Independence: logics and concurrency

J. C. Bradfield

Laboratory for Foundations of Computer Science, Division of Informatics, King's Bldgs, University of Edinburgh, Edinburgh EH9 3JZ, UK. Email jcb@dcs.ed.ac.uk

Abstract: We consider Hintikka et al.'s 'independence-friendly first-order logic'. We apply it to a modal logic setting, defining a notion of 'independent' modal logic, and we examine the associated fixpoint logics.

Key words: Independence, branching quantifiers, modal logic, fixpoints

1 Introduction

Modal and temporal logics have a long history as system specification languages in computer science, and computer scientists' study of temporal logic has generated many interesting theoretical developments, as well as many important practical advances. The logics in question typically describe properties of execution paths of systems, either explicitly, as in LTL and CTL, or via a 'next-step' operator and fixpoints, as in the modal mu-calculus.

Another major issue in computer science is concurrency. The theoretical and practical analysis of concurrent and distributed systems also has a long history. Probably the most successful approach to understanding concurrency at a fundamental semantic level is the use of 'independence models', in which certain events are stated to be independent of other events, and the associated partial order semantics, in which a partial order of causality is established between events. There have also been several major practical advances in exploiting partial order structure, such as stubborn sets and sleep sets [Val91] and unfolding techniques [McM92].

Naturally, one wishes to have modal and temporal logics for concurrent systems. There are several ways to apply the paradigm of normal temporal logic, typically by working on event structures, or similar models, and having explicit operators representing concurrency or independence: see [PK95] for a survey of such logics. In some models, and therefore in the associated logics, the notion of independence is abstract; in others, it is more concrete, perhaps derived from a notion of location; but in all cases it inheres in the model rather than the logic.

However, there is also a notion of inherent independence in logic, which seems quite natural when specifying concurrent systems, but which has not, with one notable exception, received much attention in concurrency theory.

In 1961, Leon Henkin [Hen61] introduced the quantifier $\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix}$, which is intended to mean that the choice of y depends only on x, *not* on u; and similarly the choice of v depends only on u; a formal semantics is given by using suitable Skolem functions.

The Henkin quantifier and its generalizations received some attention: the main papers were at the beginning of the 70s, by William Walkoe [Wal70] and Herbert Enderton [End70], and later also by Krynicki and Barwise. It was shown that $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ suffices to express all *partially ordered* or *branching* quantifiers (that is, quantifiers where the dependence of variables need not be simply linear). More recently, Georg Gottlob has looked at generalized quantifiers, including $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$, for capturing complexity classes.

In the last few years, logicians, in particular the (Boston and) Helsinki logicians Jaakko Hintikka, Gabriel Sandu and Jouko Väänänen, have returned with a vengeance to the study of 'independent' quantifiers as a basic logic, rather than a specialist extension, and Hintikka and Sandu [HiS96] have gone so far as to claim that their absence from Frege's logic was a 'fundamental error', a 'horror', and to claim that their reintroduction heralds 'a revolution in logic' (albeit with a question mark in the title). It is not necessary to accept this thesis to accept that 'independence-friendly first-order logic' (IF-FOL) is an interesting, and natural [Bar79,San93] logic. For example, when considering the development of systems by several designers working independently, the notion of independent choice seems natural, and may reasonably form part of any language for discussing properties of such designs. At a more concrete level, different components of a distributed system cannot be assumed to have full knowledge of other components, and so cannot make fully informed choices.

In this paper, I suggest that the use of modalities based on Henkin quantifiers gives an approach to independence in concurrency which complements the model-based independence usually used. I start by outlining the work on independent quantifiers; I then give a natural modal logic using independent quantification, and relate it to other work on logics for distributed systems; I define the obvious fixpoint extension, and consider complexity issues and the extent to which the usual theory transfers. I then return to the first-order independence logic giving our metalanguage, and consider briefly the interpretation of fixpoints there: which is not a trivial question. Finally, I raise the problem of more general modal fixpoint logics based on independence. Owing to space constraints, proofs are generally omitted or just sketched, unless they are of particular interest.

2 'Independence-friendly' first-order logic

By FOL^+ we mean first-order logic *in positive form*: that is, negation is applied only to atomic formulae, and both operators of the dual pairs $\vee/\wedge$ and $\exists/\forall$ are considered primitive; FOL denotes the usual logic where negation is not so restricted. $\mathrm{FOL}^+{+}Q$ denotes the logic where the additional operator Q appears only positively. We use similar notation for other logics.

2.1 Partial quantifiers à la *Henkin*

A *branching quantifier* Q is a set $\{x_1, \ldots, x_m, y_1, \ldots, y_n\}$ of variables, carrying a partial order $\prec$; the x_i are universal, the y_i existential. The semantics of $Q\phi$ is defined to be

that of $\exists f_1 \ldots f_n . \forall x_1 \ldots x_m . \phi[f_i(y_i\downarrow)/y_i]$, where $y_i\downarrow$ is the list of variables $\prec y_i$, and $[\cdot/\cdot]$ denotes syntactic substitution: thus f_i is a Skolem function for y_i, but it refers only to variables preceding y_i in the partial order.

In particular, the *Henkin quantifier* $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix} = \{x_1, x_2, y_1, y_2\}$ with $x_i \prec y_i$ is written $\begin{smallmatrix}\forall x_1 & \exists y_1 \\ \forall x_2 & \exists y_2\end{smallmatrix}$; thus $\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix}\phi(x, y, u, v)$ is equivalent by definition to $\exists f, g . \forall x, u . \phi(x, f(x), u, g(u))$.

Fact 1 The following are the basic properties of branching quantifiers.

(1) Any formula $Q\phi$ where Q is a branching quantifier and ϕ is FOL$^+$ is equivalent to a formula of FOL$^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$. [Wal70,End70].

(2) By definition, any FOL$^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ formula is equivalent to an existential second-order formula.

