Stochastic Behavior and Explicit Discrete Time in Concurrent Constraint Programming[*]

Jesus Aranda[1,2], Jorge A. Pérez[3], Camilo Rueda[4,5], and Frank D. Valencia[1]

[1] INRIA/CNRS and LIX, Ecole Polytechnique, France
[2] Escuela de Ingeniería de Sistemas y Computación, Universidad del Valle, Colombia
[3] Dept. of Computer Science, University of Bologna, Italy
[4] Dept. of Science and Engineering of Computing, Universidad Javeriana - Cali, Colombia
[5] IRCAM, Paris, France

Abstract. We address the inclusion of stochastic information into an *explicitly timed* concurrent constraint process language. An operational semantics is proposed as a preliminary result. Our approach finds applications in biology, among other areas.

Motivation. The study of *quantitative information* within languages for concurrency has recently gained a lot of momentum. In many applications, quantitative information becomes crucial when refining models with empirical data, and is of the essence for verification purposes. Two main models of quantitative information can be singled out from the vast literature on the subject. Given a computation that can perform different, competing actions, a *probabilistic* model provides a probability distribution over such actions. In contrast, a *stochastic* model relates each action to a random variable which determines its *duration*: given a set of competing actions, the fastest action (i.e. the one with the shortest duration) is executed. Consequently, notions not considered in a probabilistic model (e.g. speed) are fundamental in a stochastic setting. Not surprisingly, areas in which time is essential (e.g. systems biology, performance modeling) have found in languages featuring stochastic information adequate frameworks for analysis.

Concurrent constraint programming (CCP) [8] is a declarative model for concurrency with strong ties to logic. In CCP, systems are described by pieces of partial information called *constraints*. Processes interact in a shared *store*; they either add new constraints or synchronize on the already available information. Timed concurrent constraint programming (tcc) [9] is a declarative framework for reactive systems. In tcc, time is explicitly represented as discrete time units in which computation takes place; tcc provides constructs to control process execution along such units. In the light of stochastic models for quantitative information, the explicit time in tcc poses a legitimate question, that of determining to what extent the notions of stochastic duration and of discrete time unit can be harmoniously conciliated within a CCP-based framework. The question is relevant because it can give clues on clean semantic foundations for

---

quantitative information in CCP, which in turn, should contribute to the development of more effective reasoning techniques over reactive systems in many emerging applications. In this paper, we outline preliminary results on an operational semantics for a tcc language with explicit stochastic durations.

More into details, the proposed semantics aims at an explicit account of stochastically derived events using the description power of timed CCP calculi. This is a feature that in other CCP calculi (e.g. [2]) is handled at best implicitly. We define stochastic events in terms of the time units provided by the calculus: this provides great flexibility for modeling and, as mentioned before, it allows for a clean semantics. Most importantly, by considering stochastic information and adhering to explicit discrete time, it is possible to reason about processes using *quantitative* logics (both discrete and continuous), while retaining the simplicity of calculi such as ntcc [6] for deriving *qualitative* reasoning techniques (such as denotational semantics and proof systems). We consider existing qualitative reasoning techniques have a great potential for guiding/complementing the use of (usually costly) quantitative ones. In the biological context, this approach has been proposed in [4], where qualitative techniques are enough to analyze the influence graph of a reaction model as, to some extent, it is independent from kinetic parameters and kinetic expressions.

This work is part of a larger research programme aimed at developing robust CCP-based techniques for analyzing complex applications and systems in computer music, security and biology. As such, it is our objective to formalize stochastic information in tcc in such a way that resulting languages and techniques (i) remain generic enough so to fit well in the target applications, and (ii) be amenable to efficient implementations, in the form of e.g. simulators and model-checkers.

