R. RAMANUJAM

A discussion on explicit knowledge

Dedicated to Rohit Parikh on his 60^{th} birthday

Abstract

In the papers on knowledge during the last decade, a recurrent theme is the distinction between implicit knowledge ascribed to an agent and explicit knowledge available to that agent. We discuss one aspect of this distinction: even a logically omniscient reasoner's knowledge may be limited by having only a *partial view* of the system. We study a logic where an agent has explicit knowledge of only those facts which hold in worlds visible to the agent.

1 A distinction

During the last decade, there has been extensive research into the notion of ascribing knowledge to reasoning agents, particularly in the contexts of distributed systems, artificial intelligence and mathematical economics. The recent book [FHMV95] provides an excellent summary of dominant themes in this line of research.

Following Hintikka's logical treatment, knowledge is studied as a modality in (typically) a propositional modal logic, with the semantics of formulas given in Kripke frames where the accessibility relation is an equivalence. Formally, we have a countable set of propositions $P = \{p_0, p_1, \ldots\}$ and the syntax of formulas is given by:

$$\Phi \ ::= \ p \ \in \ P \mid \neg \ \alpha \mid \alpha \ \vee \ \beta \mid \mathbf{K} \ \alpha$$

The frames are Kripke structures $F = (W, \sim)$, where W is a set of possible worlds, and $\sim$ is an equivalence relation called the indistinguishability relation for the reasoning agent. Models are frames with valuations: $M = (F, V)$, where $F = (W, \sim)$ and $V : W \to 2^P$. The semantics of formulas is then given by:

- $M, w \models p$ iff $p \in V(w)$.
- $M, w \models \neg\alpha$ iff $M, w \not\models \alpha$.
- $M, w \models \alpha \vee \beta$ iff $M, w \models \alpha$ or $M, w \models \beta$.

- $M, w \models \mathbf{K}\alpha$ iff $\forall w' \sim w$, $M, w' \models \alpha$.

The satisfiability and model checking problems for the logic are decidable, the set of valid formulas can be axiomatized, and we have a decent logic on hand.

The rumblings of discontent with this logic begin with the observation that according to this logic, a reasoner knows every logical consequence of any fact she knows; further she knows all the valid formulas. Human reasoners simply do not exhibit such **logical omniscience**, and hence the logic above does not reflect reasoning about human knowledge. On the other hand, if we consider the reasoner above to be a computing agent, we come up with the problem of **resource boundedness**. The agent has only limited computational resources at its disposal, and exploring all the logical consequences of an assumption is expensive business. In particular, the validity problem for the propositional calculus is already co-NP-complete, and for the reasoner to know all tautologies requires, at the least, a co-NP machine. See [FHMV 95] for a detailed treatment of this issue.

The papers of Rohit Parikh ([Pa 87a], [Pa 87b], [Pa 90] and [Pa 94]) provide all manners of criticism which include the two observations above and much more. For instance, the logic above ignores the difference between knowledge and **information**: "it is easier to believe that he (the former U.S. President) might have had the requisite information about the Iran-Contra affair and simply failed to make the necessary deduction which could lead to *knowledge*." The logic also ignores the difference between knowledge of a proposition and knowledge of the *sentence* denoting the proposition: it is quite conceivable that an agent knows 'p' and does not know 'q' even when p and q are logically equivalent propositions. See [Pa 87a] for a series of illuminating examples describing different notions of knowledge and the need to make subtle distinctions between these notions.

From the discussion in the literature (see bibliography of Chapters 9 and 10 in [FHMV 95]), we see a central thesis emerging and subsequently, two distinct problems that need to be studied. The thesis can be stated as follows:

- The notion of knowledge defined by the logic above is an **external** or **implicit** one. It is knowledge ascribed to the reasoner, rather than one possessed by the reasoner, human or computer. The latter kind of knowledge, which we may call **explicit knowledge**, can be seen as the ability to answer questions on whatever is claimed to be known.

Attempts at formalizing and studying explicit knowledge demands solution of two distinct problems:

- a *logical* problem, namely that of modal logics in which the modalities are not consequence-closed.
- a *computational* problem, as explicit knowledge depends on the algorithms used by the reasoner and the resources available to the algorithms.

Not surprisingly, much of the research on the former has been on logics with non-standard semantics of some kind avoiding logical omniscience (Chapter 9 of [FHMV 95]) and discussion of the latter involves the invocation of an algorithm (not part of the logic) to evaluate knowledge formulas (both in [Pa 87a] and [HMV 94]).

