Online Appendix to: Algorithms for Graph-Constrained Coalition Formation in the Real World

FILIPPO BISTAFFA and ALESSANDRO FARINELLI, University of Verona
JESÚS CERQUIDES and JUAN RODRÍGUEZ-AGUILAR, IIIA-CSIC
SARVAPALI D. RAMCHURN, University of Southampton

APPENDIX

In this appendix, we first discuss the relation between GCCF and the CCF framework by Rahwan et al. [2011] in Section A.1. Then, in Sections A.2–A.4 we provide the proofs of propositions and theorems not included in our manuscript. Finally, in Section A.5 and Section A.6 we provide some additional details and experiments.

A.1. Relation with the CCF framework

Rahwan et al. [2011] also considered scenarios in which constraints enforce (or prohibit) the co-existence of agents in a coalition, introducing the problem of Constrained Coalition Formation (CCF) to adequately deal with these constraints. The authors also identify a natural, simpler subclass of CCF games, namely Basic CCF (BCCF) games, providing the state of the art algorithm for such problems.

Formally, Rahwan et al. define a CCF game as a coalitional game together with a set of feasible coalition structures $\mathcal{CS} \subseteq \Pi(A)$. The input to a CCF problem is a CCF game and its objective is to identify the most valuable coalition structure which is feasible, i.e., $CS^* = \arg\max_{CS \in \mathcal{CS}} V(CS)$. Furthermore, they identify a natural subclass of CCF games, namely Basic CCF (BCCF) games. A BCCF is a coalitional game together with: (i) a set of positive constrains $\mathcal{P} \subseteq 2^{\mathcal{A}}$, (ii) a set of negative constraints $\mathcal{N} \subseteq 2^{\mathcal{A}}$, and (iii) a set of allowed sizes $\mathcal{S} \subseteq \mathbb{N}$. A coalition C satisfies a positive constraint $P \in \mathcal{P}$ if $P \subseteq C$. A coalition C violates a negative constraint $N \in \mathcal{N}$ if $N \subseteq C$. The size of a coalition C is allowed if $|C| \in \mathcal{S}$. In a BCCF game a coalition C is considered feasible if C satisfies at least one positive constraint, does not violate any negative constraint and its size is allowed. At first sight, it appears that the GCCF and the BCCF problems may be related, since they both focus on constrained CF. Moreover, it is easy to see that such classes are not disjoint, since unconstrained CF (i.e., a CCF game on which every coalition is feasible) can be represented both as a GCCF (taking a clique as its graph G) and as a BCCF (including in the set of positive constraints all the coalitions with one player). Nonetheless, we show that GCCF and BCCF are *different* problems and hence, the algorithm for BCCF provided by Rahwan et al. cannot be applied to GCCF.

PROPOSITION A.1. *GCCF games are not a subset of BCCF games, and BCCF games are not a subset of GCCF games.*

PROOF.

A.1.1. There are GCCF games which are not BCCF games. Our first example is a GCCF game with three agents $\mathcal{A} = \{1, 2, 3\}$. The set of edges of G is $\mathcal{E} = \{(1, 2), (2, 3)\}$. Thus, the set of feasible coalitions is $\mathcal{FC}(G) = \{\{1\}, \{2\}, \{3\}, \{1, 2\}, \{2, 3\}, \{1, 2, 3\}\}$, that is, the only non-feasible coalition is $\{1, 3\}$. We will now show that this game cannot be encoded as a BCCF game. For the sake of contradiction, assume that it can. Since the single element coalitions $\{1\}, \{2\}, \{3\}$ are feasible, the set of positive constraints

DOI: http://dx.doi.org/10.1145/3040967

$\mathcal{P}$ should include $\{1\},\{2\},\{3\}$. Since the grand coalition $\{1,2,3\}$ is feasible, the set of negative constraints $\mathcal{N}$ should be empty. Since there are feasible coalitions with one, two, and three elements, the set of allowed sizes should be $\mathcal{S}=\{1,2,3\}$. It is easy to see that coalition $\{1,3\}$ is feasible on the candidate BCCF game, and thus that our example GCCF game is not a BCCF game. □

A.1.2. There are BCCF games which are not GCCF games. Our second example is a BCCF game with two agents $\mathcal{A}=\{1,2\}$. The set of positive constraints $\mathcal{P}$ is $\{\{1\}\}$, the set of negative constraints $\mathcal{N}$ is empty and the set of allowed sizes is just $\mathcal{S}=\{2\}$. Thus, the set of feasible coalitions is $\mathcal{FC}(G)=\{\{1,2\}\}$. Again assume that we can encode this game as a GCCF game. Since in every GCCF game, the single element coalitions are feasible, we reach a contradiction. □

Henceforth, the relationship between GCCF and BCCF is represented in Figure 11.

