CHAPTER 13: FORMING COALITIONS

Multiagent Systems

http://www.csc.liv.ac.uk/˜mjw/pubs/imas/

Coalitional Games

- *Coalitional games* model scenarios where agents can *benefit by cooperating*.
- Issues in coalitional games (Sandholm et al, 1999):
  - *Coalition structure generation*.
  - *Teamwork*.
  - *Dividing the benefits of cooperation*.

Coalition Structure Generation

- Deciding *in principle* who will work together.
- The basic question:

Which coalition should I join?

- The result: *partitions* agents into disjoint *coalitions*.
  The overall partition is a *coalition structure*.

Solving the optimization problem of each coalition

- Deciding *how* to work together.
- Solving the “joint problem” of a coalition
- Finding how to maximise the utility of the coalition itself.
- Typically involves joint planning etc.

Dividing the Benefits

- Deciding “who gets what” in the payoff.
- Coalition members cannot ignore each other’s preferences, because members can *defect*: if you try to give me a bad payoff, I can always walk away.
- We might want to consider issues such as *fairness* of the distribution.

Formalising Cooperative Scenarios

A coalitional game:

$$\langle Ag, \nu \rangle$$

where:

- $Ag = \{1, \ldots, n\}$ is a set of *agents*;
- $\nu : 2^{Ag} \rightarrow \mathbb{R}$ is the *characteristic function* of the game.

Usual interpretation: if $\nu(C) = k$, then coalition C can cooperate in such a way they will obtain utility k, which may then be distributed amongst team members.

Which Coalition Should I Join?

- Most important question in coalitional games:

  is a coalition stable?

  that is,

  is it rational for all members of coalition to stay with the coalition, or could they benefit by defecting from it?

- (There is no point in me trying to join a coalition with you unless you want to form one with me, and vice versa.)
- Stability is a *necessary* but not *sufficient* condition for coalitions to form.

The Core

- The *core* of a coalitional game is the set of *feasible* distributions of payoff to members of a coalition that *no* sub-coalition can reasonably object to.
- An *outcome* for a coalition C in game $\langle Ag, \nu \rangle$ is a vector of payoffs to members of C, $\langle x_1, \ldots, x_k \rangle$ which represents a *feasible distribution of payoff to members of Ag*.

  "Feasible" means:

$$\nu(C) \geq \sum_{i \in C} x_i$$

- Example: if $\nu(\{1,2\}) = 20$, then possible outcomes are $\langle 20,0\rangle$, $\langle 19,1\rangle$, $\langle 18,2\rangle$, $\ldots$, $\langle 0,20\rangle$.

  (Actually there will be infinitely many!)

Objections

- Intuitively, a coalition C *objects* to an outcome if there is some outcome *for them* that makes *all of them* strictly better off.
- Formally, $C \subseteq Ag$ objects to an outcome $\langle x_1, \ldots, x_n \rangle$ for the grand coalition if there is some outcome $\langle x'_1, \ldots, x'_k \rangle$ for C such that

$$x'_i > x_i \qquad \text{for all } i \in C$$

- The idea is that an outcome is not going to happen if somebody objects to it!

The Core

- The *core* is the set of outcomes for the *grand coalition* to which *no* coalition objects.
- If the core is *non-empty* then the *grand coalition is stable*, since nobody can benefit from defection.
- Thus, asking

is the grand coalition stable?

is the same as asking:

is the core non-empty?

Problems with the Core

- Sometimes, the core is empty; what happens then?
- Sometimes it is non-empty but isn't "fair".

  Suppose $Ag = \{1, 2\}$, $\nu(\{1\}) = 5$, $\nu(\{2\}) = 5$, $\nu(\{1, 2\}) = 20$.

  Then outcome $\langle 20, 0\rangle$ (i.e., agent 1 gets everything) is *not* in the core, since the coalition $\{2\}$ can object. (He can work on his own and do better.)

  However, outcome $\langle 15, 5\rangle$ *is* in the core: even though this seems unfair to agent 2, this agent has no objection.
- Why unfair? Because the agents are *identical!*

How To Share Benefits of Cooperation?

- The *Shapley value* is best known attempt to define how to divide benefits of cooperation fairly.

  It does this by taking into account *how much an agent contributes*.
- The Shapley value of agent i is the average amount that i is expected to contribute to a coalition.
- Axiomatically: a value which satisfies axioms:

  symmetry, *dummy player*, and *additivity*.

Shapley Defined

- Let $\delta_i(S)$ be the amount that i adds by joining $S \subseteq Ag$:

$$\delta_i(S) = \nu(S \cup \{i\}) - \nu(S)$$

... the *marginal contribution of i to S*.

- Then the Shapley value for i, denoted φ_i, is:

$$\varphi_i = \frac{\sum_{r \in R} \delta_i(S_i(r))}{|Ag|!}$$

where R is the set of all orderings of Ag and $S_i(r)$ is the set of agents preceding i in ordering r.

