characterized by $\mathcal{L}'_{\mathcal{O}}$, provided that for each p, d the set $I(p,d)$ contains elements of only finitely many equivalence classes under $R_{\mathcal{O}}$. This generalizes Lemma 1.

Acknowledgement

Thanks to two anonymous referees, for their helpful suggestions.

1 References

[1] G. Boudol and I. Castellani, Concurrency and Atomicity, *Theoretical Computer Science* **59** (1,2) (1988) 25-84.

[2] G. Ferrari and U. Montanari, The Observation Algebra of Spatial Pomsets, in: *Proc. 2nd International Conference on Concurrency Theory*, Lecture Notes in Computer Science **527** (Springer, Berlin, 1991) 188-202.

[3] G. Ferrari, U. Montanari, and M. Mowbray, On Causality Observed Incrementally, Finally, in: *Proc. International Joint Conference on Theory and Practice of Software Development*, Lecture Notes in Computer Science **493** (1) (Springer, Berlin, 1991) 26-41.

[4] G. Ferrari, U. Montanari, and M. Mowbray, Causal Streams: Tracing Causality in Distributed Systems, in: *Proc. 3rd Workshop on Concurrency and Compositionality*, GMD-Studien **191** (1991) 99-108.

[5] M. Hennessy, R. Milner, Algebraic Laws for Nondeterminism and Concurrency, *Journal of A.C.M* **32** (1985) 137-161.

[6] R. Milner, A Calculus of Communicating Systems, Lecture Notes in Computer Science **92** (Springer, Berlin 1980)

[7] R. Milner, Communication and Concurrency (Prentice Hall, New Jersey, 1989)

Note that for fixed N, a, d, the set $\{b \in I(a,d), \mathcal{N}(b) \leq N\}$ is *not* always finite. For instance, if $N = 1$, $\mathcal{O}$ is weak Concurrent History observation, $a = recx.(\alpha.NIL + \tau.x[\Phi])$, and $d = \mathcal{O}(\alpha.NIL+ > (\tau \rightsquigarrow a[\Phi]))$, then the set includes $a[\Phi]$, $a[\Phi][\Phi]$, $a[\Phi][\Phi][\Phi]$, and so on.

If A is infinite rather than finite then $col(c)$ is not well defined for all computations c, and so condition 1 on $\mathcal{O}$ does not make sense. However, $col(c)$ *is* defined for all computations from an expression in CG', where CG' is the set of $b \in CG$ for which each permutation occurring syntactically in b fixes all but finitely many elements of $A \cup \overline{A}$. If CG is replaced by CG' in all the conditions and definitions, then the resulting lemma holds even if A is infinite.

Theorem 3

If $\mathcal{O}$ is a CCS observation map satisfying conditions 1,2,3, then $a\ R_{\mathcal{O}}\ b$ iff $\mathcal{L}'_{\mathcal{O}}(a) = \mathcal{L}'_{\mathcal{O}}(b)$.
Proof
Let R be the equivalence relation on the set of computations satisfying
$(c : a \stackrel{d}{\Rightarrow} b)\ R\ (c' : a' \stackrel{d'}{\Rightarrow} b')$ iff $a\ R_{\mathcal{O}}\ a'$, $b\ R_{\mathcal{O}}\ b'$, and $d = d'$.
Let the map $\mathcal{O}'$ from equivalence classes under R to $\mathcal{D}$ send the class containing c to $\mathcal{O}(c)$. By Lemma 2, $\mathcal{O}'$ is image finite.

Say that $[a] \models F$ iff $a \models F$, where $[a]$ is the equivalence class under $R_{\mathcal{O}}$ containing a, and $F \in \mathcal{L}'_{\mathcal{O}}$. This definition is consistent, since members of the same equivalence class under $R_{\mathcal{O}}$ satisfy the same formulae of $\mathcal{L}_{\mathcal{O}}$. By definition, $\mathcal{L}'_{\mathcal{O}'}([a]) = \mathcal{L}'_{\mathcal{O}}(a)$ for all a.

