

# Belief Manipulation Through Propositional Announcements

**Aaron Hunter**  
 BC Institute of Technology  
 Burnaby, Canada  
 aaron\_hunter@bcit.ca

**Francois Schwarzenruber**  
 ENS Rennes  
 Bruz, France  
 francois.schwarzenruber@ens-rennes.fr

**Eric Tsang**  
 BC Institute of Technology  
 Burnaby, Canada  
 surplus.et@gmail.com

## Abstract

Public announcements cause each agent in a group to modify their beliefs to incorporate some new piece of information, while simultaneously being aware that all other agents are doing the same. Given a set of agents and a set of epistemic goals, it is natural to ask if there is a single announcement that will make each agent believe the corresponding goal. This problem is known to be undecidable in a general modal setting, where the presence of nested beliefs can lead to complex dynamics. In this paper, we consider not necessarily truthful public announcements in the setting of AGM belief revision. We prove that announcement finding in this setting is not only decidable, but that it is simpler than the corresponding problem in the most simplified modal logics. We then describe an implemented tool that uses announcement finding to control robot behaviour through belief manipulation.

## 1 Introduction

We are concerned with the manner in which the beliefs of a group of agents can be manipulated by making public announcements of propositional formulas. Suppose that we have  $n$  agents, each with a set of current beliefs and an epistemic goal  $\psi_i$ . We are interested in finding a single formula  $\phi$  such that each agent  $A_i$  will believe the corresponding goal  $\psi_i$  after a public announcement of  $\phi$ . We address two problems related to this *propositional announcement problem*:

1. Is there always a solution, even when the goals are collectively inconsistent?
2. What is the complexity of finding a solution?

We also describe an implemented announcement finding tool, and demonstrate how it can be used as a robot controller.

Public announcements have been addressed in modal logic [Plaza, 2007]. The problem that we address roughly corresponds to the notion of an *arbitrary public announcement* [Balbani *et al.*, 2007] in the tradition of Dynamic Epistemic Logic (DEL) [Ditmarsch *et al.*, 2008], except that we allow non-truthful announcements. The important distinction in our work is that we are interested in the simpler setting of propo-

sitional logic with AGM-style belief revision operators<sup>1</sup>. We focus on the AGM setting because we expect announcement finding to be computationally easier, as we avoid difficult problems due to nested belief. We argue that our approach is also appropriate for practical applications where the focus is on giving orders in the most efficient way possible.

This paper is an extension of [Hunter and Schwarzenruber, 2015], and it makes several contributions to existing work on the theory of belief change. First, we define arbitrary public announcements in an AGM setting. We demonstrate that the question is meaningful in this context, as there are natural applications where a reasonable treatment of announcements need not be concerned with nested beliefs. This paper can therefore be seen as part of the ongoing effort to connect the study of belief change between the AGM and DEL communities. The second contribution is a detailed analysis of the complexity of arbitrary announcements for a particular AGM revision operator, namely Dalal’s revision operator [Dalal, 1988]. Finally, we make a practical contribution, as we introduce an implemented tool for finding announcements with respect to a range of different belief revision operators.

**Motivating Example** Consider a domain involving a robot controller, and  $n$  robots that act independently. The only way the controller can communicate with the robots is through broadcast messaging<sup>2</sup>. If there are constraints on messaging in terms of cost or timing, then the controller may seek to minimize the number of messages sent to ensure all robots have the correct beliefs. Towards this end, it is useful to send messages that will impact the beliefs of each robot differently.

As a concrete example, suppose that there are two surveillance robots  $R_1$  and  $R_2$  that are tasked with patrolling a certain area. If there is a breach at the gate, we would like one robot to go check the gate while the other continues patrolling. In order to achieve the desired behaviour, the controller would like to broadcast a single alarm message “breach” that will lead each robot to simultaneously have the appropriate beliefs about how they should behave. In this paper, we formalize this kind of reasoning in an abstract setting.

<sup>1</sup>AGM stands for the highly influential approach to belief revision due to Alchourrón, Gärdenfors and Makinson [Alchourrón *et al.*, 1985].

<sup>2</sup>Assumed here to be synchronous.

