Towards a Concrete Semantics for Announcements

Riccardo Pucella and Mehrnoosh Sadrzadeh
riccardo@ccs.neu.edu, mehrs@comlab.ox.ac.uk
College of Computer and Information Science, Northeastern University,
Laboratoire Preuves Programmes et Systèmes, Université Paris 7*

The story of this paper is inspired by the existence of two schools of thought in modeling and reasoning about information flow in interactive multi-agent systems where agents aquire new information after communicating with each other. One of these schools is the *interpreted systems* of, e.g. [6, 8] and the other is the *dynamic epistemic logic* (DEL) approach of, e.g. [1, 2, 3, 5, 7]. The former is more algorithmic and inclined towards applications to distibuted computing and artificial intelligence, whereas the latter is more logical and focuses on formalizing concepts such as rational agents. One key difference between the two schools is the semantic models used to represent a situtation of interest, about which we may want to reason. Fagin et al use a very concrete semantics based on sequences of local states, representing the possible executions of the system under consideration, which can be obtained from an algorithmic description of the system in question, while DEL uses a more abstract semantics in terms of Kripke structures.

The aim of this paper, which is very much work in progress, is to bring together these two schools by developing an *interpreted systems* semantics for DEL. A full development would involve devising a model for public and private announcements, which can be both honest and dishonest. For the time being, however, we only present a system for logics of honest public and private announcements, leaving the dishonest announcements for future work. Our development consists of (1) a set of axioms describing runs of a system modeling public and private announcements, (2) an LTL-style logic enriched with epistemic modalities, and (3) a translation from dynamic epistemic logic to our logic showing that we can interpret the existing logics of announcements in our setting. It is not hard to see that our logic is sound and complete with respect to the particular class of interpreted systems we consider, and that our translation is sound, that is, its image indeed forms a DEL.

We believe that our semantics, because of its more concrete nature, provides new insights into the nature of dynamic epistemic logic, and will allow us to study announcements in real systems, as well as more easily accomodate extensions (such as fact-changing actions and parallel announcements), explore design space for informative belief in the context of cheating and lying, and calculate the complexity of logics of public and private announcements. Our semantics in general, and the calculation of its complexity in particular, should go along recent interesting results of [4]. Studying these connections is also left to future work. One nice feature of our approach is that we use standard primitives, such as time and knowledge, with an S5 interpretation for knowledge that comes from a completely natural

*Visiting Computing Laboratory, University of Oxford accessibility relation for every agent, namely indistinguishability of the agent's local state. Using these relations, we define a new modality, a form of belief, that seems to be required to capture the state of uncertainty of agents as a result of private announcements. This new modality helps explain the effect of private announcements on the knowledge of agents.

Finally, note that the relationship between our logic and DEL is similar to the relationship between temporal logic and dynamic logic: endogenous versus exogenous logics; difference in complexity; often interested in reasoning about a *particular* protocol, rather than all protocols. They indeed follow same ideas, but have different points of view.

1 Honest Public and Private Announcements

We start by defining a class of structures that correspond somewhat closely to execution of protocols or programs. Here, we use the runs and systems framework of Fagin et al, which roughly takes as models sets of traces of execution.

An interpreted system is a pair $(\mathcal{R}, \pi)$, where $\mathcal{R}$ is a set of runs, each run representing a possible execution of the system, and π is an interpretation for the primitive positions (or atoms). A run is a map from time to global states, where $r(m)$ is the global state in run r at time m. (For simplicity, we take time to range over the natural numbers.) We call a pair (r, m) a point of r. Thus, each point corresponds to a global state of the system.

A global state is of the form $(s_e, s_1, \ldots, s_n)$, where s_e is the state of the environment, and $s_1, \ldots, s_n$ are local states for each agent. If $r(m) = (s_e, s_1, \ldots, s_n)$, we write $r_e(m)$ for s_e and $r_i(m)$ for s_i. Intuitively, the local state for an agent records the observations that the agent has made. Here, we consider observations to be announcements, public or private, honest or dishonest, that the agent has received. We also record in the local state possible initial observations made by the agent. For instance, in the muddy children puzzle, initial observations for an agent include which other children are dirty. The local state of the environment records information which is not available to the agents. For example, in a coin tossing scenario, the result of the coin toss and also the sequence of all announcements that have been made can be stored in the environment state. The announcements that an agent has received will be a subsequence of this sequence of announcements.

