On independence-friendly fixpoint logics

J. C. Bradfield
University of Edinburgh

Résumé : Nous introduisons une extension aux points fixes de la logique IF (faite pour l'indépendance) de Hintikka et Sandu. Nous donnons des résultats sur sa complexité et son pouvoir expressif. Nous la relions aux jeux de parité à information imparfaite, et nous montrons une application à la définition d'un mu-calcul modal fait pour l'indépendance.

Abstract: We introduce a fixpoint extension of Hintikka and Sandu's IF (independence-friendly) logic. We obtain some results on its complexity and expressive power. We relate it to parity games of imperfect information, and show its application to defining independence-friendly modal mu-calculi.

Philosophia Scientiæ, 8 (2), 2004, 125–144.

1 Introduction

The topic of this issue, independence-friendly logic, is a logic introduced by Sandu and Hintikka [Sandu 1993, Hintikka & Sandu 1996] which gives an alternative account of branching quantifiers (Henkin quantifiers) in terms of games of imperfect information. It allows the expression of quantifiers where the choice must be independent of specified earlier choices; it has existential second-order power. As well as its interest for philosophical and mathematical logicians, it also has some natural resonances with the theory of concurrency in computer science. Specifically, in earlier work, we have argued that the modal analogues of IF logic have a role to play in concurrency theory, partly inspired by previous work by Alur, Henzinger and Kupferman [Alur et al. 1997], in which a temporal logic using imperfect information is studied. (See [Bradfield 2000] and [Bradfield & Fröschle 2002] for discussions of this role and its relation to other work in concurrency theory.)

Given a first-order logic, or a logic like IF that is supposed to *look* first-order (even though it isn't), it is natural for modal logicians of a certain bent to want to add fixpoint operators. One motivation is just the mathematical interest of studying inductive definability in many contexts; a more computer-science-based motivation is the desire to be able to produce an IF analogue of Kozen's [Kozen 1983] modal mu-calculus, a popular and interesting temporal logic – see [Bradfield & Stirling 2001] for an introductory survey of modal mu-calculus.

In [Bradfield 2000], we asserted that using the semantics given to IF by Hodges [Hodges 1997], it was possible to define an IF fixpoint logic. In this article, we give a detailed definition of IF least fixpoint logic (which, typically of IF logics, is a little more subtle than one first thinks), and then study it.

In section 2, we deal with the preliminaries, the existing syntax and semantics of IF logic. Sections 3, 4 and 5 are the main part of the paper; in section 3 we give the detailed definitions of IF fixpoint logic and its semantics; in section 4 we give a couple of interesting examples; and in section 5 we establish some partial results on complexity and expressive power. Then in section 6 we return to the game-theoretic roots of IF by giving a suitable notion of parity game of imperfect information, which gives an alternative semantics for IF fixpoint logic. Finally, in section 7 we briefly sketch the application to IF modal mu-calculus that was one of the original motivations for looking at IF with fixpoints.

This article is a revised version of [Bradfield 2003], presented at *Computer Science Logic 2003* in Vienna. I thank the referee for some helpful

suggestions.

2 IF-FOL syntax and semantics

First of all, we state one important **notational convention**: we take the scope of all quantifiers and fixpoint operators to extend as far to the right as possible.

For the purposes of this article, we will use only a sublanguage of IF-FOL (sometimes just IF for short). The full languages advocated by Hintikka and analysed by Hodges and others include the possibility of conjunctions and disjunctions that are independent of previous quantifiers. These operators do not introduce inherently new problems, but they do introduce some additional complexity (and space) in defining the semantics. We will therefore ignore them, and consider only the independent quantifiers; the interested reader can use [Hodges 1997] to put back the independent junctions.

One of the more tedious features of IF-FOL is the need to be more pedantic than usual in keeping track of free variables etc., as not all the things one takes for granted in usual logic are true in IF-FOL. When introducing fixpoint operators, even more care is needed, and we shall therefore give the semantics even more pedantically than Hodges did.

Definition 1. *Assume the usual FOL set of proposition (P, Q etc.), relation (R, S etc.), function (f, g etc.) and constant (a, b etc.) symbols, with given arities. Assume also the usual* variables v, x etc. *We write $\vec{x}, \vec{v}$ etc. for tuples of variables, and similarly for tuples of other objects; we use concatenation of symbols to denote concatenation of tuples with tuples or objects.*

For formulae ϕ and terms t, the (meta-level) notations $\phi[\vec{x}]$ and $t[\vec{x}]$ mean that the free variables of ϕ or t are included in the variables $\vec{x}$, without repetition.[1]

The terms of IF-FOL are as usual constructed from variables, constants and function symbols. The free variables of a term are as usual; the free variables of a tuple of terms are the union of the free variables of the terms.

We assume equality $=$ is in the language, and atomic formulae are defined as usual. The free variables of the formula $R(\vec{t})$ are those of $\vec{t}$.

[1][Hodges 1997] writes $\phi(\vec{x})$, but we wish to distinguish the meta-notation for free variables from the object-level syntax for atomic formulae and the meta-notation for assigning values to variables.

The compound formulae are given as follows:

Conjunction and disjunction. *If* $\phi[\vec{x}]$ *and* $\psi[\vec{y}]$ *are formulae, then* $(\phi \vee \psi)[\vec{z}]$ *and* $(\phi \wedge \psi)[\vec{z}]$ *are formulae, where* $\vec{z}$ *is the union of* $\vec{x}$ *and* $\vec{y}$*.*

Quantifiers. *If* $\phi[\vec{y}, x]$ *is a formula,* x *a variable, and* W *a finite set of variables, then* $(\forall x/W.\, \phi)[\vec{y}]$ *and* $(\exists x/W.\, \phi)[\vec{y}]$ *are formulae. If* W *is empty, we write just* $\forall x.\, \phi$ *and* $\exists x.\, \phi$*.*

Game negation. *If* $\phi[\vec{x}]$ *is a formula, so is* $(\sim\phi)[\vec{x}]$*.*