Description. We consider a variant of tcc in which certain processes are annotated with a function λ, which represents the stochastic information in the language (see below). Annotated processes are tell, when and unless. Unless considers functions as we see it as a counterpart of when processes. A careful definition of unless in the stochastic context, however, is yet to be completely determined. We do not discard that different applications (biological systems, computer music) need different unless definitions.

$$P, Q ::= \ \mathsf{tell}_\lambda(c) \ \big| \ \mathsf{when}\, c\, \mathsf{do}\, (P, \lambda) \ \big| \ P \parallel Q \ \big| \ \mathsf{local}\, x\, \mathsf{in}\, P \ \big| \ !P \ \big| \ \mathsf{next}\, (P) \ \big| \ \mathsf{unless}_\lambda\, c\, \mathsf{next}\, (P)$$

Operational Semantics. We use the same notion of discrete time as in ntcc and tcc. That is, there are discrete time units of uniform size, each of them having its own constraint store. At each time unit, some stimuli are received from the environment; the process then executes with such stimuli as input. At the end of the time unit, some output is produced in the form of responses to the environment, and a residual process to be executed in the next time unit is scheduled. Remarkably, information does not automatically transfer from one time unit to the following.

The operational semantics, given in Table 1, is defined over process-store configurations. We use γ, γ' to range over configurations, and assume a structural congruence $\equiv$ to identify processes with minor syntactic differences.

The semantics defines rules that carry both a *probability value* (denoted p) and a global *rate value* (denoted r). The rules decree two kinds of process execution, *immediate* (probability value equal to 1 and rate value max), and *stochastic*. In this sense, processes can be either immediate or stochastic. The idea of the semantics is to schedule immediate processes first, and then move to stochastic processes, whose execution involves a certain duration.

Rules for immediate execution resemble analogous rules in tcc and ntcc. Rule IMMTELL adds a constraint to the store as soon as possible. Rule IMMREP specifies that process $!P$ produces a copy P at the current time unit and then persists in the next time unit. There is no risk of infinite behavior within a time unit. In Rule IMMUNLESS, process P is precluded if c is entailed by the current store d. Rule IMMINT allows for compositional extension.

Rules for stochastic executions consider the aforementioned function λ. Using the current store as parameter, λ describes how the global rate of the whole process varies. We use $\delta^m(P)$ to denote a *delay process P* with duration m: P will be executed at the m-th time unit from the current one. Given probability and rate values for a process, function Δ determines its duration. The duration can be thus seen as an exponentially distributed random variable that depends on a probability and a rate.

Rule STOTELL defines stochastic tell actions. Rule STOCHOICE defines a choice over a number of guarded processes. Only those *enabled* processes, i.e., those whose guards entail from the current store, are considered. Rule STOINT defines the simultaneous occurrence of stochastic actions. As usual, the probability value is calculated assuming independence of the actions. Notice that the current store is not affected by stochastic actions; their influence is only noticeable in the following time units. Rules STOUNLESS and STOREP define unless and stochastic replicated actions, resp. Rule NEXT extends stochastic actions to next processes. In rule LOCAL, local in P behaves like P, except that all the information on x produced by P can only be seen by P and the information on x produced by other processes cannot be seen by P. Finally, rule STRCONG says that structurally congruent configurations have the same actions.

The described rules define behavior within a time unit; internal behavior takes place until reaching a configuration where no further computation is possible (*quiescence*, denoted by $\nrightarrow$). We need to define the *residual process* to be executed in the following time unit. We start by conjecturing that quiescent configurations have a "standard" form:

Conjecture 1 *If $\gamma \nrightarrow$ then it holds that*

$$\gamma \equiv \langle \prod_{j \in J} \text{next } (P_j) \parallel \prod_{k \in K} \text{unless } c_k \text{ next } (Q_k) \parallel \prod_{i \in I} \delta^{m_i}(P_i), \, d \rangle.$$

We use A to denote the set of (delayed) processes in the right-hand side of the "normal" quiescent configuration.

Definition 1 (Future function) *Given a quiescent configuration γ (with the form given in Prop. 1), its residual process is given by function F:*

$$F(\gamma) = \prod_{j \in J} P_j \parallel \prod_{k \in K} Q_k \parallel F'(A)$$

$$\text{IMMTELL} \quad \frac{}{\langle \text{tell}_1(d), c\rangle \longrightarrow_{1,\max} \langle \text{skip}, c \wedge d\rangle} \qquad \text{IMMREP} \quad \frac{\langle P, c\rangle \longrightarrow_{1,\max} \langle P', c'\rangle}{\langle \,!\,P, c\rangle \longrightarrow_{1,\max} \langle P \parallel \text{next}\,(!\,P), c'\rangle}$$

