Implementing Temporal Logics: Tools for Execution and Proof

Michael Fisher

Logic and Computation Group,
Department of Computer Science,
University of Liverpool, Liverpool L69, UK
M.Fisher@csc.liv.ac.uk
http://www.csc.liv.ac.uk/∼michael

1 Introduction

In this article I will present an overview of a selection of tools for execution and proof based on temporal logic, and outline both the general techniques used and problems encountered in implementing them. This selection is quite subjective, mainly concerning work that has involved researchers I have collaborated with at Liverpool (and, previously, Manchester). The tools considered will mainly be theorem-provers and (logic-based) agent programming languages. Specifically:

- clausal temporal resolution [21, 28] and several of its implementations, namely `Clatter` [14], `TRP++` [42] and `TeMP` [44], together with its application to verification [35];
- executable temporal logics [24, 4] and its implementation as both METATEM [3] and Concurrent METATEM [22, 49], together with its use as a programming language for both individual agents [23, 26, 29] and multi-agent systems [33, 30, 32].

In addition, I will briefly mention work on induction-based temporal proof [5], temporal logic programming [1], and model checking [7].

Rather than providing detailed algorithms, this presentation will concentrate on general principles, outlining current problems and future possibilities. The aim here is to give the reader an overview of the ways we handle temporal logics. In particular how we use such logics as the basis for both programming and verification.

2 What Is Temporal Logic?

2.1 Some History

Temporal logic was originally developed in order to represent tense in natural language [56]. Within Computer Science, it has achieved a significant role in a number of areas, particularly the formal specification and verification of concurrent and distributed systems [55]. Much of this popularity has been achieved as a number of useful concepts, such as *safety*, *liveness* and *fairness* can be formally, and concisely, specified using temporal logics [20, 52].

F. Toni and P. Torroni (Eds.): CLIMA VI, LNAI 3900, pp. 129–142, 2006.

2.2 Some Intuition

In their simplest form, temporal logics can be seen as extensions of classical logic, incorporating additional operators relating to time. These operators are typically: '$\bigcirc$', meaning "in the next moment in time", '$\Box$', meaning "at every future moment", and '$\Diamond$', meaning "at some future moment". These additional operators allow us to express statements such as

$$\Box(send \Rightarrow \Diamond received)$$

to characterise the statement

> "it is always the case that if we send a message then, at some future moment it *will* be received".

The flexibility of temporal logic allows us to use formulae such as

$$\Box(send \Rightarrow \bigcirc(received \vee send))$$

which is meant to characterise

> "it is always the case that, if we send a message then, at the next moment in time, either the message will be received or we will send it again"

and

$$\Box(received \Rightarrow \neg send)$$

meaning

> "it is always the case that if a message is received it cannot be sent again".

Thus, given formulae of the above form then, if we try to send a message, i.e. '$send$', we *should* be able to show that it is *not* the case that the system continually re-sends the message (but it is never received) i.e. the statement

$$\Box send \wedge \Box\neg received$$

should be inconsistent.

2.3 Some Applications

The representation of dynamic activity via temporal formalisms is used in a wide variety of areas within Computer Science and Artificial Intelligence, for example Temporal Databases, Program Specification, System Verification, Agent-Based Systems, Robotics, Simulation, Planning, Knowledge Representation, and many more. While I am not able to describe all these aspects here, the interested reader should see, for example, [52, 53, 7, 62, 47]. With respect to multi-agent systems, temporal logics provide the formalism underlying basic dynamic/distributed activity, while this temporal framework is often extended to incorporate rational agent aspects such as beliefs, goals and knowledge [27].

There are many different temporal logics (see, for example [20]). The models of time which underlie these logics can be discrete, dense or interval-based, linear, branching or partial order, finite or infinite, etc. In addition, the logics can have a wide range of operators, such as those related to discrete future-time (e.g: $\bigcirc$, $\Diamond$, $\Box$), interval future-time (e.g: $\mathcal{U}$, $\mathcal{W}$), past-time (e.g: $\bullet$ $\blacklozenge$, $\blacksquare$, $\mathcal{S}$, $\mathcal{Z}$), branching future-time (e.g: **A**, **E**), fixed-point generation (e.g: μ, ν) and propositional, quantified propositional or full first-order variations. Even then, such temporal logics are often combined with standard modal logics. For example, typical combinations involve TL + S5 modal logic (often representing 'Knowledge'), or TL + KD45 (Belief) + KD (Desire) + KD (Intention).

