

Generalized Arrow Update Logic

Barteld Kooi
University of Groningen
B.P.Kooi@rug.nl

Bryan Renne
University of British Columbia
bryan@renne.org

ABSTRACT

This paper presents a logic for reasoning about information change in multi-agent settings based on epistemic arrow deletion in Kripke models.

Categories and Subject Descriptors

I.2.0 [Artificial Intelligence]: General—*Philosophical foundations*; I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods—*Modal logic, Representation languages*; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Theory

Keywords

Public announcement logic, Dynamic Epistemic Logic, update expressivity, update succinctness

1. INTRODUCTION

There are two main lines of work on modal logics for public announcements. The first, due to Plaza [10, 11], defines the public announcement of φ as the operation $[\varphi]_p$ on Kripke models that deletes all $\neg\varphi$ -worlds (along with any epistemic arrows pointing to or from these worlds). The second, due to Gerbrandy and Groeneveld (GG) [6], defines the public announcement of φ as the operation $[\varphi]_g$ on Kripke models that deletes only the epistemic arrows that lead to $\neg\varphi$ -worlds. Plaza's line of work led to the development of the popular approach to *Dynamic Epistemic Logic* (DEL) due to Baltag, Moss, and Solecki [1, 2, 13]. The BMS approach defines operations on Kripke models that, intuitively speaking, perform a finite number of Plaza announcements, each of which an agent may entertain as one of the possible candidates for the one Plaza announcement that did in fact occur. BMS updates can be used to reason about a variety

ACM COPYRIGHT NOTICE. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or permissions@acm.org. TARK 2011, July 12–14, 2011, Groningen, The Netherlands. Copyright ©2011 ACM. ISBN 978-1-4503-0707-9, \$10.00.

of communicative types, including public and private communications, those with and without deception, and many others [1, 12, 13].

Plaza announcements and their BMS generalizations are based on the principle of deleting worlds, whereas GG announcements are based on the principle of deleting arrows. In a previous paper [8], we the authors argued that it is sometimes more convenient to delete arrows than it is to delete worlds. This led us to a generalization of GG announcements based on our notion of *arrow updates*. Arrow updates specify a partition of epistemic arrows into two categories: those we want to keep and those we want to delete. This specification consists of a finite number of triples (φ, a, φ') indicating that the a -arrows pointing from φ -worlds to φ' -worlds are to be kept; arrows that do not meet the conditions of any triple are to be deleted. Arrow updates define operations on Kripke models in which, intuitively speaking, agents respond to new information according to a common arrow-deletion policy that can be identified with a certain multi-agent belief change. While we showed that every arrow update produces a model-change operation equivalent to the change produced by a BMS update [8], common knowledge of the arrow update arrow-deletion policy seems to make it impossible for arrow updates to express private communications familiar from the BMS framework.

In this paper, we extend our previous work by developing GAUL, the theory of *Generalized Arrow Update Logic* (Section 2). GAUL arrow updates generalize our previous work by dropping the assumption of a common update policy, allowing us to prove that arrow updates are just as expressive as BMS updates (Section 3). There is therefore a perfect match between the BMS generalization of Plaza announcements and our arrow update generalization of GG announcements. Further, we show that GAUL arrow updates are at worst poly-exponentially less succinct than BMS updates, though this improves to being at worst polynomially less succinct if the Plaza announcements making up the BMS updates do not themselves contain additional BMS updates or if we allow arrow updates to contain BMS updates (Section 3). We also show that GAUL arrow updates are sometimes exponentially more succinct than BMS updates (Section 3). We then provide some concrete examples of scenarios in which we believe it is more convenient to use GAUL than it is to use DEL (Section 4), and we indicate some directions for further study (Section 5).

2. GAUL SYNTAX AND SEMANTICS

The starting point of GAUL is multi-modal epistemic logic,

which is interpreted over Kripke models. We fix a finite nonempty set \mathcal{A} of *agents* and a nonempty set \mathcal{P} of *propositional variables*. A *Kripke model* M is a tuple (W^M, R^M, V^M) consisting of a nonempty set W^M of *worlds*, an epistemic possibility function $R^M : \mathcal{A} \times W^M \rightarrow \wp(W^M)$ (notation: $R_a^M(w) \stackrel{\text{def}}{=} R^M(a, w)$), and a propositional valuation $V^M : \mathcal{P} \rightarrow \wp(W^M)$. A Kripke model M and world $w \in W^M$ form a *pointed Kripke model* (M, w) with *point* w .

The central notion in the semantics of GAUL is the *pointed arrow update* (U, o) , which consists of an *arrow update* U and the *point* o . Each pointed arrow update prescribes an operation on Kripke models that may be identified with a multi-agent belief change. When (U, o) occurs at a pointed Kripke model (M, w) , the arrow update U specifies a set O^U of possible *outcomes*. One of these is the *actual outcome* $o \in O^U$ —the point of (U, o) —though there may be other outcomes in O^U as well. Each outcome $o' \in O^U$ generates an o' -indexed copy (w, o') of world w ; this copy is to be thought of as the way world w would come to be if o' were the actual outcome. Outcomes $o' \in O^U$ also generate copies of the other worlds $v \in W^M$ in the same way. For each agent $a \in \mathcal{A}$ and outcome $o_1 \in O^U$, the arrow update U specifies a finite set $\mathbf{a}_a^U(o_1)$ of triples $(\varphi_1, o_2, \varphi_2)$ indicating that an existing a -arrow in M from world v_1 to world v_2 will bring about an a -arrow from copy (v_1, o_1) to copy (v_2, o_2) if and only if v_1 satisfies the *source condition* φ_1 and v_2 satisfies the *target condition* φ_2 . In this way, agent uncertainty between worlds in M is carried over to outcome-indexed copies of those worlds if and only if the original worlds meet at least one source-and-target specification for the corresponding outcomes. Hence we may specify a different arrow-deletion policy for each agent and each pair of outcomes. This allows us to drop the assumption of common arrow-deletion policy that was present in our previous work on *single-outcome arrow updates* (where $|O^U| = 1$) [8].¹