Applying Lemma 1, $\mathcal{L}'_{\mathcal{O}'}([a]) = \mathcal{L}'_{\mathcal{O}'}([b])$ iff $[a]\ R_{\mathcal{O}'}\ [b]$. By construction, $R_{\mathcal{O}'}$ is the identity relation on equivalence classes under $R_{\mathcal{O}}$. The result follows.
□

If the CCS observation map $\mathcal{O}$ does not satisfy conditions 1,2,3, then the result may not hold. For example, let $\mathcal{O}(t) = \nu$ whenever $t : a \stackrel{\nu}{\rightarrow} b$, $\mathcal{O}(t) =$ the empty string whenever $t : a \stackrel{\tau}{\rightarrow} b$, and $\mathcal{O}(a_1, t_1, \ldots, t_m, a_{m+1}) =$ the string concatenation of $\mathcal{O}(t_1)$, ..., $\mathcal{O}(t_m)$. Then $R_{\mathcal{O}}$ is Milner's weak bisimulation, and this is not the same as the relation induced by the logic $\mathcal{L}'_{\mathcal{O}}$. For example, if $a_1 = rec\ x.(\tau.(\alpha.NIL \mid x))$ and $a_2 = rec\ x.(\alpha.x)$, then a_1 and $a_1 + a_2$ satisfy the same formulae of $\mathcal{L}'_{\mathcal{O}}$ but are not related by $R_{\mathcal{O}}$.
If instead $\mathcal{O}$ is any of the observation maps satisfying conditions 1,2,3 which have been mentioned by name in this paper, then this counterexample does not work, because $a_1 + a_2$ satisfies the formula

$$< \mathcal{O}(a_1 + a_2, a_1+ > (\alpha \rightsquigarrow a_2), a_2, (\alpha \rightsquigarrow a_2), a_2) > TRUE$$

but a_1 does not.

The proof of Theorem 3 can be applied directly to the general case, where $\mathcal{O}$ is an observation map on the computations of *any* process algebra. In this case it shows that $R_{\mathcal{O}}$ is completely

Lemma 2

Suppose $\mathcal{O}$ is a CCS observation map satisfying conditions 1,2,3, $a \in CG$, and $d \in \mathcal{D}$. Then $I(a,d)$ contains elements of only finitely many equivalence classes under $R_{\mathcal{O}}$.

Proof (Sketch)
Suppose first that a is collapsed. Then a is the solution for x_0 in some set of guarded form equations

$$x_i = C_i(x_1, \ldots, x_m) : \ 0 \leq i \leq m$$

such that each C_i is *collapsed.*

Define the sets S_i, S_i', $i \geq 1$ of expressions as follows. S_1 is the finite set of expressions $C(a_1, \ldots a_m)$ such that C is a subcontext of C_i for some $0 \leq i \leq m$, a_i is the solution for x_i in the set of equations above, and $\mathcal{N}(C(a_1, \ldots, a_m))$=1. For each $i \geq 1$, S_i' is the set

$$\{s\backslash\alpha_1 \ldots \backslash\alpha_n : s \in S_i, n \geq 0, \alpha_1, \ldots \alpha_n \in A \text{ (not necessarily distinct)}\}$$

For each $i \geq 2$, S_i is the set

$$\{s_1 \mid s_2 : s_1 \in S_{i-j}, s_2 \in S_i \text{ for some } 1 \leq j \leq i-1\}$$

Let $S = \bigcup_{i \geq 1} S_i'$. S contains a.

Condition 2 ensures that for any b the equivalence relation $R_{\mathcal{O}}$ contains $(b\backslash\alpha, b\backslash\alpha\backslash\alpha)$ and $(b\backslash\alpha\backslash\beta, b\backslash\beta\backslash\alpha)$. Hence if S_i contains elements of only finitely many equivalence classes, so does S_i'. It is straightforward to show that if $b_1 \ R_{\mathcal{O}} \ b_1'$ and $b_2 \ R_{\mathcal{O}} \ b_2'$ then $b_1 \mid b_2 \ R_{\mathcal{O}} \ b_1' \mid b_2'$. It follows that for each $i > 1$, S_i' contains elements of only finitely many equivalence classes under $R_{\mathcal{O}}$.

Suppose that there is $t : a_1 \xrightarrow{\mu} a_2$ such that $a_1 \in S$, $a_2 \notin S$. Choose such a t whose derivation using the inference rules uses a minimal number of steps. A contradiction follows, by examination of each possibility for the last step.

Therefore S contains the targets of all computations from a. Condition 3 on $\mathcal{O}$ ensures that there is some integer N such that $\mathcal{N}(b) \leq N$ for all $b \in I(a,d)$, so $I(a,d)$ is contained in $\cup_{i \leq N} S_i'$, which contains representatives of only finitely many equivalence classes. The result follows.

