A Note on a Generalization of the Muddy Children Puzzle

Nina Gierasimczuk[*]
Department of Artificial Intelligence
Postbus 407
9700AK Groningen
Nina.Gierasimczuk@gmail.com

Jakub Szymanik[†]
Department of Artificial Intelligence
Postbus 407
9700AK Groningen
Jakub.Szymanik@gmail.com

ABSTRACT

We study a generalization of the Muddy Children puzzle by allowing public announcements with arbitrary generalized quantifiers. We propose a new concise logical modeling of the puzzle based on the number triangle representation of quantifiers. Our general aim is to discuss the possibility of epistemic modeling that is cut for specific informational dynamics. Moreover, we show that the puzzle is solvable for any number of agents if and only if the quantifier in the announcement is positively active (satisfies a form of variety).

Categories and Subject Descriptors

I.2.4 [**Computing Methodologies**]: Artificial Intelligence—*Knowledge Representation Formalisms and Methods*; I.2.6 [**Computing Methodologies**]: Artificial Intelligence—*Learning*; H.4 [**Information Systems Applications**]: Miscellaneous; F.4 [**Mathematical Logic and Formal Languages**]: Miscellaneous

General Terms

Theory

Keywords

Muddy Children puzzle; generalized quantifiers; number triangle; epistemic logic; abstraction techniques

1. INTRODUCTION

In this paper we propose a generalization of the popular Muddy Children puzzle by allowing public announcements with arbitrary quantifiers. We focus on the solvability issue, i.e., on the possibility of the convergence to knowledge

for all the agents. We prove that the puzzle is solvable for any number of agents if and only if the quantifier in the announcement is positively active (satisfies a form of variety) (see Van Benthem, 1984). For example, announcing 'Most of you are muddy' always leads to the successful reasoning while announcing 'At most 5 of you are muddy' does not. Relating to that we explore a possibility of a concise logical modeling of dynamic epistemic situations, like the Muddy Children puzzle. Our general aim is to discuss the possibility of epistemic modeling that is cut for specific informational dynamics.

A considerable amount of philosophical and logical literature has been devoted to the epistemic inferences of the Muddy Children scenario (Littlewood, 1953)[1]. In particular, the framework of dynamic epistemic logic allows a clear and comprehensive explanation of the underlying phenomena (see Van Ditmarsch et al., 2007; Gerbrandy, 1999; Moses et al., 1986). Let us briefly recall the puzzle and its classical modeling. The scenario involves a father and his three children—truthful, perfect logical reasoners. As a result of playing outside some of the children have mud on their foreheads. The father says: (A1) 'At least one of you has mud on your forehead'. Then, he asks the children: (**I**) 'Can you tell for sure whether you have mud on your forehead? If yes, announce your status'. Each child can see the mud on others but cannot see his or her own forehead. Nothing happens. The father insists—he repeats (**I**). Still no reaction. But after he repeats the question for the second time suddenly all muddy children know that they have mud on their forehead. How is that possible?

The problem is usually modeled with the help of Kripke structures describing agents' uncertainty. Let us call the three children a, b and c, and assume that, in fact, all of them are muddy. Three propositional letters m_a, m_b and m_c express that the corresponding child is muddy. The standard epistemic modeling is depicted in Figure 1, with the initial model of the situation on the left.

In the model, possible worlds correspond to the 'distribution of mud' on children's foreheads, e.g., $w_5 : m_a$ stands for a being muddy and b and c being clean in world w_5. Two worlds are joined with an edge labelled with x, if the two worlds are in the uncertainty range of agent x (i.e., if agent x cannot distinguish between the two worlds; for clarity we drop the reflexive arrows for each state). The boxed state stands for the actual world. Now, let us recall how the solution process can be modeled in this set-
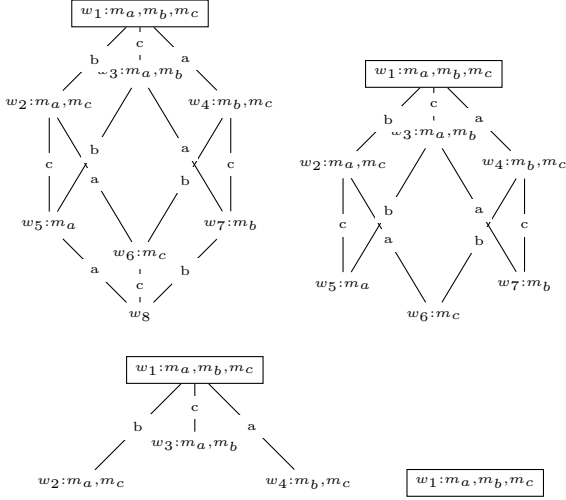
Figure 1: The Classical Muddy Children Modelling