Flattening. *If* $\phi[\vec{x}]$ *is a formula, so is* $(\downarrow \phi)[\vec{x}]$*.*

(Negation. $\neg\phi$ *is an abbreviation for* $\sim \downarrow \phi$*.***)** $\lhd$

Definition 2. *IF-FOL*$^+$ *is the logic in which* $\sim$*,* $\downarrow$ *and* $\neg$ *are applied only to atomic formulae.* $\lhd$

In the independent quantifiers the intention is that the choices of the player are independent of the values of the variables in the set W. In terms of imperfect information, the player does not know the values of the W-variables at the choice point. Hence the Henkin quantifier $\begin{smallmatrix}\forall x \, \exists y \\ \forall u \, \exists v\end{smallmatrix}$ can be written as $\forall x/\varnothing.\, \exists y/\varnothing.\, \forall u/\{x, y\}.\, \exists v/\{x, y\}$. If one then plays the usual model-checking game with this additional condition, which can be formalized by requiring strategies to be uniform in the 'unknown' variables, one gets a game semantics which characterizes 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 untrue. For example, $\begin{smallmatrix}\forall x \\ \exists y\end{smallmatrix}.x = y$ (or $\forall x.\, \exists y/\{x\}.\, x = y$) is untrue in any structure with more than one element, but Abelard has no winning strategy.

The trump semantics of Hodges [Hodges 1997], with variants by others, gives a Tarski-style semantics for this logic, equivalent to the imperfect information game semantics given by Hintikka and Sandu. The semantics is as follows:

Definition 3. *Let a structure* A *be given, with constants, propositions and relations interpreted in the usual way. A* deal $\vec{a}$ *for* $\phi[\vec{x}]$ *or* $\vec{t}[\vec{x}]$ *is an assignment of an element of* A *to each variable in* $\vec{x}$*. Given a deal* $\vec{a}$ *for a tuple of terms* $\vec{t}[\vec{x}]$*, let* $\vec{t}(\vec{a})$ *denote the tuple of elements obtained by evaluating the terms under the deal* $\vec{a}$*.*

If $\phi[\vec{x}]$ *is a formula and* W *is a subset of the variables in* $\vec{x}$*, two deals* $\vec{a}$ *and* $\vec{b}$ *for* ϕ *are* $\simeq_W$-equivalent *(*$\vec{a} \simeq_W \vec{b}$*) iff they agree on the variables not in* W*. A* $\simeq_W$-set *is a non-empty set of pairwise* $\simeq_W$*-equivalent deals.*

The interpretation $[\![\phi]\!]$ *of a formula is a pair* (T, C) *where* T *is the set of* trumps*, and* C *is the set of* cotrumps*.*

- *If* $(R(\vec{t}))[\vec{x}]$ *is atomic, then a non-empty set* D *of deals is a trump iff* $\vec{t}(\vec{a}) \in R$ *for every* $\vec{a} \in D$*;* D *is a cotrump iff it is non-empty and* $\vec{t}(\vec{a}) \notin R$ *for every* $\vec{a} \in D$*.*

- D *is a trump for* $(\phi \wedge \psi)[\vec{x}]$ *iff* D *is a trump for* $\phi[\vec{x}]$ *and* D *is a trump for* $\psi[\vec{x}]$*;* D *is a cotrump iff there are cotrumps* E, F *for* ϕ, ψ *such that every deal in* D *is an element of either* E *or* F*.*

- D *is a trump for* $(\phi \vee \psi)[\vec{x}]$ *iff it is non-empty and there are trumps* E *of* ϕ *and* F *of* ψ *such that every deal in* D *belongs either to* E *or* F*;* D *is a cotrump iff it is a cotrump for both* ϕ *and* ψ*.*

- D *is a trump for* $(\forall y/W.\,\psi)[\vec{x}]$ *iff the set* $\{\, \vec{a}b \mid \vec{a} \in D, b \in A \,\}$ *is a trump for* $\psi[\vec{x}, y]$*.* D *is a cotrump iff it is non-empty and there is a cotrump* E *for* $\psi[\vec{x}, y]$ *such that for every* $\simeq_W$*-set* $F \subseteq D$ *there is a* b *such that* $\{\, \vec{a}b \mid \vec{a} \in F \,\} \subseteq E$*.*

- D *is a trump for* $(\exists y/W.\,\psi)[\vec{x}]$ *iff there is a trump* E *for* $\psi[\vec{x}, y]$ *such that for every* $\simeq_W$*-set* $F \subseteq D$ *there is a* b *such that* $\{\, \vec{a}b \mid \vec{a} \in F \,\} \subseteq E$*;* D *is a cotrump iff the set* $\{\, \vec{a}b \mid \vec{a} \in D, b \in A \,\}$ *is a cotrump for* $\psi[\vec{x}, y]$*.*

- D *is a trump for* $\sim\phi$ *iff* D *is a cotrump for* ϕ*;* D *is a cotrump for* $\sim\phi$ *iff it is a trump for* ϕ*.*

- D *is a trump (cotrump) for* $\downarrow \phi$ *iff* D *is a non-empty set of members (non-members) of trumps of* ϕ*.* $\triangleleft$

A sentence is true in the usual sense if $\{\langle\rangle\} \in T$ (the empty deal is a trump set), and false in the usual sense if $\{\langle\rangle\} \in C$; this corresponds to Eloise or Abelard having a uniform winning strategy. Otherwise, it is undetermined.

Note that the game negation $\sim$ provides the usual de Morgan dualities.

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 most intricate part of the above definition is the clause for $\exists y/W.\,\psi$: it says that a trump for $\exists y/W.\,\psi$ is got by adding a witness for y, uniform in the W-variables, to trumps for ψ.

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).

The following definition is for later convenience: a set T of sets of deals is *well-dealt* if for every $D \in T$, D is non-empty and $D' \in T$ for every non-empty $D' \subseteq D$. A formula has well-dealt semantics (T, C) if T and C are well-dealt; the above semantics ensures that all IF-FOL formulae have well-dealt semantics.