There is *yet another* way in which explicit knowledge differs from implicit knowledge, one which is especially important in the context of distributed computing. A reasoner operating in an environment has only a limited view of the world, and her explicit knowledge is determined by the **visibility** of world states to her, whereas ascribed implicit knowledge of the agent depends on her behaviour in all possible worlds. In terms of the indistinguishability relation $\sim$ above, computing $\sim$ in itself may necessitate some effort. We can conceive of a situation where $s_0 \sim s_1 \sim s_2$, but at state s_0 only s_0 and s_1 are visible to the agent and not s_2. That is, an observer who has access to complete information about all the world states may declare that the agent would behave in the same way in all these three states, but the agent may require computation to realise this. If the agent makes the effort and computes the world state s_2, it may also realise that s_2 is in the equivalence class. However a resource limited agent might need to reuse its resources, and thus in computing the state s_2, it might 'forget' the information about s_0, for instance.

Thus, even among states within the same externally ascribed indistinguishability class, some may be 'farther' and some 'nearer' in terms of visibility to an agent. In the context of distributed systems, this happens routinely. A component of a system behaves in the same manner in all *global* system states in which its *local* state is the same. Hence all these states would be ascribed the same equivalence class for that agent. However, this component would typically communicate only with immediate neighbours in the network, and may not even be aware of the existence of many other components in the system. In this case, for the agent to compute its knowledge, the first task of computing these indistinguishable system states is well beyond its capabilities.

Moreover, as the system changes state, the view of the agent also changes, and this is also dependent on the computational resources available to the agent. For instance, a computational agent with bounded memory would forget events in the distant past. Even in static situations, computation may result in learning thereby enlarging visibility. The example about learning the prime factors of 143 in [Pa 87b] is relevant here. Thus, *logically indistinguishable states may well be computationally distinguishable*. Therefore, even granting logical omniscience to a view-limited reasoner, ascribed and computed knowledge can be quite distinct. This is one aspect of the distinction between implicit and explicit knowledge which seems to have been relatively less studied in the literature.

In this paper, we concentrate on this distinction, and to make matters simple, simply ignore the logical omniscience problem ! Our formulation of visibility-limited knowledge is inspired by the algorithm-based notion of knowledge studied in [FHMV 95] (Chapter 10) and [Pa 87a]. The latter works as follows: at any state s, when asked whether a formula ϕ holds, the reasoner evokes an algorithm available at s. If the

algorithm returns 'Yes', we say that the reasoner explicitly (or algorithmically) knows ϕ. Depending on whether the algorithm gives only 'Yes/No' outputs or whether it also has the possibility of a '?' output meaning 'I cannot find out within the resources available to me', we get the notions studied in [FHMV 95] or [Pa 87a], respectively. [FHMV 95] restricts the algorithm to be local, in the sense that when $s \sim s'$, the algorithm invoked by the reasoner at s is the same as at s'.

This is a very interesting formulation of explicit knowledge and avoids many of the philosophical pitfalls which cause the criticism above. In particular, we can see the invoked algorithm as a *model checking algorithm*. This means that the reasoner explicitly knows all valid formulas (as they definitely hold at the states being checked in any model), but also knows many non-valid formulas which happen to hold at that state and cannot tell which ones among the known formulas are valid. (See [HV 91] for arguments advocating model checking as the 'right' way to think about this problem.) The framework facilitates the study of knowledge based on probabilistic algorithms ([KNP 90]) and action based on knowledge ([HF 89], [R96]).

Unfortunately, the framework is also very general and unless we place restrictions on the class of algorithms available, the notion does not have interesting properties. For instance, if the system moves from a state s to state s', how should the algorithms available at the two states be related ? Further, the invocation of an algorithm is an extra-logical notion and as it is, we have no way of studying algorithmic knowledge *logically*.

In this paper, we study a framework in which a form of algorithmic knowledge is studied logically, and which takes into account the limited visibility of agents. The idea is simple: an agent claims that ϕ holds at a state if ϕ holds in the substructure visible to the agent at that state. This corresponds to running the model checking algorithm on the substructure and the given formula. We can then say that the agent claims to know ϕ if ϕ holds in the substructure at all the visible indistinguishable states. Thus, resource bounds (which limit visibility) as well as formula structure determine the ease or difficulty of explicit knowledge of formulas.