We will want to interpret knowledge over such systems, and to do so, we define, for each agent i, a relation over points of the system capturing which points agent i cannot distinguish. Intuitively, two points will be indistinguishable to agent i if i has the same local state in both points.[1] We define $(r, m) \sim_i (r', m')$ if $r_i(m) = r'_i(m)$. Note that this makes $\sim_i$ an equivalence relation. In this section, we consider only honest public and private announcements, meaning that the sequence of announcements seen by each agent differs

[1] One subtlety here is that we may not want to distinguish announcements that are logically equivalent. For instance, the announcement $(p \wedge q)!_\beta$ can be thought of as the same announcement as $(q \wedge p)!_\beta$. In DEL, announcements are taken to be sets of states, therefore there is no distinction between announcing $p \wedge q$ and announcing $q \wedge p$. But it is easy to imagine scenarios where the fact that an announcement is "presented" differently can be relevant; for instance, an announcement can be represented as a bit-string sent over a network link, and an agent may be able to distinguish two different bit-strings corresponding to two different presentations of the same announcement. For this paper, we consider announcements to be the same when they are structurally equal, with the understanding that much of our development could be done with an arbitrary equivalence relation over announcements.

from each other and also from that of the environment: the insiders observe it and the outsiders don't. But in any case what is announced is truthful.

We need a syntax for announcements. We write announcements in an epistemic propositional language. (In fact, this is a sublanguage of the logic we introduce immediately after.) Start with a set Φ_0 of primitive propositions, representing the basic facts that we care about, and form the language $\mathcal{L}_n^K$ of announcements by closing Φ_0 under $\wedge$, $\neg$, and the modal operator K_i. (As usual, we define $\vee$ and $\Rightarrow$ as abbreviations.) We use θ to range over formulas in $\mathcal{L}_n^K$. An announcement is written $\theta!_\beta$ and for θ a formula in $\mathcal{L}_n^K$ and β a subset of agents $\beta \subseteq \{1, \ldots, n\}$, it is interpreted as the announcement of θ to all agents in group β. An announcement is *public* if it is of the form $\theta!_{\{1,\ldots,n\}}$, otherwise, it is *private*.

We define a simple logic $\mathcal{L}_n^{K,A}$ for reasoning about announcements in that setting. The logic has essentially only epistemic and temporal operators. Again, we start with the set Φ_0 of primitive propositions, and add propositions of the form $\mathsf{ann}(\theta, \beta)$ that reads as "an announcement of θ has just happened to agents in β". Formulas of $\mathcal{L}_n^{K,A}$ include $\varphi_1 \wedge \varphi_2$, $\neg\varphi$, $K_i\varphi$ (read "agent i knows φ"), and $\bigcirc\varphi$ (read "φ is true at the next round").

$$\varphi \;::=\; p \mid \mathsf{ann}(\theta, \beta) \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid K_i\varphi \mid \bigcirc\varphi$$

Note that $\mathcal{L}_n^K$ is a sublanguage of $\mathcal{L}_n^{K,A}$, so that θ is really a formula of $\mathcal{L}_n^{K,A}$.

We interpret $\mathcal{L}_n^{K,A}$ with respect to systems. We say formula φ is true at a point (r, m) of system $\mathcal{A}$, written $(\mathcal{A}, r, m) \models \varphi$, defined inductively as

$(\mathcal{A}, r, m) \models p$ iff $\pi(r, m)(p) = \mathbf{true}$

$(\mathcal{A}, r, m) \models \mathsf{ann}(\theta, \beta)$ iff $\mathit{last}(r_e(m + 1)) = \theta!_\beta$

$(\mathcal{A}, r, m) \models \varphi_1 \wedge \varphi_2$ iff $(\mathcal{A}, r, m) \models \varphi_1$ and $(\mathcal{A}, r, m) \models \varphi_2$

$(\mathcal{A}, r, m) \models \neg\varphi$ iff $(\mathcal{A}, r, m) \not\models \varphi$

$(\mathcal{A}, r, m) \models K_i\varphi$ iff for all $(r', m') \sim_i (r, m)$, $(\mathcal{A}, r', m') \models \varphi$

$(\mathcal{A}, r, m) \models \bigcirc\varphi$ iff $(\mathcal{A}, r, m + 1) \models \varphi$.

Consider the following possible constraints on the runs of a system.

A1. There is at most one announcement made at each round. Thus, a run r satisfies: for all $m \geq 0$, if $r(m) = (s_e, s_1, \ldots, s_n)$ and $r(m + 1) = (s'_e, s'_1, \ldots, s'_n)$, then s'_e is just s_e with either a new announcement $\theta!_\beta$ or $\surd$, a marker indicating no announcement has been made, and similarly for each s'_i.

