOL programs presented in this paper satisfy these assumptions (and we will avoid stating which invariant set they have).

Notation. Differential equivalences are partitions of IDOL variables. Whenever convenient, for a program p and a given partition \mathcal{H} of \mathcal{V}_p , we write $H = \{x_{H,1}, \dots, x_{H,|H|}\}$ for any $H \in \mathcal{H}$. As usual, we denote by $\psi[t/s]$ the term that arises by replacing each occurrence of t in ψ by s .

2.2 Forward Differential Equivalence

With FDE one can write an IDOL program with one variable for each equivalence class, representing the sum of the trajectory solutions of its members. A partition \mathcal{H} is induced by an FDE if the *aggregated*

drift $\sum_{x_i \in H} f_i$ of any block $H \in \mathcal{H}$ can be written in terms of the sums of the variables within the block $\{\sum_{x_i \in H} x_i : H \in \mathcal{H}\}$. For instance, in the IDOL program

$$\begin{aligned} \dot{x}_1 &= -2x_1 - 3x_2 - 4x_3 & \dot{x}_2 &= -3x_1 - 4x_2 - 5x_3 \\ \dot{x}_3 &= -6x_1 - 4x_2 - 2x_3 & \dot{x}_4 &= x_1 + x_2 + x_3 - 2x_4 \end{aligned}$$

the aggregated drifts for the partition $\{\{x_1, x_2, x_3\}, \{x_4\}\}$ are

$$\begin{aligned} f_1 + f_2 + f_3 &= -11 \cdot (x_1 + x_2 + x_3) + 0 \cdot x_4, \\ f_4 &= 1 \cdot (x_1 + x_2 + x_3) - 2 \cdot x_4. \end{aligned}$$

Clearly they depend only on the values of $x_1 + x_2 + x_3$ and x_4 .

Answering the question whether sums of variables can be factored out from the aggregated drifts means finding new appropriate functions with arity equal to the number of equivalence classes. In this example, we would have drifts g_1 and g_2 defined as

$$g_1 = -11 \cdot y_1 \quad g_2 = 1 \cdot y_1 - 2 \cdot y_2$$

where y_1 and y_2 represent blocks $\{x_1, x_2, x_3\}$ and $\{x_4\}$, respectively. We avoid synthesizing these functions directly by exploiting an alternative characterization that involves reasoning on properties concerning only the original variables: The evaluation of the aggregated drift must be invariant under any change of assignment of the variables that preserves the sum of values across each block.

To do this formally, we rewrite each variable as a scaling of the corresponding sums-of-variables of its block, such that all scaling factors are nonnegative and sum to one; in the example, we rewrite x_1 with $s_1(x_1 + x_2 + x_3)$, x_2 with $s_2(x_1 + x_2 + x_3)$, and $s_3(x_1 + x_2 + x_3)$ with scaling factors s_1 , s_2 , and s_3 . The alternative characterization consists in proving that the aggregated drifts do not depend on the assignments of the scaling factors.

Importantly, following [74] it can be shown that if such rewriting does not change the values of the aggregated drifts for *some* choice of the scaling factors, then *any* choice will enjoy this property. The notion of FDE checks this using the *uniform* scaling that gives equal weight to every variable in the block (for instance $s_1 = s_2 = s_3 = 1/3$ in the example above). By assumption **A1**, the uniform scaling ensures that terms are always rewritten into terms that give rise to well-defined functions.

We encode this property in first order logic with function symbols from IDOL that are interpreted in the standard way, having \mathcal{V}_p as free variables. We denote by $\Theta(p)$ the logical formula that encodes $E(p)$ (e.g., if $E(p) = \mathbb{R}_{>0}^{\mathcal{V}_p}$ then $\Theta(p) := \bigwedge_{x_i \in \mathcal{V}_p} x_i > 0$).

Definition 4 (FDE). Let p be an IDOL program and \mathcal{H} a partition of \mathcal{V}_p . Then, \mathcal{H} is a forward differential equivalence if the following formula is valid:

$$\Theta(p) \rightarrow \bigwedge_{H \in \mathcal{H}} \left(\sum_{x_i \in H} f_i = \sum_{x_i \in H} f_i \left[x_j / \frac{\sum_{x_k \in H'} x_k}{|H'|} : H' \in \mathcal{H}, x_j \in H' \right] \right) \quad (\Phi^{\mathcal{H}})$$

The next definition provides the quotient IDOL program with respect to an FDE.

Definition 5 (FDE Quotient). Let p be an IDOL program and \mathcal{H} an FDE partition. Then, the forward quotient of p with respect to \mathcal{H} , denoted by $\overrightarrow{p}_{\mathcal{H}}$, is given by:

$$\dot{y}_H = \sum_{x_i \in H} f_i \left[x_j / \frac{y_{H'}}{|H'|} : H' \in \mathcal{H}, x_j \in H' \right], \quad \text{for all } H \in \mathcal{H}.$$

The uniform scaling and **A1** ensure that $\overrightarrow{p}_{\mathcal{H}}$ is not ill-posed.

We now state a crucial *dynamical characterization* theorem: A partition of IDOL variables is FDE if and only if the ODEs of the

quotient program preserve the sums of the original trajectories in each equivalence class. Hence the largest FDE represents the best possible aggregation that can be obtained in this sense.

Theorem 1 (Dynamical FDE Characterization). *Let p be an IDOL program, $T > 0$ a time horizon and \mathcal{H} a partition of \mathcal{V}_p . Then, \mathcal{H} is an FDE partition with forward quotient $\vec{p}_{\mathcal{H}}$ if and only if*

$$\llbracket y_H \rrbracket_{\vec{c}}^{\vec{p}_{\mathcal{H}}}(t) = \sum_{x_i \in H} \llbracket x_i \rrbracket_c^p(t)$$

for all $\hat{\sigma} \in E(p)$, $H \in \mathcal{H}$ and $t \in [0; T]$, where $c := (T, \hat{\sigma})$, $\vec{c} := (T, \hat{\sigma}_{\mathcal{H}})$ and $\hat{\sigma}_{\mathcal{H}}(y_H) := \sum_{x_i \in H} \hat{\sigma}(x_i)$ for all $H \in \mathcal{H}$.¹

Let us remark that FDE is stated in terms of a partition and is thus consistent with the notion of lumpability for ODEs [74]. This has the advantage that the above theorem is a direct consequence of the theory presented in [74], hence we omit the proof here. However, this is not in a form that enables an algorithm for computing the largest FDE using partition refinement, because an assignment that falsifies the FDE conditions $\Phi^{\mathcal{H}}$ does not provide information about which variables to tell apart in the refinement step.

We tackle this problem by providing a characterization of FDE in terms of *binary* checks, i.e., involving two variables only at a time. Intuitively, for each block $H \in \mathcal{H}$ and any pair $x_i, x_j \in H$, such characterization allows to check if the fact that x_i and x_j belong to the same block prevents \mathcal{H} from being an FDE.

More precisely, an equivalence relation \mathcal{R} over \mathcal{V}_p is FDE if and only if for all $(x_i, x_j) \in \mathcal{R}$ it holds that the aggregated drifts of all blocks in \mathcal{H} are invariant under a scaling of the sum-of-variables which involves *only two* variables belonging to the same block, rather than *all of them* as per Definition 4. The intuition is that any scaling considered in Definition 4 can be equivalently achieved as a composition of such “binary” scalings. To our knowledge, such a binary characterization is provided here for the first time.

Theorem 2 (Binary FDE characterization). *Let p be an IDOL program, \mathcal{R} be an equivalence relation on \mathcal{V}_p , and $\mathcal{H} = \mathcal{V}_p/\mathcal{R}$. Then \mathcal{H} is an FDE if and only if for all distinct $x_i, x_j \in \mathcal{V}_p$ we have that $(x_i, x_j) \in \mathcal{R}$ implies that the following formula is valid:*

$$\Theta(p) \rightarrow \bigwedge_{H \in \mathcal{H}} \left(\sum_{x_k \in H} f_k = \sum_{x_k \in H} f_k [x_i/s \cdot (x_i + x_j), x_j/(1-s) \cdot (x_i + x_j)] \right) \quad (\Phi_{x_i, x_j}^{\mathcal{H}})$$

Proof. Note that s is not an ODE variable from \mathbb{V} but a variable of the first order logic. However, we still denote the interpretation of any variable $\tilde{s} \notin \mathbb{V}$ by $\sigma(\tilde{s})$. Moreover, to increase readability in the proof we shall write $\llbracket \cdot \rrbracket$ instead of $\llbracket \cdot \rrbracket_c^p$.

Let us assume that $\Phi_{x_i, x_j}^{\mathcal{H}}$ is valid for all $H \in \mathcal{H}$ and $x_i, x_j \in H$. We have to show that $\Phi^{\mathcal{H}}$ is valid. For this, we fix arbitrary $\sigma \in E(p)$ and $H \in \mathcal{H}$ and assume without loss of generality that $H = \{x_1, \dots, x_m\}$ and that $\sigma(x_i) \geq \sigma(x_{i+1})$ for all $1 \leq i \leq m-1$. Together with $\mu_H := \frac{1}{|H|} \sum_{1 \leq i \leq m} \sigma(x_i)$ we then define $\sigma^i \in \mathbb{R}^{\mathcal{V}_p}$, where $1 \leq i \leq m-1$, as

$$\sigma(x_k)^{i+1} := \begin{cases} \mu_H & , k = i \\ (1 - \sigma(s_i)) \cdot (\sigma^i(x_i) + \sigma^i(x_{i+1})) & , k = i + 1 \\ \sigma^i(x_k) & , \text{otherwise} \end{cases}$$

with $\sigma^1 := \sigma$ and $\sigma(s_i) := \mu_H / (\sigma^i(x_i) + \sigma^i(x_{i+1}))$ for all $1 \leq i \leq m-1$. Since $\sigma(x_i) \geq \sigma(x_{i+1})$ for all $1 \leq i \leq m-1$, it holds that $0 < \sigma(s_i) \leq 1$ for all $1 \leq i \leq m-1$. Thus, since $\Phi_{x_i, x_j}^{\mathcal{H}}$

is valid for all $x_i, x_j \in H$, we infer that $\llbracket f \rrbracket(\sigma^{i+1}) = \llbracket f \rrbracket(\sigma^i)$ because both terms are equal to

$$\llbracket f [x_i/s_i(x_i + x_{i+1}), x_{i+1}/(1-s_i)(x_i + x_{i+1})] \rrbracket(\tilde{\sigma}^i),$$

where $\tilde{\sigma}^i := \sigma^i \cup \{(s_i, \sigma(s_i))\} \in \mathbb{R}^{\mathcal{V}_p \cup \{s_i\}}$, for all $1 \leq i \leq m-1$. Hence, $\llbracket f \rrbracket(\sigma^{|H|}) = \llbracket f \rrbracket(\sigma^{|H|-1}) = \dots = \llbracket f \rrbracket(\sigma^1) = \llbracket f \rrbracket(\sigma)$. Note also that $\sigma^{|H|}$ satisfies $\sigma^{|H|}(x_k) = \sigma^{|H|}(x_l)$ for all $x_k, x_l \in H$. By applying the above argument to the remaining blocks $\mathcal{H} \setminus \{H\}$ (i.e., in second step we would consider a block $H' \neq H$ and the vector $\sigma^{|H|}$), we infer that $\Phi^{\mathcal{H}}$ is true under the assignment σ . Since σ was chosen arbitrarily, \mathcal{H} is an FDE.

