# Qualitative Modelling via Constraint Programming

Thomas W. Kelsey[1], Lars Kotthoff[2], Christopher A. Jefferson[1], Stephen A. Linton[1], Ian Miguel[1], Peter Nightingale[1], and Ian P. Gent[1]

[1] School of Computer Science,
University of St Andrews, KY16 9SX, UK
[2] INSIGHT Centre for Data Analytics
University College Cork
Cork, Ireland

**Abstract.** Qualitative modelling is a technique integrating the fields of theoretical computer science, artificial intelligence and the physical and biological sciences. The aim is to be able to model the behaviour of systems without estimating parameter values and fixing the exact quantitative dynamics. Traditional applications are the study of the dynamics of physical and biological systems at a higher level of abstraction than that obtained by estimation of numerical parameter values for a fixed quantitative model. Qualitative modelling has been studied and implemented to varying degrees of sophistication in Petri nets, process calculi and constraint programming. In this paper we reflect on the strengths and weaknesses of existing frameworks, we demonstrate how recent advances in constraint programming can be leveraged to produce high quality qualitative models, and we describe the advances in theory and technology that would be needed to make constraint programming the best option for scientific investigation in the broadest sense.

**Keywords:** Constraint Programming, Qualitative Models

## 1 Introduction

The standard approach for non-computer scientists when investigating dynamic scientific systems is to develop a quantitative mathematical model. Differential equations are chosen in the belief that they best represent (for example) convection-diffusion-reaction or population change, and parameter values are estimated from empirical data. This approach suffers from several limitations which are widely documented, and which we summarise with examples in Section 2. In a standard modelling text [24, Chapter 5], qualitative model formulation is described as

> *... the conversion of an objective statement and a set of hypotheses and assumptions into an informal, conceptual model. This form does not contain explicit equations, but its purpose is to provide enough detail and*

> *structure so that a consistent set of equations can be written. The qualitative model does not uniquely determine the equations, but does indicate the minimal mathematical components needed. The purpose of a qualitative model is to provide a conceptual frame-work for the attainment of the objectives. The framework summarizes the modeler's current thinking concerning the number and identity of necessary system components (objects) and the relationships among them.*

Kuipers (1993) is more succinct: *All models are abstractions of the world. Qualitative models are related to ordinary differential equations, but are more expressive of incomplete knowledge* [33]. We note that our view of qualitative modelling (QM) is not based on order-sorted logic representations of dynamical systems (as described in [43]). Instead of running quantitative model simulations within the system that was used to specify a qualitative model, we explore the QM space in order to select quantitative models that are most suitable for simulations.

In QM, the dynamics of a system under investigation are described in a formal language, but with no (or few) *a priori* assumptions made about the specific mathematical model that may be produced. This means working at a higher level of abstraction than usual, it requires the formalisation of complex system behaviour, and it involves searching a large space of candidate models for those to be used to generate numerical models. Computer scientists are, in general, trained to be able to identify and work at the most suitable levels of abstraction; they also design and use highly formal languages, and routinely develop algorithms for $NP$-hard problem classes. Hence the computer scientist is ideally qualified to undertake qualitative modelling. This is by no means a new observation, and in Section 3 we give a critical evaluation of existing computer science approaches to this problem. We focus on three particular approaches, constraint programming (CP), temporal logics and process calculi. In our view, historic CP approaches were hindered by both struggles to accommodate temporality into constraints, and by limitations in the CP languages and tools available at the time. The process calculus and temporal logic approaches have been more successful, modelling important systems arising in molecular and cell biology.

The CP approach has been recently revisited, using languages and tools developed as part of the Constraint Solver Synthesiser research project at St Andrews. However, fundamental problems remain. In particular, our exploration of solution spaces is neither truly stochastic nor targeted enough to reduce non-useful search effort. Nor do we have any organised way to investigate the tradeoff between realism of qualitative model and computational complexity of quantitative model. We explore these and other limitations in Section 5, and present them as research opportunities for the CP community. Successful research activity would be beneficial to the scientific community in the widest sense. A system under investigation would be described in qualitative terms, such as:

- behaviour A is required and/or behaviour B is forbidden;
- if C happens, it happens after D;
- the second derivative of E has exactly two minima in timescale F;
- the rate of decline of G is less than the rate of change in the increase in H.

