Coalition Structure Generation in Multi-Agent Systems with Mixed Externalities

Bikramjit Banerjee
School of Computing
The University of Southern Mississippi
118 College Dr. # 5106
Hattiesburg, MS 39406, USA
Bikramjit.Banerjee@usm.edu

Landon Kraemer
School of Computing
The University of Southern Mississippi
118 College Dr. # 5106
Hattiesburg, MS 39406, USA
Landon.Kraemer@eagles.usm.edu

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Algorithms, Theory, Performance

Keywords

Coalition Structure Generation, Branch and Bound

ABSTRACT

Coalition structure generation (CSG) for multi-agent systems is a well-studied problem. A vast majority of the previous work and the state-of-the-art approaches to CSG assume a characteristic function form of the coalition values, where a coalition's value is independent of the other coalitions in the coalition structure. Recently, there has been interest in the more realistic *partition function* form of coalition values, where the value of a coalition is affected by how the other agents are partitioned, via *externalities*. We argue that in domains with externalities, a distributed/adaptive approach to CSG may be impractical, and that a centralized approach to CSG is more suitable. However, the most recent studies in this direction have focused on cases where all externalities are either *always positive or always negative*, and results on coalition structure generation in more general settings (in particular, *mixed externalities*) are lacking. In this paper we propose a framework based on *agent-types* that incorporates mixed externalities and demonstrate that it includes the previous settings as special cases. We also generalize some previous results in anytime CSG, showing that those results are again special cases. In particular, we extend the existing branch and bound algorithm to this new setting and show empirically that significant pruning can be achieved when searching for the optimal coalition structure. This extends the state-of-the-art in CSG for multi-agent systems.

1. INTRODUCTION

Coalition formation is an important problem in multi-agent systems, where several intelligent, autonomous computational agents must be partitioned into teams so that their global utility –the sum of the team utilities – is maximized. Many practical applications can be modeled by this problem, for instance, customer coalitions in electronic marketplaces to extract group discounts [16], node coalitions in grid computing for coordinated resource sharing [3], distribution of targets in sensor networks [1], and in distributed task allocation [13] in general.

An important subproblem in coalition formation is optimal coalition structure generation (CSG) [13, 12, 2, 10]. In this paper we do not address the other subproblems viz., performance optimization and payoff division [12]. Coalition structure generation has been predominantly studied in the characteristic form where the value of any coalition depends on that coalition alone. However, in practice, the value of a coalition may additionally depend on the other coalitions in the coalition structure. In economic theory, these dependencies are referred to as *externalities* [4]. For an example on a global scale, consider the coalition C, a manufacturing firm, and the complementary coalition C' of everyone else. The decisions made by C could affect the level of pollution – an effect that C' must also share. Therefore, it can be argued that the value of C' depends partly on the formation of the coalition C, and not just on the intrinsic valuation of C' itself. To see that a coalition may impact different other coalitions differently, consider the recently attempted merger of Yahoo with Microsoft. This coalition would have impacted Google and its partners in a different manner (possibly negatively) than it would have impacted, say, a coalition of stock traders on the Wall Street (possibly positively). When the goal is to form the optimal coalition structure in an agent-based system, such externalities must be taken into account to reflect the true value of a coalition structure. Until recently, CSG algorithms did not take externalities into account, and therefore, could return coalition structures where the individual coalitions may be intrinsically high-valued, but that impact each other negatively, leading to a poor coalition structure.

Recently, externalities have been taken into account for CSG [7, 9], under certain restrictions (that the externalities are either always positive or always negative) which we seek to relax in this paper. But we also seek to emphasize a necessary departure from the traditional adaptive coalition formation process in multi-agent systems with externalities. Even though agents do not possess global information, in domains without externalities they can still form the optimal (or satisficing [15]) coalition structure through repeated interaction, communication/negotiation, and adaptation [14]. This is because the valuation of a coalition, when discovered by its constituents, is independent of the behavior of the other agents. But in domains with externalities, autonomous agents with limited local view may find that reaching a "good" coalition structure is an extremely difficult proposition, since any local structure that a group

Cite as: Coalition Structure Generation in Multi-Agent Systems with Mixed Externalities, Bikramjit Banerjee and Landon Kraemer, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 175-182

of agents might discover would have differing (and an intractable number of) valuations depending on how the other agents partition themselves. A coalition structure that looks "good" in the local context may be arbitrarily "poor" in the global context. In an agent based economy with collaborators and competitors, we believe that the presence of externalities makes it extremely inefficient for agents to attempt coalition formation in an adaptive or distributed way, with the potential for a disastrous beginning. Even in human economies, laws (e.g., anti-trust laws) need to be continually evolved to encourage appropriate partnerships; therefore it is unreasonable to expect that the rules of encounter in agent systems will be perfect from the very beginning. In such a circumstance, inappropriate coalition structures (or those evolving from poor initial structures) can be arbitrarily damaging to agent economies, at least in the beginning.

Against this backdrop, we emphasize a centralized approach to computing the optimal coalition structure in domains with externalities, based on estimates of the worth of individual coalitions and their impacts on each other, as a more viable alternative. Even if the coalition values are based on inaccurate information, the solution thus produced might provide a strong basis for the agents to *initiate* their activities in a fruitful way, as they negotiate and realign their coalition structure with new knowledge, experience and goals over time.

In this paper, we develop a framework for mixed externalities (i.e., where positive and negative externalities can co-exist, as in the Yahoo/Microsoft example) based on the notions of competition and complementation, and show that this framework includes the previous settings as special cases. We also generalize some previous results in this new framework. In particular, we extend the existing branch and bound algorithm to this new setting and show empirically that significant pruning can be achieved when searching for the optimal coalition structure. In the next section we present the basic notations and definitions that will be used to formalize our framework.