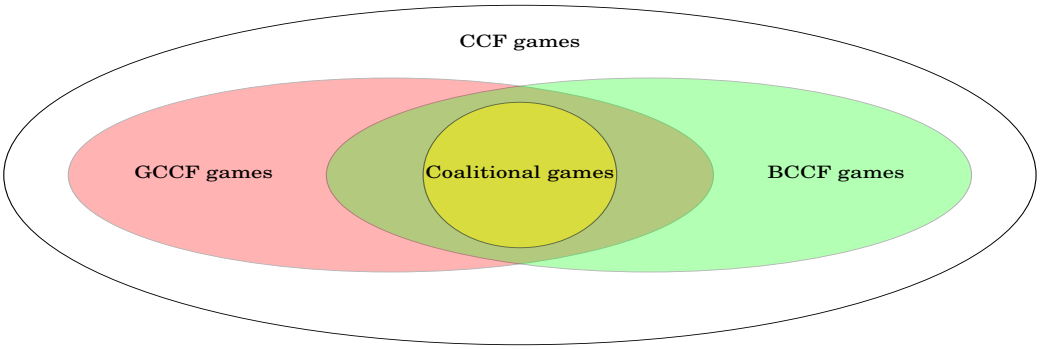

Fig. 11: Relationships between the CCF families.

A.2. Proof of Proposition 4.3

Let G_c be a 2-coloured graph, in which each vertex is labelled with a coalition. Recall that a red edge e between coalitions C_1 and C_2 encodes a constraint imposing that each pair of elements (c_1,c_2), where $c_1\in C_1$ and $c_2\in C_2$, does not lie in the same coalition.

DEFINITION 9. *Given a 2-coloured graph G_c, we say that a coalition structure is compatible with G_c if it can be obtained by contracting green edges and without breaking any of the constraints imposed by red edges.*

DEFINITION 10. *A green edge e is red-colourable with respect to CS if CS satisfies the constraint imposed by e when we colour it red.*

DEFINITION 11. *Given a coalition structure CS and an edge e, we say that e is contractible with respect to CS if the coalition formed by contracting e is included into one of the coalitions of CS.*

Note that if CS is compatible with G_c each edge should be either red-colourable or contractible with respect to CS. Based on these definitions we can prove Proposition 4.3.

PROPOSITION 4.3. *Given a 2-coloured graph G_c, the tree rooted at G_c contains all the coalition structures compatible with G_c, and each of them appears only once.*

PROOF. By induction on the number of green edges. If there is no green edge, then the tree has just one element which corresponds to the only coalition structure compatible with G_c. Assume that the statement is true for $n-1$ green edges. Let G_c have n

green edges and CS be a coalition structure compatible with G_c. If no edge in G_c is contractible with respect to CS, then CS is the coalition represented by G_c, and it cannot be in any of its children, because each of them contracts an edge in G_c. Thus CS appears in the tree rooted at G_c only once (at the root). Assume then that there is at least one green edge in G_c contractible with respect to CS. Then CS cannot be the coalition structure at the root. We would like to identify a child G' such that CS is compatible with G'. The first child of the root contracts an edge e. If e is contractible with respect to CS, then the first child of G_c is compatible with CS. Otherwise e is red-colourable with respect to CS. The same procedure goes on with the remaining children. Thus, by construction, the root has three kind of children with respect to CS: some which contract a red-colourable edge, a single child G' that contracts a contractible edge and red-colours some red-colourable edges, and from there on some that red-colours a contractible edge. It is easy to see that CS is compatible only with one child, namely G'. Now G' has at most $n-1$ green edges and by induction CS must appear in that subtree only once. Thus, it appears in the tree rooted at G_c only once. □

A.3. Proofs in Section 6.1

In this section we provide the proofs of Propositions 6.1 and 6.2.

DEFINITION 6. *Given a graph G, a function $V : \mathcal{CS}(G) \to \mathbb{R}$ is an $m + a$ function if it is the sum of a superadditive function $V^+ : \mathcal{CS}(G) \to \mathbb{R}$ and a subadditive function $V^- : \mathcal{CS}(G) \to \mathbb{R}$.*

PROPOSITION 6.1. *The collective energy purchasing function*

$$V(CS) = \underbrace{\sum_{C \in CS} \left[\sum_{t=1}^{T} q_S^t(C) \cdot p_S + T \cdot q_F(C) \cdot p_F \right]}_{V^+(CS)} + \underbrace{\sum_{C \in CS} \kappa(C)}_{V^-(CS)} \tag{7}$$

is an $m + a$ function.

