TOWARDS A COMPUTATIONAL THEORY OF GSS: A CASE FOR DOMAIN-SPECIFIC LANGUAGES

PATRIK JANSSON, CEZAR IONESCU, AND NICOLA BOTTA

1. Between narratives and code

It is a trivial observation that the emerging science of global systems relies essentially on computer models. All science nowadays relies on computer models (for an interesting analysis of just how much, see "Simulation and its Discontents" [Tur09]). However, a distinguishing feature is that in GSS these models are usually the only kind we have. In physics, for example, computer models are usually secondary to mathematical ones. It is the latter which embody the theories and serve as object of discussion between scientists: computer code rarely makes its way in scientific publications. The computer model is judged to be correct if it faithfully implements the mathematical model. While it might occasionally be difficult to formulate exactly what counts as "faithful", we have, at least in principle, a "golden standard" for correctness.

Consider a multi-agent model which is supposed to inform politicians about the effect of green building subsidies on the evolution of unemployment. What is the "golden standard" against which to judge the correctness of such a model? Usually, such models are accompanied by a narrative, an informal description of the ideas behind their development, but a narrative is too blunt an instrument to help us decide whether a computer implementation of it is correct or not, whether the results of a simulation are trustworthy or flawed by programming errors. In fact, in absence of an external criterion for the correctness of the model, it is not clear that we can talk about programming errors at all.

Developing the traditional kind of mathematical models for the complex heterogeneous socio-ecological systems involved in GSS has not been very successful so far. GSS is massively inter-disciplinary, and even when the various components of a model have mathematical representations within the disciplines involved, there is no clear way of coming up with a mathematical model of the interactions between the components. The almost automatic way that scientists solve this problem is by implementing the components and their interactions in code. This provides a formal description of the model and allows exploring its consequences by simulations, and *only* by simulations, because the description is too low level to allow us to reason about it the way we do with traditional mathematical models.

This unsatisfactory state of affairs can be summed up as follows: either we limit ourselves to informal narratives, or we use simulations of computer models which we do not understand and whose correctness we cannot guarantee. Either way, the scientific status of GSS is in question.

The solution appears to be the creation of an intermediate, mathematical layer between narratives and simulations, similar to that which exists in other established sciences. This mathematical layer cannot just be, say, the theory of partial differential equations underlying the physics of climate, or the functional analysis accounting for the general equilibrium models of economics. The formal language of GSS is computer code, therefore the mathematical layer has to be part of the mathematics of general programs, that is, computer science.

To put it in a somewhat pointed fashion: **computer science should play for GSS the same role that mathematics plays for physics**.

More concretely, we need to start by writing *specifications* for the kind of models used in GSS, which will involve choosing, adapting, and extending one or more of the formal languages for specifying and reasoning about programs. At the moment, the main candidate for such a formal language is that of constructive type theory, due to its ability to express both (functional) programs and classical mathematical results. But, just as physics has influenced the development of mathematics, providing fruitful problems and intuitions, we expect that GSS will also influence the kind of formal languages and the results of computer science.

Specifications will allow checking the correctness of implementations, but in the long run, we can do better: we can implement high-level domain-specific languages (DSLs), such that the distance between specification and implementation will be as short as possible. Ideally, the specification should be expressive enough that its compatibility with the narrative which motivated it can be seen "by inspection" and that it can serve for the communication of scientific ideas, and at the same time it should be part of the programming language used to implement it, so that the implementation is correct by construction.

Such DSLs are not "just" important for the correctness of models, they are, in fact, essential to the policy-science interface.

In the sequel, we shall illustrate this with three examples: project GRACeFUL (under grant preparation for the FETPROACT-1-2014 call on GSS); formalizing *avoidability* in the context of climate change; and developing a *theory of policy advice* in GSS.

Date: 2014-10-03 (for the 3rd Open GSS Conference).

