Coalition Formation among Bounded Rational Agents

Tuomas W. Sandholm and Victor R. Lesser *
{sandholm, lesser}@cs.umass.edu
University of Massachusetts at Amherst
Computer Science Department
Amherst, MA 01003

Abstract

This paper analyzes coalition formation among *self-interested* agents that need to solve combinatorial optimization problems to operate efficiently in the world. By colluding (coordinating their actions by solving a joint optimization problem), the agents can sometimes save costs compared to operating individually. A model of *bounded rationality* is adopted, where computation resources are costly. It is not worth solving the problems optimally: solution quality is decision-theoretically traded off against computation cost. A normative theory of coalitions among bounded rational (BR) agents is devised. The optimal coalition structure and its stability are significantly affected by the agents' algorithms' performance profiles (PPs) and the cost of computation. This relationship is first analyzed theoretically. A domain classification including rational and BR agents is introduced. Experimental results are presented in the distributed vehicle routing domain using real data from 5 dispatch centers; the optimal coalition structure for BR agents differs significantly from the one for rational agents. These problems are NP-complete and the instances are so large that, with current technology, any agent's rationality is bounded by computational complexity.

1 Introduction

In many domains, self-interested real world parties (e.g. companies) need to solve combinatorial optimization problems to operate efficiently. Often they can save costs by coordinating their activities with other parties. Such settings occur for example in distributed manufacturing among multiple companies and in distributed vehicle routing among dispatch centers. When the planning activities are automated, it is useful to also automate the coordination activities via a negotiating software agent representing each party. In such automated negotiations among self-interested agents, the question of coordination arises: what coalitions should the agents form, are they stable, and how should costs be divided within each coalition? Coalition formation includes three activities. One is *coalition structure generation:* formation of coalitions by the agents such that agents within each coalition coordinate their activities, but agents do not coordinate between coalitions. The second is the *solving of the combinatorial optimization problem* of each coalition. Conceptually this involves distributing the tasks of the coalition among the member agents and solving the optimization problem of each agent given its resources and the tasks it was distributed. The coalition's objective is to maximize monetary value: money received from outside the system for accomplishing tasks minus the cost of using resources.[1] Third, agents within each coalition have to agree on how to *divide this value* of the generated solution. These activities interact. For example, the coalition that an agent wants to join depends on the portion of the value that the agent would be allocated in each potential coalition.

Coalition formation has been widely studied [Kahan and Rapoport, 1984; van der Linden and Verbeek, 1985; Raiffa, 1982; Shechory and Kraus, 1995; Zlotkin and Rosenschein, 1994; Ketchpel, 1994], but to our knowledge, only among *rational agents.* Let us call the entire set of agents *A.* Say, that the lowest cost achievable by agents 5 C *A* working together, but without any other agents, is c^R_S. This is the minimum cost to handle the tasks of agents 5 with the resources of agents S. A coalition game is defined by a characteristic function v^R_S, which defines the *value* of each coalition S:

$$v^R_b = -C^R_S. \qquad (1)$$

The superscript *R* emphasizes that we mean the *rational* value of the coalition, i.e. the maximum value that is reachable by the coalition given its optimization problem. A rational agent can solve this combinatorial problem optimally without any deliberation costs such as CPU time costs or time delay costs.

If the problem is hard and the instance is large, it is unrealistic to assume that it can be solved without deliberation costs. This paper adopts a model of *bounded rationality* [Simon, 1982; Good, 197l], where each agent has to pay for the computational resources (CPU cycles) that it uses for deliberation. A fixed computation cost $c_{Comp} > 0$ per CPU time unit is assumed.[2] The domain cost associated with coalition 5 is denoted by

[1] In some problems, not all tasks have to be handled. This can be incorporated by associating a cost with each omitted task. Then problem solving also involves the selection of tasks to handle. The theory of this paper applies to such cases but in our example application, all tasks have to be handled, and no payments from outside the system are received for them.

[2] In practice, CPU time can already be bought on supercomputers. The market for CPU time is assumed to be so large that the demand of the agents we are studying does not impact the price of a CPU time unit. It is also assumed that

'Supported by ARPA contract N00014-92-J-1698.

$c_S(r_S) > 0$, i.e. it depends on (decreases with) the allocated computation resources r_S, Fig. 1. The functions $c_S(r_S)$ can be viewed as performance profiles (PPs) of the problem solving algorithm. They are used to decide how much CPU time to allocate to each computation. With this model of bounded rationality, the value of a coalition with BR agents can be defined. Each coalition minimizes the sum of solution cost and computation cost:

$$v_S(c_{comp}) = -\min_{r_S}[c_S(r_S) + c_{comp} \cdot r_S]. \qquad (2)$$

The coalition value decreases as the CPU time unit cost c_{comp} increases, Fig. 1. Our model also incorporates a second form of bounded rationality: the base algorithm may be incomplete, i.e. it might never find the optimal solution. If it is complete, the BR value of a coalition when $c_{comp} = 0$ equals the rational value ($v_S(0) = v_S^R$).

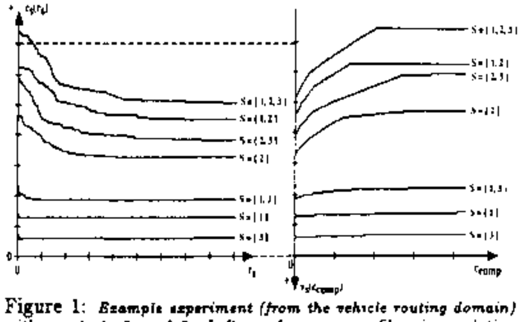

Figure 1: *Example experiment (from the vehicle routing domain) with agents 1, 2, and 3. Left: performance profiles, i.e. solution cost as a function of allocated computation resources. Right: BR coalition value as a function of computation cost.*