## 2 Preliminaries

### 2.1 Belief Revision

One highly influential approach to belief revision is the AGM approach, in which revision is captured by an operator that satisfies a particular set of rationality postulates [Alchourrón *et al.*, 1985]. The beliefs of an agent are represented by a *belief set*, which is a logically closed set of propositional formulas over a vocabulary  $P$ . In this paper, we assume  $P$  is finite; hence  $K$ , the belief set, is equivalent to a single formula. An AGM belief revision operator takes a belief set  $K$  and formula  $\phi$  as input, and returns a new belief set that incorporates  $\phi$  while keeping as much of  $K$  as consistently possible. We write  $|K|$  to denote the set of interpretations where  $K$  is true.

Every AGM revision operator  $*$  has the property that, for any belief set  $K$ , there is an underlying total pre-order  $\prec_K$  over interpretations of  $P$  such that  $|K * \phi| = \min_{\prec_K}(\phi)$  [Katsuno and Mendelzon, 1992]. So AGM revision involves finding  $\prec_K$ -minimal states consistent with a particular formula. One important example is the Dalal operator [Dalal, 1988], denoted by  $*_d$ . For this operator, the underlying ordering is defined by the *Hamming distance* between interpretations. The Hamming distance between interpretations  $v, w$  is the number of propositional variables assigned different truth values by  $v$  and  $w$ . So then  $|K *_d \phi|$  is the set of models of  $\phi$  that have minimal Hamming distance from a model of  $K$ .

*Belief manipulation* is concerned with the ability of one agent to convince another to hold certain beliefs. With some revision operators, agents can be manipulated to believe a target formula  $\phi$  exclusively through indirect statements or evidence [Hunter and Booth, 2015]. The work in this paper can be seen as an exploration of multi-agent belief manipulation.

### 2.2 Dynamic Epistemic Logic

Dynamic Epistemic Logic (DEL) provides an alternative approach to reasoning about beliefs. In DEL, modal logics of knowledge and belief are combined with dynamic logic to model epistemic change. Standard modalities include  $K$  (knowledge),  $B$  (belief), and the parametrized announcement modality  $[\phi]$ . The semantics of  $[\phi]$  is defined as a restriction to worlds where  $\phi$  holds. For an overview of DEL, we refer the reader to [Ditmarsch *et al.*, 2008].

There have been versions of DEL that incorporate belief revision operators over worlds in the AGM sense; see, for example ([Benthem, 2007; Baltag and Smets, 2006; 2008]). In these logics, each agent  $a$  has an ordering over possible worlds and the semantics of  $[\phi]$  is defined by model transformations that adjust the orderings.

The most obvious distinction between the AGM approach and the DEL approach to belief revision is the fact that the DEL approach permits the representation of nested beliefs. As a result, it is much more expressive. Also, some of the properties we expect in the AGM setting are no longer valid. For example, the *Success* postulate for AGM revision states that  $\phi \in K * \phi$ . In a modal setting, this is equivalent to the schema  $[\phi]B\phi$ . This is not valid in most multi-agent epistemic logics, because an announcement of  $\phi$  often changes the world such that  $\phi$  should not actually be believed.<sup>3</sup>

<sup>3</sup>This is the case in so-called Moore sentences, for example.

The fact that announcing  $\phi$  does not ensure belief in  $\phi$  leads to a natural question: how can we get another agent to believe something through announcement? This question is addressed in DEL via the logic of arbitrary public announcements [Balbiani *et al.*, 2007]. Arbitrary announcement logic is similar, but not identical, to the problem addressed in this paper. First of all, we do not have belief revision operators in public announcement logic; we have only hard updates. Furthermore, in DEL, announcements are required to be truthful.

### 2.3 Known Complexity Results

In this section, we give some known complexity results related to public announcements. The first result is really the starting point for this work.

#### Theorem 1 ([French and van Ditmarsch, 2008])

*Satisfiability in the logic of arbitrary announcements is undecidable.*

Hence, if we have nested knowledge and a rich domain of Kripke structures, we can not hope to find public announcements to change the beliefs of agents in the desired manner. A simpler logic called *DLPA-APAL* (Dynamic Logic of Propositional Assignments with Arbitrary Public Announcement Logic) is defined in [Charrier and Schwarzenrüber, 2015]. In *DLPA-APAL*, Kripke structures are restricted to include just one world with each valuation, and epistemic relations are represented by programs over propositional assignments.

#### Theorem 2 ([Charrier and Schwarzenrüber, 2015])

*Satisfiability in DLPA-APAL is NEXPTIME-complete for the single announcement case.*

This result shows that restricting to reasoning over beliefs on valuations is decidable. Our goal in this paper is to show that if we restrict further, by removing nested beliefs, we can get a better complexity. This is true even if we consider willful manipulation, where announcements can be false.

