CMU-CS-83-112

Behavioral Equivalence Relations

Induced By Programming Logics

Stephen D. Brookes
Carnegie-Mellon University

William C. Rounds
University of Michigan

March 1983

A version of this paper will appear in *Proceedings of ICALP 1983.*

The research of the first author was supported in part by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-81-K-1539. The second author was partially supported by the National Science Foundation under NSF Grant No. MCS-8102286. The views and conclusions contained in this report are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the NSF, or the US Government.

BEHAVIORAL EQUIVALENCE RELATIONS INDUCED BY PROGRAMMING LOGICS

Stephen D. Brookes
Carnegie-Mellon University
Pittsburgh, Pa.

William C. Rounds
University of Michigan
Ann Arbor, Mi.

1.0. Abstract.

In this paper we compare the descriptive power of three programming logics by studying the elementary equivalence relations which the logics induce on nondeterministic state-transition systems. In addition, we compare these relations with other natural state-equivalence relations for nondeterministic systems. We find that the notions of *bisimilarity* (Park [P], Ogden [O]) and *observation equivalence* (Milner [M]) are very strong equivalences compared with those induced by the logics. These three comprise *regular trace logic* (RTL), *propositional dynamic logic* (PDL), and *Hennessy-Milner logic* (HML). Regular trace logic is a new logic which can be used to give behavioral specifications for concurrent systems (*e.g.* Wolper [W], but with significant differences). It is a way of formalising those properties of programs which have been given informally in terms of path expressions [CH]. The model theory and axiomatics of this logic are interesting in their own right. Propositional dynamic logic is well-known; our treatment differs from the standard one only in that we regard the modalities as specifying intended behavior instead of being programs. Hennessy-Milner logic is a simplified modal logic which those authors used as a characterisation of their notion of observation equivalence, which we call weak observation equivalence in this paper. We also include a brief treatment in this context of two other natural equivalences for nondeterministic systems: *failure* equivalence [HBR] and *trace* equivalence [H], both of which are weaker than the relations induced by the logics but can be characterised using appropriate logical subsets.

1.1. Introduction.

In this paper we generalise the notion of state-equivalence, familiar from the theory of sequential machines, to the case of nondeterministic transition systems, and use it to investigate some issues in the semantics of parallel processes. We are motivated by several recent studies of parallelism and concurrency. In particular we are interested in modelling systems which can be controlled through interactions with a surrounding environment, but which are also capable of making internal or hidden moves, in a way which cannot be influenced by an outside agent. This sort of behavior naturally demands a nondeterministic model, especially when no probabilities can reasonably be attached to the internal actions.

Hennessy and Milner [HM] seem to have been the first to consider this modelling problem explicitly. They give a general definition of *observational equivalence* and prove a basic characterization of the property in terms of a simple modal logic. Milner subsequently considered several variations of the definition in his CCS [M]. These studies rely on the idea of a *transition system* with explicitly named actions and hidden actions. A similar model was used by Keller [K] in his early work on formalizing the notion of concurrent computations. The present work differs from Keller's in that the result of a single transition may be nondeterministic, and in having implicit or hidden transitions as well.

Our work can be seen as a direct extension of the Hennessy-Milner results. We introduce several alternatives to the definition of observational equivalence. The first was suggested to us by W. Ogden, and we subsequently learned from R. Milner that the same definition was independently given by David Park, who called it *bisimulation.* This mathematically appealing relation is the strongest of our equivalence relations. Next we consider a class of equivalences based on *distinguishability by logical formulae.* These formulae are drawn from three logics: the first is the original Hennessy-Milner logic (HM); the second, is a version of Pratt-Fischer-Ladner's *propositional dynamic logic* (PDL); and the third, called *regular trace logic,* is a new logic intended to express the nondeterministic capabilities of transition systems in a way not possible with PDL. We do not treat temporal logic in this respect since our semantics deals only with finite sequences of states and actions. Indeed, because our logics talk only about what properties of a system can be observed during the occurrence of a finite sequence of transitions, this treatment cannot cope with eventualities and fairness conditions as can temporal logics [Pn] and process logics [HKP]. We hope that our techniques can be extended to these systems. Milner and Hoare, however, have emphasised that the concept of *finite observability* is all that one could ever reasonably use in practice to detect behavioral differences between concurrent systems.

We compare the strength of these various equivalence relations, including the one already proposed by Hennessy and Milner. Under the hypothesis that all transitions have a finite number of possible outcomes, the definitions all coincide; a theorem amounting to this was proved by Hennessy and Milner. In the absence of this hypothesis, however, significant differences emerge. Perhaps the most interesting of these is the one which states that observational equivalence is a strictly stronger relation than the relation of indistinguishability by formulae of regular trace logic. The proof of this fact uses a rather delicate argument involving the pumping lemma for regular sets.

1.2. Fundamental definitions.

Our fundamental model of computation is the *labelled transition system.* This very general model was first used by Keller to study concurrent systems, although its origins can be traced back to Petri nets and to the nondeterministic automata of Rabin and Scott [RS]. Formally, a labelled transition system is a tuple

$$(Q, \Sigma, q_0, \rightarrow)$$

where Q is a set of ***states*** (at most countable), Σ is a finite ***alphabet,*** $q_0 \in Q$ is the ***initial state*** and $\rightarrow$ is a transition relation, *i.e.* a ternary relation on $Q \times (\Sigma \cup \{\tau\}) \times Q$. The special symbol τ

(not in Q) is used to denote a ***hidden action.*** We will write Δ for the set $\Sigma \cup \{ \tau \}$. The variable μ will range over Δ. If $(q, \mu, q') \in \rightarrow$, we write $q \xrightarrow{\mu} q'$.