Conceptually the agents use *design-to-time algorithms* [Garvey and Lesser, 1993; Zilberstein, 1993]: once an agent has decided how much CPU time r_S it will allocate to a computation, it can design an algorithm that will find a solution of cost $c_S(r_S)$. The design-to-time framework is used instead of the *anytime* framework [Dean and Boddy, 1988; Sandholm and Lesser, 1994; Zilberstein, 1993] because to devise a theory of self-interested agents, the possibility that they design their algorithms to time has to be accounted for. With deterministic PPs, for any desired computation time allocation or solution quality, a noninterruptible design-to-time algorithm can be constructed that performs no worse than an interruptible anytime algorithm. We assume that the PPs exactly predict the solution cost attained for a given CPU time allocation. So, we have relaxed the assumption that the base level algorithm is optimal (complete and costless), but instead we assume that the meta-level deliberation controller is optimal (exact and costless). Assuming optimality of the meta-level is more realistic than assuming optimality of the base level, but it still does not match reality exactly. In practice there is uncertainty in each PP: the meta-level is not exact.[3] Secondly, the PP depends on several features of the problem instance, and computing the mapping from the instance to the PP [Sandholm and Lesser, 1994] may take considerable time, thus making the meta-level itself costly. In the limit, the base algorithm would be run at the meta-level to determine what it would achieve for a given time setting. Assuming an optimal meta-level enables analyzing bounded rationality at the base level in isolation from uncertainty of the PPs. It also allows us to sidestep the problem of having a meta-meta-level controlling the meta-level, a meta-meta-meta-level controlling the meta-meta-level, and so on *ad infinitum.*

We assume that the problem instances (tasks and resources) of all agents are common knowledge. This is somewhat unrealistic in open environments with a large number of agents. In practice it is often necessary to learn the other agents' characteristics from previous encounters. Alternatively, the agents can be made to explicitly declare their tasks and resources, but they may lie in order to gain monetarily. [Rosenschein and Zlotkin, 1994] analyze when rational agents are motivated to declare truthfully. Unfortunately that work assumes only two agents and that they can optimally solve exponentially many NP-complete problems without computation costs. Even under these assumptions, in most cases, truth-telling is not achieved. The effect of bounded rationality on truthful revelation is unknown.

For now—this is relaxed in Section 5—we assume that the agents solve the combinatorial optimization problems equally well and that this is common knowledge. For any coalition's problem and for any setting of CPU time, the cost of the solution potentially generated by each agent is the same. The agents need not generate the same solutions, only the same quality.

With such shared deterministic PPs, each agent knows the value $v_S(c_{comp})$ of each potential coalition S upfront. Therefore coalition formation will take place before any computation. After collusion, each coalition computes its solution using the optimal amount of CPU time r_S as defined by Equation 2. Because in our model, rationality is bounded by CPU time cost, it costs the same for one agent to use *nt* CPU time units as it costs n agents to use *t* units. Therefore, it is best if a coalition's optimization problem is solved by a single agent. This is trivially true since an agent could simulate distributed problem solving among *n* agents for time *t* by using a local algorithm for *nt*. Conversely, it is not always possible (due to redundancy etc.) for n agents solving the problem for time *t* to reach a solution of the same quality as one agent using *nt* can reach. The computing agent can be arbitrarily chosen from within the coalition, and the coalition pays that agent its true cost for computing. This cost along with the domain solution cost contribute to $v_S(c_{comp})$, which is divided among the agents in the

this price is common to all agents, which corresponds to an open CPU cycle market.

[3] If the PPs are only probabilistically known, anytime algorithms may be desirable due to their flexibility with respect to termination time. In general, for optimal meta-reasoning, the remaining part of a probabilistic PP should be conditioned on the algorithm's performance on that problem instance on previous CPU time steps [Sandholm and Lesser, 1994; Zilberstein, 1993]. Such conditioning, anytime algorithms, and their integration to coalition formation are part of our current research and are too long to be presented here [Sandholm and Lesser, 1995].

coalition as will be presented later.

The value of a coalition may depend on the actions of non-member agents due to positive and negative interactions of the agents' solutions. Such settings can be modeled as normal form games (NFGs), Fig. 2. Coalition formation is usually studied in characteristic function games (CFGs), where the value of each coalition S is given by the characteristic function v_S^R, and is thus not a function of the actions of non-members. CFGs are a strict subset of NFGs.[4] The equivalent of CFGs among BR agents are BRCFGs (Fig. 2) where the value of each coalition S is defined by $v_S(c_{comp})$. This paper mainly studies BRCFGs. Non-BRCFGs are addressed in Sec. 5.

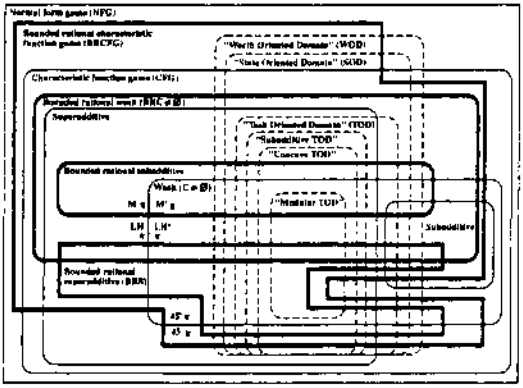

Figure 2: *Venn diagram of negotiation domains. Normal lines show the classification for rational agents. Bold lines show our new classification for BR agents, and how it relates to the rational case. Dotted lines show the rational agent domain classification of [Rosenschein and Zlotkin, 1994]. They use "Subadditive" to mean that an agent's cost for handling tasks is subadditive in tasks. We use subadditive to refer to coalition value functions that are subadditive in agents. The figure does not reflect the fact that Rosenschein and Zlotkin do not allow sidepayments.*

The paper is organized as follows. Section 2 studies the optimal coalition structure for BR agents, and Section 3 analyses its stability. Section 4 presents experimental results in the distributed vehicle routing domain with real data. Section 5 discusses agents with different problem solving capabilities. Section 6 presents related research, and 7 concludes and describes future research.

2 Optimality: BR superadditivity

Any outcome of a game can be analyzed with respect to *social welfare*, which is defined as the sum of the agents' payoffs. The payoff that agent i gets is called $x_i \in \mathcal{R}$. The sum of the agents' x_i's has to equal the sum of the values of the coalitions in the coalition structure (CS) that formed: no wealth is generated from nothing and no wealth disappears. With bounded rational agents, these coalition values incorporate the computation costs.