$$\text{IMMUNLESS} \quad \frac{}{\langle \text{unless}\, c\, \text{next}\,(P), d\rangle \longrightarrow_{1,\max} \langle \text{skip}, d\rangle} \,\text{if}\, d \models c \quad \text{IMMINT} \quad \frac{\langle P, c\rangle \longrightarrow_{1,\max} \langle P', c'\rangle}{\langle P \parallel Q, c\rangle \longrightarrow_{1,\max} \langle P' \parallel Q, c'\rangle}$$

$$\text{STOTELL} \quad \frac{}{\langle \text{tell}_\lambda(d), c\rangle \longrightarrow_{1,\lambda(c)} \langle \delta^m(\text{tell}(d)), c\rangle} \,\text{with}\, m = \Delta(1, \lambda(c))$$

$$\text{STOCHOICE} \quad \frac{}{\langle \sum_{i \in I} \text{when}\, c_i\, \text{do}\,(P_i, \lambda_i), c\rangle \longrightarrow_{p,r} \langle \delta^m(P_j), c\rangle} \,\text{if}\, c \models c_j$$
$$\text{with}\, r = \sum_{i \in \{j \,|\, c \models c_j\}} \lambda_i(c); \; p = \lambda_j(c)/r; \; m = \Delta(p, r).$$

$$\text{STOINT} \quad \frac{\langle P, c\rangle \longrightarrow_{p_1, r_1} \langle P', c\rangle \quad \langle Q, c\rangle \longrightarrow_{p_2, r_2} \langle Q', c\rangle}{\langle P \parallel Q, c\rangle \longrightarrow_{p', r'} \langle P' \parallel Q', c\rangle} \,\text{with}\, p' = p_1 \times p_2; \; r' = r_1 + r_2.$$

$$\text{STOUNLESS} \quad \frac{}{\langle \text{unless}_\lambda\, c\, \text{next}\,(P), d\rangle \longrightarrow_{p,r} \langle \delta^m(\text{unless}\, c\, \text{next}\,(P)), d\rangle} \,\text{with}\, m = \Delta(p, r).$$

$$\text{STOREP} \quad \frac{\langle P, c\rangle \longrightarrow_{p,r} \langle \delta^m(P), c'\rangle}{\langle \,!\,P, c\rangle \longrightarrow_{p,r} \langle \delta^m(P) \parallel \text{next}\,(!\,P), c'\rangle} \qquad \text{NEXT} \quad \frac{\langle P, c\rangle \longrightarrow_{p,r} \langle P', c\rangle}{\langle P \parallel \text{next}\,(Q), c\rangle \longrightarrow_{p,r} \langle P' \parallel \text{next}\,(Q), c\rangle}$$

$$\text{LOCAL} \quad \frac{\langle P, c \wedge \exists_x d\rangle_g \longrightarrow_{p,r} \langle P', c'\rangle_g}{\langle (\text{local}\, x, c)P, d\rangle \longrightarrow_{p,r} \langle (\text{local}\, x, c)P', d \wedge \exists_x c'\rangle} \qquad \text{STRCONG} \quad \frac{\gamma_1 \longrightarrow_{p,r} \gamma_2}{\gamma_1' \longrightarrow_{p,r} \gamma_2'} \,\text{if}\, \gamma_i \equiv \gamma_i' \,(i \in \{1, 2\})$$

Table 1. Operational semantics: internal transition rules.

where function F' is defined as

$$F'(\delta^{m_1}(P_1) \parallel \ldots \parallel \delta^{m_n}(P_n)) = G(\delta^{m_1}(P_1)) \parallel \ldots \parallel G(\delta^{m_n}(P_n))$$

and where G is defined as

$$G(\delta^m(P)) = \begin{cases} \delta^{m-1}(P) & \text{if}\, m > 1 \\ P & \text{if}\, m = 1. \end{cases}$$

Unlike other languages like the stochastic π-calculus [7] or sCCP [2], it is worth noticing that in our semantics stochastic actions can evolve simultaneously; they are not simulated by choosing one of them to be performed first. This way, e.g., $\text{tell}_{\lambda_1}(c_1) \parallel \text{tell}_{\lambda_2}(c_2)$ evolves into $\delta^{m_1}(\text{tell}(c_1)) \parallel \delta^{m_2}(\text{tell}(c_2))$ and in the next unit time, the configuration is $\delta^{m_1-1}(\text{tell}(c_1)) \parallel \delta^{m_2-1}(\text{tell}(c_2))$ (assuming $m_1, m_2 > 0$). This allows to naturally represent the evolution of different components in parallel.