ting. The first public announcement has the following form: $(1')$ $m_a \vee m_b \vee m_c$, and after the announcement $(1')$ becomes common knowledge among children. As a result the children perform an update, i.e., they eliminate world w_8 in which $(1')$ is false. The result is depicted in the second part of Figure 1. Then the father asks for the first time, who of them knows his status (**I**). The agents' reasoning can be as follows. In world w_6 agent c knows that he is dirty (there is no uncertainty of agent c between this world and another in which he is clean). Therefore, if the actual world was w_6, agent c would know his state and announce it. The situation is similar for a and b in w_5 and w_7, respectively. The silence of the children may also be interpreted as the announcement that none of them know whether they are muddy: $\neg(K_a m_a \vee K_a \neg m_a) \wedge \neg(K_b m_b \vee K_b \neg m_b) \wedge \neg(K_c m_c \vee K_c \neg m_c)$. Hence, all agents eliminate those worlds that do not make such announcement true: w_5, w_6, w_7. The epistemic model of the next stage is smaller by three worlds. Then it is again clear that if one of the w_2, w_3, or w_4 was the actual state, the respective agents would have announced their knowledge. The children still do not respond so, in the next round, everyone knows that the actual situation cannot be any of w_2, w_3, and w_4. Hence, they all eliminate the three states, which leaves them with just one possibility, w_1. All uncertainty disappears and they all know that they are dirty at the same time.

One of the features that make the Muddy Children puzzle surprising is that a very simple *quantitative* public announcement carries powerful *qualitative* information; agents must reason about their properties on the basis of some general quantitative statement. It seems that the possibility of convergence to knowledge depends somehow on the trade-off between the internal structure of the epistemic scenario and the amount of information provided by the public announcement. For instance, the truthful announcement 'At least one of you has mud on your forehead' allows epistemic reasoning that solves the puzzle for any configuration, while announcing 'An even number of you have mud on your forehead' leads to an *immediate* one-step solution that does not involve any epistemic reasoning. In what follows we investi-

gate those phenomena.

2. GENERALIZED MUDDY CHILDREN

Father's first announcement has the following form:

A1 **At least one** of you has mud on your forehead.

Sentence (A1) can be seen as a *background assumption* that makes the epistemic multi-agent inferential process possible. The quantifier announcement prepares the ground for epistemic reasoning, and enforces some structure on the situation. The natural question is: Does every announcement trigger the successful reasoning?

A simple but crucial observation is that the information provided by the father has the following form:

Q of you have mud on your forehead,

where Q may be substituted by various quantifiers, like 'At least one', 'An even number', 'One third' and so on. As expected, the informational power of an announcement depends on the quantifier. To see this, let us think of the Muddy Children situation as $M = (U, A)$, where U is the set of children and $A \subseteq U$ is the set of children that are muddy.[2] Of course, after father's announcement some models are no longer possible. Only those satisfying the quantifier sentence, i.e., $M \models Q_U(A)$, should be still considered. Therefore, the model of a given Muddy Children scenario consists of the structures satisfying the quantifier sentence. The agent's goal is to pinpoint one of them—the actual world. To explain this idea in more detail let us start with introducing the notion of generalized quantifiers.

DEFINITION 1 (MOSTOWSKI 1957). *A generalized quantifier Q of type (1) is a class of structures of the form $M = (U, A)$, where A is a subset of U. Additionally, Q is closed under isomorphism, i.e., if M and M' are isomorphic, then $(M \in Q \iff M' \in Q)$.*

Now, the classical Muddy Children puzzle with the father saying 'At least one of you has mud on your forehead' involves the existential generalized quantifier: $\exists = \{(U, A) : A \subseteq U \ \& \ A \neq \emptyset\}$. The variations with the father using different quantifiers may lead to other classes of possible situations, e.g., $\mathsf{Most} = \{(U, A) : A \subseteq U \ \& \ |A| > |U - A|\}$.

Isomorphism closure gives rise to the *number triangle* representation of quantifiers proposed by Van Benthem (1986). Every model belonging to a generalized quantifier of type (1) may be represented as a pair of natural numbers (k, n), where $k = |U - A|$ and $n = |A|$. In other words, the first number stands for the cardinality of the complement of A and the second number stands for the cardinality of A. The following definition gives a formal counterpart of this notion.

DEFINITION 2. *Let Q be a type (1) generalized quantifier. For any numbers $k, n \in \mathbb{N}$ we define a quantifier relation: $\mathsf{Q}(k, n)$ iff there are $U, A \subseteq U$ such that $|U| = n+k$, $|A| = n$, and $Q_U(A)$.*

PROPOSITION 1. *If Q is a type (1) generalized quantifier, then for all U and all $A \subseteq U$ we have: $Q_U(A)$ iff $\mathsf{Q}(|U - A|, |A|)$.[3]*

---

[2]In this paper we will restrict ourselves to finite models.
[3]For the proof see, e.g., (Peters and Westerståhl, 2006, p.96).