2. PRELIMINARIES AND DEFINITIONS

Let A be a set of n agents. A *coalition*, C, is a non-empty subset of A, i.e., $C \in 2^A \setminus \emptyset$. A *non-overlapping coalition structure*, CS, is a set of coalitions, $CS = \{C_1, C_2, \ldots C_k\}$, subject to the constraints

- $\bigcup_{j=1}^{k} C_j = A$,
- $C_p \bigcap C_j = \emptyset \; \forall p, j = 1 \ldots k, p \neq j$

In a setting with externalities, the value of a coalition C depends on the other coalitions in the coalition structure, and hence, is specified in context of the coalition structure as $v(C, CS)$. The problem of optimal CSG in such a setting seeks a coalition structure CS^*, such that

$$V(CS^*) = \sum_{C \in CS^*} v(C, CS^*)$$

is optimized. We represent the set of all partitions of the agent set $S \subseteq A$ as $\mathcal{P}(S)$, whereby $\mathcal{P}(A)$ is the set of all possible coalition structures. We also use v for values of coalitions, and V for the values of partitions.

The following definition specifies two kinds of partitions of any set of agents S, that will be used to define our type-based setting in section 4 and again in section 6.1 for the search algorithm.

Definition 1. *Given any set of agents, S, we define two specific partitions of S: $P_1(S) \in \mathcal{P}(S)$ where all agents are in singleton coalitions, and $P_\infty(S) \in \mathcal{P}(S)$ where all agents in S are in a single coalition. That is*

$$P_1(S) = \{\{a_i\} | a_i \in S\}$$

$$P_\infty(S) = \{S\}$$

The following defines a function α that aggregates agents out of any partition, and will be used for notational convenience throughout the paper.

Definition 2. *Given any partition, P, of any set of agents, we define $\alpha(P)$ as simply the set of all agents that can be found in any coalition within P. That is*

$$\alpha(P) = \bigcup_{p \in P} p = \{a_i \in A | \exists p \in P : a_i \in p\}.$$

The following definition explicates the externality function E for both coalitions and partitions, on which our setting crucially relies.

Definition 3. *Given a coalition C, and two partitions $P_a, P_b \in \mathcal{P}(S)$ for $S \subset A \setminus C$, the externality induced on C for switching from partition P_b to P_a in a coalition structure CS is given by*

$$E(C, P_a, P_b; \bar{P}) = v(C, P_a \cup \bar{P} \cup \{C\}) - v(C, P_b \cup \bar{P} \cup \{C\})$$

where $\bar{P} \in \mathcal{P}(A \setminus (C \cup S))$ and $CS = \{C\} \cup P_b \cup \bar{P}$. By extension, the externality on a partition P is given by

$$E(P, P_a, P_b; \bar{P}) = \sum_{p \in P} E(p, P_a, P_b; \bar{P}).$$

3. RELATED WORK

The problem of optimal coalition structure generation (CSG) has received significant attention in the past. Ketchpel [5] presented a holistic approach that addressed both both CSG and payoff division among the agents. Shehory and Kraus [13] present an algorithm for coalition formation among cooperating agents, but limit the number of agents that can belong to any single coalition. Sandholm et. al. [12] present perhaps the first algorithm for optimal CSG with *anytime* guarantees, i.e., their algorithm can produce a solution of bounded quality after some minimal search, and the quality of the solution provably improves with further search. This algorithm was also shown to be effective in average case studies [6]. Subsequently, more anytime algorithms with improved bounds were introduced [2, 10, 11]. These algorithms have significantly advanced the state-of-the-art in CSG. However, a common characteristic of all of this work is that they assume the coalition values in the *characteristic function form* where coalitions do not impact each other's values. This is a limiting assumption that is often violated in reality.

More recently, there has been interest in the more realistic *partition function form* of coalition values, where a coalition's value differs from one coalition structure to another, due to (different) *externalities* induced by the other coalitions. Various representational schemes have been proposed and compared, for such coalitional games [8]. Michalak et. al. [7] have also studied a limited setting of such games where either the externalities are always positive and the game is weakly sub-additive (PF^+_{sub}), or the externalities are always negative and the game is weakly super-additive (PF^-_{sup}). Rahwan et. al. [9] have relaxed the assumptions of sub and super-additivities in such games, and proposed an optimal CSG algorithm for the more general PF^+ and PF^- settings, where only the externalities are either always positive or always negative. However, even this setting is not general enough since in reality some externalities are negative while others are positive in the *same* game (e.g., consider the example of acquisition of Yahoo by Microsoft). We offer a further generalization over the PF^+ and PF^- settings that allows such *mixed* externalities, and derive a branch and bound algorithm for this more general setting.

4. AGENTS WITH TYPES

We assume that each agent has a *type*. In particular, a set of types $T = \{1, \ldots, t\}$, with $1 \leq t \leq n$, and a type function $\tau : A \mapsto T$ are assumed. Let n_i represent the number of agents in A that are of type $i \in T$. Therefore, $\sum_{i=1}^{i=t} n_i = n$. We also refer to the set of types of agents belonging to a coalition C as $T_C \subseteq T$. The following definition helps us extract agents of some given types, from a given coalition.

Definition 4. *Given any set of agents S (which may or may not be a coalition), and a set of types T, a subset of agents from S that have matching types from T is given by*

$$A(T, S) = \{a \in S | \tau(a) \in T\}$$

Next we define two disjoint subsets of $\bar{C}$ for any given coalition C, that will form the foundation of our setting.

