Expressiveness for highly nested Expressions in Linear Temporal Logic

William P.R.Mitchell,
Department of Computer Science,
University of Manchester,
Manchester M13 9PL
bill@cs.man.ac.uk

March 2, 1997

Abstract

In the case where one only uses the modalities always, sometime and next it is shown that nesting modalities do not constrain a model in the limit any more than a non-nested formula. It is also proved that when the until modality operator is included then nesting modalities does constrain the limit of a model more than any non-nested formula.

1 Introduction

Branching time logics have been long established as important tools for verifying finite state concurrent systems [3], [4]. The expressiveness of a temporal logic is a lively area of study. In branching time temporal logic a hierarchy of CTL^* sub-logics has been discovered [8], [9] in respect of their expressive capabilities. A concise definition of the various sub-logics and related issues can be found in the survey paper [11].

To write logic A $<$ logic B intuitively means that all properties expressible in A can also be expressed in B, and that some property can be expressed in logic B but not in logic A. To be precise A $<$ B means that for every formula g of A there is an equivalent h of B which includes only variables occurring in g, and there is a formula f of B such that there is no formula of logic A involving only the variables of f which is logically equivalent over all possible models to f.

It is not really surprising that the various sub-logics of CTL^* surveyed in [11] can be differentiated through the $<$ relation, even though the various proofs needed to establish the hierarchy are ingenious and complex. Roughly each one in the hierarchy is a syntactic generalisation of the previous one, for example by adding a new fixed point operator or allowing new syntactic combinations of fixed point operators.

The context for discussion in this paper is restricted to linear time propositional temporal logic. The model of time for these logics is taken to be the natural numbers $\mathbf{N}$. Consider two such logics, BTL which has modalities $\Box$ always, $\Diamond$ sometime and $\bigcirc$ next, and the logic PTL which is BTL plus the modality $\mathcal{U}$ until. To say that PTL is more expressive than BTL the traditional view as exhibited above is to prove that some PTL formula is not equivalent to any BTL formula. Though it is important to prove that $BTL < PTL$ this fact does not illuminate how deep the differences of expressivity are between the two logics. This paper illustrates how on a global level they differ, that is how some expressive property of all BTL formula is fundamentally different from that of PTL.

Intuitively when one says how expressive a logic is one would like some feel for all the properties which a logic is capable of describing, which is different from the comparison of two logics in the sense described above. For linear time propositional logic PTL it is known that PTL is equivalent to the First Order Language of Linear Order, see the survey paper [7] for further details.

A model M is a sequence of states or worlds $\sigma_0\sigma_1\cdots$, where each σ_i defines the truth values of the propositional variables at moment i. This generates an infinity of other models $M^n = \sigma_n\sigma_{n+1}\cdots$. Write $M \models f$ to mean that M satisfies temporal formula f.

Given a formula f and a model M satisfying f it is easy to understand how f constrains the first few moments in M, that is what constraints are imposed on σ_n for small values of n. Unfolding the modalities determines the initial behaviour of a temporal formula. An interesting f will always constrain the events along the whole length of M. What kind of constraints can a formula from BTL or PTL place on the whole of M? A question related to this is: given a model M and a formula f, at which moments t is f true in M? In particular as t becomes arbitrarily large when is f true at time t in M, when does $M^t \models f$?

A formula f is defined to be *flat* if $\bigcirc$ is the only modality to appear within the scope of any other modality appearing in f. So $\Box(\bigcirc x \vee y) \wedge \Diamond \bigcirc \bigcirc x$ is flat but $\Box(\bigcirc y \vee \neg\Diamond x)$ is not flat. A flat formula is conceptually easy to understand, it is the nested formulae where deeper expressiveness is possible. Given $M \models f$ in what way do nested modalities in f constrain M^t as t becomes arbitrarily large?

It is possible to define a finite directed graph $G(f)$ which provides a template for all models of f, see [15] for a description. The graph is given by unfolding the outer modalities of f in a particular way to derive new formulae which are then similarly unfolded. These formulae form the nodes of $G(f)$. This graph is known as the tableaux for f.

Let $g \leq f$ denote that g is a sub-term of f. The nodes of $G(f)$ are labeled with particular temporal formulae called states, and each of these states are of the form $\wedge_i w_i$ where either $w_i \leq f$ or $w_i = \bigcirc h_i$ and $h_i \leq f$. Any model for f corresponds to what is known as a good path P in $G(f)$. There is not necessarily a unique good path which gives a particular model. The precise details of what constitutes a good path are not relevant to the discussion here, it is only important that some set of paths in $G(f)$ correspond to the models which satisfy f.

If M is derived from a good path P, by definition the formula a_t which labels the t-th node of P must be true in M^t. For this discussion the constraint which f places on a model M is regarded as the set of good paths to which M corresponds. For a sequence of temporal formulae $Q = q_0 q_1 \cdots$, write $M \models Q$ to mean $M^n \models q_n$ for each n. Note that in this notation $M \models f$ if and only if for some good path P, $M \models P$.