## 3 The Propositional Announcement Problem

### 3.1 The Basic Problem

Given  $n$  agents and  $n$  AGM belief revision operators  $*_i$  as input; we are looking for the existence of a consistent formula  $\phi$  such that

$$\begin{aligned} K_1 *_1 \phi &\models \psi_1 \\ K_2 *_2 \phi &\models \psi_2 \\ &\vdots \\ K_n *_n \phi &\models \psi_n \end{aligned} \tag{1}$$

where  $K_i$  and  $\psi_i$  represent the belief set and the goal of the agent  $i$ , respectively. Obviously, if  $\bigwedge_i \psi_i$  is consistent, then we can just revise by this. But this is not always the case.

**Example** Consider the robot controller example over the vocabulary  $\{patrol, checkgate\}$ . We think of this vocabulary as defining a state machine, where each interpretation represents a state; the robot can have actions that are triggered by transitions to given states. This is a standard control mechanism for simple agents in a video game setting, and it can

function as a control mechanism for our simple robot agents as well. In our example, the *patrol* variable is true when the robot should be patrolling their area and the *checkgate* variable is true when the robot is supposed to check the gate.

Suppose we have two robots  $R_1$  and  $R_2$  with initial belief states defined as follows:

$$\begin{aligned} Bel(R_1) &= \{\neg patrol \vee checkgate\} \\ Bel(R_2) &= \{\neg patrol \vee \neg checkgate\} \end{aligned}$$

The controller believes there is a problem at the gate. Is there a formula that can be broadcast to immediately get  $R_1$  to check the gate while  $R_2$  patrols the grounds? In other words, is there a formula  $\phi$  such that:

$$\begin{aligned} \{\neg patrol \vee checkgate\} * \phi &\models checkgate \\ \{\neg patrol \vee \neg checkgate\} * \phi &\models \neg checkgate \wedge patrol \end{aligned}$$

The answer is yes; we can set  $\phi = patrol$ .

The preceding example is framed in the context of the robot controller, but it also demonstrates an important case for propositional announcement. In particular, it shows that there are cases where the goals are inconsistent, yet a solution is possible.

### 3.2 Existence of Solutions

In this section, we give some basic results related to the existence of solutions for the propositional announcement problem. The following result is immediate.

**Proposition 1** *For all  $n > 1$ , there are instances of (1) with no solution.*

**Proof** Assume the vocabulary  $\{p\}$ , let  $K_1 = \{p\}$ ,  $K_2 = \{\neg p\}$ ,  $\psi_1 = \neg p$  and  $\psi_2 = p$ . It follows immediately that any solution  $\phi$  satisfies  $\phi \models p \wedge \neg p$ , so there is no consistent solution.  $\square$

The following result may be less obvious, but it is similarly straightforward. The idea is that, if there are “enough” propositional variables, we can set up a set of inconsistent goals that can be guaranteed by an appropriate revision.

**Theorem 3** *Let  $\mathbf{P}$  be a vocabulary and let the number of agents  $n$  satisfy  $n < 2^{|\mathbf{P}|}$ . Then there is an instance of (1) over  $\mathbf{P}$  with inconsistent goals that has a solution.*

**Proof** Let  $P_0, \dots, P_n$  be a list of mutually inconsistent maximal conjunctions over literals of  $\mathbf{P}$ ; we know that such a list exists, because there are  $2^{|\mathbf{P}|}$  distinct maximal conjunctions of literals, and we know that  $2^{|\mathbf{P}|} > n$  by assumption. For each  $i$ , we set:

$$\begin{aligned} K_i &= (\neg P_0 \vee P_i) \wedge \bigwedge_{j \neq i} \neg P_j \\ \psi_i &= P_i \wedge \bigwedge_{j \neq i} \neg P_j \end{aligned}$$

Note that the conjunction  $\bigwedge_i \psi_i$  is clearly inconsistent. If we set  $\phi = P_0$ , then it is easy to verify that  $K_i * \phi \models \psi_i$  for each  $i$ , which is what we wanted to show.  $\square$

These results demonstrate that the general propositional announcement problem is non-trivial. There are instances with

no solution and there are instances with solutions, and we can not easily identify which is the case without considering the possible revisions. This leads to the following question: How hard is it to determine if there is a solution to a given instance of the problem? We turn to this problem in the next section.