Definition 5. *Given a coalition C of agents from A, we define two disjoint subsets of agents in $A \setminus C$ (i.e., $\bar{C}$) : the complementers and the competitors, given respectively as*

$$A_C = \{a \in A \setminus C | \nexists b \in C, \tau(a) = \tau(b)\} = A(T \setminus T_C, \bar{C})$$

$$A_C^{'} = \{a \in A \setminus C | \exists b \in C, \tau(a) = \tau(b)\} = A(T_C, \bar{C})$$

The set A_C stands for the set of all agents outside C, which have distinct types from all agents in C. Complementarily, $A_C^{'}$ stands for the set of all agents outside C, each of which matches the type of at least one agent in C. We argue that the maximum value of any coalition C (over all coalition structures) is achieved when all agents in $A_C^{'}$ are separated (i.e., singletons) while all agents in A_C are in a single coalition. This is because agents in $A_C^{'}$ can be considered as competitors (e.g., they offer the same products or services, being of a matching type), which gives the best value to C when $A_C^{'}$ is maximally fractured. On the other hand, the agents in A_C being of distinct (or complementary) types, the best value to C is achieved when they offer a united front of interaction to C. For instance, retail giants such as Wal-mart (representing a coalition of various types of merchandise manufacturers) are advantageous to non-retailers since they offer single-stop shopping for a wide variety of goods. However, they do compete with other retailers that offer similar merchandise. In terms of our notations,

$$v(C, CS) \leq v(C, \{C\} \cup P_1(A_C^{'}) \cup P_\infty(A_C)) \quad (1)$$

By the same token, the minimum value of a coalition C is achieved when the agents in $A_C^{'}$ (i.e., competitors) are in a single coalition, while the agents in A_C are all singletons. As a running example, consider $A = \{a_1, a_2, a_3, a_4\}$, and $T = \{1, 2\}$, with $\tau(a_1) = \tau(a_2) = 1$, $\tau(a_3) = \tau(a_4) = 2$, as shown in Table 1. Then the coalition $\{a_1\}$ has its maximum value in the coalition structure $\{\{a_1\}, \{a_2\}, \{a_3, a_4\}\}$, but its minimum value in the coalition structure $\{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}\}$. Similarly, the coalition $\{a_1, a_3\}$ has its maximum value in the coalition structure $\{\{a_1, a_3\}, \{a_2\}, \{a_4\}\}$, but its mimimum value in $\{\{a_1, a_3\}, \{a_2, a_4\}\}$. Table 1 shows all the coalition structures where a coalition has its maximal value. Formally,

$$v(C, CS) \geq v(C, \{C\} \cup P_1(A_C) \cup P_\infty(A_C^{'})) \quad (2)$$

Table 1: Running example, showing the 12 distinct optimal coalition structures for the 15 possible coalitions. The boxed coalition structures are repeated.

$A = \{a_1, a_2, a_3, a_4\}$
$T = \{1, 2\}$
$\tau(a_1) = \tau(a_2) = 1, \tau(a_3) = \tau(a_4) = 2$

Coalition (C)	Coalition structure where C has its maximum value
$\{a_1\}$	$\{\{a_1\}, \{a_2\}, \{a_3, a_4\}\}$
$\{a_2\}$	$\boxed{\{\{a_1\}, \{a_2\}, \{a_3, a_4\}\}}$
$\{a_3\}$	$\{\{a_1, a_2\}, \{a_3\}, \{a_4\}\}$
$\{a_4\}$	$\boxed{\{\{a_1, a_2\}, \{a_3\}, \{a_4\}\}}$
$\{a_1, a_2\}$	$\{\{a_1, a_2\}, \{a_3, a_4\}\}$
$\{a_1, a_3\}$	$\{\{a_1, a_3\}, \{a_2\}, \{a_4\}\}$
$\{a_1, a_4\}$	$\{\{a_1, a_4\}, \{a_2\}, \{a_3\}\}$
$\{a_2, a_3\}$	$\{\{a_1\}, \{a_2, a_3\}, \{a_4\}\}$
$\{a_2, a_4\}$	$\{\{a_1\}, \{a_2, a_4\}, \{a_3\}\}$
$\{a_3, a_4\}$	$\boxed{\{\{a_1, a_2\}, \{a_3, a_4\}\}}$
$\{a_1, a_2, a_3\}$	$\{\{a_1, a_2, a_3\}, \{a_4\}\}$
$\{a_1, a_2, a_4\}$	$\{\{a_1, a_2, a_4\}, \{a_3\}\}$
$\{a_1, a_3, a_4\}$	$\{\{a_1, a_3, a_4\}, \{a_2\}\}$
$\{a_2, a_3, a_4\}$	$\{\{a_1\}, \{a_2, a_3, a_4\}\}$
$\{a_1, a_2, a_3, a_4\}$	$\{\{a_1, a_2, a_3, a_4\}\}$

In [9], two special settings with externalities are defined, viz., PF^+ and PF^-, where all externalities are positive and negative respectively. The consequence of such monotonic externalities is that a coalition attains its higest value when all other agents (outside the coalition) are singletons (in PF^-), or united in a single coalition (in PF^+). It is instructive to note that our setting above reduces to these two special settings under the following conditions:

PF^+: when $t = n$, and $n_i = 1$, $\forall i$, i.e., all agents are of distinct types.

PF^-: when $t = 1$, $n_1 = n$, i.e., when all agents are of a single type.

In the above special cases, the total externality exerted by all other coalitions on any given coalition is either always negative (PF^-), or always positive (PF^+), and all results from [9] should extend to these cases unchanged. However, in the general case where externalities can have mixed signs, the *partition* values cannot be upper and lower bounded in *single* coalition structures, unlike [9] (Theorem 1). In other words, the upper and lower bounds of individual coalitions in any partition may occur in different coalition structures, as evident from equations 1, 2. However, the upper and lower bounds on the value of any partition can still be established as

$$V(P, CS) \leq \sum_{p \in P} \max_{CS} v(p, CS) \quad (3)$$

$$V(P, CS) \geq \sum_{p \in P} \min_{CS} v(p, CS) \quad (4)$$

both of which can be computed without any combinatorial search, exploiting equations 1 and 2.

We make the following two basic assumptions that are justified in the type-based setting. Most importantly, equations 1– 4 follow from these assumptions as consequences, therefore these assumptions are central to our setting.

