

# Challenges in Quantitative Abstractions for Collective Adaptive Systems

Mirco Tribastone

IMT School for Advanced Studies Lucca, Italy

`mirco.tribastone@imtlucca.it`

Like with most large-scale systems, the evaluation of quantitative properties of collective adaptive systems is an important issue that crosscuts all its development stages, from design (in the case of engineered systems) to runtime monitoring and control. Unfortunately it is a difficult problem to tackle in general, due to the typically high computational cost involved in the analysis. This calls for the development of appropriate quantitative abstraction techniques that preserve most of the system's dynamical behaviour using a more compact representation. This paper focuses on models based on ordinary differential equations and reviews recent results where abstraction is achieved by aggregation of variables, reflecting on the shortcomings in the state of the art and setting out challenges for future research.

## 1 Introduction

Collective adaptive systems (CAS) generally consist of a large number of entities that evolve according to local interactions, across multiple time and space scales, in order to achieve their own objectives. Their behaviour is influenced by such interactions as well as by changes in the environment. As a result, one typically deals with a highly complex system where the overall dynamics cannot be directly inferred from the analysis of the individual entities taken in isolation. This poses a serious problem to our capability of reasoning about CAS effectively. Indeed, for analysis purposes the availability of a model (e.g., an analytical description or discrete-event simulation) greatly helps understand the behaviour under a variety of conditions which would be otherwise difficult, or even impossible, to exercise on the concrete system under consideration. However, the impossibility of studying isolated CAS entities in a model requires the construction of the *product space* of the system that explicitly accounts for all possible interactions between agents. Because of the large scale involved, this clearly incurs the well-known problem of state explosion, affecting both qualitative and quantitative analysis efforts. In the latter case, it is worth to note that state explosion is not merely a computational problem: higher level of detail in the model implies a larger parameter space, which requires substantial effort for measurement, calibration, and validation.

By *abstraction* we mean a range of techniques where the basic idea is to construct a smaller CAS model, with fewer variables and/or parameters, that can yet preserve most of the dynamics of the original one. Such a more compact representation may be useful, for example, to provide a more intelligible description that ignores original low-level details, or to perform computational analyses more efficiently. Many branches of science and engineering that deal with CAS-like systems have acknowledged the importance of developing abstraction techniques, for instance:

- Ecological systems are a prototypical example of (natural) CAS with massive populations of entities interacting with and adapting to the environment. And ecology is a discipline where quantitative abstractions have long been understood (e.g., [18]), motivated by the clear presence of

multiple levels of hierarchy—from molecules to whole organisms in their ecosystem—and a high degree of heterogeneity, e.g., individuals of different age and metabolism, at various locations in space. Here, a large body of work has considered abstractions through *aggregations* of large-scale dynamical models of ecosystems, e.g., based on differential or difference equations. Aggregation is a coarsening of the state space that leads to a self-consistent smaller dynamical system where each “macro-variable” appropriately represents a group of “micro-variables” in the original model. Aggregations have been developed to ignore, e.g., age differences in population models [18] or exploit the separation of time scales typical of ecological systems (see [2] for a survey).

- In physics, cellular automata are a basic model to describe complex phenomena such as chaos and fractals, arising from simple interactions between agents with small state spaces. *Coarse graining* methods identify groups of neighbouring cells for which an aggregate dynamics can still describe the overall behaviour of the original model [17].
- In computational systems biology, a wealth of abstraction techniques have been developed to cope with the combinatorial explosion of the state space of mechanistic dynamical models based on ordinary differential equations (ODEs) for the modelling of protein interaction networks (e.g. [5, 14]). Here, available methods consider a *covering* of the state space (where a variable may appear in more than one group) [14]; *quotienting*, induced by a partition of the ODE variables (e.g., [15]); and aggregations exploiting time-scale separation (e.g., [23]).
- Large-scale dynamical systems are routinely encountered in control engineering. Here, starting with Aoki [1], numerous approaches have considered the transformation of the original model into a reduced one that preserves *controllability*, i.e., the capability of bringing the system to a desired state by an appropriate choice of control inputs [26].