(3) Moreover, any existential second-order formula is equivalent to a formula of FOL$^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$. (This is because the formula can be 'unskolemized' using $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$: for example, the assertion that there exists an injective function which never takes a particular value c, which being true only on infinite structures is not first-order, can be written as $\exists f . \forall x, u . (f(x) = f(u) \Rightarrow x = u) \wedge (f(x) \neq c)$; and this can be written using $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ as $\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix} . (x = u \Rightarrow y = v) \wedge (y = v \Rightarrow x = u) \wedge (y \neq c)$.)

(4) It follows via Fagin's Theorem that on finite structures FOL$^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ expresses NP-hard properties [BlG86]; in fact, as one would expect, FOL$+\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ captures L$^{\mathrm{NP}}$ [Got97].

2.2 Partial knowledge games and quantifiers à la *Hintikka*

An alternative way of giving semantics to branching quantifiers is via games. Recall the Hintikka model-checking game for FOL$^+$: given a formula ψ and a structure M, a position is a subformula $\phi(\vec{x})$ of ψ together with a *deal* for ϕ, that is, an assignment of values $\vec{v}$ to its free variables $\vec{x}$. At a position $(\forall x . \phi_1, \vec{v})$, Abelard chooses a value v for x, and play moves to the position $(\phi_1, \vec{v} \cdot v)$; similarly Eloise moves at $\exists x . \phi$. At $\phi_1 \wedge \phi_2$, Abelard chooses a conjunct; and at $\phi_1 \vee \phi_2$, Eloise chooses a disjunct. A play of the game terminates at (negated) atoms $P(\vec{x})$ (resp. $\neg P(\vec{x})$), and is won by Eloise (resp. Abelard) iff $P(\vec{x})$ is true with the current deal. Then it is standard that $M \models \phi$ exactly if Eloise has a winning strategy in this game, where a strategy is a function from sequences of legal positions to moves.

These games have *perfect information*; both players know everything that has happened, and in particular when one player makes a choice, they know the other player's previous choices. Game semantics for the Henkin quantifiers, following [HiS96], use games of imperfect information: in the game for $\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix}\phi$, when Eloise chooses for v, she does not know what Abelard chose for x. To make this explicit, the logic is written with a more general syntax which is linear rather than two dimensional. We here use Hodges' [Hod97] syntax, which addresses certain flaws in Hintikka's; and for reasons of space and simplicity, we omit some operators in this version.

IF-FOL$^+$ is obtained from FOL$^+$ by modifying the syntax of quantifiers to be $\forall x / W .$ and $\exists x / W .$, where W is a set of variables. The intention is that W is the set of independent variables, whose values the player is not allowed to know at this choice point. Thus the Henkin quantifier $\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix}$ can be written as $\forall x / \emptyset . \exists y / \emptyset . \forall u / \{x, y\} . \exists v / \{x, y\}$. Henceforth

we freely omit set braces and write just $\forall x.$ for $\forall x/\emptyset$. We shall also, again for reasons of space and simplicity, consider only the independent $\exists$, leaving the dualization for independent $\forall$ to the reader (or see [Hod97]). The adapted model-checking game, where the information is restricted, can be shown to characterize the Skolem function semantics in the sense that Eloise has a winning strategy iff the formula is true. However, these games are not determined, so it is *not* true that Abelard has a winning strategy iff the formula is false. For example, $\begin{smallmatrix}\forall x\\ \exists y\end{smallmatrix} x = y$ (or $\forall x.\, \exists y/x.\, x = y$) is false in any structure with more than one element, but Abelard has no winning strategy.

3 Henkin modal logic

For a concurrency theorist, it is natural to see the model-checking game for $\begin{smallmatrix}\forall x & \exists y\\ \forall u & \exists v\end{smallmatrix}\phi$ not as a game of imperfect information, but as a *concurrent* game: at $\begin{smallmatrix}\forall x & \exists y\\ \forall u & \exists v\end{smallmatrix}\phi$ the game splits into two independent concurrent components: in one component play proceeds with $\forall x.\, \exists y.\,$, in the other with $\forall u.\, \exists v.\,$, and then the components join to proceed with ϕ. This also seems a natural statement to want to make about concurrent systems: a choice in one component should not even be able to depend on a concurrent choice in another! However, until recently all modal or temporal logics enforced logical dependence of choices.

Example 2 In the children's game Scissors–Paper–Stone, the two players (say Abe and Elly) each put one hand behind their back, and make it either open (Paper), a fist (Stone), or a V-sign (Scissors). The two players then *simultaneously* bring forward their hands. The round is won according to the rules: Scissors cut Paper, Paper wraps Stone, Stone blunts Scissors (if both players choose the same object, the round is drawn). Can we ask the question 'can Elly always win?'?

If we formalize the game by seeing Abe as choosing between the three actions $A = \{sc_A, pa_A, st_A\}$, and similarly for Elly, and viewing the game as their independent concurrent composition with final states that satisfy *Ewins* when Elly wins, we can approximate the question by $[A]\langle E\rangle \mathit{Ewins} \wedge \langle E\rangle[A]\mathit{Ewins}$. This expresses that Elly can win on all interleavings; it is (correctly) false, but it is false for the wrong reason: that $\langle E\rangle[A]\mathit{Ewins}$ fails, means that after Elly chooses, Abe can, with knowledge of that choice, make her lose, which mis-models the situation.

If we formalize the simultaneity by using a synchronous concurrent composition, there is, in normal modal or temporal logics, no way to ask the question at all, since only one action (the simultaneous choice) happens.

Of course, with an eye on Henkin quantifiers, the obvious answer is to define a modality $\begin{smallmatrix}[]\\ \langle\rangle\end{smallmatrix}$ so one can write $\begin{smallmatrix}[A]\\ \langle E\rangle\end{smallmatrix}\mathit{Ewins}$, with the intended meaning.