PROOF. As shown in Equation 7, such function can be seen as an $m + a$ function, being the sum of a superadditive function, consisting of the cost of the energy necessary to fulfil the aggregated consumption profiles of the coalitions, and a subadditive one (i.e., the sum of the coalition management costs). On the one hand, since the baseline (i.e., the minimum) of the aggregate energy profile of a coalition $C_{12} = C_1 \cup C_2$ is no less than the sum of the baselines of the energy profiles of C_1 and C_2, the members of C_{12} can buy from the forward market an amount of energy which is no less than the sum of the amounts that could have been bought by C_1 and C_2 separately.[10] Therefore, the $energy(\cdot)$ function is superadditive. On the other hand, it is trivial to verify that $\kappa(\cdot)$ is a subadditive function. □

PROPOSITION 6.2. *The edge sum with coordination cost function*

$$V(CS) = \underbrace{\sum_{C \in CS} \left[\sum_{e \in edges(C)} w^+(e) \right]}_{V^+(CS)} + \underbrace{\sum_{C \in CS} \left[\sum_{e \in edges(C)} w^-(e) + \kappa(C) \right]}_{V^-(CS)}. \tag{8}$$

is an $m + a$ function.

[10] A more detailed discussion is provided by Vinyals et al. [2012].

PROOF. Equation 8 highlights the $V^+(\cdot)$ and $V^-(\cdot)$ components of this function. On the one hand, $v^+(C) = \sum_{e \in edges(C)} w^+(e)$ is clearly superadditive, since a coalition $C_{12} = C_1 \cup C_2$ contains an amounts of positive edges which is no less than the total amount of positive edges in C_1 and C_2 taken separately, hence $v^+(C_{12}) \geq v^+(C_1) + v^+(C_2)$. Similarly, it is easy to prove that $\sum_{e \in edges(C)} w^-(e)$ is a subadditive function. Hence, the edge sum with coordination cost function is an $m + a$ function. □

A.4. Proof of Theorem 5.1

We now provide the proof of Theorem 5.1, detailing some properties of our domain.

LEMMA A.5. *$\mathcal{CS}(G)$ is a lattice, i.e., a partially ordered set, in which every two elements CS_i and CS_j have a supremum ($CS_i \vee CS_j$) and an infimum ($CS_i \wedge CS_j$).*

In particular, we define the following partial order over coalition structures.

DEFINITION 13. *Given any two coalition structures CS_i and CS_j, we say that $CS_i \leq CS_j$ if every element of CS_i is a subset of some element of CS_j.*

As an example, $\{A, B\}\{C\} \leq \{A, B, C\}$, but the order between $\{A, B\}\{C\}$ and $\{A\}\{B, C\}$ is not defined. It is well known that with this partial order the set of partitions forms a complete lattice (see Section V.4 in [Grätzer 2011]), called the partition lattice or equivalence lattice. It is easy to see that our domain of interest is the sublattice generated by the set of feasible coalition structures and thus it is a lattice. Furthermore, in our scenario, the grand coalition represents a supremum of any two elements, while the coalition structure of all singletons represents an infimum.

LEMMA A.6. *The elements of $\mathcal{CS}(G)$ can be arranged in an order-preserving tree: whenever CS_j is a descendant of CS_i in the tree, then $CS_j \geq CS_i$. Thus, CS_i is the infimum of the subtree rooted at CS_i, i.e., $CS_i = \bigwedge ST(CS_i) = \inf ST(CS_i)$.*

In the search tree defined in Section 4 each child is the result of contracting an edge in the parent. As a consequence of the contraction, two of the coalitions in the parent are merged, making the child partition coarser than that of the parent. Hence, by direct application of Lemma A.5, the above statement holds.

DEFINITION 7. *Given a coalition structure CS_i represented by the 2-coloured graph G_c, the coalition structure $\overline{CS_i}$ can be obtained by removing all red edges from G_c and then contracting all the remaining green edges (which is equivalent to find the connected components in the graph after the removal of all red edges).*

LEMMA A.7. *Given a node CS_i, $\overline{CS_i}$ is bigger than any of the elements of the subtree, i.e., $\overline{CS_i} \geq \bigvee ST(CS_i) = \sup ST(CS_i)$.*

Since $\overline{CS_i}$ represents the connected components in the graph after the removal of all red edges, it can be interpreted as the coarsest partition forgetting that we decided not to contract some edges. Clearly, any partition in the subtree will be at most as coarse as this one. Given the above, we can now provide the proof of Theorem 5.1.

