Privacy Loss in Classical Multiagent Planning

Roman van der Krogt*
Cork Constraint Computation Centre
Department of Computer Science, University College Cork
roman@4c.ucc.ie

Abstract

Privacy is often cited as the main reason to adopt a multiagent approach for a certain problem. This also holds true for multiagent planning. Still, papers on multiagent planning hardly ever make explicit in what ways their systems protect their users' privacy, nor do they give a quantitative analysis. The reason for this is that a theory of privacy loss in multiagent planning is virtually non-existent so far.

This paper proposes a measure for privacy loss based on Shannon's theory of information. To illustrate our approach, we apply this metric to an existing multiagent planning system to assess its merits when it comes to privacy on two domains. For this, we compare its plans to centrally generated solutions (by a trusted third party) for the same problems. The results clearly establish the need for such an analysis: even though the multiagent planner seemingly exchanges little information, its overall performance with respect to privacy is less than that of the centralised system (not taking into account the privacy loss with respect to the central planner, of course).

1. Introduction

The literature names a great many reasons why to pursue multiagent planning. One of the reasons that often comes up is that of *privacy*. Especially in circumstances where the agents represent (possibly competing) companies, sharing data with other parties is considered undesirable. At the same time, it is well recognised that cooperation (which requires some degree of information sharing) may be mutually beneficial to all parties.

Multiagent planning is one of the tools that can help agents in these situations. It offers a way to cooperate while being in control of which information is shared and with whom. Yet it is impossible to cooperate without sharing *any* information. At the very least, a pair of cooperating agents has to agree on which subtasks are being carried out by one agent on behalf of the other [12]. Several approaches, such as (Generalised) Partial Global Planning (GPGP; see, for example, [2]), go even further and share detailed parts of their plans. In the latter approach, more information is exchanged. Clearly, this must lead to a poorer performance when it comes to privacy. This raises the obvious question of how to measure this performance. Just how much privacy is lost by exchanging certain information? How can we evaluate which method is better then another when privacy is concerned?

This paper introduces a measure for privacy loss in multiagent planning based on Shannon's Theory of Information [7]. We show how the concept of "uncertainty" that underpins Shannon's work can be interpreted in the context of plans, and how we can derive a measure from this for the information that is gained during negotiations on plan construction or coordination. We then apply this measure to an existing multiagent planner, called MPOPR [12], and show its privacy behaviour on two domains, compared to a centralised system using a trusted third party. In the latter case, agents extract information about the other agents' plans from the plan they are given from the central planner.

The remainder of the paper is organised as follows. In the next section we will introduce our measure for privacy loss. Then we introduce the multiagent planning system that we applied our work to and the two domains for which we did our analysis. After examining our findings, we discuss some related work from the field of distributed constraint satisfaction. Finally, we draw some initial conclusions and expand on future work.

2. Privacy loss in multiagent planning

If we want to speak about a loss in privacy, we have to obtain a measure to quantify this loss. Information theory, as established by Shannon [7], can give us this measure. It is based on the idea that the amount of information contained in a message exchange can be measured by the amount with

*Supported by an Irish Research Council for Science, Engineering and Technology (IRCSET) Postdoctoral Fellowship. This paper is an improved version of [11], relaxing one of the key assumptions.

which the uncertainty about certain facts decreases. We therefore establish a measure for uncertainty with respect to plans first.