Other authors [10] define $G(f)$ in such a way that there is a unique good path to each satisfying model of M. It is beyond the scope of this paper to discuss the details of such graphs. The details given above are sufficient to illustrate the results of this paper.

There are clearly only finitely many distinct states in $G(f)$. Wolper uses this fact and that the states are made from sub-terms of f to derive a bound on the total size for $G(f)$ [15]. Thus in an infinite good path $P = a_1 a_2 \cdots$ there are only finitely many distinct a_n, which means there is some N such that for each $n \geq N$, a_n occurs infinitely often along the path P.

This has implications for our understanding of the constraint that f places on M. As $M \models f$ there is a good path $P = a_1 a_2 \cdots$ where $M^t \models a_t$ for each t. For any $n \geq N$ it follows that $M \models \Box\Diamond a_n$. Thus there is a finite subset of the a_i which controls how events occur in M after some initial period, and $M \models \Box\Diamond a_i$ for each of these a_i.

In this paper we will prove the following for any BTL formula f.

Theorem 1.1 *For each $f \in BTL$ there is a flat BTL formula* f *such that*

$$\vdash \Diamond \Box (f \leftrightarrow \mathsf{f})$$

and the variables in f *are a subset of the variables in f.*

Theorem 1.1 implies that nesting of modalities in BTL does not put any more constraint on a model in the limit than a flat formula can. This is explained as follows.

For a good path P of f, as above, let $\mathsf{a_i}$ be a flat formula such that $\vdash \Diamond \Box (a_i \leftrightarrow \mathsf{a_i})$. The $\mathsf{a_i}$ exist because of Theorem 1.1. A model M of f is governed after some initial period by a finite set of the a_i, and for these a_i, $M \models \Box \Diamond a_i$. After some initial period it is also true that $a_i \equiv \mathsf{a_i}$. The model M after some initial period is therefore governed by the flat $\mathsf{a_i}$ and $M \models \Box \Diamond \mathsf{a_i}$.

Let $P[n]$ be the sequence P with each a_i replaced by $\mathsf{a_i}$ for $i \geq n$. An immediate consequence of theorem 1.1 is

Corollary 1.2 *For each good path P and model M there is a number N such that*

$$M \models P \Leftrightarrow M \models P[N]$$

As was stated earlier $M \models f$ if and only if for some good path P, $M \models P$. Intuitively corollary 1.2 implies that after some point we may flatten all the formulae in a good path without altering those models which are derived from it. This is what is meant by an eventually flat constraint.

Theorem 1.1 also helps investigate the problem of at what times t is it true that $M^t \models f$. From the theorem there is some T such that $\forall t > T$, $M^t \models f \Leftrightarrow M^t \models \mathsf{f}$.

The rest of the paper shows that the above simplifications are not possible in general for PTL formula. In particular none of the formulae

$$\alpha_n = p_1 \,\mathcal{U}\, (p_2 \,\mathcal{U}\, \cdots (p_{n-1} \,\mathcal{U}\, p_n) \cdots)$$

have the above property for $n \geq 3$. For each α_i and each flat formula w of PTL, $\Diamond \Box (\alpha_i \leftrightarrow \mathsf{w})$ is false. So when we replace BTL with PTL in theorem 1.1 above it fails for each α_n for $n \geq 3$. One interpretation of this is that nesting of until does create genuine additional constraints not given by flat formulae alone. Moreover corollary 1.2 is also untrue in the PTL case.

These results illustrate an interesting difference between BTL and PTL. Nesting of modalities in BTL only complicates the behaviour of the initial part of a model, the effects degenerating into flat constraints in the limit. Whereas in PTL nested constraints add extra complexity along the whole length of a model.

Work related to this paper can be found in [2] and [5], where first order temporal formulae are written in a flat normal form. Unlike the work here those results introduce new variables in the flattened formulae to embed fixed point iteration explicitly into the formula. By introducing new variables the strict equivalence necessary for $\vdash$ is violated.

To the surprise of the author work closely related to that presented here can also be found in work of the Polish Logician Wajsberg, who in 1933 defined an idea which is essentially the same as flatness for modal logics, and proved that in $S5$ every formula is equivalent to a flat formula. This work can be found in chapter five of Fittings book [6]. One could in fact use proposition 13.1 from Fitting to shorten the proof of Theorem 1.1, but it seems simplest to keep the proof intact.

2 Flat Sub-terms

This section provides a proof of 1.1. Before embarking on a proof of theorem 1.1 various special cases are proved in the form of a series of lemmas. When writing formulae like $(a \wedge b) \vee c$ the convention will be to write $a \wedge b \vee c$. The symbol $\wedge$ will always be assumed to have priority over $\vee$.