Assumption 1. *Given two coalitions $C_1, C_2 \in CS$ such that $T_{C_1} \cap T_{C_2} \neq \emptyset$, and any partition $P \in \mathcal{P}(C_2)$,*

$$E(C_1, P, \{C_2\}; CS) \geq 0$$

Since C_1 and C_2 contain some agents of common types, these coalitions are competitive to each other. As a result, *any* partitioning of one (here C_2) is likely to weaken it and benefit the other (here C_1), i.e., exert a positive externality on the latter. The above assumption formalizes this intuition. Clearly, it is biased toward weakening coalitions that are *known to be* competitive (even if the C_2 consists of just a few types competitive to C_1), rather than strenghtening the complementary types in C_2 since these agents are only *potential* collaborators of C_1.

Assumption 2. *Given a coalition $C \in CS$ and a partition $P \subseteq CS \setminus C$, if $T_C \cap T_{\alpha(P)} = \emptyset$ then*

$$E(C, \{\alpha(P)\}, P; CS) \geq 0$$

The above assumption formalizes the complementary intuition that when a (sub)partition of $\bar{C}$ (here P) is complementary to C in types, then the (sub)partition constitutes *potential collaborators* for C. Hence uniting that (sub)partition will strengthen the potential collaborators of C, and thus exert a positive externality on C.

5. WORST-CASE INITIAL BOUND

In this section we establish a minimal set of coalition structures, $\mathcal{P}_0 \subset \mathcal{P}(A)$, that must be searched to establish a bound

$$\beta = \frac{\max_{P \in \mathcal{P}(A)} V(P)}{\max_{P \in \mathcal{P}_0} V(P)} \tag{5}$$

on the quality of the best solution in $\mathcal{P}_0$. As established before [12], in CSG problems there exists a minimal set of coalition structures such that unless all of these are seen, no bound β can be established, since the optimal coalition structure can be arbitrarily better than the best of a smaller set. Such a bound may be useful in at least two ways. Firstly, it establishes an initial bound that can be reduced with further search in an *anytime* fashion [12, 9]. Another application would be to use the value of the best solution in $\mathcal{P}_0$ as the initial estimate in a branch and bound search for the optimal coalition structure. Such initial estimates are often produced through greedy search or other relaxations, without any bound on their quality. Instead, we would search a fixed set of coalition structures, $\mathcal{P}_0$, and know that the initial estimate thus generated has a non-trivially upper-bounded β (equation 5). In this section we produce $\mathcal{P}_0$ and the corresponding β for precisely this purpose.

The following theorem that was proved in [9] is useful in the current setting as well. Therefore, we state this theorem below.

Theorem 1. (Rahwan et. al. 2009) *Let X be a set of elements, and let Y_s be a set containing subsets of X such that: $\forall y \in Y_s, |y| \leq s$. Moreover, for all $x \in X, y \in Y_s$, let us define $v(x,y) \geq 0$ as the value of x in y, and let us define $V(y) = \sum_{x \in y} v(x,y)$ as the value of y. Then for any $Y_s^{'} \subseteq Y_s$, if*

$$\forall x \in X, \exists y' \in Y_s^{'} : x \in y' \text{ and } v(x, y') = \max_{y \in Y_s} v(x,y)$$

then the following holds:

$$\max_{y \in Y_s} V(y) \leq s \cdot \max_{y \in Y_s^{'}} V(y)$$

Next we state and prove the main theorem of this section that stipulates $\mathcal{P}_0$ and the corresponding β in settings with mixed externalities based on agent-types, utilizing the above theorem in the proof.

Theorem 2. *To establish an initial bound β on the best value of a coalition structure returned by any search algorithm, a minimal set of coalition structures, $\mathcal{P}_0$, of size $|\mathcal{P}_0| = 2^n - n + t - 2^{t-1}$ must be searched, with the resulting bound being*

$$\beta = t + \left\lfloor \frac{n-t}{2} \right\rfloor$$

Proof: We follow the same line of reasoning as [9], requiring that the maximum possible value of each coalition C must be observed (in some coalition structure), to establish a bound. This means one coalition structure would have to be observed for each coalition ($2^n - 1$ such), unless some of these coalition structures are common among multiple coalitions, which is indeed the case. Therefore the number of coalition structures that need to be searched will be $\leq 2^n - 1$. The following two disjoint cases capture all such commonalities:

Case 1: $|C| = 1$, i.e., the maximum value of a singleton is being observed. Let $C = \{a\}$. Then all other $n_{\tau(a)} - 1$ agents (of the same type as a) must be separated, and all $\sum_{j \neq \tau(a)} n_j$ agents must be together in the coalition structure where C has its maximum value. But any other singleton of type $\tau(a)$ will have its maximum value in the same coalition structure as well. Thus, only 1 coalition structure needs to be observed for n_i singleton coalitions of type i, instead of n_i. Therefore, in this case, the adjustment to $2^n - 1$ needs to be

$$\begin{aligned} & \sum_1^t (-n_i + 1) \\ = \quad & -n + t. \end{aligned}$$

Case 2: C *covers* some types, i.e., for some $T' \subset T$, C contains all agents of types in T'. Then $A \setminus C$ contains agents that are of complementary types to C, and therefore these agents must be in a single coalition for C to achieve its maximum value. Thus the optimal coalition structure for C will be $\{C, A \setminus C\}$. But by the same argument, the coalition $A \setminus C$ will have its maximum value in the same coalition structure as well. Therefore, only 1 coalition structure needs to be observed for the optimal values of both coalitions C and $A \setminus C$, instead of 2. So in this case the adjustment to $2^n - 1$ needs to be

$$\begin{aligned} & -\frac{1}{2} \sum_1^{t-1} \binom{t}{k} \\ = \quad & -2^{t-1} + 1. \end{aligned}$$

Combining the above adjustments, the size of the set of coalition structures, $\mathcal{P}_0$, to be observed is

$$2^n - n + t - 2^{t-1}$$

In the running example in Table 1, the 12 distinct coalition structures that constitute $\mathcal{P}_0$ are shown (unboxed in the right column).