Discussion. Since variables in tcc are logic (i.e. they can be defined at most once in each time unit), a potential source of inconsistencies is the simultaneous execution of several stochastic actions involving the same variables. This could represent a limitation in modeling. Consider for instance the kind of systems in which it is required to deal with quantities of elements of a certain type (as in biological reactions). In such systems, variables could be part of several actions, which would represent the changes over the elements in consideration. An inconsistency caused by two actions simultaneously altering the value of the same variable is clearly an undesirable feature. Therefore, there is the need for enhancing the semantics with a mechanism that imposes some kind of

order over those actions related with potential inconsistencies. This would also presuppose modifications over rules calculating duration of stochastic actions, as concurrent actions would be simulated in a specific order. The formal definition of such a consistency mechanism is part of ongoing work.

Applications in Biology. We think our language and semantics have applications in the biological domain. This is supported by the fact that CCP-based calculi have shown to be convenient for modelling, simulating and verifying several kinds of biological systems [5,1,2]. In [2], stochastic concurrent constraint programming (sCCP) is used to model biochemical reactions and gene regulatory networks. Functional rates in sCCP give considerable flexibility to formulate reactions. However, sCCP does not include an explicit notion of time and does not exploit the logic nature of CCP for verification. Also, sCCP lacks a means of expressing absence of information, which has proven most useful in the biological context [1]. Both a proof system for properties and a means of representing absence of information are available in the explicitly timed ccp language `ntcc`. In fact, `ntcc` was used in [5,1] to model different biological systems using two kinds of partial information: *behavioral* (e.g. the unknown relative speeds on which a system evolves) and *quantitative* (e.g. the set of possible values that a variable can take). It must be noticed that `ntcc` does not allow for stochastic or probabilistic information.

Based on the above, we think the extension to `tcc` here proposed could serve several purposes in the biological context. The most immediate use is the definition of enhanced models of the systems already modeled in `ntcc` (the Sodium-Potassium pump, regulation and mutation processes in genetic regulatory networks). Also, although it is not evident that every sCCP process can be translated into our language (the tell operator in sCCP has continuation), we are confident we can model most of the biological systems described there. We also plan to analyse the model in [3], which describes the cycle of Rho GTP-binding proteins in the context of phagocytosis.

References

1. A. Arbeláez, J. Gutiérrez, and J. A. Pérez. Timed Concurrent Constraint Programming in Systems Biology. *Newsletter of the ALP*, 19(4), November 2006.
2. L. Bortolussi. *Constraint-based approaches to stochastic dynamics of biological systems*. PhD thesis, University of Udine, 2007.
3. L. Cardelli, P. Gardner, and O. Kahramanogullari. A process model of rho gtp-binding proteins in the context of phagocytosis. *Electr. Notes Theor. Comput. Sci.*, 194:87–102, 2006.
4. F. Fages and S. Soliman. Formal cell biology in biocham. In *Formal Methods for Computational Systems Biology*, volume 5016 of *LNCS*, pages 54–80. Springer, 2008.
5. J. Gutiérrez, J. A. Pérez, C. Rueda, and F. Valencia. Timed Concurrent Constraint Programming for Analysing Biological Systems. *Electr. Notes Theor. Comp. Sci.*, 171(2):117–137, 2007.
6. M. Nielsen, C. Palamidessi, and F. Valencia. Temporal concurrent constraint programming: Denotation, logic and applications. *Nordic Journal of Computing*, 9(2):145–188, 2002.
7. C. Priami. Stochastic pi-calculus. *Comput. J.*, 38(7):578–589, 1995.
8. V. Saraswat. *Concurrent Constraint Programming*. The MIT Press, Cambridge, MA, 1993.
9. V. Saraswat, R. Jagadeesan, and V. Gupta. Foundations of timed concurrent constraint programming. In *Proc. of LICS*. IEEE, 1994.