DEFINITION 2.1 (ARROW UPDATE). *Let L be a language. An arrow update U for L is a pair (O^U, \mathbf{a}^U) consisting of a finite nonempty set O^U of outcomes and an arrow function $\mathbf{a}^U : \mathcal{A} \times O^U \rightarrow L \times O^U \times L$ having a finite graph (notation: $\mathbf{a}_a^U(o) \stackrel{\text{def}}{=} \mathbf{a}^U(a, o)$). The tuple $(\varphi, o', \varphi') \in \mathbf{a}_a^U(o)$ is an a -arrow with source o , source condition φ , target o' , and target condition φ' . An arrow update U for L and an outcome $o \in O^U$ form a *pointed arrow update* (U, o) for L with point o . Let $\mathcal{U}(L)$ be the set of arrow updates for L and $\mathcal{U}_*(L)$ be the set of pointed arrow updates for L .*

To obtain the language of GAUL, we extend the language of multi-modal logic by adding pointed arrow updates as modal operators.

¹Agent a 's arrow-deletion policy is common knowledge in a single-outcome arrow update U with outcome set $O^U = \{o\}$ because a has exactly one arrow-deletion policy $\mathbf{a}_a^U(o)$ for the one and only possible outcome $o \in O^U$, the only outcome the agents commonly consider possible. A single-outcome arrow update U with outcome set $O^U = \{o\}$ may be identified with the finite set $\bigcup_{a \in \mathcal{A}} \{(\varphi, a, \varphi') \mid (\varphi, o, \varphi') \in \mathbf{a}_a^U(o)\}$. A triple (φ, a, φ') in this set indicates that a -arrows in a Kripke model M from φ -worlds to φ' -worlds are to be kept; arrows in M that do not satisfy the conditions of any triple in the set are to be deleted. We used finite sets of this form in our previous work on single-outcome arrow updates [8].

AXIOM SCHEMES

- CL. Classical Propositional Logic
BK. $\Box_a(\varphi \rightarrow \psi) \rightarrow (\Box_a\varphi \rightarrow \Box_a\psi)$
U1. $[U, o]p \leftrightarrow p$ for $p \in \mathcal{P}$
U2. $[U, o]\neg\varphi \leftrightarrow \neg[U, o]\varphi$
U3. $[U, o](\varphi \wedge \psi) \leftrightarrow ([U, o]\varphi \wedge [U, o]\psi)$
U4. $[U, o]\Box_a\varphi \leftrightarrow \bigwedge_{(\psi, o', \psi') \in \mathbf{a}_a^U(o)} (\psi \rightarrow \Box_a(\psi' \rightarrow [U, o']\varphi))$

RULES

$$\frac{\varphi \rightarrow \psi \quad \varphi}{\psi} \text{ (MP)} \quad \frac{\varphi}{\Box_a\varphi} \text{ (BN)} \quad \frac{\varphi}{[U, o]\varphi} \text{ (UN)}$$

Table 1: The theory GAUL

DEFINITION 2.2 (LANGUAGE OF GAUL). *Let \mathcal{L}^0 be the language of multi-modal logic built up from atoms in \mathcal{P} using negation \neg , conjunction \wedge , and an epistemic modal \Box_a for each $a \in \mathcal{A}$. Other Boolean and modal connectives are defined as usual. By induction on i , define language \mathcal{L}^{i+1} to be the set of formulas φ formed by the grammar*

$$\varphi ::= \psi \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [U, o]\varphi,$$

where $\psi \in \mathcal{L}^i$, $a \in \mathcal{A}$, and $(U, o) \in \mathcal{U}_*(\mathcal{L}^i)$. Finally, define the set $\mathcal{L} \stackrel{\text{def}}{=} \bigcup_{i \in \mathbb{N}} \mathcal{L}^i$ of formulas of GAUL, the set $\mathcal{U} \stackrel{\text{def}}{=} \mathcal{U}(\mathcal{L})$ of arrow updates, and the set $\mathcal{U}_* \stackrel{\text{def}}{=} \mathcal{U}_*(\mathcal{L})$ of pointed arrow updates. To say that $\varphi \in \mathcal{L}$ is reduced means that $\varphi \in \mathcal{L}^0$.

DEFINITION 2.3 (SEMANTICS OF GAUL). *The binary truth relation \models between pointed Kripke models and formulas is defined by induction on formula construction. The classical cases are defined as usual; the modal cases are defined as follows.*

- $M, w \models \Box_a\varphi$ means $M, w' \models \varphi$ for each $w' \in R_a^M(w)$.
- $M, w \models [U, o]\varphi$ means $M * U, (w, o) \models \varphi$, where $M * U$ is defined as follows.

$$\begin{aligned} W^{M * U} &\stackrel{\text{def}}{=} W^M \times O^U \\ R_a^{M * U}((w, o)) &\stackrel{\text{def}}{=} \{(w', o') \in R_a^M(w) \times O^U \mid \\ &\quad \exists (\varphi, o', \varphi') \in \mathbf{a}_a^U(o) : \\ &\quad (M, w \models \varphi \ \& \ M, w' \models \varphi')\} \\ V^{M * U}(p) &\stackrel{\text{def}}{=} V^M(p) \times O^U \end{aligned}$$

Validity in a model $M \models \varphi$ means $M, w' \models \varphi$ for each $w' \in W^M$. Validity $\models \varphi$ means $M', w' \models \varphi$ for each pointed Kripke model (M', w') .

The axiomatic theory GAUL is defined in Table 1.

THEOREM 2.4. $\vdash \varphi$ implies $\models \varphi$ for each $\varphi \in \mathcal{L}$.