THEOREM 5.1. *Given an $m+a$ function $V : \mathcal{CS}(G) \to \mathbb{R}$, then $M(CS_i) = V^-(CS_i) + V^+(\overline{CS_i})$ is an upper bound for the value assumed by such function in every coalition structure of the subtree $ST(CS_i)$ rooted at CS_i, i.e., $M(CS_i) = V^-(CS_i) + V^+(\overline{CS_i}) \geq \max\{V(CS_j) \mid CS_j \in ST(CS_i)\}$.*

PROOF. Consider that, for the subtree rooted at CS_i, the maximum of a subadditive function will be achieved at CS_i (Lemma A.6), i.e., $V^-(CS_i) \geq \max\{V^-(CS_j) \mid CS_j \in ST(CS_i)\}$. On the other hand, the maximum of a superadditive

function will be reached at one of the leaves. However, since assessing the supremum $\overline{CS_i}$ of the subtree is computationally efficient (Lemma A.7), we can bound $V^+(\cdot)$ in the subtree as $V^+\left(\overline{CS_i}\right) \geq \max\left\{V^+(CS_j) \mid CS_j \in ST\left(CS_i\right)\right\}$. Finally, since we assume that $V(\cdot)$ is an $m+a$ function, then we can provide an upper bound for such a function by composing these two results, i.e., $M\left(CS_i\right) = V^-\left(CS_i\right) + V^+\left(\overline{CS_i}\right)$. □

A.5. Network topologies

In what follows, we provide some details on the two network topologies we adopted in our experimental evaluation, i.e., *scale-free networks* and *community networks*.

A.5.1. Scale-free networks. Scale-free networks are a widely adopted network topology thanks to their ability of adequately describing the features of many real-world scenarios. One of the simplest properties of a network that can be measured directly is the degree distribution, or the fraction $P(k)$ of nodes having k connections (degree k). Direct measurements of the degree distribution for networks of the Internet [Barabási and Albert 1999], WWW (where hypertext links constitute directed edges) [Broder et al. 2000], citations of scientific articles [Redner 1998], metabolic networks [Jeong et al. 2000], and many more, show that such a distribution does not follow the rules of random graphs. Rather, these nets often exhibit a scale-free degree distribution, i.e.,

$$P(k) = ck^{-\lambda}, \quad k = m, \ldots, K,$$

where $c \approx (\lambda - 1)\, m^{\lambda - 1}$ is a normalisation factor and m and K are the lower and upper cutoff for the connectivity of the node [Cohen et al. 2003].

Several models have been presented for the evolution of scale-free networks, each of which may lead to a different ensemble. The first suggestion was the *preferential attachment* model by Albert and Barabási [2002], which generates a scale-free network starting with an initial connected network of m_0 nodes, and adding each subsequent node one at a time with a probability that is proportional to the number of links that the existing nodes already have. Formally, the probability p_i that the new node is connected to node i is $p_i = \frac{k_i}{\sum_j k_j}$, where k_i is the degree of node i and the sum $\sum_j k_j$ is made over all pre-existing nodes j. Heavily linked nodes ("hubs") tend to quickly accumulate even more links, while nodes with only a few links are unlikely to be chosen as the destination for a new link. The new nodes have a "preference" to attach themselves to the already heavily linked nodes.

Scale-free networks provide a good description of many realistic scenarios. Nonetheless, recent studies have developed a different network topology (i.e., community networks) to model such phenomena, as described hereafter.

A.5.2. Community networks. Many realistic scenarios in which GCCF can be successfully applied exhibit an underlying structure determined by sparse synergies which naturally results from interactions between agents. These fields of application include social networks, authorships, collective energy purchasing and carpooling. One fundamental feature of these networks of agents is represented by their mesoscopic structure, characterised by the presence of groups of nodes, called *communities* or *modules*, with a high density of links between nodes of the same group and a comparatively low density of links between nodes of different groups. This compartmental organisation of networks is very common in systems of diverse origin, hence analysing the performance of our approach with the adoption of this particular graph topology is very important. In fact, many of the above-mentioned real-world networks, believed to be scale-free networks, exhibit a community network structure [Kolda et al. 2014b].

