Flexible social laws*

Will Briggs and Diane Cook
Department of Computer Science and Engineering
University of Texas at Arlington
Arlington, TX 76019 U. S. A.
wbriggs@sirius.uta.edu, ocok@centauri.uta.edu

Abstract

Although communication is generally considered to dominate over processing cost in distributed systems, the problem of communication cost in multiagent planning has not been sufficiently addressed. One method for reducing both communication cost and planning time is the use of social laws. Social laws, however, can be too restrictive, limiting soundness. Flexible social laws can enable multiagent systems to reap the benefits of reduced communication cost and planning time (except in the worst case), without limiting soundness (although optimality may be degraded). By analyzing the performance, we show that this model can make multiagent planning exponentially more efficient without limiting its applicability.

1 Introduction

As automation of intelligent tasks increases, the need will arise for heterogeneous agents working in a common environment. As a result, agents will need to plan and coordinate plans in real time.

[Cook, 1994] has described three methods of handling multiple agents: central control, distributed control, and local control (no communication). She notes that central control is effective when communication is reliable, and local control is effective if no communication is needed; otherwise, distributed control is necessary. This is the protocol generally used with multiagent systems [Durfee and Montgomery, 1991; von Martial, 1992; Gmytrasiewicz and Durfee, 1994].

Unfortunately, planning is intractable, especially if plans must be coordinated: the amount of communication (worst-case) grows as the square of the search space, which is itself exponential for deliberative planning ([Briggs and Cook, 1995a]). Multiagent planning work so far has mostly provided mechanisms for interaction [Georgeff, 1986; von Martial, 1992] without attention to this problem, although [Gmytrasiewicz and Durfee, 1994] shows a model in which agents use decision theory to decide when to communicate, and [Durfee and Montgomery, 1991] reduces the cost of communication by sending abstract messages. Previous work on *social laws* [Moses and Tennenholtz, 1990; 199l], that is, constraints on what actions an agent may take, has shown a method for reducing communication and planning time by reducing the options an agent has at any point. (In this paper, we relax the definition of social law to allow each agent to have a distinct set of constraints.)

*This research was supported in part by the National Science Foundation, under grant IRI-9308308.

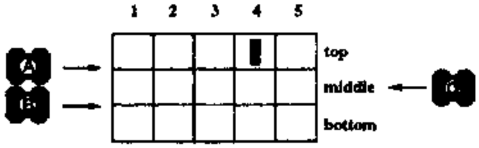

Figure 1: Robot cars negotiate over use of a 3-lane road. There is a roadblock in the westbound (top) lane.

Our work extends the social law paradigm, providing a specific analysis of the costs, and a way of controlling these costs without reducing the applicability of planning.

1.1 Social laws

Social laws provide a way to constrain the actions of agents in multiagent systems, so as to reduce both the branching factor within an agent's search space and the chance of interaction with another agent's plan.

For example, consider a domain in which mobile robots must share a three-lane highway (Figure 1). One possible social law is a common traffic regulation: "Drive on the right." This law will prune a portion of the search space, since agents will not need to consider moving into other lanes. It will also reduce possible contention, since agents moving in one direction will never share a lane with those moving the opposite direction.

Reducing possible contention over states of resources is crucial, since the detection and resolution of such interactions is the dominant communication cost in multiagent planning [Briggs and Cook, 1995a]. Specifically, agents need to communicate about the states of resources that their operators require (in a precondition list) or can provide (in a postcondition list). In Figure 1, for example, each mobile robot needs to know that a section of the highway is clear before it can safely enter.

Unfortunately, social laws can prevent us from finding a solution. Figure 1 shows a roadblock in the topmost lane. If robot C is required to drive on the right, it cannot move into columns 1-3.

Social laws, therefore, can limit soundness. Determining if this is a possibility for any particular social law is in general NP-complete [Shoham and Tennenholtz, 1992].