PROOF. We will only show soundness of Axiom U4. From left to right. Take an arbitrary pointed Kripke model (M, w) and suppose that $M, w \models [U, o]\Box_a\varphi$. Also take an arbitrary a -arrow $(\psi, o', \psi') \in \mathbf{a}_a^U(o)$ and suppose that $M, w \models \psi$. All that remains to be shown at this point is that $M, w \models \Box_a(\psi' \rightarrow [U, o']\varphi)$. In order to show this, take an arbitrary $w' \in R_a^M(w)$ and suppose that $M, w' \models \psi'$. Now it

follows that $(w', o') \in R_a^{M*U}((w, o))$. From our initial assumption, we can infer that $M*U, (w, o) \models \Box_a \varphi$. Therefore $M*U, (w', o') \models \varphi$. Hence, $M, w \models \bigwedge_{(\psi, o', \psi') \in \mathfrak{a}_a^U(o)} (\psi \rightarrow \Box_a(\psi' \rightarrow [U, o']\varphi))$.

From right to left. Take an arbitrary pointed Kripke model (M, w) and suppose that $M, w \models \psi \rightarrow \Box_a(\psi' \rightarrow [U, o']\varphi)$ for each $(\psi, o', \psi') \in \mathfrak{a}_a^U(o)$. Consider the model $M*U$ and take an arbitrary $(w', o') \in R_a^{M*U}((w, o))$. There must be an a -arrow $(\psi, o', \psi') \in \mathfrak{a}_a^U(o)$ such that $M, w \models \psi$ and $M, w' \models \psi'$. So it must be that $M, w \models \Box_a(\psi' \rightarrow [U, o']\varphi)$, and therefore $M, w' \models [U, o']\varphi$. Therefore $M*U, (w', o') \models \varphi$. Hence, $M, w \models [U, o]\Box_a \varphi$. \square

The proof that GAUL is complete with respect to the semantics of the previous section is typical for Dynamic Epistemic Logic [5, 7, 10, 13]. First, we prove the *Reduction Theorem*: each formula $\varphi \in \mathcal{L}$ can be translated to a provably equivalent, arrow-update-free “reduced” formula $\varphi^\circ \in \mathcal{L}^0$. This is proved by induction on i with $\varphi \in \mathcal{L}^i$ using modal reasoning and the so-called *reduction axioms* U1–U4. We then use this result, soundness, and completeness of the underlying modal logic (in this case, multi-modal **K**) to prove completeness for GAUL. Since these arguments are standard in Dynamic Epistemic Logic [5, 7, 10, 13], we omit proofs here.

THEOREM 2.5 (GAUL REDUCTION). *For each $\varphi \in \mathcal{L}$, there is a reduced $\varphi^\circ \in \mathcal{L}^0$ such that $\vdash \varphi \leftrightarrow \varphi^\circ$.*

THEOREM 2.6. *$\models \varphi$ implies $\vdash \varphi$ for each $\varphi \in \mathcal{L}$.*

3. CONNECTIONS WITH DYNAMIC EPISTEMIC LOGIC

Since GAUL is just as expressive as ordinary epistemic logic (Theorems 2.4 and 2.5), and the same has been established for DEL (see Theorem 3.5 below), it follows that GAUL and DEL are also equally expressive. But while this notion of expressivity concerns the *formulas* of these logics, we shall focus in this section on the expressivity of the *updates* of these logics. By establishing this connection between the updates in these logics, one can be confident that both approaches will remain equally expressive even when we extend these logics by adding other operators such as common knowledge. We will first present DEL and then prove that DEL and GAUL can express the same updates. We shall then study the relative succinctness with which these two frameworks express updates.

3.1 DEL Syntax and Semantics

This subsection contains a brief review of DEL. The central notion in DEL is that of an *action model*. The idea, originally put forward by Baltag, Moss, and Solecki [1, 2], is that one can model epistemic actions by an action model in the same way as one can model epistemic situations by a Kripke model. An action model is like a finite Kripke model except that the worlds are called *events* and the propositional valuation is replaced by a *precondition function* that assigns a formula, called a *precondition*, to each event.

DEFINITION 3.1 (ACTION MODEL). *Let L be a language. An action model A for L is a tuple (E^A, R^A, pre^A) consisting of a finite nonempty set E^A of events, an epistemic possibility function $R^A : \mathcal{A} \times E^A \rightarrow \wp(E^A)$ (notation: $R_a^A(e) \stackrel{\text{def}}{=} \dots$*

$R^A(a, e)$), and a precondition function $\text{pre}^A : E^A \rightarrow L$ assigning a precondition $\text{pre}^A(e) \in L$ to each event $e \in E^A$. An action model A for L and an event $e \in E^A$ form a pointed action model (A, e) for L with point e . Let $\mathfrak{A}(L)$ be the set of action models for L and $\mathfrak{A}_*(L)$ be the set of pointed action models for L .

To obtain the language of DEL, we extend the language of multi-modal logic by adding pointed action models as modal operators.

DEFINITION 3.2 (LANGUAGE OF DEL). *Set $\mathcal{L}_{\text{DEL}}^0 \stackrel{\text{def}}{=} \mathcal{L}^0$. By induction on i , we define the language $\mathcal{L}_{\text{DEL}}^{i+1}$ to be the set of formulas φ formed by the grammar*

$$\varphi ::= \psi \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a \varphi \mid [A, e]\varphi,$$

where $\psi \in \mathcal{L}_{\text{DEL}}^i$, $a \in \mathcal{A}$, and $(A, e) \in \mathfrak{A}_*(\mathcal{L}_{\text{DEL}}^i)$. Finally, define the set $\mathcal{L}_{\text{DEL}} \stackrel{\text{def}}{=} \bigcup_{i \in \mathbb{N}} \mathcal{L}_{\text{DEL}}^i$ of DEL-formulas, the set $\mathfrak{A} \stackrel{\text{def}}{=} \mathfrak{A}(\mathcal{L}_{\text{DEL}})$ of action models, and the set $\mathfrak{A}_* \stackrel{\text{def}}{=} \mathfrak{A}_*(\mathcal{L}_{\text{DEL}})$ of pointed action models. To say that $\varphi \in \mathcal{L}_{\text{DEL}}$ is reduced means that $\varphi \in \mathcal{L}_{\text{DEL}}^0 = \mathcal{L}^0$.