CP technology would then be used to converge iteratively on suitable models for use by the global scientific community. This would, in our opinion, represent an important transfer of CP expertise, languages and search to our colleagues working in other scientific fields.

## 2 Quantitative mathematical models

Successful computer modelling in the physical, biological and economic sciences is a difficult undertaking. Domains are often poorly measured due to ethical, technical and/or financial constraints. In extreme instances the collection of accurate longitudinal data is simply impossible using current techniques. This adversely affects the production and assessment of hypothetical quantitative models, since the incompleteness of the domain data necessitates the making of assumptions that may or may not reflect ground truths. A second category of assumptions are involved in the choice of quantitative modelling framework. Hypothetical solutions can be ruled out by restricting the complexity of models, and unrealistic models can be allowed by over-complex models. For both types of *a priori* assumption, mutually exclusive assumptions must be kept separate, sometimes with no scientific justification.

A motivating example involves the modelling of human cell populations. The human ovary contains a population of primordial (or non-growing) follicles. Some of these are recruited towards maturation and start to grow. Many of these die off through atresia, but some become primary follicles. Again, a proportion of these die off with the remainder growing into secondary follicles. This continues until a very small proportion become eggs that are released from the ovary for potential fertilisation. For the purposes of this study, we consider only the dynamics of follicle progression (primordial to primary to secondary). There are limitations to the quantitative, compartmental model approach. Empirical data is scarce for primordial follicles [46], is calculated by inference for primary follicles [30], and simply does not exist for secondary follicles [19]. As a direct result of these limitations, two entirely different and contradictory compartmental models have been published in the literature [6, 18]. A third research group investigating the same cell dynamics but with its own empirical data and modelling assumptions would be highly likely to produce a third quantitative model being fundamentally different to those already published. So there is an obvious problem: which (if any) of these models should be used by the wider research community to describe and account for changes in cell populations over time?

A second example (adapted from a paper by Degasperi and Calder presented at a workshop on Process Algebra and Stochastically Timed Activities [13]) of the limitations of starting the modelling process by selecting a mathematical model involves modelling nitric oxide (NO) bioavailability in blood vessels. Models of this scenario aim to determine the diffusion distance of NO along the radius of a vessel, where NO is produced in a narrow region on the internal wall of the vessel. Numerous models have been developed over the last decade and most share underlying assumptions and use the similar diffusion governing equations.

In particular, a vessel is modelled as a cylinder with partial differential equations (PDEs), using Fick's law of diffusion in cylindrical coordinates. Compartments define areas such as endothelium (where NO is produced), vascular wall, and lumen (i.e. where the blood flows). Another common assumption is that the diffusion operates only in the radial direction, while it can be considered negligible in other directions. A complete review and critical evaluation of these models is given in [45]. The author concludes: **Advances** *in both the experimental methodologies and* **in the theoretical models** *are required to further elucidate NO's roles in the vasculature* [45, our emphases].

We have demonstrated that these compartmental dynamics can be qualitatively modelled using finite element constraints [31]. Each of our qualitative models represents a class of CSPs. In general, there are many more solutions to the CSPs than realistic models, and many more realistic models than models that accurately reflect what happens in nature. Moreover, the resulting quantitative models can be graded by their complexity – linear ODE, piecewise-linear ODE, quadratic ODE, . . . , non-linear PDE, thereby providing insights on the computational effort needed to obtain solutions [32].

## 3 Existing approaches to QM

Qualitative modelling is a mature computer scientific technique, with existing methods and results for qualitative compartmental models [36, 35, 39] and for the use of CSPs to describe and solve qualitative models [12, 17]. However, these latter studies either reported incomplete algorithms [12] or described complicated algebras with no associated CSP modelling language or optimised CSP solver [17]. In 2002, a hybrid approach was presented in which concurrency was described in terms of CP constraints [5]. Since these studies were published 10–20 years ago, it appears that the limitations of CP technology at the time were collectively sufficient to stifle development.