1.2 Flexible social laws

Previous work by the authors [Briggs and Cook, 1995b] has analyzed what may be viewed as a special case of social laws, that is, restrictions on access to resources. This work extends the analysis to a more general concept *of flexible social laws.*

In this framework, agents prefer to obey the "laws" but are able to relax them. Assume a ranking of sets of social laws, from strictest to most lenient, where a more lenient set is one that allows a greater choice of operators. Also assume a limit on the depth of search. Each agent will try to generate a plan within the strictest set of laws. If this fails, the agent will try again, using the next set in the ranking. This continues until the agent has a successful plan or has failed with the most lenient set of laws.

Flexible social laws can reduce both planning time and communication. Soundness is not sacrificed, because if no solution is found within limits set, the agent plans with more flexible laws.

Of course, when an agent finds a solution while using a social law, the solution may not be optimal. But (especially if the plan is to be used only once) the cost of plan generation may be high enough that a solution quickly found is better than an optimal one.

If an agent has to try two or more sets of social laws before finding a solution, some redundant planning is performed. We will find that this extra cost is not great, and (for deliberative planning at least) makes the total cost greater than the cost of planning without social laws only in the extreme worst case.

This model should apply to groups of homogeneous or heterogeneous agents, which share subgoals or work on their own goals exclusively. The model should also apply with varying degrees of reactivity, and to anytime planning.

2 An example

Mobile robots A, B, and C are to share a highway (see Figure 1, previous page). A and B want to go right; C is to go left. There is a roadblock in the top lane.

Within one timeslice, a robot can move forward or back, straight ahead or diagonally; or it can wait. Each lane is 5 units long, and only one car can safely occupy the same unit in any timeslice.

Here are some possible planning scenarios, using different social law schemes. In the first scenario, we will have no social laws. For the second, we attempt to reduce communication and planning costs with a strict social law. Since this will not allow all agents to complete their plans, the third uses a more lenient social law, which shows no advantage over no social law at all. Finally, we use a ranking of social laws to reduce cost while still allowing all agents to find a solution. The plans shown are fully coordinated.

No social laws

	A's plan	B's plan	C's plan
1.	To (top, 1)	To (middle, 1)	To (top, 5)
2.	To (top, 2)	To (middle, 2)	To (middle, 4)
3.	To (top, 3)	To (middle, 3)	To (top, 3)
4.	To (middle, 4)	To (bottom, 4)	To (top, 2)
5.	To (middle, 5)	To (bottom, 5)	To (top, 1)

We have found an optimal plan, but with some communication cost. In each timeslice, A and B have a possibility of contention, since each could swerve into the other's lane; in fact, in timeslice 4, A has to, to avoid the roadblock. C has possible contention with both A and B at timeslice 3.

Now we will try planning with a strict social law:

Social law: "Drive on the right"

	A's plan	B's plan	C's plan
1.	To (bottom, 1)	Wait	(no solution)
2.	To (bottom, 2)	To (bottom, 1)	
3.	To (bottom, 3)	To (bottom, 2)	
4.	To (bottom, 4)	To (bottom, 3)	
5.	To (bottom, 5)	To (bottom, 4)	
6.		To (bottom, 5)	

In this non-optimal plan, we have communication only between A and B, which must share a lane; but C cannot complete its task.

A more lenient social law might help:

Social law: "Drive on the right *or* the center"

	A's plan	B's plan	C's plan
1.	To (middle, 1)	To (bottom,1)	To (top,5)
2.	To (middle, 2)	To (bottom, 2)	To (middle,4)
3.	To (middle, 3)	To (bottom, 3)	To (top, 3)
4.	To (middle, 4)	To (bottom, 4)	To (top, 2)
5.	To (middle, 5)	To (bottom, 5)	To (top, 1)

We have a solution for each agent again; but we have not reduced the need for communication from the first scenario, which had no social laws.

A flexible application of social laws could get us the benefit of reduced communication without sacrificing soundness. Here, agents try planning with social law 1, and on failure, try again with social law 2.