2. Project GRACeFUL and DSLs for the foundations of GSS

GRACeFUL aims to build a framework for rapid assessment tools, in order to support political decision-making in the typical context of GSS, involving multiple stakeholders faced with inter-disciplinary, global challenges. The core idea is to combine the methodology of *group model building*, used to help stakeholders develop a systemic description of the situation and of the goals to be achieved, with the technology of *constraint programming*, which will "package" the available scientific models and data and use them in an efficient manner to propose plans for achieving those goals.

The key difficulty is that the stakeholders describe problems in high-level terms such as *hazard*, *risk*, *resilience*, *vulnerability*, *protection measures*, *financial contract*, *adaptation* and *mitigation*, *desirability* versus *undesirability*, etc., which are rarely the terms which describe the available data or scientific models.

In GRACeFUL, the translation from the language of stakeholders to that of data and models is carried out by a DSL. The main components of the DSL will be the high-level terms used (within graphical tools) by the stakeholders in group model building sessions, which the DSL will"compile down" to the level of constraint programming.

Translations generate meaning, and this one is no exception: high level concepts are interpreted in the rigorous (but lower-level) language of the constraint programming layer. The mechanism of rapid assessment tools envisioned by GRACeFUL will also help in refining this translation. In the initial stages, the rapid problem-solution cycle will be used to ensure that the solutions are not polluted by errors of interpretation of the high-level terms.

However, once this validation step is passed, we will be in possession of a *formalization* of important high-level notions of GSS such as those cited above. The use of DSLs to give precise meanings to and to allow reasoning about complex concepts is not new. Perhaps the most successful example so far is that of DSLs for financial contracts initially proposed by Peyton-Jones, Eber, and Seward [PJES00] which have been adopted and extended by several companies in the financial sector. Closer to GSS, the formalization of vulnerability developed at PIK and extended in cooperation with Chalmers [IKH+05, IO06, Ion09, IKH+09, IJ13b, IJ13a] has unified various approaches to vulnerability assessment in the climate change and development studies communities.

In the next section, we explore another related effort currently underway at PIK and Chalmers, aiming to formalize the notion of *avoidability* in the context of climate change. In the long run, we hope that such systematic exploitation of the *theory-building* aspects of DSLs will lead to solid foundations for GSS, including a theory of *policy advice*, the topic of Section 4.

3. Formalizing high-level concepts: the example of "avoidability"

In most situations related to climate policy, decision making has to take place under fairly weak assumptions. In international environmental agreements (IEA), for instance, decision makers are typically faced with sequential decision processes over a certain number of steps [FvID03, Hel03, Hei13]. Each step represents a finite period of time, for instance 10 years. The kind of information available to a decision maker in a single decision step and the kind of options available in that state can often be described fairly rigorously. For example, the information available for decision making could consist of some measure of greenhouse gas (GHG) concentration and gross domestic product. The options could be perhaps greenhouse gas abatements and investments.

But even if the decision processes can often be described in terms of fairly simple terms, the "dynamics" of the system underlying the decision process is usually affected by various kinds of uncertainty, whether releated to the complexity of the system or the behaviour of other decision makers. Therefore, the decision problems are often genuinely non-deterministic.

Compounding this non-determinism is the fact that the approach usually taken for solving decison problems relies on notions of "optimality" ([Bel57], for a DSL treatment see [BIB13, BJI+14b]) requiring decision makers to estimate the value of policies in terms of payoffs associated with the decision process. Besides the "technical" difficulties of such estimates, ethical concerns have also been raised against this approach.

Such criticisms have lead a number of authors to argue that, for policy advice on climate impacts, it would be more sensible to shift the focus from the attempt to maximize questionable costs-benefits estimates towards policies that provably avoid future possible states which are known or thought to be harmful.

This is the approach exemplified in [RBH07] but also in more positive formulations such as the "tolerable window" approach proposed in [Sch98]. Such notions of "avoidability" are also at the root of the definitions of *mitigation* and *adaptation* which are at the core of IPCC's Working Group III research.

But what does it precisely mean for possible future states to be avoidable?

Since the notion of avoidability seems to play a decisive role in IEAs and climate impact discussion, avoidability has to be one of the key notions of a DSL for supporting decisions in these areas. It is to be expected that a number of notions of avoidability needs to be formulated[1].

