Coalition Formation among Bounded Rational Agents

Tuomas W. Sandholm and Victor R. Lesser
Computer Science Department
University of Massachusetts at Amherst

CMPSCI Technical Report 95-71
September 7, 1995

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 coalitions 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, protocol-independent 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 unit 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]

[1]Supported by ARPA contract N00014-92-J-1698. The content does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred. T. Sandholm also supported by the Finnish Culture Foundation, Honkanen Foundation, Ella and George Ehrnrooth Foundation, Finnish Science Academy, Leo and Regina Wainstein Foundation, Finnish Information Technology Research Foundation, and Jenny and Antti Wihuri Foundation. A short early version of this paper appeared in [22].

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 deciding how to distribute 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.[2] 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 [12, 27, 18, 25, 30, 13], 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 $S \subseteq A$ working together, but without any other agents, is c_S^R. This is the minimum cost to handle the tasks of agents S with the resources of agents S. A coalition game is defined by a characteristic function v_S^R, which defines the *value* of

[2] 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.

each coalition S:

$$v_S^R = -c_S^R. \tag{1}$$

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* [26, 10], where each agent has to pay for the computational resources (CPU cycles) that it uses for deliberation. A fixed computation cost $c_{comp} \geq 0$ per CPU time unit is assumed.[3] The domain cost associated with coalition S is denoted by $c_S(r_S) \geq 0$, 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:

$$v_S(c_{comp}) = -\min_{r_S}[c_S(r_S) + c_{comp} \cdot r_S].^4 \tag{2}$$

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 ($v_S(0) = v_S^R$). In all, the bounded rational value of a coalition is determined by three factors:

[3]In practice, CPU time can already be bought on supercomputers. Similarly, the developing infrastructure for remotely executing agents provides an equivalent setting. For example in Telescript [9], the remotely executing agents pay Teleclicks for CPU time to the owner of the host machine. In this paper, 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 this price is common to all agents, which corresponds to an open CPU cycle market.

[4]Throughout the paper, *min*-operators are used due to their familiarity, although strictly speaking the value of such a *min*-operator may be undefined because $c_S(r_S)$ need not be continuous. Thus, to be precise, *inf*-operators should be used.

- The *domain problem*: tasks and resources of the agents. Among rational agents this is the only determining factor.
- The *execution architecture* on which the problem solving algorithm is run. Specifically, the architecture determines c_{comp}.
- The *problem solving algorithm*. We make no restrictive assumptions as to how effectively the algorithm uses the execution architecture. This is realistic because in practise it is often hard to construct algorithms that optimally (in some sense) use the architecture.

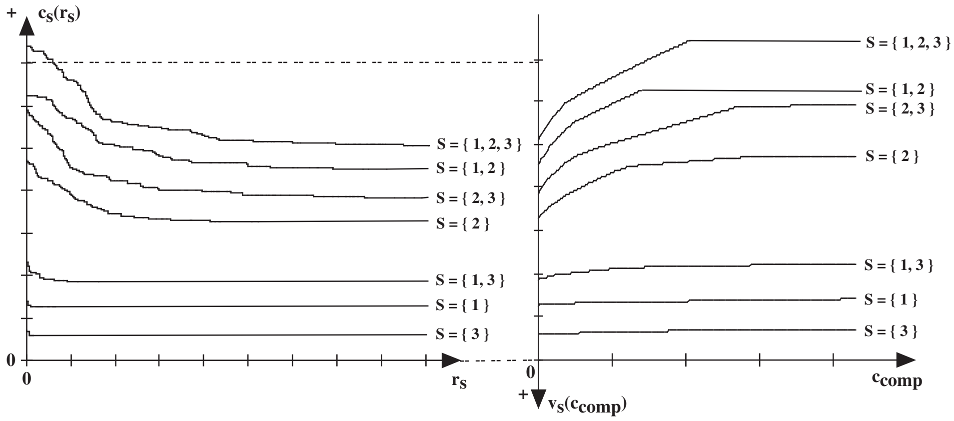

Figure 1: *Example experiment (from the vehicle routing domain) with agents 1, 2, and 3. Left: performance profiles, i.e. solution cost as a function of allocated computation resources. The curves become flat when the algorithm has reached a local optimum. Right: BR coalition value as a function of computation unit cost. The value of each coalition is negative because costs are positive. The curves become flat at a c_{comp} that is so high that it is not worth to take any iterative refinement steps: the initial solutions are used (their computation requirements are assumed negligible).*

Conceptually the agents use *design-to-time algorithms* [7, 29, 8]: 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 [21, 6, 4, 11, 29] 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.[5] Secondly, the PP depends on several features of the problem instance, and computing the mapping from the instance to the PP [21] 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 [19] 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

[5] 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 [21, 29]. Such conditioning, anytime algorithms, and their integration to coalition formation are part of our current research.

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 coalition as will be presented later.

In general, the value of a coalition may depend on the actions of nonmember 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 nonmembers. CFGs are a strict subset of NFGs. 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 [27]. 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 Section 5. There exist BRCFGs that are not CFGs. This is due to the fact that one can construct games where the domain cost of the *actual solution* (for any coalition) attained by the algorithm of a BR agent may be independent of the actions of nonmembers even though the domain cost of the *best solution* attained by a rational agent depends on the actions of nonmembers. For example, in some domains it is possible to restrict oneself to using algorithms