As usual, Σ^* (resp. Δ^*) denotes the set of finite strings over Σ (resp. Δ); the empty sequence is ϵ, and we use s, t, u to range over Σ^*, and v, w over Δ^*. The operation $\backslash\tau$ of deleting all occurrences of τ in a sequence converts $w \in \Delta^*$ into $w\backslash\tau \in \Sigma^*$. The transition relation extends in the usual way to a relation on $Q \times \Delta^* \times Q$. For $s \in \Sigma^*$ we define the relation $\xRightarrow{s}$ on $Q \times Q$ by:

$$q \xRightarrow{s} q' \quad \Leftrightarrow \quad \exists w \in \Delta^*. \, q \xrightarrow{w} q' \ \& \ w\backslash\tau = s.$$

Notice that the symbol τ does not appear in any of the strings s used in the previous definition. This is because we are only interested in the aspects of behavior which can be inferred from "experimenting" with externally visible events. We need to retain the possibility of uncontrolled behavior in the underlying system, however; this behavior might be caused, for example, by an unfair scheduler or some other unfriendly agent.

A transition system can be "unrolled" into a tree in the usual way. The initial state labels the root, states label the nodes, and elements of $\Sigma \cup \{ \tau \}$ label the arcs. The resulting tree is called a *synchronisation tree* (ST). Often, where the set of states is implicit, we will identify this tree with the transition system. When we do this, the transition relation on trees is as given by Milner [M]. A transition system (or ST) has *finite branching* iff for each $q \in Q$ and each $s \in \Sigma^*$ the set $\{ q' \mid q \xRightarrow{s} q' \}$ is finite. This property obviously holds if Q is finite, but note that it is not implied by the condition that the set $\{ q' \mid q \xrightarrow{\mu} q' \}$ is finite for all $\mu \in \Sigma \cup \{ \tau \}$. This latter condition is known as *image-finiteness* of the transition system.

We write aS for the synchronisation tree with a unique initial branch labelled $a \in \Delta$ and subtree S attached. If $\{ S_i \mid i \in I \}$ is a family of STs we denote by

$$\sum_{i \in I} S_i$$

the tree obtained by identifying the root nodes of all the S_i. The trivial tree, with a single node and no arcs, is *NIL*.

Finally, we will use α to denote a regular subset of Σ^*. We will not generally distinguish between a regular expression and the language it denotes.

1.3. Equivalences.

(i) Bisimulation.

Let $S = (Q, \Sigma, q_0, \rightarrow)$ be a transition system. A relation $\gamma \subseteq Q \times Q$ is ***invariant*** if whenever $p\gamma q$, $a \in \Delta$ and $p \xrightarrow{a} p'$, then there is a q' such that $q \xrightarrow{a} q'$ and $p'\gamma q'$. A ***bisimulation*** is a relation γ such that both γ and its inverse relation γ^{-1} are invariant. Two states p and q are ***bisimilar*** if there is a bisimulation γ such that $p\gamma q$. We write pBq in this case. This definition can be used to extend bisimulation to a relation on synchronisation trees, with initial states, in the obvious way. Note that if γ is invariant, then whenever $p\gamma q$ and $s \in \Sigma^*$ any transition sequence $p \overset{s}{\Longrightarrow} p'$ corresponds to a sequence $q \overset{s}{\Longrightarrow} q'$ for which $p'\gamma q'$. This can be shown by induction on the length of the transition sequence. The following properties of bisimulation are elementary.

PROPOSITION 1.3.1. *The relation of bisimulation is an equivalence relation.*

Proof. The composition of two bisimulations is again a bisimulation. ∎

PROPOSITION 1.3.2. *The relation B is itself a bisimulation, and any bisimulation on Q is a subset of B.*

Example 1. Consider the two STs $a(b+c)$ and $a(b+c)+a(b+c)$. If the root nodes of these two trees are states in some transition system, then there is an obvious bisimulation. However, there is no bisimulation between $ab+ac$ and $a(b+c)$.

(ii) Observational equivalence.

DEFINITION 1.3.3. Let p and q be states of some transition system. We define a sequence $\approx_n$ of equivalence relations as follows:

$$\begin{aligned} &p \approx_0 q \quad \text{always} \\ &p \approx_{n+1} q \Leftrightarrow \forall s \in \Sigma^*. \\ &\qquad (i) \quad p \overset{s}{\Longrightarrow} p' \Rightarrow \exists q'.q \overset{s}{\Longrightarrow} q' \;\&\; p' \approx_n q' \\ &\qquad (ii) \quad q \overset{s}{\Longrightarrow} q' \Rightarrow \exists p'.p \overset{s}{\Longrightarrow} p' \;\&\; p' \approx_n q'. \end{aligned}$$

We define $pOq \quad \Leftrightarrow \quad \forall n.p \approx_n q$.

Example 1. We present a sequence (S_k, T_k) of pairs of trees in which the n^{th} pair is n-equivalent but not $(n+1)$-equivalent, for all n.

$$\begin{aligned} S_0 &= aNIL & T_0 &= NIL \\ S_{n+1} &= aS_n + aT_n & T_{n+1} &= aS_n \end{aligned}$$

It is easy to check that these pairs have the desired properties.