In *superadditive* games the agents are best off by forming the grand coalition—from a social welfare viewpoint. Superadditivity means that the value of one coalition plus the value of another coalition is never more than the value of these coalitions joined into one coalition:

Definition. 2.1 [5] *A game is superadditive if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}^R \geq v_S^R + v_T^R$. *See Fig. 2.*

When computation cost is ignored, this is almost always the case, because at worst, the agents in the composite coalition can use the solutions that they had when they were in separate coalitions. Thus, the agents would almost always be better off from a social welfare point of view by forming the grand coalition, i.e. $CS^{R*} = \{A\}$. A game can be non-superadditive only if the collusion process itself involves some cost, e.g. anti-trust penalties. Some non-superadditive games are *subadditive*, Fig. 2:

Definition. 2.2 *A game is subadditive if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}^R < v_S^R + v_T^R$.

In subadditive games, the agents are best off by operating alone, i.e. $CS^{R*} = \{\{a_1\}, \{a_2\}, ..., \{a_{|A|}\}\}$. Some games are neither superadditive nor subadditive, because the characteristic function fulfills the condition of superadditivity for some coalitions and the condition of subadditivity for others. In such cases, the social welfare maximizing coalition structure varies.

Now we present a new concept for BR agents that is analogous to superadditivity among rational agents. From a social welfare viewpoint, BR agents are best off forming the grand coalition ($CS^* = \{A\}$) if the game is *bounded rational superadditive* (BRS), Fig. 2. This requires that the best value that one coalition can reach given the computation cost plus the best value that another coalition can reach given the computation cost is never greater than the best value that these coalition can reach as a composite coalition given the computation cost:

Definition. 2.3 *A game is bounded rational superadditive (BRS) for computation cost* c_{comp} *if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}(c_{comp}) \geq v_S(c_{comp}) + v_T(c_{comp})$.

BR superadditivity does not always coincide with superadditivity. In general, for a given $c_{comp} \geq 0$, a coalition game can be superadditive, BRS, both, or neither, Fig. 2. Only some non-BRS games are BR subadditive, Fig. 2:

Definition. 2.4 *A game is bounded rational subadditive for computation cost* c_{comp} *if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}(c_{comp}) < v_S(c_{comp}) + v_T(c_{comp})$.

BR superadditivity depends on the performance profiles and the unit cost of computation. The next theorem states a natural condition on the PPs. If the condition holds, the game is BRS for any $c_{comp} \geq 0$.

Theorem 2.1 ***BRS (sufficient condition).*** $[(\forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S \geq 0, \forall r_T \geq 0), c_{S \cup T}(r_S + r_T) \leq c_S(r_S) + c_T(r_T)] \Rightarrow$ *Game is BRS* $\forall c_{comp} \geq 0$.[6]

The condition states that the domain cost for coalition S after allocating a certain amount r_S of computation plus the domain cost to another coalition T after allocating a certain amount r_T of computation is never less than the

[4] The two are equivalent in constant-sum games with unrestricted side-payments and perfect communication. In such games, the characteristic function value of a coalition is its minimax value from the normal form game [van der Linden and Verbeek, 1985].

[5] Definitions 2.1, 2.2 and 3.1 are from game theory.

[6] Proofs are presented in [Sandholm and Lesser, 1995].

domain cost of these coalitions combined after allocating $r_S + r_T$. This is always achievable if the base algorithm is intelligent enough, because in the worst case, the algorithm can allocate r_S on the problem of S and then do the problem of T using r_T separately. Given a large coalition, it is difficult to intelligently guess an efficient decomposition of this type. Usually, the algorithm that is used on the composite problem does not apply this type of a separate solving. Whether the algorithm's performance actually satisfies the condition without using a separate solving approach depends on the problem, the specific instances under study, and the algorithm itself. In general, the game can be BRS $\forall c_{comp} \geq 0$ even if the above condition does not hold on the PPs:

Theorem 2.2 $[(\forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S \geq 0, \forall r_T \geq 0), c_{S\cup T}(r_S + r_T) \leq c_S(r_S) + c_T(r_T)] \not\Leftarrow$ *Game is BRS* $\forall c_{comp} \geq 0$.

It is reasonable to assume that the PP $c_S(r)$ is decreasing in r if the agent can inexpensively store the best solution it has arrived at so far. Furthermore, $c_S(r)$ is often convex in r: greater savings are achieved in the early stages of computation and the savings per time unit decrease as problem solving proceeds. We conjecture that PPs of design-to-time algorithms are almost always convex. On the other hand, PPs of anytime algorithms are typically not convex at points where the base algorithm switches from one approach to another. One example is completing an iterative refinement algorithm by running an exhaustive complete algorithm after the refinement phase. Another example is switching from using one refinement operator (e.g. 2-swap in TSP [Lin and Kernighan, 1971; Sandholm, 1993]) to using another refinement operator (e.g. 3-swap in TSP). Furthermore, refinements often decrease solution cost in a step-wise manner rendering the PPs locally nonconvex—as in our experiments (Fig. 1 left)[7]. The PPs in our experiments exhibited an overall convex nature, but also had true local nonconvexities (because the design-to-time algorithms were constructed from anytime algorithms, and were not tailored for each time setting separately, Sec. 4). Convexity is significant because with convex PPs, a domain is BRS for all computation costs if *and only if* the condition of Theorem 2.1 on the PPs holds:

Theorem 2.3 ***BRS (necessary and sufficient condition).*** *Let us restrict ourselves to such performance profiles that* $\forall U \subseteq A, c_U(r)$ ***is decreasing*** *and* ***convex*** *in* r. *Now,* $[(\forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S \geq 0, \forall r_T \geq 0), c_{S\cup T}(r_S + r_T) \leq c_S(r_S) + c_T(r_T)] \Leftrightarrow$ *Game is BRS* $\forall c_{comp} \geq 0$.