[Hodges 1997] shows that every well-dealt set is the semantics of some IF formula (given suitable atomic relations), giving us

Proposition 4. *On a structure A with n elements, IF formulae of length m require space exponential in n^m to represent their semantics.*

Proof. The set of tuples for m free variables has n^m elements; Given a k element set, there are 2^k subsets, but not all sets of subsets are well-dealt; however, there are about $2^k/\sqrt{k}$ sets of size $k/2$, and hence at least $2^{2^k/\sqrt{k}}$ well-dealt sets of subsets. (Cameron and Hodges [Cameron & Hodges 2001] look in more detail at the combinatorics of trumps.) $\square$

We can record the easy loose upper bounds on the time complexity of IF-FOL operations:

Proposition 5. *In a structure A of size n, the trump components of the IF operators can be calculated in the following times on formulae with m free variables, where $k = n^m$: $\vee$ and $\wedge$ in $2^{k+1} \cdot k^2$; $\forall x$ in $2^k \cdot k^3 n$; $\exists x/W$ in $2^{k+k \lg n}$.*

Proof. A crude analysis of the cost of computing the trump semantics more or less directly from the definitions. Note that the computation for $\exists$ has further exponential factors above the 2^k from the number of possible trumps, effectively due to the computation of choice functions. $\square$

In the case of IF, these exponential upper bounds are much worse than is really required for determining whether a deal satisfies (i.e. is a singleton trump for) an IF formula, since IF expressible properties are in NP (because we can guess values for choice functions).

3 Adding fixpoint operators.

The prime motivation for considering fixpoint extensions is in the modal setting, where it is a standard way to produce temporal logics from modal logics. However, fixpoint extensions to IF logics raise a number of issues, and it is useful to recall briefly the first-order case.

In the classical settings, fixpoint operators are added to allow sets or relations to be inductively defined by formulae: $\mu(x, X).\phi(x, X)$, where X is a set variable, is the least set A such that $A = \{ x \mid \phi(x, A) \}$, and the syntax of formulae is extended to allow terms of the form $t \in X$ or $t \in \mu(x, X).\phi(x, X)$ (among set theorists) or $X(t)$ and $(\mu(x, X).\phi(x, X))(t)$ (among finite model theorists).

In applying this directly to IF-FOL, there is the obvious problem that we no longer have a simple notion of an element satisfying a formula, so the usual definition no longer type-checks. There are two possible approaches, depending on how one views the use of fixpoint terms. If one takes the view that their purpose is to define sets, and the logic is a means to this end, then it is natural to retain the use of set variables, and work out how to make $\phi(x, X)$ reduce to a boolean. On the other hand, if one views fixpoint operators as a means of introducing recursion into the logical formulae, it is more natural to decide that fixpoint terms should have the same semantics as other formulae, namely sets of trumps, and that therefore the variables X range over trump sets rather than sets. We then have to decide the meaning of $X(t)$. This is the approach we suggested in [Bradfield 2000], and will now pursue.

Definition 6. *IF-LFP extends the syntax of IF-FOL as follows:*

- *There is a set* $\mathrm{Var} = \{X, Y, \ldots\}$ *of fixpoint variables. Each variable* X *has an arity* $(\mathrm{ar}_1(X), \mathrm{ar}_2(X))$*;* $\mathrm{ar}_1(X)$ *is the arity of the fixpoint, and* $\mathrm{ar}_2(X)$ *is the number of free parameters of the fixpoint.*

- *If* X *is a fixpoint variable, and* $\vec{t}$ *an* $\mathrm{ar}_1(X)$*-vector of terms then* $X(\vec{t})$ *is a formula.*

- *The notation* $\phi(X)$ *indicates that* X *is among the free fixpoint variables of* ϕ*. If* $\phi(X)[\vec{x}, \vec{z}]$ *is a formula with* $\mathrm{ar}_1(X)$ *free individual variables* $\vec{x}$ *and* $\mathrm{ar}_2(X)$ *free individual variables* $\vec{z}$*, and* $\vec{t}$ *is a sequence of* $\mathrm{ar}_1(X)$ *terms with free variables* $\vec{y}$*, then* $(\mu(X, \vec{x}).\phi)(\vec{t})[\vec{z}, \vec{y}]$ *is a formula;* **provided that** ϕ *is IF-FOL*$^+$*.*

- *Similarly for* $\nu(X, \vec{x}).\phi$*.* $\triangleleft$

The process of extending the trump semantics to fixpoint formulae is not entirely straightforward. First we define valuations for free fixpoint variables.

Definition 7. *A fixpoint valuation $\mathscr{V}$ maps each fixpoint variable X to a pair*

$$(\mathscr{V}_T(X), \mathscr{V}_C(X)) \in (\wp(\wp(A^{\mathrm{ar}_1(X)+\mathrm{ar}_2(X)})))^2.$$

Let D be a non-empty set of deals for $X(\vec{t})[\vec{x},\vec{z},\vec{y}]$, where $\vec{y}$ are the free variables of $\vec{t}$ not already among $\vec{x},\vec{z}$. A deal $d = \vec{a}\vec{c}\vec{b} \in D$, where $\vec{a},\vec{c},\vec{b}$ are the deals for $\vec{x},\vec{z},\vec{y}$ respectively, determines a deal $d' = \vec{t}(d)\vec{c}$ for $X[\vec{x},\vec{z}]$. Let $D' = \{\, d' \mid d \in D \,\}$. D is a trump for $X(\vec{t})$ iff $D' \in \mathscr{V}_T(X)$; it is a cotrump iff $D' \in \mathscr{V}_C X$. $\lhd$

Then we define a suitable complete partial order on denotations:

Definition 8. *If (T_1, C_1) and (T_2, C_2) are elements of $(\wp(\wp(A^n)))^2$, define $(T_1, C_1) \preceq (T_2, C_2)$ iff $T_1 \subseteq T_2$ and $C_1 \supseteq C_2$.* $\lhd$

Lemma 9. *If $\phi(X)[\vec{x},\vec{z}]$ is an IF-FOL$^+$ formula and $\mathscr{V}$ is a fixpoint valuation, the map on $(\wp(\wp(A^{\mathrm{ar}_1(X)+\mathrm{ar}_2(X)}))^2$ given by*

$$(T, C) \mapsto [\![\phi]\!]_{\mathscr{V}[X:=(T,C)]}$$

is monotone with respect to $\preceq$; hence it has least and greatest fixpoints, with ordinal approximants defined in the usual way.