Recently, Alur, Henzinger and Kupferman [AHK97] have implicitly taken this approach in their Alternating Temporal Logic, apparently without awareness of Henkin quantifiers: they express, rather, their logic directly in terms of games and strategies. We shall now define a simple logic of Henkin modalities; a fragment of this provides a generalization of ATL, and also includes several other distributed logics.

3.1 A distributed system model

First let us define a notation for system models that, although by no means the most general possible, is sufficient for all our examples. Given an algebra Act of basic actions, sequential components P are built from non-deterministic choice $P_1 + P_2$, action prefix $a.P_1$, and mutual recursive definitions $P = Q(P)$. Systems are defined as a parallel composition $\|_i^{\mathscr{S}} P_i$ of n components; $Act_\perp = Act \cup \{\perp\}$ where $\perp$ is the 'non-action' or idle action, and $\mathscr{S} \subseteq Act_\perp^n \times Act$ gives the synchronization rules by $\|_i^{\mathscr{S}} P_i \xrightarrow{a} \|_i^{\mathscr{S}} P_i'$ (where $a \in Act$) iff each $P_i \xrightarrow{\alpha_i} P_i'$ (for $\alpha_i \in Act_\perp$) and also $\mathscr{S}(\alpha_1, \ldots, \alpha_n, a)$. We also write $\|_i^{\mathscr{S}} P_i \xrightarrow{\bigotimes_i \alpha_i} \|_i^{\mathscr{S}} P_i'$, particularly in the case when $\mathscr{S}$ is a function $Act_\perp^n \to Act$.

This model gives a convenient notation for various distributed automata formalisms, or for finite state CCS or CSP (by adjusting the synchronization); and a slightly less convenient notation for arbitrary finite 1-safe Petri nets, by taking each place as a two state component and using $\mathscr{S}$ to code the transitions. The alternating transition systems of [AHK97] can also be coded into this notation.

3.2 Henkin modalities

We now define Henkin modal logic (${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$) on such a system thus: in addition to ML^+, that is, basic modal logic in positive form, we have the *concurrent modalities* of the form $\bigotimes_{i=1,\ldots,n} \mathsf{Q}_i^1(A_i^1) \ldots \mathsf{Q}_i^m(A_i^m)$, where $\mathsf{Q}_i^j(A_i^j)$ is either $[A_i^j]$ or $\langle A_i^j \rangle$, and $A_i^j \subseteq Act_\perp$; the *length* of the modality is m. We shall sometimes use '$-$' to mean 'Act' in modalities. These modalities are given a semantics in terms of the corresponding first-order Henkin quantifier. This is notationally tedious to write out, so we just give the semantics of the length 2 Henkin modality $\begin{smallmatrix}[\,]\\[\,]\end{smallmatrix}\!\!\begin{smallmatrix}\langle\,\rangle\\\langle\,\rangle\end{smallmatrix}$ on a 2-component system, by way of example. (We write $\begin{smallmatrix}[A]\\[C]\end{smallmatrix}\!\!\begin{smallmatrix}\langle B\rangle\\\langle D\rangle\end{smallmatrix}$ in Henkin style for the formal $([A]\langle B\rangle \otimes [C]\langle D\rangle)$.)

$$P_1 \|^{\mathscr{S}} P_2 \vDash \begin{smallmatrix}[A_1]\\[A_2]\end{smallmatrix}\!\!\begin{smallmatrix}\langle B_1\rangle\\\langle B_2\rangle\end{smallmatrix} \phi$$

holds iff

$$\begin{array}{l}\forall \alpha_1 \in A_1, P_1' \, \exists \beta_1 \in B_1, P_1'' \\ \forall \alpha_2 \in A_2, P_2' \, \exists \beta_2 \in B_2, P_2''\end{array} (P_1 \|^{\mathscr{S}} P_2) \xrightarrow{\alpha_1 \otimes \alpha_2} (P_1' \|^{\mathscr{S}} P_2') \xrightarrow{\beta_1 \otimes \beta_2} (P_1'' \|^{\mathscr{S}} P_2'')$$

Note that by definition ${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$ does not include the duals of the Henkin modalities.

Let ${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_n$ denote the sublogic where the modalities have length at most n. ${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_1$ with the addition of fixpoints (see later) includes many existing distributed modal logics.

Example 3 The indexed transition systems of Andersen [And93] are systems $\|_{1 \le i \le n}^{\mathscr{S}} P_i$, where $\mathscr{S}(\alpha_1, \ldots, \alpha_n, a)$ iff there is some k such that $\alpha_k = a$ and all other α_i are $\perp$. (The systems move one component at a time.) Then in Andersen's 'polyadic mu-calculus', the indexed modality $[a]_k$ is in our notation $\bigotimes_i [\alpha_i]$ where $\alpha_k = a$ and the other α_i are all $\perp$.

Example 4 The distributed net systems used by Huhn, Niebert and Wallner [HNW97] are systems $\|_{1 \le i \le n}^{\mathscr{S}} P_i$ where each P_i is a sequential Petri net, and $\mathscr{S}$ defines the synchronization of common transitions. The transitions are labelled, which we may express either by using instead sets of transitions, or by letting $\mathscr{S}$ map local transitions to global labels

rather than to global transitions. Their logic is defined locally via an event structure style semantics, so that the basic modality is $\langle a \rangle_J \phi$, where $J \subseteq \{1, \ldots, n\}$; this is true at a *local state* of P_i, for $i \in J$, if a (which must share some location with J) can fire as an immediate action of the J processes, possibly with non-J actions happening first, and then ϕ holds.

Such a modality is not expressible in ${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_1$ without fixpoints, since we have an interleaving, global semantics, but with fixpoints it is just $\mu Z. \langle a \rangle \phi \vee (\bigotimes_i \langle A_i \rangle) Z$, where A_i is $\bot$ for $i \in J$ and $Act_\bot$ otherwise.