A2. Facts do not change during a run. Thus, for all primitive propositions p and all times $m, m' \geq 0$, $\pi(r, m)(p) = \mathbf{true}$ if and only if $\pi(r, m')(p) = \mathbf{true}$.

A3. Every announcement $\theta!_\beta$ is put in the local state of the agents in β only, as well as in the local state of the environment. For all r, m, and i, if the announcements in $r_i(m)$ are $\langle a_1, \ldots, a_k\rangle$ and the announcements in $r_e(m)$ are $\langle b_1, \ldots, b_k\rangle$, then if b_j is $\theta!_\beta$, then a_j is $\theta!_\beta$ if $i \in \beta$, and $\surd$ if $i \notin \beta$.

A4. For all r, m, θ, if $(\mathcal{A}, r, m) \models \mathsf{ann}(\theta, \beta)$, then $(\mathcal{A}, r, m) \models \theta$.

Definition 1.1 A *private announcement system* is a system satisfying A1–3. A private announcement system is *honest* if it also satisfies A4. A private announcement system in which the announcements in the local state of all agents are the same as in the state of the environment is called a *public announcement system*. It is honest if the private announcement system is.

Example 1.2 The announcement system for the muddy children puzzle. *Suppose there are n children and k of them are dirty. Propositions (for the announcements) are $D_1, \ldots, D_n$ (interpreted as "agent i is dirty"). The system has 2^n runs, one per subset of $\{1, \ldots, n\}$ indicating which children are dirty. Local state of each agent records which other agent is dirty. Initially, there are no announcements. At the first step, if at least one child is dirty, the announcement (by the father) is*

$$(\bigvee_{i=1}^{n} D_i)!$$

This is followed by $k-1$ announcements (where k is the number of dirty children)

$$(\bigwedge_{i=1}^{n} \neg K_i D_i \wedge \neg K_i \neg D_i)!$$

Note that the run where there are no dirty children has no announcement. (We will see later why we want to do things that way—we want to make sure that we do not announce false things.)

We say φ is valid in $\mathcal{A}$ if $(\mathcal{A}, r, m) \models \varphi$ for all r and m. A formula φ is valid with respect to an honest private announcement systems if $\mathcal{A} \models \varphi$ for all honest private announcement systems $\mathcal{A}$.

Proposition 1.3 $\mathcal{L}_n^{K,A}$ is sound and complete with respect to honest announcement systems.

2 Translation to the Honest Fragment of DEL

We interpret DEL formulas in our system. The basic DEL formula is $[\alpha]\varphi$, which says "after announcing α (a complex announcement), φ is true". In DEL, complex announcements are sequences and nondeterministic choices of more primitive announcements: $\alpha_1; \alpha_2$, and $\alpha_1 \vee \alpha_2$. Primitive honest announcements are of the form $\theta!_\beta$. Intuitively, we interpret $[\theta!_\beta]\varphi$ as $\mathsf{ann}(\theta, \beta) \Rightarrow \bigcirc\varphi$. More complex announcements can be obtained similarly; for instance, $[\theta_1!_{\beta 1}; \theta_2!_{\beta 2}]\varphi$ corresponds to $\mathsf{ann}(\theta_1, \beta 1) \wedge \bigcirc \mathsf{ann}(\theta_2, \beta 2) \Rightarrow \bigcirc\bigcirc\varphi$, and so on.

There is a slight discrepancy in the interpretation, in the sense that $[\alpha]\varphi$ in DEL talks about the possibility of doing α, while in our setting $\mathsf{ann}(\theta, \beta)$ talks about θ being just announced. This is a consequence of our choice of using linear time models - using branching time models would allow us to express an interpretation for $[\alpha]\varphi$ which is slightly closer to DEL at the cost of more complex semantic models.

We denote the translation of a DEL formula φ as $(\varphi)^T$.

$$\begin{aligned}
(p)^T &= p \\
(\varphi \wedge \psi)^T &= (\varphi)^T \wedge (\psi)^T \\
(\neg\varphi)^T &= \neg(\varphi)^T \\
(\Box_i\varphi)^T &= K_i(\varphi)^T \\
([\alpha]\varphi)^T &= \mathit{after}(\alpha, \varphi)
\end{aligned}$$

where $\mathit{after}(\alpha, \varphi)$ is inductively defined:

$$\begin{aligned}
\mathit{after}(\psi!_\beta, \varphi) &= \mathsf{ann}(\psi, \beta) \Rightarrow \bigcirc\varphi \\
\mathit{after}(\alpha_1 \vee \alpha_2, \varphi) &= \mathit{after}(\alpha_1, \varphi) \vee \mathit{after}(\alpha_2, \varphi) \\
\mathit{after}(\alpha_1; \alpha_2, \varphi) &= \mathit{after}(\alpha_1, \mathit{after}(\alpha_2, \varphi))
\end{aligned}$$