For the rest of the paper we will abbreviate $\Diamond \Box(f \leftrightarrow g)$ to $f \equiv_\infty g$. Under this notation theorem 1.1 states that for every BTL formula f there is a flat BTL formula f such that $\vdash f \equiv_\infty \mathsf{f}$. Also $\Diamond \Box(f \rightarrow g)$ is abbreviated to $f \rightarrow_\infty g$, and $\Diamond \Box(f \leftarrow g)$ is abbreviated to $f \leftarrow_\infty g$. Notice that $f \equiv_\infty g$ if and only if $f \rightarrow_\infty g$ and $f \leftarrow_\infty g$.

Lemma 2.1 *Let* $f = \Box(\Diamond a \wedge b \vee c)$ *and* $\mathsf{f} = \Box c \vee \Diamond a \wedge \Box(b \vee c)$, *then*

$$\vdash f \equiv_\infty \mathsf{f}$$

Proof *First consider* $\mathsf{f} \rightarrow_\infty f$. *Let* M *be any model. Note that* $\Diamond \Box p \vee \Box\Diamond\neg p$ *is a tautology, so that* $\Diamond \Box c \vee \Box\Diamond\neg c$ *holds in* M.

If $M \models \Diamond \Box c$ *then clearly* $M \models \Diamond \Box f$ *and that proves the implication. Suppose then that* $M \models \Box\Diamond\neg c$. *By the above remark* $M \models \Box\Diamond\mathsf{f} \vee \Diamond \Box\neg\mathsf{f}$. *If* $M \models \Diamond \Box\neg\mathsf{f}$ *then the implication is proved since for some* n, $M^n \models$

$\Box\neg\mathsf{f}$. *For every* $m \geq n$ *it follows that* $M^m \models \mathsf{f} \rightarrow f$ *because* f *reduces to* $False$ *in* M^m *and* $False \rightarrow f$ *is always true.*

Suppose then that $M \models \Box\Diamond\mathsf{f}$. *Thus*

$$\begin{array}{ll}
M \models \Box\Diamond\mathsf{f} & \Rightarrow \\
M \models \Box\Diamond(\Diamond a \wedge \Box(b \vee c)) & \Rightarrow \\
M \models \Box\Diamond a \wedge \Diamond\Box(b \vee c) & \Rightarrow \\
M \models \Diamond(\Box\Diamond a \wedge \Box(b \vee c)) & \Rightarrow \\
M \models \Diamond\Box f & \Rightarrow \\
M \models \Diamond\Box(\mathsf{f} \rightarrow f) &
\end{array}$$

This completes the first part of the proof. That is $\vdash \mathsf{f} \rightarrow_\infty f$.

Next consider $f \leftarrow_\infty \mathsf{f}$. *Let* M *be any model. As before if* $M \models \Diamond\Box c$ *then the proof is done. Suppose then that* $M \models \Box\Diamond\neg c$. *If* $M \models \Diamond\Box\neg f$ *then as before the implication is proved. Suppose then that* $M \models \Box\Diamond f$. *That implies* $M \models \Box\Diamond(\Diamond a \wedge b)$ *and so* $M \models \Box\Diamond a$. *If for some* n *we have* $M^n \models f$ *then* $M^n \models \Box(b \vee c)$. *Also* $M^n \models \Box\Diamond a$, *from which it follows* $M^n \models \Diamond a$, *thus* $M^n \models \Box\Diamond a \wedge \Box(b \vee c)$ *so* $M^n \models \mathsf{f}$. *That proves* $M \models \Diamond\Box(f \rightarrow \mathsf{f})$ *which completes the proof.*

The rest of the lemmas have a similar pattern of proof so they are not written out in quite such detail.

Lemma 2.2 *Let* $f = \Box(\Box a \wedge b \vee c)$ *and* $\mathsf{f} = \Box c \vee \Box a \wedge \Box(b \vee c)$, *then*

$$\vdash f \equiv_\infty \mathsf{f}$$

Proof *In this case* $\vdash \mathsf{f} \rightarrow f$, *so it only remains to show that* $\vdash f \rightarrow_\infty \mathsf{f}$. *Let* M *be any model. If* $M \models \Diamond\Box c$ *then the proof is done. Suppose then that* $M \models \Box\Diamond\neg c$.