DEFINITION 1.3.4. (Weak observational equivalence) Let p and q be states in Q as above. Say that p is weakly equivalent to q iff $p \approx_n^W q$ for all n, where $\approx_n^W$ is defined inductively by

$$\begin{aligned} &p \approx_0^W q \quad \text{always} \\ &p \approx_{n+1}^W q \Leftrightarrow \forall a \in \Sigma \cup \{\epsilon\}. \\ &\qquad (i) \quad p \overset{a}{\Longrightarrow} p' \Rightarrow \exists q'.q \overset{a}{\Longrightarrow} q' \;\&\; p' \approx_n q' \\ &\qquad (ii) \quad q \overset{a}{\Longrightarrow} q' \Rightarrow \exists p'.p \overset{a}{\Longrightarrow} p' \;\&\; p' \approx_n q'. \end{aligned}$$

Weak observational equivalence is obtained by restricting attention at each stage in the construction to sequences of visible actions of length at most one. The relation W of weak observational equivalence is in general weaker than observation equivalence, although, as Milner states, the two relations coincide on image-finite trees. We are not making the assumption that all transition systems are image-finite.

(iv) Failure equivalence.

In the failure set semantics of nondeterministic communicating processes [HBR], the behavior of a process is described in terms of so-called ***failures:*** each failure is a pair (s, X) in which s records a possible sequence of visible transitions and X is a set of transitions which the process ***may***, as the result of a nondeterministic decision, be incapable of performing on the next step. This leads naturally to a failure equivalence relation on synchronisation trees (see also [B1]).

DEFINITION 1.3.5. The failure set of a synchronisation tree S is

$$\text{failures}(S) = \{\,(s, X) \mid \exists S'.S \overset{s}{\Longrightarrow} S' \;\&\; \forall x \in X \; S' \overset{x}{\not\Longrightarrow}\}.$$

We use the abbreviation $S \overset{x}{\not\Longrightarrow}$ for $\neg\exists S'(S \overset{x}{\Longrightarrow} S')$. The failure equivalence relation is defined as follows:

DEFINITION 1.3.6. Trees S and T are failure equivalent iff failures(S) = failures(T). An extensive discussion of the properties enjoyed by this equivalence relation, and its use in giving a semantics to concurrent processes, can be found in [HBR] and [RB]; fuller accounts are provided in [B2] and [R].

(v) Trace equivalence.

An early model for process semantics was based on the notion of ***traces*** [H]. A trace is a finite sequence of visible transitions, and two synchronisation trees are trace-equivalent iff they have the same set of possible traces:

DEFINITION 1.3.7. Two trees S and T are trace-equivalent iff

$$\forall s \in \Sigma^*.(\exists S'.S \overset{s}{\Longrightarrow} S') \Leftrightarrow (\exists T'.T \overset{s}{\Longrightarrow} T').$$

This is a very simple equivalence relation, and coincides with Milner's $\approx_1$ above.

(vi) Logical equivalences.

We present three logics which can be interpreted in state transition systems, and which describe aspects of the behavior of such systems. The behavioral equivalences induced by these logics are exactly the elementary equivalences for the corresponding models.

These logics are all specializations of ***regular trace logic*** (Rounds and Gurevich), with which we begin our definitions.

Syntax. The formulas ϕ, ψ, of regular trace logic (RTL) are built up from ***constants*** T and F, by means of boolean combinations using &, $\vee$ and $\neg$, and by modal combinations $\forall\alpha\langle\phi\rangle$ and $\forall\alpha[\phi]$. Here α is a regular expression over Σ.

Semantics. We interpret RTL formulas in state transition systems. A ***structure*** for RTL consists of a transition system $S = (Q, \Sigma, q_0, \rightarrow)$ and a state $q \in Q$. The system S will usually be understood. We define the *satisfaction* relation as follows, by structural induction on the formulas:

$$\begin{array}{rcl}
q \models T & \text{always} & \\
q \models F & \text{never} & \\
q \models \phi \,\&\, \psi & \Leftrightarrow & q \models \phi \text{ and } q \models \psi \\
q \models \phi \vee \psi & \Leftrightarrow & q \models \phi \text{ or } q \models \psi \\
q \models \neg\phi & \Leftrightarrow & \text{not } q \models \phi \\
q \models \forall\alpha(\phi) & \Leftrightarrow & \forall s \in \alpha . \exists q' \; q \xRightarrow{s} q' \;\; \& \;\; q' \models \phi \\
q \models \forall\alpha[\phi] & \Leftrightarrow & \forall s \in \alpha . \forall q' \; q \xRightarrow{s} q' \Rightarrow q' \models \phi
\end{array}$$

Remarks:

1. α is a regular expression for the syntax, and denotes a regular subset of Σ^* for the semantics.

2. Propositional variables have been omitted. We are interested in purely behavioral properties, and not in the conditions holding of a state.

3. We may define two further modalities:

$$\begin{array}{rcl}
\exists\alpha(\phi) & \equiv & \neg(\forall\alpha[\neg\phi]) \\
\exists\alpha[\phi] & \equiv & \neg(\forall\alpha(\neg\phi)).
\end{array}$$

Clearly,

$$\begin{array}{rcl}
q \models \exists\alpha(\phi) & \Leftrightarrow & \exists s \in \alpha . \exists q' \; q \xRightarrow{s} q' \;\; \& \;\; q' \models \phi, \\
q \models \exists\alpha[\phi] & \Leftrightarrow & \exists s \in \alpha . \forall q' \; q \xRightarrow{s} q' \; \Rightarrow \; q' \models \phi.
\end{array}$$

DEFINITION 1.3.8. PDL (propositional dynamic logic) is obtained from RTL by omitting the modality $\forall\alpha(\cdot)$.

DEFINITION 1.3.9. HML (Hennessy-Milner logic) is obtained from RTL by restricting the sets α used in modalities to be singletons.