Definition 10. *$[\![\mu(X,x).\phi(X)[\vec{x},\vec{z}]]\!]$ is the least fixpoint of the map just defined; $[\![\nu(X,x).\phi(x)[\vec{x},\vec{z}]]\!]$ is the greatest fixpoint. $\mu^\zeta(X,x).\phi$ means the ζth approximant of $\mu(X,x)$, defined by*

$$[\![\mu^\zeta(X,x).\phi]\!] = [\![\phi(\bigcup_{\zeta'<\zeta} \mu^{\zeta'}(X,x).\phi)]\!];$$

we may also write X^ζ or ϕ^ζ when convenient. $\lhd$

The following lemma records the usual basic properties (which have to be checked again in this setting), and one new basic property, particular to the IF case.

Lemma 11.

1. *The trump and cotrump components of $[\![\mu(X,x).\phi]\!]$ are well-dealt.*
2. *If Y is free in ϕ, then $[\![\mu(X,x).\phi]\!]$ is monotone in Y; hence the definition extends to further fixpoints in the usual way, as does this lemma.*

3. *μ and ν are dual: T is a trump for $\mu(X, x).\phi(X)$ iff it is a cotrump for $\nu(X, x).{\sim}\phi({\sim}X)$ (with the outer negation pushed in by duality).*

Proof. (1) by induction on approximants; (2) as usual; (3) from definitions. □

A distinctive feature of the definition, compared to the normal LFP definition, is the way that free variables are explicitly mentioned. Normally, one can fix values for the free variables, and then compute the fixpoint, but because of independent quantification this is not possible in the IF setting. For example, consider the formula fragment

$$\forall z.\ \ldots \mu(X, x). \ldots \vee \exists y/\{z\}.\, X(y)$$

The independent choice of y means that the trumps for the fixpoint depend on the possible deals for z, not just a single deal.

Another point is that the trump set of a least fixpoint is the union of the trump sets of its approximants; but the interpretation of logical disjunction is not union of trump sets, but union of trumps (applied pointwise to the trump sets). Thus the usual view of a least fixpoint as a transfinite disjunction is not valid in general. The following explains why, despite this, the IF-LFP semantics is consistent with classical LFP semantics.

Proposition 12. *Call a set T of trumps or cotrumps* full *iff it is the set of non-empty subsets of $\bigcup T$. Call a formula ϕ of IF-LFP* classical *iff it is in IF-FOL$^+$ and it contains no independent quantification (i.e. all quantifiers are $\exists x/\varnothing$ and $\forall x/\varnothing$). Then*

1. *if ϕ is fixpoint free, then $[\![\phi]\!]_{\mathrm{IF}} = (T, C)$ is full, $\bigcup T = [\![\phi]\!]_{\mathrm{FO}}$, and $\bigcup C = [\![\neg\phi]\!]_{\mathrm{FO}}$;*
2. *if T_ζ is a (transfinite) sequence of full well-dealt deal sets, then $\bigcup_\zeta T_\zeta$ and $\bigcap_\zeta T_\zeta$ are full well-dealt sets;*
3. *hence (1) is true for any classical IF-LFP formula.*

4 Examples of IF-LFP

IF logic is not entirely easy to understand and mu-calculi are also traditionally hard to understand, so we now consider some examples that demonstrate interesting features of the combination. For convenience, we introduce the abbrevation $\phi \Rightarrow \psi$ for $\psi \vee \sim\phi$ provided that ϕ is atomic.

Let $G = (V, E)$ be a directed graph. The usual LFP formula

$$R(y,z) \stackrel{\text{def}}{=} (\mu(X,x).z = x \vee \exists w.\, E(x,w) \wedge X(w))(y)$$

asserts that the vertex z is reachable from y. Hence the formula

$$\forall y.\, \forall z.\, R(y,z)$$

asserts that G is strongly connected. Now consider the IF-LFP formula

$$\forall y.\, \forall z.\, (\mu(X,x).z = x \vee \exists w/\{y,z\}.\, E(x,w) \wedge X(w))(y).$$

At first sight, one might think this asserts not only that every z is reachable from every y, but that the path taken is independent of the choice of y and z. This is true exactly if G has a directed Hamiltonian cycle, a much harder property than being strongly connected.

Of course, the formula does not mean this, because the variable w is fresh each time the fixpoint is unfolded. In the trump semantics, the denotation of the fixpoint will include all the possible choice functions at each step, and hence all possible combinations of choice functions. Thus the formula reduces to strong connectivity.

It may be useful to look at the approximants of this formula in a little more detail, to get some intuitions about the trump semantics. Considering just

$$H \stackrel{\text{def}}{=} (\mu(X,x).z = x \vee \exists w/\{y,z\}.\, E(x,w) \wedge X(w))[x,y,z],$$

we see that in computing each approximant, the calculation of $[\![\exists w/\{y,z\}.\, \ldots]\!]$ involves generating a trump for every possible value of a choice function $f\colon x \mapsto w$. This is a feature of the original trump semantics, and can be understood by viewing it as a second-order semantics: just as the compositional Tarskian semantics of $\exists x.\, \phi(x)$ involves computing all the witnesses for $\phi(x)$, so computing the trumps of $\exists x/\{y\}.\, \phi$ involves computing all the Skolem functions; and unlike the first-order case, it is necessary to work with functions (as IF can express existential second-order logic). Consequently, the nth approximant includes all states such that $x \to f_1(x) \to f_2 f_1(x) \to \ldots \to f_n \ldots f_1(x) = z$ for any sequence of successor-choosing functions f_i. Thus we see that the cumulative effect is the same as for a normal $\exists w$, and the independent choice has indeed not bought us anything.