## 4 Complexity

### 4.1 The Decision Problem

In this section, we restrict the problem slightly by requiring that all agents have the same revision operator. Given  $n$  agents and an AGM revision operator  $*$ , we are looking for the existence of a consistent formula  $\phi$  such that

$$\begin{aligned} K_1 * \phi &\models \psi_1 \\ K_2 * \phi &\models \psi_2 \\ &\vdots \\ K_n * \phi &\models \psi_n \end{aligned} \tag{2}$$

where  $K_i$  and  $\psi_i$  represent the belief set and the goal of the agent  $i$ , respectively. In principle,  $*$  could be any shared revision operator.

In practice, we will often be interested in the complexity of *finding* announcements; but we must first consider the corresponding decision problem.

#### The Propositional Announcement Problem ( $PAP(*)$ )

##### Input:

- An integer  $n$
- A list  $K_1, \dots, K_n$  of formulas (initial beliefs).
- A list  $\psi_1, \dots, \psi_n$  of formulas (goals).

##### Output:

- Yes, if there exists  $\phi$  satisfying (2)
- No, otherwise.

We refer to this problem as  $PAP(*)$  to emphasize that it depends on some given operator  $*$  on belief sets. We normally assume that  $*$  is an AGM revision operator, but this need not be the case in general.

### 4.2 Announcement Existence

For the moment, we are interested in analyzing a simple case to obtain the most efficient algorithm possible. Hence, we consider Dalal’s well-known revision operator based on Hamming distance [Dalal, 1988]. As in §2.1, we let  $*_d$  denote Dalal’s revision operator.

We present a non-deterministic algorithm that decides  $PAP(*_d)$ . In the algorithm, we use non-deterministic choice to select the minimum distance  $d_i$  between  $K_i$  and  $\psi_i$ . Also note that  $d(K, v)$  denotes the minimum Hamming distance between the set of models of formula  $K$  and interpretation  $v$ .

#### EXIST-ANN ( $K_1, \dots, K_n, \psi_1, \dots, \psi_n$ )

0. Let  $m$  be the size of the underlying input vocabulary of  $K_1, \dots, K_n, \psi_1, \dots, \psi_n$
1. Guess  $d_1, \dots, d_n \in \{0, 1, \dots, m\}$
2. Guess valuations  $v_1, \dots, v_n$
3. For all  $i, j \in \{1, \dots, n\}$   
If  $d(K_j, v_i) < d_j$ , reject.

4. For all  $i \in \{1, \dots, n\}$   
If  $d(K_i, v_i) > d_i$ , reject.
5. For all  $i, j \in \{1, \dots, n\}$   
If  $d(K_j, v_i) = d_j$  and  $v_i \not\models \psi_j$ , reject.
6. Accept.

We need to prove that  $\text{EXIST\_ANN}(\bar{K}, \bar{\psi})$  actually produces the desired result. In the proof, we write  $\min_{\leq K_i}(\phi)$  for the minimal Hamming distance from a model of  $\bar{\phi}$  to a model of  $K_i$ .

**Theorem 4** Let  $\bar{K} = K_1, \dots, K_n$  and let  $\bar{\psi} = \psi_1, \dots, \psi_n$  be sequences of formulas. Then  $\text{EXIST\_ANN}(\bar{K}, \bar{\psi})$  accepts if and only if there exists  $\phi$  such that  $K_i *_{\bar{d}} \phi \models \psi_i$  for each  $i$ .

**Proof** Suppose that  $\text{EXIST\_ANN}(\bar{K}, \bar{\psi})$  accepts, and let us consider an accepting run. Let  $\phi$  be a formula with  $\|\phi\| = \bigcup_i \{v_i\}$ , where the valuations  $v_i$  are those from an accepting run guessed on line 2. By line 3,

$$\min_{\leq K_i}(\phi) \geq d_j$$

because if this is not the case, the algorithm would have rejected the input. But then, by line 4,

$$\min_{\leq K_i}(\phi) = d_j.$$

In other words, for each  $i$ , the models of  $\phi$  that are models of  $K_i$  are of distance  $d_i$  from  $K_i$ . Finally, since the conditions of all lines are false, it follows that for every  $v$  such that  $v \models \phi$ , it must be the case that  $d(K_j, v) = d_j$  implies  $v \models \psi_j$ . Hence, the minimal models of  $\phi$  with respect to  $K_i$  are also models of  $\psi_i$ . So  $K * \phi \models \psi_i$ , which is what we wanted to show.