Analogous to Theorem 2.1, there is an easy sufficient condition on the PPs that guarantees that the game is BR subadditive for all computation costs:

Theorem 2.4 ***Bounded rational subadditivity (sufficient condition).*** $[(\forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S \geq 0, \forall r_T \geq 0), c_{S\cup T}(r_S + r_T) > c_S(r_S) + c_T(r_T)] \Rightarrow$ *Game is bounded rational subadditive* $\forall c_{comp} \geq 0$.

[7]If the algorithm is stochastic, these step-related nonconvexities are reduced as the PP is averaged over multiple runs.

If the game is BR subadditive, agents are best off alone, i.e. by colluding with nobody ($CS^* = \{\{a_1\}, \{a_2\}, ..., \{a_{|A|}\}\}$). In games that are neither BRS nor bounded rational subadditive, the optimal CS varies, and several CSs may be equally good wrt. social welfare. We denote any one of these best CSs by CS^*.

3 Stability: bounded rational core

In the previous section we presented conditions on the PPs that describe what CS the agents are best off forming from the social welfare viewpoint. In this section we analyze the *stability* of that CS. Can the social good be distributed among the agents so that each agent is motivated to stay with CS^{R*} (individual rationality)? Furthermore, can it be distributed so that every subgroup of agents is better off with CS^{R*} than by forming a coalition of their own (coalition rationality)? The core (C) is the solution concept that satisfies both of these conditions [Kahan and Rapoport, 1984; van der Linden and Verbeek, 1985; Raiffa, 1982]. The core of a game is a set of vectors $\vec{x}$, where each $\vec{x}$ is a vector of payoffs to the agents in such a manner that no subgroup (individual agents and the group of all agents are also subgroups) is motivated to depart from CS^{R*}. Given payoffs according to $\vec{x}$, the value of each subgroup is less than or equal to the sum of the payoffs that the agents of that subgroup get under CS^{R*}. Obviously, only CSs that maximize welfare can be stable in the sense of the core, because from any other CS the group of all agents would prefer to switch to a CS^{R*}. Formally,

Definition. 3.1 *Core* $C = \{\vec{x} | \forall S \subset A, \sum_{i\in S} x_i \geq v_S^R$ *and* $\sum_{i\in A} x_i = \sum_{j \in CS^{R*}} v_{S_j}^R\}$.

The core is the strongest solution concept used for coalition formation. It is often too strong: in many cases it is empty, i.e. the social good cannot be divided so that the individual and coalition rationality conditions are satisfied [Kahan and Rapoport, 1984; van der Linden and Verbeek, 1985; Raiffa, 1982]. A lesser problem is that the core may include multiple $\vec{x}$'s and the agents have to agree on one of them. An often used solution is to pick the *nucleolus* which is, intuitively speaking, the center of the core [Kahan and Rapoport, 1984; van der Linden and Verbeek, 1985; Raiffa, 1982]. Games with non-empty cores are called *weak*, Fig. 2.

Now we introduce the analog of the core for BR agents.

Definition. 3.2 *The* ***bounded rational core (BRC)*** *for computation cost* c_{comp} *is* $BRC(c_{comp}) = \{\vec{x} | \forall S \subset A, \sum_{i\in S} x_i \geq v_S(c_{comp})$ *and* $\sum_{i\in A} x_i = \sum_{j\in CS^*} v_{S_j}(c_{comp})\}$.

If the BRC is not empty, BR agents can divide the social good among themselves in a way that no subgroup is motivated to break away from CS^*. Sometimes the BRC is empty, but this does not always coincide with the core being empty. There are games, where the BRC and the core exist, games where either one of them exists separately, and games where both are empty, Fig. 2. If the agents are best off working separately, the CS with separate agents is stable, Fig. 2:

Theorem 3.1 ***Bounded rational subadditive core.*** *Game is bounded rational subadditive for some* $c_{comp} \Rightarrow BRC(c_{comp}) \neq \emptyset$.

In domains that are not BR subadditive, the BRC is sometimes empty. The condition $C \neq \emptyset$ can be converted into necessary and sufficient conditions on the v_S^R's in games where the grand coalition maximizes social welfare [Shapley, 1967; Charnes and Kortanek, 1966]. We convert the condition $BRC(c_{comp}) \neq \emptyset$ into conditions on the $v_S(c_{comp})$'s analogously. Let $B_1, ..., B_p$ be distinct, nonempty, proper subsets of A. The set $B = \{B_1, ..., B_p\}$ is called *balanced* if there are positive coefficients $\lambda_1, ..., \lambda_p$ such that $\forall i \in A, \sum_{\{j|i \in B_j\}} \lambda_j = 1$. A *minimal balanced set* includes no other balanced sets.

Theorem 3.2 ***Bounded rational core in grand coalition games.*** *In games where* $CS^* = \{A\}$ *for some* c_{comp}, $BRC(c_{comp}) \neq \emptyset$ *iff for every minimal balanced set* $B = \{B_1, ..., B_p\}, \sum_{j=1}^{p} \lambda_j v_{S_j}(c_{comp}) \leq v_A(c_{comp})$.

Example. In any 3-agent game where $CS^* = \{A\}$ for some c_{comp}, $BRC(c_{comp}) \neq \emptyset$ iff $v_{\{1\}}(c_{comp}) + v_{\{2,3\}}(c_{comp}) \leq v_{\{1,2,3\}}(c_{comp})$ and $v_{\{2\}}(c_{comp}) + v_{\{1,3\}}(c_{comp}) \leq v_{\{1,2,3\}}(c_{comp})$ and $v_{\{3\}}(c_{comp}) + v_{\{1,2\}}(c_{comp}) \leq v_{\{1,2,3\}}(c_{comp})$ and $v_{\{1\}}(c_{comp}) + v_{\{2\}}(c_{comp}) + v_{\{3\}}(c_{comp}) \leq v_{\{1,2,3\}}(c_{comp})$ and $\frac{1}{2}v_{\{1,2\}}(c_{comp}) + \frac{1}{2}v_{\{1,3\}}(c_{comp}) + \frac{1}{2}v_{\{2,3\}}(c_{comp}) \leq v_{\{1,2,3\}}(c_{comp})$. All but the last inequality are implied by the fact that $CS^* = \{A\}$.