2.1. Uncertainty with respect to plans

If we go back to the classical definition of planning, we consider a plan Δ to be a (totally-ordered) sequence of actions from a set $O = \{\overline{o}_1, \overline{o}_2, \ldots \overline{o}_m\}$ of possible actions that brings about a state change from the current state I to some goal state G. Hence, a plan is a sequence $o_1 \cdot o_2 \cdots o_n$, where $o_i \in O$ is the action executed at time point i. Actions in a plan are linked through their preconditions and effects. Therefore, the uncertainty regarding which action is executed at time i is dependent upon the actions at times $1 \ldots i-1$. In this paper, we make the simplifying assumption that the occurrence of each o_i is solely dependent upon the previous action o_{i-1}.[1] In this case, the uncertainty H about which action o_k is executed at time step k follows straightforward from the entropy of a first-order Markov source:

$$H = -\sum_{i=1}^{|O|} P_i \sum_{j=1}^{|O|} P_i(j) \log_2 P_i(j) \tag{1}$$

Here, P_i is the probability of action $\overline{o}_i$ being executed and $P_i(j)$ is the probability of action $\overline{o}_j$ being executed immediately following action $\overline{o}_i$. The equation above requires the probabilities of actions to be known. Unfortunately, we cannot establish the probability of a given action, or sequences of actions. From previous plans, or simulation, we can estimate the probability by the frequency F_i that an action $\overline{o}_i$ occurred in the past, as well as the frequency $F_i(j)$. Therefore, subsequent equations will have the probabilities P_i and $P_i(j)$ substituted with the frequencies F_i and $F_i(j)$.

Equation 1 gives us the uncertainty regarding a single action in the plan. A plan consists of a number of actions, however, for successive timesteps. Because of the additive property of the uncertainty measure, we can extend the equation to the uncertainty $H_{\Delta,T}$ of a a plan Δ over a certain timewindow $[1 \ldots T]$ as follows:

$$H_{\Delta,T} \quad = \quad -T \sum_{i=1}^{|O|} F_i \sum_{j=1}^{|O|} F_i(j) \log_2(F_i(j)) \tag{2}$$

Remark *Notice how we compose the uncertainty from the indiviual timesteps, rather than taking on the uncertainty of the plan as a whole. As we see below, this allows us to easily incorporate knowledge that we gain about particular time steps.*

[1] This assumption is too strong in many situations. However, the resulting measure gives a good indication of the uncertainty. Nevertheless, see the discussion on how the assumption can be relaxed.

2.2. Privacy loss

Now that we have defined the concept of uncertainty in the context of planning, we can define the privacy loss that agents incur by exchanging information with other agents while constructing a multiagent plan. Using the formulas above, we can express information as follows: Suppose we have an uncertainty H_{before} before an event (such as a negotiation session) and that uncertainty after the event is H_{after}. Then the information that was gained in the event equals

$$R = H_{before} - H_{after} \tag{3}$$

This is what Shannon calls the *rate of information transmission*. Thus, information always relates two points in time, and the uncertainties at those times. The information that one gains, can be conceived as the privacy that is lost by the other party. To be more precise, consider the case of two agents, entering into negotiations for some aspect of their planning problems. Before the negotiations, the uncertainty with regard to the other agent's plan is governed by Equation 2, i.e. nothing is known about the plan. During their negotiations, agents may learn about certain aspects of the other agent's plan. In particular, they may receive information about certain actions (not) being executed at a certain timepoint. Let $O_t \subseteq O$ be the set of actions that are known to be possible at timestep $t > 1$, let $O_0 = O$, and let $F_i|_{O_t}$ be the normalisation of F_i to a set of possible actions O_t and similarly let $F_i(j)|_{O_{t_i},O_{t_j}}$ be the normalisation of $F_i(j)$ to the sets O_{t_i} and O_{t_j}. Then the uncertainty *after* the negotiations equals

$$H_{\Delta,T} = -\sum_{t=1}^{T} \sum_{i=1}^{|O_{t-1}|} \left[F_i|_{O_{t-1}} \times \sum_{j=1}^{|O_t|} F_i(j)|_{O_{t-1},O_t} \log_2(F_i(j)|_{O_{t-1},O_t}) \right] \tag{4}$$

Notice how this reduces to Equation 2 when $O_i = O$ for all $0 \leq i \leq T$.