Now suppose that there exists  $\phi$  such that  $K_i *_{\bar{d}} \phi \models \psi_i$  for each  $i$ . For each  $i$ , define  $d_i = \min_{\leq K_i}(\phi)$ . For each  $i$ , let  $v_i$  denote one of the valuations of  $\psi_i$  such that  $d(K_i, v_i) = d_i$ . We know that such a  $v_i$  exists, since  $K_i *_{\bar{d}} \phi \models \psi_i$ . If we guess  $\bar{d} = d_1, \dots, d_n$  and  $\bar{v} = v_1, \dots, v_n$  on lines 1 and 2 of the algorithm, then we have an accepting run.  $\square$

So  $\text{EXIST\_ANN}(\bar{K}, \bar{\psi})$  does what we want it to do: it determines if there is a public announcement  $\phi$  such that each agent  $i$  believes  $\psi_i$  following revision by  $\phi$ . We can also say something about the complexity. The initial guessing is a polynomial number of guesses with respect to the input length. Note that in the algorithm, tests  $d(K_j, v_i) < d_j$  etc. can be answered by an NP oracle. Hence, the algorithm intuitively seems to lie in the complexity class  $\Sigma_2^P = NP^{NP}$ . The following result makes this claim precise.

**Theorem 5**  $\text{PAP}(*_{\bar{d}})$  is in  $\Sigma_2^P$ .

**Proof** After the initial guesses, the algorithm clearly runs in polynomial time other than the checks of the form  $d(K, v) < d$  at several stages. But recall that we are using the Hamming distance here, so this check can be performed as follows. Guess a set of atomic propositional variables of size less than  $d$ , and let  $v'$  be the interpretation obtained from  $v$  by switching the truth values of these variables. If  $v' \models K$ , then the minimum distance between  $K$  and  $v$  is less than  $d$ . This process is in NP.

Hence the entire algorithm runs in time NP with a polynomial number of calls to an NP oracle. This gives the complexity  $NP^{NP} = \Sigma_2^P$ , which was the desired result.  $\square$

Recall that announcement finding was *NEXPTIME*-complete in the modal case, even if we restricted possible worlds to be valuations. So this result shows that the problem is much easier in the AGM case. Note that hardness of  $\text{PAP}(*_{\bar{d}})$  is left as an open issue.

So far, we have restricted attention to Dalal's revision operator. This choice can be justified to some extent by the fact that Dalal's operator permits iterated revision, which we anticipate will be important in practice for a robot controller. Inspection of our proofs shows that the actual properties of the Hamming distance operator only occur in two places.

- When the minimum distance is guessed in the algorithm.
- In the checks of the form  $d(K, v) < d$ .

The advantage of Dalal's revision operator in both cases is that we have a polynomial number of steps, because the distance only ranges from 0 to the size of the vocabulary. However, this is obviously not true for other revision operators; we return to this point in the conclusion.

## 5 Implementation

### 5.1 Basic Approach

In this section, we are interested in a different question: can we use announcement finding to implement a feasible solution to a problem of practical interest? To answer this question, we describe a practical tool that uses announcement finding as the basis for a simulated robot controller.

We describe AnnB, an automated robot controller that calculates public announcements to manipulate the beliefs of agents and force them to perform particular tasks. The key points of operation for AnnB are as follows:

- Each agent has a set of beliefs and a revision operator.
- There is a fixed set of movement behaviours.
- Believing certain formulas causes an agent to perform behaviours.

In this environment, giving an agent information for revision can trigger changes in activity. Given a set of goal behaviours, AnnB computes the announcement required to change the beliefs of the agents in a way that gets each agent to perform a desired action.

AnnB is written in Kotlin<sup>4</sup>. It is built on the publicly available libraries of GenB [Hunter and Tsang, 2016]. Briefly, GenB is a general belief revision solver that is able to calculate the result of any AGM belief revision operation. There are built-in sample operators, but GenB also allows the user to specify any revision operator by giving the corresponding faithful assignment. Hence, building AnnB on the GenB framework gives us considerable flexibility.

<sup>4</sup>Kotlin is a variant of java, available at [www.kotlinlang.org](http://www.kotlinlang.org).