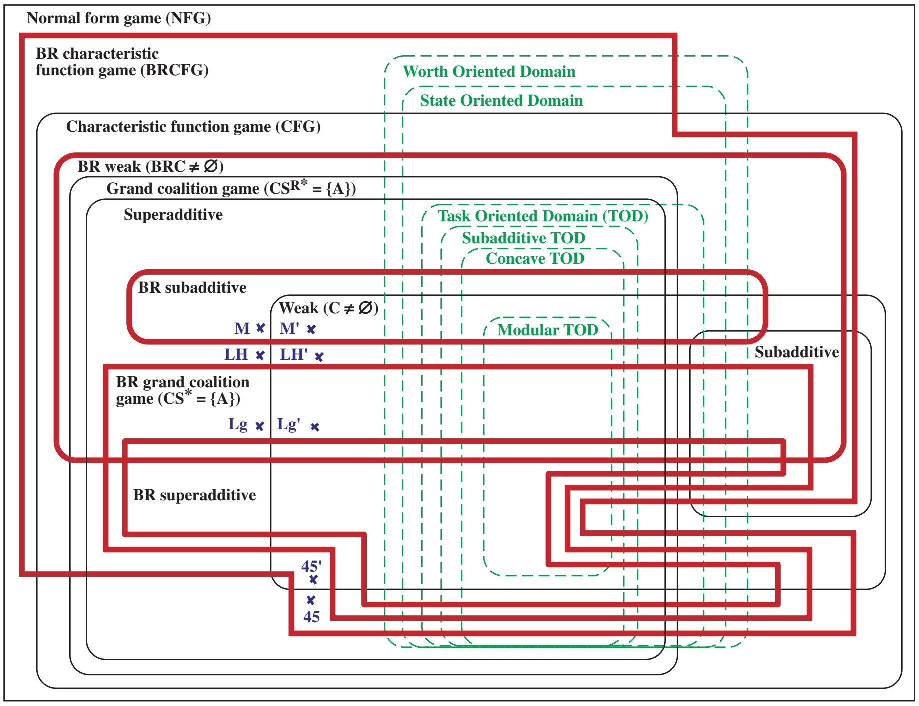

Figure 2: *Venn diagram of negotiation domains. Normal lines show the classification for rational agents. Bold lines show our new classification for BR agents, and how it relates to the rational case. Dotted lines show the rational agent domain classification of Rosenschein and Zlotkin [19]. They use "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.*

that only consider solutions whose value is not affected by nonmembers. There also exist CFGs that are not BRCFGs. For example, the agents may have different performance profiles and therefore the bounded rational value of a coalition may depend whether nonmembers are willing to do the computation for the coalition. There is also another reason why some CFGs are not BRCFGs. The algorithms that the agents use may produce solutions whose values depend on the actions of nonmembers although the value of the optimal solution would not.

The paper is organized as follows. Section 2 studies the optimal coalition structure for BR agents, and Section 3 analyzes 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 (BR) agents, these coalition values incorporate the computation costs.

A game is *superadditive* if 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 [6] *A game is* ***superadditive*** *if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v^R_{S \cup T} \geq v^R_S + v^R_T$. *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. A game can be non-superadditive only if the collusion process itself involves some cost, e.g. anti-trust penalties. All superadditive games are *grand coalition games*, i.e. the agents are best off—from a social welfare viewpoint—by forming the grand coalition ($CS^{R*} = \{A\}$). Some non-superadditive games are *subadditive*, Fig. 2:

Definition 2.2 *A game is* ***subadditive*** *if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v^R_{S \cup T} < v^R_S + v^R_T$.

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.

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

Now we present a new concept for BR agents that is analogous to superadditivity among rational agents. A game is *bounded rational superadditive* (BRS) if 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 unit cost* c_{comp} *if* $(\forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}(c_{comp}) \geq v_S(c_{comp}) + v_T(c_{comp})$.

Every BRS game is a *bounded rational grand coalition game*, Fig. 2. In such games, BR agents are best off—from a social welfare viewpoint—by forming the grand coalition ($CS^* = \{A\}$). BR superadditivity does not always coincide with superadditivity. In general, for a given c_{comp}, a game can be superadditive, BRS, both, or neither. Only some non-BRS games are BR subadditive, Fig. 2:

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

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 will denote any one of these best CSs by CS^*.

The rest of this section analyzes the relationship between the shape of the performance profiles and the class of the game.

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}.

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* $\forall c_{comp}$.

Proof. Let us analyze two arbitrary potential coalitions S and T, where $S, T \subseteq A$ and $S \cap T = \emptyset$. The conditions in the theorem state

$$\forall r_S \geq 0, \forall r_T \geq 0, c_{S \cup T}(r_S + r_T) \leq c_S(r_S) + c_T(r_T)$$

and obviously

$$\begin{aligned}&\exists r'_S, r'_T \geq 0 \text{ s.t. } c_S(r'_S) + c_{comp} \cdot r'_S + c_T(r'_T) + c_{comp} \cdot r'_T \\ &= \min_r [c_S(r) + c_{comp} \cdot r] + \min_r [c_T(r) + c_{comp} \cdot r]\end{aligned}$$

It follows that

$$\begin{aligned}& \exists r'_S, r'_T \geq 0 \text{ s.t. } c_{S\cup T}(r'_S + r'_T) + c_{comp} \cdot (r'_S + r'_T) \\ & \leq \min_r [c_S(r) + c_{comp} \cdot r] + \min_r [c_T(r) + c_{comp} \cdot r] \\ \Leftrightarrow \quad & \exists r' \geq 0 \text{ s.t. } c_{S\cup T}(r') + c_{comp} \cdot r' \\ & \leq \min_r [c_S(r) + c_{comp} \cdot r] + \min_r [c_T(r) + c_{comp} \cdot r] \\ \Leftrightarrow \quad & \min_r [c_{S\cup T}(r) + c_{comp} \cdot r] \leq \min_r [c_S(r) + c_{comp} \cdot r] + \min_r [c_T(r) + c_{comp} \cdot r] \\ \Leftrightarrow \quad & v_{S\cup T}(c_{comp}) \geq v_S(c_{comp}) + v_T(c_{comp})\end{aligned}$$

□

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 domain cost of these coalitions combined after allocating $r_S + r_T$. This is always achievable in theory 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. To be sure of BR superadditivity, the algorithm would need to solve each agent's problem separately—thus ensuring superadditivity trivially by additivity.