It is, however, possible[2] to produce a slightly more involved formula expressing the Hamiltonian cycle property in this inductively defined way, by using the standard trick for expressing functions in Henkin quantifier logics. We replace the formula H by

$$\forall s.\, \exists t/\{y,z\}.\, E(s,t) \wedge \mu(X,x).x = z \vee \forall u.\, \exists v/\{x,y,z,s,t\}.$$
$$(s = u \Rightarrow t = v) \wedge (x = u \Rightarrow X(v)).$$

This works because the actual function f selecting a successor for every node is made outside the fixpoint by $\forall s.\, \exists t/\{y,z\}.\, E(s,t) \wedge \ldots$; then inside the fixpoint, a new choice function g is made so that $X(g(x))$, and g is constrained to be the same as f by the clause $(s = u \Rightarrow t = v)$. (The reader who is not familiar with the IF/Henkin to existential second-order translation might wish to ponder why $\forall s.\, \exists t/\{y,z\}.\, E(s,t) \wedge \mu(X,x).x = z \vee (x = s \Rightarrow X(t))$ does not work.)

5 Complexity and expressive power of IF-LFP

The above examples have shown IF-LFP being used to express relatively simple NP properties. Since, as remarked, it is well known that Henkin quantifiers and IF logic express just the NP properties (Henkin logic is equi-expressive with Σ^1_1, and Σ^1_1 captures NP on finite models), and since it is also known [Gottlob 1997] that LFP plus Henkin quantifiers express P^{NP}, one might imagine that IF-LFP (which is not closed under classical negation) also expresses only NP properties, or at worst some subset of P^{NP}. This is not the case; adding fixpoints to the IF formulation gives a more significant increase in expressive power.

Firstly, we note that the approximant semantics of fixpoints gives the usual behaviour in simple upper bounds:

Proposition 13. *If $\phi(X)[x, z_1, \ldots, z_m]$ is an IF-FOL$^+$ formula, then in a structure of size n, the approximants of $\mu(X,x).\phi$ close after at most 2^{n^m} steps. Hence in an IF-LFP formula with d alternating fixpoints and m variables, 2^{dn^m} evaluations of IF formulae are required. If the formula size is l, this gives a total cost of $2^{dn^m} \cdot l \cdot 2^{n^m(1+\lg n)} = l \cdot 2^{n^m(1+d+\lg n)}$.*

[2]Since IF logic is equi-expressive with Henkin quantified logic, it is also equi-expressive with existential second-order logic, and so can express ‘Hamiltonian cycle’ without using fixpoints. Thus we are not, in this example, adding technical expressive power. However, the pure IF definition is quite complex, as it involves defining a binary relation coded via functions; so we are adding expressive convenience.

Observe, however, that the contribution from fixpoint alternation is small compared to the cost of computing independent existential quantifiers.

Despite the relative weakness of adding fixpoints, they do in some sense release the power of independent quantification. This is shown by the following theorem.

Theorem 14. *There is an IF-LFP sentence (with one least fixpoint) which is EXPTIME-hard to evaluate.*

Proof. We give a reduction from the EXPTIME-complete problem of determining whether Player 1 has a winning strategy for the game of generalized chess.

A structure for a generalized chess game between 1 and 2 of order n comprises a board R with n^2 (or any other fixed polynomial) squares r and a set P of n (or any other fixed polynomial) pieces p. A position of the game is a function $\pi\colon P \to R$. There may be some relations on P and R in the signature. The game is defined by three first-order formulae with parameter π: a formula $\phi_I(\pi)$ true only of the initial position, a formula $\phi_W(\pi)$ which is true if player 1 has won at π, and a formula $\phi_M(\pi, i, p, r)$ which is true if moving piece p to square r is a legal move for player i from position π. (Without loss of generality, we assume that a move consists of moving exactly one piece. We also assume that ϕ_W includes those positions where player 2 is due to move but cannot.)

Given a position π and a move p, r, the 'next position' formula $N(\pi, p, r, \pi')$ is defined to be $\forall p'.\, (p' = p \wedge \pi'(p) = r) \vee \pi'(p') = \pi(p')$ so that π' is the position resulting from the move.

The set X of winning positions (i.e. from which 1 can force a win) for 1 can then be inductively defined by the type 3 functional

$$
\begin{aligned}
&F(X, \pi) \;\Leftrightarrow\; \Phi \stackrel{\mathrm{def}}{=} \\
&\phi_W(\pi) \vee ((\forall p, r.\, \phi_M(\pi, 2, p, r) \Rightarrow \exists \pi'.\, N(\pi, p, r, \pi') \wedge X(\pi')) \\
&\wedge\, (\exists p, r.\, \phi_M(\pi, 1, p, r) \wedge \exists \pi'.\, N(\pi, p, r, \pi') \wedge X(\pi'))).
\end{aligned}
$$

We now show how to express this inductive definition in IF-LFP. Part of the coding is the well-known [Walkoe 1970, Enderton 1970] expression of existential second-order logic in IF or Henkin logic, which we have already seen in the Hamiltonian cycle example. The general technique is thus: assume given an ESO formula $\exists f.\, \psi$. Let $Q_1(f(\tau_1)), \ldots, Q_n(f(\tau_n))$ be the instances in ψ of applications of f occurring in atoms Q_i. Then the translation is

$\forall x^f.\, \exists y^f.\, \forall x_1^f.\, \exists y_1^f/\{x^f, y^f\}.\, \ldots \forall x_n^f.$

$\exists y_n^f/\{x^f, x_1^f, \ldots, x_{n-1}^f, y^f, y_1^f, \ldots, y_{n-1}^f,\}.\ \bigwedge_i (x_i^f = x^f \Rightarrow y_i^f = y^f) \wedge \hat{\psi}$,

where $\hat{\psi}$ is obtained from ψ by replacing $Q_i(f(\tau_i))$ with $x_i^f = \tau_i \Rightarrow Q_i(y_i^f)$.