A few issues arise, however, that prevent us from simply combining Equation 3 with Equations 2 and 4. Firstly, the uncertainty with respect to a plan $H_{\Delta,T}$ depends on the length T of that plan. However, in general agents do not know the length of the other agent's plan. This gives rise to a number of different measures of privacy loss for different choices of estimating the length of the plan, as we see below. Secondly, because of synchronisation issues, agents may have to wait in their plan, not executing any actions at all. We can resolve this by introducing an additional $idle$ action to the set O of allowed actions. This action has neither preconditions nor effects and is used to pad agents' plans, so that all plans of length T have T actions. We define $O^+ = O \cup \{idle\}$ to denote the expanded set of actions.

As indicated, different estimates of the plan length result in different measures of privacy loss. The basic equation from which these different measures derive is the following:

$$R(T) = -T \sum_{i=1}^{|O^+|} F_i \sum_{j=1}^{|O^+|} F_i(j) \log_2(F_i(j)) \quad (5)$$
$$+ \sum_{t=1}^{T} \sum_{i=1}^{|O^+_{t-1}|} \Bigg[F_i|_{O_{t-1}} \times$$
$$\sum_{j=1}^{|O^+_t|} F_i(j)|_{O_{t-1},O_t} \log_2(F_i(j)|_{O_{t-1},O_t}) \Bigg]$$

This equation is derived from Equation 3, substituting Equation 2 for H_{before} and Equation 4 for H_{after} and taking into account the $idle$ action.

To define the different measures of privacy loss, consider a group of agents $A = \{a_1, a_2, \ldots, a_N\}$ with respective plans Δ_{a_i} for $1 \leq i \leq N$. If we assume that the agents do not share their knowledge about other agents' plans, a loss of privacy always occurs between a pair of agents. Let the *target agent* τ be the owner of the plan Δ_τ under consideration, and the *invading agent* ι, $\iota \neq \tau$, be the agent that has gained some knowledge about Δ_τ. We can now distinguish the following measures of privacy loss:

internal The *internal privacy loss* can only be measured by the target agent. It takes into account the true length of Δ_τ: $R_{\text{internal}} = R(|\Delta_\tau|)$.

maximum The *maximum privacy loss* assumes that Δ_τ ends at the final timestep about which some information was gained : $R_{\text{maximum}} = R(\max\{i \,|\, O_i \neq O^+\})$.

estimated The *estimated privacy loss* assumes that $|\Delta_\tau| = |\Delta_\iota|$, unless information was gained about a later timestep: $R_{\text{estimated}} = R(\max\{|\Delta_\iota|\} \cup \{i \,|\, O_i \neq O^+\})$.

theoretical The *theoretical privacy loss* assumes that the plan is as long as that of the longest plan in the group of agents: $R_{\text{theoretical}} = R(max\{\Delta_a \,|\, a \in A\})$. It can only be computed by an omniscient entity.

The difference between these measures lies in the estimation of the length of the target plan. Obviously, this is an important aspect, as it has a profound impact. However, the invading agent has no way to determine this exact length, and has to make an estimate. The *internal* privacy loss is, in some sense, the *true* privacy loss, as it takes into account the actual length of the plan Δ_τ. The target agent can use this measure to evaluate its own plan with respect to privacy. At the other end of the scale, we have the *theoretical* loss. This is the loss over the entire horizon of the planning episode. It cannot be computed by any of the agents (as none of them know the exact lengths of the other agents' plans), but it can be used by developers to assess the privacy implications of their systems. The *maximum* and *estimated* privacy losses can be used by invading agents to estimate what information they have gained. The maximum loss takes an optimistic view, and assumes that the last timestep of which information was gained is also the last step of the target plan. The estimated loss takes a more moderate view, and uses the length of the invading agent's plan instead, if it is longer. Hence, it is an estimate of the theoretical loss, and is a more accurate measure of privacy loss over the whole episode than the maximum loss. In our next section, we show some examples of the different measures.