Usually, the algorithm that is used on the composite problem does not apply this type of problem decomposition. The real desideratum is not necessarily to generate algorithms that guarantee BR superadditivity (and thus the superiority of the grand coalition over other coalition structures), but algorithms that provide the highest social welfare (for the best coalition structure, which need not be the grand coalition). Sometimes these goals are conflicting. Whether the algorithm's PPs actually satisfy the conditions for BR superadditivity without using a decomposition method depends on the problem, the specific instances under study, and the algorithm itself.

In general, the game can be BRS $\forall c_{comp}$ 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* $\forall c_{comp}$.

Proof. **Counterexample.** Let us analyze a 2-agent game where $A = \{1, 2\}$. Let the performance profiles of the algorithms be

$$c_{\{1\}}(r) = c_{\{2\}}(r) = \begin{cases} \frac{1}{2} - \frac{1}{2}r & \text{if } 0 \leq r \leq 1 \\ 0 & \text{if } r > 1 \end{cases} \quad \text{and}$$

$$c_{\{1,2\}}(r) = \begin{cases} 1 & \text{if } 0 \leq r \leq 1 \\ 2 - r & \text{if } 1 < r \leq 2 \\ 0 & \text{if } r > 2 \end{cases}$$

Thus (see also Figure 3),

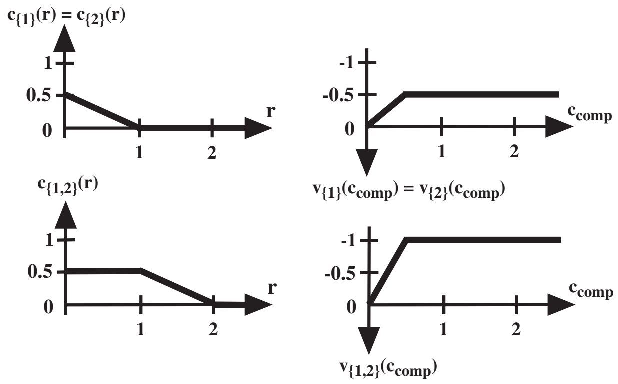

Figure 3: *Performance profiles and value functions of the counterexample.*

$$v_{\{1\}}(c_{comp}) = v_{\{2\}}(c_{comp}) = -\min_r[c_{\{2\}}(r) + c_{comp} \cdot r] = \begin{cases} -c_{comp} & \text{if } c_{comp} \leq \frac{1}{2} \\ -\frac{1}{2} & \text{if } c_{comp} > \frac{1}{2} \end{cases}$$

$$\text{and } v_{\{1,2\}}(c_{comp}) = -\min_r[c_{\{1,2\}}(r) + c_{comp} \cdot r] = \begin{cases} -2c_{comp} & \text{if } c_{comp} \leq \frac{1}{2} \\ -1 & \text{if } c_{comp} > \frac{1}{2} \end{cases}$$

So when $c_{comp} \leq \frac{1}{2}$,

$$v_{\{1,2\}}(c_{comp}) = -2c_{comp} = -c_{comp} + -c_{comp} = v_{\{1\}}(c_{comp}) + v_{\{2\}}(c_{comp})$$

and when $c_{comp} > \frac{1}{2}$,

$$v_{\{1,2\}}(c_{comp}) = -1 = -\frac{1}{2} + -\frac{1}{2} = v_{\{1\}}(c_{comp}) + v_{\{2\}}(c_{comp})$$

Thus, $(\forall c_{comp}, \forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}(c_{comp}) \geq v_S(c_{comp}) + v_T(c_{comp})$, i.e. the game is BRS for all c_{comp}. But $c_{\{1,2\}}(\frac{1}{2}+\frac{1}{2}) = 1 > \frac{1}{4}+\frac{1}{4} = c_{\{1\}}(\frac{1}{2})+c_{\{2\}}(\frac{1}{2})$. $\square$

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 [15, 20]) to using another refinement operator (e.g. 3-swap in TSP). Furthermore, refinements often decrease solution cost in a step-wise, noncontinuous manner rendering the PPs locally nonconvex—as in our experiments (Fig. 1 left). If the algorithm is stochastic, these step-related nonconvexities are reduced as the PP is averaged over multiple runs. 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 unit 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 $\forall c_{comp}$.*

The proof of Theorem 2.3 relies on the following Lemma:

Lemma 2.1 *Let $f(x)$ be a decreasing, convex function. For any given x^*, $\exists c \geq 0$ s.t.*

$$\min_x [f(x) + cx] = f(x^*) + cx^*$$

Proof. (Lemma 2.1). Let us define $x' = \operatorname{argmin}_x[f(x)+cx]$. Assume—for contradiction—that $\exists x^*$ s.t. $\forall c \geq 0$,

$$\begin{aligned}
\min_x[f(x)+cx] &\neq f(x^*)+cx^* \\
\Leftrightarrow f(x')+cx' &\neq f(x^*)+cx^*
\end{aligned}$$

Because $f(x)$ is convex,

$$\begin{aligned}
f(x^*) &\leq \frac{f(x^*-\delta)+f(x^*+\delta)}{2} \\
\Rightarrow \lim_{\delta\to 0}\frac{f(x^*)-f(x^*-\delta)}{\delta} &\leq \lim_{\delta\to 0}\frac{f(x^*+\delta)-f(x^*)}{\delta}
\end{aligned}$$