The previous examples are merely using the idea of location; to use the power of independent quantification, we need:

Example 5 The alternating transition systems of [AHK97] are an unlabelled setting. An ATS is a global state space Q, shared by n agents P_i. Each agent has a transition function $\delta_i : Q \to 2^{2^Q}$. At a state q, each agent chooses $Q_i \in \delta_i(q)$. It is required that for all such choices, $\bigcap_i Q_i$ is a singleton set $\{q'\}$, and this determines the next state of the system. Thus the intuition is that each agent chooses its desired successors, independently of the others, and the system moves to the one state they all desire.

To encode this in our setting, it is easiest to make the choice an explicit local transition happening before the move to the next state. So we take each agent P_i to be a transition system with state space $Q \cup 2^Q$, with transitions $q \xrightarrow{\tau} Q_i$ (where τ is a dummy label) for every $Q_i \in \delta_i(q)$, and $Q_i \xrightarrow{q'} q'$ for every $q' \in Q_i$ (note that q' is also being used as a label). The synchronization algebra $\mathscr{S}$ is then the diagonal on $Q \cup \{\tau\}$, and one move of the original ATS corresponds to two moves of our system $\|_i^{\mathscr{S}} P_i$. Thus the reachable states of our system have the form $(q, q, \ldots, q)$ or $(Q_1, \ldots, Q_n)$, and the requirements on the ATS transition function mean that there is exactly one transition from any reachable $(Q_1, \ldots, Q_n)$, to the unique $(q', q', \ldots, q')$ with $q' \in \bigcap_i Q_i$.

The Alternating Temporal Logic of [AHK97] is like CTL, except that the path quantifiers are not just $\forall$ and $\exists$, but have the form $\langle\langle E \rangle\rangle$, where $E \subseteq \{1, \ldots, n\}$. The interpretation is that, given a path formula ϕ (for example, $\mathbf{F}\psi$, 'eventually ψ'), $\langle\langle E \rangle\rangle \phi$ is true if the agents E can choose their successors to make ϕ hold on all paths, regardless of how the other agents $\overline{E}$ choose their successors. In other words, the E agents can win the path game defined by ϕ.

This logic is then in turn defined as a fragment of a suitable mu-calculus, just as CTL is defined in the usual modal mu-calculus. The [AHK97] mu-calculus has the basic modality $\langle\langle E \rangle\rangle \bigcirc \phi$, where $\bigcirc$ is the 'next' operator, interpreted as just described.

Now in our encoding, this modality is just $(\bigotimes_i \mathsf{Q}_i)\langle - \rangle \phi$, where Q_i is $\langle \tau \rangle$ for $i \in E$, and $[\tau]$ otherwise.

Remark 6 A referee suggests that my presentation of ${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$ on distributed systems obscures that fact that it is a natural fragment of IF-FOL just as ML is a natural fragment of FOL. There is some force in this, but there is a counter: if one defines the modalities without reference to locality, one must use a model such as ATS's in which locality is encoded by means of sets of possible next states for each agent. It is arguable that the

ATS model is less natural than the usual models, and also that the distributed modalities are themselves more natural for the user that a pure Henkin quantifier. A fuller discussion of this point requires further analysis of guarded and finite-variable fragments of IF-FOL, which, as far as I know, has not yet been done.

3.3 Adding fixpoints

We have already used fixpoint notation above; to justify this, it suffices to note the following fact:

Proposition 7 A Henkin modality defines a monotone operator on sets of global states. Therefore least and greatest fixpoint operators can be added to ${}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$ as in the normal modal mu-calculus: call this $\mu\,{}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$. □

Because we are defining interpretations as sets of global states, some of the usual theory carries through trivially: for example,

Proposition 8 The usual simple complexity upper bound applies: given a $\mu\,{}^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$ formula ϕ of length m and fixpoint alternation depth d, and a system of size n (meaning here the number of global states plus the number of transitions), the complexity of determing whether $s \vDash \phi$ is $O(m \cdot H \cdot n^d)$, where H is the cost of evaluating a basic modality or boolean. □

In normal modal mu-calculus, H is $O(n)$; we consider below what H is for Henkin modalities.

Other parts of the theory can be made to carry through in a rather uninteresting way, by ignoring the concurrency and using rather the Skolem semantics:

Proposition 9 Consider the tableau model-checking system of [StW90], and add rules for the Henkin modalities in the following form:

$$\frac{s_1 \| s_2 \vdash ([A_1]\langle B_1\rangle \otimes [A_2]\langle B_2\rangle)\phi}{s_{11} \| s_{21} \vdash \phi \ldots s_{1m} \| s_{2m} \vdash \phi}$$

where the states on the bottom are given by: for every A_1 successor s_1' of s_1 there is some B_1 successor s_{1j} of s_1', and similarly for the second component, such that $s_1 \| s_2 \stackrel{A_1 \otimes A_2}{\longrightarrow} s_1' \| s_2' \stackrel{B_1 \otimes B_2}{\longrightarrow} s_{1j} \| s_{2j}$.

Then the resulting tableau system is sound and complete. □

Other parts do not carry through so easily: for example, the relationships with automata and parity games, where one must extend the usual frameworks with either second-order moves or Henkin quantification in the winning conditions.

3.4 Cost of Henkin modalities

As we remarked above, the usual complexity analysis carries through, but depends, of course, on the cost of evaluating the modalities.

Proposition 10 The cost of evaluating a Henkin modality of length 1 is $O(n^3)$.

Proof. To check $s_1 \| s_2 \in [\![([A]\otimes\langle B\rangle)\phi]\!]$, it suffices to try in turn each of the $O(n)$ possible B-successors of s_2 and check whether all the so chosen $A \otimes B$-successors are included in $[\![\phi]\!]$. □

The result of this is that $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_1$ is not significantly worse (in theory!) than normal modal mu-calculus. However, since it is also a slight generalization of the AMC of [AHK97], the complexity hardness results there can also be applied to $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_1$, giving, for example, P-hardness (rather than the NL-hardness of modal mu-calculus).