If $M \models \Diamond\Box\neg f$ *again the proof is complete so assume* $M \models \Box\Diamond f$. *Therefore* $M \models \Box\Diamond(\Box a \wedge b)$, *and hence* $M \models \Diamond\Box a$. *That is for some* n, $M^n \models \Box a$. *If for* $r \geq n$, $M^r \models f$ *then* $M^r \models \Box(b \vee c)$, *so* $M^r \models \Box a \wedge \Box(b \vee c)$. *That is* $M^r \models \mathsf{f}$ *and the proof is done.*

Lemma 2.3 *Let* $f = \Diamond(a \wedge \Diamond b)$ *and* $\mathsf{f} = \Diamond a \wedge \Diamond b$, *then*

$$\vdash f \equiv_\infty \mathsf{f}$$

Proof *Note that* $\vdash f \rightarrow \mathsf{f}$, *so it only remains to show that* $\vdash \mathsf{f} \rightarrow_\infty f$. *Let* M *be any model. If* $M \models \Diamond\Box\neg\mathsf{f}$ *then the proof is done. Assume that* $M \models \Box\Diamond\mathsf{f}$. *Notice* $\vdash \Diamond(\Diamond a \wedge \Diamond b) \rightarrow \Diamond a \wedge \Diamond b$ *and that* $\vdash \Box(\Diamond a \wedge \Diamond b) \rightarrow \Box\Diamond(a \wedge \Diamond b)$. *Therefore* $M \models \Diamond\Box\Diamond(a \wedge \Diamond b)$. *Which completes the proof.*

Lemma 2.4 *Let* $f = \Diamond(a \wedge \Box b)$ *and* $\mathsf{f} = \Diamond a \wedge \Box b$*, then*

$$\vdash f \equiv_\infty \mathsf{f}$$

Proof *Note that* $\vdash \mathsf{f} \rightarrow f$*, so it only remains to show that* $\vdash f \rightarrow_\infty \mathsf{f}$*. Let* M *be any model. If* $M \models \Box\neg f$ *then trivially the proof is done. Assume then that* $M \models \Diamond f$*. That implies* $M \models \Diamond\Box b$*. For some* n*,* $M^n \models \Box b$*. If* $m \geq n$ *and* $M^m \models f$ *it follows that* $M^m \models \Box b \wedge \Diamond a$ *hence the proof is complete.*

This is the last lemma needed in the proof of theorem 1.1, which follows and concludes this section.

Proof of theorem 1.1.
To prove the theorem we are required to construct a flat formula f for each $f \in BTL$ such that $\vdash f \equiv_\infty \mathsf{f}$. It is enough to prove the result holds for those f where no $\bigcirc$ occurs with an $\Box$ or $\Diamond$ in its scope, and where negations only occur in literals. This is acceptable as any BTL formula can be written in this form via the equivalences $\bigcirc\Box x \equiv \Box\bigcirc x$, $\bigcirc\Diamond x \equiv \Diamond\bigcirc x$, $\bigcirc(x \vee y) \equiv \bigcirc x \vee \bigcirc y$, $\bigcirc(x \wedge y) \equiv \bigcirc x \wedge \bigcirc y$, and $\neg\Box x \equiv \Diamond\neg x$, $\neg\Diamond x \equiv \Box\neg x$, $\neg\bigcirc x \equiv \bigcirc\neg x$, $\neg(x \wedge y) \equiv \neg x \vee \neg y$, $\neg(x \vee y) \equiv \neg x \wedge \neg y$ and $\neg\neg x \equiv x$.

Let w be some BTL formula with h a sub-term of w. Suppose there is a flat h such that $\mathsf{h} \equiv_\infty h$. Let w/h be w with h replaced by h. It follows that $w/\mathsf{h} \equiv_\infty w$.

The proof of 1.1 is by induction on $|f|$, the length of f. The base case $|f| = 0$ is trivial. Assume that the theorem is true for $|f| < n$. Suppose that f is a formula with $|f| = n$.

Let μg be a sub-term of f which is not within the scope of any modality and μ is $\Box$ or $\Diamond$ and where g contains an $\Box$ or $\Diamond$. If no such sub-term exists then f is flat and the proof is finished.

Possible forms that μg can take are

1. $\Box((\Diamond a \wedge b \vee c) \wedge e)$
2. $\Box((\Box a \wedge b \vee c) \wedge e)$
3. $\Diamond((\Diamond a \wedge b \vee c) \wedge e)$
4. $\Diamond((\Box a \wedge b \vee c) \wedge e)$

Where c and/or b and/or e might be null formulae. For cases 1 through 4 one can apply lemmas 2.1 through 2.4 to move the inner modality out of the scope of the outer modality. For example in case 3 lemma 2.3 proves that

$$\vdash \Diamond((\Diamond a \wedge b \vee c) \wedge e) \equiv_\infty \Diamond a \wedge \Diamond(b \wedge e) \vee \Diamond(c \wedge e)$$

By the use of lemmas 2.1 to 2.4 applied respectively to each of these special cases we can say that

$$\vdash \mu g \equiv_\infty \bigvee_j \bigwedge_i h_{i,j} \wedge g_{i,j}$$

where $|h_{i,j}|$, $|g_{i,j}| < |g|$.

By induction there are flat $\mathsf{h}_{\mathsf{i,j}}$ and $\mathsf{g}_{\mathsf{i,j}}$ as in the theorem. Hence

$$\vdash \mu g \equiv_\infty \bigvee_j \bigwedge_i \mathsf{h}_{\mathsf{i,j}} \wedge \mathsf{g}_{\mathsf{i,j}}$$

Repeat this process for all other sub-terms like μg.