Now in order to establish the bound β we utilize theorem 1, by

Step 1: selecting X to be the set of all coalitions, augmented by partitions of singletons, *whose maximal values can be observed in $\mathcal{P}_0$*, and

Step 2: selecting Y_s to be the set of coalition structures where many structures have been compressed by coalescing partitions that appear in X.

The augmentation of X in step 1 above can be performed in the following way: for any type $i \in T$, we can select a subset (of size 2 thru n_i) of the n_i agents of type i and create a singleton partition of these agents. These partitions are guaranteed to appear at least once in $\mathcal{P}_0$ with their optimal values, by its construction, and only contain agents of the same type. However, no other partition can be guaranteed, in general, to occur in $\mathcal{P}_0$ with their optimal values, and therefore only

$$\begin{aligned} & \sum_{i=1}^{t} \binom{n_i}{2} + \ldots + \sum_{i=1}^{t} \binom{n_i}{n_i} \\ = & \sum_{i=1}^{t} 2^{n_i} - n - t \end{aligned}$$

singleton-partitions can be added to X in step 1 [1]. In the running example in Table 1, only the $2 \cdot 2^2 - 4 - 2 = 2$ singleton-partitions $\{\{a_1\}, \{a_2\}\}$ and $\{\{a_3\}, \{a_4\}\}$ can be added to X. The partition $\{\{a_1\}, \{a_3\}\}$, for instance, cannot be added to X since it may not be guaranteed to occur with its optimal value in $\mathcal{P}_0$, [2] although it does appear in $\mathcal{P}_0$.

Given the above augmentation of X, in step 2 we coalesce the singletons in any coalition structure that match any partition in augmented X. We are interested in the size of the largest coalition structure at the end of this process. We construct such a coalition structure by first picking t agents –one of each distinct type producing a singleton coalition– and then pairing the remaining $n - t$ agents. Evidently, a larger coalition structure will contain a partition matching with one in the augmented X, and thus compressed. In the running example in Table 1, $\{\{a_1\}, \{a_3\}, \{a_2, a_4\}\}$ is such a coalition structure, and the larger coalition structure $\{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}\}$ will have been compressed to $\{\{\{a_1\}, \{a_2\}\}, \{\{a_3\}, \{a_4\}\}\}$ which is only of size 2. Also note that if $n - t$ is odd, then the one agent left-over from the pairing will be coalesced with one of the t agents selected first, since it must have a matching type. Therefore, the set of coalition structures, Y_s, that we can limit our focus to, contains no coalition structure with more than

$$s = t + \left\lfloor \frac{n-t}{2} \right\rfloor$$

(compressed) coalitions. Now, given the augmented X from step 1, and Y_s determined in step 2, theorem 1 applies with $Y_s^{'} = \mathcal{P}_0$. Therefore,

$$\beta = t + \left\lfloor \frac{n-t}{2} \right\rfloor$$

□

It is straightforward to verify that the results in the above Theorem reduce to the results in Theorem 3 in [9] in the special cases of PF^+ and PF^-. For instance, in PF^- (where $t = 1$ in the type-based setting), Rahwan et. al (2009) proved $\beta = \lceil \frac{n}{2} \rceil$, which matches the above expression of β upon the substitution $t = 1$. Furthermore, the above proof not only establishes minimum $|\mathcal{P}_0|$ but also stipulates how to construct $\mathcal{P}_0$, which will be used for branch and bound search, as described in the next section.

[1] If $n_p = 1$ for some type p, then it will contribute 0 to this sum, since that agent already appears in X as a coalition.

[2] In this particular example, actually the partition $\{\{a_1\}, \{a_3\}\}$ can be shown to occur with its optimal value in $\mathcal{P}_0$, by elimination. Since there are only two possible coalition structures where this partition can occur, viz., $\{\{a_1\}, \{a_2, a_4\}, \{a_3\}\}$ and $\{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}\}$, but $\{\{a_1\}, \{a_2\}, \{a_3\}, \{a_4\}\}$ happens to be the coalition structure where the partition occurs with its *minimum* value, the other must have its maximum value. Note that such guarantees are instance-dependent, and may not hold in general.

6. BRANCH AND BOUND SEARCH

One important feature of the branch and bound search in [9] was the use of pre-pruning in the integer partition space. This might indicate that an analogous partitioning of the type-space would benefit the current setting in the same way. However, a type-space partitioning of the multi-agent system in a PF^+ setting would degenerate to an agent-space partitioning (since all agents are of different types) and pruning partitions in this space would be as hard as pruning partitions in the agent-space. Therefore, we do not perform any search in the type-space and instead search directly in the agent space, hoping that with the appropriate bounding criteria we can still perform significant pruning.

The following theorem generalizes Theorem 4 in [9] for our type-based setting and is an important bounding criterion in our branch and bound search.

Theorem 3. *Given a coalition $C \subseteq A$ and a partition $P \in \mathcal{P}(C)$, a coalition structure containing P can be pruned from the search space if there exists another partition $P' \in \mathcal{P}(C)$ such that $UB_P \leq LB_{P'}$ and $\forall p' \in P'$ at least one of the following is true:*

1. $\exists p \in P : p' = p$
2. $\exists p \in P : p' \subset p$ and $T_p \supseteq T_{\bar{C}}$
3. $T_{p'} \cap T_{\bar{C}} = \emptyset$ and $\exists S \subseteq P : p' \subseteq \alpha(S)$ and
   $\forall p \in S, T_p \cap T_{\bar{C}} = \emptyset \implies p \subset p'$ and
   $\forall p \in S, T_p \cap T_{\bar{C}} \neq \emptyset \implies T_p \supseteq T_{\bar{C}}$