From a GSS perspective, a theory of avoidability and, in particular, generic decision procedures for assessing avoidability could be easily instantiated for

[1] In much the same way as different notions of equilibrium turns out to be necessary to formulate problems in economics and game theory.

other domains than climate change. In financial markets and after two decades of Financial Stability Reviews, for instance, unambiguous notions (let apart operational tests) of stability are still elusive [Goo04].

Here GSS could play a decisive role in pioneering more reliable (and more accountable) approaches towards stability assessment. Sound rules for financial markets have to cope with systems which are, as in the case of IEAs, imperfectly known (and, therefore, genuinely non-deterministic or, at most, stochastic) and for which assuming the availability of credible estimates of the value of policies would be questionable. Here too, it could be more realistic to focus on policies that focus on the avoidance of potentially harmful future states.

4. Towards a theory of policy advice

Consider again the sequential decision problems related to international environmental agreements in the previous section. There, decision makers needed to select, at each time step, between a few different rates of abatement of CO_2 emissions or, perhaps, some emission cap.

It might seem that, in this case, the ideal advice a decision maker could receive would be a list of, e.g., "optimal" rates of abatement, one for each step. However, this first impression is misleading. As we mentioned, the decision process in the case of IEAs (and in most realistic GSS related situtations) is genuinely non-deterministic.

Therefore, the advisor cannot know exactly what will happen after the first decision step, and the second element of the list could very well be non-applicable (not to mention sub-optimal). This is especially easy to see in the case where decisions are framed in terms of investments: the investment is constrained by the actual capital and, possibly, by some measure of loan availability; if the level of this capital (loan) is unknown, the "optimal" investment cannot be known, either.

In such situations, a list of future decisions to be made is inadequate. What is needed is, instead, a set of rules which take into account the possible unfoldings of the process.

Decision making that takes into account the facts as they unfold during a particular realization of the decision process is not only much more flexible than decision making based on some fixed control plan. It is also, in general, more effective. It makes the difference between planned economy and adaptive allocation of resources and allows decision making to take advantage of the information that becomes available as time advances.

What seems to be an obvious observation is often neglected in policy advice. In fact, most of the tools currently used for integrated assessment studies, impact research and IEAs are based on deterministic systems and "advice" refers to static decision plans. But, if policy advice cannot be about recommending static decision plans and delivering scenarios according to such plans what should then be the content of policy advice? The answer is both obvious and compelling:

In control theory, sets of rules which take into account the possible unfoldings of a decision process are called policies and we argue that the main content of policy advice – what advisors are to provide to decision maker – are policies.

These informal notions of policy are conceptually correct but, as it turns, too simplistic. If policy advice has to be accountable, these notions have to be refined and carefully formalized. We cannot expand here but we have outlined a theory of decision making under uncertainty in [BJI+14b]. In [BJI14a], we will present the preliminaries of a theory of policy advice and avoidability.

5. Beyond foundations

In this document, we have focused mainly on the foundational role of computer science in GSS. This does not exhaust by any means the role of ICT in GSS. For example, we have not touched so far on the empirical side of GSS. In common with almost every other scientific discipline, in GSS we also have the problem of the massive amounts of data to be collected and analyzed, but here we also face the *heterogeneity* of the data. As an integrative science, GSS must develop the concepts and techniques for dealing with data coming from a wide spectrum of scientific disciplines. Efficiently dealing with this kind of data must lead to establishing correlations between disciplines, developing indicators to summarize otherwise unmanageable quantities of information, suggesting new concepts for the theory of GSS. These concepts should also allow us to deal better with the new sources of data available from social networking tools like Facebook and Twitter.

These networking tools have been developed in an ad-hoc fashion, raising questions about their security, reliability, use and misuse. We should expect a science of global systems to enable us to understand these issues better, and design a next generation of communication tools. Also in view of e-Governance issues, we need scientific theories that can tell us about how to increase the reliability of information, how to counteract disinformation, how governments can encourage a broad democratic participation, enabling change and building trust.