A common feature shared by these examples is that CAS can be effectively described using dynamical systems. Even when the original description is a Markov chain, the underlying behaviour is characterised by a (large-scale) system of difference/differential equations (i.e., the forward equations of motion of the probability distribution). This paper briefly reviews recently developed techniques that consider the abstraction problem from an algorithmic viewpoint [7, 9, 8], with the intent of computing reduced dynamical systems that enjoy the following properties:

- P1.** The abstraction should come with formal guarantees on the relationship between the abstract dynamics and the original one. This enables the modeller to use the abstract model with full confidence in the results of the analysis.
- P2.** The construction of the abstract model should be fully automatic, since the original model is likely to be unintelligible due to size.
- P3.** The method should be generic in order to be applicable to as wide a range of CAS models as possible.
- P4.** The abstract model should preserve user-defined observables of the original system. For instance, it should be possible to fully recover the dynamics of selected variables of the original model.

The techniques discussed here revolve around the notion of *differential equivalence*, an equivalence relation over the variables of a dynamical system that induces a reduced model where each macro-variable represents the aggregate dynamics of an equivalence class. Although the theory considers first-order ODEs, it straightforwardly carries over to their discrete-time analogues. After a brief introduction in Section 2, some challenges for future research in this area are discussed in detail in Section 3.

## 2 Differential Equivalences

According to the terminology above, differential equivalences induce a quotienting of the ODE variables. Two distinct flavours have been provided in [9]. *Forward differential equivalence* (FDE) is such that a macro-variable describes the sum of the variables of an equivalence class. For instance, consider:

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

where  $k_1$  and  $k_2$  are constants and the ‘dot’ operator denotes the derivative operator (with respect to time). Then,  $\{\{x_1\}, \{x_2, x_3\}\}$  is an FDE because

$$\dot{x}_1 = -x_1, \quad (\dot{x}_2 + \dot{x}_3) = \dot{x}_2 + \dot{x}_3 = (k_1 + k_2) \cdot x_1 - (x_2 + x_3). \quad (2)$$

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

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

In this quotient model, whenever the *initial condition* (at time point 0) satisfies  $y(0) = x_2(0) + x_3(0)$  we get that  $y(t) = x_2(t) + x_3(t)$  at all time points  $t$ .

*Backward differential equivalence* (BDE) equates variables that have the same solutions at all time points. In (1),  $\{\{x_1\}, \{x_2, x_3\}\}$  is also a BDE provided that  $k_1 = k_2$ . In this case, we obtain a quotient ODE by removing either equation between  $x_2$  and  $x_3$ , say  $x_3$ , and rewriting every occurrence of  $x_3$  as  $x_2$ :

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

Both FDE and BDE satisfy **P1** because the relationship between the original model and the abstract one is exact; there is, however, loss of information when FDE is applied because the individual traces of the members of an equivalence class may not be recovered in general. Differential equivalences are closely related to the notion of exact ODE lumpability, very well understood in the chemistry literature (e.g., [27, 24, 22]). However this approach lacks of an automatic way of identifying lumping schemes (e.g., [31]). To cope with this, i.e., to satisfy **P2**, restrictions are imposed to be able to develop minimisation algorithms.

**Symbolic minimisation.** In [9] each ODE variable is treated explicitly as a real function and a differential equivalence is encoded in a logical formula over ODE variables. Thus, checking whether a candidate partition is BDE/FDE can be done *symbolically* using an encoding into satisfiability modulo theories (SMT) [4]. In fact, differential equivalences belong to the quantifier-free fragment of first-order logic. It is possible to restrict the admissible ODE systems to those for which an SMT solver for nonlinear real arithmetic — e.g., Z3 [10] — is a decision procedure. This can be done by, roughly speaking, excluding trigonometric functions (somewhat satisfying **P3**).

Let us consider the example (1), assuming  $k_1 = k_2 = 1$ . The condition for  $\{\{x_1\}, \{x_2, x_3\}\}$  to be a BDE can be shown to correspond to requiring that related variables with equal assignments always have equal derivatives. This can be encoded in a logical formula  $\phi$  thus:

$$\phi := x_2 = x_3 \Rightarrow k_1 \cdot x_1 - x_2 = k_2 \cdot x_1 - x_3.$$

The SMT check  $\text{sat}(\neg\phi)$  looks for an assignment of the variables  $x_1$ ,  $x_2$ , and  $x_3$  for which  $\neg\phi$  holds. Thus, the partition is a BDE if and only if the procedure returns “unsat” (as it is obviously the case in

this example). More interestingly, it is possible to exploit the ability of the solver to return a *witness* in case of satisfiability. This can be interpreted as a counterexample that distinguishes variables originally supposed to be equivalent. For instance, an SMT check for the candidate BDE partition  $\{\{x_1, x_2, x_3\}\}$  might return the witness  $(x_1 = 1, x_2 = 1, x_3 = 1)$ , which yields derivatives that are not equivalent ( $\dot{x}_1 = -1, \dot{x}_2 = 0, \dot{x}_3 = 0$ ). This suggests the implementation of an algorithm that *splits* the partition in such a way to preserve the equalities in the witness. That is, at the next iteration the candidate BDE partition would be  $\{\{x_1\}, \{x_2, x_3\}\}$ . It turns out that such an algorithm does iteratively compute the largest BDE that refines a given initial partition of ODE variables. This algorithm also meets **P4**: indeed each variable that should be treated as an *observable* can be put in a singleton initial block.