Representing Coalitional Games

- It is important for an agent to know (eg) whether the core of a coalition is non-empty . . .

  so, how hard is it to decide this?
- Problem: naive, obvious representation of coalitional game is *exponential* in the size of Ag!
- Now such a representation is:
  - *utterly* infeasible in practice; and

– so large that it renders comparisons to this input size meaningless: stating that we have an algorithm that runs in (say) time *linear* in the size of such a representation means it runs in time *exponential* in the size of Ag!

How to Represent Characteristic Functions?

Two approaches to this problem:

- try to find a *complete* representation that is succinct in "most" cases
- try to find a representation that is not complete but is always succinct
- A common approach:

  interpret characteristic function over combinatorial structure.

Representation 1: Induced Subgraph

- Represent ν as an undirected graph on Ag, with integer weights $w_{i,j}$ between nodes $i, j \in Ag$.
- Value of coalition C then:

$$\nu(C) = \sum_{\{i,j\} \subseteq Ag} w_{i,j}$$

i.e., the value of a coalition $C \subseteq Ag$ is the weight of the subgraph induced by C.

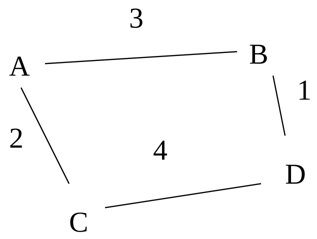

the original graph defining v

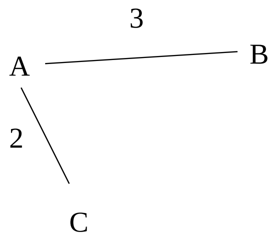

subgraph induced by {A,B,C}
giving v({A,B,C}) = 3 + 2 = 5

Representation 1: Induced Subgraph

(Deng & Papadimitriou, 94)

- Computing Shapley:
  in polynomial time.
- Determining emptiness of the core:
  NP-complete
- Checking whether a specific distribution is in the core
  co-NP-complete

But this representation is not *complete*.

Representation 2: Weighted Voting Games

- For each agent $i \in Ag$, assign a weight w_i, and define an overall *quota*, q.

$$\nu(C) = \begin{cases} 1 & \text{if } \sum_{i \in C} w_i \geq q \\ 0 & \text{otherwise.} \end{cases}$$

- Shapley value:

  #P-complete, and "hard to approximate" (Deng & Papadimitriou, 94).

- Core non-emptiness:

  in polynomial time.

Not a complete representation.

Representation 3: Marginal Contribution Nets

(Ieong & Shoham, 2005)

- Characteristic function represented as rules:

$$\text{pattern} \longrightarrow \text{value}.$$

- Pattern is conjunction of agents, a rule *applies* to a group of agents C if C is a superset of the agents in the pattern.

  Value of a coalition is then sum over the values of all the rules that apply to the coalition.

  Example:

$$
\begin{aligned}
a \wedge b &\longrightarrow 5 \\
b &\longrightarrow 2
\end{aligned}
$$

  We have: $\nu(\{a\}) = 0$, $\nu(\{b\}) = 2$, and $\nu(\{a, b\}) = 7$.

- We can also allow negations in rules (agent not present).

Representation 3: Marginal Contribution Nets

- Shapley value:
  in polynomial time
- Checking whether distribution is in the core:
  co-NP-complete
- Checking whether the core is non-empty:
  co-NP-hard.

A complete representation, but not necessarily succinct.

Qualitative Coalitional Games

- Often not interested in utilities, but in *goals* – either the goal is satisfied or not
- *QCG*s are a type of coalitional game in which each agent has a set of goals, and wants one of them to be achieved (doesn't care which)

  Agents cooperate in QCGs to achieve mutually satisfying sets of goals.

  Coalitions have *sets of choices* representing the different ways they could cooperate

  Each choice is a set of goals.

QCGs

A *Qualitative Coalitional Game* (QCG) is a structure:

$$\Gamma = \langle G, Ag, G_1, \ldots, G_n, V \rangle$$

where

- $G = \{g_1, \ldots, g_m\}$ is a set of *possible goals*;
- $Ag = \{1, \ldots, n\}$ is a set of *agents*;
- $G_i \subseteq G$ is a set of goals for each agent $i \in Ag$, the intended interpretation being that any of G_i would satisfy i;

- $V : 2^{Ag} \to 2^{2^G}$ is a *characteristic function*, which for every coalition $C \subseteq Ag$ determines a set $V(C)$ of *choices*, the intended interpretation being that if $G' \in V(C)$, then one of the choices available to coalition C is to bring about *all* the goals in G' simultaneously.

Feasible/Satisfying Goal Sets

- Goal set $G' \subseteq G$ *satisfies* an agent i if $G_i \cap G' \neq \emptyset$.
  Goal set $G' \subseteq G$ satisfies a coalition $C \subseteq Ag$ if
  $$\forall i \in C, G_i \cap G' \neq \emptyset$$
- A goal set G' is *feasible* for C if $G' \in V(C)$.