Notice that if $\alpha = \{ s \}$ is a singleton set then the existential and universal quantifiers have the same effect. We abbreviate $\forall\{ s \}[\phi]$ and $\exists\{ s \}[\phi]$ to $s[\phi]$ and similarly for the other modality.

We will be interested in the possibility of distinguishing between states of a transition system by the sets of formulas which they satisfy. If L is a subset of the class of RTL formulas, we define L-equivalent states to be states which satisfy precisely the same formulas from L. In the sequel we will use the name L to denote these equivalence relations, where L can be HML, PDL or RTL.

2. Classifying equivalence relations.

In this section we compare the various equivalence relations introduced so far. We consider the following relations on transition systems:

B	bisimulation
O	observational equivalence
W	weak observational equivalence
RTL	
PDL	equivalence w.r.t. formulas of a logic
HML	
F	failure equivalence
T	trace equivalence

We obtain the following diagram, in which the strongest equivalences are at the top:

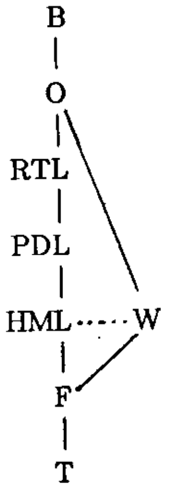

We will show that the equality $HML = W$ holds when Σ is finite. Hennessy and Milner's results established that $B = O = HML = W$ in case all relations are of finite image.

THEOREM 2.1. $B \subseteq O \subseteq RTL \subseteq PDL \subseteq HML \subseteq F \subseteq T$.

Proof. We only prove in detail the inclusions $B \subseteq O \subseteq RTL$. First we indicate briefly the reasons for the final two inclusions in our diagram. The other cases are also straightforward.

($HML \subseteq F \subseteq T$) Trace equivalence is clearly the elementary equivalence induced by the set of formulas of the form $s\langle T\rangle$ of HML, and failure equivalence is characterised similarly by the set of formulas of the form $s\langle\psi\rangle$, where ψ has the form $a_1[F] \,\&\, \ldots a_n[F]$ for some set $\{a_1, \ldots, a_n\} \subseteq \Sigma$. (We adopt the convention that an empty conjunction denotes T.) Thus it follows easily that $HML \subseteq F \subseteq T$.

($B \subseteq O$) To prove that bisimilarity implies observation equivalence, let γ be a bisimulation. We show that $p\gamma q \Rightarrow p \approx_n q$, for all n. We proceed by induction on n. The case $n = 0$ is trivial, since $\approx_0$ is the universal relation. Suppose the result holds for $n = k$. Let $p\gamma q$. If $p \xRightarrow{s} p'$, then by bisimilarity of p and q there is a state q' such that $q \xRightarrow{s} q'$ and $p'\gamma q'$. By the inductive hypothesis, this implies $p' \approx_k q'$, which is the conclusion we need for $k+1$-equivalence of p and q. The converse (for q) goes through similarly. Thus we have proved by induction that $p\gamma q \Rightarrow \forall n. p \approx_n q$. Thus, $p\gamma q \;\Rightarrow\; pOq$.

($O \subseteq RTL$) For the inclusion $O \subseteq RTL$, we argue by induction on the length of formulas, showing that

$$\forall k.(p \approx_k q \Rightarrow \forall\phi(\text{length}(\phi) \leq k \Rightarrow (p \models \phi \Leftrightarrow q \models \phi)))$$

Again the base case is trivial. Assume the result for length k. Let ϕ have length $k+1$, and assume that $p \approx_{k+1} q$. There are several cases, depending on the structure of ϕ.

Case 1: ϕ is $\neg\psi$, with length$(\psi) \leq k$. In this case we use the fact that $p \approx_{k+1} q \Rightarrow p \approx_k q$, easily established by induction. Then we have by inductive hypothesis that $p \models \psi \Leftrightarrow q \models \psi$, giving also $p \models \neg\psi \Leftrightarrow q \models \neg\psi$, as required.

Case 2: ϕ is $\forall\alpha\langle\psi\rangle$, with length$(\psi) \leq k$. Suppose $p \models \phi$ and let $s \in \alpha$. Then there is a state p' such that $p \overset{s}{\Longrightarrow} p'$ and $p' \models \psi$. Since we are assuming that $p \approx_{k+1} q$, there is also a q' such that $q \overset{s}{\Longrightarrow} q'$ and $p' \approx_k q'$. By the inductive hypothesis, we must have $q' \models \psi$. Since we chose s arbitrarily in α this argument shows that $q \models \forall\alpha\langle\psi\rangle$, as required. The converse $(q \models \phi \Rightarrow p \models \phi)$ is similar.

Case 3: ϕ is $\forall\alpha[\psi]$, with length$(\psi) \leq k$. Suppose $p \models \phi$. We want $q \models \phi$. Let $s \in \alpha$ and suppose $q \overset{s}{\Longrightarrow} q'$, for some q'. Then we know there is a state p' such that $p \overset{s}{\Longrightarrow} p'$ and $q' \approx_k p'$, because p and q are assumed to be $k+1$-equivalent. Since $p \models \phi$, we know that $p' \models \psi$. By the inductive hypothesis, $q' \models \psi$. This gives $q \models \forall\alpha[\psi]$,b as required. The converse is again similar. ∎

The rest of this section is devoted to establishing that the inclusions of Theorem 2.1 are proper.