Other approaches include process calculi and temporal logics, which are successful at the molecular level [7] and the protein network level [8, 42], but not as yet at inter- and intra-cellular levels. Despite this, the process calculus and temporal logic communities are engaging in active current research to improve techniques and widen access to other scientific areas. Of particular note are BIOCHAM (temporal logic) and BioPEPA (process calculus).

BIOCHAM [9] consists of two languages (one rule-based, the other based on either the CTL or LTL temporal logic languages) that allows the iterative development of quantitative models from qualitative ones. This answers the obvious question posed by newcomers to qualitative modelling: "given a good qualitative model, how do I derive a model that I can use for numeric studies?" BIOCHAM has sophisticated tool support and is under active current development (version 3.4 released in September 2012).

BioPEPA [11] is a process algebra for the modelling and the analysis of biochemical networks. It is a modification of PEPA (originally defined for the performance analysis of computer systems), in order to handle the use of general kinetic laws. The Edinburgh-based BioPEPA research group has received substantial funding to improve the accessibility of their framework by researchers at all levels of systems biology. A cloud-based architecture is under development, as is improved translation to and from SBML (System Biology Markup Language) formats, thereby supporting easier exchange and curation of models.

In summary, from the competing candidates for a computer science basis for successful qualitative modelling, CP has – as it were – fallen by the wayside, while temporal logics and process calculi are providing real support, at least to the biomedical modelling communities. We see no obvious reason for this: clearly time is a variable in all dynamical modelling, and therefore notions of "liveness", "before" and "after" needed to be incorporated into the qualitative modelling framework. But this is perfectly possible in CP, as demonstrated in [31].

## 4 Case study: cell dynamics QM using constraints

Our case study is the compartmental modelling of NGFs described in Section 2, where cell populations grow to a peak of unknown size and location, then decline with increasing age. We use the Savile Row tool that converts constraint problem models formulated in the solver-independent modelling language Essence$'$ [20] to the input format of the Constraint Satisfaction Problem (CSP) solver Minion [22].

We expect our candidate qualitative models to be implemented as differential equations or by non-linear curve-fitting. In both case we need to specify the notions of rate of change and smoothness. Suppose that $X[0,\ldots,n]$ is a series of variables representing a follicle population at different ages. Then we can approximate first derivatives by $X'[1,\ldots,n]$ where $X'[i] = X[i] - X[i-1]$, and second derivatives by $X''[1,\ldots,n-1]$ where $X''[j] = X'[j+1] - X'[j]$. These definitions allow us to post qualitative constraints about peak populations

$$\exists p \in [1,\ldots,n] \text{ such that } \forall i > p, X'[i] < 0 \wedge \forall i < p, X'[i] > 0.$$

We can require or forbid smoothness by restricting the absolute value of the $X''$ variables, and constrain rates of population growth by restrictions on the $X'[i]$. By having three sets of variables (one for each cell type) each with up to two derivative approximations, we can model interactions between the populations at different ages. To further abstract away from quantitative behaviour, populations can be defined in terms of proportion of peak rather than absolute numbers of cells, different time scales can be used for different age ranges (e.g. neonatal vs post-menopausal), and we can model the qualitative behaviour of values that are normally log-adjusted in quantitative studies. Table 1 gives an illustrative example of a model involving one type of follicle.

Any solution of such a model is a candidate for the basis of a quantitative model of actual cell dynamics, once boundary conditions and scale conditions are supplied. Using a combination of facts and quantitative information, a range of quantitative models can be produced for later empirical validation.

| Essence′ statement | Qualitative description |
|---|---|
| find $x$ : $[int(0..max)]$ of $int(0..100)$ | percentage of peak population |
| find $y$: $[int(1..max)]$ of $int(-r \cdots r)$ | 1st deriv. variables |
| find $z$ : $int(1..max - 1)]$ of $int(-r \cdots r)$ | 2nd deriv. variables |
| forAll $i : int(1..max).y[i] = x[i] - x[i-1]$ | 1st deriv.definition |
| forAll $j : int(1..max-1).z[j] = y[j+1] - y[j]$ | 2st deriv. definition |
| exists $k, j : int(2..birth)$. | |
|   forall $i : int(birth..max)$. | |
|     $i < k \Rightarrow y[i] > 0$ | positive 1st deriv. pre-peak |
|     $i > k \Rightarrow y[i] < 0$ | negative 1st deriv. post-peak |
|     $x[k] = 100 \wedge y[k] = 0$ | it is a peak |
|     $i > birth \Rightarrow \lvert z[i] \rvert < \lfloor \sqrt{r} \rfloor$ | smooth post-gestation |