258

As we restrict ourselves to finite universes, we can represent all that is relevant for type (1) generalized quantifiers in the structure called number triangle, which simply enumerates all finite models of type (1). The node labeled (k, n) stands for a model in which $|U - A| = k$ and $|A| = n$. Now, every generalized quantifier of type (1) can be represented by putting '+' at those (k, n) that belong to Q and '−' at the rest. For example, the quantifier 'At least one' in number triangle representation is shown in Figure 2. Number triangle plays a crucial role in Generalized Quantifier Theory and it also comes handy in our study, as we can now interpret the pairs (k, n) as possible worlds.

$$(0,0)$$
$$(1,0) \quad (0,1)$$
$$(2,0) \quad (1,1) \quad (0,2)$$
$$(3,0) \quad (2,1) \quad (1,2) \quad (0,3)$$
$$(4,0) \quad (3,1) \quad (2,2) \quad (1,3) \quad (0,4)$$

....

$$-$$
$$- \quad +$$
$$- \quad + \quad +$$
$$- \quad + \quad + \quad +$$
$$- \quad + \quad + \quad + \quad +$$

....

Figure 2: Number triangle and the representation of 'At least 1'

3. EPISTEMIC MODELING BASED ON NUMBER TRIANGLE

How can we combine all introduced elements into an epistemic modeling in order to characterize the solvability of generalized Muddy Children puzzle? To answer that question let us analyze a concrete Muddy Children scenario. As before, we take agents a, b, and c. All possibilities with respect to the size of the set of muddy children are enumerated in the third level of the number triangle. Let us also assume at this point that the actual situation is that agents a, b are muddy and c is clean. Therefore, with respect to our representation the real world is $(1, 2)$, one child is clean and two are muddy:

$$(3,0) \qquad (2,1) \qquad \boxed{(1,2)} \qquad (0,3)$$

Now, let us focus on what the agents observe. Agent a sees one muddy child and one clean child. The same holds for agent b, in this sense they are perfectly symmetric. Their observational state can be encoded as $(1,1)$. Accordingly, the observational state of c is $(0,2)$. In general, if the number of agents is n, each agent can observe $n - 1$ agents. As a result what agents observe is encoded in the second level of the number triangle.

Figure 3: The Number-Triangle Muddy Children Modelling; the labelled arrows indicate that agents a, b are in the observational state $(1, 1)$, and agent c is in $(0, 2)$.

$$(2,0) \qquad (1,1) \qquad (0,2)$$
$$(3,0) \qquad (2,1) \qquad \boxed{(1,2)} \qquad (0,3)$$

The question that each of the agents is facing is whether he is muddy. For example, agent a has to decide whether he should *extend* his observation state, $(1, 1)$, to the left state $(2, 1)$ (a decides that he is clean) or to the right state $(1, 2)$ (a decides that he is muddy). The same holds for agent b. The situation of agent c is similar, his observational state is $(0, 2)$ and it has two potential extensions $(1, 2)$ and $(0, 3)$. In general, note that every observational state has two possible successors.

Given this representation, we can now analyze what happens in the Muddy Children scenario. Figure 3 represents the process, with the initial model at the top. First, the announcement is given: '**At least one** of you is muddy'. According to the number triangle representation (see Figure 2 on the right), this allows eliminating those factual states representing finite models that are not in the quantifier. In this case it is $(3, 0)$. The resulting model is the second from the top. Then the father asks: 'Can you tell for sure whether or not you have mud on your forehead?' In our graph, this question means: 'Does any of you have only one successor?' All agents know that $(3, 0)$ has just been eliminated. Agent a considers it possible that the actual state is $(2, 1)$, i.e., that two agents are clean and one is muddy, so that he himself would have to be clean. But then he knows that there would have to be an agent whose observational state is $(2, 0)$—there has to be a muddy agent that observes two clean ones. For this hypothetical agent the uncertainty disappeared just after the quantifier announcement (for $(2, 0)$ there is only one successor left). So, when it becomes clear that no one knows and the father asks the question again, the world $(2, 1)$ gets eliminated and the only possibility for agent a is now $(1, 2)$ via the right successor, and this indicates that he has to be muddy. Agent b is in exactly the same situation. They both can announce that they know. And since c witnessed the

whole process he knows that the only way for them to know was to be in $(1,1)$ and decides on $(1,2)$.