Thus $c \geq 0$ is well-defined when chosen as follows:

$$\lim_{\delta\to 0}\frac{f(x^*)-f(x^*-\delta)}{\delta} \leq -c \leq \lim_{\delta\to 0}\frac{f(x^*+\delta)-f(x^*)}{\delta}$$

Now there are two cases:
Case 1: $x' < x^*$:

$$\begin{aligned}
x' &< x^* \\
\Leftrightarrow \operatorname*{argmin}_x[f(x)+cx] &< x^* \\
\Leftrightarrow f(\operatorname*{argmin}_x[f(x)+cx]) + c\cdot\operatorname*{argmin}_x[f(x)+cx] &< f(x^*)+cx^* \\
\Leftrightarrow f(x')+cx' &< f(x^*)+cx^* \\
\Leftrightarrow f(x^*-\epsilon)+c\cdot(x^*-\epsilon) &< f(x^*)+cx^* \\
\Leftrightarrow f(x^*)-f(x^*-\epsilon) &> -c\epsilon \\
\Leftrightarrow \frac{f(x^*)-f(x^*-\epsilon)}{\epsilon} &> -c \\
\Rightarrow \frac{f(x^*)-f(x^*-\epsilon)}{\epsilon} &> \lim_{\delta\to 0}\frac{f(x^*)-f(x^*-\delta)}{\delta}
\end{aligned}$$

This violates convexity. Contradiction.
Case 2: $x' > x^*$:

$$\begin{aligned}
x' &> x^* \\
\Leftrightarrow \operatorname*{argmin}_x[f(x)+cx] &> x^*
\end{aligned}$$

$$\begin{aligned}
\Leftrightarrow f(\operatorname*{argmin}_x[f(x)+cx]) + c\cdot\operatorname*{argmin}_x[f(x)+cx] &< f(x^*)+cx^* \\
\Leftrightarrow f(x')+cx' &< f(x^*)+cx^* \\
\Leftrightarrow f(x^*+\epsilon)+c\cdot(x^*+\epsilon) &< f(x^*)+cx^* \\
\Leftrightarrow \frac{f(x^*+\epsilon)-f(x^*)}{\epsilon} &< -c \\
\Rightarrow \frac{f(x^*+\epsilon)-f(x^*)}{\epsilon} &< \lim_{\delta\to 0}\frac{f(x^*+\delta)-f(x^*)}{\delta}
\end{aligned}$$

This also violates convexity. Contradiction. Because both cases lead to a contradiction, the original assumption is false. □

Proof. (Theorem 2.3). The *if*-part was proven in Theorem 2.1. Now the *only if*-part is proven.

$$\begin{aligned}
& \text{Game is BRS } \forall c_{comp} \\
\Leftrightarrow\quad & (\forall c_{comp}, \forall S,T\subseteq A, S\cap T=\emptyset), v_{S\cup T}(c_{comp}) \geq v_S(c_{comp}) + v_T(c_{comp}) \\
\Leftrightarrow\quad & (\forall c_{comp}, \forall S,T\subseteq A, S\cap T=\emptyset), \\
& \min_r[c_{S\cup T}(r)+c_{comp}\cdot r] \leq \min_r[c_S(r)+c_{comp}\cdot r] + \min_r[c_T(r)+c_{comp}\cdot r] \\
\Leftrightarrow\quad & (\forall c_{comp}, \forall S,T\subseteq A, S\cap T=\emptyset, \forall r_S, r_T \geq 0), \\
& \min_r[c_{S\cup T}(r)+c_{comp}\cdot r] \leq c_S(r_S)+c_{comp}\cdot r_S + c_T(r_T)+c_{comp}\cdot r_T
\end{aligned}$$

Now, by Lemma 2.1, for any $r_S + r_T \geq 0$, $\exists c_{comp} \geq 0$ s.t. $\min_r[c_{S\cup T}(r) + c_{comp}\cdot r] = c_{S\cup T}(r_S+r_T) + c_{comp}\cdot(r_S+r_T)$. Thus,

$$\begin{aligned}
& (\forall S,T\subseteq A, S\cap T=\emptyset, \forall r_S, r_T\geq 0, \exists c_{comp}\geq 0), \\
& c_{S\cup T}(r_S+r_T) + c_{comp}\cdot(r_S+r_T) \\
& \leq c_S(r_S)+c_{comp}\cdot r_S + c_T(r_T) + c_{comp}\cdot r_T \\
\Leftrightarrow\quad & (\forall S,T\subseteq A, S\cap T=\emptyset, \forall r_S, r_T\geq 0, \exists c_{comp}\geq 0), \\
& c_{S\cup T}(r_S+r_T) \leq c_S(r_S) + c_T(r_T) \\
\Leftrightarrow\quad & (\forall S,T\subseteq A, S\cap T=\emptyset, \forall r_S, r_T\geq 0), c_{S\cup T}(r_S+r_T) \leq c_S(r_S)+c_T(r_T)
\end{aligned}$$

□

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 unit 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* $\forall c_{comp}$.

Proof.

$$
\begin{array}{ll}
 & (\forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S, r_T \geq 0), \\
 & c_{S \cup T}(r_S + r_T) > c_S(r_S) + c_T(r_T) \\
\Leftrightarrow & (\forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S, r_{S \cup T} \geq 0), \\
 & c_{S \cup T}(r_{S \cup T}) > c_S(r_S) + c_T(r_{S \cup T} - r_S) \\
\Leftrightarrow & (\forall c_{comp}, \forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S, r_{S \cup T} \geq 0), \\
 & c_{S \cup T}(r_{S \cup T}) + c_{comp} \cdot r_{S \cup T} \\
 & > c_S(r_S) + c_{comp} \cdot r_S + c_T(r_{S \cup T} - r_S) + c_{comp} \cdot (r_{S \cup T} - r_S) \\
\Rightarrow & (\forall c_{comp}, \forall S, T \subseteq A, S \cap T = \emptyset, \forall r_S, r_{S \cup T} \geq 0), \\
 & \min_r [c_{S \cup T}(r) + c_{comp} \cdot r] \\
 & > c_S(r_S) + c_{comp} \cdot r_S + c_T(r_{S \cup T} - r_S) + c_{comp} \cdot (r_{S \cup T} - r_S) \\
 & \geq \min_r [c_S(r) + c_{comp} \cdot r] + \min_r [c_T(r) + c_{comp} \cdot (r)] \\
\Leftrightarrow & (\forall c_{comp}, \forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}(c_{comp}) < v_S(c_{comp}) + v_T(c_{comp}) \\
\Leftrightarrow & \text{Game is bounded rational subadditive } \forall c_{comp}
\end{array}
$$