Thus for each non-flat sub-term $\mu_i g_i$ of f there is a flat w_i for which $\vdash \mathsf{w}_\mathsf{i} \equiv_\infty \mu_i g_i$. Define f to be f with each μg_i replaced by the corresponding w_i.

That completes the induction step and so completes the proof.

3 PTL does not Flatten

Let α_n be as defined in the introduction. In this section it is proved that for $n = 3$ the PTL version of theorem 1.1 does not hold. That is there is no flat formula w such that $\mathsf{w} \equiv_\infty a\,\mathcal{U}\,(b\,\mathcal{U}\,c)$.

The proof for α_3 generalises to the α_n case, but the corresponding proof is cumbersome. Thus for the sake of approaching readability the proof is only given for the α_3 case. The proof will be by contradiction. Consider $\alpha_3 = a\,\mathcal{U}\,(b\,\mathcal{U}\,c)$ which we shall abbreviate to w.

If $a_0 a_1 \cdots a_n$ is a sequence of temporal formulae and M is a model, we write $N = a_0 a_1 \cdots a_n M$ to denote that N is a model where, for $0 \leq i \leq n$, $N^i \models a_i$ and $N^{n+1} = M$. For a formula f write f^n to denote the sequence $ff \cdots f$ of f repeated n times. For two sequences a and b, ab denotes their concatenation.

A finite model is a finite sequence of truth valuations $\sigma_0 \sigma_1 \cdots \sigma_n$. For two models M, a finite model, and N, possibly infinite, let MN be the obvious concatenation model. If M is a finite model let M^∞ be the model given by infinitely repeating M. That is $M^\infty = MMM\cdots$. A *cyclic model* is by definition any model of the form M^∞.

For an arbitrary model M let $M[a,b]$ denote the finite model of length $b-a+1$ specified by the truth valuations given by M from time a to time b. If $M = \sigma_0\sigma_1\cdots$, where σ_i is the truth valuation function for the i-th moment of M, then $M[a,b] = \sigma_a\cdots\sigma_b$. For a finite model M we will write $M \models f$ to denote that for all (infinite) models N, $MN \models f$.

Suppose that M is a cyclic model of w. If theorem 1.1 were true for w it would follow that $M \models w \equiv_\infty \mathsf{w}$, where w is some flat PTL formula. Since M is cyclic that means $M \models \Box(w \leftrightarrow \mathsf{w})$. Thus theorem 1.1 implies that in any cyclic model $\Box(w \leftrightarrow \mathsf{w})$.

In this section we will show that for an arbitrary flat w there is a cyclic model M in which $\Box(w \leftrightarrow \mathsf{w})$ is false. This is of course a contradiction which refutes theorem 1.1. The proof for α_3 can be generalised to α_n.

Define $\mathbf{T}\bigcirc$ to be those PTL formulae where the only modality which may occur is $\bigcirc$. Note that all flat formulae can be formed as disjuncts and conjuncts of formulae of the form a, $\Box a$, $\Diamond a$, $a\,\mathcal{U}\,b$ where a, $b \in \mathbf{T}\bigcirc$.

Lemma 3.1 *There is no flat PTL formula* w *equivalent over all cyclic models to* $a\,\mathcal{U}\,(b\,\mathcal{U}\,c)$*, where* a*,* b *and* c *are propositional variables.*

Proof.
The proof is by contradiction. Suppose that $\mathsf{w} \in PTL$ is equivalent to w on every cyclic model. Write w as a disjunct of terms which is of the form

$$\bigvee_i \Box(z_i) \wedge \bigwedge_r \Diamond q_{r,i} \wedge h_i \wedge \bigwedge_j a_{i,j}\,\mathcal{U}\,b_{ij},$$

where each z_i, $q_{r,i}$, h_i, $a_{i,j}$ and $b_{i,j}$ are formulae of $\mathbf{T}\bigcirc$.

Let M be a model of the form

$$(a^n b^m M_0)^\infty$$

Note this does not stipulate what the value of c is at any time, nor the value of b at times 0 to $n-1$, nor the value of a at times n to $n+m-1$. Gradually the form of M is to be restricted using the structure of w, until finally we are left with a contradiction.

Let $|\mathsf{w}|$ be the syntactic length of w. In the proof let $\Box^r g$ denote $g \wedge \bigcirc g \wedge \cdots \bigcirc^r g$, that is g is true now and at each of the next r moments. We will often write $\Box^{|\mathsf{w}|} g$ in the proof which means g is true now and for the next $|\mathsf{w}|$ moments.