This epistemic reasoning took two steps. If the actual world was $(2,1)$ some agent's observation would be $(2,0)$, and this agent would know his status after the first announcement, and the rest of the agents would follow. Accordingly, for $(0,3)$ this would have taken three steps. This can be summed up in the following way: the quantifier breaks the perfect 'uncertainty structure' of the model, and the farther the actual state is from this break, the longer it takes to solve the puzzle (as will become clear in Section 5.1).

In general, if there are n agents, we take the nth level of the triangle, i.e., finite models with $|U| = n$, enumerating all possible settings (up to isomorphism). This level will be called the *factual level* and it always consists of $n+1$ states. It is an analogue of the initial uncertainty domain of the children in the classical modeling. Moreover, in the puzzle every child sees all other children, but not himself, so every possible observation consists of $n-1$ children. Therefore, level $n-1$ of the number triangle can be interpreted as enumerating every possible observation of the agents. We will call it the *observational level*. Each observation can be extended to one of the two factual states that are the closest below—to the left if the observer in question is clean or to the right if he is muddy.

4. RELATING TO KRIPKE EPISTEMIC MODELS

Let us now focus on clarifying the relation of our approach to the standard epistemic modeling via Kripke structures. First of all, note that every agent is in one of the two groups: either among the muddy children or among the clean ones. Moreover, in any situation there are at most two possible observational states that the agents might be distributed among. Every clean child observes the same quantitative situation as other clean children, and every muddy child observes quantitatively the same situation as all other muddy children.

FACT 1. *Every agent's observation is encoded by one of at most two states in the observational level. Those two are neighboring states.*

PROOF. For the first part. Assume that the total number of children is n, the number of muddy children is m. Let us pick an agent and call him a. There are two possibilities, either a is muddy, then a's observational state is $(n-m, m-1)$, or a is clean, then a's observational state is $((n-m)-1, m)$. The two states neighbor each other in the model because they are the only two states that can be extended to the actual state $(n-m, m)$. $\square$

Moreover, it is easy to notice that the actual world determines the number of agents perceiving the same situation.

OBSERVATION 1. *Given a certain situation (c, m), where $c, m > 0$, there are c children that are in the observational state $(c-1, m)$ and m children in the observational state of $(c, m-1)$. In case $m = 0$ all children are in the observational state $(c-1, 0)$. In case $c = 0$ all children are in the observational state $(0, m-1)$.*

Therefore, it is possible to formalize our structures in the following way.

DEFINITION 3. Muddy Children model *for n children is a quadruple $\mathcal{M}_n^{MC} = (S_o, S_f, R_m, R_{\bar{m}})$, where:*

- $S_o = \{(c, m) \mid c+m = n-1\}$ *(the observational states),*

- $S_f = \{(c, m) \mid c+m = n\}$ *(the factual states),*

- $R_m \subseteq S_o \times S_f$, *such that $R_m((c_1, m_1), (c_2, m_2))$ iff $m_2 = m_1 + 1$,*

- $R_{\bar{m}} \subseteq S_o \times S_f$, *such that $R_{\bar{m}}((c_1, m_1), (c_2, m_2))$ iff $c_2 = c_1 + 1$.*

In other words, our model is a two-row fragment of the number triangle with a double successor relation. An agent having access from an observational state (c, m) to the factual state $(c, m+1)$ corresponds to the possibility that he is muddy. Every such two states are in the relation R_m. The analogous situation holds for (c, m), $(c+1, m)$ and $R_{\bar{m}}$. Therefore, it is easy to observe that the size of such models is linear with respect to the number of children:

OBSERVATION 2. *If $n \in \mathbb{N}$ is the number of children, then the Muddy Children model has $2n+1$ states.*

This is a significant improvement with respect to the classical modeling, which requires an exponential number of states (cf. Van Ditmarsch et al., 2007; Fagin et al., 1995).

In our setting generalized quantifiers can be interpreted as propositional letters evaluated over the factual states of the Muddy Children model. For any generalized quantifier of type (1) we take a propositional letter q. Now, let $\mathcal{M}_n^{MC} = (S_o, S_f, R_m, R_{\bar{m}})$ be a Muddy Children model, and $(c, m) \in S_f$, then the semantics of q can be defined in the following way:

$$\mathcal{M}_n^{MC}, (c, m) \models q \text{ iff } (c, m) \in \mathsf{Q}.$$

The quantifier cuts the initial uncertainty range. This cut is in fact an *update* of the factual level of the Muddy Children model with a corresponding 'quantifier' letter q. Below we define what happens to the general Muddy Children model when a quantifier is introduced.