Now suppose that $a \in CG$ is not collapsed. There is $c' : col(a) \stackrel{d'}{\Rightarrow} b'$ iff there is $c : a \stackrel{d}{\Rightarrow} b$ such that $c' = col(c)$. If $c' = col(c)$, then $d' = d$ by condition 1, and $b' = col(b)$ by definition of $col(c)$. It follows from the definition of $R_{\mathcal{O}}$ that $b \ R_{\mathcal{O}} \ col(b)$ for all $b \in CG$. Hence each equivalence class containing an element of $I(a,d)$ also contains an element of $I(col(a), d)$. The result follows by the collapsed case.
□

Write $\mathcal{L}_{\mathcal{O}}(p)$ for $\{F \in \mathcal{L}_{\mathcal{O}} : p \models F\}$. The proof of proposition 6, section 10.5 of [7] can be applied to show that $p\ R_{\mathcal{O}}\ q$ iff $\mathcal{L}_{\mathcal{O}}(p) = \mathcal{L}_{\mathcal{O}}(q)$.

The disadvantage of $\mathcal{L}_{\mathcal{O}}$ is that some of its formulae cannot be written as a finite string of symbols. Let $\mathcal{L}'_{\mathcal{O}}$ be the sublanguage consisting of expressions in which the logical operator $\wedge_{i \in I}$ is only ever used with finite index sets I. Each formula in $\mathcal{L}'_{\mathcal{O}}$ can be written as a finite string of symbols.

Theorems 2.1 and 2.2 of [5] give the following, which will be used in the proof of Theorem 3.

Lemma 1

If $\mathcal{O}$ is image finite, then $p\ R_{\mathcal{O}}\ q$ iff $\mathcal{L}'_{\mathcal{O}}(p) = \mathcal{L}'_{\mathcal{O}}(q)$.

Conditions

Theorem 3 will use the following conditions on the CCS observation map $\mathcal{O}$.
1. $\mathcal{O}(col(c)) = \mathcal{O}(c)$ for all CCS computations c.
2. $\mathcal{O}(t\backslash\alpha) = \mathcal{O}(t)$ for all transitions t and $\alpha \in A$
3. $\{\mathcal{N}(b) : b \in I(a,d)\}$ is finite for all $a \in CG$ and $d \in \mathcal{D}$.
Informally, these conditions state that $\mathcal{O}$ interprets the restriction and relabelling operators in a similar way to the usual interpretation of CCS, and that the observation of a computation together with its source state gives a bound on the top-level parallelism of its target state. Theorem 3 will show that these conditions on $\mathcal{O}$ are sufficient (although they are not necessary) for the bisimulation $R_{\mathcal{O}}$ to be characterized by the finitary logic $\mathcal{L}'_{\mathcal{O}}$.

If $\mathcal{O}$ is any CCS observation map which satisfies condition 1, then
$(\alpha.NIL)[\Phi]\ R_{\mathcal{O}}\ \Phi(\alpha).NIL$, so the identity relation on CG is not equal to $R_{\mathcal{O}}$ for any such $\mathcal{O}$. However the author knows of no other proposed observational bisimulation for CCS whose observation map does not satisfy conditions 1 and 2.

If $\mathcal{O}$ is image finite then condition 3 certainly holds. Condition 3 uses information about the internal structure of the target expression, which is not available in any semantics using a monolithic global state. In particular, the observation map associated with the ordinary weak interleaving semantics for CCS does not satisfy this condition.

The observation maps for the weak versions of of Concurrent History Equivalence, Causal Stream Equivalence, and Spatial Pomset Equivalence [3, 4, 2] satisfy all three conditions, but are not image finite.

- *col* commutes with operations yielding transitions except those of the form $[\Phi]$; that is, $col(\mu \rightsquigarrow \alpha) = \mu \rightsquigarrow \alpha$, $col(t < +a) = col(t) < +col(a)$, $col(a+ > t) = col(a)+ > col(t)$, $col(t \rfloor a) = col(t) \rfloor col(a)$, $col(a \lfloor t) = col(a) \lfloor col(t)$, $col(t_1 \mid t_2) = col(t_1) \mid col(t_2)$, $col(t \backslash \alpha) = col(t) \backslash \alpha$
- $col(t[\Phi])$ is the result of (syntactically) replacing each occurrence of the operator $\nu \rightsquigarrow$ in $col(t)$ by $\Phi(\nu) \rightsquigarrow$, and each occurrence of $\backslash\alpha$ by $\backslash(A \cap \{\Phi(\alpha), \Phi(\overline{\alpha})\})$, for each ν, α
- If c is the sequence $a_1, t_1, \ldots, a_m, t_m, a_{m+1}$ then $col(c)$ is the sequence $col(a_1), col(t_1), \ldots, col(a_m), col(t_m), col(a_{m+1})$

If c is a CCS computation then $col(c)$ is a CCS computation. In fact, it can be shown by structural induction that a sequence c' with first entry $col(a)$ is a CCS computation iff there is some CCS computation c from a such that $c' = col(c)$.