3. Illustration

To illustrate the metric that we propose, we apply it to evaluate the privacy aspect of a multiagent planner. For this analysis, we have used the results for the benchmark set of the MPOPR planning system [10, 12].

3.1. The system and the data sets

The idea behind the MPOPR system is to combine a dynamic planning method for each agent with an auction for delegating (sub)tasks. The system consists of a number of agents that first concurrently plan for a single goal. As some goals may involve subgoals that the agent cannot achieve itself, each agent comes equipped with some high-level information about the services of others. They can use this information to reason about which subgoals they should auction. After the first planning phase, the agents take part in an auction (if there is any) to exchange some of these unattainable subgoals. Then, they apply a plan repair technique to add another goal to their plan, and take part in the next auction. They continue to alternatingly perform these steps of adapting a plan using plan repair and taking part in an auction until a complete and valid plan is computed. When an agent gets a task assigned on which others depend, a heuristic is employed that lets the agent schedule it early in its plan to prevent cyclic dependencies.

Notice that the information exchanged is very little: the auctioneer only specifies a fact that it wants to see achieved, the bidders issue a single value as their bid, and the winner communicates when it starts working on the goal and and at which point it is finished. No information is exchanged about actions that are being undertaken, as is for example done in the partial global planning approach [2]. This may lead one to assume that the system has favourable privacy properties. Indeed, this is what the authors suggest [12]:

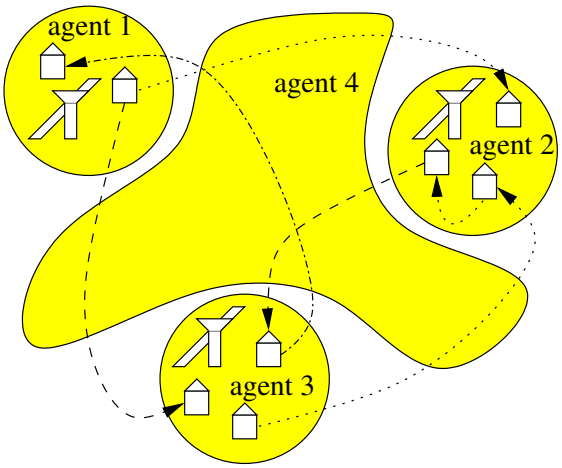

Figure 1. The multiagent logistics domain.

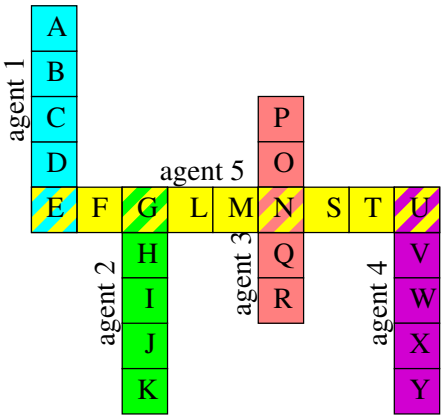

Figure 2. Graphical representation of the robots domain. Depicted are the 25 rooms, and the areas in which each agent may travel.

"*However, it allows us to create valid multiagent plans without exchanging details about the plans, which is very important for self-interested agents*."

The system is evaluated on two sets of benchmark problems. In [12], the authors present a multiagent variant on the well-known logistics domain. As can be seen in Figure 1, the standard problem is divided over a number of agents: one for each city, responsible for the intra-city transport of goods, and an additional agent that handles the inter-city transport. For our analysis, we have used the B benchmark set, which consists of 45 plan repair problems.[2]

In [10], an additional domain is introduced called robots, cf. Figure 2. This is a multiagent variant on the gripper domain, in which a moving robot has to deliver objects. In the multiagent variant, each agent is confined to a single corridor and the agents have to coordinate to bring objects between corridors. This benchmark set consists of a number of random problems generated for different problem sizes (measured by the total number of goals). For our analysis, we have averaged the results over 3 problems of each size.