For the proof of the converse, we first consider the case $E(p) = \mathbb{R}_{>0}^{\mathcal{V}_p}$ and define, for any partition \mathcal{H} of \mathcal{V}_p , the matrix $M_{\mathcal{H}} \in \{0, 1\}^{\mathcal{H} \times \mathcal{V}_p}$ by setting $(M_{\mathcal{H}})_{H, x_k}$ to 1 if $x_k \in H$, and 0 otherwise. The matrix $M_{\mathcal{H}}$ can be thought of as an “aggregation” matrix. In particular, the rows of the matrix $M_{\mathcal{H}}$ encode the blocks of \mathcal{H} . Then, for any positive generalized right inverse of $M_{\mathcal{H}}$, i.e. a matrix $\bar{M}_{\mathcal{H}} \in (0, 1]^{\mathcal{V}_p \times \mathcal{H}}$ that satisfies $M_{\mathcal{H}} \bar{M}_{\mathcal{H}} = \mathbb{I}$, the function $\sigma \mapsto M_{\mathcal{H}} \bar{M}_{\mathcal{H}} \sigma$ defines a scaling on \mathcal{H} . Since the entries of $\bar{M}_{\mathcal{H}}$ are positive, we infer $M_{\mathcal{H}} \bar{M}_{\mathcal{H}} \sigma \in \mathbb{R}_{>0}^{\mathcal{V}_p}$ for all $\sigma \in \mathbb{R}_{>0}^{\mathcal{V}_p}$. As pointed out at the beginning of Section 2.2, any scaling that yields well-defined terms is equivalent to the notion of FDE [74]. Consequently, \mathcal{H} is an FDE if and only if there exists a generalized right inverse $\bar{M}_{\mathcal{H}} \in (0, 1]^{\mathcal{V}_p \times \mathcal{H}}$ of $M_{\mathcal{H}}$ such that $M_{\mathcal{H}}(\llbracket f \rrbracket(\sigma)) = M_{\mathcal{H}}(\llbracket f \rrbracket(\bar{M}_{\mathcal{H}} M_{\mathcal{H}} \sigma))$ for all $\sigma \in \mathbb{R}_{>0}^{\mathcal{V}_p}$. With this in mind, let us now assume that \mathcal{H} is an FDE partition and fix arbitrary $\sigma \in E(p)$, $H_0 \in \mathcal{H}$, $x_i, x_j \in H_0$ and $\sigma(s) \in (0, 1]$. We next show that $\Phi_{x_i, x_j}^{\mathcal{H}}$ is true for the assignment σ . Fix the generalized right inverse

$$(\bar{M}_{\mathcal{H}})_{x_k, H} = \begin{cases} \frac{\sigma(s)(\sigma(x_i) + \sigma(x_j))}{\sum_{x_l \in H_0} \sigma(x_l)} & , x_k = x_i \\ \frac{(1 - \sigma(s))(\sigma(x_i) + \sigma(x_j))}{\sum_{x_l \in H_0} \sigma(x_l)} & , x_k = x_j \\ \sigma(x_k) / (\sum_{x_l \in H} \sigma(x_l)) & , x_k \notin \{x_i, x_j\} \end{cases}$$

It is straightforward to show that

$$(M_{\mathcal{H}} \bar{M}_{\mathcal{H}} \sigma)(x_k) = \begin{cases} \sigma(s)(\sigma(x_i) + \sigma(x_j)) & , x_k = x_i \\ (1 - \sigma(s))(\sigma(x_i) + \sigma(x_j)) & , x_k = x_j \\ \sigma(x_k) & , \text{otherwise} \end{cases}$$

Since \mathcal{H} was assumed to be an FDE, the above discussion implies that $M_{\mathcal{H}}(\llbracket f \rrbracket(\sigma)) = M_{\mathcal{H}}(\llbracket f \rrbracket(\bar{M}_{\mathcal{H}} M_{\mathcal{H}} \sigma))$. This, however, implies that $\Phi_{x_i, x_j}^{\mathcal{H}}$ holds true for the assignment σ . Since $\sigma \in E(p)$, $H_0 \in \mathcal{H}$ and $x_i, x_j \in H_0$ were chosen arbitrarily, we infer the claim in the case where $E(p) = \mathbb{R}_{>0}^{\mathcal{V}_p}$.

The case $E(p) \in \mathbb{R}^{\mathcal{V}_p}$ (resp. $E(p) \in \mathbb{R}_{\geq 0}^{\mathcal{V}_p}$) follows by generalizing our argumentation. In particular, the generalized right inverse has to be chosen from $\mathbb{R}^{\mathcal{V}_p \times \mathcal{H}}$ (resp. $[0, 1]^{\mathcal{V}_p \times \mathcal{H}}$). Moreover, note that $(\bar{M}_{\mathcal{H}})_{x_k, H}$ is well-defined only if all aggregated variables underlying the fixed $\sigma \in \mathbb{R}^{\mathcal{V}_p}$ are nonzero. However, such assignments build a (Lebesgue) zero set of $E(p)$ and the vector field $\llbracket f \rrbracket$ is continuous on $E(p)$, which yields the claim. \square

2.3 Backward Differential Equivalence

BDE relates IDOL variables having the same semantics whenever they are given the same initial assignment. This property is characterized by the following implication: if the variables in each block of the partition have the same equal assignments, then the drifts of any two variables of a block have equal values. Similarly to FDE, we formalize this in first order logic.

Definition 6 (BDE). *Let p be an IDOL program and \mathcal{H} a partition of \mathcal{V}_p . Then \mathcal{H} is a backward differential equivalence if the following formula is valid:*

¹ Proofs not included in this paper are provided in a technical report available at the authors’ web pages.

$$\begin{aligned} \Theta(p) &\rightarrow \left(\bigwedge_{H \in \mathcal{H}} (x_{H,1} = \dots = x_{H,|H|}) \right. \\ &\quad \left. \rightarrow \bigwedge_{H \in \mathcal{H}} (f_{H,1} = \dots = f_{H,|H|}) \right) \quad (\Psi^{\mathcal{H}}) \end{aligned}$$

For instance, let us consider the IDOL program

$$\dot{x}_1 = -\min(x_1, 1) + x_2 \quad \dot{x}_2 = -\min(x_2, 1) + x_1 \quad (2)$$

We seek to verify that $\{\{x_1, x_2\}\}$ is a BDE partition. Indeed this holds since $\Psi^{\mathcal{H}}$ becomes

$$x_1 = x_2 \rightarrow -\min(x_1, 1) + x_2 = -\min(x_2, 1) + x_1$$

Definition 7 (BDE Quotient). *Let p be an IDOL program and \mathcal{H} a BDE partition of \mathcal{V}_p . The backward quotient of p with respect to \mathcal{H} , denoted by $\overleftarrow{p}_{\mathcal{H}}$, is given by*

$$\dot{y}_H = f_{H,1}[x_{H',1}/y_{H'}, \dots, x_{H',|H'|}/y_{H'} : H' \in \mathcal{H}], \text{ for } H \in \mathcal{H}.$$

Similarly to FDE, the BDE quotient is not ill-posed. For instance, the BDE quotient of (2) with respect to $\{\{x_1, x_2\}\}$ is

$$\dot{y} = (-\min(x_1, 1) + x_2)[x_1/y, x_2/y], \text{ i.e., } \dot{y} = -\min(y, 1) + y$$

The next characterization result is analogous to Theorem 1.

Theorem 3 (Dynamical BDE Characterization). *Let p be an IDOL program, $T > 0$ a time horizon and \mathcal{H} a partition of \mathcal{V}_p . Then, \mathcal{H} is a BDE partition with backward quotient $\overleftarrow{p}_{\mathcal{H}}$ if and only if $\hat{\sigma}_{\mathcal{H}}(y_H) = \hat{\sigma}(x_{H,1}) = \dots = \hat{\sigma}(x_{H,|H|})$ for all $H \in \mathcal{H}$ implies*

$$\llbracket y_H \rrbracket_{\tilde{c}}^{\overleftarrow{p}_{\mathcal{H}}}(t) = \llbracket x_{H,1} \rrbracket_c^p(t) = \dots = \llbracket x_{H,|H|} \rrbracket_c^p(t)$$

for all $H \in \mathcal{H}$ and $t \in [0; T]$, with $c := (T, \hat{\sigma})$ and $\tilde{c} := (T, \hat{\sigma}_{\mathcal{H}})$.

The statement is shown by using the same strategy as in the proof of Theorem 6 in [20].

Let us point out that the notions of FDE and BDE are not comparable. For instance, the partition $\{\{x_1, x_2\}\}$ is not an FDE of (2) because the formula

$$\begin{aligned} &-\min(x_1, 1) + x_2 - \min(x_2, 1) + x_1 = \\ &-\min\left(\frac{x_1 + x_2}{2}, 1\right) + \frac{x_1 + x_2}{2} - \min\left(\frac{x_1 + x_2}{2}, 1\right) + \frac{x_1 + x_2}{2} \end{aligned}$$

is not true for the assignment $\sigma(x_1) = 2$ and $\sigma(x_2) = 0$. Conversely, $\{\{x_1\}, \{x_2, x_3\}\}$ is an FDE partition of (1) for any choice of k_1 and k_2 , but a BDE only if $k_1 = k_2$.

3. Computing Differential Equivalences

We now discuss how to compute differential equivalences and how to implement this using SMT. We first consider the problem of checking if a given partition is a differential equivalence. Then we focus on computing the largest differential equivalence for an IDOL program using partition refinement.

3.1 Checking Differential Equivalences

Tarski's famous result ensures that one can decide whether $\Phi^{\mathcal{H}}$, $\Phi_{x_i, x_j}^{\mathcal{H}}$ and $\Psi^{\mathcal{H}}$ are valid because the functions supported by IDOL can be expressed in the theory of reals $(\mathbb{R}, +, -, \cdot, 0, 1, <, =)$.

For instance, terms with roots like $y = x^{\frac{1}{2}}$ can be encoded as $\exists y(y^2 = x)$, while the encoding of fractions is straightforward. However, no efficient computation is possible in general.

Proposition 1. *Deciding a differential equivalence is coNP-hard.*

Despite this, in many cases the computation is feasible in practice. We provide examples in Section 4.2.2. Here we briefly discuss how an SMT solver can be used for this purpose. The

validity of the quantifier-free formulae $\Phi^{\mathcal{H}}$, $\Phi_{x_i, x_j}^{\mathcal{H}}$ and $\Psi^{\mathcal{H}}$ can be encoded, as usual, into the unsatisfiability problem of their negation, i.e., by invoking $\text{sat}(\neg\Phi^{\mathcal{H}})$, $\text{sat}(\neg\Phi_{x_i, x_j}^{\mathcal{H}})$, and $\text{sat}(\neg\Psi^{\mathcal{H}})$. These can be decided using the decision procedure `nlSAT` [50], which is implemented in Z3 v4.0 [30]. Thus checking differential equivalences is sound and complete using state-of-the-art SMT technology: A partition \mathcal{H} is FDE (resp., BDE) if and only if $\text{sat}(\neg\Phi^{\mathcal{H}})$ (resp., $\text{sat}(\neg\Psi^{\mathcal{H}})$) returns “unsatisfiable”. As a concrete example, consider the IDOL program (1) and the partition $\mathcal{H} = \{\{x_1\}, \{x_2, x_3\}\}$, which, as discussed, is a BDE if and only if the parameters k_1 and k_2 are equal. The executable Z3 encoding of $\neg\Psi^{\mathcal{H}}$ for both the cases $k_1 \neq k_2$ and $k_1 = k_2$ is available at <http://rise4fun.com/Z3/PmJtS>.

3.2 Partition Refinement

We compute the largest differential equivalence for an IDOL program using a partition refinement algorithm. First, however, we show that this is a well-posed problem.

