Measuring the Expressiveness of Rewriting Systems through Event Structures

Part I: Event Structures

Damiano Mazza
Laboratoire d'Informatique de Paris Nord
CNRS–Université Paris 13

CONCERTO final workshop
Torino, 10 June 2010

Motivations

- Interaction nets (Lafont, 1990) are a model of deterministic computation, born as a generalization of linear logic proof nets (Girard, 1987).

- How expressive are they? They are Turing-complete. . . but this means nothing! What about parallelism?

- In addition, there are several non-deterministic variants:
  - *multiwire* (Alexiev 1999, Beffara-Maurel 2006);
  - *multiport* (Alexiev 1999, Khalil 2003, Mazza 2005);
  - *multirule* (Alexiev 1999, Ehrhard-Regnier 2006).

- How do these relate to each other? Can they model concurrency?

- We are not only interested in *what* we compute, but also *how*.

Computational dynamics: what is (discrete) time?

- As Winskel (1980), we consider the structure of time to be given by:

  causality: there is time as soon as there is a "before" and and "after", which in turn can be defined in terms of *causal relationship*;

  conflict: there must be an idea of *parallelism* or, dually, of *conflict*, telling whether two non-causally related events may happen together.

- Conflict also encompasses the notion of non-determinism.

- The (complexity of the) structure of time may be considered as a measure of expressiveness: if we describe computational models in terms of their "dynamic structures", we may try to use these structures to compare their intensional expressiveness.

Event structures

- An event structure (Winskel, 1980) is a triple $E = (|E|, \leq, \smile)$ such that:
  - $|E|$ is a set of *events*, called *web*;
  - $\leq$ is a partial order on $|E|$, called *causal order*, such that, for all $a \in |E|$, $\downarrow a = \{b \in |E| \mid b \leq a\}$ is finite;
  - $\smile$ is an anti-reflexive symmetric relation on $|E|$, called *conflict relation*, such that, for all $a, b, c \in |E|$, $a \smile b \leq c$ implies $a \smile c$.

- Let $u \subseteq |E|$. We say that u is a *configuration* iff

  causality: $a \in u$ and $b \leq a$ implies $b \in u$.
  coherence: $a, b \in u$ implies $a \frown b$;

  The set of finite configurations of E is denoted by $\mathcal{C}(E)$.
  A configuration $u \in \mathcal{C}(E)$ *enables* $a \in |E|$ if $a \notin u$ and $u \cup \{a\} \in \mathcal{C}(E)$.
  The smallest configuration *enabling* $a \in |E|$ is $\lceil a \rceil = \downarrow a \setminus \{a\}$.

Configuration posets

Let $(X, \leq)$ be a poset, and let $x, y \in X$.

- We write $x \uparrow y$ (*compatibility*) iff $\exists z \in X$ s.t. $x, y \leq z$. We say that $u \subseteq X$ is a *clique* iff $x, y \in u$ implies $x \uparrow y$. We say that X is *coherent* iff, whenever u is a clique, $\bigvee u$ exists.

- We say that y *covers* x iff $x < y$ and there is no z s.t. $x < z < y$. We say that an element of X is *prime* if it covers exactly one element, and we set $\mathfrak{p}(x) = \{a \in X \mid a \leq x,\ a \text{ prime}\}$. We say that X is *prime algebraic* iff, $\forall x \in X$, we have $\bigvee \mathfrak{p}(x) = x$.

- X is a *configuration poset* iff it is coherent, prime algebraic, and, $\forall x \in X$, $\downarrow x$ is finite.

- Example: the compact elements of a coherent dI-domain (Berry, 1979).

Configuration posets and event structures

- The following is adapted from Nielsen, Plotkin, Winskel (1981):

  Theorem 1. [Representation] *The groupoids* $\mathbf{E}_{\mathbf{Grp}}$ *of event structures and their isomorphisms and* $\mathbf{Conf}_{\mathbf{Grp}}$ *of configuration posets and their isomorphisms are equivalent.*

- More precisely, we have:
  - if E is an event structure, then $\Phi(E) = (\mathcal{C}(E), \subseteq)$ is a configuration poset;
  - if $(X, \leq)$ is a configuration poset, $\mathfrak{P}(X)$ the set of its prime elements, then $\Psi(X, \leq) = (\mathfrak{P}(X), \leq, \not\uparrow)$ is an event structure;
  - $\Psi(\Phi(E))$ is an event structure isomorphic to E;
  - $\Phi(\Psi(X, \leq))$ is a configuration poset isomorphic to $(X, \leq)$.

Transitions in event structures

- Let E, E' be event structures, and let $R \subseteq |E| \times |E'|$. If $u \in \mathcal{C}(E)$, we write $\mathrm{supp}_R(u) = u \cap \pi_1(R)$.

- Let $u, v \in \mathcal{C}(E)$ with $v = u \cup \{a\}$, $a \notin u$. We define:
  - $u \xrightarrow{a}_R v$ if $a \in \pi_1(R)$ (*computational transition*);
  - $u \longrightarrow_R v$ if $a \notin \pi_1(R)$ (*administrative transition*).

- We denote by $\Longrightarrow_R$ the reflexive-transitive closure of $\longrightarrow_R$.

- We write $u \stackrel{a}{\Longrightarrow}_R v$ iff $\exists u', v'$ s.t. $u \Longrightarrow_R u' \xrightarrow{a}_R v' \Longrightarrow_R v$.

- We do the same for E', with π_2 instead of π_1.