Restrict M so that $M \models \Box^{n+m-1} \neg c$. Also choose n larger than $|h_i|$ for each h_i in w (we can pick any $n \geq |\mathsf{w}|$). Let $\alpha = \max\{|h_i|\}$. Arbitrarily choose a value for b for each of the first α moments. Each h_i is made either true or false by this choice. Thus with regards to the model M each h_i in w can be appropriately replaced by either true or false. If h_i is replaced by false then the whole conjunct to which it belongs can be removed from w, otherwise just remove h_i from the conjunct.

It is possible at this point that either w has been reduced to true, or to false. If the first is the case we can trivially choose the rest of M so that $M \models \neg w$, which contradicts that $M \models \Box(w \leftrightarrow \mathsf{w})$. In the second case we can choose M so that $M \models w$ which is also a contradiction. We can then suppose that in a suitable M with value of b, a and c chosen as indicated w is equivalent to a disjunct of formulae of the form

$$\bigvee_i \Box(z_i) \wedge \bigwedge_r \Diamond q_{r,i} \wedge \bigwedge_j a_{i,j} \,\mathcal{U}\, b_{ij}$$

For each $q_{r,i}$ in w let $Q_{r,i}$ be a finite model to satisfy $q_{r,i}$. Note that any $q_{r,i}$ which is not satisfiable makes its containing conjunct redundant. Assume that all the different $Q_{r,i}$ in w can be numbered from 1 to N, that is they can be re-indexed as Q_1, Q_2, …, Q_N.

Restrict M_0 to be of the form $W_0 Q_1 Q_2 \cdots Q_N W_1$. Where later restrictions will be placed on W_0 and W_1. This new restriction makes all the $\Diamond q_{r,i}$ valid in M, so w is equivalent in M to

$$\bigvee_i \Box(z_i) \wedge \bigwedge_j a_{i,j} \,\mathcal{U}\, b_{ij}$$

As before the truth value for w may be determined by these new restrictions. If so there is a contradiction, derived in the same way as previously. Assume then this does not occur.

If in a given conjunct $z_i \equiv \mathrm{TRUE}$ delete $\Box(z_i)$ from its containing conjunct. If in a conjunct z_i remains it must be unsatisfiable on some finite model Z_i. Suppose that Z_1, …, Z_l falsify respective conjuncts containing a z_i. Each remaining $\Box(z_i)$ will be false in the finite model $Z_1 Z_2 \cdots Z_l$. Restrict W_1 to be of the form $Z_1 Z_2 \cdots Z_l W_2$. With this new restriction on

M, w is equivalent in M to a disjunct of the form

$$\bigvee_{i=1}^{u} \bigwedge_{j=1}^{v} a_{i,j} \,\mathcal{U}\, b_{i,j}$$

where v depends on i, for simplicity it has not been written as v_i. The model M has been restricted so that it is of the form $(a^n b^m W_0 N)^{\infty}$ where

1. the truth values of a, b and c are fixed for time t s.t. $0 \leq t \leq \alpha - 1$
2. $M \models \Box^{n+m-1} \neg c$, and n is much larger than α.
3. N is a finite model of the form $Q_1 Q_2 \cdots Q_N Z_1 Z_2 \cdots Z_l W_2$

Note that no restrictions are yet in place on W_0.

It may be at this point $M[0, \alpha - 1]$ determines if $M \models \mathsf{w}$. This clearly leads to a contradiction since we can still choose to make $M \models \bigcirc^{n+m} \neg(c \vee a \vee b)$ or $M \models \bigcirc^{n+m} c$. We can suppose then defining $M[0, \alpha - 1]$ in item 1 above has not determined the truth value of w.

Choose n so that

$$n > \sum_{i,j} |a_{i,j}| + \alpha$$

Divide $M[\alpha, n]$ into $u \cdot v$ parts with the i,j-th part denoted $M^a_{i,j}$ of length $\geq |a_{i,j}|$.

Consider the formula $\Box^{|a_{i,j}|}(\neg c \wedge a) \longrightarrow a_{i,j}$ for a particular i and j. If it is false then we can pick the value of b in $M^a_{i,j}$ so that $a_{i,j}$ is false in $M^a_{i,j}$. This would mean one of two things.

Firstly $b_{i,j}$ could be true at or before the point where $a_{i,j}$ is false. If so then $a_{i,j} \,\mathcal{U}\, b_{i,j}$ becomes true. This can then safely be removed from w. Secondly $b_{i,j}$ could be false at the point where $a_{i,j}$ is false. In which case $a_{i,j} \,\mathcal{U}\, b_{i,j}$ is false, so remove the conjunct containing it from w.

In consequence w is reduced to a formula where

$$\vdash \Box^{|a_{i,j}|}(\neg c \wedge a) \longrightarrow a_{i,j}$$

holds for each i and j.