DEFINITION 3.3 (SEMANTICS OF DEL). *We extend the binary truth relation \models from Definition 2.3 by adding the following inductive clause.*

- $M, w \models [A, e]\varphi$ means we have that $M, w \not\models \text{pre}^A(e)$ or $M[U], (w, e) \models \varphi$, where $M[U]$ is defined as follows.

$$W^{M[U]} \stackrel{\text{def}}{=} \{(v, f) \in W^M \times E^A \mid M, v \models \text{pre}^A(f)\}$$

$$R_a^{M[U]}((v, f)) \stackrel{\text{def}}{=} \{(v', f') \in W^{M[U]} \mid$$

$$v' \in R_a^M(v) \ \& \ f' \in R_a^A(f)\}$$

$$V^{M[U]}(p) \stackrel{\text{def}}{=} \{(v, f) \in W^{M[U]} \mid M, v \models p\}$$

Verifying that $M[U]$ is a Kripke model whenever $M, w \models \text{pre}^A(e)$ is straightforward.

The core semantic definition (due to Baltag, Moss, and Solecki [1, 2]) is the so-called “product update” operation $M \mapsto M[A]$ on Kripke models. We note that if $M, w \not\models \text{pre}^A(e)$, then a formula of the form $[A, e]\varphi$ is vacuously true at (M, w) . Said informally, if (M, w) does not satisfy the precondition of e , then we need not “execute” the product update in order to determine the truth of the formula $[A, e]\varphi$.

DEFINITION 3.4. *Let (M, w) be a pointed Kripke model. To say that $(A, e) \in \mathfrak{A}_*$ is executable at (M, w) means $M, w \models \text{pre}^A(e)$. Given $E \subseteq E^A$, we define the set*

$$(M, w)^{A, E} \stackrel{\text{def}}{=} \{e \in E \mid M, w \models \text{pre}^A(e)\}$$

of events in E executable at (M, w) .

Since a pointed action model (A, e) is not necessary executable at every pointed Kripke model (M, w) , we think of the operation $(M, w) \mapsto (M[A], (w, e))$ as a *partial function* operating on pointed Kripke models. A pointed Kripke model (M, w) is in the domain of the partial function induced by pointed action model (A, e) if and only if (A, e) is executable at (M, w) .

The following is proved by the standard *reduction axiom method* of Dynamic Epistemic Logic [3, 5, 7, 10, 13].

THEOREM 3.5 (DEL REDUCTION). *For each $\varphi \in \mathcal{L}_{\text{DEL}}$, there is a reduced $\varphi^\bullet \in \mathcal{L}^0$ such that $\models \varphi \leftrightarrow \varphi^\bullet$.*

3.2 Update Expressivity

In this subsection, we prove that the class of Kripke model-changing updates expressible using action models is the same as those expressible using arrow updates. To define the expressivity of updates, we present two adaptations of a notion of action model equivalence due to van Eijck, Ruan, and Sadzik [14].

DEFINITION 3.6. *Let $A \in \mathfrak{A}$ and $U \in \mathfrak{U}$. To say that A and U are partial-update equivalent means that there exists a pair $(e, o) \in E^A \times O^U$, called a partial-witness, satisfying the property that for each pointed Kripke model (M, w) at which (A, e) is executable, $(M[A], (w, e))$ and $(M * U, (w, o))$ are bisimilar. To say that A and U are update equivalent means that there is an $(E, o) \in \wp(E^A) \times O^U$, called a witness, satisfying the property that for each pointed Kripke model (M, w) , we have that $(M, w)^{A, E} \neq \emptyset$ and that $(M[A], (w, e))$ and $(M * U, (w, o))$ are bisimilar for each $e \in (M, w)^{A, E}$.*

Partial-update equivalence says that the arrow update must agree up to bisimulation with the action model only on those pointed Kripke models on which the event making up the partial-witness is executable. Intuitively, we look for partial-update equivalence when we start with an action model and are asked whether a given arrow update induces a model-change operation yielding bisimilar output. Since the events in an action model induce a partial (and not necessarily total) functional operation on pointed Kripke models, we only require agreement up to bisimulation in those cases where execution is possible. Update equivalence, on the other hand, requires agreement up to bisimulation on all pointed Kripke models. Intuitively, we look for update equivalence when we start with an arrow update and are asked whether a given action model induces a model-change operation yielding bisimilar output. Since an arrow update is always “executable,” agreement up to bisimulation on all pointed Kripke models is required for (full) update equivalence to hold. We formalize these intuitive descriptions by way of the validities of the following theorem.

THEOREM 3.7. *Let $A \in \mathfrak{A}$, $U \in \mathfrak{U}$, and $\varphi \in \mathcal{L} \cup \mathcal{L}_{\text{DEL}}$.*

1. *If A and U are partial-update equivalent with partial-witness (e, o) , then $\models [A, e]\varphi \leftrightarrow (\text{pre}^A(e) \rightarrow [U, o]\varphi)$.*
2. *If A and U are update equivalent with witness (E, o) , then $\models [U, o]\varphi \leftrightarrow \bigwedge_{e \in E} [A, e]\varphi$.*

PROOF. For $\varphi \in \mathcal{L}^0$, use Definition 3.6; for $\varphi \in \mathcal{L} - \mathcal{L}^0$, use Theorems 2.4 and 2.5; for $\varphi \in \mathcal{L}_{\text{DEL}} - \mathcal{L}^0$, use Theorem 3.5. \square

We now prove that every model-change operation describable by a DEL action model is also describable by a GAUL arrow update.

THEOREM 3.8 (DEL TO GAUL). *Given a pointed action model (A, e) , the pair (e, e) is a partial-witness to the partial-update equivalence between A and the arrow update $U(A)$ defined by*

$$\begin{aligned} O^{U(A)} &\stackrel{\text{def}}{=} E^A \\ \mathfrak{a}^{U(A)} &\stackrel{\text{def}}{=} \{((a, f), (\top, f', \text{pre}^A(f')^\bullet) \mid \\ &\quad (a, f, f') \in \mathcal{A} \times E^A \times R_a^A(f)\} . \end{aligned}$$

PROOF. Let (M, w) be a pointed Kripke model at which (A, e) is executable. We prove that the identity $\iota : W^{M[A]} \rightarrow W^{M * U(A)}$ is a bisimulation between $(M[A], (w, e))$ and $(M * U(A), (w, e))$. Note that ι is well-defined because $(v, f) \in W^{M[A]}$ implies

$$(v, f) \in W^M \times E^A = W^M \times O^{U(A)} = W^{M * U(A)} .$$

Further, we have $v \in V^{M[U]}(p)$ iff $v \in V^M(p)$ iff $v \in V^{M * U(A)}(p)$, and $\iota((w, e)) = (w, e)$ by definition.