Definition 8 (Refinement). *Let S be a set, and $\mathcal{H}_1, \mathcal{H}_2$ two partitions of S . Then, \mathcal{H}_1 is a refinement of \mathcal{H}_2 if for any block $H_1 \in \mathcal{H}_1$ there exists a block $H_2 \in \mathcal{H}_2$ such that $H_1 \subseteq H_2$.*

Theorem 4. *Let p be an IDOL program and \mathcal{G} be a partition of \mathcal{V}_p . Then, there exists a unique coarsest FDE/BDE partition refining \mathcal{G} .*

Proof. For a given partition \mathcal{H} of \mathcal{V}_p , write $x_i \sim_{\mathcal{H}} x_j$ whenever there exists some $H \in \mathcal{H}$ such that $x_i, x_j \in H$.

FDE case. Set $\sim_{\Phi}^{\mathcal{H}} := \{(x_i, x_j) : x_i = x_j \text{ or } \Phi_{x_i, x_j}^{\mathcal{H}} \text{ is valid}\}$ and note that $\sim_{\Phi}^{\mathcal{H}}$ is an equivalence relation on \mathcal{V}_p . We fix FDE partitions $\mathcal{H}_1, \dots, \mathcal{H}_n$ of \mathcal{V}_p and set for the sake of brevity $\sim_i := \sim_{\mathcal{H}_i}$ and $\sim_* := \sim_{\mathcal{H}^*}$ where $\mathcal{H}^* := \mathcal{V}_p / (\bigcup_{i=1}^n \sim_i)^*$ and the asterisk denotes transitive closure. Thanks to Theorem 2, it suffices to prove that $y_1 \sim_{\Phi}^{\mathcal{H}^*} y_2$ for all $H^* \in \mathcal{H}^*$ and $y_1, y_2 \in H^*$. Thus, let us fix some $H^* \in \mathcal{H}^*$ and $y_1, y_2 \in H^*$. Since $\sim_{\Phi}^{\mathcal{H}^*}$ is transitive and $y_1 = x_0 \sim_{i_0} x_1 \sim_{i_1} \dots \sim_{i_{k-1}} x_k = y_2$ for some $x_0, \dots, x_k \in \mathcal{V}_p$ and $i_0, \dots, i_{k-1} \in \{1, \dots, n\}$, it suffices to show that $x_j \sim_{\Phi}^{\mathcal{H}^*} x_{j+1}$ for all $0 \leq j \leq k-1$. For this, let us fix some arbitrary $G^* \in \mathcal{H}^*$. Then, it can be easily seen that there exist (unique) subsets $\{G_1^i, \dots, G_{m_i}^i\} \subseteq \mathcal{H}_i$ such that

$$\biguplus_{l=1}^{m_i} G_l^i = G^*$$

for all $1 \leq i \leq n$. Since $x_j \sim_{i_j} x_{j+1}$ implies that

$$\begin{aligned} \sum_{x_i \in G^*} f_i &= \sum_{l=1}^{m_{i_j}} \sum_{x_i \in G_l^{i_j}} f_i \\ &= \sum_{l=1}^{m_{i_j}} \sum_{x_i \in G_l^{i_j}} f_i[x_j/s(x_j + x_{j+1}), x_{j+1}/(1-s)(x_j + x_{j+1})] \\ &= \sum_{x_i \in G^*} f_i[x_j/s(x_j + x_{j+1}), x_{j+1}/(1-s)(x_j + x_{j+1})], \end{aligned}$$

we infer that $x_j \sim_{\Phi}^{\mathcal{H}^*} x_{j+1}$.

BDE case Define $\sim_{\Psi}^{\mathcal{H}} := \{(x_i, x_j) : \Psi_{x_i, x_j}^{\mathcal{H}} \text{ is valid}\}$, where $\Psi_{x_i, x_j}^{\mathcal{H}} := \Theta(p) \rightarrow (\bigwedge_{H \in \mathcal{H}} (x_{H,1} = \dots = x_{H,|H|}) \rightarrow f_i = f_j)$ and note that $\sim_{\Psi}^{\mathcal{H}}$ is an equivalence relation on \mathcal{V}_p . Then, we fix BDE partitions $\mathcal{H}_1, \dots, \mathcal{H}_n$ of \mathcal{V}_p and, by applying the very same reasoning as in the case of FDE, we have to show that $x_j \sim_{\Psi}^{\mathcal{H}^*} x_{j+1}$ for all $0 \leq j \leq k-1$. Since $x_j \sim_{i_j} x_{j+1}$ implies

Algorithm 1 Construction of the largest FDE and BDE.

Require: Program p , context c , partition \mathcal{G} of \mathcal{V}_p and $\chi \in \{F, B\}$.
 $\mathcal{H} \leftarrow \mathcal{G}$
while true do
 $\mathcal{H}' \leftarrow \text{refine}_\chi(\mathcal{H})$
 if $\mathcal{H}' = \mathcal{H}$ **then**
 return \mathcal{H}
 else
 $\mathcal{H} \leftarrow \mathcal{H}'$
 end if
end while

Algorithm 2 Routine refine_F

Require: Program p , context c and a partition \mathcal{H} of \mathcal{V}_p .
 $\mathcal{H}' \leftarrow \emptyset$
for all $H \in \mathcal{H}$ **do**
 $\mathcal{R} \leftarrow \{(x_i, x_j) : x_i, x_j \in H \text{ and } (x_i = x_j \text{ or } \Phi_{x_i, x_j}^{\mathcal{H}} \text{ is valid})\}$
 $\mathcal{H}' \leftarrow \mathcal{H}' \cup (H/\mathcal{R})$
end for
return \mathcal{H}'

that $x_j \sim_{\Psi}^{\mathcal{H}_{i,j}} x_{j+1}$ and any block of \mathcal{H}^* is a union of blocks of $\mathcal{H}_{i,j}$, it can be easily seen that $x_j \sim_{\Psi}^{\mathcal{H}_{i,j}} x_{j+1}$ implies $x_j \sim_{\Psi}^{\mathcal{H}^*} x_{j+1}$.

So far, we have shown that the coarsening $\mathcal{V}_p / (\bigcup_{i=1}^m \sim_i)^*$ of FDE (BDE) partitions $\mathcal{H}_1, \dots, \mathcal{H}_n$ is again an FDE (BDE) partition. The claim follows then by noting that Lemma 26 in [20] ensures that $\mathcal{V}_p / (\bigcup_{i=1}^m \sim_i)^*$ is a refinement of \mathcal{G} if each \mathcal{V}_p / \sim_i is a refinement of \mathcal{G} . \square

The main difference with respect to the classical partition-refinement algorithms developed for discrete-state transition systems (e.g., [6, 33, 51]) is that each IDOL variable represents a continuous (uncountable) state space. To tackle this problem we build a variant which performs a symbolic evaluation at each iteration that checks the validity of the FDE/BDE conditions. As usual, the algorithm returns the coarsest FDE/BDE partition that refines a given input partition: This is the trivial partition $\{\mathcal{V}_p\}$ when computing the largest differential equivalence. We remark that the freedom in choosing the initial partition can be useful. For FDE, it allows to single out variables to be preserved in the aggregated program. These are the variables for which the modeler is interested in obtaining distinct ODE solutions. BDE requires equivalent variables to be initialized with same initial conditions. In this case, an appropriate \mathcal{G} can be used to tell apart variables having different initial conditions. This is similar to the pre-partitioning for the largest bisimulation of a labeled Markov chain (e.g., [5]), where states with different sets of atomic propositions are told apart.

The outer loop of the algorithm, shown in Algorithm 1, is a classic fixed-point iteration. The specific refinement depends on an inner procedure, parameterized by the notion of differential equivalence that is considered ($\chi = F$ and $\chi = B$).

FDE Partition Refinement. Routine refine_F , shown in Algorithm 2, refines each block of the *current* partition of IDOL variables according to FDE. As discussed, we use the binary characterization of FDE in Theorem 2. Specifically, for each block $H \in \mathcal{H}$, routine refine_F computes an equivalence relation \mathcal{R} on H relating variables x_i, x_j of H respecting $\Phi_{x_i, x_j}^{\mathcal{H}}$, and adds to \mathcal{H}' blocks of \mathcal{R} -equivalent variables of H . (As discussed, the SMT solver is used when computing \mathcal{R} to check the validity of $\Phi_{x_i, x_j}^{\mathcal{H}}$ for each pair of variables $x_i, x_j \in H$.) The algorithm is correct, since \mathcal{H}' is a

Algorithm 3 Routine refine_B

Require: Program p , context c and a partition \mathcal{H} of \mathcal{V}_p .
if $\Psi^{\mathcal{H}}$ is valid **then**
 $\mathcal{H}' \leftarrow \mathcal{H}$
else
 $\sigma_w \leftarrow \text{getWitness}(\text{sat}(\neg\Psi^{\mathcal{H}}))$
 $\mathcal{H}' \leftarrow \emptyset$
 for all $H \in \mathcal{H}$ **do**
 $\mathcal{R} \leftarrow \{(x_i, x_j) : x_i, x_j \in H \text{ and } \llbracket f_i \rrbracket_c^p(\sigma_w) = \llbracket f_j \rrbracket_c^p(\sigma_w)\}$
 $\mathcal{H}' \leftarrow \mathcal{H}' \cup (H/\mathcal{R})$
 end for
end if
return \mathcal{H}'

refinement of \mathcal{H} , and two variables x_i, x_j for which $\Phi_{x_i, x_j}^{\mathcal{H}}$ is not valid cannot belong to the same block of an FDE partition.

Theorem 5. *If p is an IDOL program, \mathcal{G} a partition of \mathcal{V}_p and $\chi = F$, Algorithm 1 returns the coarsest FDE partition refining \mathcal{G} .*

BDE Partition Refinement. Routine refine_B , shown in Algorithm 3, refines the given current partition for computing BDE. Differently from the FDE case, Definition 6 can be directly used for this. Furthermore we fully exploit the SMT technology: If $\Psi^{\mathcal{H}}$ is not valid then $\neg\Psi^{\mathcal{H}}$ is satisfiable, hence by invoking $\text{sat}(\neg\Psi^{\mathcal{H}})$ the SMT solver provides us with a witnessing assignment σ_w for which $\neg\Psi^{\mathcal{H}}$ holds. We use such witness as a “counter-example” to refine \mathcal{H} : Each block $H \in \mathcal{H}$ is split in sub-blocks of variables $x_{H,i}$ whose drifts $f_{H,i}$ have same value if evaluated according to σ_w . The algorithm is correct because the obtained partition \mathcal{H}' is a refinement of \mathcal{H} , and two variables whose drifts have different values for σ_w cannot belong to the same block of a BDE partition.

Theorem 6. *If p is an IDOL program, \mathcal{G} a partition of \mathcal{V}_p and $\chi = B$, Algorithm 1 returns the coarsest BDE partition refining \mathcal{G} .*

As with other SMT-based partition refinement algorithms [9, 32], Algorithm 1 has a whole can be implemented in a standard way as a routine in a general purpose programming language which calls the SMT solver when required.

4. Applications

In this section we relate IDOL to CTMCs, CRNs, and the FPA process algebra. We show that differential equivalences include the already available notions of equivalence developed in those domains. To do this in a self-contained manner we present the definitions of the semantics as well as of the original equivalences, while we refer to the literature for the intuitions and motivations behind the languages themselves. In all cases, the encoding of the original semantics into IDOL is straightforward, hence we omit this formal step and directly give the underlying IDOL program.

4.1 Continuous-time Markov Chains

Let us consider a CTMC with states $\{1, \dots, n\}$ that is given in terms of its generator matrix $Q = (q_{i,j})_{1 \leq i, j \leq n}$ where $q_{i,j} \in \mathbb{Q}$. That is, for $i \neq j$, the entry $q_{i,j} \geq 0$ defines the rate at which the CTMC moves from state i into state j , whereas we set $q_{i,i} = -\sum_{i \neq j} q_{i,j}$ for all $1 \leq i \leq n$. Then, the corresponding IDOL program is given by the Kolmogorov forward equations.