Here, I will mainly concentrate on one very popular variety, namely discrete linear temporal logic, which has an underlying model of time isomorphic to the Natural Numbers (i.e. an infinite sequence with distinguished initial point) and is also linear, with each moment in time having at most one successor. (Note that the infinite and linear constraints ensure that each moment in time has *exactly* one successor, hence the use of a single '$\bigcirc$' operator.)

3 Where's the Difficulty?

Temporal Logics tend to be complex. To give some intuition why this is the case, let us look at a few different ways of viewing (initially propositional) temporal logic.

Propositional temporal logic can be thought of as

1. *A specific decidable (PSPACE-complete) fragment of classical first-order logic.*
   For example, the semantics of (discrete, linear) propositional temporal logic can be given by translation to first-order logic as follows. Here, we interpret a temporal formula at a moment in time (indexed by i), and encode this index as an argument to the first-order formula. For simplicity, we consider just propositional symbols, such as p. Then, the question of whether the formula p is satisfied at moment i in a temporal model is translated to the question of whether $p(i)$ is satisfied in a classical first-order logic model:
   $i \models_{TL} p \quad \rightarrow \quad \models p(i)$
   $i \models_{TL} \bigcirc p \quad \rightarrow \quad \models p(i+1)$
   $i \models_{TL} \Diamond p \quad \rightarrow \quad \models \exists j.\ (j \geq i) \wedge p(j)$
   $i \models_{TL} \Box p \quad \rightarrow \quad \models \forall j.\ (j \geq i) \Rightarrow p(j)$
   However, this first-order translation can be a problem as proof/execution techniques often find it hard to isolate exactly this fragment.
2. *A multi-modal logic, comprising two modalities,* $[1]$ *and* $[*]$, *which interact closely.*
   The modality $[1]$ represents a move of one step forwards, while $[*]$ represents all future steps.
   Thus, the induction axiom in discrete temporal logic
   $$\vdash\ \Box(\varphi \Rightarrow \bigcirc\varphi)\ \Rightarrow\ (\varphi \Rightarrow \Box\varphi)$$
   can be viewed as the *interaction* axiom between modalities
   $$\vdash\ [*](\varphi \Rightarrow [1]\varphi)\ \Rightarrow\ (\varphi \Rightarrow [*]\varphi)$$
   Usually, $[1]$ is represented as '$\bigcirc$', while $[*]$ is represented as '$\Box$'.
   However, while mechanising modal logics is relatively easy, multi-modal problems become complex when interactions occur between the modalities; in our case the interaction is of an *inductive* nature, which can be particularly complex.

3. *A characterisation of simple induction.*
   The induction axiom in discrete temporal logic

$$\vdash \ \square(\varphi \Rightarrow \bigcirc\varphi) \ \Rightarrow \ (\varphi \Rightarrow \square\varphi)$$

   can alternatively be viewed as

$$[\forall i.\, \varphi(i) \Rightarrow \varphi(i+1)] \ \Rightarrow \ [\varphi(0) \Rightarrow \forall j.\, \varphi(j)]$$

   Reformulated, this becomes

$$[\varphi(0) \ \wedge \ \forall i.\, \varphi(i) \Rightarrow \varphi(i+1)] \ \Rightarrow \ \forall j.\, \varphi(j)$$

   which should be familiar as a version of arithmetical induction, i.e. if φ is true of 0 and if φ being true of i implies it is true of $i+1$, then we know φ is true of all Natural Numbers.
   However, this use of induction can again cause problems for first-order proof techniques.
4. *A logic over sequences, trees or partial-orders (depending on the model of time).*
   For example, a sequence-based semantics can be given for discrete linear temporal logic:

   $s_i, s_{i+1}, \ldots, s_\omega \models \bigcirc p$ if, and only if, $s_{i+1}, \ldots, s_\omega \models p$
   $s_i, s_{i+1}, \ldots, s_\omega \models \lozenge p$ if, and only if, there exists a $j \geq i$ such that $s_j, \ldots, s_\omega \models p$
   $s_i, s_{i+1}, \ldots, s_\omega \models \square p$ if, and only if, for all $j \geq i$ then $s_j, \ldots, s_\omega \models p$

   This shows that temporal logic can be used to characterise a great variety of, potentially complex, computational structures.
5. *A syntactic characterisation of finite-state automata over infinite words (ω-automata).*
   For example
   - formulae such as $p \Rightarrow \bigcirc q$ give constraints on possible state transitions,
   - formulae such as $p \Rightarrow \lozenge r$ give constraints on accepting states within an automaton, and
   - formulae such as $p \Rightarrow \square s$ give global constraints on states in an automaton.