Social law 1: "Drive on the right";
Social law 2: "Drive on the right *or* the center"

	A's plan	B's plan	C's plan
1.	To (bottom, 1)	Wait	To (top,5)
2.	To (bottom, 2)	To (bottom, 1)	To (middle,4)
3.	To (bottom, 3)	To (bottom, 2)	To (top, 3)
4.	To (bottom, 4)	To (bottom, 3)	To (top, 2)
5.	To (bottom, 5)	To (bottom, 4)	To (top, 1)
6.		To (bottom, 5)	

In this case, as in the case with the strictest social law, we sacrificed optimality, and reduced interactions to the possible contention of A and B for the bottom lane; however, by relaxing the law, C was able to find a solution.

3 The model

This section contains definitions and assumptions.

For our analysis, we will assume cooperative agents, that is, agents which are honest and are willing to change their plans to resolve harmful interactions. We also assume that agents plan correctly, that is, that they communicate whenever there exists a possibility of interactions between plans.

3.1 Communication needs

During plan generation, these two types of communication occur between agents: *interaction resolution*,[1] which is the most costly (demonstrated in [Briggs and Cook, 1995a]), and *fact sharing*, in which one agent gives another information about the state of some resource, possibly from sensor data. If agents plan for each other, they must also send *operator sets* so they will know each other's capabilities, as well as information about goals to achieve and returned information about completed plans, which can be collectively called *load balancing* information. Since interaction resolution tends to dominate the other types of communication, we will consider it only.

3.2 Definitions

The *leniency* of a set of social laws is defined as $\sum_{i \in states}[|operators\ allowed\ in\ state\ i| / |operators|] / |states|$; this is the average proportion of operators allowed in any state. For convenience, we define *strictness* as the opposite of leniency, or strictness = 1− leniency.

Let there be $|A|$ agents. Define for each agent A_i a ranking of sets of social laws, $LawSets_i$, monotonically increasing in leniency, ranging from $LawSets_{i,1}$ (the least lenient) to $LawSets_{1,|LawSets_i|} = \{\}$ (by definition, the most lenient). We can ensure that for $j < k$, $LawSets_{i,j}$ is no more lenient than $LawSets_{i,k}$ by requiring that $LawSets_{i,k} \subseteq LawSets_{i,j}$. Each agent may have its own set of laws; this could be helpful in heterogeneous environments. For this reason we define $max(|LawSets|)$ to be the maximum number of sets of laws for any agent.

At each point in the agent's search space, adhering to a social law means restricting the search to a subset of all possible operator instances. Using flexible social laws requires the establishment of a ranking of such sets. Call this ranking $OpSets_i$. Note that if $j \leq k$, $OpSets_{i,j} \subseteq OpSets_{i,k}$.

Let $P_{i,j}$ be the set of all predicates that appear as pre- or post-conditions of the operators in $OpSets_{i,j}$; if there are any derived effects, they are included as well. $P_{i,j}$ is therefore the set of all predicates which agent A_i may reference or affect while using operators from $OpSets_{i,j}$. There is therefore for each ranking $OpSets_i$ a corresponding ranking P_i. Note that if $j \leq k$, $P_{i,j} \subseteq P_{i,k}$.

Define θ_i to be a number in the range $1..|LawSets_i|$, such that $LawSets_{i,\theta}$ is the set of social laws agent A_i is currently using. (The subscript will be omitted when it is clear from context.)

With each new value of θ, an agent may need to tell all the others what predicates it may be referencing (that is, the predicates in $P_{i,\theta}$), so as to know who to negotiate with. For deliberative planning, the cost will prove to be insignificant. For reactive planning, we will assume that each agent has the same *LawSets*, such that there is no need to inform each other of what predicates one may be referencing. (Further work is needed to examine the effects of relaxing this assumption.) In either case, we may omit the cost of this information sharing in our analyses.

Here is the algorithm for multiagent planning with flexible social laws.