The second part is passing a function through a fixpoint. This is fairly simple to do: one just passes the domain and codomain as normal parameters, and relies on the quantification outside forcing them to represent a function. In this case, the classical type 2 relation $X(\pi)$ is replaced by a binary IF type 1 relation $Y(x^\pi, y^\pi)$, so that the classical $\exists \pi.\, (\mu X.\varPhi)(\pi)$ becomes $\forall x^\pi.\, \exists y^\pi.\, (\mu Y.\hat{\varPhi})(x^\pi, y^\pi)$, where $\hat{\varPhi}$ is obtained from $\varPhi$ by applying the ESO–IF translation using x^π, y^π for π (etc.) and replacing $X(\pi')$ with $Y(x^{\pi'}, y^{\pi'})$.

We now show by an inductive argument on ζ that a function π satisfies $(\mu^\zeta X.\varPhi)(\pi)$ iff the corresponding functional deal for x^π, y^π satisfies $(\mu^\zeta Y.\hat{\varPhi})(x^\pi, y^\pi)$. The base case is trivial. Now suppose that the lemma holds for all $\zeta' < \zeta$. Then by definition of approximants there is π such that $(\mu^\zeta X.\varPhi)(\pi)$ iff there is π such that

$$\begin{aligned}\phi_W(\pi) \quad \vee \quad & ((\forall p, r.\, \phi_M(\pi, 2, p, r) \Rightarrow \exists \pi'.\, N(\pi, p, r, \pi') \wedge (\bigcup_{\zeta' < \zeta} X^{\zeta'})(\pi')) \\ & \wedge (\exists p, r.\, \phi_M(\pi, 1, p, r) \wedge \exists \pi'.\, N(\pi, p, r, \pi') \wedge (\bigcup_{\zeta' < \zeta} X^{\zeta'})(\pi'))).\end{aligned}$$

If π satisfies ϕ_W then (x^π, y^π) satisfies $\hat{\phi}_W$ and conversely. If π satisfies $\forall p, r.\, \phi_M(\pi, 2, p, r) \Rightarrow \exists \pi'.\, N(\pi, p, r, \pi') \wedge (\bigcup_{\zeta' < \zeta} X^{\zeta'})(\pi')$, then for those p, r such that $\phi_M(\pi, 2, p, r)$, there is (since N gives π' as a function of π, p, r) a unique π' satisfying the consequent. By induction, the corresponding functional deal for $(x^{\pi'}, y^{\pi'})$ satisfies $(\bigcup_{\zeta' < \zeta} Y)(x^{\pi'}, y^{\pi'})$; and thence $\forall x^{\pi'}.\, \exists y^{\pi'}/\ldots.\, \hat{N}(\ldots) \wedge (\bigcup_{\zeta' < \zeta} Y)(x^{\pi'}, y^{\pi'})$ holds, and thence the entire translation; and conversely. Similarly for the existential clause.

Finally, if we wish to determine whether the initial position is winning for 1, we evaluate $\exists \pi.\, \phi_I(\pi) \wedge (\mu X.\varPhi)(\pi)$

(We should note that we have extended the IF abbrevation $\phi \Rightarrow \psi$ to the case where ϕ is classical, not just atomic. This is acceptable because game negation coincides with classical negation for classical formulae.)

□

The above argument was applied to the case of finite structures, but there is nothing in it that depends on finiteness. We can therefore obtain the following theorem, which refutes our conjecture in [Bradfield 2000] that a fixpoint extension of IF would be within Δ^1_2.

Theorem 15. *Let $F(X, \alpha)$ be a positive Σ^1_1 type 3 functional in the language of arithmetic. Then a set of integers definable from the set of reals inductively defined by F can be expressed in IF-LFP. It follows that that IF-LFP (even with just one fixpoint) over the natural numbers can express Σ^1_2 properties.*

Proof. F is defined by a Σ^1_1 formula $\phi(X, \alpha)$. Use the technique of the previous proof to express $\exists\alpha.\,(\mu X.\phi)(\alpha) \wedge \psi(\alpha, n)$, where ψ is first-order. Cenzer [Cenzer 1976] showed that any Σ^1_2 set of reals is the closure of a Σ^1_1 positive inductive definition over the reals. Since if α is a Σ^1_2 real, the set $\{\alpha\}$ is also Σ^1_2, we also have the stated consequence. □

Cenzer's results also allow us to obtain an improvement (for those who don't believe CH) on the closure ordinal for a single IF fixpoint over ω. The usual cardinality argument for fixpoints tells us merely that an IF fixpoint over ω must close by $2^{\aleph_0}$. The improvement is

Theorem 16. *If $\phi(X)$ is an IF-FOL$^+$ formula (i.e. with $\sim$ and $\downarrow$ applied only to atoms), then $\mu X.\phi$ has closure ordinal $\leq \aleph_1$.*

Proof. Seen as operations on $\wp(2^\omega)$, the semantics of the IF boolean operators and quantifiers are Σ^1_1. (This is not immediately apparent from the definitions as presented above, but a small amount of rearrangement reveals it.) Cenzer showed that the closure ordinal of a Σ^1_1 monotone inductive definition over the reals is $\leq \aleph_1$. □

It remains to investigate lower bounds on the complexity of multiple IF fixpoints. We remark only that the absence of classical negation makes this less easy than it otherwise would be.

6 IF parity games

We briefly recall the game semantics of first-order logic and of IF logic.