3.2. Evaluation

We evaluate the privacy aspect of the MPOPR system by comparing it to a centralised generation of plans by the VHPOP [15] system.[3] To do so, we compute the information that an agent can gather from the final plan that is computed. This information is present in the form of inter-agent dependencies. Such dependencies exist when one agent achieves a subgoal for another agent. Moreover, information can be gained by relating this information. For example, in the logistics domain, if an agent is informed that particular subgoal of moving a package within a city is started at time i and finishes at time $i+2$, we can not only infer that a *load* action takes place at time i and an *unload* takes place at time $i+2$, but due to the workings of the domain, we can also deduct that a *move* action is executed at step $i+1$. In this case, the uncertainty for three timesteps of the plan is removed, which is the information we gained. For one agent, this does not hold, however: the inter-city transport agent in the logistics domain controls two airplanes. Hence, for every timestep, there is uncertainty about *two* actions. If we look at the dependencies, we see that still a large degree of uncertainty remains. For example, if we can infer a *load* action for a certain timestep, we do not know which of the two airplanes is involved. Thus, the uncertainty in this case is only slightly reduced.

Remark *Notice, that even in the case when no deductions can be made, it follows from Equation 4 that if we gain information on a step t, the uncertainty regarding step $t+1$ decreases too. This improvement in accuracy of the measure is the main difference with our earlier work [11].*

Logistics. Notice that in the logistics domain, interactions only occur between the inter-city agents and the intra-city agent. We therefore look at the average privacy loss that the inter-city agent incurs to each of the other agents, and *vice versa*. Figure 3 shows the *theoretical* privacy loss that the intra-city agents incur with respect to the intra-city agent for the 45 different problems in this set. As we can see, the privacy loss is worse for the MPOPR system compared to the centralised result. The main reason for this is that the centralised approach combines more orders. The anti-cycle heuristic employed by MPOPR has the effect that its plans usually contain consecutive *load*, *drive*, and *unload*

[2] For all problems, we computed solutions from scratch, without taking into consideration a previously existing plan.

[3] The reason for choosing VHPOP is that MPOPR is ultimately derived from this planner.