   This shows some of the power of temporal logic as a variety of different ω-automata can be characterised in this way.

The decision problem for a simple propositional (discrete, linear) temporal logic is already PSPACE-complete [58]; other variants of temporal logic may be worse! When we move to *first-order* temporal logics, things begin to get unpleasant. It is easy to show that first-order temporal logic is, in general, incomplete (i.e. not recursively-enumerable [59, 2]). In fact, until recently, it has been difficult to find *any* non-trivial fragment of first-order temporal logic that has reasonable properties. A breakthrough by Hodkinson *et al.* [39] showed that *monodic* fragments of first-order temporal logic could be complete, even decidable. Monodic first-order temporal logics add quantification to temporal logic but only allow at most one free variable in any temporal subformula. Thus,

$$\forall x.\, a(x) \;\Rightarrow\; \Box b(x)$$
$$\forall x.\, a(y) \;\Rightarrow\; \Diamond c(y, y)$$
$$\forall z.\, (\exists w.\, d(w, z)) \;\Rightarrow\; \bigcirc(\forall v.\, e(z, v))$$

are all *monodic* formulae, whereas

$$\forall x.\, \forall y.\, f(x, y) \;\Rightarrow\; \bigcirc g(y, x)$$

is not. The monodic fragment of first-order temporal logic is recursively enumerable [39] and, when combined with a decidable first-order fragment, often produces a decidable first-order temporal logic [38, 10, 9, 45].

4 What Tools?

The main tools that we are interested in are used to carry out *temporal verification*, via resolution on temporal formulae, and *temporal execution*, via direct execution of temporal formulae. In our case, both of these use temporal formulae within a specific normal form, called *Separated Normal Form (SNF)* [25].

4.1 SNF

A temporal formula in Separated Normal Form (SNF) is of the form

$$\Box \bigwedge_{i=1}^{n} (P_i \;\Rightarrow\; F_i)$$

where each of the '$P_i \Rightarrow F_i$' (called *clauses* or *rules*) is one of the following

$$\mathbf{start} \;\Rightarrow\; \bigvee_{k=1}^{r} l_k \qquad \text{(an initial clause)}$$

$$\bigwedge_{j=1}^{q} m_j \Rightarrow \bigcirc \bigvee_{k=1}^{r} l_k \qquad \text{(a step clause)}$$

$$\bigwedge_{j=1}^{q} m_j \Rightarrow \Diamond l \qquad \text{(a sometime clause)}$$

where each l, l_k or m_j is a literal and '**start**' is a formula that is only satisfied at the "beginning of time".

Thus, the intuition is that:

- initial clauses provide *initial* constraints;
- step clauses provide constraints on the *next* step; and
- sometime clauses provide constraints on the *future*.

We can provide simple examples showing some of the properties that might be represented directly as SNF clauses.

- Specifying initial conditions: **start** $\Rightarrow$ $reading$
- Defining transitions between states: $(reading \wedge \neg finished) \Rightarrow \bigcirc reading$
- Introducing new eventualities (goals): $(tired \wedge reading) \Rightarrow \Diamond \neg reading$
  $reading \Rightarrow \Diamond finished$
- Introducing permanent properties:

$$(increasing \wedge (value = 1)) \Rightarrow \bigcirc \Box (value > 1)$$

which, in SNF, becomes

$$\begin{aligned} (increasing \wedge (value = 1)) &\Rightarrow \bigcirc (value > 1) \\ (increasing \wedge (value = 1)) &\Rightarrow \bigcirc r \\ r &\Rightarrow \bigcirc (value > 1) \\ r &\Rightarrow \bigcirc r \end{aligned}$$

Translation from an arbitrary propositional temporal formula into SNF is an operation of polynomial complexity [25, 28].

We also need the concept of a *merged* SNF clause: any SNF clause is a merged SNF clause and, given two merged SNF clauses $A \Rightarrow \bigcirc B$ and $C \Rightarrow \bigcirc D$, we can generate a new merged SNF clause $(A \wedge C) \Rightarrow \bigcirc (B \wedge D)$.