Proof: Given $C \subseteq A$ and two partitions $P, P' \in \mathcal{P}(C)$ such that $UB_P \leq LB_{P'}$ and $\forall p' \in P', \exists p \in P$ such that at least one of the previous statements is true, we follow the same line of reasoning as [9] to show that, for any coalition structure $CS \supseteq P$, there exists another coalition structure CS' such that $V(CS) \leq V(CS')$. Letting $\bar{P} = CS \setminus P$ we have:

$$\begin{aligned} V(CS) &= V(P, CS) + V(\bar{P}, CS) \\ V(CS') &= V(P', CS') + V(\bar{P}, CS') \end{aligned}$$

Since $UB_P \leq LB_{P'}$, we know that

$$V(P, CS) \leq V(P', CS') \tag{6}$$

Therefore we have only left to show that

$$V(\bar{P}, CS) \leq V(\bar{P}, CS') \tag{7}$$

In order to prove equation 7, we must show that the externalities exerted on $\bar{P}$ by P' are greater than or equal to those exerted on $\bar{P}$ by P. That is, we must show that $E(\bar{P}, P', P; \{\}) \geq 0$. To prove this, we must verify that each p' is the result of some operation on P that does not decrease the externalities exerted on $\bar{P}$. We verify this for the three cases outlined in the theorem:

Case 1: p' equals some $p \in P$.
This case is trivial because p' appears in both P and P' and $E(\bar{P}, \{p'\}, \{p'\}, P \setminus \{p\}) = 0$.

Case 2: p' is the result of partitioning (splitting) some $p \in P$. Since $T_p \supseteq T_{\bar{C}}$, then $\forall \bar{p} \in \bar{P}$, $T_{\bar{p}} \cap T_p \neq \emptyset$. Thus, by Assumption 1, $E(\bar{P}, P^{''}, \{p\}; P \setminus \{p\}) \geq 0$ for all $P^{''} \in \mathcal{P}(p)$ with $p' \in P^{''}$.

Case 3: In this case, p' contains agents from multiple $p \in P$. That is, the process by which p' is created entails merging and may also include splitting. We show first that any splitting required to create p' cannot decrease the externalties exerted on $\bar{P}$. We assume without loss of generality that $\forall p \in S, p' \cap p \neq \emptyset$. Let $S' = \bigcup_{p \in S} \{p \cap p', p \setminus p'\}$. Notice that for p that do not need to be split (i.e. $p \setminus p' = \emptyset$), $p \cap p' = p$. Since $T_p \cap T_{\bar{C}} = \emptyset \implies p \subset p'$, we know that $p \setminus p'$ is only non-empty when $T_p \cap T_{\bar{C}} \neq \emptyset$. Furthermore, since $T_p \cap T_{\bar{C}} \neq \emptyset \implies T_p \supseteq T_{\bar{C}}$ we know that $\forall p \in S$, if p was split, $T_p \cap T_{\bar{p}} \neq \emptyset$ for all $\bar{p} \in \bar{P}$. Then, by Assumption 1, $\forall p \in S : T_p \cap T_{\bar{C}} \neq \emptyset, E(\bar{P}, \{p \cap p', p \setminus p'\}, \{p\}; P \setminus \{p\}) \geq 0$. Because s' that are not the result of splitting remain constant from S to S' and therefore do not change the externality, it follows then that $E(\bar{P}, S', S; P \setminus S)) \geq 0$.

Now, we must show that the merging required to create p' from S' cannot decrease the externalities exerted on $\bar{P}$. Let $S^{''} = \{s' \in S | s' \cap p' \neq \emptyset\}$. Because all s' in S' which intersect with p' are subsets of p', and since $S^{''}$ only contains the partitions in S' which intersect with p', S'' is a partition of p'. Since $T_{p'} \cap T_{\bar{C}} = \emptyset$, by Assumption 2, $E(\bar{P}, \{p'\}, S''; P \setminus S) \geq 0$.

We have shown that if a p' satisfies any of the three cases, the process by which p' is created exerts a positive externality on $\bar{P}$. Since each p' must satisfy one of these cases, it follows that

$$V(\bar{P}, CS') \geq V(\bar{P}, CS).$$

This proves equation 7 and therefore completes the proof. □

6.1 The Search Algorithm

We perform branch and bound search under the guidance of the integer partition space, in a manner similar to [9]. Specifically, for any given integer partition of n, $I = [i_1, i_2, \ldots, i_{|I|}]$, we generate non-overlapping coalitions of size i_1, i_2, etc., and at any point $k < |I|$, we check

$$UB_{\{C_1, C_2, \ldots, C_k\}} + Max_{i_{k+1}} + \ldots + Max_{i_{|I|}} < V(CS_{best}) \quad (8)$$

where

- $|C_j| = i_j,\ j = 1 \ldots k$
- $Max_{i_j} = \max_{C \subseteq A : |C| = i_j} v(C, \{C\} \cup P_1(A_C^{'}) \cup P_\infty(A_C))$ (based on equation 1)
- $UB_{\{C_1, C_2, \ldots, C_k\}}$ is the upper bound of partition $\{C_1, C_2, \ldots, C_k\}$, and is calculated from equation 3.
- $V(CS_{best})$ is initialized based on the search in Theorem 2 (with the associated guarantees), and updated during the branch and bound search.

If an agent-partition satisfies equation 8, then it is pointless to expand that partial solution, so that branch is pruned. We also use Theorem 3 as an additional pruning criterion in the following way:

- For any partition $P \in \mathcal{P}(S)$ where $S \subset A$, such that P constitutes the partial solution at the current location in the branch and bound tree, we construct P' as $X \cup Y \cup Z$ where

$$\begin{aligned} X &= \{p \in P | T_{\bar{S}} \setminus T_p \neq \emptyset \text{ and } T_{\bar{S}} \cap T_p \neq \emptyset\} \\ Y &= \bigcup_{p \in P \setminus X} P_1(A(T_{\bar{S}}, p)) \\ Z &= \bigcup_{p \in P \setminus X} \{A(T \setminus T_{\bar{S}}, p)\} \end{aligned}$$

  the set of coalitions based on S that contain some but not all types in $\bar{S}$, and hence cannot be changed to improve the value of $\bar{S}$. Y is built by extracting all agents from the remaining coalitions that have matching types with $\bar{S}$ and turning them into singletons. Z simply keeps all remaining (after X and Y) coalitions in tact since these are potential collaborators of $\bar{S}$. The resulting P' guarantees that the conditions of Theorem 3 (except the bounds) apply and thus form a feasible candidate to prune P by that theorem.