**Syntax-driven minimisation.** A more efficient minimisation algorithm can be provided for ODEs with derivatives that are multivariate polynomials of degree at most two [8]. This covers the ubiquitous linear systems as well as chemical reaction networks, at the basis of the aforementioned dynamic models in systems biology. At the basis of this approach is a finitary representation of an ODE system as a so-called *reaction network* (RN), consisting of species/variables interacting by means of reactions parameterised by a real value. On this representation two bisimulation equivalences, the forward and backward RN bisimulations, are related to FDE and BDE, respectively [7]. This makes such bisimulations similar in spirit to quantitative equivalences on labelled transition systems, e.g., Larsen and Skou’s probabilistic bisimulation [21]. In particular, the computation of the largest RN bisimulations that refine a given partition can be computed using an appropriate variant of Paige and Tarjan’s famous algorithm [25]. In [8] a partition refinement algorithm is developed along the lines of efficient analogues for Markov chain lumping such as [12] and [32], and for probabilistic transition systems [3]. This computes the largest FB/BB refining a given partition of variables in  $O(mn \log n)$  time, where  $m$  is the number of monomials in the ODE system and  $n$  is the number of variables.

### 3 Outlook

The differential equivalences reviewed in this paper allow automatic reductions of ODE systems. The benchmarks in [8] show that the algorithms for reduction up to forward and backward bisimulations can scale to systems with millions of variables and monomials, terminating in a few seconds also in some challenging models. There are, however, further challenges ahead which we wish to tackle.

Forward bisimulation is only a sufficient condition for FDE; while it has been shown to yield significant reductions in practice [7, 8], some other examples from the literature demonstrate that the algorithm may miss some FDE reductions (see [9]). Ongoing work is aiming to develop a new variant of forward bisimulation that characterises FDE (for multivariate polynomial derivatives of degree at most two). Backward bisimulation, on the other hand, does characterise BDE. In this case, the specialised partition-refinement algorithm of [8] is to be preferred over the symbolic SMT-based approach of [9] for computational reasons, since we observed runtimes generally separated by two/three orders of magnitude in our experiments. This is an unsurprising fact because the SMT-based approach does not exploit the finitary RN representation of [8], where no symbolic computation is performed. We plan to exploit the advantage of the RN representation, extending [8] to polynomial derivatives with arbitrary degree.

The whole SMT technology has been treated as a black box in our proof-of-concept experiments, despite the well-known fact that appropriate heuristics for SMT can sensibly affect the runtimes. Developing problem-specific solution strategies for differential equivalences is an interesting line of research; a natural question to investigate is the possibility of parallelising the SMT checks at every iteration of

the partition refinement. Improving the scalability of the SMT approach will be advantageous for all such models that do not feature degree-two polynomial derivatives. Examples are: third-body reactions (giving degree-three polynomials), appearing frequently in chemical engineering [31]; chemical reaction networks that are based on different kinetics than mass-action, such as Hill's [33]; and ODE performance models of computing systems, such as those based on minimum-rate semantics (e.g., stochastic process algebra [29] and queuing networks [28]). In addition, also the models that admit the RN representation may benefit from improvements made in our SMT approach. For instance, the symbolic checks may be used in models with uncertainties in rates (a well-known issue in mathematical biology): here the SMT framework can already be easily extended to compute partitions that are differential equivalences under all possible assignments of such uncertain parameters, left as free variables in the satisfiability problem (similarly to the SMT-based parametric minimisation approach of [11] for probabilistic models written in PRISM [19]).