Next choose m so large that $m > \sum_{i,j} |a_{i,j}|$. Divide $M[n, n+m-1]$ into $u \cdot v$ parts denoted $M^b_{i,j}$ of length $\geq |a_{i,j}|$. In an analogous way to the $M^a_{i,j}$ we can alter the value of a in $M^b_{i,j}$ so that the only conjuncts left in w are ones where

$$\vdash \Box^{|a_{i,j}|}(\neg c \wedge b) \longrightarrow a_{i,j}.$$

for each i and j.

At this point we have fixed $M[0, n+m-1]$. Suppose this determines the truth of $M \models \mathsf{w}$. If $M \models \mathsf{w}$, then just as before choose $M^{n+m} \models \neg(c \vee b \vee a)$ to obtain a contradiction. If $M \models \neg\mathsf{w}$ then choose $M \models \bigcirc^{n+m} c$ for a contradiction. Continue the proof assuming the first $n+m$ moments have not determined the truth of w.

Restrict W_0 to be of the form $(a \wedge b \wedge \neg c)^{\gamma} W_3$ where $\gamma > |\mathsf{w}|$, and W_3 is not restricted in any way yet. Just as before if $M[0, n+m+\gamma-1]$ determines the truth of w then there is a contradiction, continue assuming this is not the case. Since $\gamma > |\mathsf{w}|$ and $b_{i,j}$ and $a_{i,j} \in \mathbf{T}\bigcirc$ there must be some $S \subseteq \{1, \ldots, u\}$ and subsets $R \subseteq \{1, \ldots v\}$ such that

$$M \models \mathsf{w} \leftrightarrow \bigcirc^{n+m} \left(\bigvee_{i \in S} \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j} \right)$$

The R subsets depend on i of course, but this is suppressed to simplify notation.

Let M' be the model M with W_3 replaced by $c(\neg c \wedge \neg a \wedge \neg b)B$, where B is an arbitrary finite model. For the rest of the proof write (x, y, z) as shorthand for the formula $(x \wedge y \wedge z)$. So M' is of the form

$$(a^n b^m (a, b, \neg c)^{\gamma} c(\neg c, \neg a, \neg b) BN)^{\infty}$$

with the same restrictions applied to it as to M. Note $w = a \,\mathcal{U}\, (b \,\mathcal{U}\, c)$ is true in M' so by the assumption that $\mathsf{w} \leftrightarrow w$ in cyclic models, $M' \models \mathsf{w}$. All the facts that are assumed about w in M carry over to w in M'. So w is a disjunct of the form above in the model M'. Therefore

$$M' \models \bigcirc^{n+m} \left(\bigvee_{i \in S} \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j} \right)$$

By definition $M' = a^n b^m (a, b, \neg c)^{\gamma} c(\neg c, \neg a, \neg b) BNM'$ so that

$$(a, b, \neg c)^{\gamma} c(\neg c, \neg a, \neg b) BNM' \models \bigvee_{i \in S} \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j}$$

The finite model B is arbitrary and $a_{i,j}, b_{i,j} \in \mathbf{T}\bigcirc$, therefore for an arbitrary infinite model W

$$(a, b, \neg c)^{\gamma} c(\neg c, \neg a, \neg b) W \models \bigvee_{i \in S} \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j}.$$

So for some i,

$$(a, b, \neg c)^{\gamma} c(\neg c, \neg a, \neg b) W \models \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j}.$$

Using this information from M' the construction of M can be finalised. Restrict W_3 to be of the form

$$(\neg a, b, \neg c)^{\gamma}(a, b, \neg c)^{\gamma}(a, \neg b, \neg c)^{\gamma}(a, b, \neg c)^{\gamma}\ c\ (\neg c, \neg a, \neg b)$$

So that M is of the form

$$(a^n b^m (a, b)^{\gamma}(\neg a, b)^{\gamma}(a, b)^{\gamma}(a, \neg b)^{\gamma}(a, b)^{\gamma}\ c\ (\neg c, \neg a, \neg b) N)^{\infty}$$

and in M, $\neg c$ is true up to the moment as shown. That is $M \models \square^{n+m+5\gamma-1} \neg c$. In this model $M \models \neg w$. To obtain a contradiction we must show $M \models \mathsf{w}$.

Model M is defined so that

$$M \models \mathsf{w} \leftrightarrow \bigcirc^{n+m}(\bigvee_{i \in S} \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j}).$$

By definition $M = a^n b^m W_0 N M$, therefore

$$M \models \mathsf{w} \Leftrightarrow W_0 N M \models \bigvee_{i \in S} \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j}.$$

This will provide the required contradiction.