**Table 1.** An example of a simple qualitative model specified in Essence′. When supplied with values for *max*, *r*, and *birth*, Savile Row will construct a Minion instance, the solutions of which are all hypothetical models that respect the qualitative description.

## 5 Future directions for CP

Currently, applying constraint technology to a large, complex problem requires significant manual tuning by an expert. Such experts are rare, so a natural aim is to improve dramatically the scalability of constraint technology, while simultaneously removing its reliance on manual tuning by an expert. It is our view that here are many techniques in the literature that, although effective in a limited number of cases, are not suitable for general use. Hence, they are omitted from current general solvers and remain relatively undeveloped. QM is an excellent example. There have been many important recent advances in CP technology. However, we are at the proof-of-concept stage for QM, having shown the ability in principle to produce useful results, rather than extensive research output. We now present specific avenues of research that would allow not only the production of high quality qualitative models, but also a robust schema for deriving a suitable quantitative model from the space of solutions of a CSP that represents a QM. The research areas are given in order of realisability: the first version of Savile Row (Section 5.1) was released in July 2012 and is under current active development, whereas the systematic search for models that are both realistic and lead to computationally inexpensive differential equations (Section 5.3) is a completely unexplored research topic.

### 5.1 Essence′ and Savile Row

Savile Row [44] is a modelling assistant tool that reads the language Essence′ and transforms it into the input format of a number of solvers (currently Minion [23], Gecode [21] and Dominion [2]). It was designed from the start to be solver-independent and easily extended with new transformation rules. It is also straightforward to add new output languages supported by an alternate sequence of transformations. At present Savile Row is at an early stage of development compared to other tools such as MiniZinc [38]. However it has some features that are particularly relevant to qualitative modelling, and its extensibility makes it suitable for the future work we describe below.

Uniquely Savile Row can produce Minion and Dominion's logical metaconstraints for conjunction and disjunction. This is highly relevant to qualitative modelling because disjunctions arise from exists statements, and conjunctions from forAll statements (when they are nested inside exists or some logical operator). Exists and forAll will be extensively used in qualitative modelling to model time. Minion's logical metaconstraints can be much more efficient than other methods [29].

Savile Row also implements common subexpression elimination (CSE) [40]. This replaces two or more equivalent expressions in a model with a single auxiliary variable. The auxiliary variable is then constrained to be equal to the common expression. In many cases CSE will strengthen propagation. CSEs tend to arise when quantifiers are unrolled, so we expect this feature to be relevant to QM because of the extensive use of quantifiers to model time. To fully exploit CSE for QM, we would firstly study the types of expressions that arise in QM for semantic equivalences and eliminate these through active reformulation CSE [40]. Secondly we would investigate sequences of reformulations of the model (or parts of it) to reveal CSEs and other implied constraints. A particular sequence of reformulations may hinder CSE with one model and help with another, therefore there is no generic solution and it is likely QM would benefit from research in this area.

### 5.2 Exploring Search Spaces I

With the exception of MDD- and BDD-based solvers [26], current CP solvers are tailored towards finding a single solution to a problem, or proving no solution exists. The solution found can be either the first one discovered, or the "best" solution under a single optimisation condition. For QM, this is insufficient, as users want to be able to understand and reason about *all* solutions to their problem. While current CP technologies allow the exploration of the entire search space, they provide no facilities to compare and reason about different solutions. Furthermore, many techniques are specialised for finding only a single solution and are ineffective when looking for all solutions. We believe CP solvers must be extended to be able to solve such problems, while maintaining and improving the efficiency and ease-of-use of existing CP tools.