4.2 Clausal Resolution

Given a set of clauses in SNF, we can apply resolution rules, such as

Initial Resolution:
$$\frac{\begin{array}{l}\textbf{start} \Rightarrow (A \vee l) \\ \textbf{start} \Rightarrow (B \vee \neg l)\end{array}}{\textbf{start} \Rightarrow (A \vee B)}$$

Step Resolution:
$$\frac{\begin{array}{r}P \Rightarrow \bigcirc (A \vee l) \\ Q \Rightarrow \bigcirc (B \vee \neg l)\end{array}}{(P \wedge Q) \Rightarrow \bigcirc (A \vee B)}$$

Temporal Resolution (simplified)[1]:
$$\frac{\begin{array}{l}A \Rightarrow \bigcirc \Box \neg l \\ Q \Rightarrow \Diamond l\end{array}}{Q \Rightarrow (\neg A)\, \mathcal{W}\, l}$$

As we will see later, it is this *temporal resolution* rule that causes much of the difficulty.

It should be noted here that the above is a basic explanation of clausal temporal resolution. A number of refinements, both in terms of what resolution rules are used, and the form of SNF, have been developed [11, 8].

There has also been considerable work on extending the clausal resolution approach to handle logics formed by combining temporal logic with one or more modal logic. In particular, resolution for a *temporal logic of knowledge* (i.e. temporal logic combined with an S5 modal logic of knowledge) have been developed [19]. More recent work in this area has concerned extending resolution to cope with more complex interactions between the knowledge and time dimensions [18, 54] and application of such logics in verification [17, 16].

[1] $(\neg A)\, \mathcal{W}\, l$ is satisfied either if $\neg A$ is always satisfied, or if $\neg A$ is satisfied up to a point when l is satisfied.

4.3 Executable Temporal Logics

In executing temporal logic formulae, we use the *Imperative Future* approach [4]:

- transforming the temporal specification into SNF;
- from the initial constraints, *forward chaining* through the set of temporal rules representing the agent; and
- constraining the execution by attempting to satisfy goals, such as $\Diamond g$ (i.e. g eventually becomes true).

Since some goals might not be able to be satisfied immediately, we must keep track of the outstanding goals and reconsider them later. The basic strategy used is to attempt to satisfy the oldest outstanding eventualities first and keep a record of the others, retrying them as execution proceeds.

Example. Imagine a search agent which can *search*, *speedup* and *stop*, but can also run out of resources (*empty*) and *reset*.

The agent's internal definition might be given by a temporal logic specification in SNF, for example,

$$\begin{aligned} \textbf{start} &\Rightarrow \neg searching \\ search &\Rightarrow \Diamond searching \\ (searching \wedge speedup) &\Rightarrow \bigcirc (empty \vee reset) \end{aligned}$$

The agent's behaviour is implemented by *forward-chaining* through these formulae.

- Thus, *searching* is false at the beginning of time.
- Whenever *search* is made true, a commitment to eventually make *searching* true is given.
- Whenever both *speedup* and *searching* are made true, then either *reset* or *empty* will be made true in the next moment in time.

This provides the basis for temporal execution, and has been extended with execution for combinations with *modal* logics, deliberation mechanisms [26], resource-bounded reasoning [29] and a concurrent operational model [22].

5 Implementations

5.1 Clausal Temporal Resolution

The essential complexity in carrying out clausal temporal resolution is implementing the temporal resolution rule itself. First, let us note that the Temporal Resolution rule outlined earlier is not in the correct form. The *exact* form of this rule is

$$\text{Temporal Resolution (full):} \quad \begin{array}{c} A_1 \Rightarrow \bigcirc B_1 \\ \ldots \Rightarrow \ldots \\ A_n \Rightarrow \bigcirc B_n \\ Q \Rightarrow \Diamond l \\ \hline Q \Rightarrow (\bigwedge_{i=1}^{n} \neg A_i)\, \mathcal{W}\, l \end{array}$$

where each $A_i \Rightarrow \bigcirc B_i$ is a merged SNF clause and each B_i satisfies $B_i \Rightarrow (\neg l \wedge \bigvee_{j=1}^{n} A_j)$.

Thus, in order to implement this rule, a *set* of step clauses satisfying certain properties must be found; such a set is called a *loop*. This process has undergone increasing refinement, as has the implementation of clausal temporal resolution provers in general:

1. The original approach proposed was to conjoin all sets of step clauses to give, so called, *merged* SNF clauses and then treat these merged clauses as edges/transitions in a graph. Finding a loop is then just a question of extracting a *strongly connected component* from the graph, which is a linear operation [60].

   The problem here is explicitly constructing the large set of merged SNF clauses.
2. Dixon [12, 13, 14] developed an improved (breadth-first) search algorithm, which formed the basis of the '`Clatter`' prover. This search approach effectively aimed to generate only the merged SNF clauses that were *required* to find a loop, rather than generating all such clauses.