When we move to real Henkin modalities, things become more expensive. As we know that first-order Henkin modalities are NP-complete, we might expect this; but we might also hope that the restricted quantification involved in modal logic reduces the complexity. Unfortunately, this is not the case:

Proposition 11 Model-checking $^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_2$ is NP-hard (and obviously NP).

Proof. There is a direct reduction (with thanks to Perdita Stevens) from CNF-SAT, which illustrates the use of the Henkin modalities quite nicely, and so is worth giving in full.

Consider an instance Φ of CNF-SAT: it has the form $\bigwedge_{1\le i\le m} C_i$, where $C_i = \bigvee_{1\le j\le n_i} \ell_{ij}$, and each ℓ_{ij} is v_k or $\neg v_k$ for one of the r variables v_k.

We now define a system with two components: the first represents the formula, the second a choice of variable assignment. Let $D = \sum_{1\le i\le m} c.C_i$, and $C_i = \sum_{1\le j\le n_i} d.L_{ij}$, and $L_{ij} = v_k.t$ if $\ell_{ij} = v_k$, or $L_{ij} = v_k.f$ if $\ell_{ij} = \neg v_k$

Let $A = \sum_{1\le k\le r} \overline{c}.V_k$, and $V_k = \overline{d}.\overline{v_k}.\overline{t} + \overline{d}.\overline{v_k}.\overline{f}$.

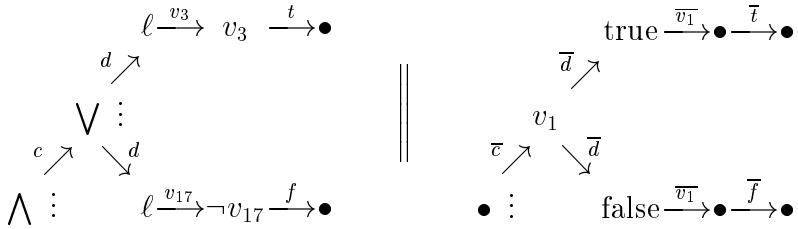

Now let $\mathscr{S}$ be the synchronization algebra given by $\alpha \otimes \overline{\alpha} = \tau$, and symmetrically, for every $\alpha \in \{c, d, v_k, t, f\}$, and let $P = D\|^{\mathscr{S}} A$. (In standard CCS notation, this would be $(D \mid A)\backslash\{c, d, v_k, t, f\}$.) Note that the size of this system is at most quadratic in the size of Φ.

Now we claim that $P \vDash \left[\begin{smallmatrix}-\\-\end{smallmatrix}\right]\left\langle\begin{smallmatrix}-\\-\end{smallmatrix}\right\rangle[\tau]\langle\tau\rangle\mathrm{tt}$ iff Φ is satisfiable.

Firstly, suppose Φ has a satisfying assignment S. In the game for $\left[\begin{smallmatrix}-\\-\end{smallmatrix}\right]\left\langle\begin{smallmatrix}-\\-\end{smallmatrix}\right\rangle$, Eloise plays the following strategy: in the top half, in response to Abelard's choice of C_i, she chooses any L_{ij} made true by S – such a literal exists, by definition of satisfying assignment. In the bottom half, in response to Abelard's choice of V_k, she chooses $\overline{v_k}.\overline{t}$ or $\overline{v_k}.\overline{f}$ according as v_k is true or false in S. The resulting process is then $v_{k'}.(t/f)\|^{\mathscr{S}}\overline{v_k}.(\overline{t}/\overline{f})$ according as ℓ_{ij} is $v_{k'}$ or $\neg v_{k'}$, and as v_k is true or false. If $k' \neq k$, then the process is deadlocked, and so satisfies $[\tau]\langle\tau\rangle\mathrm{tt}$. If $k' = k$, then since Eloise chose L_{ij} to be true in S, the left hand process has t iff the right hand process has $\overline{t}$, and so again the process satisfies $[\tau]\langle\tau\rangle\mathrm{tt}$.

Conversely, if $P \vDash \left[\begin{smallmatrix}-\\-\end{smallmatrix}\right]\left\langle\begin{smallmatrix}-\\-\end{smallmatrix}\right\rangle[\tau]\langle\tau\rangle\mathrm{tt}$, then Eloise's strategy for $\left[\begin{smallmatrix}-\\-\end{smallmatrix}\right]\left\langle\begin{smallmatrix}-\\-\end{smallmatrix}\right\rangle$ defines, in the bottom half, an assignment S, and in the top a selection of one literal L_{ij_i} for each conjunct C_i, such that the literal is made true by the assignment; so S is a satisfying assignment. □

So it appears that the Henkin modality is exponential to check. This has two apparent consequences: it means that one might reasonably argue it is useless, even if there are natural properties to be expressed with it; and it means that the fixpoint alternation depth is no longer the dominant factor in the combined complexity of model-checking.

The first consequence is in any case rather dubious – worst-case complexity may or may not have any relevance to practical complexity, as demonstrated by the practical utility of the theoretically non-elementary Mona system – but both consequences can be mitigated if systems have certain structure.

Given a system in our framework, define the local size to be the maximum of the sizes of the individual components.

Proposition 12 In a k-component system with local size d, evaluating a Henkin modality of length 2 costs $O(d^{d+2k}) = O(2^{(d+2k)\log d})$.

Proof. Brute force exploration of all the possible 'Skolem functions': in each component there are at most d^d possible local strategies in the local $[]\langle\rangle$ game; to check that a candidate k-tuple of local strategies satisfies the formula costs at most $d^k \cdot d^k$. Thus the total cost is d^{d+2k} □

In a loosely coupled system, the global state space is of size d^k; thus as the number of components increases, the exponential of the Henkin modality is absorbed by the exponential of the state space explosion, rather than adding to it.

The fact that the Henkin modality in the worst case dominates fixpoint complexity raises the question of the fixpoint alternation hierarchy for finite models. In the *infinite* case, this question can be solved by extending previous techniques.

Proposition 13 For $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_2$ on infinite models (and even on 2-component systems), the fixpoint alternation hierarchy is strict.