Algorithm Plan-with-flexible-social-laws
mark all agents "active"
for $\theta = 1$ to $max(|LawSets|)$ do
    for each "active" agent A_i do, in parallel
        tell all other agents what
            predicates A_i may access
        plan, using only operators in $OpSets_{i,\theta}$
        if a solution is found
            mark agent A_i "successful"
        if $\theta = |LawSets_i|$
            mark agent A_i "failed"
    endfor
endfor

We define $b(\theta)$ to be the branching factor within the search space of any agent, given θ. For simplicity we assume that $\forall\theta$, $b(\theta)$ is the same for all agents, and is constant for all choice points in an agent's search space. Further, we assume that $b(\theta) = \theta b_0$ for some constant b_0; that is, every increase in b, as we increment any θ to $\theta + 1$, is close to the average increase.

We assume a lookahead limit L; we will not try to generate plans longer than L. We also assume worst-case search space size $Size(A_i, \theta) = b(\theta)^L = {b_0}^L \theta^L$ for all θ. (Search space is measured in terms of operator instances considered.) We define $Size(\theta) \equiv \sum_{i=1}^{|A|} Size(A_i, \theta)$, where $|A|$ is the number of agents. We also assume that we lack sufficient memory to store these exponential search spaces.

We also define cost function $Cost_{comm}(A_i, \theta)$ to be the number of messages sent for interaction resolution during the θ^{th} iteration of the algorithm. $Cost_{comm}(\theta) \equiv \sum_{i=1}^{|A|} Cost_{comm}(A_i, \theta)$. We assume worst-case cost[2] for any θ.

4 Theoretical analysis

This section contains an analysis of optimality; savings in communication and planning costs, accounting for the penalty from redundant planning work; and best- and worst-case analyses.

[1]Interactions may be harmful (conflicts) or helpful (favors).

[2]This means that each node in one agent's search space has a possible conflict with each node in any other agent's. This simplifying assumption will increase our results for communication cost, both with and without social laws; further work is needed to determine how this affects our savings.

4.1 Optimality

Branching factor $b(\theta)$ is necessarily monotonically increasing with θ, but the length (or cost, if we use operator cost) of each agent's final plan (if any is found) is necessarily monotonically decreasing with θ, since a more lenient social law can't possibly make a plan longer.[3] (This is true for any domain.)

Therefore a solution for some $\theta < max(|LawSets|)$ may not be optimal.

4.2 Communication

Here we calculate $Cost_{comm}(A_i, \theta)$, the maximum number of messages used for interaction resolution between any two agents, at any iteration in the algorithm. This is proportional to the number of interacting pairs of operators belonging to separate agent's search spaces.

The maximum number of interacting pairs for any θ is the number of nodes in one agent's search space times the number of nodes in the other's search space. If both agents are still searching (that is, neither found a solution at a previous level), then this is $(b(\theta)^L)^2$. For all pairs of all $|A|$ agents, there are $|A|(|A|-1)b(\theta)^{2L}$ interacting pairs. If at most C messages are required to resolve interactions between such a pair, there will be $C(|A|-1)b(\theta)^{2L}$ such messages per agent.

By our assumptions, if we are performing deliberative planning, at each iteration of the algorithm an agent will need to tell all the others what predicates it may be referencing (that is, the predicates in $P_{i,\theta}$), so as to know who to negotiate with. The cost per agent A_i is $(|A|-1)|P_{i,\theta}|$. This is insignificant compared to the exponential cost $C(|A|-1)b(\theta)^{2L}$ and can thus be ignored.

If planning with $b(\theta)$ fails, planning with $b(\theta+1)$ must either search the same space again, or store the space generated for θ. It is not enough to store the leaves of the search tree; a final solution may interleave operators in $OpSets_{i,\theta}$ with operators in $OpSets_{i,\theta+1} - OpSets_{i,\theta}$.