   The problem with the `Clatter` family of provers was the relatively slow link to the classical resolution system (which was used to carry out the step resolution operations).
3. Hustadt then developed `TRP` [46]. The idea here was to use arithmetical translations to translate as much as possible of the process to classical resolution operations and then use an efficient classical resolution system. In addition, `TRP` used a translation of the breadth-first loop search algorithm into a series of classical resolution problems suggested in [15]. (`TRP` is also able to deal with the combination of propositional temporal logic with various modal logics including KD45 and S5.) The resulting system was evaluated against other decision procedures for this form of temporal logic and was shown to be very competitive [46, 43].

   The main problems with `TRP` were that it was implemented in `Prolog` and that the data/term representation/indexing techniques could be improved.
4. A more recent variety of clausal resolution system for propositional temporal logic is `TRP++`, implemented by Hustadt and Konev [42]. Here, `TRP` is re-implemented in `C++` and is refined with a number of contemporary data representation and indexing techniques.

   `TRP++` currently performs very well in comparison with other provers for propositional temporal logic.
5. Finally, `TeMP` [44] is a clausal resolution prover specifically designed for monodic first-order temporal logic [8, 50]. This utilises the Vampire [57] system to handle much of the internal first-order proving.

Both `TRP++` and `TeMP` are available online[2].

5.2 Executable Temporal Logics

The Imperative Future style of execution provides a relatively simple approach to executing temporal logic formulae. As described above, beginning at the initial conditions

[2] See `http://www.csc.liv.ac.uk/~konev`

we simply forward chain through the step clauses/rules generating a model, all the time constraining the execution with formulae such as '$\Diamond g$'.

The development in this area has not primarily been concerned with speed. As we will see below, the developments have essentially involved refining and extending the internal capabilities of the programs and allowing for more complex interactions between programs.

Thus, the implementations of this approach, beginning with METATEM, proceeded as follows.

1. The first approach, reported in [34], essentially used a `Prolog` meta-interpreter to implement the system. The forward chaining aspect is relatively standard, and the management of outstanding eventualities (i.e. formulae such as '$\Diamond g$') was handled with a queue structure.

   In order to maintain completeness (in the propositional case) a form of *past-time loop checking* had to be employed. This involved retaining a large proportion, and sometimes *all*, of the history of the computation and checking for loops over this as every new computation state was constructed. (Note that this loop-checking aspect is usually omitted from the later languages.)

   The main problems with this approach were the lack of features, particularly those required for programming rational agents, such as internal reasoning, deliberation and concurrency.
2. In [22], Concurrent METATEM was developed. This allowed for multiple asynchronous, communicating METATEM components and provided a clean interaction between the internal (logical) computation and the concurrent operational model.

   Concurrent METATEM was implemented in `C++` but was relatively slow and static (i.e. process scheduling was implemented directly).
3. Kellett, in [48, 49], developed more refined implementation techniques for Concurrent METATEM. Here, individual METATEM programs were compiled into (linked) pairs of I/O automata [51], one to handle the internal computation, the other to handle the interaction with the environment. Such automata can then, potentially, be cloned (for process spawning) and moved (for load balancing and mobility).

   While Concurrent METATEM provides an interesting model of simple multi-agent computation, work was still required on the internal computation mechanism for each individual agent.
4. More recently, the internal computation has been extended by providing a *belief* dimension, allowing meta-control of the deliberation[3], allowing resource-bounded reasoning and incorporating agent abilities [26, 29, 30].

   This has led to a programming language in which rational agents can be implemented, and in which complex multi-agent organisations can be developed. Recent work by Hirsch [37] has produced a `Java` implementation of both individual and group aspects, and has applied this to multi-agent and pervasive computing applications [31, 32, 36].

[3] Deliberation here means the process of deciding in which order to attempt outstanding eventualities at each computation step.

5.3 Other Techniques

In this section, I will briefly mention a few other systems related to temporal logic.

Induction-based Temporal Proof. As mentioned above, first-order temporal logics are complex. In particular, full first-order temporal logic is not recursively-enumerable. However, as we still wish to prove theorems within such a logic, we have been developing techniques to support this. Such a system is described in [5], where an induction-based theorem-prover is enhanced with heuristics derived from the clausal temporal resolution techniques (see above) and implemented in λClam/λProlog.