We are given that $\square^{|a_{i,j}|}(\neg c \wedge a) \longrightarrow a_{i,j}$ and $\square^{|a_{i,j}|}(\neg c \wedge b) \longrightarrow a_{i,j}$. Therefore, as $\gamma > |a_{i,j}|$, $(W_0 N M)^t \models a_{i,j}$ for $0 \leq t \leq 4\gamma - 1$. We are also given that for some i and any W,

$$(a, b, \neg c)^{\gamma} c(\neg c, \neg a, \neg b) W \models \bigwedge_{j \in R}^{v} a_{i,j} \,\mathcal{U}\, b_{i,j}.$$

Hence for some i

$$W_0 N M \models \square^{4\gamma-1}(\bigwedge_{j \in R} a_{i,j}) \wedge \bigcirc^{4\gamma}(\bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j})$$

that is

$$W_0 N M \models \bigwedge_{j \in R} a_{i,j} \,\mathcal{U}\, b_{i,j}$$

Which proves that $M \models \mathsf{w}$. This completes the proof.

Conclusion

Consider the problem: given model M and formula f at which times t does $M^t \models f$? Investigating this problem for the logics BTL and PTL illustrates a difference in their expressivity. For BTL one can find a flat formula f such that $\vdash f \equiv_\infty \mathsf{f}$. Therefore for large t, $M^t \models f \Leftrightarrow M^t \models \mathsf{f}$, i.e.

$$\lim_{t \to \infty} (M^t \models f \Leftrightarrow M^t \models \mathsf{f})$$

From this it is possible to argue, (as in the introduction), that formulae in BTL with nested modalities are no more expressive than flat formulae in their ability to govern the events at the tail end of a model. For PTL nesting modalities does add extra expressive power to govern a model in the limit. Even the simplest possible PTL nested formula fails to have a flat counterpart. To express these ideas at an intuitive level, if you wait long enough the modalities of a BTL formula will float to the outer level, whereas in PTL the $\mathcal{U}$ modality locks in place.

Acknowledgements

I would like to express my sincere thanks to David Carlisle for his invaluable help with this work.

References

[1] M.Abadi 'Temporal-Logic Theorem Proving' Phd Thesis, Stanford University, Report No. STAN-CS-87-1151.

[2] H. Barringer, M. Fisher, D. Gabbay, G. Gough and R. Owens, 'METATEM: A Framework for Programming in Temporal Logic' REX Workshop on Stepwise Refinement of Distributed Systems, Ed. Mook, Lecture Notes in Computer Science 430, Springer Verlag 1989.

[3] Clarke, E. M., E. A. Emerson, 'Design and Synthesis of Synchronisation Skeletons Using Branching Time Temporal Logic', Proc. of Workshop on Logics of Programs, Yorktown Heights, Springer Verlag LNCS 131, editor D. Kozen, pp 52-71, 1981.

[4] Clarke, E. M., E. A. Emerson, A. P. Sistla, 'Automatic Verification of FInite State Concurrent Systems Using Temporal Logic Specifications:

A Practical Approach', Proceedings 10th Annual ACM Symposium on Principals of Programming Languages, Austin, pp 117-126, 1983.

[5] M. Fisher 'A Resolution Method for Temporal Logic' in Proceedings of the International Joint conference on Artificial Intelligence', Sydney Australia, Ed. J. Mylopoulos and R. Reiter, August 1991.

[6] M. Fitting, 'Proof Methods for Modal and Intuitionistic Logics', (Synthese Library Vol 169), published by D. Reidel , 1983.

[7] E. A. Emerson 'Temporal and Modal Logics' in Handbook of Theoretical Computer Science, volume B, Ed Jan Van Leeuwen, MIT press 1990.

[8] E. A. Emerson, J. Y. Halpern, 'Decision Procedures an Expressiveness in the Temporal Logic of Branching Time' Journal of Computer and System Sciences, Vol 30, no. 1, pp 1-24, 1985.

[9] E. A. Emerson, J. Y. Halpern, '"Sometime" and "Not Never" Revisited', Journal ACM, Vol 33, no. 1, pp 151-178, 1986.

[10] E. A. Emerson, A. P. Sistla, 'Deciding full Branching Time Logic', journal of Information and Control, Vol 61, No 3, 1984.

[11] E. A. Emerson, J. Srinivasan, 'Branching Time Temporal Logic', Lecture Notes in Computer Science No 354, pp 123-172, Spinger-Verlag, 1988.

[12] G.Gough 'Decision Procedures for Temporal Logic' MSc Thesis, Manchester University October 1984.

[13] T. Hafer, W. Thomas, 'CTL^* is equivalent to Monadic Second Order Theory of Two Successors over Infinite Binary Trees, with Set Quantification restricted to Infinite Paths', in Proceedings of 14th Annual International Colloquium on Automata, Languages and Programming, LNCS No. 267, pp 269-279, Springer-Verlag, 1987.

[14] O. Lichtenstein, A. Pnueli 'Checking Finite State Concurrent Programs Satisfy their Linear Specifications', 'Proceedings of the twelfth ACM Symposium on the Principles of Programming Languages', January 1985, 97-107.

[15] P. Wolper, 'Temporal Logic can be more Expressive', Information and Control, 56, pp 72-99, 1983