- UB_P and $LB_{P'}$ are calculated using equations 3 and 4 respectively.

Lastly, the integer partitions of n are ordered by decreasing values of their upper bounds, UB_I, given by

$$UB_I = UB_{[i_1, i_2, \ldots, i_{|I|}]} = \sum_{k=1}^{|I|} Max_{i_k},$$

where Max_{i_k} is calculated as stated above. Furthermore, an integer partition whose upper bound is lower than the current $V(CS_{best})$ is also pruned. We call this the 3rd pruning criterion.

6.2 Data Generation

We follow a similar strategy to [9] to generate the values $v(C, CS)$ for experimentation, except that unlike [9], where the externalities are exclusively added or subtracted, we must perform both. In particular, we need to consider the relationship between the type signature of C and each coalition in $CS \setminus \{C\}$. We do this by defining two partitions of $CS \setminus \{C\}$:

$$C_L = \{C' \in CS \setminus \{C\} | T_{C'} \cap T_C = \emptyset\}$$

$$C_C = \{C' \in CS \setminus \{C\} | T_{C'} \cap T_C \neq \emptyset\}$$

where C_L is the set of coalitions which are potential collaborators for C, and therefore exert positive externalities on C. Likewise C_C is the set of coalitions which compete with C and exert negative externalities on it.

As in [9], we randomly generate non-negative numbers v_C and e_C and then produce the sequence $e_{C,1}, e_{C,2}, \ldots, e_{C,|\bar{C}|}$, such that $\sum_{j=1}^{|\bar{C}|} e_{C,j} = e_C$. These enunciate the contribution of each agent in $\bar{C}$ to the total externality on C. Based on these values, we calculate the total externality of any partition $P \in \mathcal{P}(\bar{C})$ on C as:

$$e(P) = \sum_{a_i \in \alpha(P)} e_{C, l(a_i, \bar{C})} * (1 - \frac{l(a_i, P) - 1}{|\alpha(P)|})$$

where $l(a_i, \bar{C})$ gives the location of agent a_i in the set $\bar{C}$ when enumerated lexicographically (e.g., $l(a_3, \{a_2, a_3, a_5, a_8\}) = 2$), and similarly $l(a_i, P)$ gives the location of the (lexicographically listed) coalition containing a_i in the partition P (e.g., $l(a_5, \{\{a_1, a_3\}, \{a_2, a_6\}, \{a_5, a_8\}\}) = 3$). The function l is directly adopted from [9]. Using the e function as defined above, our coalition value function can be written as:

$$v(C, CS) = v_C + e(C_L) - e(C_C)$$

This technique produces coalition values consistent with our mixed externality setting (in particular, satisfying Assumptions 1 and 2), since for every $P \in \mathcal{P}(\bar{C})$, the following necessarily holds:

$$l(a_i, \bar{C}) \geq l(a_i, P) \geq 1.$$

Table 2: Relative (percentage of the total number of coalition structures) prunings by various bounding criteria in branch and bound search in various scenarios.

$\|A\|$	Type partition	Pruning by Cri1	Pruning by Cri2	Pruning by IP	#complete seen
8	[8]	48.92 ± 13.06	15.16 ± 5.53	34.46 ± 22.2	1.46 ± 0.87
	[4,4]	27.57 ± 12.91	19.72 ± 8.1	51.7 ± 22.8	1.01 ± 0.94
	[2,2,2,2]	0.73 ± 0.79	33.49 ± 11	64.33 ± 13.56	1.45 ± 1.26
	[1,1,1,1,1,1,1,1]	1.3 ± 1.4	41.9 ± 10	54.3 ± 16.81	2.5 ± 1.52
10	[10]	45.32 ± 12.9	14.78 ± 10.8	39.55 ± 26.46	0.35 ± 0.41
	[5,5]	27.23 ± 11	17.78 ± 8.7	54.85 ± 18	0.15 ± 0.1
	[2,2,2,2,2]	0.15 ± 0.15	34.52 ± 15.2	64.7 ± 19.71	0.63 ± 0.27
	[1,1,1,1,1,1,1,1,1,1]	1.17 ± 0.89	37.66 ± 11.3	60.21 ± 16.9	0.96 ± 0.55
12	[12]	36.08 ± 9	27.13 ± 10.3	36.74 ± 20.5	0.05 ± 0.05
	[4,4,4]	4.82 ± 1.87	30.65 ± 12.3	64.5 ± 14.2	0.03 ± 0.02
	[2,2,2,2,2,2]	0.114 ± 0.08	49.79 ± 14.4	49.75 ± 17.73	0.35 ± 0.24
	[1,1,1,1,1,1,1,1,1,1,1,1]	1.89 ± 0.96	41.63 ± 13.3	56.19 ± 14.87	0.29 ± 0.25
14	[14]	31.17 ± 18.12	15.51 ± 10.2	53.32 ± 21.23	0.003 ± 0.003
	[2,2,2,2,2,2,2]	0.034 ± 0.013	36.9 ± 12.9	63.04 ± 14.6	0.043 ± 0.03
	[1,1,1,1,1,1,1,1,1,1,1,1,1,1]	1.25 ± 0.4	25.7 ± 7.2	73 ± 0.6	0.048 ± 0.02

6.3 Results

Table 2 shows the number of coalition structures (as percentages of the total number of possible coalition structures in each case) that are pruned by the three pruning criteria, viz., Theorem 3 (criterion 1, called "Cri1" in the table), equation 8 (criterion 2, called "Cri2" in the table), and the IP pruning criterion (criterion 3 in the previous subsection, called "IP" in the table). Note that any branch that can be pruned by both criteria 1 and 2 are counted under criterion 1. We generate different data sets based on 8, 10, 12 and 14 agents, and allocate types to the agents based on the type partitioning as shown in Table 2. Since the type space is large, the type partitions were chosen with no particular motivation other than that we wanted to study uniform partitions in this paper. The last column of the table shows the percentage of complete solutions (i.e., complete coalition structures) seen during the branch and bound search. Each figure in the table is averaged over 10 runs, and the means and standard deviations are shown. The standard deviations are often comparable to the mean since these are averaged over multiple instances for a type setting.