□

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 [12, 27, 18]. 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 [12, 27, 18]. 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 [12, 27, 18]. 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 unit 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 $c_{comp} \Rightarrow BRC(c_{comp}) \neq \emptyset$.*

Proof. Let us analyze a game that is bounded rational subadditive for some c_{comp}, i.e. $(\forall S, T \subseteq A, S \cap T = \emptyset), v_{S \cup T}(c_{comp}) < v_S(c_{comp}) + v_T(c_{comp})$. Let us study a coalition structure $CS^* = \{\{1\}, \{2\}, ..., \{|A|\}\}$. Let us choose $\vec{x}$ s.t. $\forall i \in A$, $x_i = v_{\{i\}}(c_{comp})$. Now,

$$\sum_{i \in A} x_i = \sum_{i \in A} v_{\{i\}}(c_{comp}) = \sum_{j \in CS^*} v_{S_j}(c_{comp})$$

and

$$\forall S \subset A, \sum_{i \in S} x_i = \sum_{i \in S} v_{\{i\}}(c_{comp}) \geq v_S(c_{comp})$$

Thus $\vec{x} \in BRC(c_{comp})$ which implies $BRC(c_{comp}) \neq \emptyset$. □

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 [24, 5]. 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 $\mathcal{B} = \{B_1, ..., B_p\}$ 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 $\mathcal{B} = \{B_1, ..., B_p\}, \sum_{j=1}^{p} \lambda_j v_{B_j}(c_{comp}) \leq v_A(c_{comp})$.*

Proof. Shapley [24] proved the following fact (his Theorem 2) for rational agents. In games where $CS^{R*} = \{A\}$, $C \neq \emptyset$ iff for every minimal balanced set $\mathcal{B} = \{B_1, ..., B_p\}, \sum_{j=1}^{p} \lambda_j v_{B_j}^R \leq v_A^R$. Theorem 3.2 follows by analogy. □

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 $CS^* = \{A\}$.

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*

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).*

$\mathcal{B} = \{B_1, ..., B_p\}, \sum_{j=1}^{p} \lambda_j v_{B_j}(c_{comp}) \leq v_A(c_{comp})$. *Furthermore, this set of inequalities is minimal: no smaller set is sufficient.*

Proof. Shapley [24] proved the following fact (his Theorem 3) for rational agents. In a superadditive game, $C \neq \emptyset$ iff for every proper minimal balanced set $\mathcal{B} = \{B_1, ..., B_p\}, \sum_{j=1}^{p} \lambda_j v_{B_j}^R \leq v_A^R$. Charnes and Kortanek [5] proved that this set of inequalities is minimal. Theorem 3.3 follows by analogy. □

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}, *[for every minimal balanced set* $\mathcal{B} = \{B_1, ..., B_p\}$, $(\forall B \in \mathcal{B}, \forall r_B \geq 0) \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) \geq c_A(\sum_{j=1}^{p} \lambda_j r_{B_j})] \Rightarrow$ $BRC(c_{comp}) \neq \emptyset$.

Proof. Let us analyze an arbitrary minimal balanced set $\mathcal{B} = \{B_1, ..., B_p\}$.

$$
\begin{aligned}
& (\forall B \in \mathcal{B}, \forall r_B \geq 0), \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) \geq c_A(\sum_{j=1}^{p} \lambda_j r_{B_j}) \\
\Rightarrow \quad & (\forall c_{comp}, \forall B \in \mathcal{B}, \forall r_B \geq 0, \exists r_A \geq 0), \\
& \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) + c_{comp} \cdot (-r_A + \sum_{j=1}^{p} \lambda_j r_{B_j}) \geq c_A(r_A) \\
\Leftrightarrow \quad & (\forall c_{comp}, \forall B \in \mathcal{B}, \forall r_B \geq 0, \exists r_A \geq 0), \\
& \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) + c_{comp} \cdot \sum_{j=1}^{p} \lambda_j r_{B_j} \geq c_A(r_A) + c_{comp} \cdot r_A \\
\Leftrightarrow \quad & (\forall c_{comp}, \forall B \in \mathcal{B}, \forall r_B \geq 0), \\
& \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) + c_{comp} \cdot \sum_{j=1}^{p} \lambda_j r_{B_j} \geq \min_r [c_A(r) + c_{comp} \cdot r] \\
\Leftrightarrow \quad & (\forall c_{comp}, \forall B \in \mathcal{B}, \forall r_B \geq 0), \\
& \sum_{j=1}^{p} \lambda_j [c_{B_j}(r_{B_j}) + c_{comp} \cdot r_{B_j}] \geq \min_r [c_A(r) + c_{comp} \cdot r] \\
\Leftrightarrow \quad & (\forall c_{comp}), \\
& \sum_{j=1}^{p} \lambda_j \min_{r'_{B_j}} [c_{B_j}(r'_{B_j}) + c_{comp} \cdot r'_{B_j}] \geq \min_r [c_A(r) + c_{comp} \cdot r] \\
\Leftrightarrow \quad & (\forall c_{comp}), \\
& \sum_{j=1}^{p} \lambda_j v_{B_j}(c_{comp}) \leq v_A(c_{comp})
\end{aligned}
$$