Assume $(v, f) \in \text{domain}(\iota)$ and $(v', f') \in R_a^{M[A]}((v, f))$. By DEL Reduction (Theorem 3.5), it follows that $M, v' \models \text{pre}^A(f')^\bullet$, that $v' \in R_a^M(v)$, and that $f' \in R_a^A(f)$. By the definition of $U(A)$, the conjunction of the previous sentence implies $(v', f') \in R_a^M(v) \times O^{U(A)}$ and $(\top, f', \text{pre}^A(f')^\bullet) \in \mathfrak{a}_a^{U(A)}(f)$ with $M, v' \models \top$ and $M, v' \models \text{pre}^A(f')^\bullet$. But the latter implies $(v', f') \in R_a^{M * U(A)}((v, f))$.

Assume $(v, f) \in \text{image}(\iota)$ and $(v', f') \in R_a^{M * U(A)}((v, f))$. It follows that $M, v \models \text{pre}^A(f)$, that $(v', f') \in R_a^M(v) \times O^{U(A)}$, and that $(\top, f', \text{pre}^A(f')^\bullet) \in \mathfrak{a}_a^{U(A)}(f)$ with $M, v' \models \text{pre}^A(f')^\bullet$. By the definition of $U(A)$ and DEL Reduction (Theorem 3.5), we have that $M, v \models \text{pre}^A(f)$, that $M, v' \models \text{pre}^A(f')$, that $v' \in R_a^M(v)$, and that $f' \in R_a^A(f)$. But the latter is what it means to have $(v', f') \in R_a^{M[A]}((v, f))$. \square

Theorems 3.7(1) and 3.8 yield the following corollary.

COROLLARY 3.9. $\models [A, e]\varphi \leftrightarrow (\text{pre}^A(e) \rightarrow [U(A), e]\varphi)$ for each $(A, e) \in \mathfrak{A}_*$ and $\varphi \in \mathcal{L} \cup \mathcal{L}_{\text{DEL}}$.

We now prove that every model-change operation describable by a GAUL arrow update is also describable by a DEL action model. We begin with a preliminary definition.

DEFINITION 3.10. *Let U be an arrow update and $\Phi(U)$ be the set of GAUL-formulas that are a source or target condition in U . Set $\Phi^\pm(U) \stackrel{\text{def}}{=} \Phi(U) \cup \{\neg\varphi \mid \varphi \in \Phi(U)\}$. To say that a set Γ of GAUL-formulas is U -maxcons means that $\Gamma \subseteq \Phi^\pm(U)$, Γ is GAUL-consistent (i.e., $\Gamma \not\vdash \perp$), and no $\Gamma' \subseteq \Phi^\pm(U)$ strictly containing Γ is GAUL-consistent. Let $\text{mc}(U)$ denote the collection of U -maxcons sets of GAUL-formulas. Define action model $A[U]$ as follows.*

$$\begin{aligned} E^{A[U]} &\stackrel{\text{def}}{=} \text{mc}(U) \times O^U \\ R_a^{A[U]}((\Gamma, o)) &\stackrel{\text{def}}{=} \{(\Gamma', o') \in E^{A[U]} \mid \exists (\varphi, o', \varphi') \in \mathfrak{a}_a^U(o) : \\ &\quad (\varphi \in \Gamma \ \& \ \varphi' \in \Gamma')\} \\ \text{pre}^{A[U]}((\Gamma, o)) &\stackrel{\text{def}}{=} (\bigwedge \Gamma)^\circ \end{aligned}$$

THEOREM 3.11 (GAUL TO DEL). *Let M be a Kripke model and U be an arrow update.*

1. *For each $w \in W^M$, there is a unique $\Gamma_w \in \text{mc}(U)$ satisfying $M, w \models (\bigwedge \Gamma_w)^\circ$.*
2. *The function $f_M^U : W^{M * U} \rightarrow W^{M[A[U]}}$ defined by $f_M^U((w, o)) \stackrel{\text{def}}{=} (w, (\Gamma_w, o))$ is a total isomorphism between the Kripke models $M * U$ and $M[A[U]]$.*
3. *For each $o \in O^U$, we have that U and $A[U]$ are update equivalent with witness $(\text{mc}(U) \times \{o\}, o)$.*

PROOF. (1) and (2) follow by adapting proofs of these results for the case of single-outcome arrow updates appearing in a previous paper by the authors [8]. (3) follows by (1) and (2). \square

COROLLARY 3.12. *Let $(U, o) \in \mathcal{U}_*$ and $\varphi \in \mathcal{L} \cup \mathcal{L}_{\text{DEL}}$. For each pointed Kripke model (M, w) , we have $M, w \models [U, o]\varphi \leftrightarrow [A[U], (\Gamma_w, o)]\varphi$. Further, we have*

$$\models [U, o]\varphi \leftrightarrow \bigwedge_{\Gamma \in \text{mc}(U)} [A[U], (\Gamma, o)]\varphi .$$

PROOF. Given a pointed Kripke model (M, w) , we have $M, w \models (\bigwedge \Gamma_w)^\circ$ by Theorem 3.11(1). By Theorem 3.11(3), it follows that $(M * U, (w, o))$ and $(M[A[U]], (w, (\Gamma_w, o)))$ are bisimilar. By Reduction (Theorems 2.4, 2.5, and 3.5), it follows that $M, w \models [U, o]\varphi \leftrightarrow [A[U], (\Gamma_w, o)]\varphi$. Finally, to see that $\models [U, o]\varphi \leftrightarrow \bigwedge_{\Gamma \in \text{mc}(U)} [A[U], (\Gamma, o)]\varphi$, apply Theorems 3.11(3) and 3.7(2). \square

3.3 Update Succinctness

We showed in the previous subsection that DEL and GAUL describe the same class of model-changing operations. However, it is not necessarily the case that they express these operations with comparable succinctness. In fact, as we prove below, arrow updates are at worst poly-exponentially less succinct than action models, though this improves to being at worst polynomially less succinct if the action models have purely epistemic preconditions (i.e., preconditions in \mathcal{L}^0) or if we allow arrow updates to have target conditions in \mathcal{L}_{DEL} . We also show that arrow updates are sometimes exponentially more succinct than action models. To prove these results, we begin by defining the following notions of length and size.