Given a FO formula ψ (in positive form) and a structure A, 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 $(\phi_1 \wedge \phi_2, \vec{v})$, Abelard chooses a conjunct ϕ_i, and play moves to $(\phi_i(\vec{x}'), \vec{v}')$, where $\vec{x}', \vec{v}'$ are $\vec{x}, \vec{v}$ restricted to the free variables of ϕ_i; and at $(\phi_1 \vee \phi_2, \vec{v})$, Eloise similarly chooses a disjunct. A play of the game terminates at (negated) atoms $(P(\vec{x}), \vec{v})$ (resp. $(\neg P(\vec{x}), \vec{v})$), and is won by Eloise (resp. Abelard) iff $P(\vec{v})$ is true. Then it is standard that $M \vDash \phi$ 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 IF logic [Hintikka & Sandu 1996] use games of imperfect information: at the position $\exists x/W.\, \phi$, when Eloise chooses a value v for x, she does not know what Abelard chose for the values of the independent variables W. A *uniform* Eloise strategy for the game is one in which her choice of v is indeed uniform in the values of W, and we say a formula is true if Eloise has a uniform winning strategy.

Now recall that in a parity game the positions are assigned ranks $0, \ldots, r$, and if a run of the game is infinite, Eloise wins if the highest rank appearing infinitely often is even. The model-checking game for FOL extends to a model-checking game for LFP by assigning even ranks to maximal fixpoints and odd to minimal, such that the rank of an inner fixpoint is less than the rank of its enclosing fixpoints. Then the formula is true iff Eloise has a winning strategy for the defined parity game.

Combining these two concepts, a general parity game of imperfect information is given by a usual parity game together with imperfect information requirements at each position, requiring a player to move uniformly in some part of the game history. The winning runs are those given by the usual parity winning conditions; a player wins the game if she has winning strategy for the parity game that is uniform as required by the imperfect information requirements.

In general, infinite imperfect information games are undecidable even on finite structures, since they require players to keep arbitrary knowledge (and lack of knowledge) of the history of the game. To obtain a class of decidable imperfect parity games, we will first give a parity game semantics for IF-LFP, and then define a class of imperfect parity games characterized by IF-LFP.

Definition 17. *The model-checking game for an IF-LFP formula is defined by adding the following clauses to the Hintikka–Sandu game for IF. The moves are extended by the usual fixpoint unfolding rule: at a position* $((\mu(X,x).\phi)(t),\vec{u})$*, play moves to* $(\phi,\vec{u}v)$*, where* v *is the value of* t*; at a position* $(X(t),\vec{u}v\vec{w})$*, where* $\vec{u}$ *is the deal for the free variables of* X*,* v *for* x*, and* $\vec{w}$ *for the variables bound inside* ϕ*, play moves to* $(\phi,\vec{u}v')$ *where* v' *is the value of* t*. Parities are assigned to positions in the usual way, and the usual infinite parity winning condition is added.*

The independence requirements are that at a quantifier $\exists x/W.$ *(and dually), Eloise must choose* x *without knowing the values of the* W *variables* and without knowing the values of any variables bound in some currently enclosing fixpoint but chosen before the most recent unfolding of that fixpoint. *(In other words, she does not remember choices that have gone out of scope and have no value in the current deal.)*

Correspondingly, a uniform strategy in the parity game is a strategy where the choice function is uniform in the independent variables and the out-of-scope variables. ◁

Theorem 18. *If* ϕ *is an IF-LFP sentence, then Eloise has a uniform winning strategy for the model-checking game if and only if* $\{\langle\rangle\}$ *is a trump for* ϕ*. Moreover, the strategy can be history-free – that is, the choice of move depends only on the current position in the game.*

Proof. The argument relating parity conditions to alternating fixpoints relates any set of monotone operators with fixpoints added, to a parity game where moves correspond to operator application; not just in the case of FOL or modal logic. If there is a trump for ϕ, and $\vec{u}$ is in the trump, then the trump gives Eloise a strategy to follow from $(\phi,\vec{u})$ up to the next fixpoint unfolding, and so on, *ad infinitum* for a greatest fixpoint, or *ad finem* for a least fixpoint, according to the approximant semantics – and this strategy is history-free, as it depends only on the trump and the formula. Conversely, if Eloise has a uniform winning strategy for $(\phi,\vec{u})$, then the strategy choices in the initial portion of the game tree up to the first fixpoint unfolding on each branch define trumps. The details are as usual. □

This game account of the IF-LFP semantics brings out the key factor, which may have been less obvious in the trump semantics, that keeps model-checking decidable. This is that passing through a fixpoint variable throws away all information about choices made within the body of the fixpoint, unless they are explicitly passed as parameters. Of course, this is also true in usual LFP, but in the IF case knowledge of previous choices is explicitly part of the semantics.

This suggests the following definition:

Definition 19. *An imperfect information parity game on a structure A is* finite-memory *if each player is equipped with a finite memory in which they can remember previous moves. A player's choice at a move is required to depend only on the current position and memory, with additional imperfect information requirements imposed by the game on the memory (i.e. a player may have to temporarily forget things).*

A player wins the game if they have a uniform history-free winning strategy. ◁

The expected theorem is

Theorem 20. *Given a finite-memory imperfect parity game on A, the statement 'Eloise wins the game' is expressible by an IF-LFP formula whose fixpoint alternation depth is the parity rank of the game.*

Proof. The finite memory is modelled by parameters of fixpoints. We will use fixpoints X which carry one parameter p for the position in the game, and parameters m_i for the memory 'cells'. The inner loop of an inductive definition of winning positions is the usual expression of 'it is Eloise's move and there exists a move such that the next position is in X, or it is Abelard's move and all next moves are in X', as in the formula we used earlier for generalized chess. The quantifiers are made explicitly independent of the memory items required to be unknown (which may require a case analysis of the moves of the game).

To deal with the parities, we use the first-order version of the usual 'parity game formula' from parity automata and modal mu-calculus (see [Bradfield 1999] for a detailed explanation of the parity game formula): for each rank $j = 0, \ldots, r$, there is a fixpoint variable X_j. Then the inner loop is enclosed by $\nu X_0.\mu X_1.\ldots\mu/\nu.X_r.$, and the formula $X(p, \vec{m})$, where p and $\vec{m}$ are the position and memory after the next move, is conjoined with

$$\bigwedge_{0\leq j\leq r}(R_j \Rightarrow X_j(p, \vec{m}))$$

where R_j is the formula expressing that the next position has rank j.