Since this holds for an arbitrary minimal balanced set, it has to hold for every minimal balanced set. Thus, by Theorem 3.2, $BRC(c_{comp}) \neq \emptyset$. □

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 $\mathcal{B} = \{B_1, ..., B_p\}, (\forall B \in \mathcal{B}, \forall r_B \geq 0) \sum_{j=1}^{p} \lambda_j c_{B_j}(r_{B_j}) \geq c_A(\sum_{j=1}^{p} \lambda_j r_{B_j})] \Rightarrow BRC(c_{comp}) \neq \emptyset$.*

Proof. Analogous to the proof of Theorem 3.4, except that now an arbitrary *proper* minimal balanced set is considered. Furthermore, the reference to Theorem 3.2 should be changed to a reference to Theorem 3.3. □

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 $\forall c_{comp}$, $[(\forall r_{\{1,2\}} \geq 0, \forall r_{\{1,3\}} \geq 0, \forall r_{\{2,3\}} \geq 0), \frac{1}{2}c_{\{1,2\}}(r_{\{1,2\}}) + \frac{1}{2}c_{\{1,3\}}(r_{\{1,3\}}) + \frac{1}{2}c_{\{2,3\}}(r_{\{2,3\}}) \geq c_{\{1,2,3\}}(\frac{1}{2}r_{\{1,2\}} + \frac{1}{2}r_{\{1,3\}} + \frac{1}{2}r_{\{2,3\}})] \Rightarrow \forall c_{comp}, BRC(c_{comp}) \neq \emptyset$.

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 [16] are unusable.

The rational value (v_S^R) of each coalition S is defined by the tasks and the resources (vehicles, depots) of the agents in the coalition. Specifically, v_S^R is independent of how nonmembers solve their optimization problems. Therefore our problem is a characteristic function game (CFG), Fig. 2.

Our problem is outside the domain classification of Rosenschein and Zlotkin [19], Fig. 2, because agents do not have symmetric capabilities due to heterogeneous fleets. If their definition were extended to allow asymmetric capabilities, our domain would be in SOD \ TOD. Our domain would not be a TOD because any one agent is not necessarily able to individually handle all tasks of all agents. If we further dropped the maximum route length constraint (this experiment will also be presented), and restricted 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 would be a TOD. The following simple example shows that it would not be a "Subadditive TOD" because the depots are geographically distributed. Let us look at a game with just two agents (A1 and A2), two delivery tasks (T1 and T2), and two identical vehicles—one for each agent. Say that the pickup site and the drop-off site of T1 are close to A1's depot, and T2's pickup and drop-off are close to A2's depot. Now say that the depots are far from each other. Thus the sum of the route lengths when A1 manages T1 and A2 manages T2 is lower than when either agent individually manages both tasks.

To analyze 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 analyzed all of the $\binom{5}{3} = 10$ 3-agent games that can be acquired by choosing 3 of the 5 dispatch centers. 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 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. 4. 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. In some of these mixed games, for low c_{comp}, the grand coalition was the best coalition structure (point Lg in Fig. 2). Existence of the core for rational agents is unknown for our games: the points M, LH, and Lg might really be M', LH', and Lg'. 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 low c_{comp}'s), which are always stable. We also reran the experiments without the maximum route length restriction, and these results prevailed, Fig. 4.

Centers 2, 3 and 5 were located near 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.

Next we analyzed 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

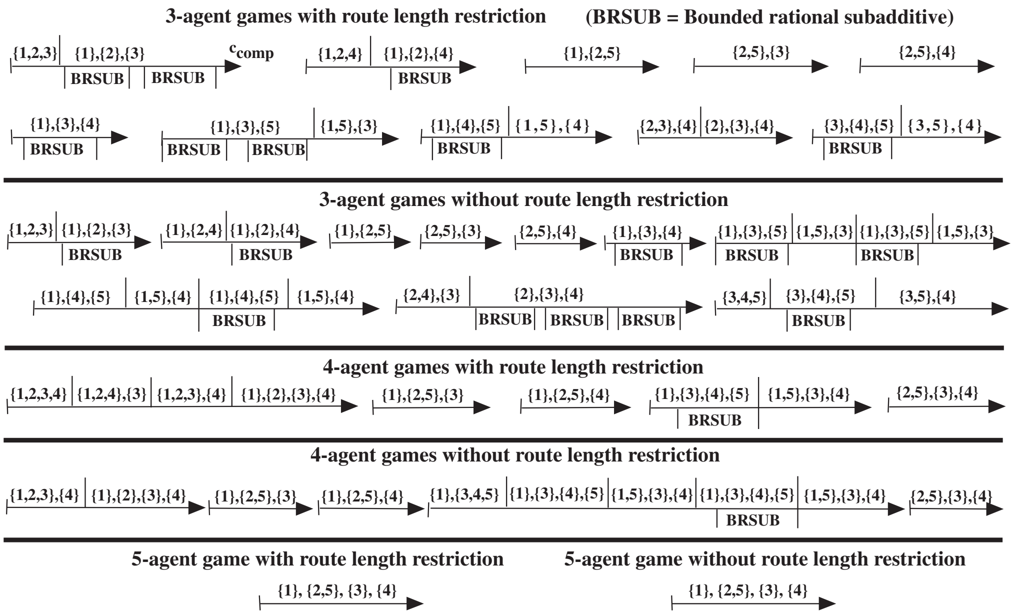

Figure 4: *Optimal coalition structure (CS^*) and bounded rational subadditivity as a function of c_{comp}. Tested by evaluating all possible coalition structures and super/subadditivity at varying points of c_{comp} chosen from a grid where c_{comp} is always incremented by 1%.*