THEOREM 2.2. *Observation equivalence properly includes bisimulation, $O \supset B$.*

Proof. We find U and V which are observation equivalent but not bisimilar. To begin, let us define some auxiliary trees. Let (S_n, T_n) be a sequence of pairs in which the n^{th} pair is $\approx_{n-1}$ equivalent but not $\approx_n$ equivalent, for each n. An example of such a sequence was given in the previous section. If S is an ST and $s \in \Sigma^*$, we write sS for the tree obtained by prefixing a path s to the root of S. Let $c \in \Sigma$. Define for each $i \geq 0$,

$$U_i = \sum_{j=1}^{\infty} c^j T_j + \sum_{j=1}^{i-1} c^j S_j$$

and let

$$U_\omega = \sum_{j=1}^{\infty} c^j T_j + \sum_{j=1}^{\infty} c^j S_j.$$

Let $U = \sum_{i=1}^{\infty} aU_i$ and $V = U + aU_\omega$. Now we show that

$$(\dagger) \qquad \forall i.(U_i \not\approx_{i+1} U_\omega \;\&\; U_{i+1} \approx_i U_\omega)$$

This is easy to establish. For each i, $U_\omega \overset{c^i}{\Longrightarrow} S_i$, but the subtree S_i is missing from U_i. In fact, if $U_i \overset{c^i}{\Longrightarrow} T$ and $T \approx_i S_i$ then the only possibility would be that $T = T_i$, which cannot happen by construction. On the other hand, for $j \geq i$, we know that $S_j \approx_i T_j$. Thus, $(c^j S_j + c^j T_j) \approx_i c^j T_j$

for each such j. It follows that

$$U_\omega \approx_i (\sum_{j=1}^{\infty} c^j T_j + \sum_{j=1}^{i} c^j S_j) = U_{i+1}.$$

This shows that (†) holds. To see why (†) implies the desired conclusion, suppose for a contradiction that U and V are bisimilar. Let γ be a bisimulation such that $U \gamma V$. Since $V \overset{a}{\Longrightarrow} U_\omega$, by construction of U there must be an i such that $U \overset{a}{\Longrightarrow} U_i$ and $U_i \gamma U_\omega$. But this implies $U_i B U_\omega$, which contradicts Theorem 1, since we have $U_i \not\approx_i U_\omega$, by (†). Thus U and V are not bisimilar. On the other hand, $U \approx_i V$ for each i. Indeed, let i be fixed and let $V \overset{s}{\Longrightarrow} V'$ for some $s \in \Sigma^*$. We show that U has a similar transition sequence to some U' with $U' \approx_{i-1} V'$. The only interesting case is when the transition of V enters the subtree U_ω, clearly. In this case, s must have the form at, and $U_\omega \overset{t}{\Longrightarrow} V'$. By (†), $U_{i+1} \approx_i U_\omega$. Therefore $U \overset{a}{\Longrightarrow} U_{i+1} \overset{t}{\Longrightarrow} U'$ for some subtree U' of U_{i+1} with $U' \approx_{i-1} V'$. This is the conclusion needed for $\approx_i$-equivalence of U and V. So U and V are observation equivalent but not bisimilar, as required. ∎

Next we investigate the logical equivalences.

THEOREM 2.3. *RTL properly contains O: $RTL \supset O$.*

Proof. The proof of this theorem is quite intricate, and is deferred until we have established some definitions and lemmas.

For any set $L \subseteq \Sigma^*$, we let S^L be the synchronisation tree determined by L : having one branch for each $s \in L$;

$$S^L = \sum_{s \in L} sNIL.$$

When L is a singleton $\{ s \}$ it determines the tree $sNIL$ with a single branch; we will identify this tree with the string s where there is no possibility of confusion.

LEMMA 2.4. *Let $\phi \in RTL$. Then the set $L(\phi) = \{ s \in \Sigma^* \mid sNIL \models \phi \}$ is regular.*

Proof. Induction on the structure of RTL formulas. The base case, when ϕ is either T or F, is trivial: $L(T) = \Sigma^*$, $L(F) = \emptyset$. The boolean combinations go through because regular sets are closed under complement, intersection and union:

$$\begin{aligned} L(\neg\phi) &= \Sigma^* - L(\phi), \\ L(\phi \,\&\, \psi) &= L(\phi) \cap L(\psi), \\ L(\phi \vee \psi) &= L(\phi) \cup L(\psi). \end{aligned}$$

If α denotes an infinite regular set, then $L(\forall\alpha[\phi]) = L(\forall\alpha\langle\phi\rangle) = \emptyset$, because each tree under consideration here is finite. Otherwise, let α denote the finite set $\{ s_1, \ldots, s_k \}$. Then $L(\forall\alpha\langle\phi\rangle) = \bigcap_{i=1}^{k} s_i L(\phi)$. The same holds for the other modality, because the trees here have only a single branch. That concludes the proof of Lemma 1. ∎

LEMMA 2.5. *Suppose α denotes a regular set not containing the empty sequence. Let ϕ be of the form $\forall\alpha[\psi]$ or $\exists\alpha[\psi]$. Then for all L, if $S(L) \models \phi$ then $S(L') \models \phi$, for all $L' \subseteq L$.*

Proof. Routine application of the definitions. ▮

Now consider the language

$$K = \{\, a^{n^2} b^{n+j} \mid n \geq 1, 0 \leq j \leq n \,\}.$$

Let $W = S^K$, and for each pair n, j let W^n_j be the tree obtained from W by removing the branch $a^{n^2} b^{n+j}$. Notice that K is not a regular language. We will use the properties of this set in constructing two trees which cannot be distinguished by any RTL formula, but which are not observation equivalent. First we establish an important property of the synchronisation tree W determined by K.