DEFINITION 4. *Having the Muddy Children model $\mathcal{M}_n^{MC} = (S_o, S_f, R_m, R_{\bar{m}})$ and a generalized quantifier Q of type (1), we define the quantifier update of $\mathcal{M}_n^{MC}$ with the quantifier Q as resulting in the Q-Muddy Children model $\mathcal{M}_n^{MC}|q = (S'_o, S'_f, R'_m, R'_{\bar{m}})$ in the following way:*

- $S'_o = \{(c, m) \mid (c, m) \in S_o \ \& \ (\mathsf{Q}(c+1, m) \lor \mathsf{Q}(c, m+1))\}$,

- $S'_f = \{(c, m) \mid (c, m) \in S_f \ \& \ \mathsf{Q}(c, m)\}$,

- $R'_m = R_m \restriction (S'_o \times S'_f)$,

- $R'_{\bar{m}} = R_{\bar{m}} \restriction (S'_o \times S'_f)$.

Now, the epistemic information can be expressed with formulae evaluated in the observational states. We can say that an agent in an observational state (c, m) knows that φ_m (that he is muddy) if and only if the only successor of (c, m) is $(c, m+1)$. Moreover, if an epistemic announcement eliminates an observational state it also eliminates all its successors. As our modeling is Muddy-Children specific, we do not give here a full epistemic language.

As we have seen, our model is significantly smaller then the representation needed for the classical modeling. Why

is that? Does it imply some different kind of epistemic reasoning for achieving the solution? The trick is that in our modeling we make a heavy use of three facts.

First of all, the only thing that an agent needs to be able to decide at every round, except the knowledge of his own status, is if other agents know whether they are muddy.

Secondly, provided the information that each agent possesses in the Muddy Children scenario, explicitly representing the specific information about the distribution of muddiness is not necessary. To see this in full light, let $n \in \mathbb{N}$ be the number of agents. For each i, $0 \le i \le n$ we define $m_i :=$ 'Exactly i of the children are muddy'. It is easy to verify that for each agent a it is the case that at any state of the classical epistemic model $K_a m_0 \vee K_a m_1 \vee \ldots \vee K_a m_n$ is true if and only if $K_a m_a \vee K \neg m_a$.

Thirdly, within certain group (the group of muddy children or the group of clean ones) all the agents converge to knowledge in the same way and simultaneously. To make it more specific, take a pointed epistemic muddy children model (M, w) for n agents. Note, that w determines two groups of agents those that are muddy, and those that are clean. Then, for any two agents a, b from one of the two groups it is the case that if $K_a m_a \vee K_a \neg m_a$ then $K_b m_b \vee K_b \neg m_b$.

This three observations lead to the following epistemic representation:

$$t_0 : m_0 \xrightarrow{\;0\;} t_1 : m_1 \xrightarrow{\;1\;} t_2 : m_2 \xrightarrow{\;0\;} \cdots \xrightarrow{\text{n-1 mod 2}} t_n : m_n$$

The choice of the actual world t_k determines the distribution of the agents among the two relations R_0 and R_1 in the following way, for $i, j \in \{0, 1\}$ such that $i \ne j$: if $R_i(t_k, t_{k+1})$ or $R_j(t_i, t_{i-1})$ then $R_i := R_{\bar{m}}$ and $R_j := R_m$. The intuition is that R_m (resp. $R_{\bar{m}}$) reflects the relevant uncertainty of all muddy (resp. clean) children. The above structure does not represent all epistemic facts of the classical modeling, but it gives all that is relevant for the convergence to knowledge for all the agents. Looking at the above model, one can observe a striking similarity to the dynamic epistemic logic modeling of the Consecutive Numbers puzzle (see e.g. Van Ditmarsch et al., 2007; Parikh, 1991).

Intuitively, the way in which we obtain the above concise model from the classical epistemic representation is by grouping the agents with respect to their informational similarity and introducing a new set of propositions that express the cardinality of the set of muddy children. Moreover, the states in our model merge the isomorphic worlds in the classical model. This leads to a conclusion that our modeling is in fact obtained by means of a merge applied to the classical Kripke model.[4]

5. MUDDY CHILDREN SOLVABILITY

---

[4]Our transformation is similar to abstraction-refinement techniques that have been developed for model-checking in temporal logics, modal transition systems (Godefroid et al., 2001). We are also aware of the similar attempts formulated in (Cohen et al., 2009). Recently abstraction-refinement techniques have been applied to dynamic epistemic logic (Wang, 2010). However, the above model does not comply to the rules of the abstraction-refinement technique defined in the latter. In fact, in line of Wang (2010) a similar model can be obtained, but it makes use of 3-valued Public Announcement Logic.

In this section, applying the epistemic modeling described above, we answer our initial question. Namely, we characterize the solvability of the generalized Muddy Children puzzle by determining the number of reasoning steps needed to reach the solution in a concrete setting. In particular, we characterize public announcements with arbitrary generalized quantifiers that trigger successful epistemic reasoning.