As described in Section 2, each solution from a class of CSPs that describes a QM will not only capture the qualitative nature specified in the constraints, but also return specific information that may form the basis of a quantitative model. Hence we need to be able to explore the space of all CSP solutions. We currently have good techniques for recording numbers of solutions, but have only recently started work on compiling (representations of) solutions into data libraries that can be used by interested research groups [15, 14, 16]. Our current area of interest is finite algebraic structures; by curating and releasing comprehensive data libraries we allow users to search the solution space for “interesting” exemplars, obviating the need for them to re-solve the original CSPs.

We plan on extending CP so it can generate efficient compact representations of the solutions to problems, and allow users to explore and understand these solutions. This will allow CP to be used to tackle many new classes of problems, of interest to many different types of user.

### 5.3 Exploring Search Spaces II

In Section 5.2 we described issues to do with the efficient search of large solution spaces. However, even if efficiency is assured, there are two further problems to overcome if high quality QM is to be achieved. The first is the organisation of search in a controlled and stochastic way – i.e. using the mathematical theory of probability to express and utilise the inherent degrees of uncertainty in which qualitative model solutions are likely to lead to “good” quantitative models. Existing CP search heuristics allow the user to specify the order in which the variables and/or values are selected during search. This order can be randomised, but this is far from a fully stochastic exploration of the search space. Both BIOCHAM and BioPEPA (described in Section 3) fully support iterative stochastic simulation allowing convergence to preferred numeric models.

The second issue relates to the tradeoff between scientific accuracy and plausibility of a QM (as determined by testing generalisation to empirical data) and the mathematical and computational complexity of the preferred quantitative model. Qualitative models can be ranked in terms of realism in a continuum ranging from highly unrealistic to a highly accurate simulation of what we understand the system in question to be. The models can also be ranked in terms of the type of differential equations needed to implement a numeric simulation. Many simple systems of linear ODEs are solvable in polynomial time and space. Others are not (depending on Lipschitz conditions and whether or not $P = PSPACE$ [32]). Nonlinear ODEs are strictly harder to solve as a class, and most PDEs have no closed form solution. The complexity of obtaining approximate solutions follows the same scale, in general. It is clear that given two qualitative models that are roughly equivalent in terms of assessed realism, the one that leads to the differential equations that are easier to solve should normally be selected. The CP technology needed to make these decisions does not exist, and its development is a completely unexplored avenue of future research.

It may be the case that instead of exploring the solution space of pre-defined CSPs, progress could be made by exploring the space of (models of) CSPs that capture the required qualitative behaviour. This would build on exciting new developments in this area [4].

### 5.4 Solver Generation and Automatic Tuning

Our final future direction is not obviously directly associated with QM. However, we believe that a key component of constraints-based QM is the investigation of the realism–complexity tradeoff described in Section 5.3. The tradeoff situation is clearly improved when constraint solvers are improved in terms of efficiency and instance size. A major challenge facing constraints research is to deliver constraint solving that scales easily to problems of practical size. Current constraint solvers, such as Choco [34], Eclipse [1], Gecode [21], Ilog Solver [28], and Minion [23] are monolithic in design, accepting a broad range of models. This convenience comes at the price of a necessarily complex internal architecture, resulting in significant overheads and inhibiting efficiency and scalability. Each solver may thus incorporate a large number of features, many of which will not be required for most constraint problems. The complexity of current solvers also means that it is often prohibitively difficult to incorporate new techniques as they appear in the literature. A further drawback is that current solvers perform little or no analysis of an input model and the features of an individual model cannot be exploited to produce a more efficient solving process.

To mitigate these drawbacks, constraint solvers often allow manual tuning of the solving process. However, this requires considerable expertise, preventing the widespread adoption of constraints as a technique for solving the most challenging combinatorial problems. The components of a constraint solver are also usually tightly coupled, with complex restrictions on how they may be linked together, making automated generation of different solvers difficult.

Initial results from comparing solvers generated by Dominion with an existing solver are positive and indicate this approach is promising [3]. Dominion is in fact expected to make bigger gains in the cases where there are many interdependent decisions to be made from a large number of components, where traditional solvers are limited by having to cater for the generic problem.