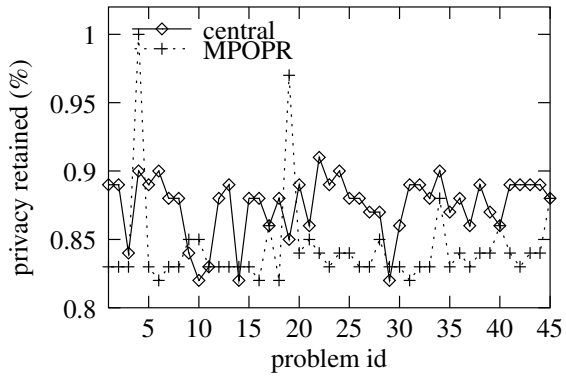

Figure 3. Average (theoretical) privacy loss of the intra-city (trucking) agents

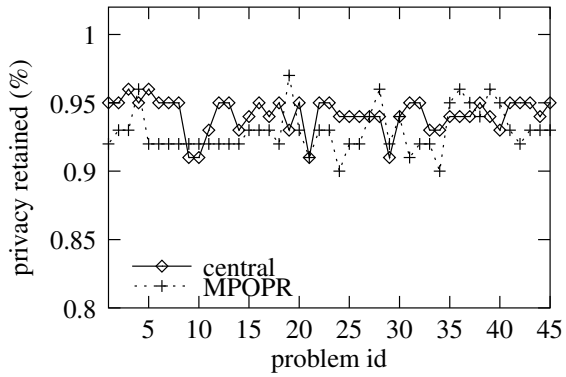

Figure 4. Average (theoretical) privacy loss of the inter-city (flying) agent

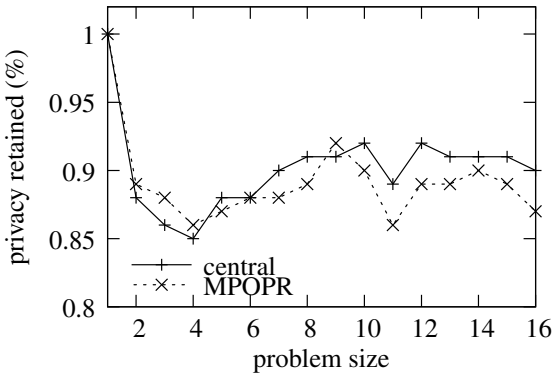

Figure 5. Average internal privacy loss agents 1-4

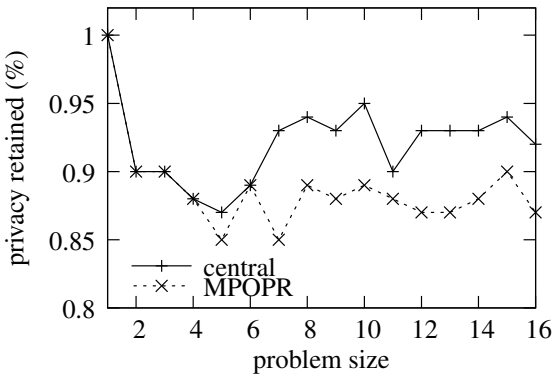

Figure 6. Average estimated privacy loss agents 1-4

actions for tasks that it agrees to do for another agent. As discussed before, this leads to a decrease in uncertainty for three timesteps. Compare this to the central approach, which often has (also more efficient) sequences as $load$, $load$,$drive$, $unload$, $unload$. Depending on the order of the loads and unloads, the $drive$ action may or may not be deduced, which on average leads to a lower loss in privacy.

Figure 4 shows a quite different picture for the inter-city agent. As we noted earlier in this section, the information that agents can gain from interacting with this agent is lower, because of the two airplanes this agent can use. As a result, the loss in privacy is very little for both methods. Here, too, MPOPR fares worse than the centralised method on most problems.

Robots. In the robots domain, we also have a single agent that interacts with all other agents, as the corridors all intersect with the central corridor, and with no others. Hence, we look at the privacy that the agents lose from interacting with this central agent. Here, we shall also look into the internal privacy loss and the estimated loss.

The average loss that the agents incur from interacting with the agent in the central corridor is shown in Figure 5 (internal privacy loss), Figure 6 (estimated loss) and Figure 7 (theoretical loss). Again, we observe that MPOPR often gives a greater loss in privacy than the centralised method. Interestingly, if we look at the average theoretical privacy loss (and to a lesser extent the estimated loss), this seems to remain relatively stable over different problem sizes. We conjecture that this is due to the fact that while the plans do get bigger (as there are more goals to achieve), there are also more opportunities for learning about the other agent's actions. Together these two effects counter each other, leading to the stable behaviour we see. This behaviour is not noticeable in the internal privacy loss. Often one of the agents only has one or two packages to handle. If this involves a coordination with the central agent, almost all of the plan is exposed. However, since the central agent does not know the lengths of the plans, the estimate does not reflect this. The observer simply doesn't know that it has learned so

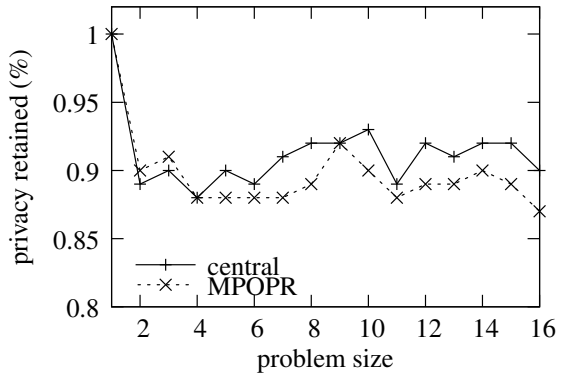

Figure 7. Average theoretical privacy loss agents 1-4

much information.