2 The basic logic

We now present the basic logic of view-based knowledge. This is simply the usual logic of knowledge augmented with a new modality: $\mathbf{X}\alpha$ will denote that the formula α holds explicitly (or visibly) for the reasoner. As we will see, this is more like a notion of belief, in the sense that $\mathbf{X}\alpha$ may hold, even when α does not. For explicit knowledge, we will use formulas like $\mathbf{XK}\alpha$ or $\mathbf{KX}\alpha$, depending upon what we require. (As [FHMV 95] observe in Chapter 10, this should be seen as a reasoner claiming to know α rather than knowing α.)

Formally, we have a countable set of propositions $P = \{p_0, p_1, \ldots\}$ as before and the syntax of formulas is given by:

$$\Psi \ ::= \ p \ \in \ P \mid \neg \ \alpha \mid \alpha \ \vee \ \beta \mid \mathbf{K} \ \alpha \mid \mathbf{X}\alpha$$

The frames are Kripke equivalence structures augmented with view functions.

Definition 2.1 *A* **view frame** *is a tuple $F = (W, \sim, \nu)$, where W is a set of worlds, $\sim \subseteq W \times W$ is an equivalence on W, and $\nu : W \to 2^W$ is a map such that $\forall w \in W$, $w \in \nu(w)$.*

A model is a tuple $M = (F, V)$, where $F = (W, \sim, \nu)$ is a frame, and $V : W \to 2^P$ is the valuation.

When $w' \in \nu(w)$, we say that w' is *visible* from w. Note that every world state is visible to the reasoner at that state in view frames. While this is a reasonable assumption, it is easy to conceive of situations where this may be unacceptable. We will discuss such situations later. When we have $w \sim w'$ and $w' \notin \nu(w)$, we can interpret this as w and w' being logically indistinguishable, but computationally distinguishable. (This happens in situations where discovering that w and w' constitute the same evidence requires computation).

Definition 2.2 *Given a view frame $F = (W, \sim, \nu)$ and $W' \subseteq W$, the* restriction *of F to W' is given by: $F\lceil W' \stackrel{\text{def}}{=} (W', \sim', \nu')$, where $\sim' \stackrel{\text{def}}{=} \sim \cap (W' \times W')$ and $\forall w \in W'$, $\nu'(w) \stackrel{\text{def}}{=} \nu(w) \cap W'$. For a model $M = (F, V)$, we define $M\lceil W' \stackrel{\text{def}}{=} (F\lceil W', V')$, where for $w \in W'$, $V'(w) = V(w)$.*

The intended meaning of $\mathbf{X}\alpha$ at w should be clear now: it simply asserts that α holds in the substructure induced by $\nu(w)$ at w. We assume that the semantics of the other connectives being as before.

$$M, w \models \mathbf{X}\alpha \ \text{ iff } \ M\lceil \nu(w), w \models \alpha$$

(The technical convenience in having $w \in \nu(w)$ is obvious now.) As usual, we say α is satisfiable if there exists a model $M = (F, V)$, where $F = (W, \sim, \nu)$ is a view frame and for some $w \in W$, we have $M, w \models \alpha$. In addition, we say α is *explicitly satisfiable* if there exists a model $M = (F, V)$, where $F = (W, \sim, \nu)$ is a view frame and for some $w \in W$, we have $W = \nu(w)$ and $M, w \models \alpha$. A formula is (explicitly) valid if and only if its negation is not (explicitly) satisfiable.

The first thing to note is that the modality does not collapse — in the sense that there are formulas α such that $\alpha \wedge \mathbf{X}\neg\alpha$ is satisfiable. To see this, consider the view frame given by: $W = \{w_1, w_2\}$, $w_1 \sim w_2$ and $\nu(w_j) = \{w_j\}$ for $j \in \{1, 2\}$. Let M be a model on this frame with the valuation: $V(w_1) = \{p\}$ and $V(w_2) = \emptyset$. Clearly, we see that $M, w_1 \models (\mathbf{X}\mathbf{K}p \wedge \neg\mathbf{K}p)$.

Secondly, $\mathbf{X}$ is its own dual. Further, though $\mathbf{X}\alpha$ does not generally imply α, it does imply $\mathbf{X}\mathbf{X}\alpha$, and $\neg\mathbf{X}\alpha$ implies $\mathbf{X}\neg\mathbf{X}\alpha$. Thus we have a modality which is not quite like a knowledge modality and somewhat stronger than a belief modality.