Temporal Logic Programming. Standard logic programming techniques were transferred to temporal logic in [1]. However, because of the incompleteness of first-order temporal logics, the language was severely restricted. In fact, if we think of SNF above then the fragment considered is essentially that consisting of initial and step clauses. Thus, implementation for such a language is a small extension of classical logic programming techniques and constraint logic programming techniques.

Model Checking. Undoubtedly the most popular application of temporal logic is in *model checking*. Here, a finite-state model capturing the executions of a system is checked against a temporal formula. These finite state models often capture hardware descriptions, network protocols or complex software [40, 7]. While much model-checking technology was based on automata-theoretic techniques, advances in *symbolic* [6] and *on-the-fly* [41] techniques have made model checking the success it is. Current work on abstraction techniques and Java model checking, such as [61], promise even greater advances.

6 Summary

I have overviewed a selection of tools for execution and proof within temporal logic. While this selection has been heavily biased towards those in which I have been involved, several of the techniques are at the forefront of their areas. Although these tools are generally prototypes, they are increasingly used in realistic scenarios, and more sophisticated versions appear likely to have a significant impact in both Computer Science and Artificial Intelligence.

Acknowledgments. Much of this work has been supported by funding from the UK's Engineering and Physical Sciences Research Council (EPSRC)[4], and has involved collaboration with a variety of outstanding colleagues who work (or have worked) within the Logic and Computation group at Liverpool[5]. I thank them all for their endeavours. I would also like to thank the anonymous referees of the CLIMA VI post-proceedings for their helpful comments.

[4] `http://www.epsrc.ac.uk`

[5] `http://www.csc.liv.ac.uk/research/logics`

References

1. M. Abadi and Z. Manna. Temporal Logic Programming. *Journal of Symbolic Computation*, 8: 277–295, 1989.
2. M. Abadi. The Power of Temporal Proofs. *Theoretical Computer Science*, 64:35–84, 1989.
3. H. Barringer, M. Fisher, D. Gabbay, G. Gough, and R. Owens. METATEM: An Introduction. *Formal Aspects of Computing*, 7(5):533–549, 1995.
4. H. Barringer, M. Fisher, D. Gabbay, R. Owens, and M. Reynolds, editors. *The Imperative Future: Principles of Executable Temporal Logics*. Research Studies Press, Chichester, United Kingdom, 1996.
5. J. Brotherston, A. Degtyarev, M. Fisher, and A. Lisitsa. Implementing Invariant Search via Temporal Resolution. In *Proceedings of LPAR-2002*, pages 86–101. Springer Verlag, 2002. Lecture Notes in Computer Science, 2514.
6. J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang. Symbolic Model Checking: 10^{20} states and beyond. In *Proceedings of the 5th Symposium on Logic in Computer Science, Philadelphia*, June 1990.
7. E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, December 1999.
8. A. Degtyarev, M. Fisher, and B. Konev. Monodic Temporal Resolution. *ACM Transactions on Computational Logic* 7(1), January 2006.
9. A. Degtyarev, M. Fisher, and B. Konev. Handling equality in monodic temporal resolution. In *Proceedings of 10th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, volume 2850 of *Lecture Notes in Computer Science*, pages 214–228, Almaty, Kazakhstan, 2003. Springer-Verlag.
10. A. Degtyarev, M. Fisher, and A. Lisitsa. Equality and Monodic First-Order Temporal Logic. *Studia Logica*, 72(2):147–156, November 2002.
11. A. Degtyarev, M. Fisher, and B. Konev. A simplified clausal resolution procedure for propositional linear-time temporal logic. In U. Egly and C. G. Fermüller, editors, *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 2381 of *Lecture Notes in Computer Science*, pages 85–99. Springer-Verlag, July 30–August 1 2002.
12. C. Dixon. *Strategies for Temporal Resolution*. PhD thesis, Department of Computer Science, University of Manchester, Manchester M13 9PL, U.K., December 1995.
13. C. Dixon. Search Strategies for Resolution in Temporal Logics. In *Proceedings of the Thirteenth International Conference on Automated Deduction (CADE)*, volume 1104 of *Lecture Notes in Computer Science*, New Jersey, July/August 1996. Springer-Verlag.
14. C. Dixon. Temporal Resolution using a Breadth-First Search Algorithm. *Annals of Mathematics and Artificial Intelligence*, 22:87–115, 1998.
15. C. Dixon. Using Otter for Temporal Resolution. In *Advances in Temporal Logic*. Kluwer Academic Publishers, 1999.
16. C. Dixon. Using Temporal Logics of Knowledge for Specification and Verification — a Case Study. *Journal of Applied Logic*, 2005. To appear.
17. C. Dixon, M-C. Fernández Gago, M. Fisher, and W. van der Hoek. Using temporal logics of knowledge in the formal verification of security protocols. In *Proceedings of International Symposium on Temporal Representation and Reasoning (TIME)*. IEEE CS Press, 2004.
18. C. Dixon and M. Fisher. Clausal Resolution for Logics of Time and Knowledge with Synchrony and Perfect Recall. In *Proceedings of Joint International Conference on Temporal Logic and Advances in Modal Logic (AiML-ICTL)*, Leipzig, Germany, October 2000.
19. C. Dixon, M. Fisher, and M. Wooldridge. Resolution for Temporal Logics of Knowledge. *Journal of Logic and Computation*, 8(3):345–372, 1998.
20. E. A. Emerson. Temporal and Modal Logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 996–1072. Elsevier, 1990.