Example. In any 4-agent game where $CS^* = \{A\}$ for some c_{comp}, $BRC(c_{comp}) \neq \emptyset$ iff the 41 inequalities of Table 1 hold. Constraints 1, 2, 3 and 5 correspond to partitions of A (all λ's are 1). They are thus implied by the fact that $CS^* = \{A\}$.

Id	Constraint	#
1	$v_{\{1,2\}}(c_{comp}) + v_{\{3,4\}}(c_{comp}) \leq v_{\{1,2,3,4\}}(c_{comp})$	3
2	$v_{\{1,2,3\}}(c_{comp}) + v_{\{4\}}(c_{comp}) \leq v_{\{1,2,3,4\}}(c_{comp})$	4
3	$v_{\{1,2\}}(c_{comp}) + v_{\{3\}}(c_{comp}) + v_{\{4\}}(c_{comp})$ $\leq v_{\{1,2,3,4\}}(c_{comp})$	6
4	$\frac{1}{2}v_{\{1,2,3\}}(c_{comp}) + \frac{1}{2}v_{\{1,2,4\}}(c_{comp}) + \frac{1}{2}v_{\{3,4\}}(c_{comp})$ $\leq v_{\{1,2,3,4\}}(c_{comp})$	6
5	$v_{\{1\}}(c_{comp}) + v_{\{2\}}(c_{comp}) + v_{\{3\}}(c_{comp}) + v_{\{4\}}(c_{comp})$ $\leq v_{\{1,2,3,4\}}(c_{comp})$	1
6	$\frac{1}{2}v_{\{1,2\}}(c_{comp}) + \frac{1}{2}v_{\{1,3\}}(c_{comp}) + \frac{1}{2}v_{\{2,3\}}(c_{comp})$ $+v_{\{4\}}(c_{comp}) \leq v_{\{1,2,3,4\}}(c_{comp})$	4
7	$\frac{1}{2}v_{\{1,2,3\}}(c_{comp}) + \frac{1}{2}v_{\{1,4\}}(c_{comp}) + \frac{1}{2}v_{\{2,4\}}(c_{comp})$ $+\frac{1}{2}v_{\{3\}}(c_{comp}) \leq v_{\{1,2,3,4\}}(c_{comp})$	12
8	$\frac{2}{3}v_{\{1,2,3\}}(c_{comp}) + \frac{1}{3}v_{\{1,4\}}(c_{comp}) + \frac{1}{3}v_{\{2,4\}}(c_{comp})$ $+\frac{1}{3}v_{\{3,4\}}(c_{comp}) \leq v_{\{1,2,3,4\}}(c_{comp})$	4
9	$\frac{1}{3}v_{\{1,2,3\}}(c_{comp}) + \frac{1}{3}v_{\{1,2,4\}}(c_{comp}) + \frac{1}{3}v_{\{1,3,4\}}(c_{comp})$ $+\frac{1}{3}v_{\{2,3,4\}}(c_{comp}) \leq v_{\{1,2,3,4\}}(c_{comp})$	1

Table 1: *Conditions for existence of the BRC in a 4-agent grand coalition game. Last column shows the number of constraints generated from that constraint by permuting the agents (including the presented permutation).*

In BRS games, a subset of the above inequalities suffices. Let us call a minimal balanced set *proper* if no two of its elements are disjoint.

Theorem 3.3 ***BRS bounded rational core.*** *In a game that is BRS for some* c_{comp}, $BRC(c_{comp}) \neq \emptyset$ *iff for every proper minimal balanced set* $B = \{B_1, ..., B_p\}, \sum_{j=1}^{p} \lambda_j v_{S_j}(c_{comp}) \leq v_A(c_{comp})$. *Furthermore, this set of inequalities is minimal: no smaller set is sufficient.*

Example. In a 3-agent game that is BRS for some c_{comp}, $BRC(c_{comp}) \neq \emptyset$ iff $\frac{1}{2}v_{S_{\{1,2\}}}(c_{comp}) + \frac{1}{2}v_{S_{\{1,3\}}}(c_{comp}) + \frac{1}{2}v_{S_{\{2,3\}}}(c_{comp}) \leq v_{S_{\{1,2,3\}}}(c_{comp})$.

Example. In a 4-agent game that is BRS for some c_{comp}, $BRC(c_{comp}) \neq \emptyset$ iff the 11 conditions acquired from Table 1's constraints 4, 8 and 9 are satisfied.

Next we present conditions on the PPs that are *sufficient* to guarantee that the BRC exists. According to Theorem 3.1, the conditions on the PPs that guarantee BR subadditivity (Theorem 2.4) form one such set of conditions. The following set suffices for games where $CS^* = \{A\}$:

Theorem 3.4 ***BRC in grand coalition games (sufficiency).*** *In games where* $CS^* = \{A\}$ *for some* $c_{comp} \geq 0$, *[for every minimal balanced set* $B = \{B_1, ..., B_p\}, (\forall B \in B, \forall r_B \geq 0) \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) \geq c_A(\sum_{j=1}^{p} \lambda_j r_{B_j})] \Rightarrow BRC(c_{comp}) \neq \emptyset$.

If $CS^* = \{A\}$ for *all* $c_{comp} \geq 0$, the above conditions guarantee existence of the $BRC(c_{comp})$ for *all* $c_{comp} \geq 0$. In BRS games, fewer conditions suffice:

Theorem 3.5 ***BRC in BRS games (sufficiency).*** *In a game that is BRS for some* $c_{comp} \geq 0$, *[for every proper minimal balanced set* $B = \{B_1, ..., B_p\}, (\forall B \in B, \forall r_B \geq 0) \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) \geq c_A(\sum_{j=1}^{p} \lambda_j r_{B_j})] \Rightarrow BRC(c_{comp}) \neq \emptyset$.

Again, if the game is BRS for *all* $c_{comp} \geq 0$, the above conditions guarantee existence of the $BRC(c_{comp})$ for *all* $c_{comp} \geq 0$.