Lastly, it has long been argued that *exact* reduction techniques such as the ones presented here are too restrictive because they are sensitive to the values of the parameters. It is natural to consider approximate variants (e.g., [6, 13, 20]) as weaker notions that, for instance, allow variability in the parameters, considering the exact versions as a degenerate case in which no such variability is needed. We remark that, at least for the numerical benchmarks considered so far with realistic models, the exact reductions can already be quite effective. Of course, approximate ones might be able to provide even coarser descriptions. In this case, however, the main challenge is to be able to relate the variability in the parameters tolerated by the coarsening procedure with the error incurred when considering an approximate, smaller model, instead of the original one. This issue has been attempted by some previous work which has provided an asymptotic result of correctness of the approximation for small enough perturbations [16]. This has been further improved upon using differential inequalities in [30], providing more usable bounds that however tend to degrade for larger perturbations. In this respect, two orthogonal lines of research may look into the problem of improving the quality of the bounds and automatically detecting near-symmetries to synthesise candidate approximately reduced models.

**Acknowledgement.** This work is partially supported by the EU project QUANTICOL (<http://www.quanticol.eu>), 600708. The author is indebted to Luca Cardelli, Max Tschaikowski, and Andrea Vandin for their collaboration in [7, 9, 8] and the numerous fruitful discussions.

## References

- [1] M. Aoki (1968): *Control of large-scale dynamic systems by aggregation*. *IEEE Trans. Autom. Control* 13(3), pp. 246–253, doi:10.1109/TAC.1968.1098900.
- [2] P. Auger, R. de la Parra, J. Poggiale, E. Sánchez & T. Nguyen-Huu (2008): *Aggregation of Variables and Applications to Population Dynamics*. In Pierre Magal & Shigui Ruan, editors: *Structured Population Models in Biology and Epidemiology, Lecture Notes in Mathematics* 1936, Springer, pp. 209–263, doi:10.1007/978-3-540-78273-5\_5.
- [3] C. Baier, B. Engelen & M. E. Majster-Cederbaum (2000): *Deciding Bisimilarity and Similarity for Probabilistic Processes*. *J. Comput. Syst. Sci.* 60(1), pp. 187–231, doi:10.1006/jcss.1999.1683.
- [4] C.W. Barrett, R. Sebastiani, S.A. Seshia & C. Tinelli (2009): *Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications*, chapter Satisfiability Modulo Theories. IOS Press.
- [5] N.M. Borisov, N.I. Markevich, J.B. Hoek & B.N. Kholodenko (2005): *Signaling through Receptors and Scaffolds: Independent Interactions Reduce Combinatorial Complexity*. *Biophysical Journal* 89(2), pp. 951–966, doi:10.1529/biophysj.105.060533.