Observational bisimulations and logics

Suppose $\mathcal{O}$ is a observation map from a set of computations (this is not necessarily the set of CCS computations, it could be the set of computations of any process algebra) to a domain $\mathcal{D}$. Use $c : p \stackrel{d}{\Rightarrow} q$ as shorthand for "a computation c from p to q with $\mathcal{O}(c) = d$." Let P be the set of source states of computations, and let $I(p,d) = \{q : \text{there exists } c : p \stackrel{d}{\Rightarrow} q\}$. $\mathcal{O}$ is *image finite* if $I(p,d)$ is finite for all $p \in P$, $d \in \mathcal{D}$.

Given an observation map $\mathcal{O}$, define the bisimulation $R_{\mathcal{O}}$ on P to be the maximal equivalence relation R such that whenever $p_1 R p_2$ and $q_1 \in I(p_1, d)$ there is some $q_2 \in I(p_2, d)$ such that $q_1 R q_2$.

Let $\mathcal{L}_{\mathcal{O}}$ be the following Hennessy-Milner type modal logic [5]. The syntax of the logic is

$$F ::= TRUE \mid \neg F \mid \wedge_{i \in I} F_i \mid < d > F$$

where d ranges over $\mathcal{D}$ and I is any index set. The satisfaction relation $\models \subseteq P \times \mathcal{L}_{\mathcal{O}}$ is given by:

- $p \models TRUE$ for all $p \in P$
- $p \models \neg F$ iff $p \not\models F$
- $p \models \wedge_{i \in I} F_i$ iff $p \models F_i$ for all $i \in I$
- $p \models < d > F$ iff there is some $q \in I(p,d)$ such that $q \models F$

The equation $\mathcal{N}(a+b) = 1$ comes from an assumption that applying the choice operator gives a *global* state with just one autonomous component, that is, there is a centralized mechanism for nondeterministic choice. An alternative definition of $\mathcal{N}$ would have $\mathcal{N}(a+b) = max\{\mathcal{N}(a), \mathcal{N}(b)\}$, and the results of this paper can also be proven for this alternative definition.

A *CCS computation* of length $m \geq 0$ from a_1 to a_{m+1} is a sequence $a_1, t_1, a_2, \ldots, a_m, t_m, a_{m+1}$ such that for each $1 \leq i \leq m$ there is some $\mu(i)$ such that $t_i : a_i \stackrel{\mu(i)}{\rightarrow} a_{i+1}$.

A CCS context or expression is *collapsed* if it does not involve any operators $[\Phi]$.

For each $\Phi \in Perm$, define a renaming function f_Φ on recursion-free terms with variables in the set $\{x_{i,\Psi} : i \geq 1, \Psi \in Perm\}$, as follows.

- $f_\Phi(NIL) = NIL$
- $f_\Phi(x_{i,\Psi}) = x_{i,\Psi\circ\Phi}$ where $\Psi \circ \Phi(\mu) = \Phi(\Psi(\mu))$
- $f_\Phi(\mu.E) = \Phi(\mu).f_\Phi(E)$
- $f_\Phi(E_1 + E_2) = f_\Phi(E_1) + f_\Phi(E_2)$
- $f_\Phi(E_1 \mid E_2) = f_\Phi(E_1) \mid f_\Phi(E_2)$
- $f_\Phi(E\backslash\alpha) = f_\Phi(E)\backslash(A \cap \{\Phi(\alpha), \Phi(\overline{\alpha})\})$
- $f_\Phi(E[\Psi]) = f_{\Psi\circ\Phi}(E)$

If a is the solution for x_0 in the guarded form equations

$$x_i = C_i(x_1, \ldots, x_m) : \ 0 \leq i \leq m$$

then define $col(a)$ to be the solution for $x_{0,I}$ (where I is the identity permutation) in the equations

$$x_{i,\Phi} = f_\Phi(C_i(x_{1,I}, \ldots, x_{m,I})) : \ 0 \leq i \leq m, \Phi \in Perm$$

For example, if $\Phi \circ \Phi = I$, then $col(rec\ x.(\mu.x[\Phi])) = rec\ x.(\mu.\Phi(\mu).x)$.

$col(a)$ is a collapsed expression in CG, and it can be shown that it is independent of the choice of guarded form equations for which a is the solution.

Define the collapsing function on CCS transitions and computations, also called col, as follows.