Definition 9. *The IDOL program p_Q of a CTMC $(q_{i,j})_{1 \leq i, j \leq n}$ is*

$$\dot{x}_i = - \sum_{j \neq i} q_{i,j} \cdot x_i + \sum_{j \neq i} q_{j,i} \cdot x_j, \quad \text{for all } 1 \leq i \leq n.$$

Meaningful contexts for p_Q are such that the initial condition $\hat{\sigma}$ is a probability distribution, i.e., $\sum_{1 \leq i \leq n} \hat{\sigma}(x_i) = 1$, with $\hat{\sigma}(x_i) \geq 0$

for all $1 \leq i \leq n$. For such a context c , $\llbracket x_i \rrbracket_c^p(t)$ gives the probability of being in state i at time t .

We next provide the notions of lumpability for CTMCs [12].

Definition 10 (Ordinary and Exact Lumpability). *Let \mathcal{Z} be a partition of $\{1, \dots, n\}$ and set*

$$q[i, Z] := \sum_{j \in Z} q_{i,j} \quad \text{and} \quad q[Z, i] := \sum_{j \in Z} q_{j,i},$$

where $1 \leq i \leq n$ and $Z \subseteq \{1, \dots, n\}$.

- \mathcal{Z} is called *ordinarily lumpable* if $q[i, Z'] = q[i', Z']$ for all $Z, Z' \in \mathcal{Z}$ and $i, i' \in Z$.
- \mathcal{Z} is called *exactly lumpable* if $q[Z', i] = q[Z', i']$ for all $Z, Z' \in \mathcal{Z}$ and $i, i' \in Z$.

This definition motivates our terminology. Ordinary lumpability is a “forward” criterion because it relates states according to their outgoing transitions (toward equivalence classes); exact lumpability is a “backward” criterion since it relates states according to incoming transitions (from predecessor equivalence classes). On the domain of CTMCs, FDE and BDE turn out to be equivalent to ordinary lumpability and exact lumpability, respectively.

Theorem 7. *Fix a CTMC $Q = (q_{i,j})_{1 \leq i, j \leq n}$ and let \mathcal{Z} be a partition of $\{1, \dots, n\}$. Then, \mathcal{Z} is ordinarily lumpable (resp., exactly lumpable) if and only if the partition $\mathcal{H}_{\mathcal{Z}} = \{\{x_i : i \in Z\} : Z \in \mathcal{Z}\}$ of \mathcal{V}_{p_Q} is an FDE (resp., BDE) of p_Q .*

Proof. FDE case. As observed already by Proposition 1 in [78], for all $Z \in \mathcal{Z}$ and $i \in Z$ it holds that

$$q[i, Z] = q(i, i) + \sum_{\substack{j \in Z \\ j \neq i}} q(i, j) = - \sum_{\substack{Z' \in \mathcal{Z} \\ Z' \neq Z}} q[i, Z'].$$

That is, \mathcal{Z} is ordinarily lumpable if and only if $q[i, Z'] = q[i', Z']$ for all $Z, Z' \in \mathcal{Z}$ and $i, i' \in Z$ where $Z' \neq Z$. In the following, we will use this alternative formulation of ordinary lumpability to establish the equivalence with FDE. Since [12] and the proof of Theorem 2 ensure that any ordinarily lumpable partition \mathcal{Z} induces an FDE $\mathcal{H}_{\mathcal{Z}}$, let us assume that $\mathcal{H}_{\mathcal{Z}}$ is an FDE and pick some arbitrary $Z, Z' \in \mathcal{Z}$ and $i, i' \in Z$ with $Z' \neq Z$. We have to show that $q[i, Z'] = q[i', Z']$. Thanks to Theorem 2, we know that $\Phi_{x_i, x_{i'}}^{\mathcal{H}_{\mathcal{Z}}}$ holds true. This, however, implies that the value of

$$\sum_{j \in Z'} q_{i,j} \cdot s \cdot (x_i + x_{i'}) + \sum_{j \in Z'} q_{i',j} \cdot (1-s) \cdot (x_i + x_{i'})$$

does not depend on the assignment $0 < \sigma(s) \leq 1$, meaning that $q[i, Z'] = q[i', Z']$.

BDE case. Assume without loss of generality that $\mathcal{Z} = \{\bar{1}, \dots, \bar{M}\}$ with $M = |\mathcal{Z}|$ and $\bar{I} = \{(\bar{I}, 1), \dots, (\bar{I}, |\bar{I}|)\}$ for any $\bar{I} \in \mathcal{Z}$. We note that

$$\dot{x}_k = - \sum_{j \neq k} q_{k,j} \cdot x_k + \sum_{j \neq k} q_{j,k} \cdot x_j = \sum_j q_{j,k} \cdot x_j,$$

which yields, for any $\bar{I} \in \mathcal{Z}$ and $1 \leq l \leq |\bar{I}|$,

$$\begin{aligned} \dot{x}_{\bar{I},l} &= \sum_{\bar{J} \in \mathcal{Z}} \sum_{k=1}^{|\bar{J}|} q_{(\bar{J},k),(\bar{I},l)} \cdot x_{\bar{J},k} \\ &= \sum_{\bar{J} \in \mathcal{Z}} \left(\sum_{k=1}^{|\bar{J}|} q_{(\bar{J},k),(\bar{I},l)} \right) \cdot x_{\bar{J},1} =: \wp_{\bar{I},l} \end{aligned}$$

in the case of $\bigwedge_{\bar{J} \in \mathcal{Z}} (x_{\bar{J},1} = \dots = x_{\bar{J},|\bar{J}|})$. Since $\mathcal{H}_{\mathcal{Z}}$ is a BDE and the real polynomials $\wp_{\bar{I},l}$ and $\wp_{\bar{I},l'}$, where $\bar{I} \in \mathcal{Z}$

and $1 \leq l, l' \leq |\bar{I}|$, coincide if and only if they have the same coefficients, we infer that

$$\sum_{k=1}^{|\bar{J}|} q_{(\bar{J},k),(\bar{I},l)} = \sum_{k=1}^{|\bar{J}|} q_{(\bar{J},k),(\bar{I},l')}$$

for all $\bar{J} \in \mathcal{Z}$, thus closing the proof. \square

This result explains why differential equivalences can be seen as a somewhat natural generalization of more traditional notions of equivalence for discrete-state stochastic systems. In principle, the coarsest lumpable partition of a CTMC could be computed using the partition refinement algorithm in Section 3.2. However, in practice, one would use the efficient algorithms specialized for CTMCs, which run in polynomial time and space [33]. Still, an SMT-based approach to computing CTMC lumpability can be useful to handle uncertainty in rate values, by treating them symbolically (e.g., as suggested in [32] for PRISM [55]).

As a side product, we remark that Theorem 7 provides a characterization of ordinary and exact lumpability by means of real calculus, based on the Kolmogorov ODEs, instead of the classical argument [12] that combines the well-known concept of *uniformization* (e.g., [72]) with the characterization of lumpability for discrete time Markov chains [52].

4.2 Chemical Reaction Networks

A CRN is a set of rules (reactions) describing interactions between *species*. For instance, the reaction $A + B \xrightarrow{\alpha} 2C$ states that one element (e.g., molecule) of species A interacts with one element of species B to form two elements of species C . The label α decorates the reaction with information about the speed at which the reaction occurs; its signature depends on the chosen kinetics.

Formally, let S be a finite set of species. Either side of a reaction is a multiset of S , i.e., a function in \mathbb{N}_0^S associating each species with its multiplicity (the *stoichiometry*) as a reactant or product. The stoichiometry of a species A in multiset ρ is denoted by ρ_A . A reaction r over S is a triple $(\rho, \pi, \alpha) \in R_S \subseteq \mathbb{N}_0^S \times \mathbb{N}_0^S \times \mathcal{L}$, where \mathcal{L} is the label set, represented usually with $\rho \xrightarrow{\alpha} \pi$.

Mass-action CRNs. We now give the semantics of CRNs according to standard mass action kinetics. In this case the labels are *rates*, i.e., positive real numbers; the speed of the reaction is proportional by such rates to the product of the amounts of the reactant species.

Definition 11. *A mass-action CRN is a pair (S, R_S) where R_S is a finite set of reactions over S , with $R_S \subseteq \mathbb{N}_0^S \times \mathbb{N}_0^S \times \mathbb{Q}_{>0}$.*

For mass-action reactions, set $\phi(A, \rho \xrightarrow{\alpha} \pi) := \alpha(\pi_A - \rho_A)$.

Definition 12. *The IDOL program p_S of a mass-action CRN (S, R) is*

$$\dot{x}_A = f_A := \sum_{\rho \xrightarrow{\alpha} \pi \in R_S} \phi(A, \rho \xrightarrow{\alpha} \pi) \prod_{B \in S} x_B^{\rho_B}, \quad \text{for all } A \in S.$$

Hill CRNs. We discuss the semantics of CRNs according to the Hill kinetics (e.g., [80]) in the case of *catalytic* reactions, i.e., reactions which are in the form $B + C \xrightarrow{l} D + C$ with $B \neq D$. Here, C plays the role of a catalyst, a species promoting the reaction but which is not affected by it. Species B is the *substrate* that is modified, becoming D , when the reaction occurs. Each reaction is labeled with a triple $(\beta_1, \beta_2, \nu) \in \mathbb{Q}_{>0}^3$.

Definition 13. *A Hill CRN is a pair (S, R_S) where R_S is a finite set of catalytic reactions with $R_S \subseteq \mathbb{N}_0^S \times \mathbb{N}_0^S \times \mathbb{Q}_{>0}^3$.*

Definition 14. The IDOL program p_S of a Hill CRN is

$$\dot{x}_A = h_A := \sum_{\substack{(\beta_1, \beta_2, \nu) \\ \rho = B+C, \pi = D+C}} (\pi_A - \rho_A) \frac{\beta_1 x_B^\nu}{\beta_2 + x_B^\nu}, \quad \text{for all } A \in S.$$

In both semantics, reasonable contexts for CRNs are such that $\hat{\sigma}(x_A) \geq 0$ for all $A \in S$, since the IDOL variables represent concentrations of species, i.e., molecular counts divided by the volume of the environment where the reactions take place.

4.2.1 CRN Emulation

Emulation is a recently developed notion of comparison between mass-action CRNs [18]. The definition is presented below, slightly simplified from [18] and directly stated in IDOL terms.

Definition 15. Let (S, R_S) and $(\tilde{S}, \tilde{R}_{\tilde{S}})$ denote two mass-action CRNs, with vector fields denoted by $\llbracket f \rrbracket_c^{p_S}$ and $\llbracket \tilde{f} \rrbracket_{\tilde{c}}^{p_{\tilde{S}}}$ and contexts denoted by c and \tilde{c} , respectively. A species morphism from (S, R_S) to $(\tilde{S}, \tilde{R}_{\tilde{S}})$ is a function $\mu_S : S \rightarrow \tilde{S}$. It is an emulation when $\llbracket \tilde{f} \rrbracket_{\tilde{c}}^{p_{\tilde{S}}}(\tilde{\sigma} \circ \mu_S) = (\llbracket f \rrbracket_c^{p_S}(\tilde{\sigma})) \circ \mu_S$ for all $\tilde{\sigma} \in \mathbb{R}^{\tilde{S}}$.