largest values of c_{comp}. No game was BRS for any c_{comp}, but some games were bounded rational subadditive for interior values, Fig. 4. In only one game (with agents 1, 2, 3, and 4, and the route length restriction), for low c_{comp}, the best coalition structure was the grand coalition. When this occurred, $BRC(c_{comp})$ happened to be non-empty (point Lg (or Lg') in Fig. 2). In none of the experiments was the $BRC(c_{comp})$ empty when the best coalition structure was the grand coalition. Thus, depending on c_{comp}, the games were at the points M, LH, Lg, or 45 (or M', LH', Lg', 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 unit 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. These intuitions suggest that with refinement steps of superlinear complexity, *higher computation unit costs often promote smaller coalitions, and lower computation unit costs promote larger ones.* Thus it was not surprising that in games where the grand coalition was optimal, it was optimal for very small computation unit 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 unit costs often 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 unit 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. 4.

5 Different performance profiles, different computation unit costs, and domain solution interactions

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 nonmembers. 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.

Games where the agents have different unit costs (c_{comp}'s) for computation–e.g. due to different execution architectures—are also in general not BRCFGs. Actually such games are analogous to games with a global c_{comp} but agents with different PPs. Namely, games where agents have different computation unit costs (c_{comp}'s) can be modeled as games with a uniform computation unit cost after the c_{comp}-axis of each $v_S(c_{comp})$ function is appropriately rescaled based on the real c_{comp} of the corresponding coalition S.

Interactions between domain solutions of different coalitions may also exclude some problems from the class BRCFG. In general, the rational value of a coalition may depend on the actions of nonmember agents due to positive and negative interactions of the agents' solutions. Such games are normal form games (NFGs), but not characteristic function games (CFGs), Fig. 2. For the same reason, the value of some BR coalition's domain solution—computed by a BR agent—may depend on the actions of nonmembers. Negative interactions are often caused by shared resources of finite capacity. Once nonmembers are using the resource to a certain extent, not enough of that resource is available to agents in the coalition to carry out the planned solution at the minimum cost. Negative interactions can also be caused by conflicting goals. In satisfying their goals, nonmembers may actually move the world further from the coalition's goal state(s) [19]. Positive interactions are often caused by partially overlapping goals. In satisfying their goals, nonmembers may actually move the world closer to the coalition's goal state(s), from where the coalition can reach its goals less expensively than it could have without the actions of nonmembers.

In the distributed vehicle routing domain of this paper, there were no shared resources, because all of the resources—vehicles and depots—in the domain were exclusively and exhaustively distributed among the agents (and thus among coalitions). Secondly, each agent (and thus each coalition) had its own goal: delivering all of the parcels at the lowest possible cost. The deliveries of one coalition are unaffected by the deliveries of nonmember agents. Thus, as stated earlier, the domain is a CFG. For the same reason, domain solution interactions do not preclude the problem from belonging to the class BRCFG. Yet if the agents had different PPs or computation unit costs, the problem would not necessarily be within BRCFG.

In non-CFGs, superadditivity, subadditivity, and the core are undefined, Fig. 2. Thus, other solution concepts are necessary. One alternative is the *Nash equilibrium* [17, 14], which guarantees stability in the sense that no agent alone is motivated to deviate from the solution given that others in the game do not deviate. Often this solution concept is too weak because subgroups of agents can deviate in a coordinated manner. The *Strong Nash equilibrium* [1] is a solution concept that guarantees more stability in the sense that it requires that there is no subgroup that can deviate in a manner that increases the payoff of all of its members given that nonmembers do not deviate from the original solution. The Strong Nash equilibrium is often too strong a solution concept because in many games no such equilibria exist. Recently, the *Coalition-Proof Nash equilibrium* [2, 3] has been suggested to remedy this problem. This solution concept requires that there is no subgroup that can make a mutually beneficial deviation (keeping the strategies of nonmembers fixed) *in a way that the deviation itself is stable according to the same criterion.* A problem with this solution concept is that the deviation may be stable within the deviating group, but the solution concept ignores the possibility that some of the agents that deviated may prefer to deviate again with agents that did not originally deviate. Clearly, there is room for further research on coalition formation even among rational agents.

Similar problems arise with BR agents. In non-BRCFGs, BR superadditivity, BR subadditivity, and the BRC are undefined, Fig. 2. Again, other solution concepts are necessary, e.g. the Nash equilibrium or some of its refinements. This is part of our current research.

6 Related DAI research on collusion

Coalition formation has been widely studied in game theory [12, 2, 3, 1, 27, 18]; only the most relevant concepts were presented here. This section compares our work to other recent DAI research on coalition formation.

Zlotkin and Rosenschein [30] 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. In STODs this is optimal 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 (symmetric cost functions for task sets). Their method guarantees each agent an expected value that equals its Shapley value [12, 18]. The Shapley value motivates individual agents to stay with the coalition structure (individual rationality) and the group of all agents to stay (group rationality). Unlike the core, the Shapley value does not in general motivate every subgroup of agents to stay with the coalition structure (coalition rationality). In a subset of STODs, "Concave Task Oriented Domains" (Fig. 2), the Shapley value also satisfies coalition rationality, i.e. the vector of Shapley value payoffs is in the core.

A naive method that guarantees an expected value equal to the Shapley value has exponential complexity in the number of agents, but Zlotkin and Rosenschein present a novel cryptographic method for achieving this with linear complexity in the number of agents. Yet each one of these linearly many problems involving the agents' tasks needs to be solved optimally. In combinatorial problems such as the vehicle routing problem of this paper (and the Postmen Domain of Zlotkin and Rosenschein for that matter), this is clearly intractable if the problem instances are large.