5.1 Number of epistemic iterations

By reinterpreting the Muddy Children puzzle within the semantics of quantifiers we can associate every finite model with the number of steps needed to solve the puzzle, if it is solvable at all.

DEFINITION 5. *An* epistemic quantifier *is a pair* $Q^{MC} = (Q, f_Q)$, *where Q is a quantifier and $f_Q : Q \to \mathbb{N}$ is a function that assigns to a pair of numbers representing $M \in Q$ the number of steps needed to solve an epistemic situation with quantifier Q (in particular, the Muddy Children puzzle with the background assumption containing quantifier Q).*

Now, we need to know how to determine values of f_Q for a given quantifier.

PROPOSITION 2. *Let Q be a generalized quantifier, and n be the number of children. Then the corresponding epistemic quantifier is $Q^{MC} = (Q, f_Q)$, where the partial function $f_Q : Q \rightharpoonup \mathbb{N}$ is defined in the following way.*

$$f_Q((n-m, m)) =$$

$\min(\mu_{x \le n-m}\, (n-m-x, m+x) \notin \mathsf{Q}, \mu_{y \le m}\, (n-m+y, m-y) \notin \mathsf{Q}),$

where function $\mu x\, \varphi(x)$ finds the smallest x that satisfies φ.

PROOF. Observe that the function assigns a value x to $(u-k, k)$ in the level u of the number triangle if $(u-k, k) \in Q$ and there is $(u - \ell, \ell)$ in the level u such that $(u - \ell, \ell) \notin Q$. Moreover, the value x encodes the distance from the nearest $(u - \ell, \ell)$ such that $(u - \ell, \ell) \notin \mathsf{Q}$. The proposition now follows from the merged modeling in Section 4. $\square$

Concerning the assignment of the number of steps needed for solving the puzzle, we can also ask what is the structure of those steps. Namely, we can characterize situations in which some agents infer their status from the announcements of other agents, in contrast to the cases in which it happens simultaneously (we use '+'-superscripts to identify those situations). The definition of the partial function $f_{\mathsf{Q}}^{+} : Q \rightharpoonup \{+\}$ can be then given in the following way. $f_Q^{+}((n - m, m)) = +$ iff:

(1) $f_Q((n - m, m))$ is defined, and

(2) $m \ne 0$ and $m \ne n$ and some agent considers two factual worlds possible.

For shaping the intuitions, let us give a few examples of epistemic quantifiers in the number triangle representation. First let us consider the quantifier 'At least k'. It is easy to observe that increasing k causes the downward triangle to move down along the $(0, 0)$–$(0, n)$ axis.

This quantifier allows solving the Muddy Children puzzle for any configuration of 'muddiness'. However, within a certain level, the farther from a minus the longer it takes.

Now let us have a look at the quantifier 'At most k'.

The question-marks occur in place of models that satisfy the quantifier, but for which it is impossible to solve

```
              —
           —     1
        —     1⁺     2
     —     1⁺     2⁺     3
  —     1⁺     2⁺     3⁺     4
  .............................

              —
           —     —
        —     —     1
     —     —     1⁺     2
  —     —     1⁺     2⁺     3
  .............................
```

Figure 4: Increasing muddy-quantifiers 'At least 1' and 'At least 2'

```
              —
           ?     ?
        2     1⁺     —
     2     1⁺     —     —
  2     1⁺     —     —     —
  .............................

              —
           ?     ?
        ?     ?     ?
     3     2⁺     1⁺     —
  3     2⁺     1⁺     —     —
  .............................
```

Figure 5: Decreasing muddy-quantifiers 'At most 1' and 'At most 2'

the Muddy Children puzzle. For example, if one child is clean and one child is muddy (the actual world is $(1, 1)$) the Muddy Children situation does not lead to a solution if the announcement is: 'At most two of you are muddy'. Again, the farther from a minus the longer it takes to solve the puzzle.

Parity quantifiers in the Muddy Children setting do not involve much inference—every situation is solvable in one step (see Figure 6) and all answers are given simultaneously by all the agents.

5.2 Characterization

The above discussion leads to the observation that solving the Muddy Children Puzzle is possible if the announcement of the quantifier leaves one observational state with just one successor. Therefore the solvability of the particular Muddy Children scenario can be characterized in the following way:

THEOREM 1 (MUDDY CHILDREN SOLVABILITY). *Let n be the number of children, $m \leq n$ the number of muddy children, and Q be the background assumption. A Muddy Children situation is solvable iff $(n - m, m) \in Q$ and there is an $\ell \leq n$ such that $(n - \ell, \ell) \notin Q$.*

PROOF. Let us fix n—the number of children and $m \leq n$—the number of muddy children, Q is the quantifier background assumption.