Representing QCGs

- So, how do we represent the function $V : 2^{Ag} \to 2^{2^G}$?
- We use a formula Ψ_V of propositional logic over propositional variables Ag, G, such that:

$$\Psi[C, G'] = \top \text{ if and only if } G' \in V(C)$$

- "Often" permits *succinct* representations of V.
- Note that given Ψ_V, C, G', determining whether $G' \in V(C)$ can be done in time polynomial in size of C, G', Ψ_V.

Fourteen QCG Decision Problems (AIJ, Sep 2004)

Problem	Description	Complexity	q^{mono}
SC	SUCCESSFUL COALITION	NP-complete	NP-complete
SSC	SELFISH SUCCESSFUL COALITION	NP-complete	NP-complete
UGS	UNATTAINABLE GOAL SET	NP-complete	NP-complete
MC	MINIMAL COALITION	co-NP-complete	co-NP-complete
CM	CORE MEMBERSHIP	co-NP-complete	co-NP-complete
CNE	CORE NON-EMPTINESS	D^p-complete	D^p-complete
VP	VETO PLAYER	co-NP-complete	-
MD	MUTUAL DEPENDENCE	co-NP-complete	-
GR	GOAL REALISABILITY	NP-complete	P
NG	NECESSARY GOAL	co-NP-complete	-
EG	EMPTY GAME	co-NP-complete	co-NP-complete
TG	TRIVIAL GAME	Π_2^p-complete	Π_2^p-complete
GU	GLOBAL UNATTAINABILITY	Σ_2^p-complete	NP
IG	INCOMPLETE GAME	D_2^p-complete	-

Coalitional Resource Games (CRGs)

- Problem:

  where does characteristic function come from?

- One answer provided by *Coalitional Resource Games (CRGs)*.
- Key ideas:
  - achieving a goal requires *expenditure of resources*;
  - each agent *endowed* with a profile of resources;
  - coalitions form to pool resource so as to achieve mutually satisfactory set of goals.

CRGs

A *coalitional resource game* Γ is an $(n+5)$-tuple:

$$\Gamma = \langle Ag, G, R, G_1, \ldots, G_n, \mathbf{en}, \mathbf{req} \rangle$$

where:

- $Ag = \{a_1, \ldots, a_n\}$ is a set of *agents*;
- $G = \{g_1, \ldots, g_m\}$ is a set of *possible goals*;
- $R = \{r_1, \ldots, r_t\}$ is a set of *resources*;
- for each $i \in Ag$, $G_i \subseteq G$ is a set of goals, as in QCGs;
- $\mathbf{en} : Ag \times R \to \mathbb{N}$ is an *endowment function*,
- $\mathbf{req} : G \times R \to \mathbb{N}$ is a *requirement function*.

Nine Decision Problems for CRGs

Problem	Complexity
SUCCESSFUL COALITION	NP-complete
MAXIMAL COALITION	co-NP-complete
NECESSARY RESOURCE	co-NP-complete
STRICTLY NECESSARY RESOURCE	D^p-complete
(C, G', r)-OPTIMAL	NP-complete
R-PARETO OPTIMALITY	co-NP-complete
SUCCESSFUL COALITION WITH RESOURCE BOUNDS	NP-complete
CONFLICTING COALITIONS	co-NP-complete
ACHIEVABLE GOAL SET	in P

QCG and CRG Equivalence

- We can define a notion of "equivalence" ($\equiv$) between QCGs and CRGs:

  $\Gamma_1 \equiv \Gamma_2$ means that QCG Γ_1 and CRG Γ_2 agree on the goal sets that are feasible for coalitions

- Given a QCG Γ_1 and CRG Γ_2, the problem of determining whether $\Gamma_1 \equiv \Gamma_2$ is co-NP-complete.

Can we translate between QCGs and CRGs?

Four questions suggest themselves:

1. Given a *crg*, Γ, is there always a QCG, Q_Γ such that $Q_\Gamma \equiv \Gamma$?
2. Given a *qcg*, Q, is there always a CRG, Γ_Q such that $\Gamma_Q \equiv Q$?
3. How "efficiently" can a given CRG be expressed as an equivalent QCG in those cases where such an equivalent structure exists?
4. How "efficiently" can a given QCG be expressed as an equivalent CRG in those cases where such an equivalent structure exists?

Translating CRGs $\rightarrow$ QCGs

- We can always translate a CRG into an equivalent QCG.
- More interestingly, we can do this *efficiently*:

  for every CRG Γ_1 there exists an equivalent QCG Γ_2 such that $|\Gamma_2| \leq |\Gamma_1|^2$.

Translating QCGs to CRGs

- We *cannot* always translate QCGs to equivalent CRGs.
- Moreover, even when we *can* translate, we can't always do it efficiently:

  there exist QCGs Γ for which equivalent CRGs exist but for which the size of the *smallest* equivalent CRG is at least $2^{|\Gamma|}$