The syntax for CCS expressions is

$$E ::= NIL \mid x \mid \mu.E \mid E + E \mid E|E \mid E\backslash\alpha \mid E[\Phi] \mid rec\, x.E$$

A closed guarded CCS expression a is an expression with no free variables, such that x is guarded in E (i.e. each occurrence of x in E is within some subexpression of the form $\mu.E'$) whenever $rec\ x.E$ is a subexpression of a.
rec is treated as a fixpoint operator. Suppose x is guarded in E, and E has no other free variables. Then $rec\ x.E$ is defined to be a solution for x in the equation $x = E$, and in this paper *all* solutions for x in this equation are identified with $rec\ x.E$.
a, b (and a_1, b_1, etc.) will denote elements of the set CG of closed guarded CCS expressions, modulo the identifications for recursion.

The identifications ensure that there is a unique solution in CG for x_0 in the equations

$$x_i = C_i(x_1, \ldots x_m)\ 0 \leq i \leq m$$

where $C_0, C_1, \ldots C_m$ are recursion-free contexts, and x_j is guarded in $C_i(x_1, \ldots x_m)$ for $i, j \geq 1$. Say that such a set of equations has *guarded form*. Each $a \in CG$ is the solution of some guarded form equations.

The dynamic behaviour of CCS is given by the following inference rules.

$$\frac{}{\mu \rightsquigarrow a : \mu.a \xrightarrow{\mu} a} \qquad \frac{t : a \xrightarrow{\mu} b}{t < +c : a + c \xrightarrow{\mu} b} \qquad \frac{t : a \xrightarrow{\mu} b}{c+ > t : c + a \xrightarrow{\mu} b}$$

$$\frac{t : a \xrightarrow{\mu} b}{t\rfloor c : a \mid c \xrightarrow{\mu} b \mid c} \qquad \frac{t : a \xrightarrow{\mu} b}{c\lfloor t : c \mid a \xrightarrow{\mu} c \mid b} \qquad \frac{t_1 : a_1 \xrightarrow{\nu} b_1,\ t_2 : a_2 \xrightarrow{\overline{\nu}} b_2}{t_1 \mid t_2 : a_1 \mid a_2 \xrightarrow{\tau} b_1 \mid b_2}$$

$$\frac{t : a \xrightarrow{\mu} b,\ \mu \notin \{\alpha, \overline{\alpha}\}}{t\backslash\alpha : a\backslash\alpha \xrightarrow{\mu} b\backslash\alpha} \qquad \frac{t : a \xrightarrow{\mu} b}{t[\Phi] : a[\Phi] \xrightarrow{\Phi(\mu)} b[\Phi]}$$

Some notation

Define $\mathcal{N}(a)$, "the amount of parallelism in the top level of a", by induction on the structure of a as follows.

- $\mathcal{N}(NIL) = \mathcal{N}(a + b) = \mathcal{N}(\nu.a) = 1$
- $\mathcal{N}(a\backslash\alpha) = \mathcal{N}(a[\Phi]) = \mathcal{N}(a)$
- $\mathcal{N}(a \mid b) = \mathcal{N}(a) + \mathcal{N}(b)$

Finitary logics for some CCS observational bisimulations

Miranda Mowbray
Hewlett Packard Pisa Science Centre, Corso Italia 115, Pisa, Italy
mjfm@hplb.hpl.hp.com

November 7, 1995

Keywords: concurrency, formal semantics, modal logic, finitary

Introduction

This paper proves that observational bisimulations for CCS which satisfy certain conditions have a *finitary* associated modal logic.

The conditions involve information on the amount of top-level parallelism of states. This result may add fuel to the interleaving-versus-true-concurrency debate, since it suggests that the ability to observe internal structure of states (which is what distinguishes non-interleaving semantics from interleaving semantics) can lead not only to semantics which are more expressive, but also to semantics which are in some way simpler – their associated modal logics are finitary.

Some cases of the result (for instance, Strong Equivalence [7] and Pomset Bisimulation Equivalence [1]) were already known, due to Lemma 1, which was proved in [5] and is generalized here.

The language CCS

The process description language CCS was introduced in [6]. Its syntax and behaviour are recalled here.

Throughout this paper α will range over the finite set A, ν over $A \cup \overline{A}$, μ over $A \cup \overline{A} \cup \{\tau\}$, and x over an infinite set of variables. Φ will range over $Perm$, the set of permutations of $A \cup \overline{A} \cup \{\tau\}$ such that $\Phi(\tau) = \tau$ and $\Phi(\overline{\nu}) = \overline{\Phi(\nu)}$ for all ν.