Example. In a 3-agent game that is BRS $\forall c_{comp} \geq 0$, $[(\forall r_{\{1,2\}} \geq 0, \forall r_{\{1,3\}} \geq 0, \forall r_{\{2,3\}} \geq 0), \frac{1}{2}c_{\{1,2\}}(r_{\{1,2\}}) + \frac{1}{2}c_{\{1,3\}}(r_{\{1,3\}}) + \frac{1}{2}c_{\{2,3\}}(r_{\{2,3\}}) \geq c_{\{1,2,3\}}(\frac{1}{2}r_{\{1,2\}} + \frac{1}{2}r_{\{1,3\}} + \frac{1}{2}r_{\{2,3\}})] \Rightarrow \forall c_{comp} \geq 0, BRC(c_{comp}) \neq \emptyset$.

4 Experimental results: vehicle routing

BR coalition formation was tested in the vehicle routing domain using one week real-world vehicle and order data from 5 geographically distributed dispatch centers. Each center had its own vehicles and delivery tasks. In all, they had 771 deliveries to make with 77 vehicles. Each vehicle had to begin and end its tour at the depot of its center, but neither the pickup nor the drop-off locations of the orders were at the depot. The vehicles had heterogeneous maximum load weight and maximum load volume constraints. All vehicles had the same maximum route length. The domain cost $c_S(r_S)$ for a coalition S was the sum of the route lengths of the vehicles of that coalition (while handling all of its orders) in the solution that had been reached after computation r_S. The problem is NP-hard, because ΔTSP can be trivially reduced to it. It is in NP, because the cost and feasibility of a solution can easily be checked in polynomial time. Thus, the problem is NP-complete. Moreover, the problem instances in our example are so large that even the smallest ones are too hard to solve optimally. Therefore, rational coalition formation algorithms for the vehicle routing problem [Lundgren *et al.*, 1992] are unusable.

Our problem is outside the domain classification of [Rosenschein and Zlotkin, 1994], Fig. 2, because agents do not have symmetric capabilities due to heterogeneous fleets. If we extend their definition to allow asymmetric capabilities, our domain is in SOD \ TOD. If we further drop the maximum route length constraint (this experiment will also be presented), and restrict ourselves to domains where each center has at least one sufficient vehicle to satisfy the weight/volume constraints of any order of any center (not true in our data), then the domain is a TOD. It is not a "Subadditive TOD", because the depots are geographically distributed.

To analyse a game we ran the same algorithm on the vehicle routing problem of each subgroup of agents separately and thus acquired a PP for each potential coalition. The algorithm first generates an initial solution by giving each vehicle one long delivery and then, in order, giving each vehicle the delivery that can be added to its route with the least cost without violating the constraints. The second phase of the algorithm is based on iterative refinement. At each step, a delivery (chosen from a randomly ordered circular list) is removed from the routing solution and inserted back to the solution, but into the least expensive place while not violating the constraints. The drop-off location of the delivery has to be inserted after the pickup location into the same vehicle's route, but not necessarily into the same leg. We ran the refinement algorithm until no remove-insert operation enhanced the solution: a local optimum was reached. In the PPs we ignored the time to construct the initial solution, and only viewed how the solution cost decreased with more CPU seconds of iterative refinement, Fig. 1 left. The refinement algorithm is an anytime algorithm, but because the PPs are exact (as explained, they are precomputed for experimental purposes by running the base algorithm itself), the agents do not gain information from execution on that instance so far. Therefore the algorithm is equivalent to a design-to-time algorithm for our purposes.

We analysed all of the $\binom{5}{3} = 10$ 3-agent games that can be acquired by choosing 3 of the 5 dispatch centers.[8] Figure 1 shows the PPs with agents 1, 2 and 3. Each of our games is superadditive for reasons that were explained in Section 2. Thus rational agents would be best off by forming the grand coalition. Surprisingly, none of the games were BRS for any c_{comp}, Fig. 3. For c_{comp}'s in the mid-range, the 3-agent games were often BR subadditive (point M in Fig. 2), while in the low and high ranges (point LH in Fig. 2), they were often neither BRS nor bounded rational subadditive. Existence of the core for rational agents is unknown for our games: the points M and LH might really be M' and LH'. The BRC was non-empty in all 3-agent games for all values of c_{comp}. So, rational agents would be best off forming the possibly unstable grand coalition, while BR agents should form varying coalition structures (the grand coalition for some c_{comp}'s), which are always stable. We also reran the experiments without the maximum route length restriction, and these results prevailed, Fig. 3.

Centers 2, 3 and 5 were located near to each other, while 1 and 4 were far from each other and the other centers. Centers 1, 3, 4 and 5 transported heavy low volume items, while 2 transported light voluminous items. Centers 1..5 had 65, 200, 82, 124, and 300 deliveries, and 10, 13, 21, 18, and 15 vehicles respectively. Both with and without the route length restriction, 2 and 5 were best off by only mutually colluding for any c_{comp}. Their deliveries have considerable areal overlap due to adjacency, and the light voluminous items and heavy low volume items can be profitably joined into the weight and volume constrained vehicles. Centers 2 and 3 did not collude as much as 2 and 5 because 3's vehicles had tighter volume constraints than 5's—hindering the transport of 2's goods. No other two centers besides 2 and 5 were always best off in a 2-agent coalition independent of the third agent of the game. Relaxing the route length constraint increased collusion between the distant 2 and 4 while demoting collusion of the adjacent 2 and 3.

[8] There are 7 subgroups of the 3 agents: {1}, {2}, {3}, {1,2}, {2,3}, {3,1}, {1,2,3} and 5 coalition structures: {{1}, {2}, {3}}, {{1}, {2,3}}, {{2}, {1,3}}, {{3}, {1,2}}, {{1,2,3}}.