We have assumed that there was insufficient memory to store these search spaces; instead we will search again each time afresh, forgetting what we generated for previous values of θ. With this assumption, we will need to communicate old information for each new θ, increasing $Cost_{comm}(A_i,\theta)$ to $C(|A|-1)\sum_{\psi=1}^{\theta} b(\psi)^{2L}$. The total savings for communication per agent, that is, the cost without any social law minus the cost using flexible social laws, is

$$Savings_{comm}(A_i,\theta)$$

$$= C(|A|-1)b(|LawSets_i|)^{2L} - C(|A|-1)\sum_{\psi=1}^{\theta} b(\psi)^{2L}$$

$$= C(|A|-1)(b_0|LawSets_i|)^{2L} - C(|A|-1)\sum_{\psi=1}^{\theta}(b_0\psi)^{2L}$$

$$= C(|A|-1)b_0^{2L}[|LawSets_i|^{2L} - \sum_{\psi=1}^{\theta}\psi^{2L}]$$

$$\approx C(|A|-1)b_0^{2L}[|LawSets_i|^{2L} - \frac{\theta^{2L+1}}{2L+1}]$$

This is positive when $\theta^{2L+1} < |LawSets|^{2L}(2L+1)$. For large L, this is when $\theta < |LawSets|$; that is, when we are able to use social laws at all.

If, when planning is done, all agents have the same θ — that is, if they all find their solutions in the same iteration — and if $\forall i, |LawSets_i| = |LawSets|$, then total savings for all $|A|$ agents is:

$$Savings_{comm}(\theta) \approx C|A|(|A|-1)b_0^{2L}[|LawSets|^{2L} - \frac{\theta^{2L+1}}{2L+1}]$$

In general, savings is polynomial in θ, but with a very large exponent $(2L)$ if planning is deliberative.

4.3 Planning time

The savings for planning time is also degraded by the redundant planning, so that

$$Savings_{planning}(A_i,\theta)$$

$$= Size(A_i, |LawSets_i|) - \sum_{\psi=1}^{\theta} Size(A_i,\psi)$$

$$= b(|LawSets_i|)^L - \sum_{\psi=1}^{\theta} b(\psi)^L$$

$$\approx b_0^L[|LawSets_i|^L - \frac{\theta^{L+1}}{L+1}]$$

This is positive, for large L, when $Savings_{comm}$ is positive — in all cases when a solution can be found using any social laws.

If all agents have the same θ, and $\forall i, |LawSets_i| = |LawSets|$, then:[4]

$$Savings_{planning}(\theta) \approx |A|b_0^L[|LawSets|^L - \frac{\theta^{L+1}}{L+1}]$$

[3] If we use breadth-first search. If we use some variation of depth-first search, reducing our options may lead us straight to the solution; or it may make finding a solution impossible

[4] If we also consider the savings in terms of constraints imposed by the agent upon its own operators, we will find $Savings_{constraints}(\theta) \approx C|A|b_0^{2L}[|LawSets|^{2L} - \frac{\theta^{2L}}{2L+1}]$; this is proportional to $Savings_{comm}(\theta)$.

4.4 Best- and worst-case bounds

In the best case, each agent finds a solution with its strictest set of laws. For large L, $Savings_{comm} = C(|A| - 1)b_0{}^{2L}[|LawSets_i|^{2L} - 1] \approx C(|A| - 1)b_0{}^{2L}|LawSets_i|^{2L}$, which is just the cost of communication without social laws. Similarly, $Savings_{planning}$ is approximately the cost of planning without social laws. Optimality of the resulting plan may be degraded.

In the worst case, for all A_i, $\theta_i = |LawSets_i|$. The communication penalty for trying flexible social laws (per agent) is $C(|A| - 1)b_0{}^{2L}\frac{|LawSets_i|^{2L+1}}{2L+1}$; that is, $\frac{|LawSets_i|}{2L+1}$ times the cost without social laws; so we have an increase in communication linear with the number of law sets, but inversely proportional to plan length. If $L \gg |LawSets_i|$, this penalty is not significant. By similar reasoning, we find a linear increase in planning time. Optimality is not degraded.

5 Reactive and anytime planning

As agents become more reactive, we are more likely to be planning real time, so optimality may be less important. Also, although the exponent L (which represents the lookahead limit) is reduced, we will see that flexible social laws can still reduce cost. Since anytime planning has a smaller depth limit than deliberative planning, this discussion also applies to anytime planners.