This allows us to interpret DEL formulas as formulas in our logic. In order to prove that the interpretation of DEL formulas in our logic are sound, we need the following result:

Proposition 2.1 *An honest public announcement system satisfies the following property:*

$$(r', m'+1) \sim_i (r, m+1) \quad \Rightarrow \quad (r, m) \sim_i (r', m')$$

Proposition 2.2 Our translation is faithful to DEL with respect to *honest public announcement systems*.

Proof. We need to show that DEL's axioms (other than the usual modal logic ones) are valid in our system. These are

D1. $[\theta!_\beta]p \Leftrightarrow (\theta \Rightarrow p)$.

D2. $[\theta!_\beta]\neg\varphi \Leftrightarrow (\theta \Rightarrow \neg[\theta!_\beta]\varphi)$

D3. $[\theta!_\beta]\Box_i\psi \Leftrightarrow (\theta \Rightarrow \Box_i[\theta!_\beta]\psi)$.

D1 represents *preservation of facts* – that epistemic actions do not affect the truth value of primitive propositions. This is true in our setting via the validity

$$p \Leftrightarrow \bigcirc p.$$

D2 represents *partial functionality*, it can be expressed in our setting as follows

$$(ann(\theta, \beta) \Rightarrow \bigcirc\neg\varphi) \Leftrightarrow (\theta \Rightarrow \neg(ann(\theta, \beta) \Rightarrow \bigcirc\varphi))$$

and follows from the standard semantics of our temporal modality.

D3 represents the *action-knowledge axiom* – that agents know the consequences of actions: they know a fact after an action exactly when they know that performing that action yields that fact. This can be expressed in our setting as follows:

$$K_i(\mathsf{ann}(\varphi, \beta) \Rightarrow \bigcirc\psi) \Leftrightarrow \mathsf{ann}(\varphi, \beta) \Rightarrow \bigcirc K_i\psi.$$

and follows from proposition 2.1. $\Box$

3 Private Announcements and Belief

Consider a system of three agents where agent 1 tosses a coin, witnessed by no one. He then announces privately to himself and to agent 2 the result of the toss. Consider primitive propositions H and T, true at states where H is the result of the coin toss, and T is the result of the coin toss. The corresponding system has two runs, the initial state of the environment in the first run records that the coin landed heads, while the second run records that the coin landed tails. After time 0, agent 1 announces the result of the coin toss, which is recorded in the state of the agents at time 1. In the system corresponding to this scenario, the runs are r^T, r^H, the states of the environment is $r^T_e(1) = (T, \langle T!_{\{1,2\}}\rangle)$, and the local states of agents are $r^T_1(1) = r^T_2(1) = \langle T!_{\{1,2\}}\rangle$, but $r^T_3(1) = \langle \surd \rangle$, and similarly for r^H.

What is true in this system?

$$(r^H, 0) \models \neg K_1(H) \wedge \neg K_1(T)$$

$$(r^H, 1) \models K_1(H) \wedge K_2(H) \wedge \neg K_3(H) \wedge \neg K_3(T)$$

$$(r^H, 1) \models \neg K_3 K_2 H.$$

In DEL, the knowledge operator seems to capture something sligtly different. In particular, they establish that at time 1 on run r^H, $\Box_3 \neg \Box_2(H)$. However, this says that agent 3 knows something false. Since knowledge is generally taken to satisfy axiom T, that is $K\varphi \Rightarrow \varphi$, the DEL modality $\Box_i$ could be more appropriately interpreted as some sort of belief.

How do we capture the notion of $\Box_i$ in our setting? Clearly, our K_i operator does not let us know false formulas. We therefore introduce a new operator, $B_i\varphi$, and appropriate semantic machinery to interpret it.

What is going on in the example above, is that agent 3 *believes* that agent 2 does not know H because, having not seen the private announcement, he assumes that there was no announcement whatsoever.[2] This is the source of the wrong belief. In order to model this, we need to define operator B_i in such a way that agent i considers possible other states where other agents have in fact not received the purported private announcement. Doing this requires adding runs to the system to capture these non-real executions.