DEFINITION 3.13 (LENGTH & SIZE). *The length of $\varphi \in \mathcal{L}^0$, written $\text{len}(\varphi)$, is the number of symbols occurring in φ . The size of arrow update $U \in \mathcal{U}(\mathcal{L}^0)$, written $s(U)$, is defined by $s(U) \stackrel{\text{def}}{=} \sum_{((a,o),(\varphi,o'),\varphi') \in \mathfrak{a}^U} (\text{len}(\varphi) + \text{len}(\varphi'))$. Define $n(U) \stackrel{\text{def}}{=} |O^U|$ and $m(U) \stackrel{\text{def}}{=} |\mathfrak{a}^U|$. The size of action model $A \in \mathfrak{A}(\mathcal{L}^0)$, also written $s(A)$, is defined by $s(A) \stackrel{\text{def}}{=} |R^A| + \sum_{e \in E^A} \text{len}(\text{pre}^A(e))$. Define $n(A) \stackrel{\text{def}}{=} |E^A|$ and $m(A) \stackrel{\text{def}}{=} |R^A|$.*

It is possible to carry out a recursion that defines sizes for all arrow updates and action models and lengths for all GAUL- and DEL-formulas. The general scheme: stage $i + 1$ size is defined in terms of stage i length (as in Definition 3.13), and the stage $i + 1$ length of φ is equal to the number of symbols occurring in φ , except that pointed arrow update modals $[U, o]$ and pointed action model modals $[A, e]$ are counted a number of times equal to the stage $i + 1$ size of U or A , respectively. It then suffices to show that the length or size of something at one stage is equal to its length or size at all later stages.

We observe that every formula is of length at least 1. It follows that $s(U) > n(U)$ and $s(U) > m(U)$ for each arrow update U ; that is, the size $s(U)$ of U is strictly greater than both the number $n(U)$ of outcomes in U and the number $m(U)$ of arrows in U . Similarly, $s(A) \geq n(A)$ and $s(A) > m(A)$ for each action model A ; that is, the size $s(A)$ of A is no less than the number $n(A)$ of events in A and is strictly greater than the number $m(A)$ of arrows in A .

Our concern here is to address the question of *update succinctness*. In particular, we report that arrow updates are at worst poly-exponentially less succinct than action models, though this improves to being at worst polynomially less succinct if the action models have purely epistemic preconditions (i.e., preconditions all in \mathcal{L}^0) or if we allow arrow

updates to have target conditions in \mathcal{L}_{DEL} . These results, presented below in Theorem 3.14, are new. We also report that arrow updates are sometimes exponentially more succinct than action models. This result, presented below in Theorem 3.15, is a straightforward adaptation of an update succinctness result for single-outcome arrow updates [8].

THEOREM 3.14 (WORST-CASE UPDATE SUCCINCTNESS). *Given action model A , define $K \stackrel{\text{def}}{=} \max_{e \in E^A} \text{len}(\text{pre}^A(e))$.*

1. *Arrow update $U(A)$, which is partial-update equivalent to A (Theorem 3.8), has size $O(|\mathcal{A}| \cdot n(A)^2 \cdot 2^K)$.*

2. *If $A \in \mathcal{U}(\mathcal{L}^0)$, then arrow update $U(A)$ has size $O(|\mathcal{A}| \cdot n(A)^2 \cdot K)$.*

3. *If we allow \mathcal{L}_{DEL} target conditions in arrow updates and modify the definition of $U(A)$ (in Theorem 3.8) by omitting the \bullet operation, then arrow update $U(A)$ has size $O(|\mathcal{A}| \cdot n(A)^2 \cdot K)$.*

PROOF. Let $K' \stackrel{\text{def}}{=} \max_{e \in E^A} \text{len}(\text{pre}^A(e)^\bullet)$. By inspection of the definition of $U(A)$ from Theorem 3.8, we have $s(U(A)) \leq |\mathcal{A}| \cdot n(A)^2 \cdot (1 + K')$. If $A \in \mathcal{L}^0$, then $\text{pre}^A(e)^\bullet = \text{pre}^A(e)$ for each $e \in E^A$ and hence $K' = K$. In general, we have by inspection of the standard DEL reduction axioms [13] that $\text{len}(\text{pre}^A(e)^\bullet) = O(2^{\text{len}(\text{pre}^A(e))})$. It also follows by a result of Lutz [9] that we might have $\text{len}(\text{pre}^A(e)^\bullet) \geq 2^{\text{len}(\text{pre}^A(e))}$ for some $e \in E^A$. We therefore have $K' = 2^K$ in the general case. Finally, we observe that the argument for (3) is like the argument for (1). \square

Letting $x \stackrel{\text{def}}{=} \max\{|\mathcal{A}|, n(A), K\}$, the first upper-bound in Theorem 3.14 is bounded above by the poly-exponential $x^3 \cdot 2^x$ and second upper-bound in Theorem 3.14 is bounded above by the polynomial x^4 .

The following theorem uses results from a previous paper by the authors [8].

THEOREM 3.15 (EXPONENTIAL UPDATE SUCCINCTNESS). *Fix an agent $a \in \mathcal{A}$ and assume $\mathcal{P} = \{p_i \mid i \in \mathbb{N}\}$. For each $k \in \mathbb{N}$, define arrow update*

$$U_k \stackrel{\text{def}}{=} (\{o\}, \{((a, o), (\neg p_i, o, p_i)) \mid i \leq k\})$$

and action model $A_k \stackrel{\text{def}}{=} A[U_k]$.