Bisimulations

- Let E, E' be event structures, and let $R \subseteq |E| \times |E'|$. A *$R$-bisimulation* between E and E' is a relation $\mathcal{B} \subseteq \mathcal{C}(E) \times \mathcal{P}_{\mathrm{fin}}(R) \times \mathcal{C}(E')$ such that $(\emptyset, \emptyset, \emptyset) \in \mathcal{B}$ and, whenever $(u, \phi, u') \in \mathcal{B}$, we have:

  i. ϕ is a poset isomorphism between $(\mathrm{supp}_R(u), \leq)$ and $(\mathrm{supp}_R(u'), \leq')$;
  ii. $u \xrightarrow{a}_R v$ implies $u' \overset{a'}{\Longrightarrow}_R v'$ with $(v, \phi \cup \{(a, a')\}, v') \in \mathcal{B}$;
  iii. $u \longrightarrow_R v$ implies $u' \Longrightarrow_R v'$ with $(v, \phi, v') \in \mathcal{B}$;
  iv. $u' \xrightarrow{a'}_R v'$ implies $u \overset{a}{\Longrightarrow}_R v$ with $(v, \phi \cup \{(a, a')\}, v') \in \mathcal{B}$;
  v. $u' \longrightarrow_R v'$ implies $u \Longrightarrow_R v$ with $(v, \phi, v') \in \mathcal{B}$.

- If such a bisimulation exists, we write $E \approx_R E'$.

- This is a generalization of *history-preserving bisimulations* (Rabinovitch and Traktenbrot, 1988; van Glabeek and Goltz, 1989).

Bisimilar embeddings

- Let E, E' be event structures. A *bisimilar embedding* of E into E' is a relation $\iota \subseteq |E| \times |E'|$ such that:

  totality: $\pi_1(\iota) = |E|$;
  injectivity: for all $a, b \in |E|$, $\iota(a) \cap \iota(b) \neq \emptyset$ implies $a = b$;
  bisimilarity: $E \approx_\iota E'$; a ι-bisimulation proving this is said to be *associated with* ι.

- We write $E \overset{\iota}{\hookrightarrow} E'$ to denote the fact that ι is an embedding of E into E', or simply $E \hookrightarrow E'$ to state the existence of an embedding.

- Embeddings compose: if $E \overset{\iota'}{\hookrightarrow} E'$ and $E' \overset{\iota''}{\hookrightarrow} E''$, then $E \overset{\iota'' \circ \iota'}{\hookrightarrow} E''$.

- Morally, if computational processes P, P' are described by E, E', $E \hookrightarrow E'$ should mean that "P' faithfully simulates/is at least as expressive as P".

Confusion

Let $E = (|E|, \leq, \smile)$ be an event structure.

- Given $a, b \in |E|$, we write $a \# b$ (*immediate conflict*) iff $a \smile b$ and $\exists u \in \mathcal{C}(E)$ enabling both a and b.
- Note that every conflict is either immediate or inherited: $a \smile b$ implies $\exists a_0, b_0$ s.t. $a_0 \# b_0$ and $a_0 \leq a$, $b_0 \leq b$.
- An event structure is *confusion-free* (Varacca et al., 2006), iff:
  - the reflexive closure of $\#$ is transitive;
  - $a \# b$ implies $\lceil a \rceil = \lceil b \rceil$.
- Non-deterministic Turing machines are confusion-free. Here are two **non**-confusion-free structures:

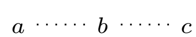

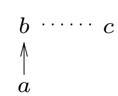

Two separation results

Theorem 2. *Let E, E' be event structures, with E' confusion-free. Then, $E \hookrightarrow E'$ implies E confusion-free.*

- We say that an embedding $E \overset{\iota}{\hookrightarrow} E'$ *introduces divergence* if, $\forall$ $\mathcal{B}$ associated with ι, $\exists$ $(u, \phi, u') \in \mathcal{B}$ and an infinite sequence of administrative transitions $u' \longrightarrow_\iota u'_1 \longrightarrow_\iota u'_2 \longrightarrow_\iota \cdots$

- An *anticlique* is a finite set of events in pairwise conflict, with a finite configuration enabling all of them. The *degree of non-determinism* of E is the least ordinal $\alpha \leq \omega$ such that, $\forall$ anticlique A of E, $\sharp A \leq \alpha$.

Theorem 3. *Let E, E' be event structures of degree of non-determinism $\alpha < \alpha'$, and let $E \overset{\iota}{\hookrightarrow} E'$. Then, ι introduces divergence.*

Discussion

- How meaningful is all this? In other words, how sensible is our notion of bisimilar embedding?

- Some well known encodings induce bisimilar embeddings (we will see this in Part II). That's good `:-)`

- However, there are surprises: the folklore encoding of non-deterministic Turing machines into non-deterministic Turing machines of degree of non-determinism 2 is problematic.

Non-deterministic Turing machines

- The degree of non-determinism of Turing machines coincides with the degree of non-determinism we defined for event structures. Hence, Theorem 3 applies, in striking contrast with the folklore encoding: simulate a branching of degree $n > 2$ with $n - 1$ successive branchings of degree 2. Such encoding only slows down the machine by a multiplicative factor, it does not introduce divergence in any reasonable sense.

- The problem is in how bisimulations treat non-determinism: *the only way to say "no" to somebody, is to say "yes" to someone else*.

- This is because bisimulations were conceived to deal with "open" systems. Turing machines are "closed". How do we deal with non-determinism in closed systems?