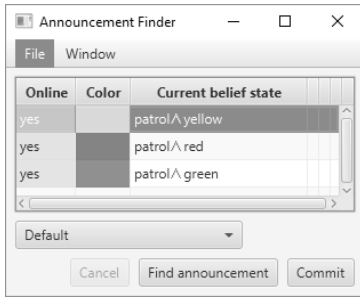


Figure 1: Defining Agents

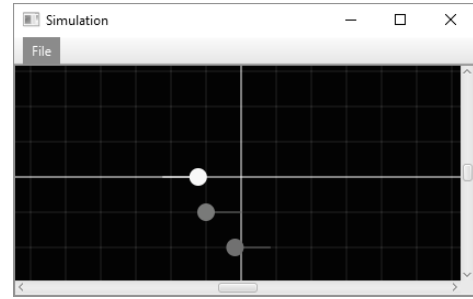


Figure 2: Simulation View

## 5.2 Defining Behaviours

The first step is to specify the *agents*. This is done through a simple dialog box in which agents are added and given an initial belief state, as in Figure 1. By convention, we introduce a propositional variable that matches with the colour of the agent. This variable will initially only be believed by an agent of that colour. Figure 2 gives the simulation view.

After setting up the set of agents, we define behaviours through the *behaviour definition table*. Essentially, the behaviour definition table associates the truth of propositional formulas with movement on screen. Built-in algorithms define standard movement patterns such as *wander*, *patrol*, and *find-path*. The behaviour of an agent is determined by checking if it believes each formula in the behaviour window, starting from the top and working down. An example is given in Figure 3. In this example, the yellow agent is initially patrolling at position  $y = 0$ . The green agent is initially patrolling at position  $y = 2$ , because it does not believe the first two propositions. The final line in the behaviour table is the default behaviour that any agent should display if it does not believe any of the previous propositions.

## 5.3 Algorithm

The algorithm used by AnnB is based on **EXIST\_ANN**; however, there are several changes. First, the input includes a plausibility function  $P_i$  for each agent  $i$ . Informally,  $P_i(v) < P_j(w)$  indicates that agent  $i$  considers the valuation  $v$  to be more plausible than  $w$ . This function defines a faithful assignment, which allows any revision operator to be specified. The second difference is that the algorithm is deterministic, so it loops through possible values rather than guessing.

The following is a high level description of the algorithm. In practice, we iterate over the  $d_i$ 's starting each at 0 and incrementing one position at a time. We write  $form(\bar{v})$  to mean the disjunctive normal form sentence obtained from a sequence of valuations in a the natural way.

**FIND\_ANN\_DET** ( $K_1, \dots, K_n, P_1, \dots, P_n, \psi_1, \dots, \psi_n$ )

0. Let  $m = \max_{i,s} P_i(s)$
1. For each  $\bar{d} = (d_1, \dots, d_n)$  with  $0 \leq d_i \leq m$ :
2. For each  $n$ -tuple  $\bar{v} = (v_1, \dots, v_n)$  of valuations:
  3. If  $P_j(v_i) < d_j$  for any  $i, j$ , continue.
  4. If  $P_i(v_i) > d_i$  for any  $i$ , continue.
  5. If  $P_j(v_i) = d_j$  and  $v_i \not\models \psi_j$  for any  $i, j$ , continue.
6. Return  $\phi = form(\bar{v})$
7. Return "no solution."

We can prove that this algorithm is correct.

**Theorem 6** Let  $\bar{K} = K_1, \dots, K_n$ , and  $\bar{\psi} = \psi_1, \dots, \psi_n$  be sequences of formulas. Let  $P_1, \dots, P_n$  be a sequence of plausibility functions. If  $\phi$  is the formula returned by **FIND\_ANN\_DET**( $\bar{K}, \bar{P}, \bar{\psi}$ ) then  $K_i *_i \phi \models \psi_i$  for each  $i$ , where  $*_i$  is the AGM revision operator obtained from the plausibility function  $P_i$ .

**Proof** Parallel to the proof of Theorem 4.  $\square$

In Figure 4, we give an illustration of the way this algorithm is used. Suppose that we would like the yellow agent to stop patrolling, and guard the north entrance. At the same time, we would like the red agent to continue patrolling. We can specify this behaviour in a declarative manner by indicating that *yellow* should no longer believe patrol, whereas *red* should continue to believe it. We can then find a suitable announcement by clicking the Find Announcement button. Clicking on Commit changes the beliefs of each agent, and causes the new behaviour to be exhibited in the simulation window.