21. M. Fisher. A Resolution Method for Temporal Logic. In *Proc. Twelfth International Joint Conference on Artificial Intelligence (IJCAI)*, Sydney, Australia, 1991. Morgan Kaufman.
22. M. Fisher. Concurrent METATEM — A Language for Modeling Reactive Systems. In *Parallel Architectures and Languages, Europe (PARLE)*, Munich, Germany, June 1993. (Published in *Lecture Notes in Computer Science*, volume 694, Springer-Verlag).
23. M. Fisher. Representing and Executing Agent-Based Systems. In M. Wooldridge and N. R. Jennings, editors, *Intelligent Agents*. Springer-Verlag, 1995.
24. M. Fisher. An Introduction to Executable Temporal Logics. *Knowledge Engineering Review*, 11(1):43–56, March 1996.
25. M. Fisher. A Normal Form for Temporal Logic and its Application in Theorem-Proving and Execution. *Journal of Logic and Computation*, 7(4):429–456, August 1997.
26. M. Fisher. Implementing BDI-like Systems by Direct Execution. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. Morgan-Kaufmann, 1997.
27. M. Fisher. Temporal Development Methods for Agent-Based Systems. *Journal of Autonomous Agents and Multi-Agent Systems*, 10(1):41–66, January 2005.
28. M. Fisher, C. Dixon, and M. Peim. Clausal Temporal Resolution. *ACM Transactions on Computational Logic*, 2(1):12–56, January 2001.
29. M. Fisher and C. Ghidini. Programming Resource-Bounded Deliberative Agents. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. Morgan Kaufmann, 1999.
30. M. Fisher and C. Ghidini. The ABC of Rational Agent Programming. In *Proc. First International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 849–856. ACM Press, July 2002.
31. M. Fisher, C. Ghidini, and B. Hirsch. Organising Computation through Dynamic Grouping. In *Objects, Agents and Features*, volume 2975 of *Lecture Notes in Computer Science*, pages 117–136. Springer-Verlag, 2004.
32. M. Fisher, C. Ghidini, and B. Hirsch. Programming Groups of Rational Agents. In *Computational Logic in Multi-Agent Systems (CLIMA-IV)*, volume 3259 of *849–856*. Springer-Verlag, November 2004.
33. M. Fisher and T. Kakoudakis. Flexible Agent Grouping in Executable Temporal Logic. In *Proceedings of Twelfth International Symposium on Languages for Intensional Programming (ISLIP)*. World Scientific Press, 1999.
34. M. Fisher and R. Owens. From the Past to the Future: Executing Temporal Logic Programs. In *Proceedings of Logic Programming and Automated Reasoning (LPAR)*, St. Petersberg, Russia, July 1992. (Published in *Lecture Notes in Computer Science*, volume 624, Springer-Verlag).
35. M-C. Fernández Gago, U. Hustadt, C. Dixon, M. Fisher, and B. Konev. First-Order Temporal Verification in Practice. Accepted for publication in *Journal of Automated Reasoning*.
36. B. Hirsch, M. Fisher, C. Ghidini, and P. Busetta. Organising Software in Active Environments. In *Computational Logic in Multi-Agent Systems (CLIMA-V)*, volume 3487 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
37. B. Hirsch. *Programming Rational Agents*. PhD thesis, Department of Computer Science, University of Liverpool, United Kingdom, May 2005.
38. I. Hodkinson. Monodic Packed Fragment with Equality is Decidable. *Studia Logica*, 72(2):185–197, November 2002.
39. I. Hodkinson, F. Wolter, and M. Zakharyashev. Decidable fragments of first-order temporal logics. *Annals of Pure and Applied Logic*, 2000.
40. G.J. Holzmann. *Design and Validation of Computer Protocols*. Prentice-Hall, Englewood Cliffs, New Jersey, 1991.