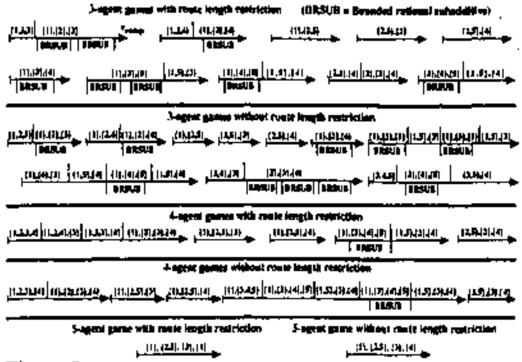

Figure 3: *Optimal coalition structure (CS^*) and bounded rational subadditivity as a function of c_{comp}. Tested by evaluating all possible coalition structures and super/subadditivity at varying points of c_{comp} chosen from a grid where c_{comp} is always incremented by 1%.*

Next we analysed the $\binom{5}{4} = 5$ 4-agent games and the 5-agent game with and without the route length restriction. In every game, the existence of $BRC(c_{comp})$ varied many times as a function of c_{comp}, but it existed for the largest values of c_{comp}. No game was BRS for any c_{comp}, but some games were bounded rational subadditive for interior values, Fig. 3. Thus, depending on c_{comp}, the games were at the points M, LH, or 45 (or M', LH', or 45') in Figure 2. The best coalition structure varied despite the fact that rational agents would be best off forming the grand coalition due to superadditivity. Again, whenever both agents 2 and 5 participated, they were best off by mutually colluding for all computation costs. In those games no other agents colluded.

Each step of the refinement algorithm takes $\Theta(vd^2)$ time, where v is the number of vehicles and d is the number of deliveries. Because this is superlinear in deliveries, a larger coalition can make fewer refinement steps in a given time than the agents in partitions of that coalition can. To compensate, a refinement step of the larger coalition would need to reduce solution cost more than a refinement step of a smaller coalition. The size of the saving has to be averaged over all refinement steps in the optimal time allocation. If c_{comp} is low, more time is allocated, and small coalitions will often run out of profitable refinements. If c_{comp} is high, less time is allocated, and all coalitions will have profitable refinements, though the larger coalition will have time to make fewer of them. Thus it was not surprising that in games where the grand coalition was optimal, it was optimal for very small computation costs only. Surprisingly, two agents colluding was often better than all agents working separately even for large c_{comp}'s. The result that *higher computation costs promote smaller coalitions* is somewhat deemphasized by our choice of not including the initial solution construction phase in the PPs. Shifting the PPs right to begin at the time when the initial solution was finished (instead of 0) would shift the PPs of small coalitions less than the PPs of large coalitions because the initial solution construction is superlinear both in tasks and vehicles. Thus small coalitions would gain an advantage—that is most significant for large c_{comp}. If the time of initial solution generation is discarded, the best coalition structure for the greatest computation costs depends only on the quality of the initial solutions of the different coalitions because no refinement steps are beneficial. For example, coalitions {1,3} (Fig. 1), {1,5} and {2,5} achieved a better initial solution cost than the sum of the initial solution costs of the two agents separately, Fig. 3.

5 Different performance profiles

So far games where each agent has the same PP for a given coalition were presented. In general, domains where the agents have different PPs—due to different algorithms—are not characteristic function games for BR agents (BRCFGs), because the value of a coalition sometimes depends on the actions of non-members. The value of a coalition can depend on whether an outside agent is willing to compute the solution for the coalition (for a payment) if its algorithm is better than any of the algorithms of the agents in the coalition. Also, interactions between domain solutions of different coalitions may exclude some problems from the class BR-CFG. In non-BRCFGs, BR superadditivity, BR subadditivity, and the BRC are undefined, Fig. 2. Instead, the Nash equilibrium may be a reasonable solution concept—although only individual agents are motivated to pertain to it: coalitions may prefer to deviate. These issues are discussed in [Sandholm and Lesser, 1995].

6 Related DAI research on collusion

Coalition formation has been widely studied in game theory [Kahan and Rapoport, 1984; van der Linden and Verbeek, 1985; Raiffa, 1982]; only the most relevant concepts were presented here. This section compares our work to other recent DAI work on coalition formation.

[Zlotkin and Rosenschein, 1994] analyze rational agents that cannot make side payments, while our agents do. Their analysis is limited to "Subadditive Task Oriented Domains" (STODs), which are a strict subset of CFGs, Fig. 2. In their solution concept, one agent handles all the tasks, because STODs never exhibit diseconomies of scale. We do not assume that one agent can take care of all the agents' tasks. Unlike our work, they also assume that all agents have the same capabilities. With exponential computation they guarantee each agent an expected value that equals its Shapley value [Kahan and Rapoport, 1984; Raiffa, 1982]. In a subset of STODs, "Concave Task Oriented Domains" (Fig. 2), the computational complexity is reduced to linear (in agents) using an encryption scheme. Yet at least one (intractable) combinatorial problem involving all tasks of all agents needs to be solved optimally.

[Ketchpel, 1994] presents a coalition formation method for rational agents which have different expectations of coalition values. The (computational) origin of these expectations is not addressed. His assumption of imperfect information differs from our setting, where the agents have perfect information, but cannot perfectly deduce. Ketchpel's coalition formation algorithm runs in cubic time in the number of agents, but does not guarantee stability. His protocol is based on mutual offers. , In practice it is hard to prevent out-of-protocol offers such as multiagent offers. In our approach, if the agents' payoff vector is chosen from within the BRC, the coalition structure is stable against *all* offers. Finally, his 2-agent auction is manipulable and computationally inefficient. He approaches the coalition formation and the payoff division problems simultaneously.

This is closely related to the contracting protocol of Sandholm [Sandholm, 1993] (TRACONET), where agents construct the global solution by contracting a small number of tasks at a time, and payments are made regarding each contract before new contracts take place. An agent updates its approximate solution after each task transfer. In general equilibrium approaches such as WALRAS [Wellman, 1992], non-manipulative agents iterate over the allocation of resources and tasks, and payments are made only after a final solution is reached.

[Shechory and Kraus, 1995] analyze coalition formation among rational agents with perfect information in domains that are not necessarily superadditive. Their protocol guarantees that if agents follow it, a certain stability criterion (K-stability) is met. This requires the solution of an exponential number of optimization problems. Their other protocol guarantees a weaker form of stability (polynomial K-stability), but only requires the solution of a polynomial number of optimization problems. Unfortunately, each one of these may be intractable. Their algorithm switches from one coalition structure to another guaranteeing improvements at each step: coalition structure formation is an anytime algorithm, although each domain problem is solved optimally. In our approach, each domain problem is solved using an approximation (design-to-time) algorithm.