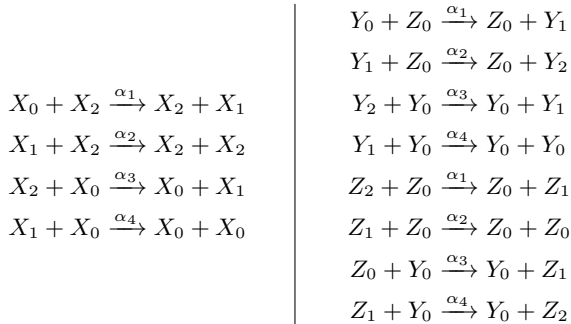
The emulation condition, stated in terms of function composition, can be checked syntactically on the CRN structure by using the notions of reactant morphism and stoichiomorphism presented in [18]. Here we formally relate emulation to BDE.

Proposition 2. If μ_S is an emulation from (S, R_S) to $(\tilde{S}, \tilde{R}_{\tilde{S}})$ then:

- i) $\{\mu_S^{-1}(\tilde{A}) : \tilde{A} \in \tilde{S}\}$ is a BDE partition of (S, R_S) .
- ii) Assume $S \cap \tilde{S} = \emptyset$. Then, $\{\mu_S^{-1}(\tilde{A}) \cup \{\tilde{A}\} : \tilde{A} \in \tilde{S}\}$ is a BDE partition of the CRN $(S \cup \tilde{S}, R_S \cup \tilde{R}_{\tilde{S}})$.

We observe that with i) BDE allows to relate species within the same CRN. By ii), we note that emulation relates species essentially like BDE: whenever all species are initialized with the same conditions as the target species to which they are mapped, then such species have the same ODE traces. We note that the assumption on disjoint sets of species in ii) is without loss of generality since it is always possible to rename species of one CRN with fresh names.

Example 1. The following two mass-action CRNs describe the behavior of AM, a basic biological switch (left) and MI, a mutual inhibition mechanism (right) [18]:



The following species morphism can be shown to be an emulation:

$$\begin{aligned} \mu_S(Y_0) = X_0, \quad \mu_S(Y_1) = X_1, \quad \mu_S(Y_2) = X_2, \\ \mu_S(Z_0) = X_2, \quad \mu_S(Z_1) = x_1, \quad \mu_S(Z_2) = X_0 \end{aligned}$$

Since emulation is a particular BDE partition, with BDE it is possible to automatically check whether such correspondence of traces carries over to non-mass-action kinetics. The possibility of reasoning using different hypotheses for the reaction kinetics is of biological relevance because in different situations one may find mass-action mechanisms (e.g., phosphotransfers) or Hill-type

mechanisms (e.g., enzymes). For instance, much of the utility of Hill kinetics is owed to supporting non-integer exponents. Famously, this ranges in 2.3-3.0 for haemoglobin. Furthermore, biologists often consider exponents less than 1 in order to describe ‘‘anticooperative’’ behavior [61]. Any rational exponent can be expressed in IDOL.

In Example 1 it is possible to show that replacing equal mass-action rates with equal (and arbitrary) Hill triples, a BDE partition that is related to an emulation in the sense of item ii) of Proposition 2 is still BDE for the resulting Hill CRN. This suggests a structural relationship between CRNs with different biological functionality, which is insensitive (to some extent) to underlying kinetics that is considered. Indeed, through BDE it is possible to show that *all* of the thirteen species morphisms found in [18] do enjoy this property. This is particularly interesting because, at the same time, Hill kinetics can never be exactly matched by mass-action kinetics, and vice versa. (This holds because the drift for Hill kinetics has partial derivatives of arbitrary high order that are not identical to zero, whereas the drift for mass-action kinetics does not.)

4.2.2 Forward and Backward Bisimulations for CRNs

Forward and backward bisimulations for CRNs have been recently introduced in [20]. These are equivalence relations over species for elementary mass-action reactions, where at most two reactants (possibly of the same species) can appear in the left hand side. A specialized partition-refinement algorithm is provided which computes the largest bisimulations in polynomial time and space.

Forward bisimulation. Forward bisimulation depends on the computation of the following quantities from the CRN syntax.

Definition 16. Let (S, R_S) be a CRN, $A, A' \in S$, and $\rho \in S \cup \{\emptyset\}$. The ρ -reaction rate of A , and the ρ -production rate of A -elements by A are defined respectively as

$$\begin{aligned} \mathbf{crr}[A, \rho] &:= (\rho_A + 1) \sum_{A+\rho \xrightarrow{\alpha} \pi \in R_S} \alpha, \\ \mathbf{pr}(A, \rho, A') &:= (\rho_A + 1) \sum_{A+\rho \xrightarrow{\alpha} \pi \in R_S} \alpha \cdot \pi_{A'} \end{aligned}$$

Finally, for $Z \subseteq S$ we define $\mathbf{pr}[A, \rho, Z] := \sum_{A' \in Z} \mathbf{pr}(A, \rho, A')$.

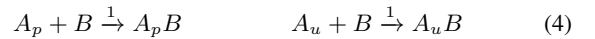
Definition 17. Let (S, R_S) be a CRN, \mathcal{R} an equivalence relation over S and $\mathcal{Z} = S/\mathcal{R}$. Then, \mathcal{R} is a forward bisimulation if for all $(A, A') \in \mathcal{R}$, all ρ , and all blocks $Z \in \mathcal{Z}$ it holds that

$$\mathbf{crr}[A, \rho] = \mathbf{crr}[A', \rho] \text{ and } \mathbf{pr}[A, \rho, Z] = \mathbf{pr}[A', \rho, Z]. \quad (3)$$

Proposition 3 (Forward bisimulation implies FDE). *Let (S, R_S) be an elementary mass-action CRN and \mathcal{R} be a forward bisimulation. Then, the IDOL program p underlying (S, R_S) is such that $\mathcal{V}_p/\{(x_A, x_{A'}) : (A, A') \in \mathcal{R}\}$ is an FDE partition.*

We offer an intuitive explanation as to why forward bisimulation is only a sufficient condition for FDE. For this, we observe that it *separately* considers the negative and the positive contribution by reactions to the concentration of a species, captured by the ρ -reaction rate and the ρ -production rate, respectively. Instead, FDE considers the net contribution to the species’ concentration. We illustrate this difference using an example.

Example 2. Consider the mass-action CRN



This is a basic instance of CRNs resulting from higher-level approaches using rule-based languages such as κ [26] and BioNetGen [11]. It involves *agents*, A and B , which may exhibit different

Table 1. FDE reduces more than forward bisimulation (FB).

Original model	Largest FB		Largest FDE				
	Model	R	S	Red.(s)	Size	Red.(s)	Size
M1 [34, 70]	8620	745	6.54E-1	745	7.85E+3	105	
M2 [34, 70]	3680	354	2.81E-1	354	3.22E+3	105	
M3 [1]	4944	411	1.29E-1	411	6.46E+2	47	
M4 [8]	3447	348	2.46E-1	348	5.22E+3	215	

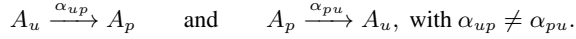
internal states (e.g., being phosphorylated or unphosphorylated, subscripts p and u , respectively). Line 4 shows *complexation* reactions, where agents can combine into a single compound, modelled as a different species (A_pB and A_uB). Line 5 are *decomplexation* reactions, where the agents unbind. By Definition 12 we have that the IDOL program of this CRN is

$$\begin{aligned} \dot{x}_{A_p} &= -x_{A_p} \cdot x_B + 2x_{A_pB} & \dot{x}_{A_pB} &= x_{A_p} \cdot x_B - 2x_{A_pB} \\ \dot{x}_{A_u} &= -x_{A_u} \cdot x_B + 2x_{A_uB} & \dot{x}_{A_uB} &= x_{A_u} \cdot x_B - 2x_{A_uB} \\ \dot{x}_B &= -x_{A_p} \cdot x_B + 2x_{A_pB} - x_{A_u} \cdot x_B + 2x_{A_uB} \end{aligned}$$

It holds that $\mathcal{Z}_C = \{\{A_p, A_u\}, \{B\}, \{A_pB, A_uB\}\}$ is a forward bisimulation, so $\mathcal{H}_C = \{\{x_{A_p}, x_{A_u}\}, \{x_B\}, \{x_{A_pB}, x_{A_uB}\}\}$ is an FDE for the corresponding IDOL program. For example, the aggregated drift of the block $\{x_{A_p}, x_{A_u}\}$ is

$$f_{A_p} + f_{A_u} = -(x_{A_p} + x_{A_u}) \cdot x_B + 2(x_{A_pB} + x_{A_uB}) \quad (6)$$

Let us now add reactions



Biologically, this is a classical model of phosphorylation/unphosphorylation of a molecule. With these new reactions, \mathcal{Z}_C is not a forward bisimulation, because $\mathbf{crr}[A_u, \emptyset] = \alpha_{up}$ while $\mathbf{crr}[A_p, \emptyset] = \alpha_{pu}$. However, \mathcal{H}_C is still an FDE because the influence of the new reactions “disappears” in the aggregated ODEs:

$$\begin{aligned} f_{A_p} + f_{A_u} &= -(x_{A_p} + x_{A_u}) \cdot x_B + 2(x_{A_pB} + x_{A_uB}) \\ &\quad + \alpha_{up} \cdot x_{A_u} - \alpha_{pu} \cdot x_{A_p} - \alpha_{up} \cdot x_{A_u} + \alpha_{pu} \cdot x_{A_p} \end{aligned}$$

where the last line, which represents the contributions of the new reactions, evaluates to zero. Hence the new aggregated drift $f_{A_p} + f_{A_u}$ coincides with (6).

We now show that FDE can aggregate biochemical models from the literature that cannot be aggregated using forward bisimulation. For this experimental study, we developed a Java prototype that implements Algorithm 1 for IDOL programs of mass-action CRNs, using the Z3 application programming interface to evaluate the conditions on differential equivalence. Our tool is available at <http://sysma.imtlucca.it/erode> to replicate our tests.

The results are presented in Table 1. Alongside the model identifier we show the reference from which the CRN was taken; headers $|R|$ and $|S|$ give the number of reactions and species, respectively, of the original CRN. Headers $Red.(s)$ give the time in seconds to compute the largest equivalences. Measurements were taken on a 2.6 GHz Intel Core i5 with 4 GB of RAM. We obtain up to one order of magnitude fewer variables (models M1 and M3) using FDE.

The runtime comparisons show that FDE is computationally more demanding than forward bisimulation, exhibiting larger runtimes even if it performs fewer iterations on each model because it gives coarser aggregations. The reason is that the computation of \mathcal{R} at each iteration of Algorithm 2 requires in the worst case to establish the validity of $\Phi_{x_i, x_j}^{\mathcal{H}}$ for each pair of IDOL variables x_i, x_j belonging to the same block. Furthermore, each check is performed

Table 2. BDE has runtimes similar to backward bisimulation (BB).

Original model	Reduction					
	Model	R	S	BB (s)	BDE (s)	S
M5 [70]	786432	65538	3.68E+3	1.01E+3	167	
M6 [70]	172032	16386	1.77E+2	3.01E+2	122	
M7 [70]	48	18	2.00E-3	6.00E-2	12	
M8 [73]	194054	14531	1.32E+3	3.45E+3	6634	
M9 [34, 70]	187468	10734	2.71E+2	1.57E+3	5575	
M10 [22, 23]	5832	730	6.00E-1	3.22E+0	217	
M11 [53]	487	85	6.00E-3	2.71E-1	56	
M12 [18]	24	18	7.00E-3	5.20E-2	3	

symbolically using the SMT solver, while the partition refinement algorithm of [20] splits candidate partitions using Definition 16. However we stress that these tests are “unfair” to our differential equivalences because the comparison is with a specialized partition refinement algorithm which iterates using concrete values (that can be computed syntactically).

Backward bisimulation. On the domain of elementary CRNs with mass-action kinetics, BDE corresponds to backward bisimulation; hence we avoid recalling that notion (see [20, Def. 14]).