A number of avenues are open for further work. In particular learning how to automatically create high quality solvers quickly is a major open problem. This is essentially an instance of the Algorithm Selection Problem [41]. A lot of research has investigated ways of tackling this problem, but veritable challenges remain. A prime example for new challenges in Algorithm Selection are the issues related to contemporary machine architectures with a large number of computing elements with diverse capabilities (e.g. multiple CPU and GPU cores in modern laptops). Research to date has largely focussed on using a single processor, with some research into parallelisation on homogeneous hardware. Being able to run several algorithms at once has a significant impact on how algorithms should be selected. In particular, constraints on the type of algorithms that be run at the same time, for example because only one of them can use the GPU, as well as collaboration between the algorithms pose promising directions for research.

Existing monolithic constraint solvers can be very successful on the benchmarks that are commonly used in the research community, however QM problems are quite different to the existing benchmarks, and require new approaches (described in Sections 5.2 and 5.3). Rather than add the required functionality to a monolithic solver and then tuning the added components manually (constrained by architectural decisions made for other problem classes), we would add components to the solver synthesiser and tune automatically for QM. Automatic tuning can produce much better results than manual tuning [27], and solver synthesis gives a much larger scope for tuning than a monolithic solver [3]. Hence advances in both these fields are highly relevant to QM and would, in our opinion, accelerate QM research.

## 6 Conclusions

In this paper we describe an area of use for CP technologies that has fallen into neglect, for no apparent good reason. The temporal logic and process calculus research communities are achieving success in qualitative modelling by publishing papers, being awarded grants, and by having the fruits of their research efforts used to solve real problems in systems biology. But dynamic systems can be perfectly well described in terms of finite difference relationships that obviate the need for temporal and process components in the underlying system description language. All finite difference methods rely on discretising a function on a grid, and the discretisation can be readily expressed in terms of CP variables and values with simple arithmetic constraints: in [31] we described the standard backward-difference approximation of a derivative, using unit step-length in order to maintain integer value domains. Forward and central differences can be approximated using the same technique, as can derivatives to any required higher order. (It should be noted this approach is therefore limited to linear dynamics, and will fail to capture chaotic dynamics of a system.) Our integer domain requirement could be relaxed, since CP can solve problems over the reals using interval methods [25]. The fact that time is the dependent variable in our models is unimportant: the discretisation works for arbitrary choice of variable representation. In addition, it is our view that the CP framework is inherently more attractive than temporal and process frameworks, since the ability to formally reason about a timeline in terms of "until", "since", etc. is not needed, and, if present, makes searching for solutions harder than necessary due to well-documented problems with state-space explosion. Moreover, recent advances in temporal CP tools and techniques [10] are specifically addressing these issues.

However, current CP technology is not well enough developed to compete with (and ideally replace) the areas of computer science that have dedicated more research effort and resource to this area of study. CP research effort into qualitative modelling faltered in the early years of this century, and has not yet recovered. The specific areas identified in Section 5 are a non-exhaustive set of future research directions for the CP community that, if successful, would allow our languages and tools to be routinely used by researchers from the physical, biological and economic sciences.

**Acknowledgments.** The authors are supported by United Kingdom EPSRC grant EP/H004092/1. LK is supported by EU FP7 grants 284715 (ICON) and 288147 (ePolicy).