Proof. If one considers arithmetic with fixpoints and the Henkin quantifier, the fixpoint alternation hierarchy is strict – see the next section. Then a natural extension of the techniques of [Bra97], in which a Henkin modality is used to encode a first-order Henkin quantifier, transfers this hierarchy to $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_2$ on 2-component systems.

However, for this result to transfer down to finite models, we would need the finite model property for $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_2$ *and* its closure under negation. One can see that:

Proposition 14 Given a finite action set, fixed number of components, and given synchronization algebra, then if a $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}_2$ formula has a countably infinite model, it has a finite model.

Proof. The elegant proofs of the modal mu-calculus finite model property do not easily transfer. However, the brute force construction of a finite model by surgery on the unravelling, does transfer: essentially one builds an infinite tableau (using the tableau rule

given in Proposition 9 for the Henkin modality); then removes all branches that are not required to exist by a diamond modality; and then looks for repetitions of the same set of subformulae annotating a state (after defining a suitable notion of subformula for the Henkin modality). Closing repetitions to form loops gives a finite (albeit large) model. □

Unfortunately, I cannot see how to obtain this result for the dual of the Henkin modality, if indeed it holds.

4 IF-FOL and fixpoints

4.1 Fixpoints and the Henkin quantifier

We return now to the general setting of first-order logic. As we remarked in the last proposition, there is no great difficulty in combining a simple Henkin quantifier with fixpoints: $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ is just another monotone operator. Getting a handle on the expressive power is less trivial. Consider $\mathrm{LFP}^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$, that is, first-order logic with fixpoints in positive form, plus the Henkin quantifier occurring only positively, and consider the structure of arithmetic. The Henkin quantifier itself has Σ^1_1 power, so a formula of the form $\mu(z, Z).\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix}\phi$, for first-order ϕ, is at worst a fixpoint over Σ^1_1, which is at worst Σ^μ_2 in the normal fixpoint hierarchy (because $\Sigma^1_1 = \Pi^\mu_1$: Kleene's theorem). However, it is not immediately obvious that, say, $\begin{smallmatrix}\forall x & \exists y \\ \forall u & \exists v\end{smallmatrix}\mu(z, Z).\,\phi$ is well-behaved; naively, it might be Σ^1_2. Of course, it is not, because the fixpoint is parametrized on the variables x, y, u, v, and cannot actually refer to the defined Skolem functions. The key to analysing the expressive power is to extend the normal form results of [Lub93,Bra97] with the

Lemma 15 In arithmetic (or other structures with suitable coding power), a Henkin quantifier can be pushed inside a fixpoint operator. □

As corollaries of the normal form and existing theorems [Bra99a,Bra99b] on the power of fixpoints, one obtains

Corollary 16 The fixpoint alternation hierarchy for $\mathrm{LFP}^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ (on arithmetic) is strict. □

Corollary 17 $\mathrm{LFP}^+ + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ is no more expressive than LFP^+ (on arithmetic). □

and in fact both these apply to $\mathrm{LFP} + \begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ as well; adding the dual is not problematic. The latter result contrasts sharply with the position on finite models, where a single $\begin{smallmatrix}\forall & \exists \\ \forall & \exists\end{smallmatrix}$ is stronger than fixpoints, unless P = NP.

4.2 Fixpoints and full IF-FOL

Now consider the full independence-friendly logic, with the linear syntax allowing arbitrary specification of independence in quantifiers. We described the Skolem function semantics, and the game semantics of Hintikka. These semantics have the problem that they are not compositional, and in particular they give no meaning to a formula such

as $\exists y/x.\,\phi$, which occurs as a subformula of $\forall x.\,\exists y/x.\,\phi$ (alias $\begin{smallmatrix}\forall x\\ \exists y\end{smallmatrix}\phi$). Hintikka and Sandu [HiS96] thought that this was an inevitable fact, and Hintikka even went so far as to say that 'no perverse ingenuity' could produce a compositional semantics. So Wilfrid Hodges [Hod97] promptly gave a compositional semantics. This allows us to add fixpoints freely in the usual way (call the result IF-LFP, 'independence-friendly least fixpoint logic'), but at an (apparently) considerable price: the interpretation of a formula with free variables is no longer just a set of value tuples, but a set of sets of tuples. For the formal semantics, recall the definition of the syntax in section 2.2 – which is, for simplicity, a fragment of the full logic. We will give, using Hodges' terminology, the semantics for this fragment.

Let $\phi(\vec{x})$ mean that $\vec{x}$ is all the free variables of ϕ without repetition. A *deal* for ϕ is an assignment of elements of a structure A to each variable in $\vec{x}$. The interpretation of a formula is the set of its *trumps*, defined as follows.

- If $P(\vec{x})$ is atomic, then a non-empty set X of deals is a trump iff every deal in X satisfies P.
- X is a trump for $(\phi \wedge \psi)(\vec{x})$ iff X is a trump for $\phi(\vec{x})$ and X is a trump for $\psi(\vec{x})$.
- X is a trump for $(\phi \vee \psi)(\vec{x})$ iff it is non-empty and there are trumps U of ϕ and V of ψ such that every deal in X belongs either to U or V.
- X is a trump for $\forall y.\,\psi(\vec{x}, y)$ iff the set $\{\, \vec{a}b \mid \vec{a} \in X, b \in A \,\}$ is a trump for ψ.
- The interesting case is the existential quantifier (and, in the presence of negation, the universal quantifier). Given $\phi(\vec{x}) = \exists y/W.\,\psi(\vec{x}, y)$ (where W is a subset of the variables $\vec{x}$), say that two deals are W-equivalent iff they agree on all variables not in W. Say that a non-empty set X of deals is a W-set if its members are pairwise W-equivalent. Then a set X of deals for ϕ is a trump iff: there is a trump U for ψ such that for every W-set $Y \subseteq X$ there is a b such that $\{\, \vec{a}b \mid \vec{a} \in Y \,\} \subseteq U$.