Proposition 4 (Backward bisimulation corresponds to BDE). *Let (S, R_S) be an elementary mass-action CRN and \mathcal{R} be an equivalence relation over S . Then, the IDOL program p underlying (S, R_S) is such that $\mathcal{V}_p / \{(x_A, x_{A'}) : (A, A') \in \mathcal{R}\}$ is a BDE partition if and only if \mathcal{R} is a backward bisimulation.*

Yet, we still compare BDE and backward bisimulation to show, surprisingly to us, that our prototypical SMT-based implementation of BDE can already compete with the polynomial mass-action specific partition refinement algorithm for backward bisimulation. The results are presented in Table 2, using models also studied in [20]. For instance, the reduction of M5 is obtained three times faster using BDE, while M6 and M8 registered runtimes of the same order. This shows that the SMT theory and technology are well suited for BDE: a single iteration allows to perform as many partition splits as the number of distinct evaluations of the drifts with the counterexample assignment. In the largest benchmark, M5, at each iteration of the partition refinement algorithm the SMT solver evaluated equivalences involving ca. 786,000 nonlinear monomials and 1,500,000 linear monomials, from binary and unary reactions, respectively.

Finally, we note that the formula $\Psi^{\mathcal{H}}$ used at each iteration to refine the current partition \mathcal{H} is much simpler than the ones considered for FDE, as just single drifts are compared rather than cumulative ones over blocks. For instance, it can be shown that in model M11 the coarsest FDE and BDE partitions coincide, but it took 9.90E+1 s to compute the largest FDE, as opposed to the reported 2.71E-1 s for the largest BDE.

4.3 Process Algebra

Lastly we consider a fragment of Fluid Process Algebra (FPA) presented in [49], which corresponds to the process algebra studied in [77]. The grammar of FPA considers parallel composition of sequential processes with synchronization over shared actions. Let \mathcal{A} denote the set of actions and \mathcal{K} the set of constants. Each process $P \in \mathcal{K}$ is defined as $P \stackrel{\text{def}}{=} \sum_{i \in I_P} (\alpha_i, r_i) \cdot P_i$, where I_P is an index set, $\alpha_i \in \mathcal{A}$, $r_i \in \mathbb{Q}_{>0}$ is a rate, and $P_i \in \mathcal{K}$.

Using the obvious standard operational semantics for the choice and prefix operator, we let $\mathcal{B}(P)$ be the states of the underlying

LTS, with transitions denoted by $P \xrightarrow{(\alpha_i, r_i)} P_i$. Furthermore, we let $\mathcal{A}(P)$ denote the set of actions labeling transitions from P .

Definition 18. An FPA model \mathcal{M} is generated by

$$\mathcal{M} ::= P \mid \mathcal{M} \parallel_L \mathcal{M}, \quad \text{with } L \subseteq \mathcal{A} \text{ and } P \in \mathcal{K}.$$

We let $\mathcal{G}(\mathcal{M})$ be the set of sequential components appearing in \mathcal{M} and $\mathcal{B}(\mathcal{M})$ for $\bigcup_{P \in \mathcal{G}(\mathcal{M})} \mathcal{B}(P)$. For any two $P, Q \in \mathcal{G}(\mathcal{M})$, we assume $\mathcal{B}(P) \cap \mathcal{B}(Q) = \emptyset$.

We introduce the following elementary concepts that will be needed to define the semantics.

Definition 19. Let $Z \subseteq \mathcal{K}$ and $\alpha \in \mathcal{A}$. Then

$$r_\alpha(P) := \sum_{P \xrightarrow{(\alpha, r)} P'} r \quad \text{and} \quad q[P, Z, \alpha] := \sum_{P' \in Z} \sum_{P \xrightarrow{(\alpha, r)} P'} r$$

Also, we say that an action α is enabled in an FPA model \mathcal{M} if for any submodel $\mathcal{M}_1 \parallel_L \mathcal{M}_2$ of \mathcal{M} with $\alpha \in L$ there exist $P_1 \in \mathcal{B}(\mathcal{M}_1)$ and $P_2 \in \mathcal{B}(\mathcal{M}_2)$ with $r_\alpha(P_1) > 0$ and $r_\alpha(P_2) > 0$.

An IDOL variable (hence, an ODE) is associated with each LTS state of every sequential process appearing in an FPA model \mathcal{M} .

Definition 20. The IDOL program of an FPA model \mathcal{M} is given by

$$\begin{aligned} \dot{x}_P &= \sum_{\alpha \in \mathcal{A}} \sum_{P' \in \mathcal{B}(\mathcal{M})} x_{P'} \cdot q[P', \{P\}, \alpha] \cdot \mathcal{R}_\alpha^*(\mathcal{M}, P') \\ &\quad - \sum_{\alpha \in \mathcal{A}} x_P \cdot r_\alpha(P) \cdot \mathcal{R}_\alpha^*(\mathcal{M}, P), \quad \text{for all } P \in \mathcal{B}(\mathcal{M}), \end{aligned}$$

where $r_\alpha^*(\mathcal{M})$ is recursively defined as

$$\begin{aligned} r_\alpha^*(\mathcal{M}_1 \parallel_L \mathcal{M}_2) &:= \begin{cases} r_\alpha^*(\mathcal{M}_1) + r_\alpha^*(\mathcal{M}_2) & , \alpha \notin L \\ \min(r_\alpha^*(\mathcal{M}_1), r_\alpha^*(\mathcal{M}_2)) & , \alpha \in L \end{cases} \\ r_\alpha^*(P) &:= \sum_{P' \in \mathcal{B}(P)} x_{P'} \cdot r_\alpha(P') \end{aligned}$$

$$\begin{aligned} \mathcal{R}_\alpha^*(\mathcal{M}_1 \parallel_L \mathcal{M}_2, P) &:= \begin{cases} \mathcal{R}_\alpha^*(\mathcal{M}_i, P) & , P \in \mathcal{B}(\mathcal{M}_i) \wedge \alpha \notin L \\ \mathcal{R}_\alpha^*(\mathcal{M}_i, P) \frac{r_\alpha^*(\mathcal{M}_1 \parallel_L \mathcal{M}_2)}{r_\alpha^*(\mathcal{M}_i)} & , P \in \mathcal{B}(\mathcal{M}_i) \wedge \\ & \alpha \in L \text{ is enabled in } \mathcal{M}_i \\ 0 & , \text{otherwise} \end{cases} \\ \mathcal{R}_\alpha^*(P, P') &:= \begin{cases} 1 & , P' \in \mathcal{B}(P) \\ 0 & , \text{otherwise} \end{cases} \end{aligned}$$

We refer the reader to the literature (e.g., [49, 77]) for a more detailed discussion on the semantics of FPA. Here we stress that the crucial definition is $r_\alpha^*(\mathcal{M}_1 \parallel_L \mathcal{M}_2)$ for a synchronization action $\alpha \in L$. Intuitively, it provides a contribution to the drift that establishes a threshold-based contention between the capacities (i.e., rates) of the operands, dictated by the minimum function.

4.3.1 Differential Bisimulation

The first equivalence for FPA that we study is differential bisimulation [49]. It is a relation over the set of constants of an FPA model, defined in terms of conditions on the sequential behavior and on the compositional structure of processes. The latter is captured by collecting actions which affect the sequential behavior.

Definition 21. Let \mathcal{M} be an FPA model, and $P \in \mathcal{B}(\mathcal{M})$. Then

$$\mathcal{D}(P, \mathcal{M}) := \begin{cases} L \cup \mathcal{D}(P, \mathcal{M}_i) & , \mathcal{M} = \mathcal{M}_1 \parallel_L \mathcal{M}_2, P \in \mathcal{B}(\mathcal{M}_i) \\ \emptyset & , \text{otherwise} \end{cases}$$

and

$$\mathcal{I}(P, \mathcal{M}) := \mathcal{D}(P, \mathcal{M}) \cap \mathcal{A}(P).$$

For any $Z \subseteq \mathcal{B}(\mathcal{M})$, we set $\mathcal{D}(Z, \mathcal{M}) = \bigcup_{P \in Z} \mathcal{D}(P, \mathcal{M})$ and $\mathcal{I}(Z, \mathcal{M}) = \bigcup_{P \in Z} \mathcal{I}(P, \mathcal{M})$.

Definition 22. Let \mathcal{M} be an FPA model, and $P, Q \in \mathcal{B}(\mathcal{M})$. Then we write $P \stackrel{s.i.}{=}_{\mathcal{M}} Q$ if

- (i) $\mathcal{A}(P) = \mathcal{A}(Q)$, and
- (ii) if there exists an $\overline{\mathcal{M}} = \mathcal{M}_1 \parallel_L \mathcal{M}_2$ within \mathcal{M} with $P \in \mathcal{B}(\mathcal{M}_1)$, and $Q \in \mathcal{B}(\mathcal{M}_2)$ (or vice versa), then $\mathcal{I}(P, \overline{\mathcal{M}}) = \mathcal{I}(Q, \overline{\mathcal{M}}) = \emptyset$.

Definition 23 (Differential bisimulation). Let \mathcal{M} be an FPA model, \mathcal{R} an equivalence relation over $\mathcal{B}(\mathcal{M})$, and $\mathcal{Z} = \mathcal{B}(\mathcal{M})/\mathcal{R}$. We say that \mathcal{R} is a differential bisimulation for \mathcal{M} if for all $(P, P') \in \mathcal{R}$ we have:

- (i) $q[P, Z, \alpha] = q[P', Z, \alpha]$, for all $Z \in \mathcal{Z}$ and $\alpha \in \mathcal{A}$,
- (ii) $P \stackrel{s.i.}{=}_{\mathcal{M}} P'$.

We observe that differential bisimulation can be checked in a non-symbolic fashion. This is because the aggregate rates $q[P, B, \alpha]$ only depend on the rates that label the prefix operator. The equivalence relation $\stackrel{s.i.}{=}_{\mathcal{M}}$ involves parsing the FPA model and appropriately collecting the action types that decorate the compositional operator of FPA. It turns out that differential bisimulation is a sufficient condition for FDE on the corresponding IDOL program.

Proposition 5. Let \mathcal{M} be an FPA model and \mathcal{R} a differential bisimulation. Then, the IDOL program p underlying \mathcal{M} is such that $\mathcal{V}_p / \{(x_P, x_Q) : (P, Q) \in \mathcal{R}\}$ is an FDE partition.

We discuss why it is not a necessary condition using the example taken from [49].