Note that $\alpha \equiv \mathbf{X}\alpha$ is explicitly valid, as we would expect.

A logician may shudder in horror at this point: there are no conditions at all on view functions and determining the truth of a formula in any structure seems to involve arbitrary substructures and induced partitions ! However, the following lemma asserts that things are not so bad.

Definition 2.3 *A view frame $F = (W, \sim, \nu)$ is said to be an* **equivalence** *frame, if for all $w \in W$, $\nu(w) \subseteq [w]$. It is said to be* **universal** *if for all $w, w' \in W$ such that $w \sim w'$.*

Lemma 2.4 *Every satisfiable formula is satisfiable in a model based on a universal equivalence frame.*

Proof: Given a formula α_0 satisfied at w_0 in model $M = ((W, \sim, \nu), V)$, define the new model $M' = M\lceil[w_0]$. We can then show by induction (on the structure of α) that, for all α, $\forall w \in [w_0]$, $M, w \models \alpha$ if and only if $M', w \models \alpha$; and, $\forall S \subseteq W$ such that $w \in S$, $M\lceil S, w \models \alpha$ if and only if $M'\lceil(S \cap [w_0]), w \models \alpha$. □

Unfortunately the logic is still hard to analyse. Consider a universal equivalence structure $W = \{w_0, w_1, \ldots\}$, where for all j, $\nu(w_j) = \{w_j, w_{j+1}\}$ and some model M based on it. Then for any modal formula α, the truth of α at w_j $(j > 0)$ in the structure $M\lceil\nu(w_j)$ is in general independent of the truth of the same formula at the same state in the structure $M\lceil\nu(w_{j-1})$. This is what contributes to the complexity of any decision procedure for the logic.

Theorem 2.5 *The satisfiablity of a formula α can be decided in time $2^{2^{O(k)}}$, where $k = |\alpha|$.*

The proof is filtration-based, but the quotienting is not only by identifying worlds which satisfy the same subformulas of the formula being decided. We also need to check whether the equated worlds induce the same substructures with respect to these sets of (sub)formulas. The details are a bit involved and we omit them here for the sake of brevity. The upper bound above may well be improved upon, we leave that for future work.

We can also get a solution to the model checking problem for the logic, where the procedure is linear in the size of the given model and singly exponential in the size of the given formula. This is deceptive, since there is already a hidden exponent in the presentation of the given frame (in the view function), and it is the guessing of this function that contributes to the double exponent in the decision procedure for satisfiability.

While we can easily think of many sound axioms, coming up with a complete set seems hard, mainly because of the plethora of structures involved. All this hardly

seems to be good advertisement for the logic. In our opinion, this is only because of the generality of the framework offered, and points to the need to study reasonable subclasses. In the next section, we consider a restricted class of frames, on which we have better decision procedures as well as an axiomatization.

3 Monotone frames

In the example discussed in the last section, we considered a situation where $w_0 \sim w_1 \sim w_2$ where $\nu(w_0) = \{w_0, w_1\}$ and $\nu(w_1) = \{w_1, w_2\}$. We might see this as the agent 'forgetting' the information about w_0 at w_1; for instance, this may be a case of reusing memory to compute the state w_2. We now consider the subclass of frames where this cannot happen, and in which visibility preserves logical indistinguishability.

Definition 3.1 *A view frame $F = (W, \sim, \nu)$ is said to be* **monotone** *if and only if $\forall w, w_1, w_2 \in W$, if $\{w_1, w_2\} \subseteq \nu(w)$ and $w_1 \sim w_2$ then $\nu(w) \cap \nu(w_1) = \nu(w) \cap \nu(w_2)$.*

In monotone frames, once a logically indistinguishable world becomes visible from another, they both contain the same information. The following properties of view functions make this clearer.