Acknowledments: The work presented in this paper heavily relies on free software, among others on hugs, GHC, vi, the GCC compiler, Emacs, LaTeX and on the FreeBSD and Debian GNU/Linux operating systems. It is our pleasure to thank all developers of these excellent products.

References

[Bel57] Richard Bellman. *Dynamic Programming.* Princeton University Press, 1957.

[BIB13] Nicola Botta, Cezar Ionescu, and Edwin Brady. Sequential decision problems, dependently-typed solutions. In *Proceedings of the Conferences on Intelligent Computer Mathematics (CICM 2013), "Programming Languages for Mechanized Mathematics Systems Workshop (PLMMS)"*, volume 1010 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2013.

[BJI14a] Nicola Botta, Patrik Jansson, and Cezar Ionescu. Towards a computational theory of policy advice and avoidability. In preparation, 2014.

[BJI+14b] Nicola Botta, Patrik Jansson, Cezar Ionescu, David R. David R. Christiansen, and Edwin Brady. Sequential decision problems, dependent types and generic solutions. Submitted to LMCS (Logical Methods in Computer Science), August 2014, 2014.

[FvID03] Michael Finus, Ekko van Ierland, and Rob Dellink. Stability of climate coalitions in a cartel formation game. FEEM Working Paper No. 61.2003, 2003.

[Goo04] Charles Goodhart. Some New Directions for Financial Stability? Per Jacobsson lecture, Zurich, 27 June 2004, 2004.

[Hei13] Jobst Heitzig. Bottom-up strategic linking of carbon markets: Which climate coalitions would farsighted players form? Working Papers 2013.48, Fondazione Eni Enrico Mattei, 2013.

[Hel03] Carsten Helm. International emissions trading with endogenous allowance choices. *Journal of Public Economics*, 87:2737–2747, 2003.

[IJ13a] Cezar Ionescu and Patrik Jansson. Dependently-typed programming in scientific computing: Examples from economic modelling. In Ralf Hinze, editor, *24th Symposium on Implementation and Application of Functional Languages (IFL 2012)*, volume 8241 of *LNCS*, pages 140–156. Springer, 2013.

[IJ13b] Cezar Ionescu and Patrik Jansson. Testing versus proving in climate impact research. In *Proc. TYPES 2011*, volume 19 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 41–54, Dagstuhl, Germany, 2013. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[IKH+05] C. Ionescu, R. J. T. Klein, J. Hinkel, K. S. Kavi Kumar, and R. Klein. Towards a formal framework of vulnerability to climate change. *NeWater Working Paper 2 and FAVAIA Working Paper 1*, 2005.

[IKH+09] C. Ionescu, R. J. T. Klein, J. Hinkel, K. S. Kavi Kumar, and R. Klein. Towards a formal framework of vulnerability to climate change. *Environmental Modelling and Assessment*, 14(1):1–16, 2009.

[IO06] Cezar Ionescu and Karen O'Brien. A mathematical interpretation of the distinction between "outcome" and "contextual" vulnerability. Personal Communication, 2006.

[Ion09] Cezar Ionescu. *Vulnerability Modelling and Monadic Dynamical Systems*. PhD thesis, Freie Universität Berlin, 2009.

[PJES00] Simon Peyton Jones, Jean-Marc Eber, and Julian Seward. Composing contracts: An adventure in financial engineering (functional pearl). In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming*, ICFP '00, pages 280–292, New York, NY, USA, 2000. ACM.

[RBH07] Peter Raven, Rosina Bierbaum, and John Holdren. Confronting climate change: Avoiding the unmanageable and managing the unavoidable. UN-Sigma Xi Climate Change Report, 2007.

[Sch98] Hans Joachim Schellnhuber. Discourse: Earth system analysis - the scope of the challenge. In Hans Joachim Schellnhuber and Volker Wenzel, editors, *Earth System Analysis: Integrating Science for Sustainability*, pages 3–195. Springer, Berlin/Heidelberg, 1998.

[Tur09] Sherry Turkle. *Simulation and its discontents*. Simplicity: design, technology, business, life. The MIT Press, 2009.