Let N = the length of the *total* plan for any agent. (It may be greater for a reactive planner than for deliberative planning, of course.) N/L is the number of plans an agent must generate and execute to accomplish the task. For a totally reactive planner, $L = 1$.

The savings for both communication and planning are N/L times the values given in the previous section; but the exponent L is now smaller, so the savings will be less. We note that deliberative planning gets a greater savings from the use of social laws than does reactive planning — but reactive planning, being closer to tractability to start with, will become tractable more quickly with decreasing θ.

For completely reactive planning ($L = 1$), we find that $Savings_{comm}(\theta)$ is positive when $\theta < 1.4|LawSets|^{2/3}$ (approximately). $Savings_{planning}(\theta)$ is positive when $\theta < 1.4\sqrt{|LawSets|}$ (approximately).

For large numbers of agents and law sets, then, if agents are completely reactive, problems must work out to fit our social law ranking very well (that is, to have solutions when agent use only the first $1.4|LawSets|^{2/3}$ of the social law sets). But for even moderately reactive agents (say, $L = 5$) the range is much wider ($\theta < 1.2|LawSets|^{10/11}$).

6 Conclusions

Flexible social laws are useful in multiagent planning for limiting both communication and planning time without affecting soundness. Despite redundant planning and communication, communication and planning-time savings are polynomial in the number of social law sets not considered by an agent, but with a large exponent if planning is deliberative. The tradeoff for this is a loss of optimality, and increase in communication and planning time in the worst case proportional to the number of social law sets in an agent's ranking, but inversely proportional to plan length. (For deliberative planning, this increase is only in the extreme worst case). We have shown best- and worst- case bounds, and have shown that the model is helpful for reactive and anytime planning as well as deliberative planning.

Our current work involves finding tighter theoretical limits, and testing the performance of this model. We are also developing a more restricted model which deals only with allocation of resources [Briggs and Cook, 1995b], in real domains. Another interesting area will be developing criteria for good rankings for the relaxation of social laws, like the downward refinement property [Bacchus and Yang, 1991] in hierarchical planning. *

7 Acknowledgements

We would like to thank Qiang Yang and Edmund Durfee for their advice and encouragement.

References

[Bacchus and Yang, 1991] F. Bacch us and Q. Yang. The downward refinement property. In *Proceedings of IJCAI-91.IJCAI,* 1991.

[Briggs and Cook, 1995a] Will Briggs and Diane Cook. Local planning and teamwork: minimizing communication in multiagent domains. Technical Report TR-CSE-95-001, University of Texas at Arlington, 1995.

[Briggs and Cook, 1995b] Will Briggs and Diane Cook. Scaleable modularity to reduce communication in multiagent planning. Technical Report TR-CSE-95-004, University of Texas at Arlington, 1995.

[Cook, 1994] Diane Cook. Reconfiguration of multiagent planning systems. In *Proceedings of AIPS-94,* pages 225-230. AIPS, 1994.

[Durfee and Montgomery, 1991] Edmund H. Durfee and Thomas A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics,* 21(6):1363 1378, Nov/Dec 1991.

[Georgeff, 1986] Michael P. Georgeff. Representation of events in multiagent domains. In *Proceedings of AAAI-86,* pages 70-75. AAAI, 1986.

[Gmytrasiewicz and Durfee, 1994] Piotr J. Gmytrasiewicz and Edmund H. Durfee. Rational coordination and communication in multiagent environments through recursive modeling. Submitted to *Journal of AI Research,* 1994.

[Moses and Tennenholtz, 1990] Yoram Moses and M. Tennenholtz. Artificial social systems part I: basic principles. Technical Report CS90-12, Weizmann Institute, 1990.

[Moses and Tennenholtz, 1991] Yoram Moses and M. Tennenholtz. On formal aspects of artificial social systems. Technical Report CS91-01, Weizmann Institute, 1991.

[Shoham and Tennenholtz, 1992] Yoav Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial societies (preliminary report). In *Proceedings o}AAAI-92.* AAAI, 1992.

[von Martial, 1992] F. von Martial. Coordinating Plans of Autonomous Agents. Springer-Verlag, 1992.