Ketchpel [13] 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 [20] (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 [28], 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 [25] 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, domain-independent theory of coalitions in combinatorial domains was presented, where the rationality of self-interested agents is bounded by computational complexity. This work is an extension of game theory, which classically assumes perfect rationality: algorithms that find the optimal solution, and zero computation unit cost.

A domain classification was presented for rational and bounded rational (BR) agents. The algorithms used by the agents significantly impact the coalition structure that should form as well as its stability. General theorems were presented that relate an algorithm's performance profiles (PPs) to the social welfare maximizing coalition structure. This analysis was carried out using the new concepts of BR superadditivity and BR subadditivity. General theorems were also presented that relate the PPs to the non-emptiness of the bounded rational core (BRC), which determines the stability of the coalition structure.

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 coalition structure for BR agents varied although rational agents should always form the grand coalition. Section 2 developed conditions on the PPs that guarantee

BR superadditivity. It also discussed a separate solving approach—based on a problem decomposition step—that guarantees that the base algorithm fulfills those conditions. With our reasonable deterministic iterative refinement algorithm, these conditions were—somewhat surprisingly—never met. The real desideratum is not necessarily to generate algorithms that guarantee BR superadditivity (and thus the superiority of the grand coalition over other coalition structures), but algorithms that provide the highest social welfare (for the best coalition structure, which need not be the grand coalition). Sometimes these goals are conflicting.

The observed BR subadditivity of some of the games implies a non-empty BRC: the best coalition structure in those games is stable. Even when BR subadditivity did not hold, the BRC was often non-empty—especially for large computation unit costs c_{comp}. Often with superlinear iterative refinement steps, low c_{comp} promotes large coalitions while high c_{comp} suggests smaller ones. The best BR coalition structures mostly agreed with our intuitions of what coalitions should form among rational agents 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 [21, 29]. 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. Finally, we are in the process of developing interaction protocols [23] that efficiently guide self-interested agents towards the optimal and stable (whenever possible) coalition structures—as determined by the theory developed in this paper.

References

[1] R. Aumann. Acceptable points in general cooperative n-person games. volume IV of *Contributions to the Theory of Games*. Princeton University Press, 1959.

[2] B. D. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof nash equilibria: 1. concepts. *Journal of Economic Theory*, 42(1):1–12, June 1987.

[3] B. D. Bernheim and M. D. Whinston. Coalition-proof nash equilibria: 2. applications. *Journal of Economic Theory*, 42(1):13–29, June 1987.

[4] M. Boddy and T. Dean. Solving time-dependent planning problems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 979–984, Detroit, MI, Aug. 1989.

[5] 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.

[6] T. Dean and M. Boddy. An analysis of time-dependent planning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 49–54, St. Paul, MN, Aug. 1988.

[7] A. Garvey and V. Lesser. Design-to-time real-time scheduling. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(6), 1993.

[8] A. Garvey and V. Lesser. A survey of research in deliberative real-time artificial intelligence. *Real-Time Systems*, 6:317–347, 1994.

[9] General Magic, Inc. Telescript technology: The foundation for the electronic marketplace, 1994. White paper.

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

[11] E. J. Horvitz. Reasoning about beliefs and actions under computational resource constraints. In L. Kanal, T. Levitt, and J. Lemmer, editors, *Uncertainty in Artificial Intelligence*, volume 3, pages 301–324. 1989.

[12] J. P. Kahan and A. Rapoport. *Theories of Coalition Formation.* Lawrence Erlbaum Associates Publishers, 1984.

[13] S. Ketchpel. Forming coalitions in the face of uncertain rewards. In *Proceedings of the National Conference on Artificial Intelligence*, pages 414–419, Seattle, WA, July 1994.

[14] D. M. Kreps. *A course in microeconomic theory.* Princeton University Press, 1990.

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

[16] M. G. Lundgren, K. Jörnsten, and P. Värbrand. On the nucleolus of the basic vehicle routing game. Technical Report 1992-26, Linköping Univ., Dept. of Mathematics, Sweden, 1992.

[17] J. Nash. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences*, 36:48–49, 1950.

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

[19] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter.* MIT Press, 1994.

[20] T. W. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. 11th National Conference on Artificial Intelligence (AAAI-93)*, pages 256–262, July 1993.

[21] T. W. Sandholm and V. R. Lesser. Utility-based termination of anytime algorithms. In *ECAI Workshop on Decision Theory for DAI Applications*, pages 88–99, Amsterdam, The Netherlands, 1994. Extended version: Univ. of Mass. at Amherst, Comp. Sci. Tech. Report 94-54.

[22] T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. In *Proc. 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 662–669, Montreal, Canada, Aug. 1995.

[23] T. W. Sandholm and V. R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proc. First International Conference on Multiagent Systems (ICMAS-95)*, pages 328–335, San Francisco, June 1995.

[24] L. S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14:453–460, 1967.

[25] O. Shechory and S. Kraus. Feasible formation of stable coalitions among autonomous agents in general environments. *Computational Intelligence Journal*, 1995. Submitted.

[26] H. A. Simon. *Models of bounded rationality*, volume 2. MIT Press, 1982.

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

[28] M. Wellman. A general- equilibrium approach to distributed transportation planning. In *Proc. 10th National Conference on Artificial Intelligence (AAAI-92)*, pages 282–289, San Jose, CA, July 1992.

[29] S. Zilberstein. *Operational rationality through compilation of anytime algorithms*. PhD thesis, University of California, Berkeley, 1993.

[30] G. Zlotkin and J. S. Rosenschein. Coalition, cryptography and stability: Mechanisms for coalition formation in task oriented domains. In *Proceedings of the National Conference on Artificial Intelligence*, pages 432–437, Seattle, WA, July 1994.