3.3. Preliminary conclusions

While the aim of this section is mainly to illustrate the metrics that we propose, we can draw some initial conclusions. First of all, it underlines the necessity for a privacy metric. A system that was designed to leak little information, turns out to be actually worse when it comes to privacy than a system in which a plan is computed centrally and then distributed. Even though we singled out the MPOPR planning system, other multiagent systems have neither undertaken a privacy analysis. Rather, they rely on enforcing certain properties that are thought to provide privacy for their users. Unfortunately, good intentions do not make for a secure system, as we have seen from this analysis.

A recommendation that can be drawn from these experiments for a privacy aware planner would be to allow for a certain degree of slack and randomisation. Another option would be the interleaving of actions. The MPOPR system enforces neither of those aspects, whereas the central planning has many tasks interleaved, which improves privacy. However, in light of the previous paragraph, we cannot simply assume that adhering to these properties would make for a secure system. A privacy analysis remains indispensable.

4. Related work

Recently, researchers in the field of Distributed Constraint Optimisation (DCOP) and Distributed Constraint Satisfaction (DisCSPs) have started to propose metrics for analysis of privacy loss in such systems. Most of the work in this area focuses on distributed meeting scheduling. In this type of problems, a number of agents has to schedule a number of meetings. Each meeting requires a certain set of agents to be present, and each agent has preferences or costs attached to timeslots and locations. Silaghi and Faltings [8] use a measure of privacy to drive their algorithm. Each agent has certain costs associated with the revelation of whether some tuple of values is feasible. During the exchange of messages, agents have to pay this cost if some tuple is fully determined by the other agents. Negotiations are terminated if the cost of revealing a certain tuple is greater than the potential reward for collaborating. A similar model is the basis of Silaghi and Mitra [9]. However, the privacy metric here is the size of the smallest coalition necessary to deduce an agent's costs for certain tuples. Wallace and Freuder [14] consider a measure of privacy loss that is very close to ours. Their work, like ours, is based on information entropy. However, the application of information theory is more straight-forward as they consider the uncertainty of each of the variables in the constraint satisfaction problem, rather than having to apply it to an additional datastructure (i.e. the plan) as we do. Recent work by Maheswaran *et al.* [4] proposes a general quantitative framework to analyse privacy loss. The three earlier approaches can be seen as specific instances of this framework.

Beyond DCOP and DisCSP, research on privacy is undertaken in the agent community at large. This includes work on cryptographic techniques, secure auctions and randomisation (see e.g. work by Brandt [1], Van Otterloo [13] and Naor [5]). Of particular interest to planning is the work on randomisation (e.g. Paruchi *et al*. [6] and Van Otterloo [13]). These approaches assume that actions and behaviours can be observed. By choosing actions in a randomised fashion (e.g. using policies with a high entropy) agents can try to provide minimal information on their preferences, while still attaining their goals.

5. Discussion and extensions

Although privacy is an issue that is often mentioned in work on multiagent planning (as well as multiagent approaches to different problems), heretofore this notion was neither made explicit, nor analysed. The present work shows how Shannon's Information Theory can be applied to (classical) planning to derive meaningful definitions of concepts such as uncertainty, information and privacy loss. As an illustration of our work, we applied it to an existing multiagent planning system to evaluate its performance with regard to privacy, compared to a central solution with a trusted third party. This clearly established the need for such an analysis: even though the multiagent planner seemingly exchanged little information, its overall performance with respect to privacy was less than that of the centralised system.

A number of extensions to this work seem obvious. Firstly, it seems reasonable to consider more than just the type of action. Our model is limited in the sense that we do

not distinguish between knowing that a $move$ action takes place, and knowing the precise locations. A step in this direction was made in the *logistics* domains, where we allowed uncertainty to exist with regard to which plane was used by the intercity transport agent. However, there may be domains where such a distinction is vital and requires a more rigorous assessment then our current model allows for.