- [6] F. van Breugel & J. Worrell (2006): *Approximating and computing behavioural distances in probabilistic transition systems*. *Theoretical Computer Science* 360(1–3), pp. 373–385, doi:10.1016/j.tcs.2006.05.021.
- [7] L. Cardelli, M. Tribastone, M. Tschaikowski & A. Vandin (2015): *Forward and Backward Bisimulations for Chemical Reaction Networks*. In: *26th International Conference on Concurrency Theory, CONCUR*, pp. 226–239, doi:10.4230/LIPIcs.CONCUR.2015.226.
- [8] L. Cardelli, M. Tribastone, M. Tschaikowski & A. Vandin (2016): *Efficient Syntax-driven Lumping of Differential Equations*. In: *Proceedings of Tools and Algorithms for the Construction and Analysis of Systems — 21st International Conference (TACAS)*, doi:10.1007/978-3-662-49674-9\_6.
- [9] L. Cardelli, M. Tribastone, M. Tschaikowski & A. Vandin (2016): *Symbolic computation of differential equivalences*. In: *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL*, pp. 137–150, doi:10.1145/2837614.2837649.
- [10] L. De Moura & N. Bjørner (2008): *Z3: An Efficient SMT Solver*. In: *TACAS*, pp. 337–340, doi:10.1007/978-3-540-78800-3\_24.
- [11] C. Dehnert, J.-P. Katoen & D. Parker (2013): *SMT-Based Bisimulation Minimisation of Markov Models*. In R. Giacobazzi, J. Berdine & I. Mastroeni, editors: *Proc. 14th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI’13)*, LNCS 7737, Springer, pp. 28–47, doi:10.1007/978-3-642-35873-9\_5.
- [12] S. Derisavi, H. Hermanns & W.H. Sanders (2003): *Optimal state-space lumping in Markov chains*. *Information Processing Letters* 87(6), pp. 309–315, doi:10.1016/S0020-0190(03)00343-0.
- [13] J. Desharnais, V. Gupta, R. Jagadeesan & P. Panangaden (2004): *Metrics for labelled Markov processes*. *Theor. Comput. Sci.* 318(3), pp. 323–354, doi:10.1016/j.tcs.2003.09.013.
- [14] J. Feret, V. Danos, J. Krivine, R. Harmer & W. Fontana (2009): *Internal coarse-graining of molecular systems*. *Proceedings of the National Academy of Sciences* 106(16), pp. 6453–6458, doi:10.1073/pnas.0809908106.
- [15] J. Feret, T. Henzinger, H. Koepl & T. Petrov (2012): *Lumpability abstractions of rule-based systems*. *Theoretical Computer Science* 431, pp. 137–164, doi:10.1016/j.tcs.2011.12.059.
- [16] G. Iacobelli & M. Tribastone (2013): *Lumpability of fluid models with heterogeneous agent types*. In: *DSN*, pp. 1–11, doi:10.1109/DSN.2013.6575346.
- [17] N. Israeli & N. Goldenfeld (2006): *Coarse-graining of cellular automata, emergence, and the predictability of complex systems*. *Phys. Rev. E* 73, p. 026203, doi:10.1103/PhysRevE.73.026203.
- [18] Y. Iwasa, V. Andreassen & S. Levin (1987): *Aggregation in model ecosystems. I. Perfect aggregation*. *Ecological Modelling* 37(3–4), pp. 287–302, doi:10.1016/0304-3800(87)90030-5.
- [19] M.Z. Kwiatkowska, G. Norman & D. Parker (2011): *PRISM 4.0: Verification of Probabilistic Real-Time Systems*. In: *CAV*, pp. 585–591, doi:10.1007/978-3-642-22110-1\_47.
- [20] K.G. Larsen, R. Mardare & P. Panangaden (2012): *Taking It to the Limit: Approximate Reasoning for Markov Processes*. In: *MFCS*, pp. 681–692, doi:10.1007/978-3-642-32589-2\_59.
- [21] K.G. Larsen & A. Skou (1991): *Bisimulation through probabilistic testing*. *Information and Computation* 94(1), pp. 1–28, doi:10.1016/0890-5401(91)90030-6.
- [22] G. Li & H. Rabitz (1989): *A general analysis of exact lumping in chemical kinetics*. *Chemical Engineering Science* 44(6), pp. 1413–1430, doi:10.1016/0009-2509(89)85014-6.
- [23] J.D. Murray (2002): *Mathematical Biology I: An Introduction*, 3rd edition. Springer, doi:10.1007/b98868.
- [24] M.S. Okino & M.L. Mavrovouniotis (1998): *Simplification of Mathematical Models of Chemical Reaction Systems*. *Chemical Reviews* 2(98), pp. 391–408, doi:10.1021/cr950223l.
- [25] R. Paige & R. Tarjan (1987): *Three Partition Refinement Algorithms*. *SIAM Journal on Computing* 16(6), pp. 973–989, doi:10.1137/0216062.
- [26] G.J. Pappas (2003): *Bisimilar linear systems*. *Automatica* 39(12), pp. 2035–2047, doi:10.1016/j.automatica.2003.07.003.

- [27] J. Toth, G. Li, H. Rabitz & A.S. Tomlin (1997): *The Effect of Lumping and Expanding on Kinetic Differential Equations*. *SIAM Journal on Applied Mathematics* 57(6), pp. 1531–1556, doi:10.1137/S0036139995293294.
- [28] M. Tribastone (2013): *A Fluid Model for Layered Queueing Networks*. *IEEE Transactions on Software Engineering* 39(6), pp. 744–756, doi:10.1109/TSE.2012.66.
- [29] M. Tribastone, S. Gilmore & J. Hillston (2012): *Scalable Differential Analysis of Process Algebra Models*. *IEEE Transactions on Software Engineering* 38(1), pp. 205–219, doi:10.1109/TSE.2010.82.
- [30] M. Tschaikowski & M. Tribastone (2015): *Approximate reduction of heterogeneous nonlinear models with differential hulls*. *IEEE Transactions on Automatic Control*, doi:10.1109/TAC.2015.2457172.
- [31] T. Turanyi & A. Tomlin (2014): *Analysis of Kinetic Reaction Mechanisms*. Springer, doi:10.1007/978-3-662-44562-4.
- [32] A. Valmari & G. Franceschinis (2010): *Simple  $O(m \log n)$  Time Markov Chain Lumping*. In: *TACAS*, pp. 38–52, doi:10.1007/978-3-642-12002-2\_4.
- [33] Eberhard O. Voit (2013): *Biochemical Systems Theory: A Review*. *ISRN Biomathematics* 2013, p. 53, doi:10.1155/2013/897658.