LEMMA 2.6. *Let ϕ be an RTL formula. If $W \models \phi$ then for all but finitely many pairs n, j the tree W^n_j also satisfies ϕ.*

Proof. First put ϕ into *monotonic form,* by moving all the negations through to the inside. Then ϕ has one of the forms:

$$T,\ F,\ (\psi \vee \theta),\ (\psi \,\&\, \theta),$$
$$\exists\alpha\langle\psi\rangle,\ \exists\alpha[\psi],\ \forall\alpha\langle\psi\rangle,\ \forall\alpha[\psi],$$

where ψ and θ are also in monotonic form. We proceed by induction on the structure of monotonic formulas. The base case and the boolean combinations are straightforward. The first three modalities are simple, with the help of Lemma 2.5. The final case is when ϕ is $\forall\alpha\langle\psi\rangle$. Notice that the structure of W allows us to assume without loss of generality that α does not contain the empty string, as W has no nontrivial empty transitions. We therefore assume $\epsilon \not\in \alpha$. Moreover, we claim that without loss of generality we can assume that $\alpha \subseteq a^+$. To show this, first notice that because $W \models \forall\alpha\langle\psi\rangle$, every string in α is a prefix of some string in K; thus $\alpha \subseteq a^*b^*$. Suppose that $\alpha \cap a^*b^+$ is infinite. For each n, the number of strings of the form $a^{n^2}b^j$ which are prefixes of members of K is finite, because j can be at most $2n$. Therefore, the set

$$\{\, t \in a^*b \mid \exists u(tu \in \alpha \cap a^*b^+) \,\}$$

is an infinite regular subset of $\{\, a^{n^2}b \mid n \geq 0 \,\}$, which is impossible. Therefore $\alpha \cap a^*b^+$ is finite, say $\{\, t_1, \ldots t_m \,\}$. We may therefore write

$$\forall\alpha\langle\psi\rangle \equiv \forall\beta\langle\psi\rangle \,\&\, t_1\langle\psi\rangle \,\&\, \ldots t_m\langle\psi\rangle,$$

where $\beta \subseteq \alpha^+$. The claim follows using the inductive reasoning for the & connective.

Now suppose the conclusion of Lemma 2.6 to be false for $\forall\alpha\langle\psi\rangle$, where $\alpha \subseteq a^+$. Then $W \models \forall\alpha\langle\psi\rangle$, but for infinitely many pairs n, j $W^n_j \models \neg\forall\alpha\langle\psi\rangle$. Let $s = a^{n^2}b^{n+j}$ be a trace corresponding to such a pair. For each such s there must therefore be a prefix $t \leq s$ such that $t \in \alpha$ and $W^n_j \models t[\neg\psi]$. This t must be a prefix of s because $W \models t\langle\psi\rangle$ and W^n_j differs from W only in the s-branch. We may therefore write $s = tu$, where $t \in a^*$ and $u \models \psi$. Furthermore, for any $w \neq s$ in K if $w = tv$ for some v then $v \models \neg\psi$. Recall that $L(\psi)$ is a regular language. Let k be the number of states of a FSA accepting this language. Choose an $s \in K$ as above but such that it has the form $a^{n^2}b^{n+j}$ with $n > 2k$. Decomposing s into tu as above, the number i of b's in u must satisfy

$$n \leq i \leq 2n,$$

because $0 \leq j \leq n$. Since $u \in L(\psi)$, the FSA accepting u must repeat a state while reading across the b's. Further, this happens in at most p steps, where $p \leq k$. We then have

$$u = a^r b^i,$$

for some r, and hence $a^r b^{i-p} \in L(\psi)$. If $i-p \geq n$ we have a contradiction, because then $ta^r b^{i-p} \in K$, and $a^r b^{i-p} \models \neg\psi$, which is impossible. But if $i-p < n$ we have

$$i + p < n + 2p \leq n + 2k \leq 2n,$$

because $2k \leq n$. Therefore $a^r b^{i+p} \in L(\psi)$ and $ta^r b^{i+p} \in K$. Therefore $a^r b^{i+p} \models \neg\psi$, again a contradiction. That completes the proof of Lemma 2.6. ▮

Now we are ready to prove Theorem 2.2. Define the two trees

$$U = \sum_{n,j} aW_j^n,$$
$$V = U + aW.$$

It is clear that $U \not\approx_2 V$, because $V \overset{a}{\Longrightarrow} W$, and for each n, j pair we have $W_j^n \not\approx_1 W$. So U and V are not observation equivalent. We claim that U and V satisfy precisely the same RTL formulas:

$$\forall \phi \in RTL. (U \models \phi \;\leftrightarrow\; V \models \phi).$$

The proof is again by induction on the structure of ϕ. The base case and the boolean connectives are trivial. We consider the cases $\phi = \exists\alpha\langle\psi\rangle$ and $\phi = \exists\alpha[\psi]$ in detail. The other modalities can be deduced using the argument for negation.

Case 1. Let $\phi = \exists\alpha\langle\psi\rangle$. Clearly if $U \models \phi$ then so does V. Conversely, if $V \models \phi$ then choose $t \in \alpha$ and V' such that $V \overset{t}{\Longrightarrow} V'$ and $V' \models \psi$. If $t = \epsilon$ then we must have $V = V'$, and $V \models \psi$; then by inductive hypothesis, $U \models \psi$, from which we get $U \models \phi$. The only other possibility is that $t = au$ for some u. If $V \overset{a}{\Longrightarrow} W_j^n \overset{u}{\Longrightarrow} V'$ there is no problem, because U has a similar subtree. Suppose that $V \overset{a}{\Longrightarrow} W \overset{u}{\Longrightarrow} V'$. Then $W \models u\langle\psi\rangle$, so by Lemma 2.6 there is a W_j^n also satisfying this formula. Hence, $U \models au\langle\psi\rangle$ and $U \models \phi$.