1. $s(U_k) = 3k + 3$ and $n(A_k) = 2^{k+1}$.

2. *If an action model A is update equivalent to U_k , then $s(A) \geq 2^{k+1}$.*

PROOF. (1) was proved in a previous paper by the authors [8]. Since this argument is short and sweet, we repeat it here. Proceeding, for each $i \leq k$, the pair $((a, o), (\neg p_i, o, p_i)) \in \mathfrak{a}^U$ contributes $\text{len}(\neg p_i) + \text{len}(p_i) = 3$ to the size. Since $k + 1$ pairwise distinct pairs of this form occur in \mathfrak{a}^U , it follows that $s(U_k) = 3k + 3$. As for the size of A_k , we have $\Phi^\pm(U_k) = \{p_i, \neg p_i, \neg \neg p_i \mid i \leq k\}$. Hence for each $i \leq k$, a U_k -maxcons set has exactly one of the two sets $\{p_i, \neg \neg p_i\}$ and $\{\neg p_i\}$ as a subset. Since a U_k -maxcons set is a union of $k + 1$ pairwise disjoint sets, each of which has one of these two distinct forms, it follows that $n(A_k) = 2^{k+1}$.

We now prove (2). Let C_k be the k -dimensional epistemic hypercube; that is, the worlds of C_k are the subsets of $\bar{k} \stackrel{\text{def}}{=} \{1, \dots, k\}$.

$\{p_i \mid i \leq k\}$, we have $R_i^{C_k}(w) \stackrel{\text{def}}{=} \wp(\bar{k})$ for each $i \in \mathcal{A}$ and $w \in \wp(\bar{k})$, and $w \in V^{C_k}(p)$ iff $p \in w$. In a previous paper by the authors [8], it is shown that if A is an action model with $n(A) < n(A_k)$ and $E \subseteq E^A$ is a set of events in A , then there is a reduced formula $\varphi \in \mathcal{L}^0$ and a world w in C_k such that exactly one of $\bigwedge_{e \in E} [A, e]\varphi$ and $[U_k, o]\varphi$ is true at (C_k, w) . Applying Theorem 3.7(2), it follows that A and U_k are not update equivalent. Therefore, since $s(A) \geq n(A)$ for every action model A , no action model A with $s(A) < 2^{k+1}$ is update equivalent to U_k . \square

4. EXAMPLES

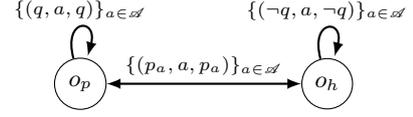
We showed in the previous section that DEL action models and GAUL arrow updates express the same model-change operations. Moreover, our succinctness results also indicate that it is sometimes more convenient to use arrow updates than it is to use action models. In this section, we provide further evidence of the utility of arrow updates by considering concrete model-change scenarios that we believe are more conveniently represented using arrow updates. To describe these scenarios, we begin with some terminology. A *semi-private announcement* of φ is a message that informs a group $G \subseteq \mathcal{A}$ of the truth value of a formula φ while telling the other agents in $\mathcal{A} - G$ that group G learned the truth about φ without saying what it is that was learned. A *private announcement* of φ is a message that informs a group $G \subseteq \mathcal{A}$ that φ is true without providing any new information to agents in $\mathcal{A} - G$, not even the information that the private announcement took place.

In this section, we will investigate semi-private and private announcements in which the membership of the recipient group G is not necessarily known. Situations in which group membership is not necessarily known are said to have *nonrigid* group membership [4]. A well-known example of nonrigid group membership is the *Simultaneous Byzantine Agreement Problem* [4, §6.3]. In this problem, a group $G \subseteq \mathcal{A}$ of loyal agents must reach common knowledge among themselves despite the malicious efforts of the group $\mathcal{A} - G$ of traitors whose identities are unknown.

4.1 Nonrigid Semi-private Announcements

Professor Fickle divides his colleagues into two disjoint groups: friends and foes. The professor changes his mind all the time about which colleague goes in which group, so a colleague generally does not know the group to which she belongs. The professor celebrates his birthday each year by having a party in his favorite pub or at his house, but he finds it very hard to decide where to throw the party. On the day before his birthday, he sends his friends the message where the party is going to be. If a colleague is a friend, then the colleague either gets the message “the party is in the pub” or the message “the party is at my house”. Naturally the professor does not send messages to his foes. We suppose that Professor Fickle has used this invitation procedure for as long as anyone can remember, and that the procedure is therefore assumed to be common knowledge among his colleagues.

Let \mathcal{A} be the set of colleagues. Fix a propositional letter q whose meaning is “the party is in the pub”. For each agent $a \in \mathcal{A}$, let p_a be a (different) propositional letter whose meaning is “agent a is a foe”. We define the arrow update U_I using the following diagram.



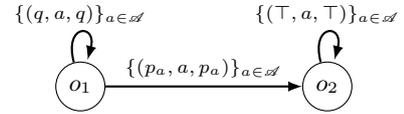
In this diagram, nodes are labeled by outcomes and each arrow from an outcome o to an outcome o' is labeled by a set of triples (φ, a, φ') indicating that $(\varphi, o', \varphi') \in \mathfrak{a}_a^{U_I}(o)$. Arrow update U_I can be used to reason about the epistemic affects of Professor Fickle’s invitation: o_p is the outcome in which the message “the party is in the pub” is sent, and o_h is the outcome in which the message “the party is at my house” is sent. Let M be the Kripke model defined by $W^M \stackrel{\text{def}}{=} \wp(\{p_a \mid a \in \mathcal{A}\} \cup \{q\})$, $R_a^M(w) \stackrel{\text{def}}{=} W^M$ for each $(w, a) \in W^M \times \mathcal{A}$, and $V^M(p) \stackrel{\text{def}}{=} \{w \in W^M \mid p \in w\}$. At each world in the updated model $M[U_I]$, we see that friends know the location of the party, foes do not, and no one knows whether someone else is a friend.