In particular, in our experiments *community networks* are provided by the Block Two-Level Erdős-Rényi (BTER) model[11] [Kolda et al. 2014b]. Such a model is based on the idea of a graph comprising communities in the form of dense ER subgraphs (i.e., generated with the Erdős-Rényi model), hence matching well real-world graphs. The generation of community networks with the BTER model is divided in two phases: in the first one, BTER builds a collection of ER subgraphs in such a way that the specified degree distribution is respected. The BTER model allows one to construct a graph with any degree distribution. Real-world degree distributions might be idealised as power laws, but it is by no means a completely accurate description. When the degree distribution is heavy tailed, then the BTER graph naturally has scale-free ER subgraphs. The internal connectivity of the ER graphs is specified by the user and can be tuned to match observed data. These communities are then interconnected in the second phase, in which a Chung-Lu model [Aiello et al. 2000] is used over the excess degrees to form the edges that connect communities.

A.6. DyCE vs CFSS on community networks

We report here the additional experiments not included in Section 7.1 due to space constraints. In particular, here we discuss the results of the comparison between CFSS and DyCE when considering the community network topology and the collective energy purchasing characteristic function. We generated community networks with the BTER model [Kolda et al. 2014b], with an average degree (denoted as $\overline{deg}$) within $\{2, 4, 6\}$. Figure 12 shows that both algorithms achieves a performance comparable to the scale-free network scenario, hence the discussion in Section 7.1 also applies here.

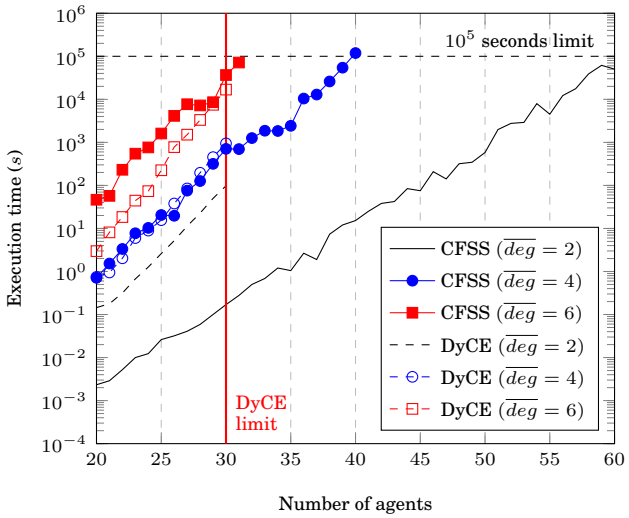

Fig. 12: Collective energy purchasing, community networks.

[11] We use the FEASTPACK v1.1 MATLAB implementation [Kolda et al. 2014a] of the BTER model.

REFERENCES

William Aiello, Fan Chung, and Linyuan Lu. 2000. A Random Graph Model for Massive Graphs. In *ACM Symposium on Theory of Computing*. 171–180.

Réka Albert and Albert L. Barabási. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics* 74 (2002), 47–97.

Albert L. Barabási and Réka Albert. 1999. Emergence of Scaling in Random Networks. *Science* 286 (1999), 509–512.

Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet Wiener. 2000. Graph Structure in the Web. *Computer Networks* 33 (June 2000), 309–320.

Reuven Cohen, Shlomo Havlin, and Daniel Ben-Avraham. 2003. Structural properties of scale-free networks. In *Handbook of Graphs and Networks*. Wiley, 85–110.

George Grätzer. 2011. *Lattice Theory: Foundation*. Springer.

Hoidong Jeong, B. Tombor, Réka Albert, Zoltan N. Oltvai, and Albert L. Barabási. 2000. The large-scale organization of metabolic networks. *Nature* 407 (05 Oct. 2000), 651–654.

Tamara G. Kolda, Ali Pinar, Todd Plantenga, and Comandur Seshadhri. 2014a. FEASTPACK. (2014). http://www.sandia.gov/~tgkolda/feastpack/

Tamara G. Kolda, Ali Pinar, Todd Plantenga, and Comandur Seshadhri. 2014b. A scalable generative graph model with community structure. *SIAM Journal on Scientific Computing* 36 (2014), 424–452.

Talal Rahwan, Tomasz P. Michalak, Edith Elkind, Piotr Faliszewski, Jacek Sroka, Michael Wooldridge, and Nicholas R Jennings. 2011. Constrained Coalition Formation. In *AAAI Conference on Artificial Intelligence*. 719–725.

Sidney Redner. 1998. How popular is your paper? An empirical study of the citation distribution. *The European Physical Journal B-Condensed Matter and Complex Systems* 4 (1998), 131–134.

Meritxell Vinyals, Filippo Bistaffa, Alessandro Farinelli, and Alex Rogers. 2012. Coalitional energy purchasing in the smart grid. In *IEEE International Energy Conference*. 848–853.