Example 3. Let $\mathcal{M}_F := P_1 \parallel_{\{\alpha\}} Q_1$, with P_1, Q_1 defined as

$$\begin{aligned} P_1 &\stackrel{\text{def}}{=} (\beta, r).P_2 + (\beta, r).P_3 & P_2 &\stackrel{\text{def}}{=} (\alpha, s).P_1 & P_3 &\stackrel{\text{def}}{=} (\alpha, s).P_1 \\ Q_1 &\stackrel{\text{def}}{=} (\gamma, 2r).Q_2 & Q_2 &\stackrel{\text{def}}{=} (\alpha, s).Q_1 \end{aligned}$$

Applying Definition 20, its IDOL program is

$$\begin{aligned} \dot{x}_{P_1} &= s \min(x_{P_2} + x_{P_3}, x_{Q_2}) - 2r x_{P_1} \\ \dot{x}_{Q_1} &= s \min(x_{P_2} + x_{P_3}, x_{Q_2}) - 2r x_{Q_1} \\ \dot{x}_{P_2} &= r x_{P_1} - s x_{P_2} \frac{\min(x_{P_2} + x_{P_3}, x_{Q_2})}{x_{P_2} + x_{P_3}} \\ \dot{x}_{Q_2} &= 2r x_{P_1} - s \min(x_{P_2} + x_{P_3}, x_{Q_2}) \\ \dot{x}_{P_3} &= r x_{P_1} - s x_{P_3} \frac{\min(x_{P_2} + x_{P_3}, x_{Q_2})}{x_{P_2} + x_{P_3}} \end{aligned} \quad (7)$$

It can be shown that $\mathcal{Z}_F = \{\{P_1\}, \{P_2, P_3\}, \{Q_1\}, \{Q_2\}\}$ is a differential bisimulation, hence the corresponding partition on the IDOL variables $\mathcal{H}_F = \{\{x_{P_1}\}, \{x_{P_2}, x_{P_3}\}, \{x_{Q_1}\}, \{x_{Q_2}\}\}$ is an FDE. Indeed, summing the variables within each equivalence class, we obtain:

$$\begin{aligned} \dot{x}_{P_1} &= s \min(x_{P_2} + x_{P_3}, x_{Q_2}) - 2r x_{P_1} \\ \dot{x}_{Q_1} &= s \min(x_{P_2} + x_{P_3}, x_{Q_2}) - 2r x_{Q_1} \\ \dot{x}_{P_2} + \dot{x}_{P_3} &= 2r x_{P_1} - s \min(x_{P_2} + x_{P_3}, x_{Q_2}) \\ \dot{x}_{Q_2} &= 2r x_{P_1} - s \min(x_{P_2} + x_{P_3}, x_{Q_2}) \end{aligned}$$

The converse, i.e., that an FDE over the IDOL variables implies a differential bisimulation for the corresponding processes, does not hold in general. For instance, by changing the definition of P_2 of Example 3 with $P_2 \stackrel{\text{def}}{=} (\delta, s).P_1$ we have that \mathcal{R}_F is not a differential bisimulation because we obtain $q[P_2, \{P_1\}, \alpha] = 0$

and $q[P_3, \{P_1\}, \alpha] = s$. Instead, \mathcal{H}_F remains an FDE, since the “domain-specific” information about action types is lost in the IDOL program.

As a further example, let us replace P_2 with the following definition: $P_2 \stackrel{\text{def}}{=} (\alpha, s).P_1 + (\delta, s).P_3$. Again, we have that \mathcal{R}_F is not a differential bisimulation, while \mathcal{H}_F is an FDE. This is because the added δ -transition from P_2 to P_3 distinguishes P_2 and P_3 (i.e., $q[P_2, \{P_2, P_3\}, \delta] = s$ and $q[P_3, \{P_2, P_3\}, \delta] = 0$), but its influence disappears in the lumped ODEs: the negative drift term $-\delta \cdot x_2$ in \dot{x}_{P_2} cancels out the positive drift term $\delta \cdot x_2$ in \dot{x}_{P_3} . This is similar to Example 2 for CRNs, and is actually an instance of a more general observation that transitions internal to an equivalence class do not interfere at the FDE level but may tell apart processes according to differential bisimulation. Indeed, it is not difficult to see that for any FPA model \mathcal{M} , the trivial partition corresponding to $\{\mathcal{B}(\mathcal{M})\}$ is always an FDE. This is, intuitively, a conservation-of-mass property due to the fact that processes are not created nor destroyed in FPA, hence all transitions are internal to the trivial partition. This remark also stresses the usefulness in having an algorithm that can refine any given initial partition, since computing the largest FDE for an FPA model always collapses to an uninteresting reduction.

4.3.2 Label Equivalence

Similarly to differential bisimulation, we next provide the notion of label equivalence that has been introduced in [77] and that describes a sufficient, but not necessary, condition for a partition of FPA constants to be a BDE.

Definition 24 (Label Equivalence). *Let \mathcal{M} be a FPA model and let $\mathcal{P} = (\mathbf{P}^1, \dots, \mathbf{P}^N)$, $\mathbf{P}^i = (P_1^i, \dots, P_{K_i}^i)$, be a tuple partition on $\mathcal{G}(\mathcal{M}) = \{Q_1, \dots, Q_n\}$, that is, for each $P \in \mathcal{G}(\mathcal{M})$ there exist unique $1 \leq i \leq N$ and $1 \leq k \leq K_i$ with $P = P_k^i$. \mathbf{P}^i and \mathbf{P}^j are said to be label equivalent, written $\mathbf{P}^i \sim_{\mathcal{P}} \mathbf{P}^j$, if $K_i = K_j$ and there exist bijections $\rho_k^{i,j} : \mathcal{B}(P_k^i) \rightarrow \mathcal{B}(P_k^j)$, where $1 \leq k \leq K_i$, such that for all $\alpha \in \mathcal{A}$ it holds that $r_\alpha(P_k^i) = r_\alpha(P_k^j)$ and*

- $\forall x_{Q_1} \dots \forall x_{Q_n} \left(\mathcal{R}_\alpha^*(\mathcal{M}, P) = \mathcal{R}_\alpha^*(\mathcal{M}, \rho_k(P))[\dots] \right)$,
- $\forall x_{Q_1} \dots \forall x_{Q_n} \left(\sum_{P' \in \mathcal{G}(\mathcal{M})} q[P', \{P\}, \alpha] \mathcal{R}_\alpha^*(\mathcal{M}, P'), \right.$
 $\left. = \sum_{P' \in \mathcal{G}(\mathcal{M})} q[P', \{\rho_k(P)\}, \alpha] \mathcal{R}_\alpha^*(\mathcal{M}, P')[\dots] \right)$,
- $\forall x_{Q_1} \dots \forall x_{Q_n} \left(\mathcal{R}_\alpha^*(\mathcal{M}, P) = \mathcal{R}_\alpha^*(\mathcal{M}, P)[\dots] \right)$ for all P in $ds(P_k^i)$ with $P_k^i \notin \mathbf{P}^i, \mathbf{P}^j$ and
- $\forall x_{Q_1} \dots \forall x_{Q_n} \left(r_\alpha^*(\mathcal{M}) = r_\alpha^*(\mathcal{M})[\dots] \right)$,

where $[\dots]$ abbreviates

$$[x_Q/x_{\rho_k(Q)}, x_R/x_{\rho_k^{-1}(R)} : Q \in \mathcal{B}(P_k^i), R \in \mathcal{B}(P_k^j), 1 \leq k \leq K].$$

It can be proven that label equivalence is an equivalence relation on the tuple partition \mathcal{P} . More importantly, the following result from [77] connects label equivalence to the notion of BDE.

Proposition 6. *Fix an FPA model \mathcal{M} , a tuple partition \mathcal{P} of $\mathcal{G}(\mathcal{M})$ and let $\sim_{\mathcal{P}}$ be a label equivalence on \mathcal{P} . In particular, let $\rho_k^{i,j} : \mathcal{B}(P_k^i) \rightarrow \mathcal{B}(P_k^j)$ denote a set of bijections that relates any two label equivalent tuples $\mathbf{P}^i, \mathbf{P}^j \in \mathcal{P}$ and that satisfies $(\rho_k^{i,j})^{-1} = \rho_k^{j,i}$. Then, $\{\{x_P : Q \in Z\} : Z \in \mathcal{B}(\mathcal{M})/\approx\}$ is a BDE partition, where $Q \approx Q'$ whenever $Q' = \rho_k^{i,j}(Q)$ for some i, j and k .*

There exist BDE partitions that are not induced by label equivalence. To see this, consider the FPA model $\mathcal{M} = P_1$ with $P_1 = (\alpha, 1).P_2$, $P_2 = (\alpha, 2).P_3$, $P_3 = (\alpha, 1).P_4$ and $P_4 = (\alpha, 2).P_1$. Then, it can be easily verified that the partition $\{\{x_{P_1}, x_{P_3}\}, \{x_{P_2}, x_{P_4}\}\}$ of the underlying IDOL model is BDE. However, it cannot be constructed using label equivalence since this notion relates FPA constants of *distinct* elements of $\mathcal{G}(\mathcal{M})$.

5. Conclusion

We have provided a generic framework for reasoning about languages that have ordinary differential equations (ODEs) as their quantitative semantics. Three main principles can be borrowed from more traditional domains based on labeled transition systems or discrete-state stochastic processes such as Markov chains: program comparison and minimization are understood in terms of equivalence relations over the states of a program; partition-refinement algorithms can be used to compute the largest equivalences; and SMT can be used for program verification. Yet the technical details involved in this transplantation are somewhat intricate: in ODE semantics, the state space is implicitly given as a continuous function. Therefore, proving programs equivalent involves a universal quantification over an uncountable domain. We developed algorithms for our differential equivalences by exploiting the possibility of reasoning over the reals symbolically using SMT.

We have worked on a basic intermediate language for ODEs. Conceptually, it can be seen as the analogous of a “bytecode” format for higher-level languages, where differential equivalences are compiler-optimization techniques that transform the original program while exactly preserving its behavior. Reasoning at such an intermediate level leads to equivalences that are more general than analogous notions developed for higher-level languages, because no domain-specific elements and issues are involved (such as action types and compositionality in process algebra). This can lead to potentially coarser minimizations, as our case studies of biochemical reaction networks have shown. However we argue that our contribution can still be useful when the modeller must work with higher-level equivalences to account for domain specificity. In this case, establishing a relationship with a differential equivalence may provide a way to automatizing checks. It appears already to be the case for label equivalence for FPA [77], which can be encoded into SMT, being a backward differential equivalence.

There are interesting avenues of future research from this work. The most direct one is that of improving the performance and scalability of the presented algorithms. For example, validity checks of different formulae can be performed independently, thus allowing to refine more blocks at a time, or to parallelize the refinement of a single block. Also, it is natural to relax the assumption of exactness in favor of approximate equivalence relations, similarly to what has been done for models with stochastic semantics [67]. Using SMT to compute differential equivalences opens a number of possibilities for symbolic computation. Our intermediate language could be extended with parameter variables in order to find, for instance, equivalences that hold under any possible assignment of such variables; or synthesize assignments for which a candidate partition is a differential equivalence. The ability to reason symbolically can be particularly useful in domains such as computational biology, where uncertainty on rate parameters is a well-known hindrance.

Acknowledgement

The authors thank Michele Loreti for helpful discussions. This work was partially supported by the EU project QUANTICOL, 600708. Luca Cardelli is partially funded by a Royal Society Research Professorship.