We think that the arrow update U_I is a much more convenient representation of this scenario than is an update equivalent action model representation. As an indication of this convenience, we note that the Theorem 3.11-equivalent action model $A[U_I]$ is large in size when compared to the size of U_I . This suggest to us that arrow updates are a more convenient means of reasoning about Professor Fickle’s invitation, and, by way of extension, about semi-private announcements to nonrigid groups.²

4.2 Nonrigid Private Announcements

Suppose there is a secret society that communicates by way of an anonymous online message board. Membership of this secret society is itself a secret, meaning that an agent knows whether she herself is a member but does not know the identities of the other members. Suppose now that the message q is posted on the message board. For simplicity, we assume that all members instantly read and understand any message posted on the board. q is therefore known by all members of the society. But since no member knows the identity of another member, no member knows of any other person that knows q .

Let q be a propositional letter. For each agent $a \in \mathcal{A}$, we let the (different) propositional letter p_a mean “agent a is not a member of the society”. Using the arrow update drawing conventions from the previous subsection, we define the arrow update U_q according to the following diagram.



²A natural update equivalent action model A suggested by a referee has

$$\begin{aligned} E^A &\stackrel{\text{def}}{=} \{(q, G), (-q, G) \mid G \subseteq \mathcal{A}\} \\ R_a^A((x, G)) &\stackrel{\text{def}}{=} \{(y, H) \mid a \in G \cap H \ \& \ x = y\} \cup \{(y, H) \mid a \notin G \cup H\} \\ \text{pre}^A((x, G)) &\stackrel{\text{def}}{=} x \wedge (\bigwedge_{a \in G} \neg p_a) \wedge (\bigwedge_{a \in \mathcal{A} - G} p_a) . \end{aligned}$$

However, we note that $s(A)$ is exponential in the number of agents, whereas $s(U_I)$ is linear in the number of agents. For this reason, we believe that arrow update U_I is a more convenient representation of the Professor Fickle scenario than is action model A .

Arrow update U_q can be used to reason about the epistemic changes brought about as a result of q being posted on the anonymous message board: o_1 is used as the actual outcome, while o_2 acts as a hypothetical outcome that the agents mistakenly entertain as possible. Let the Kripke model M be defined as in the previous subsection except that, for each agent $a \in \mathcal{A}$, the set $R_a^M(w)$ consists of the worlds $w' \in W^M$ such that $p_a \in w$ if and only if $p_a \in w'$. This ensures that an agent always knows whether she is a member. At each world in the updated model $M[U_q]$ having the form (w, o_1) , we see that members know q (and members know that someone else knows q if she is a member), non-members do not know q (nor do they consider it possible that anyone knows q), and no one knows whether someone else is a member. For similar reasons as in the previous subsection, we believe that the arrow update representation of this and other private announcements to nonrigid groups is preferable to an update equivalent action model representation.

5. CONCLUSION

We presented GAUL, our theory of generalized (i.e., multi-outcome) arrow updates. We showed that arrow updates and action models can express the same model-changing operations. But mutual expressivity does not come for free. In particular, arrow updates are at worst poly-exponentially less succinct than action models, though this improves to being at worst polynomially less succinct if the action models have purely epistemic preconditions (i.e., preconditions all in \mathcal{L}^0) or if we allow arrow updates to have target conditions in \mathcal{L}_{DEL} . Further, arrow updates are sometimes exponentially more succinct than action models. These results suggest that using arrow updates is more convenient than using action models in certain cases. We also presented two concrete examples, semi-private and private announcements to nonrigid groups, in which we believe arrow updates are preferable to action models.

One direction for further study is to adapt the authors' arrow update-based notion of common knowledge for single-outcome arrow updates [8] to the present generalized setting. This work raises a number of interesting relative expressivity questions (comparing DEL with various common knowledge notions to GAUL with these common knowledge notions) that have yet to be studied beyond our preliminary investigation in [8]. Moreover, while we showed in this paper that permitting multiple outcomes allows arrow updates to express all action model-based updates, it remains an open question as to how this result is affected if we place a fixed upper-bound on the maximum number of allowable outcomes. While we suspect that this causes the result to fail, it has yet to be proved. In particular, we do not yet understand the exact update-expressive relationship between action models and single-outcome arrow updates.

References

- [1] Alexandru Baltag and Lawrence S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.
- [2] Alexandru Baltag, Lawrence S. Moss, and Sławomir Solecki. The logic of public announcements, common knowledge, and private suspicions. In Itzhak Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK VII)*, pages 43–56, Evanston, Illinois, USA, July 22–24, 1998.
- [3] Alexandru Baltag, Lawrence S. Moss, and Sławomir Solecki. Logics for epistemic actions: completeness, decidability, expressivity. Manuscript, 2005.
- [4] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. The MIT Press, 1995.
- [5] Jelle Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999. ILLC Dissertation Series DS-1999-01.
- [6] Jelle Gerbrandy and Willem Groeneveld. Reasoning about information change. *Journal of Logic, Language and Information*, 6(2):147–169, 1997.
- [7] Barteld Kooi. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-classical Logics*, 17(2):231–253, 2007.
- [8] Barteld Kooi and Bryan Renne. Arrow update logic. Forthcoming in *Review of Symbolic Logic*; accepted May 2011.
- [9] Carsten Lutz. Complexity and succinctness of public announcement logic. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*, pages 137–144, Hakodate, Japan, May 8–12, 2006. Association for Computing Machinery (ACM), New York, New York, USA.
- [10] Jan Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems (ISMIS 1989): Poster Session Program*, pages 201–216, Charlotte, North Carolina, USA, 1989. Oak Ridge National Laboratory ORNL/DSRD-24.
- [11] Jan Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007.
- [12] Bryan Renne. Public and private communication are different: Results on relative expressivity. *Synthese*, 165(2):225–245, 2008.
- [13] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer Netherlands, 2007.
- [14] Jan van Eijck, Ji Ruan, and Tomasz Sadzik. Action emulation. Manuscript, March 2008.