Case 2. Let $\phi = \forall\alpha\langle\psi\rangle$. Consider the possibilities for $s \in \alpha$. In each case we must show that U has an s-branch leading to a subtree where ψ holds. If $s = \epsilon$ we can use the inductive hypothesis. If s is traceable into some W_j^n there is no problem, because U has a corresponding transition. If s is traceable into W, then $s = at$ for some t which must be a prefix of a string w in K. If $t = \epsilon$ then $s = a$ and $W \models \psi$; by Lemma 2.6 there is a pair n, j with $W_j^n \models \psi$. Otherwise, t is a prefix of some string in K, and for all but one pair n, j the tree W_j^n has a branch t. In each case we have shown that U has a corresponding s-branch. That completes the proof. ▮

THEOREM 2.7. *RTL is properly contained in PDL.*

Proof. We give only an example to show that the inclusion cannot be reversed. The proof that PDL cannot distinguish the two trees follows the lines of Theorem 2.2 but is much easier. Define

$$S_n = \sum_{i=1}^{n} b^i NIL,$$
$$S_\omega = \sum_{i=1}^{\infty} b^i NIL.$$

Let U and V be the trees

$$U = \sum_{n=1}^{\infty} aS_n,$$
$$V = U + aS_\omega.$$

Then the RTL formula $a\langle\forall b^+\langle T\rangle\rangle$ is satisfied by V but not U, since only V has an a-branch to a place where arbitrarily many B-transitions can be made. (Here we have used the notation b^+ for $b^*-\{\epsilon\}$.) However, all PDL formulas agree on U and V. The relevant lemma is: for any PDL formula ϕ, if $S_\omega \models \phi$ then for all but finitely many n $S_n \models \phi$. ∎

THEOREM 2.8. *PDL properly is contained in HML.*

Proof. Again we exhibit an example. Let $U = \sum_{i\geq 1} a^i bNIL$, and $V = U + a^\omega NIL$. Then the PDL formula $a[\exists a^*\langle b\langle T\rangle\rangle]$ is satisfied by U but not V, because V has an a-branch to a place where no future b-transitions are possible. However, all HML formulas agree on U and V, because if ψ is an HML formula and $a^\omega \models \psi$, then $a^n \models \psi$ also holds for all but finitely many n. ∎

Theorem 2.8 will also follow from the fact that U and V are weakly observation equivalent, once we have established that $W = HML$. Hennessy and Milner proved this result in the case when the underlying system has the *finite-image property:* for each $a \in \Delta$ and each $q \in Q$ the set

$$\{\, q' \mid q \xrightarrow{a} q' \,\}$$

is finite. They also showed the identity $HML = B$ under the finite-image hypothesis. We now show that the result $W = HML$ still holds when the finite-image hypothesis is not assumed, provided we assume that Σ is finite. Since we are allowing infinitely branching systems, finiteness of Σ does not, of course, imply the finite-image property. Our proof makes use of normal-form arguments for HML which are of independent interest. Indeed, these normal form results can be used to show that a natural pseudometric structure on synchronisation trees induces a compact metric topology on the set of W-equivalence classes (see [GR]).

THEOREM 2.9. *If Σ is finite, then $HML \subseteq W$.*

Proof. Let $\equiv$ denote the relation of logical equivalence between HML formulas:

$$\phi \equiv \psi \ \Leftrightarrow\ \forall p\,(p \models \phi \ \Leftrightarrow\ p \models \psi).$$

Define the *depth* of an HML formula as follows:

$$\begin{aligned} \text{depth}(T) &= \text{depth}(F) = 0 \\ \text{depth}(\neg\phi) &= \text{depth}(\phi) \\ \text{depth}(\phi \vee \psi) &= \text{depth}(\phi \,\&\, \psi) = \max(\text{depth}(\phi), \text{depth}(\psi)) \\ \text{depth}(a\langle\phi\rangle) &= 1 + \text{depth}(\phi). \end{aligned}$$

The depth of a formula is the maximum number of nested modalities. We let $H_k = \{ \theta \mid \text{depth}(\theta) \leq k \}$, for each $k \geq 0$. Then for each k there is an integer E_k such that $\equiv$ partitions H_k into at most E_k equivalence classes. To show this, we give an algorithm for converting an arbitrary $\phi \in H_k$ into a disjunctive normal form ϕ^* such that distinct normal forms are logically inequivalent and the number of distinct normal forms is less than or equal to E_k. We use induction on depth. Every HML formula is either *basic,* which we define to mean of the form T or F or $a\langle\psi\rangle$ for some ψ, or else a boolean combination of such basic formulas. It is easy to see that a depth 0 fomula is logically equivalent to either T or F, so that H_0 is partitioned into two distinct equivalence classes. We may, therefore, put $E_0 = 2$. For the inductive step, let $\phi \in H_{k+1} - H_k$, and suppose that H_k is partitioned into E_k equivalence classes. Let the modal subformulas of ϕ be $a_i\langle\theta_i\rangle$; each a_i can be assumed to belong to the set $\Sigma \cup \{ \epsilon \}$, and each θ_i has lower depth than ϕ. Put each θ_i into normal form. There are at most E_k possible normal forms for each θ_i, and we may replace logically equivalent formulas. We can then treat ϕ as a propositional combination of at most $m = E_k \times (\mid \Sigma \mid +1)$ variables, and as such put it into disjunctive normal form. In order to guarantee that H_{k+1} has no more than E_{k+1} equivalence classes, we may take $E_{k+1} = 2^{2^m}$. Now we prove by induction on k that

$$(1) \quad \forall k.\forall p, q\,(p \not\approx^W_k q \;\Rightarrow\; \exists\phi(\text{depth}(\phi) \leq k \;\&\; p \models \phi \;\&\; q \models \neg\phi)).$$