Secondly, our work is based on the classical definition of plans. Over the years, more advanced definitions have emerged, allowing for parallel execution of actions and durative actions. Obviously, to assess the privacy issues of modern planners built on these richer formalisms, the definitions of uncertainty, information and privacy loss will have to be extended. It is not immediately clear how our work can be generalised to these situations. One direction we are pursuing in this regard is based on the partial plan representation [3], but this still has a number of difficulties.

Thirdly, it would be interesting to investigate the impact of coalition formation. A group of agents that joins the knowledge they have gained separately, could conceivably achieve an amount of information that is bigger than the sum of the individual pieces of knowledge.

Most importantly, however, this work provides the means to evaluate multiagent planning systems on their privacy impact and allows new multiagent planning systems to be designed. Whereas before, privacy was mentioned as a driving force but not explicitly taken into account, the existence of a metric for privacy loss can direct research to new algorithms that are optimised for privacy. This would also entail research into the relation between privacy loss and other factors such as optimality and search efficiency.

References

[1] F. Brandt. Fully private auctions in a constant number of rounds. In *Proc. of the* 7^{th} *Annual Conf. on Financial Cryptography (FC-03)*, pages 223–238, 2003.

[2] K. S. Decker and J. Li. Coordinating mutually exclusive resources using GPGP. *Autonomous Agents and Multi-Agent Systems*, 3(2):113–157, 2000.

[3] S. Kambhampati. Refinement planning as a unifying framework for plan synthesis. *AI Magazine*, 18(2):67–97, 1997.

[4] R. T. Maheswaran, J. P. Pearce, E. Bowring, P. Varakantham, and M. Tambe. Privacy loss in distributed constraint reasoning: A quantitative framework for analysis and its applications. *Autonomous Agents and Multi-Agent Systems*, 13(1):27–60, 2006.

[5] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the first ACM conference on Electronic Commerce*, pages 129–139, 1999.

[6] P. Paruchi, M. Tambe, D. Dine, S. Kraus, and F. Ordonez. Safety in multiagent systems via policy randomization. In *AAMAS workshop on Safety and Security in Multiagent Systems*, 2005.

[7] C. Shannon. A mathematical theory of communication. *Bell System Tech. J.*, 27:379–423,623–656, 1948.

[8] M. Silaghi and B. Faltings. A comparison of distributed constraint satisfaction approaches with respect to privacy. In *Proc. of the* 3^{rd} *workshop on Distributed Constraints Reasoning*, pages 147–155, 2002.

[9] M. Silaghi and D. Mitra. Distributed constraint satisfaction and optimization with privacy enforcement. In *Proc. of the 2004 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT-04)*, pages 531–535, 2004.

[10] R. van der Krogt. *Plan Repair in Single-Agent and Multi-Agent Systems*. PhD thesis, Delft University of Technology, Delft, The Netherlands, 2005.

[11] R. van der Krogt. Privacy in multiagent planning: A classical definition with illustration. In *Proc. AAMAS '07 Workshop on Coordinating Agents' Plans and Schedules*, 2007.

[12] R. van der Krogt and M. de Weerdt. Coordination through plan repair. In *MICAI 2005: Advances in Artificial Intelligence*, pages 264–274, 2005.

[13] S. van Otterloo. The value of privacy: optimal strategies for privacy minded agents. In *Proc. of the Fourth Int. Conf. on Autonomous Agents and Multi- Agent Systems (AAMAS-05)*, pages 1015–1022, 2005.

[14] R. J. Wallace and E. C. Freuder. Constraint-based reasoning and privacy/efficiency tradeoffs in multi-agent problem solving. *AI*, 161:209–227, 2005.

[15] H. L. S. Younes and R. G. Simmons. VHPOP: Versatile heuristic partial order planner. *Journal of AI Research*, 20:405–430, 2003.