For left to right. Assume that the scenario ends successfully—all agents arrive to knowledge about their status. Assume towards contradiction that it is not the case that $(n - m, m) \in Q$ or it is not the case that there is $\ell \leq n$ such that $(n - \ell, \ell) \notin Q$.

- if $(n - m, m) \notin Q$ then the father's announcement is not truthful. Contradiction.

- if for all $\ell \leq n$ it is the case that $(n - \ell, \ell) \in Q$, then the public announcement of Q does not eliminate any world an thus the iterated epistemic reasoning is impossible and the convergence to knowledge fails for all the agents. Contradiction.

For the other direction, assume that $(n - m, m) \in Q$ and there is $\ell \leq n$ such that $(m - \ell, \ell) \notin Q$. Then by Proposition 2 $f_Q((n - m, m))$ is defined and hence the puzzle is solvable in $f_Q((n - m, m))$ steps. □

In fact, the solvability issue coincides with a known and very important property of generalized quantifiers.

DEFINITION 6 (VAN BENTHEM (1984)). *A quantifier Q is active (alternatively: Q satisfies variety, VAR) iff for every non-empty set U, there exists $A \subseteq U$ such that $Q_U(A)$ but there is also $B \subseteq U$ such that it is not the case that $Q_U(B)$.*

Note that VAR can be viewed as a conjunction of two weaker conditions[5], VAR$^+$ and VAR$^-$.

DEFINITION 7.

VAR$^+$ *A quantifier Q is positively active (alternatively: Q satisfies VAR$^+$) iff for every non-empty set U if there exists $A \subseteq U$ such that $Q_U(A)$, then there is also $B \subseteq U$ such that it is not the case that $Q_U(B)$.*

VAR$^-$ *A quantifier Q is negatively active (alternatively: Q satisfies VAR$^-$) iff for every non-empty set U if there exists $A \subseteq U$ such that it is not the case that $Q_U(A)$, then there is also $B \subseteq U$ such that $Q_U(B)$.*

Now, we can characterize the general Muddy Children Solvability in the following way:

COROLLARY 1 (MUDDY CHILDREN SOLVABILITY). *A Muddy Children situation with Q as the background assumption is solvable for any number of children and any distribution of muddiness iff Q is positively active.*

---

[5] Our focus on such forms of VAR is consistent with the usefulness of weaker variability assumptions in Generalized Quantifier Theory Van Benthem (1984).

```
                _
          1     _
       1     _     1
     1     _     1     _
   1     _     1     _     1
 1     _     1     _     1     _
 ..............................................


                _
          1     _
       1     _     _
     1     _     _     1
   1     _     _     1     _
 1     _     _     1     _     _
 ..............................................
```

Figure 6: Muddy-quantifiers 'Divisible by 2' and 'Divisible by 3'

6.   CONCLUSIONS AND PERSPECTIVES

In this paper we introduced a new kind of logical modeling of the puzzle based on the idea of the number triangle which is similar, in perspective, to the standard modeling of the Consecutive Numbers puzzle. In our approach the representation of the problem is exponentially smaller than in other models based on dynamic epistemic logic. In this sense, our work might be viewed as searching for more concise models that utilize some symmetry features as, for example, the recent work of Wang (2010). Our abstraction preserves properties guaranteeing that the *results* of the reasoning are in some sense equivalent to the ones we obtain in the concrete model. In our case the similarity is defined by achieving the same epistemic (cognitive) outcome. We observed that from agents local perspective solving the puzzle might be viewed as counting the number of muddy children, so we compressed our model by merging isomorphic states. As a result, we came up with concise modeling that may be attractive in all those applications where an agent's internal representation of the problem is crucial, like cognitive science or designing multi-agent systems in the domain of artificial intelligence. Moreover, the counting-like structure of the new model shows that Muddy Children puzzle might be in a way reducible to other problems, like Consecutive Numbers puzzle. This indicates that, perhaps, equivalence classes on the domain of epistemic processes could be defined and in this way their dynamics-oriented complexity could be captured.

The domain of the classical epistemic model can be partitioned by an equivalence relation, and the efficiency of the announcements depends on their 'compatibility' with this partition. In other words, all assertions that are used in the scenario either remove or retain whole partition cells. In that case, clearly, the update process will terminate with a number of steps measured by the number of equivalence classes, and not with the size of the actual model. In particular, in the Muddy Children puzzle the equivalence classes are given by permutations of individuals, and all relevant assertions, both the father's announcement and the children's

subsequent 'silence', respect that equivalence relation. This perspective on the *generic assertions* of the Muddy Children puzzle can be linked to the common assumption that information structures are partition based (see e.g. Aumann, 1999; Fagin et al., 1995) and, furthermore, to the rationality assertions of game solution procedures (see Van Benthem, 2007). Another interesting interpretation of the partition is the one with the notion of *issue* in dynamic epistemic logic of questions (Van Benthem and Minica, 2010): by choosing an optimal issue, we can speed up the learning processes dramatically.