The usual proof now applies to give the result. □

Corollary 21. *Finite-memory imperfect parity games on finite structures are decidable.*

We note that Schobbens [Schobbens 2004] has independently studied ATL (Alternating Temporal Logic) with imperfect recall. ATL games are a restricted instance of IF parity games, and the imperfect recall is our finite memory. Schobbens goes into a more detailed analysis of the differences of the complexities of ATL with or without imperfect information and recall.

7 Application to IF modal mu-calculus.

Our original motivation for looking at fixpoint extensions of IF logic was the desire to combine two threads of work. Firstly, modal mu-calculus is a well studied and widely used temporal logic. Secondly, we have argued in [Bradfield 2000] and [Bradfield & Fröschle 2002] that modal versions of Henkin quantifiers and independence logics provide a natural expression of some properties of concurrent systems. Given a concurrent modal logic, it is natural to extend it to a concurrent temporal logic by adding fixpoint operators. In [Bradfield 2000] we looked at modal analogues of Henkin quantifiers acting on systems composed of several concurrent components; since a single Henkin quantifier gives an operator on the powerset of states, there was no difficulty in adding such modalities to mu-calculus. In [Bradfield & Fröschle 2002], we designed a modal analogue of IF logic, defined on certain structures appropriate for true concurrency. The full definition of the structures and the logic is, for concurrency-theoretic technical reasons, somewhat long and complex. We refer the reader to [Bradfield & Fröschle 2002] for full definitions; here we will give an overview of the logic.

IFML extends the syntax of usual modal logic as follows. Instead of the simple 'next step' modality $\langle a \rangle \Phi$, each modality carries a *tag* α, and may be declared to be independent of previous tags β by the Hintikka slash, giving a syntax $\langle a \rangle_{\alpha/\beta} \Phi$. The intended interpretation is that the choice of a action must be independent of the action chosen in the modality tagged by β; for this to make sense, the action at β should be concurrent (in the technical sense of event structures etc.) with the action at α. The structures for this logic are not simple transition systems, but *transition systems with concurrency*[3], in which there is a concurrency relation $\mathcal{C}$ between transitions, satisfying certain axioms so that concurrent actions are not causally dependent on one another: in particular, if a and b are concurrent, and ab is a possible sequence, then so must

[3]These are normally called 'transition systems with independence', but that would be confusing in the current context

ba be. The semantics of IFML is then given in terms of runs (sequences of states) of the system, directly via an imperfect information model-checking game. We say that a run satisfies $\langle a \rangle_{\alpha/\beta}\Phi$ if Eloise can choose an a transition and move to a position satisfying Φ, and can do so 'uniformly' in the previous transitions labelled by β. Here 'uniformly' means that the choice of a transition is good also for all other runs in which the β-labelled transitions b_i are substituted by concurrent transitions b'_i.

Remark 22. *At this point, we should note that Tulenheimo in his doctoral dissertation [Tulenheimo 2004] has also defined and studied 'independence-friendly modal logic'. Tulenheimo's definition of IFML differs from ours – he sticks strictly to informational independence in strategies, and uses ordinary Kripke structures as the structures for his logic. This IFML turns out to be weaker than ours, and indeed first order expressible, whereas our IFML can express NP (and hence Σ^1_1-hard) properties. Accordingly, Tulenheimo takes issue with the claim above from [Bradfield & Fröschle 2002] that concurrency is 'necessary' for interpreting an IFML. We must concede his point; but we believe the link between concurrency and independence in our formulation is certainly of interest for computer science, and arguably of interest to philosophical logicians.*

The following remarks explain the concurrency-theoretic motivation for defining IF-LFP. The theorems formalizing these remarks will be presented in a more concurrency-theoretic forum.

Remark 23. *The game semantics of IFML given in [Bradfield & Fröschle 2002] can be equivalently expressed by translating to IF as a metalanguage (modulo the introduction of some fairly messy defined functions and relations on runs of the system) such that the main variable holding the state ranges over runs (as in the game), and auxiliary variables range over actions. Consequently, IFML has a trump semantics. The evaluation of a formula on a finite system is decidable, since the maximum length of runs that must be considered is bounded by the modal depth of the formula.*

Remark 24. *We can define an IF modal mu-calculus by adding fixpoint formulae of the form $\mu(X, \chi).\Phi$ and $X(\alpha)$, where the fixpoint variable X has not only an implicit parameter for the current 'state', but also explicit parameters χ for tags to be passed through the fixpoint.*

This can be given a semantics via IF-LFP. However, since a 'state' in the semantics is a run, not a system state, it is not obvious that decidability of model-checking is maintained for IF mu-calculus. (We conjecture *that it is, but some results from concurrency theory, such as the undecidability of hereditary history-preserving bisimulation, give some cause for doubt.)*

The IF modal mu-calculus has a model-checking game that is an IF version of the usual parity games for modal logic, as done above for IF-LFP.

8 Conclusion

We have defined a suitable fixpoint extension of independence-friendly logic, and established some results. We have related it to parity games of imperfect information, and we have shown how it may be applied to the construction of independence-friendly modal mu-calculi.

For IF-LFP itself, there are still many questions remaining. Chief among these are better upper and lower bounds on the complexity of model-checking (in the finite case) and descriptive complexity (in the infinite case). We have shown that IF-LFP is more complex than we surmised in earlier work, and it is not unlikely that it will turn out to be much more complex. For the finite case, a forthcoming article with Stephan Kreutzer will contain several results, including that model-checking is at least EXPSPACE-hard.

Once these are resolved, the question also arises, as remarked by the referee, of the complexity of IF modal mu-calculus. IFML itself is already Σ^1_1; we expect IF modal mu-calculus to have complexity similar to that of IF-LFP. Clearly it also has the usual expressive power of temporal logics over modal logics, in that it can describe infinite behaviours; but how this combines with the Henkin quantification is unclear.