7 Conclusions and future research

A normative theory of coalitions in combinatorial domains was presented, where the rationality of self-interested agents is bounded by computational complexity. A domain classification was presented for rational and BR agents. The algorithms used by the agents significantly impact the coalition structure that should form as well as its stability. Theorems were presented on the PPs guaranteeing BR superadditivity, BR subadditivity, and existence of the BRC. Although almost all domains are superadditive, BR superadditivity is surprisingly all but obvious in practice. None of the vehicle routing games of our experiments—using real data and a reasonable iterative refinement algorithm—exhibited BR superadditivity. Thus, the optimal CS for BR agents varied, although rational agents should always form the grand coalition. Section 2 developed conditions on the PPs that guarantee BR superadditivity, and it discussed a separate solving approach—based on a non-deterministic splitting step—that guarantees that the base algorithm fulfills those conditions. We are currently developing methods of constructing algorithms that satisfy the conditions without such splitting. The observed BR subadditivity of some of the games implies a non-empty BRC: the best CS in those games is stable. Even when BR subadditivity did not hold, the BRC was often non-empty—especially for large c_{comp}. Often with superlinear iterative refinement steps, low c_{comp} promotes large coalitions while high c_{comp} suggests smaller ones. The best BR CSs mostly agreed with our intuitions of what coalitions should form based on strategic domain specific considerations such as adjacency of the dispatch centers and the combinability of their loads.

Our model of bounded rationality is based on costly computation resources. Future work includes analyzing another model, where each agent has a fixed free CPU and no more CPU time can be bought. If the domain cost increases with real time due to a dynamic environment, such agents with bounded computational capabilities are often best off by distributing the computation. In the costly computation model of this paper, it is best to allocate each coalition's computation to a single agent. The models are equivalent if the domain cost increases linearly with real time and distribution does not speed up computation.

Extensions include generalizing these methods to agents with different PPs, probabilistic PPs, and anytime algorithms where PPs are conditioned on execution so far [Sandholm and Lesser, 1995; 1994; Zilberstein, 1993]. Agents with probabilistic PPs may want to reselect a coalition if the value of their original coalition is lower than expected—but sunk computation cost has already been incurred. Future research also includes agents that can refine solutions generated by others.

References

[Charnes and Kortanek, 1966] A Charnes and K O Kortanek. On balanced sets, cores, and linear programming. Technical Report 12, Cornell Univ., Dept. of Industrial Eng. and Operations Res., Ithaca, NY, 1966.

[Dean and Boddy, 1988] Thomas Dean and Mark Boddy. An analysis of time-dependent planning. In Proceedings of the National Conference on Artificial Intelligence, pages 49-54, St. Paul, MN, August 1988.

[Garvey and Lesser, 1993] Alan Garvey and Victor Lesser. Design-to-time real-time scheduling. IEEE Trans. Systems, Man, and Cybernetics, 23(6), 1993.

[Good, 1971] Irving Good. Twenty-seven principles of rationality. In V Godambe and D Sprott, editors, Foundations of Statistical Inference. Toronto: Holt, Rinehart, Winston, 1971.

[Kahan and Rapoport, 1984] James P Kahan and Amnon Rapoport. Theories of Coalition Formation. Lawrence Erlbaum Associates Publishers, 1984.

[Ketchpel, 1994] Steven Ketchpel. Forming coalitions in the face of uncertain rewards. In AAAI, pages 414-419, Seattle, WA, July 1994.

[Lin and Kernighan, 1971] S Lin and B W Kernighan. An effective heuristic procedure for the traveling salesman problem. Operations Research, 21:498-516, 1971.

[Lundgren et al., 1992] M G Lundgren, K Jornsten, and P Varbrand. On the nucleolus of the basic vehicle routing game. Technical Report 1992-26, Linkoping Univ., Dept. of Mathematics, Sweden, 1992.

[Raiffa, 1982] H. Raiffa. The Art and Science of Negotiation. Harvard Univ. Press, Cambridge, Mass., 1982.

[Rosenschein and Zlotkin, 1994] Jeffrey S Rosenschein and G Zlotkin. Rules of Encounter. MIT Press, 1994.

[Sandholm and Lesser, 1994] Tuomas W Sandholm and Victor R Lesser. Utility-based termination of anytime algorithms. In ECAI-94 Workshop on Decision Theory for DAI Applications, pp. 88-99, Amsterdam. Extended version: Univ. of Mass., Comp. Sci. TR 94-54.

[Sandholm and Lesser, 1995] Tuomas W Sandholm and Victor R Lesser. Coalition formation among bounded rational agents. Technical report, University of Massachusetts at Amherst Computer Science Department, 1995. Extended version. In preparation.

[Sandholm, 1993] Tuomas W Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In AAAI, pp. 256-262, 1993.

[Shapley, 1967] Lloyd S Shapley. On balanced sets and cores. Naval Research Logistics Quarterly, 14:453-460, 1967.

[Shechory and Kraus, 1995] Onn Shechory and Sarit Kraus. Feasible formation of stable coalitions among autonomous agents in general environments. CI Journal, 1995. Submitted.

[Simon, 1982] Herbert A Simon. Models of bounded rationality, volume 2. MIT Press, 1982.

[van der Linden and Verbeek, 1985] Wim J van der Linden and Albert Verbeek. Coalition formation: a game-theoretic approach. In Henk A M Wilke, editor, Coalition Formation, volume 24 of Advances in Psychology. North Holland, 1985.

[Wellman, 1992] Michael Wellman. A general- equilibrium approach to distributed transportation planning. In AAAI, pages 282-289, San Jose, CA, July 1992.

[Zilberstein, 1993] Shlomo ZUberstein. Operational rationality through compilation of anytime algorithms. PhD thesis, University of California, Berkeley, 1993.

[Zlotkin and Rosenschein, 1994] GUad Zlotkin and Jeffrey S Rosenschein. Coalition, cryptography and stability: Mechanisms for coalition formation in task oriented domains. In AAAI, pages 432-437, Seattle, WA, July 1994.