There are many further methodological questions concerning our logical modeling. First is that of the generality of our approach. Even though our framework is clearly compatible with the one of dynamic epistemic logic in terms of (the structure and the number of) steps needed for completion of epistemic reasoning we still wonder if it can be extended in a way that allows epistemic logic flexibility. A possible direction would be to associate explicitly our local representations with computational procedures, e.g., by viewing the representation in terms of automata theory (cf. Van der Meyden, 1996). Secondly, our work includes extension of public announcements to arbitrary generalized quantifiers. This in itself leads to a number of important issues, e.g., what is the epistemic logic of quantifier public announcements?

In section 5 we characterize solvability of the Muddy Children puzzle with arbitrary generalized quantifier announcements using a formal counterpart of the intuitive condition of announcements' non-voidance. The characterization shows how general, independent from particular settings, properties of announcements influence convergence to knowledge. Similar properties of quantifiers have been widely studied in logic, linguistics and cognitive science. The present link to the multi-agent approach leads to a significant extension of the existing approach.

Our work generates many directions of follow-up research. For instance, we could consider combinatorially interesting situations with many predicates (e.g., children having spots of different colors on their foreheads), manipulate the observational power of the children or restrict their abilities to infer higher-order epistemic states to account for well-known processing bottlenecks (see, e.g., Verbrugge, 2009). Finally, distinguishing between factual and observational states in the proposed epistemic modeling can be used to investigate other types of epistemic inferences and puzzles, for example Russian Cards or the Top-Hat puzzle. In general, we hope that this fresh view on the old puzzle will motivate new developments in the study of agents' local perspective in multi-agent intelligent interaction.

References

Aumann, R. J. (1999). Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3):263–300.

Van Benthem, J. (1984). Questions about quantifiers. *Journal of Symbolic Logic*, 49(2):443–466.

Van Benthem, J. (1986). *Essays in Logical Semantics*. D. Reidel, Dordrecht.

Van Benthem, J. (2007). Rational dynamics and epistemic logic in games. *International Game Theory Review*, 9:1:13–45. Erratum reprint, Volume 9:2, 377–409.

Van Benthem, J. and Minica, S. (2010). Questions and issue management. Toward a dynamic logic of questions. to appear in the Journal of Philosophical Logic.

Cohen, M., Dam, M., Lomuscio, A., and Qu, H. (2009). A symmetry reduction technique for model checking temporal-epistemic logic. In *Proceedings of the 21st international jont conference on Artifical intelligence*, pages 721–726, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Van Ditmarsch, H., Van der Hoek, W., and Kooi, B. (2007). *Dynamic Epistemic Logic*. Springer Netherlands.

Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1995). *Reasoning About Knowledge*. MIT Press, Cambridge.

Gerbrandy, J. (1999). *Bisimulations on planet Kripke*. PhD thesis, Universiteit van Amsterdam.

Godefroid, P., Huth, M., and Jagadeesan, R. (2001). Abstraction-based model checking using modal transition systems. In Larsen, K. G. and Nielsen, M., editors, *CONCUR*, volume 2154 of *Lecture Notes in Computer Science*, pages 426–440. Springer.

Littlewood (1953). *A mathematician's miscellany*. London: Meuthen.

Van der Meyden, R. (1996). Finite state implementations of knowledge-based programs. In *FSTTCS'96: Proceedings of the Annual Conference on Foundations of Software Technology and Theoretical Computer Science 1996*, pages 262–273.

Moses, Y., Dolev, D., and Halpern, J. Y. (1986). Cheating husbands and other stories: A case study of knowledge, action, and communication. *Distributed Computing*, 1(3):167–176.

Mostowski, A. (1957). On a generalization of quantifiers. *Fundamenta Mathematicae*, 44:12–36.

Parikh, R. (1991). Finite and infinite dialogues. In *Proceedings of a Workshop on Logic from Computer Science*, pages 481–498. Springer.

Peters, S. and Westerståhl, D. (2006). *Quantifiers in Language and Logic*. Oxford University Press, Oxford.

Rabelais (German translation 1839). *Gargantua und Pantagruel,*. Herausgabe Gottlob Regis Berlin/Leipzig.

Verbrugge, R. (2009). Logic and social cognition. *Journal of Philosophical Logic*, 38(6):649–680.

Wang, Y. (2010). *Epistemic Modelling and Protocol Dynamics*. PhD thesis, Universiteit van Amsterdam.