## References

1. Aggoun, A., Chan, D., Dufresne, P., Falvey, E., Grant, H., Harvey, W., Herold, A., Macartney, G., Meier, M., Miller, D., Mudambi, S., Novello, S., Perez, B., van Rossum, E., Schimpf, J., Shen, K., Tsahageas, P.A., de Villeneuve, D.H.: Eclipse user manual release 5.10 (2006), http://eclipse-clp.org/
2. Balasubramaniam, D., de Silva, L., Jefferson, C., Kotthoff, L., Miguel, I., Nightingale, P.: Dominion: An architecture-driven approach to generating efficient constraint solvers. In: Software Architecture (WICSA), 2011 9th Working IEEE/IFIP Conference on. pp. 228–231 (2011)
3. Balasubramaniam, D., Jefferson, C., Kotthoff, L., Miguel, I., Nightingale, P.: An automated approach to generating efficient constraint solvers. In: 34th International Conference on Software Engineering (June 2012)
4. Beldiceanu, N., Simonis, H.: A model seeker: Extracting global constraint models from positive examples. In: Milano [37], pp. 141–157
5. Bockmayr, A., Courtois, A.: Using hybrid concurrent constraint programming to model dynamic biological systems. In: 18th International Conference on Logic Programming. pp. 85–99. Springer (2002)
6. Bristol-Gould, S.K., Kreeger, P.K., Selkirk, C.G., Kilen, S.M., Mayo, K.E., Shea, L.D., Woodruff, T.K.: Fate of the initial follicle pool: empirical and mathematical evidence supporting its sufficiency for adult fertility. Developmental biology 298(1), 149–54 (Oct 2006)
7. Calder, M., Hillston, J.: Process algebra modelling styles for biomolecular processes. In: Priami, C., Back, R.J., Petre, I. (eds.) Transactions on Computational Systems Biology XI, pp. 1–25. Springer-Verlag, Berlin, Heidelberg (2009)
8. Calzone, L., Chabrier-Rivier, N., Fages, F., Soliman, S.: Machine learning biochemical networks from temporal logic properties. T. Comp. Sys. Biology pp. 68–94 (2006)
9. Calzone, L., Fages, F., Soliman, S.: BIOCHAM: an environment for modeling biological systems and formalizing experimental knowledge. Bioinformatics (Oxford, England) 22(14), 1805–7 (Jul 2006)
10. Cimatti, A., Micheli, A., Roveri, M.: Solving temporal problems using smt: Strong controllability. In: Milano [37], pp. 248–264
11. Ciocchetta, F., Hillston, J.: Bio-pepa: A framework for the modelling and analysis of biological systems. Theor. Comput. Sci. 410(33-34), 3065–3084 (Aug 2009)
12. Clancy, D.: Qualitative simulation as a temporally-extended constraint satisfaction problem. Proc. AAAI 98 (1998)
13. Degasperi, A., Calder, M.: On the formalisation of gradient diffusion models of biological systems. In: Proc. 8th Workshop on Process Algebra and Stochastically Timed Activities. pp. 139–144 (2009)
14. Distler, A., Kelsey, T., Kotthoff, L., Jefferson, C.: The semigroups of order 10. In: CP. Lecture Notes in Computer Science, vol. 7514, pp. 883–899. Springer (2012)
15. Distler, A., Kelsey, T.: The monoids of orders eight, nine & ten. Annals of Mathematics and Artificial Intelligence 56(1), 3–21 (Jul 2009)
16. Distler, A., Kelsey, T.: The semigroups of order 9 and their automorphism groups. Semigroup Forum pp. 1–20 (2013)
17. Escrig, M.T., Cabedo, L.M., Pacheco, J., Toledo, F.: Several Models on Qualitative Motion as instances of the CSP. Revista Iberoamericana de Inteligencia Artificial 6(17), 55–71 (2002)