The base case is trivial, as we may choose $\phi = T$. Assume the result for k, and suppose $p \not\approx^W_{k+1} q$. Then for some $a \in \Sigma \cup \{ \epsilon \}$ there is a p' for which $p \overset{a}{\Longrightarrow} p'$ and p' is not $\approx^W_k$ to any q' such that $q \overset{a}{\Longrightarrow} q'$. Let the set of possible a-derivatives of q be $\{ q_i \mid i \in I \}$. Notice that we are not assuming this set to be finite. By hypothesis there are distinguishing formulas θ_i, each of depth at most k, such that for each $i \in I$ we have

$$p' \models \neg\theta_i \;\&\; q_i \models \theta_i.$$

For each $i \in I$ let θ_i^* be a normal form logically equivalent to θ_i. Note that $\text{depth}(\theta_i) = \text{depth}(\theta_i^*) \leq k$. Only finitely many of these normal forms can be logically inequivalent, say $\theta_1^*, \ldots, \theta_m^*$. Let $\phi = a\langle\theta_1^* \;\& \ldots \&\; \theta_m^*\rangle$. Then $p \models \neg\phi$ and $q \models \phi$. That completes the proof. ∎

COROLLARY 2.10. *If* Σ *is finite,* $HML = W$.

Proof. The inclusion $W \subseteq HML$ follows by a straightforward induction, using the converse to the inductive hypothesis of Theorem 2.8. This does not depend on the finiteness hypothesis. ∎

3. Conclusions.

We have investigated the descriptive power of three programming logics by examining the elementary equivalence relations induced on nondeterministic state transition systems by the logics. These equivalence relations have also been examined in the context of some other natural behavioral equivalence relations from the literature. An exact characterisation of a behavioral equivalence as the elementary equivalence induced by a particular logic provides an indication of the essential semantic properties of the equivalence; equally, delineating the relationships between the various existing equivalences serves to illuminate their differences.

We have shown that in general the three logical equivalences are not as discriminating as other natural behavioral equivalences such as Milner's observation equivalence, but are themselves finer than failure equivalence and trace equivalence. These latter two relations are, in fact, characterizable as the elementary equivalences generated from restricted sets of logical formulas.

We finish with a remark on complexity. If we interpret the logics in nondeterministic finite state automata, the finite-branching condition holds, and all of the logical equivalence relations coincide with observation equivalence and the bisimulation relation; moreover, it can be shown that these equivalences are decidable in polynomial time. In contrast, failure equivalence of finite automata turns out to be a PSPACE-complete problem.

4. Acknowledgements.

We would like to thank Joyce Friedman for suggesting the proof of Theorem 2.9. Our original proof was much more complicated.

5. References.

[B1] Brookes, S.D., On the relationship of CCS and CSP, to appear in Proceedings of ICALP83

[B2] Brookes, S.D., A Model for Communicating Sequential Processes, Ph.D. thesis, University of Oxford (submitted 1983).

[CH] Campbell, R., and Habermann, N., The Specification of Process Synchronization by Path Expressions, Springer LNCS Vol. 16.

[GR] Golson, W.G., and Rounds, W.C., Connections between Two Theories of Concurrency: Metric Spaces and Synchronisation Trees, Technical Report, Computing Research Laboratory, University of Michigan (January 1983).

[H] Hoare, C.A.R., A model for Communicating Sequential Processes, Technical Report PRG-22, University of Oxford, Programming Research Group (1981).

[HBR] Hoare, C.A.R., Brookes, S.D., and Roscoe, A.W., A Theory of Communicating Sequential Processes, Technical Report PRG-16, Oxford University, Programming Research Group (1981).

[HKP] Harel, D., Kozen, D., and Parikh, R., Process Logic: Expressiveness, Decidability and Completeness, Proceedings of IEEE Symposium on Foundations of Computer Science (1980).

[HM] Hennessy, M., and Milner, R., On Observing Nondeterminism and Concurrency, Proc. 7th ICALP, Springer LNCS Vol. 85 (1980).

[K] Keller, R., Formal Verification of Parallel Programs, CACM 19, Vol. 7 (July 1976).

[M] Milner, R., A Calculus of Communicating Systems, Springer LNCS Vol. 92.

[O] Ogden, W.F., Private communication.

[P] Park, D.M.R., Concurrency and Automata on Infinite Sequences, Computer Science Department, University of Warwick.

[Pn] Pnueli, A., The Temporal Logic of Programs, Proceedings of IEEE Symposium on Foundations of Computer Science (1977).

[RS] Rabin, M.O., and Scott, D.S., Finite Automata and their Decision Problems, IBM J. Res. 3:2 (1959).

[R] Roscoe, A.W., A Mathematical Theory of Communicating Processes, Ph.D. thesis, Oxford University (1982).

[RB] Rounds, W.C., and Brookes, S.D., Possible Futures, Acceptances, Refusals, and Communicating Processes, Proc. 22^{nd} IEEE Symposium on Foundations of Computer Science (October 1981).

[W] Wolper, P., Temporal Logic can be more expressive, Proc. 22^{nd} IEEE Symposium on Foundations of Computer Science.