A trump for ϕ is essentially a set of winning positions for the model-checking game for ϕ, for a given *uniform* strategy, that is, a strategy where choices are uniform in the 'hidden' variables. The meaning $[\![\phi]\!]$ of a formula is then defined to be the set of its trumps.

It is easy to see that any subset of a trump is a trump. In the case of an ordinary first-order $\phi(\vec{x})$, the set of trumps of ϕ is just the power set of the set of tuples satisfying ϕ. To see how a more complex set of trumps emerges, consider the following formula, which has x free: $\exists y/\{x\}.\,x = y$. Any singleton set of deals is a trump, but no other set of deals is a trump. Thus we obtain that $\forall x.\,\exists y/\{x\}.\,x = y$ has no trumps (unless the domain has only one element).

Now consider adding fixpoints. In normal LFP, we form an inductive definition via a formula $\phi(x, X)$ with a relation variable X. In Hodges' semantics, the 'relation' variable should instead range over sets of (potential) trumps, and hence the fixpoint is taken over functionals $\wp(\wp(A)) \to \wp(\wp(A))$ rather than $\wp(A) \to \wp(A)$. For this, we require

Lemma 18 The operators $\vee$, $\wedge$, $\forall x.$, $\exists x/W.$ are monotone in the lattice of trump sets. □

Thus given a n-ary 'relation' variable X, and a formula $\phi(\vec{x}, X)$ with n free variables, we have in the usual way an operator on $\wp(\wp(A^n))$, and we can form the fixpoints $\mu X.\,\phi$

and $\nu X.\,\phi$. Note, however, that this is *not* the form of a general inductive definition over $\wp(A^n)$, since we do not have variables ranging directly over $\wp(A^n)$.

The usual machinery of ordinal approximants applies. However, as the fixpoints are over $\wp(\wp(A))$, the naive bound on the closure ordinal of a single fixpoint is now exponential in $|A|$. If A is countably infinite, this raises the possibility of $2^{\aleph_0}$-step approximation; and high expressive power.

Indeed, the question of the expressive power of this logic, on the structure of arithmetic, is most interesting. We know that ordinary induction on the integers (i.e. LFP) is a small fragment of Δ^1_2; but with IF-LFP, we are performing induction over the continuum, which is, unrestricted, extremely powerful: ${}^{\omega}\omega$ induction over a Π^1_1 formula gives Σ^1_2, and then induction over Δ^1_2 gives all the semihyperprojective sets [Cen76] (including, for example, the entire analytical hierarchy). In our framework, it is natural to conjecture that the expressive power of IF-LFP is much less than this (and in fact I conjecture it is in Δ^1_2), but at present I do not have the tools to analyse the power. The reasons for the conjectured weakness are that the form of induction is very restricted, as noted above.

We can, however, note the simple fact that this semantics agrees with the usual semantics for LFP when there is no independence:

Proposition 19 If an IF-LFP formula ϕ contains no slashes, then its denotation $[\![\phi]\!]^{\mathrm{IF-LFP}}$ is exactly the set of non-empty subsets of the ordinary LFP denotation $[\![\phi]\!]^{\mathrm{LFP}}$; and in particular, $\vec{x} \in [\![\phi]\!]^{\mathrm{LFP}} \Leftrightarrow \{\vec{x}\} \in [\![\phi]\!]^{\mathrm{IF-LFP}}$.

Proof. By induction on approximants and the structure of ϕ. □

5 Independence-friendly modal mu-calculus?

The fixpoint extension of IF-FOL is interesting in its own right, but it is also a tool for understanding more sophisticated Henkin modal logics than we have defined so far.

We have already used the notation $\otimes$ to describe the concurrent, independent combination of sequences of modalities. A natural extension is to take it as an operator on formulae instead, so one can write, for example, $([a](P_1 \wedge (\langle b\rangle * \wedge \langle c\rangle *))) \otimes ([\overline{a}](P_2 \wedge \langle - \rangle *))Q$, with the intended interpretation that the first component satisfies the first factor (with P_1 probably, but not necessarily, being a *local* proposition as in [HNW97]), making its existential choices independently; the second satisfies the second; and in the process, the two components synchronize on the a; and when they rendezvous at $*$, the system satisfies Q.

The simplest way to define such operators formally is to use a game semantics, in the style of Hintikka. Precisely, take the syntax of modal logic, in positive form, and add: an n-ary (for each n) operator $\otimes$ ("parallel composition"), an atomic formula $*$ ("end of parallel"), and a binary operator . (roughly, "sequential composition"). In our usual framework, only one level of $\otimes$ nesting is allowed, and $\otimes$ and $*$ may only occur on the left side of a "sequential composition".

Given a system (say 2-component, for simplicity) in our usual framework, a formula defines a game thus: the normal modalities are played as usual. At a formula $(\phi_1 \otimes \phi_2).\psi$,

the game splits into two concurrent parts: every modality move in one half must match a modality move in the other half, with the global system advancing as dictated by the synchronization algebra, but the choices in one half are independent of those in the other. When play reaches $*$ in both halves, play continues at ψ; if play stops because an atomic formula (other than $*$) is reached in either component, Eloise wins if the formula is locally satisfied (in both halves, in the event that both halves reach an internal atomic formula).

This game semantics, although not unnatural, is quite complex. It is an exercise in applying [Hod97] to see that:

Remark 20 Given appropriate predicates on systems in our framework, the above logic can be given a semantics using IF-FOL, with the IF-FOL game semantics corresponding naturally to the above modal game semantics. □

The issue of fixpoints then arises. Provided that fixpoint operators are restricted to 'system formulae', that is, are not allowed inside $\otimes$, there is no particular problem, and we obtain a logic that can express moderately complex properties concerning independence and synchronization, but which is still decidable.