41. G.J. Holzmann. The model checker spin. *IEEE Trans. on Software Engineering*, 23(5):279–295, May 1997. Special issue on Formal Methods in Software Practice.
42. U. Hustadt and B. Konev. TRP++ 2.0: A temporal resolution prover. In *Proceedings of Conference on Automated Deduction (CADE-19)*, number 2741 in *Lecture Notes in Artificial Intelligence*, pages 274–278. Springer, 2003.
43. U. Hustadt and R. A. Schmidt. Scientific benchmarking with temporal logic decision procedures. In D. Fensel, F. Giunchiglia, D. McGuinness, and M.-A. Williams, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Eighth International Conference (KR'2002)*, pages 533–544. Morgan Kaufmann, 2002.
44. U. Hustadt, B. Konev, A. Riazanov, and A. Voronkov. **TeMP**: A temporal monodic prover. In D. A. Basin and M. Rusinowitch, editors, *Proceedings of the Second International Joint Conference on Automated Reasoning (IJCAR 2004)*, volume 3097 of *Lecture Notes in Artificial Intelligence*, pages 326–330. Springer, 2004.
45. U. Hustadt, B. Konev, and R. A. Schmidt. Deciding monodic fragments by temporal resolution. In R. Nieuwenhuis, editor, *Proceedings of the 20th International Conference on Automated Deduction (CADE-20)*, volume 3632 of *Lecture Notes in Artificial Intelligence*, pages 204–218. Springer, 2005.
46. U. Hustadt and R. A. Schmidt. Formulae which highlight differences between temporal logic and dynamic logic provers. In E. Giunchiglia and F. Massacci, editors, *Issues in the Design and Experimental Evaluation of Systems for Modal and Temporal Logics*, Technical Report DII 14/01, pages 68–76. Dipartimento di Ingegneria dell'Informazione, Unversitá degli Studi di Siena, 2001.
47. M. Huth and M. Ryan. *Logic in Computer Science*. Cambridge University Press, 2000.
48. A. Kellett. *Implementation Techniques for Concurrent* METATEM. PhD thesis, Department of Computing and Mathematics, Manchester Metropolitan University, 2000.
49. A. Kellett and M. Fisher. Automata Representations for Concurrent METATEM. In *Proceedings of the Fourth International Workshop on Temporal Representation and Reasoning (TIME)*. IEEE Press, May 1997.
50. B. Konev, A. Degtyarev, C. Dixon, M. Fisher, and U. Hustadt. Mechanising first-order temporal resolution. *Information and Computation*, 199(1-2):55–86, 2005.
51. N. A. Lynch and M. R. Tuttle. An Introduction to Input/Output Automata. Technical Report MIT/LCS/TM-373, Laboratory for Computer Science, Massachusetts Institute of Technology, November 1988.
52. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1992.
53. Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer-Verlag, New York, 1995.
54. C. Nalon, C. Dixon, and M. Fisher. Resolution for Synchrony and No Learning. In *Proceedings of Advances in Modal Logic Confernece (AiML-5)*, Manchester, UK, September 2004.
55. A. Pnueli. The Temporal Logic of Programs. In *Proceedings of the Eighteenth Symposium on the Foundations of Computer Science*, Providence, USA, November 1977.
56. A. Prior. *Past, Present and Future*. Oxford University Press, 1967.
57. A. Riazanov and A. Voronkov. Vampire 1.1 (system description). In *Proc. IJCAR 2001*, pages 376–380. Volume 2083 of *Lecture Notes in Artificial Intelligence*, Springer-Verlag, 2001.
58. A. P. Sistla and E. M. Clarke. Complexity of propositional linear temporal logics. *ACM Journal*, 32(3):733–749, July 1985.

59. A. Szalas and L. Holenderski. Incompleteness of First-Order Temporal Logic with Until. *Theoretical Computer Science*, 57:317–325, 1988.
60. R. Tarjan. Depth-First Search and Linear Graph Algorithms. *SIAM Journal on Computing*, 1:146–160, 1972.
61. W. Visser, K. Havelund, G. Brat, and S. Park. Model checking programs. In *International Conference on Automated Software Engineering (ASE)*, September 2000.
62. M. Wooldridge. *Reasoning about Rational Agents*. MIT Press, 2000.