References

- [1] MAPK cascade in yeast - dimerization of Ste5. Available at <http://vcell.org/bionetgen/samples.html>.
- [2] A. Antoulas. *Approximation of Large-Scale Dynamical Systems*. Advances in Design and Control. SIAM, 2005.
- [3] U. M. Ascher and L. R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, 1988.
- [4] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. Model-checking continuous-time Markov chains. *ACM Trans. Comput. Logic*, 1(1): 162–170, 2000.
- [5] C. Baier and J.-P. Katoen. *Principles of model checking*. MIT Press, 2008. ISBN 978-0-262-02649-9.
- [6] C. Baier, B. Engelen, and M. E. Majster-Cederbaum. Deciding bisimilarity and similarity for probabilistic processes. *J. Comput. Syst. Sci.*, 60(1):187–231, 2000.
- [7] C. Barrett, R. Sebastiani, S. A. Seshia, S. A. S. Cesare Tinelli Clark Barrett, Roberto Sebastiani, and C. Tinelli. *Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications*, chapter Satisfiability Modulo Theories. IOS Press, 2009.
- [8] D. Barua and B. Goldstein. A mechanistic model of early FcεRI signaling: lipid rafts and the question of protection from dephosphorylation. *PLoS One*, 7(12), 2012.
- [9] J. Berdine and N. Bjørner. Computing all implied equalities via SMT-based partition refinement. In *Automated Reasoning*, volume 8562 of *LNCS*, pages 168–183. Springer, 2014.
- [10] M. Bernardo. A survey of Markovian behavioral equivalences. In *Formal Methods for Performance Evaluation*, volume 4486 of *LNCS*, pages 180–219. Springer, 2007.
- [11] M. L. Blinov, J. R. Faeder, B. Goldstein, and W. S. Hlavacek. BioNet-Gen: software for rule-based modeling of signal transduction based on the interactions of molecular domains. *Bioinformatics*, 20(17): 3289–3291, 2004.
- [12] P. Buchholz. Exact and ordinary lumpability in finite markov chains. *Journal of Applied Probability*, 31(1):59–75, 1994.
- [13] P. Buchholz. Exact performance equivalence: An equivalence relation for stochastic automata. *Theoretical Computer Science*, 215(1–2): 263–287, 1999.
- [14] L. Calzone, F. Fages, and S. Soliman. BIOCHAM: an environment for modeling biological systems and formalizing experimental knowledge. *Bioinformatics*, 22(14):1805–1807, 2006.
- [15] F. Camporesi and J. Feret. Formal reduction for rule-based models. *Electronic Notes in Theoretical Computer Science*, 276:29–59, 2011.
- [16] F. Camporesi, J. Feret, H. Koepl, and T. Petrov. Combining model reductions. *Electronic Notes in Theoretical Computer Science*, 265: 73–96, 2010.
- [17] L. Cardelli. On process rate semantics. *Theoretical Computer Science*, 391(3):190–215, 2008.
- [18] L. Cardelli. Morphisms of reaction networks that couple structure to function. *BMC Systems Biology*, 8(1):84, 2014.
- [19] L. Cardelli and A. Csikász-Nagy. The cell cycle switch computes approximate majority. *Sci. Rep.*, 2, 2012.
- [20] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Forward and backward bisimulations for chemical reaction networks. In *CONCUR*, pages 226–239, 2015.
- [21] F. Ciocchetta and J. Hillston. Bio-PEPA: A framework for the modelling and analysis of biological systems. *Theoretical Computer Science*, 410 (33-34):3065–3084, 2009.
- [22] J. Colvin, M. I. Monine, J. R. Faeder, W. S. Hlavacek, D. D. V. Hoff, and R. G. Posner. Simulation of large-scale rule-based models. *Bioinformatics*, 25(7):910–917, 2009.
- [23] J. Colvin, M. I. Monine, R. N. Gutenkunst, W. S. Hlavacek, D. D. V. Hoff, and R. G. Posner. Rulemonkey: software for stochastic simulation of rule-based models. *BMC Bioinformatics*, 11:404, 2010.
- [24] H. Conzelmann, J. Saez-Rodriguez, T. Sauter, B. Kholodenko, and E. Gilles. A domain-oriented approach to the reduction of combinatorial complexity in signal transduction networks. *BMC Bioinformatics*, 7(1): 34, 2006.
- [25] H. Conzelmann, D. Fey, and E. Gilles. Exact model reduction of combinatorial reaction networks. *BMC Systems Biology*, 2(1):78, 2008.
- [26] V. Danos and C. Laneve. Formal molecular biology. *Theoretical Computer Science*, 325(1):69–110, 2004.
- [27] V. Danos, J. Desharnais, F. Laviolette, and P. Panangaden. Bisimulation and confluence for probabilistic systems. *Information and Computation*, 204(4):503–523, 2006.
- [28] V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine. Abstracting the differential semantics of rule-based models: Exact and automated model reduction. In *LICS*, pages 362–381, 2010.
- [29] R. David and H. Alla. *Discrete, Continuous, and Hybrid Petri Nets*. Springer, 2005.
- [30] L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340, 2008.
- [31] R. De Nicola, D. Latella, M. Loreti, and M. Massink. A uniform definition of stochastic process calculi. *ACM Computing Surveys*, 46 (1):5:1–5:35, 2013.
- [32] C. Dehnert, J.-P. Katoen, and D. Parker. SMT-based bisimulation minimisation of Markov models. In *VMCAI*, volume 7737 of *LNCS*, pages 28–47, 2013.
- [33] S. Derisavi, H. Hermanns, and W. H. Sanders. Optimal state-space lumping in Markov chains. *Inf. Process. Lett.*, 87(6):309–315, 2003.
- [34] J. R. Faeder, W. S. Hlavacek, I. Reischl, M. L. Blinov, H. Metzger, A. Redondo, C. Wofsy, and B. Goldstein. Investigation of early events in FcεRI-mediated signaling using a detailed mathematical model. *The Journal of Immunology*, 170(7):3769–3781, 2003.
- [35] J. Feret. Fragments-based model reduction: Some case studies. *Electronic Notes in Theoretical Computer Science*, 268:77–96, 2010.
- [36] J. Feret, V. Danos, J. Krivine, R. Harmer, and W. Fontana. Internal coarse-graining of molecular systems. *Proceedings of the National Academy of Sciences*, 106(16):6453–6458, 2009.
- [37] J. Feret, T. Henzinger, H. Koepl, and T. Petrov. Lumpability abstractions of rule-based systems. *Theoretical Computer Science*, 431: 137–164, 2012.
- [38] J. Fisher and T. Henzinger. Executable cell biology. *Nature Biotechnology*, 25(11):1239–1249, 2007.
- [39] V. Galpin. Continuous approximation of PEPA models and Petri nets. *International Journal of Computer Aided Engineering and Technology*, 2:324–339, 2010.
- [40] S. Gao, S. Kong, and E. Clarke. Satisfiability modulo ODEs. In *FMCAD*, pages 105–112, 2013.
- [41] S. Gulwani, S. Jha, A. Tiwari, and R. Venkatesan. Synthesis of loop-free programs. In *PLDI*, pages 62–73, 2011.
- [42] R. A. Hayden and J. T. Bradley. A fluid analysis framework for a Markovian process algebra. *Theoretical Computer Science*, 411(22-24):2260–2297, 2010.
- [43] M. Heiner, D. Gilbert, and R. Donaldson. Petri nets for systems and synthetic biology. In *Formal Methods for Computational Systems Biology*, volume 5016 of *LNCS*, pages 215–264. Springer, 2008.
- [44] H. Hermanns and M. Rettelbach. Syntax, semantics, equivalences, and axioms for MTIPP. In *Proceedings of Process Algebra and Probabilistic Methods*, pages 71–87, Erlangen, 1994.
- [45] H. Hermanns and M. Siegle. Bisimulation algorithms for stochastic process algebras and their BDD-based implementation. In *ARTS*, pages 244–264, 1999.
- [46] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
- [47] J. Hillston. Fluid flow approximation of PEPA models. In *QEST*, pages 33–43, Sept. 2005.
- [48] D. T. Huynh and L. Tian. On some equivalence relations for probabilistic processes. *Fundam. Inform.*, 17(3):211–234, 1992.

- [49] G. Iacobelli, M. Tribastone, and A. Vandin. Differential bisimulation for a Markovian process algebra. In *MFCS*, pages 293–306, 2015.
- [50] D. Jovanovic and L. M. de Moura. Solving non-linear arithmetic. In *IJCAR*, pages 339–354, 2012.
- [51] P. C. Kanellakis and S. A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. *Inf. Comput.*, 86(1): 43–68, 1990.
- [52] J. Kemeny and J. Snell. *Finite Markov Chains*. Springer New York, Heidelberg, Berlin, 1976.
- [53] P. Kocieniewski, J. R. Faeder, and T. Lipniacki. The interplay of double phosphorylation and scaffolding in MAPK pathways. *Journal of Theoretical Biology*, 295:116–124, 2012.
- [54] A. S. Köksal, V. Kuncak, and P. Suter. Constraints as control. In *POPL*, pages 151–164, 2012.
- [55] M. Kwiatkowska, G. Norman, and D. Parker. Prism 4.0: Verification of probabilistic real-time systems. In *CAV*, pages 585–591, 2011.
- [56] M. Kwiatkowski and I. Stark. The continuous pi-calculus: A process algebra for biochemical modelling. In *CMSB*, pages 103–122, 2008.
- [57] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inf. Comput.*, 94(1):1–28, 1991.
- [58] Y. Li, A. Albarghouthi, Z. Kincaid, A. Gurfinkel, and M. Chechik. Symbolic optimization with SMT solvers. In *POPL*, pages 607–618, 2014.
- [59] S. Mover, A. Cimatti, A. Tiwari, and S. Tonetta. Time-aware relational abstractions for hybrid systems. In *EMSOFT*, pages 1–10, 2013.
- [60] M. Nagasaki, S. Onami, S. Miyano, and H. Kitano. Bio-calculus: Its concept and molecular interaction. *Genome Informatics*, 10:133–143, 1999.
- [61] D. L. Nelson and M. M. Cox. *Lehninger Principles of Biochemistry*. Palgrave Macmillan, 6th edition, 2013.
- [62] J. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.
- [63] M. S. Okino and M. L. Mavrouniotis. Simplification of mathematical models of chemical reaction systems. *Chemical Reviews*, 2(98):391–408, 1998.
- [64] R. Paige and R. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
- [65] G. J. Pappas. Bisimilar linear systems. *Automatica*, 39(12):2035–2047, 2003.
- [66] M. Pedersen and G. Plotkin. A language for biochemical systems: Design and formal specification. In *Transactions on Computational Systems Biology XII*, volume 5945 of *LNCIS*, pages 77–145. Springer, 2010.
- [67] A. D. Pierro, C. Hankin, and H. Wiklicky. Quantitative relations and approximate process equivalences. In *CONCUR*, pages 498–512, 2003.
- [68] A. Regev and E. Shapiro. Cellular abstractions: Cells as computation. *Nature*, 419(6905):343–343, 2002.
- [69] S. Sankaranarayanan and A. Tiwari. Relational abstractions for continuous and hybrid systems. In *CAV*, pages 686–702, 2011.
- [70] M. W. Sneddon, J. R. Faeder, and T. Emonet. Efficient modeling, simulation and coarse-graining of biological complexity with NFsim. *Nature Methods*, 8(2):177–183, 2011.
- [71] J. Sproston and S. Donatelli. Backward bisimulation in Markov chain model checking. *IEEE Trans. Software Eng.*, 32(8):531–546, 2006.
- [72] W. J. Stewart. *Probability, Markov Chains, Queues, and Simulation*. Princeton University Press, 2009.
- [73] R. Suderman and E. J. Deeds. Machines vs. ensembles: Effective MAPK signaling through heterogeneous sets of protein complexes. *PLoS Comput. Biol.*, 9(10):e1003278, 10 2013.
- [74] J. Toth, G. Li, H. Rabitz, and A. S. Tomlin. The effect of lumping and expanding on kinetic differential equations. *SIAM Journal on Applied Mathematics*, 57(6):1531–1556, 1997.
- [75] M. Tribastone. A fluid model for layered queueing networks. *IEEE Trans. Software Eng.*, 39(6):744–756, 2013.
- [76] M. Tribastone, S. Gilmore, and J. Hillston. Scalable differential analysis of process algebra models. *IEEE Trans. Software Eng.*, 38(1):205–219, 2012.
- [77] M. Tschaikowski and M. Tribastone. Exact fluid lumpability for Markovian process algebra. In *CONCUR*, pages 380–394, 2012.
- [78] A. Valmari and G. Franceschinis. Simple $O(m \log n)$ time Markov chain lumping. In *TACAS*, pages 38–52, 2010.
- [79] A. J. van der Schaft. Equivalence of dynamical systems by bisimulation. *IEEE Transactions on Automatic Control*, 49, 2004.
- [80] E. O. Voit. Biochemical systems theory: A review. *ISRN Biomatematics*, 2013:53, 2013. URL <http://dx.doi.org/10.1155/2013/897658>.