Several interesting trends are visible in Table 2. Firstly, the percentage of complete solutions seen decreases with increasing $|A|$, which usually translates to a diminishing growth in the corresponding run-time. Therefore, the branch and bound algorithm becomes relatively more efficient on larger problems. Secondly, pruning by IP usually dominates the other two pruning criteria, growing to about 60% with more type varieties. Thirdly, criterion 1 (Theorem 3) prunes a significant percentage of coalition structures, and even dominates criterion 2 with fewer agents as well as fewer types. However, this domination seems to wane with larger $|A|$, and criterion 2 always dominates with more types. Interestingly, the effectiveness of criterion 1 in PF^+ (i.e., type partition $[1, 1, \ldots, 1]$ meaning agents are of distinct types) seems to be fixed between 1 and 2%, which indicates a sort of degeneracy of this search algorithm for PF^+.

7. CONCLUSIONS

We have introduced a framework to represent multi-agent systems with mixed externalities, based on the notions of competition and complementation. We have shown that this framework includes the cases PF^+ and PF^- considered in previous work as special cases. We also generalized some of the previous results in anytime and branch and bound search to this setting. Our branch and bound search algorithm shows significant pruning of coalition structures, and a relative speed-up on larger problems.

We have argued that centralized search such as branch and bound may be more suitable for coalition formation in multi-agent systems *with externalities*, than the traditional distributed agent-based approaches such as adaptation with experience, or negotiation. As such, our approach is well-motivated and also extends the state-of-the-art in coalition formation in multi-agent systems with externalities. In the future we will pursue improvement in the ratio bound β with further search to establish an anytime algorithm for this setting. We also intend to investigate the possibility of pre-pruning in the type-partition space, and study its impact on the amount of pruning and the run-time, particularly in the PF^+ setting where the current pruning criteria leave space for improvement.

8. ACKNOWLEDGMENT

The authors gratefully acknowledge the constructive criticism and the suggestions from the anonymous reviewers. This work was supported in part by a start-up grant from the University of Southern Mississippi.

9. REFERENCES

[1] V. D. Dang, R. K. Dash, A. Rogers, and N. R. Jennings. Overlapping coalition formation for efficient data fusion in multi-sensor networks. In *Proceedings of AAAI*, 2006.

[2] V. D. Dang and N. R. Jennings. Generating coalition structures with finite bound from the optimal guarantees. In *Proceedings of AAMAS*, 2004.

[3] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid - enabling scalable virtual organizations. *International Journal of Supercomputer Applications*, 15:2001, 2001.

[4] I. Hafalir. Efficiency in coalition games with externalities. *Games and Economic Behaviour*, 61(2):209–238, 2007.

[5] S. Ketchpel. Forming coalitions in the face of uncertain rewards. In *Twelfth National Conference on Artificial Intelligence*, pages 414–419. MIT Press/AAAI Press, 1994.

[6] K. Larson and T. Sandholm. Anytime coalition structure generation: An average case study. *Journal of Experimental and Theoretical Artificial Intelligence*, 12(1):23–42, 2000.

[7] T. Michalak, A. Dowell, P. McBurney, and M. Wooldridge. Optimal coalition structure generation in partition function games. In *Proceedings of the 18th European Conference on Artificial Intelligence*, 2008.

[8] T. Michalak, T. Rahwan, J. Sroka, A. Dowell, M. Wooldridge, P. McBurney, and N. Jennings. On representing coalitional games with externalities. In *The Proceedings of Electronic Commerce*, Stanford, CA, 2009. ACM.

[9] T. Rahwan, T. Michalak, N. Jennings, M. Wooldridge, and P. McBurney. Coalition structure generation in multi-agent systems with positive and negative externalities. In *Proceedings of The International Joint Conference on Artificial Intelligence*, 2009.

[10] T. Rahwan, S. D. Ramchurn, V. D. Dang, A. Giovannucci, and N. R. Jennings. Anytime optimal coalition structure generation. In *Proceedings of the 22nd National Conference on Artificial Intelligence*, 2007.

[11] T. Rahwan, S. D. Ramchurn, V. D. Dang, and N. R. Jennings. Near-optimal anytime coalition structure generation. In *Proceedings of IJCAI*, pages 2365–2371, 2007.

[12] T. W. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Anytime coalition structure generation with worst case guarantees. In *Proceedings of the Fifteent National Conference on Artificial Intelligence*, pages 46–53, Menlo Park, CA, 1998. AAAI Press/MIT Press.

[13] O. Shehory and S. Kraus. Task allocation via coalition formation among autonomous agents. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 655–661, August 1995.

[14] M. Sims, C. Goldman, and V. Lesser. Self-organization through bottom-up coalition formation. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003)*, pages 867–874. ACM Press, July 2003.

[15] L.-K. Soh and C. Tsatsoulis. Satisficing coalition formation among agents. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems (AAMAS 2002)*, pages 1062–1063, New York, NY, USA, 2002. ACM.

[16] M. Tsvetovat and K. Sycara. Customer coalitions in electronic markets. In *Proceedings of AMEC 2000*, pages 121–138, 2000.