Proposition 3.2 *Let $F = (W, \sim, \nu)$ be a monotone frame.*

1. *$\forall w_1, w_2 \in W$, if $w_2 \in \nu(w_1)$ and $w_1 \sim w_2$ then $\nu(w_1) = \nu(w_2)$.*
2. *$\forall w, w_1, w_2 \in W$, if $\{w_1, w_2\} \subseteq \nu(w)$ and $w_1 \sim w_2$ then $\nu(w_1) = \nu(w_2)$.*

Proof: Since for any w, $w \in \nu(w)$, the definition of monotone frames at once implies that whenever $w' \in \nu(w)$ and $w \sim w'$, $\nu(w) = \nu(w) \cap \nu(w')$, that is, $\nu(w) \subseteq \nu(w')$. This already establishes one direction of (1). We now have $w \in \nu(w')$, and repeating the argument, we get $\nu(w') \subseteq \nu(w)$, establishing the first statement. The definition of monotonicity and the hypotheses of (2) together imply that $w_1 \in \nu(w_2)$ and $w_2 \in \nu(w_1)$, and the claim follows from (1). □

Proposition 3.3 *If $F = (W, \sim, \nu)$ is a monotone frame and $W' \subseteq W$, then so also is $F\lceil W'$.*

We say α is *monotone* satisfiable if there exists a model $M = (F, V)$, where $F = (W, \sim, \nu)$ is a monotone frame and for some $w \in W$, we have $M, w \models \alpha$. A formula is monotone valid if and only if its negation is not monotone satisfiable.

Note that the implication $\mathbf{KX}\alpha \supset \mathbf{XK}\alpha$ is monotone valid, though not valid. To see this, suppose $M, w \models \mathbf{KX}\alpha$ in a monotone model M. To show that we

also have $M\lceil\nu(w), w \models \mathbf{K}\alpha$, let $w' \in \nu(w)$ and $w \sim w'$. We need to show that $M\lceil\nu(w), w' \models \alpha$. By proposition above, $\nu(w) = \nu(w')$ and hence it suffices to show that $M\lceil\nu(w'), w' \models \alpha$, that is $M, w' \models \mathbf{X}\alpha$ which follows from our assumption that $M, w \models \mathbf{KX}\alpha$.

The advantage of monotone frames is that determining the truth of a formula in a structure of size k no longer necessitates such an exercise in 2^k substructures, as the following lemma asserts.

Lemma 3.4 *Let $M = ((W, \sim, \nu), V)$ be a monotone model. Then for every formula α and $w, w' \in W$ such that $w' \in \nu(w)$,*

$$M\lceil\nu(w), w' \models \alpha \quad \textit{iff} \quad M\lceil(\nu(w) \cap \nu(w')), w' \models \alpha.$$

Proof: The proof is by induction on the structure of α. The only interesting case is when α is of the form $\mathbf{K}\beta$. Here we use the fact that for any $w'' \in \nu(w)$ such that $w'' \sim w'$, $\nu(w') = \nu(w'')$, by Proposition 3.2 (2) above. □

The benefits bestowed by this lemma are immediate:

Theorem 3.5 *The monotone satisfiability problem is NP-complete. The monotone model checking problem is decidable in time $O(n.k)$, where k is the size of the given structure and n is the length of the checked formula.*

We now present an axiomatization of monotone valid formulas. We present it in two layers, one for explicitly valid formulas, in which the agent is supposed to be reasoning, and one for monotone valid ones, in which we ascribe knowledge to the agent. We use the notation $\vdash_X$ to denote formulas proved as theorems in the system for explicit validity and $\vdash$ for the theorems of the system for monotone validity.

Ax, The axiom system for monotone valid formulas

- $(A0)$ All the substitutional instances of the tautologies of PC
- $(A1)$ $\mathbf{K}(\alpha \supset \beta) \supset (\mathbf{K}\alpha \supset \mathbf{K}\beta)$
- $(A2)$ $\mathbf{X}(\alpha \supset \beta) \supset (\mathbf{X}\alpha \supset \mathbf{X}\beta)$
- $(A3)$ $\mathbf{K}\alpha \supset \alpha$
- $(A4)$ $\mathbf{K}\alpha \supset \mathbf{KK}\alpha$
- $(A5)$ $\neg\mathbf{K}\alpha \supset \mathbf{K}\neg\mathbf{K}\alpha$
- $(A6)$ $\neg\mathbf{X}\alpha \equiv \mathbf{X}\neg\alpha$
- $(A7)$ $\alpha \equiv \mathbf{X}\alpha$ *for boolean* α
- $(A8)$ $\mathbf{KX}\alpha \supset \mathbf{XK}\alpha$

Inference rules

(MP_i) $\dfrac{\alpha,\ \alpha \supset \beta}{\beta}$ $\quad$ (KG) $\dfrac{\alpha}{\mathbf{K}\alpha}$ $\quad$ (XG) $\dfrac{\vdash_X \alpha}{\mathbf{X}\alpha}$