Accordingly, we define an honest *belief-based private announcement system* to be an honest private announcement system augmented with the following runs:

A5. For every real run r, and time m, there exists another run r' such that if $r_e(m)$ has announcements $\langle b_1, \ldots, b_k\rangle$, then $r'_i(m)$ has last announcement $\surd$ (instead of possibly b_k) for all i, and for each $m' \geq m$, the local state of each agent has $\surd$ as the k-th announcement, instead of possibly b_k (in other words, the replacement of b_k by $\surd$ is done consistently across the run.)[3].

[2] But if he suspects that there was such an announcement, then he cannot get that belief, and in fact does not get any information out of the announcement that he does not witness, whereas if he does not suspect it, he gets some (possibly wrong) information out of something he does not even know happened.

[3] Note that we potentially change the meaning of K_i for future propositions. (For instance, $K_1 \bigcirc K_2 p$ for some p.)

Equipped with all these runs, we can interpret belief as follows:

$(\mathcal{A}, r, m) \models B_i\varphi$ iff either

- $last(r_i(m)) = \theta!_\beta$ and for all $(r', m') \sim_i (r, m)$, $(r', m') \models \varphi$, or
- $last(r_i(m)) = \surd$ and for all $(r', m') \sim_i (r, m)$ such that $last(r'_j(m')) = \surd$ for all j, $(r', m') \models \varphi$.

Definition 3.1 An honest *belief-based private announcement system* is an honest private announcement system which moreover satisfies A5.

Example 3.2 *We can check that for the coin-toss problem, we have the following which is clearly an instance of false belief:*

$$(r^H, 1) \models B_3\neg K_2(H) \wedge K_2(H),$$

This new modality makes our translation sound with regard to honest belief-based private announcement systems, but first we need a similar result on equivalence classes:

Proposition 3.3 *An honest belief-based private announcement system satisfies the following property:*

$$(r', m'+1) \sim_i (r, m+1) \quad \Rightarrow \quad (r, m) \sim_i (r', m')$$

Proposition 3.4 In our previous translation if we change $(\Box_i\varphi)^T = K_i(\varphi)^T$ to $(\Box_i\varphi)^T = B_i(\varphi)^T$, then the translation would become faithful to DEL with respect to honest belief-based private announcement systems.

Proof. The only difference with honest public announcement systems is the translation of axiom D3, which is now expressed as follows

$$B_i(\mathsf{ann}(\varphi, \beta) \Rightarrow \bigcirc\psi) \Leftrightarrow \mathsf{ann}(\varphi, \beta) \Rightarrow \bigcirc B_i\psi.$$

and follows from the interpretation of B_i and proposition 3.3. $\Box$

If we assign an *announcer agent* to each announcement, as a map from the set of all announcements to the set of all agents, then we can show the following

Proposition 3.5 *For all r, m, θ, if $(\mathcal{A}, r, m) \models ann(\theta, \beta)$, then $(\mathcal{A}, r, m) \models B_i\theta$ where i is the announcer agent of the $\theta!_\beta$ announcement.*

4 Future Work

We will define systems with dishonest announcements, we believe this is easily achieved similar to private announcements, that is by augmenting the runs and add new clauses to the definition of B_i. We also plan to prove that soundness of our translation will extend to these models. Also in this case, we believe that this is easily done by similar unification results on the equivalence classes (propositions 2.1 and 3.3).

References

[1] A. Baltag and L.S. Moss. 'Logics for epistemic programs'. *Synthese* **139**, 2004.

[2] A. Baltag, B. Coecke, and M. Sadrzadeh, 'Epistemic actions as resources', *Journal of Logic and Computation* **17 (3)**, May 2007.

[3] J. van Benthem, J. van Eijck, and B. Kooi, 'Logics of Communication and Change', *Journal of Information and Computation* **204**, 2006.

[4] J. van Benthem and E. Pacuit, 'The tree of knowledge in action: towards a common perspective', *Proceedings of Advances in Modal Logic*, 2006.

[5] W. van Der Hoek and M. Wooldridge, 'Time, Knowledge, and Cooperation: Alternating-Time Temporal Epistemic Logic', COORDINATION 2002.

[6] R. Fagin, J. Y. Halpern, Y. Moses and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[7] J. Gerbrandy, 'Dynamic Epistemic Logic', in L.S. Moss et al. (eds.) *Logic, Language, and Information* **2**, Stanford University, CSLI Publication, 1999.

[8] P. Panangaden and K. Taylor, 'Concurrent Common Knowledge: Defining Agreement for Asynchronous Systems', *Journal of Distributed Computing* **6**, 1992.