## 5.4 Performance and Functionality

One advantage of AnnB is that it is able to solve the announcement problem for *all* AGM revision operators. However, if we would like to do a second revision, we need to supply a new set of plausibility orderings. While it would be possible to do this in accordance with a particular iterated revision operator, it must be done manually in the current software.

Due to the complexity result in Theorem 5, we can not expect a general announcement solver to find solutions quickly for large instances. For the application at hand, this is not a major limitation. We are able to define behaviours using a small set of variables and patterns, inspired by standard character movement in video games. However, run times for large instances quickly become unmanageable. To analyze the efficiency, we need to define and specify suitable bench mark problems in our domain of interest; such a specification is beyond the scope of the present paper. We leave a systematic analysis of the practical efficiency of AnnB for future work.

It is not our intention to suggest that AnnB can function as a complete robot control system. But it can play an important role. In general, the coordination of autonomous robots involves *replanning* when obstacles are faced [Coltin and Veloso, 2013]. When an obstacle is encountered by one

Proposition	Behaviour
patrol ∧ yellow	Patrol waypoints=[Guard(x=5, y=0, direction=NORTH), Guard(x=-5, y=0, direction=NORTH)]
patrol ∧ red	Patrol waypoints=[Guard(x=5, y=1, direction=NORTH), Guard(x=-5, y=1, direction=NORTH)]
patrol	Patrol waypoints=[Guard(x=5, y=2, direction=NORTH), Guard(x=-5, y=2, direction=NORTH)]
1	Guard(x=0, y=2, direction=NORTH)

Figure 3: Defining Behaviours

robot, a broadcast announcement is the simplest way to inform the others. By using an announcement finder, we can select an announcement that ensures all goals are achieved.

## 6 Discussion

### 6.1 Relation with Modal Approaches

Belief revision has been addressed in several different modal frameworks, including [Benthem, 2007; Baltag and Smets, 2006; Caridroit *et al.*, 2015; 2016]. The most obvious difference in our approach is that we are unable to capture nested beliefs. However, it is also significant that we are not constrained with any notion of the “actual” world. Indeed, the announcer is just searching for a formula that suitably interacts with the revision operators; this gives us a more flexible setting to address problems of belief manipulation.

Although we work in a propositional framework, our results are closely connected to modal work on announcement finding. In particular, we remark that we can actually reformulate  $PAP(*)$  as a variant of the model checking problem of DEL with plausibility models ([Benthem, 2007; Baltag and Smets, 2006]). We create a (potentially exponentially bigger than  $K_1, \dots, K_n$ ) pointed Kripke model  $\mathcal{M}, w$  such that the valuations of the most plausible worlds in  $w$  are exactly models of  $K_i$  for all agents  $i \in \{1, \dots, n\}$ . Then we ask whether there exists a propositional formula  $\phi$  such that  $\mathcal{M}, w \models [\phi]_{BR}(B_1\psi_1, \dots, B_n\psi_n)$  where  $[\phi]_{BR}$  is the announcement operator that semantically performs the belief revision for all agents and  $B_i$  is the belief operator for agent  $i$ .

### 6.2 Future Work

We have only considered the complexity of announcement finding in the case of the Dalal’s revision operator, where the

number of levels in the plausibility ordering is at most  $|\mathbf{P}|$  (the size of the vocabulary). In general, the worst case complexity to check all levels of a total pre-order over a vocabulary of size  $\mathbf{P}$  will be  $2^{|\mathbf{P}|}$ . If we are actually required to make this many checks, then the decision problem for other revision operators would lie at a higher level of the polynomial hierarchy. This would not be surprising, as it is known that Dalal’s operator is computationally simpler than other AGM revision operators [Eiter and Gottlob, 1992]. We leave the complexity of additional AGM operators for future work.

It would also be useful to consider announcement finding for iterated revision in the tradition of Darwiche and Pearl [Darwiche and Pearl, 1997; Jin and Thielscher, 2007]. In this context, the problem takes  $n$  agents and  $n$  Darwiche-Pearl revision operators  $*_i$  as input. We are looking for the existence of a consistent formula  $\phi$  such that  $\psi_i \in \min(\prec_i *_i \phi)$  for each  $i$ , where  $\prec_i$  and  $\psi_i$  represent the epistemic state and the goal of the agent  $i$ , respectively. Following [Liberatore, 1997], we expect the complexity will be higher in this case.