18. Faddy, M.J., Gosden, R.G.: A mathematical model of follicle dynamics in the human ovary. Human reproduction (Oxford, England) 10(4), 770–5 (Apr 1995)
19. Fleming, R., Kelsey, T.W., Anderson, R.A., Wallace, W.H., Nelson, S.M.: Interpreting human follicular recruitment and antimullerian hormone concentrations throughout life. Fertil. Steril. 98(5), 1097–1102 (Nov 2012)
20. Frisch, A.M., Harvey, W., Jefferson, C., Martínez-Hernández, B., Miguel, I.: Essence: A constraint language for specifying combinatorial problems. Constraints 13(3), 268–306 (Jun 2008)
21. http://www.gecode.org/
22. Gent, I.P., Jefferson, C., Miguel, I.: Minion: A fast scalable constraint solver. In: Brewka, G., Coradeschi, S., Perini, A., Traverso, P. (eds.) The European Conference on Artificial Intelligence 2006 (ECAI 06). pp. 98–102. IOS Press (2006)
23. Gent, I.P., Jefferson, C.A., Miguel, I.: MINION: A fast scalable constraint solver. In: Proceedings of the Seventeenth European Conference on Artificial Intelligence. pp. 98–102 (2006)
24. Haefner, J.: Modeling Biological Systems. Springer-Verlag, New York (2005)
25. Hentenryck, P.V., Hentenryck, P.V., Michel, L., Michel, L.: Newton: Constraint programming over nonlinear real constraints. In: Science of Computer Programming. pp. 1–2. Numerica, MIT Press (1997)
26. Hoda, S., van Hoeve, W.J., Hooker, J.N.: A systematic approach to MDD-based constraint programming. In: Cohen, D. (ed.) CP. Lecture Notes in Computer Science, vol. 6308, pp. 266–280. Springer (2010)
27. Hutter, F., Hoos, H.H., Leyton-Brown, K., Stützle, T.: Paramils: An automatic algorithm configuration framework. Journal of Artificial Intelligence Research 36(1), 267306 (2009)
28. http://www.ilog.com/products/cp/
29. Jefferson, C., Moore, N., Nightingale, P., Petrie, K.E.: Implementing logical connectives in constraint programming. Artificial Intelligence 174, 1407–1429 (2010)
30. Kelsey, T.W., Anderson, R.A., Wright, P., Nelson, S.M., Wallace, W.H.B.: Data-driven assessment of the human ovarian reserve. Molecular human reproduction 18(2), 79–87 (Sep 2011)
31. Kelsey, T., Linton, S.: Qualitative models of cell dynamics as constraint satisfaction problems. In: Backhoven, R., Will, S. (eds.) Proc. of the Workshop on Constraint Based Methods for Bioinformatics (WCB12). pp. 16–22 (2012)
32. Ko, K.I.: On the computational complexity of ordinary differential equations. Inf. Control 58(1-3), 157–194 (Jul 1984)
33. Kuipers, B.: Reasoning with qualitative models. Artificial Intelligence 59, 125–132 (1993)
34. Laburthe, F.: Choco: a constraint programming kernel for solving combinatorial optimization problems, http://choco.sourceforge.net/
35. Menzies, T., Compton, P.: Applications of abduction: hypothesis testing of neuroendocrinological qualitative compartmental models. Artificial intelligence in medicine 10(2), 145–75 (Jun 1997)
36. Menzies, T., Compton, P., Feldman, B., Toth, T.: Qualitative compartmental modelling. AAAI Technical Report SS-92-02 (1992)
37. Milano, M. (ed.): Principles and Practice of Constraint Programming - 18th International Conference, CP 2012, Québec City, QC, Canada, October 8-12, 2012. Proceedings, Lecture Notes in Computer Science, vol. 7514. Springer (2012)
38. Nethercote, N., Stuckey, P.J., Becket, R., Brand, S., Duck, G.J., Tack, G.: Minizinc: Towards a standard cp modelling language. In: Proceedings of 13th International

Conference on Principles and Practice of Constraint Programming. pp. 529–543 (2007)

39. Radke-Sharpe, N., White, K.: The role of qualitative knowledge in the formulation of compartmental models. IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews) 28(2), 272–275 (May 1998)
40. Rendl, A., Miguel, I., Gent, I.P., Jefferson, C.: Automatically enhancing constraint model instances during tailoring. In: Proceedings of Eighth Symposium on Abstraction, Reformulation, and Approximation (SARA) (2009)
41. Rice, J.R.: The algorithm selection problem. Advances in Computers 15, 65–118 (1976)
42. Rizk, A., Batt, G., Fages, F., Soliman, S.: Continuous valuations of temporal logic specifications with applications to parameter optimization and robustness measures. Theor. Comput. Sci. 412(26), 2827–2839 (2011)
43. Robertson, D., Bundy, A., Meutzelfeldt, R., Haggith, M., Uschold, M.: Eco-Logic: Logic-Based Approaches to Ecological Modelling. MIT Press (1991)
44. http://savilerow.cs.st-andrews.ac.uk/
45. Tsoukias, N.M.: Nitric oxide bioavailability in the microcirculation: insights from mathematical models. Microcirculation (New York, N.Y. : 1994) 15(8), 813–34 (Nov 2008)
46. Wallace, W.H.B., Kelsey, T.W.: Human ovarian reserve from conception to the menopause. PloS one 5(1), e8772 (Jan 2010)