Ax_X, The axiom system for explicitly valid formulas

$(B0)$ $\alpha \equiv \mathbf{X}\alpha$

Inference rules

(MP_X) $\dfrac{\alpha,\ \ \alpha \supset \beta}{\beta}$ $\qquad$ (IE) $\dfrac{\vdash \alpha}{\alpha}$

Theorem 3.6 *Ax provides a complete axiomatization of monotone valid formulas.*

4 Discussion

We have argued that the distinction between logical distinguishability and visibility is one worth making and attempted to demonstrate that a logic of knowledge bearing such features is technically interesting. While this distinction has the sort of philosophical motivation presented in Section 1, there is also a more pragamatic reason for considering it, and that has to do with the way we model systems. One important reason for studying knowledge in the context of computing agents is that the agents' *actions* at any instant of time are based on the knowledge at that time. This interplay of knowledge, time and action is a major theme in the literature on knowledge in distributed systems (see [HF 89] and [FHMV 95]). However, knowledge based action must involve computed knowledge and it is here that we feel that the distinction made above is crucial, as the agent being resource limited, may have only limited visibility.

To study this, we need to enrich the frames to include time and action. Further, we have only talked about one reasoner so far, and this needs to be generalized to systems of many reasoners. We believe that the class of Knowledge Transition Systems introduced in [KR 94], [R 96], augmented with view functions provide a suitable framework for such a study. Correspondingly the logic would also need to be enriched with modalities talking about reachability. The advantage of such a set-up is that resource limitations can be modelled in a natural way: for instance, we can consider systems where views are uniformly bounded to have at most k states; or insist that only the adjacent states (in the graph) be visible at any state. We can then study interactions between agents one of whom does a one-step-look-ahead whereas the other always looks ahead two steps, and so on.

A technical remark: there are situations when the assumption $w \in \nu(w)$ may need to be dropped. For instance, in a distributed system, an agent waiting for some event may see at a global state w only the last state when it was active, and hence we may have $w \notin \nu(w)$. For defining a logic on such frames, we need another modality which refers to the "latest" visible state among those in $\nu(w)$. We have avoided this here in the interest of simplicity.

References

[FHMV95] Fagin, R., Halpern, J., Moses, Y. and Vardi, M., *Reasoning about knowledge*, M.I.T. Press, 1995.

[HF89] Halpern, J., and Fagin, R., "Modelling knowledge and action in distributed systems", *Distributed Computing*, vol 3, #4, 1989, 159-177.

[HMV94] Halpern, J., Moses, Y., and Vardi, M., "Algorithmic knowledge", *TARK V*, Theoretical Aspects of Reasoning about Knowledge, 1994, 255-266.

[HV91] Halpern, J., and Vardi, M., "Model checking vs theorem proving: a manifesto", in *Artificial Intelligence and Mathematical Theory of Computation*, (ed) V. Lifschitz, Acad Press, 1991, 151-176.

[KNP90] Krasucki, P., Ndjatou, G. and Parikh, R., "Probabilistic knowledge and probabilistic common knowledge", in *ISMIS 90*, International Symp. on Methodology for Intelligent Systems, 1990, 1-8.

[KR94] Krasucki, P., and Ramanujam, R., "Knowledge and the ordering of events in distributed systems", *TARK V*, Theoretical Aspects of Reasoning about Knowledge, 1994, 267-283.

[Pa87a] Parikh, R., "Knowledge and the problem of logical omniscience", *ISMIS 87*, International Symp. on Methodology for Intelligent Systems, 1987, 432-439.

[Pa87b] Parikh, R., "Some recent applications of knowledge" in *FST and TCS 7*, Foundations of Software Technology and Theoretical Computer Science, 1987, LNCS #287, 528-539.

[Pa90] Parikh, R., "Recent trends in reasoning about Knowledge" *TARK III*, Theoretical Aspects of Reasoning about Knowledge, 1990, 3-10.

[Pa94] Parikh, R., "Logical omniscience" in *Logic and Computational Complexity* LNCS # 960, 22-29.

[R96] Ramanujam, R., "Local knowledge assertions in a changing world", *TARK VI*, Theoretical Aspects of Rationality and Knowledge, 1996, pp. 1-17.

The Institute of Mathematical Sciences
C.I.T. Campus, Madras 600 113
INDIA

jam@imsc.ernet.in