However, it is natural to ask whether formula such as $((\mu Z.\ * \vee (P_1 \wedge [-]Z)) \otimes (\mu Z.\ * \vee (P_2 \wedge [-]Z))).\langle a \rangle \mathrm{tt}$ make sense. They appear to make good intuitive sense, when one thinks of μ as just meaning 'finite looping', but it is not apparent how to give a semantics to them within IF-LFP. Furthermore, an attempt to give a game semantics results in games where the two agents can proceed arbitrarily far without 'rendezvous'ing and acquiring knowledge of each other's state space. In general, such games, in which the lack of knowledge persists indefinitely, have undecidable outcome problems; for example, in [AHK97], a general game quantifier temporal logic is defined, which is undecidable on finite systems for that reason. Whether such formulae can be allowed with restrictions sufficient to retain decidability, is an issue for further investigation.

6 Summary and future work

In this paper, we have introduced the notion of 'independence-friendly' modal logic; we have laid some of the groundwork for a modal theory, and looked at the extension of the first-order theory by fixpoints. This opens up many avenues for further exploration, both in a 'computer science' setting and in a more 'logical' setting. For example:

- Does full $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$ with negation have a finite model property?
- What is the expressive power of IF-LFP on arithmetic?
- Does $\mu^{\mathrm{H}}_{\mathrm{M}}\mathrm{L}$ have a fixpoint alternation hierarchy on finite models? If so, how does this connect to the complexity?
- In this paper, we have used only a fragment of IF-FOL; in particular, we have not used independent disjunction (and conjunction), which is a subtle connector (see [Hod97]); and we have not addressed the issue of negation and duality. The latter question was traditionally ignored, by sticking to positive form; Hodges has provided an account of negation, and there is some work to do in applying it in our setting.

- What is the relation between the *logical* independence we have been studying, and the *model* independence in semantic accounts of true concurrency? There appear to be some links between logics of causality and locality, including history-sensitive logics such as the hereditary history-preserving logic of [NiC94], and these should be investigated.
- Independence-friendly logics are philosophically and mathematically interesting, but are they really useful? Putative applications such as distributed system design should be investigated, and it is hoped to pursue this in a future project.

Finally, I should like to thank the several colleagues with whom I have discussed the idea of $\overset{\mathrm{H}}{\underset{\mathrm{M}}{}}\mathrm{L}$, in particular Perdita Stevens and Juliana Küster Filipe. I also thank the referees for perceptive criticisms, which, regrettably, remain largely unanswered; a longer version of this paper will be found via my home page `http://www.dcs.ed.ac.uk/home/jcb/`.

References

[AHK97] R. Alur, T. Henzinger and O. Kupferman, Alternating-time temporal logic, in *Proc. 38th FOCS (1997)*, 100–109.

[And93] H. R. Andersen, *Verification of Temporal Properties of Concurrent Systems*, DAIMI PB – 445, Computer Science Dept, Aarhus University, 1993.

[Bar79] J. Barwise, On branching quantifiers in English, *J. Philos. Logic* **8** 47–80 (1979).

[BlG86] A. Blass and Y. Gurevich. Henkin quantifiers and complete problems. *Ann. Pure Appl. Logic* **32**:1, 1–16 (1986).

[Bra97] J. C. Bradfield, The modal mu-calculus alternation hierarchy is strict, *Theor. Comput. Sci.* **195** 133–153 (1997).

[Bra99a] J. C. Bradfield, Fixpoint alternation and the game quantifier, *Proc. CSL '99*, LNCS **1683** 350–361 (1999).

[Bra99b] J. C. Bradfield, Fixpoint alternation: arithmetic, transition systems, and the binary tree, *Theoret. Informatics Appl.* **33** 341–356 (1999).

[Cen76] D. Cenzer, Monotone inductive definitions over the continuum, *J. Symbolic Logic* **41**:1 188–198 (1976).

[End70] H. B. Enderton, Finite partially ordered quantifiers, *Z. für Math. Logik u. Grundl. Math.* bf 16 393–397 (1970).

[Got97] G. Gottlob, Relativized logspace and generalized quantifiers over finite ordered structures. *J. Symbolic Logic* **62**:2, 545–574 (1997).

[Hen61] L. Henkin, Some remarks on infinitely long formulas, *Infinitistic Methods*, Pergamon Press, Oxford and PAN, Warsaw, 167–183 (1961).

[HiS96] J. Hintikka and G. Sandu, A revolution in logic?, *Nordic J. Philos. Logic* **1**(2) 169–183 (1996).

[Hod97] W. Hodges, Compositional semantics for a language of imperfect information, *Int. J. IGPL* **5**(4), 539–563.

[HNW97] M. Huhn, P. Niebert and F. Wallner, Verification on local states, *Proc. TACAS '98*, LNCS **1384** 36–51.

[Lub93] R. S. Lubarsky, μ-definable sets of integers, *J. Symbolic Logic* **58** (1993) 291–313.

[McM92] K. L. McMillan, Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits, *Proc. CAV'92*, 161–171 (1992).

[NiC94] M. Nielsen and C. Clausen, Bisimulations, games and logic. BRICS report RS-94-6. BRICS, Aarhus (1994).

[Pen92] W. Penczek, On undecidability of propositional temporal logics on trace systems, *Inf. Proc. Let.* **43**(3) 147–153 (1992).

[PK95] W. Penczek and R. Kuiper, Traces and Logic, in: Diekert and Rozenberg, ed. *The Book of Traces*, World Scientific (1995).

[San93] G. Sandu, On the logic of information independence and its applications, *J. Philos. Logic* **22** 361-372 (1993).

[StW90] C. P. Stirling and D. J. Walker, A general tableau technique for verifying temporal properties of concurrent programs. *Proc. Int. BCS–FACS Workshop on Semantics for Concurrency* Workshops in Computing (Springer–Verlag, Berlin, 1990) (1–15).

[Val91] A. Valmari, A stubborn attack on state explosion. *Proc. CAV '90*, DIMACS vol. 3, 25–42 (1991).

[Wal70] W. J. Walkoe, Jr, Finite partially-ordered quantification. *J. Symbolic Logic* **35** 535–555 (1970).