At a practical level, improvements can be made to AnnB. One obvious direction for future work would be to improve the running time of the algorithm through suitable heuristics, or by using an efficient SAT solver for the NP calculations. We are also interested in moving past simulation to a prototype controller for real, physical robots using our tool.

### 6.3 Conclusion

We have considered arbitrary public announcements in the setting of propositional belief revision. We have proved that the decision problem for announcement finding in this setting is computationally simpler than it is in a modal setting, even under very strong simplifying assumptions. We have argued that this result is important, because revision by announcements can be used in practical applications where the notion of fully nested beliefs is not required. This has been demonstrated in practice through the development of robot control software based on announcement finding.

## Acknowledgements

Many thanks are due to Hans van Ditmarsh, who hosted the first two authors as visitors at the Lorraine Research Laboratory in Computer Science and its Applications (LORIA) while they wrote an early version of this paper.

Online	Color	Current belief st...	Target belief state	Belief revision operator	Revised belief state
yes		{patrol, yellow}	{yellow}	hamming distance	{yellow}
yes		{patrol, red}	{patrol, red}	hamming distance	{patrol, red}

Models:

Buttons: Cancel, Find announcement, Commit

Figure 4: Finding Announcements

## References

- [Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [Balbiani *et al.*, 2007] Philippe Balbiani, Alexandru Baltag, Hans van Ditmarsch, Andreas Herzig, Tomohiro Hoshi, and Tiago De Lima. What can we achieve by arbitrary announcements? a dynamic take on fitch’s knowability. In *Proceedings of the Conference on Theoretical Aspects of Rationality and Knowledge (TARK-07)*, pages 42–51, 2007.
- [Baltag and Smets, 2006] Alexandru Baltag and Sonja Smets. Dynamic belief revision over multi-agent plausibility models. In *Proceedings of LOFT*, volume 6, pages 11–24, 2006.
- [Baltag and Smets, 2008] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. *Texts in logic and games*, 3:9–58, 2008.
- [Benthem, 2007] Johan van Benthem. Dynamic logic for belief revision. *Journal of applied non-classical logics*, 17(2):129–155, 2007.
- [Caridroit *et al.*, 2015] Thomas Caridroit, Sébastien Konieczny, Tiago de Lima, and Pierre Marquis. Private revision in a multi-agent setting. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1677–1678, 2015.
- [Caridroit *et al.*, 2016] Thomas Caridroit, Sébastien Konieczny, Tiago de Lima, and Pierre Marquis. On distances between kd45n kripke models and their use for belief revision. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, pages 1053–1061, 2016.
- [Charrier and Schwarzenrüber, 2015] Tristan Charrier and François Schwarzenrüber. Arbitrary public announcement logic with mental programs. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multi-Agent Systems*, 2015.
- [Coltin and Veloso, 2013] Brian Coltin and Manuela Veloso. Towards replanning for mobile service robots with shared information. In *Proceedings of the AAMAS’10 Workshop on Autonomous Robots and Multirobot Systems (ARMS)*, 2013.
- [Dalal, 1988] Mukesh Dalal. Investigations into a theory of knowledge base revision. In *Proceedings of the National Conference on Artificial Intelligence (AAAI88)*, pages 475–479, 1988.
- [Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997.
- [Ditmarsch *et al.*, 2008] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, 2008.
- [Eiter and Gottlob, 1992] Thomas Eiter and Georg Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57(2-3):227–270, 1992.
- [French and van Ditmarsch, 2008] Tim French and Hans van Ditmarsch. Undecidability for arbitrary public announcement logic. In *Advances in Modal Logic*, pages 23–42, 2008.
- [Hunter and Booth, 2015] Aaron Hunter and Richard Booth. Trust-sensitive belief revision. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3062–3068, 2015.
- [Hunter and Schwarzenrüber, 2015] Aaron Hunter and François Schwarzenrüber. Arbitrary announcements in propositional belief revision. In *Proceedings of the International Workshop on Defeasible and Ampliative Reasoning (DARE)*, 2015.
- [Hunter and Tsang, 2016] Aaron Hunter and Eric Tsang. GenB: A general solver for AGM revision. In *Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA)*, pages 564–569, 2016.
- [Jin and Thielscher, 2007] Yi Jin and Michael Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.
- [Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Albert O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(2):263–294, 1992.
- [Liberatore, 1997] Paolo Liberatore. The complexity of iterated belief revision. In *Proceedings of the 6th International Conference on Database Theory*, pages 276–290, 1997.
- [Plaza, 2007] Jan Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007.