

COMP310

Multi-Agent Systems

Chapter 13 - Forming Coalitions

Dr Terry R. Payne
Department of Computer Science



SECOND EDITION

An Introduction to

MultiAgent Systems

MICHAEL WOOLDRIDGE

Cooperative Game Theory

- So far we have taken a game theoretic view of multi-agent interactions
 - Prisoner's Dilemma suggests that cooperation should not occur, as the conditions required are not present:
 - Binding agreements are not possible
 - Utility is given to individuals based on individual action
- These constraints do not necessarily hold in the real world
 - Contracts, or collective payments can facilitate cooperation, leading to ***Coalition Games*** and ***Cooperative Game Theory***

Coalitional Games

- Coalitional games model scenarios where agents can benefit by cooperating.
 - Sandholm (et. al., 1999) identified the following stages:

Coalitional Structure Generation

Deciding in principle who will work together. It asks 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, and how to solve the “joint problem” of a coalition. It also involves finding how to maximise the utility of the coalition itself, and 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 Characteristic Function Game (CFG) is represented as the tuple: $G = \langle Ag, \nu \rangle$

$$Ag = \{1, \dots, n\} \quad \text{A set of agents}$$
$$\nu : 2^{Ag} \rightarrow \mathbb{R} \quad \text{the characteristic function of the game}$$

- From this, we form a coalition $C \subseteq Ag$
 - Singleton: where a coalition consists of a single member
 - Grand Coalition: where $C = Ag$ (i.e. all of the agents)
- Each coalition has a payoff value, defined by the characteristic function ν
 - i.e. if $\nu(C) = k$ then the coalition will get the payoff k if they cooperate on some task

```
% Representation of a Simple  
% Characteristic Function Game
```

```
% List of Agents
```

```
1,2,3
```

```
% Characteristic Function
```

```
1 = 5
```

```
2 = 5
```

```
3 = 5
```

```
1,2 = 10
```

```
1,3 = 10
```

```
2,3 = 10
```

```
1,2,3 = 25
```


Characteristic Function Games


- The objective is to join a coalition that the agent cannot object to
 - This involves calculating the characteristic function for different games
- Sandholm (1999) showed that:
 - If the game is **superadditive**: if $v(U) + v(V) < v(U \cup V)$
 - The coalition that maximises social welfare is the **Grand Coalition**
 - If the game is **subadditive**: if $v(U) + v(V) > v(U \cup V)$
 - The coalitions that maximise social welfare are **singletons**
 - However as some games are neither subadditive or superadditive:
 - **the characteristic function value calculations need to be determined for each of the possible coalitions!**
 - This is exponentially complex

Which Coalition Should I Join?

- Assuming that we know *the characteristic function* and the *payoff vector*, what coalition should an agent join?
 - An *outcome* x for a coalition C in game $\langle Ag, \nu \rangle$ is a *vector of payoffs* to members of C , such that $x = \langle x_1, \dots, x_k \rangle$ which represents an *efficient distribution of payoff* to members of Ag
 - Where “*efficient*” means:
$$\nu(C) = \sum_{i \in C} x_i$$
 - Example: if $\nu(\{1, 2\}) = 20$, then possible outcomes are: $\langle 20, 0 \rangle, \langle 19, 1 \rangle, \langle 18, 2 \rangle \dots \langle 1, 19 \rangle, \langle 0, 20 \rangle$
- Thus, the agent should only join a coalition C which is:
 - *Feasible*: the coalition C really could obtain some payoff than an agent could not object to; and
 - *Efficient*: all of the payoff is allocated

Which Coalition Should I Join?

- However, there may be many coalitions
 - Each has a different characteristic function
 - Agents prefer coalitions that are as productive as possible
 - Therefore a coalition will only form if all the members prefer to be in it
 - I.e. they don't defect to a more preferable coalition
- Therefore:
 - “***which coalition should I join?***” can be reduced to “***is the coalition stable?***”
 - Is it rational for all members of coalition C to stay with C , or could they benefit by defecting from it?
 - There's no point in me joining a coalition with you, unless you want to form one with me, and vice versa.



Stability and the Core

- Stability can be reduced to the notion of the **core**
 - Stability is a **necessary** but not **sufficient** condition for coalitions to form
 - i.e. Unstable coalitions will never form, but a stable coalition isn't guaranteed to form
- 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
 - Intuitively, a coalition C **objects to an outcome** if there is some other outcome that makes all of them **strictly better off**
 - Formally, $C \subseteq Ag$ objects to an outcome $x = \langle x_1, \dots, x_n \rangle$ for the grand coalition if there is some outcome $x' = \langle x'_1, \dots, x'_k \rangle$ for C such that: $x'_i > x_i$ for all $i \in C$
- The idea is that an outcome is not going to happen if somebody objects to it!
 - i.e. if the core is empty, then no coalition can form

The Core and Fair Payoffs

- Sometimes the core is non-empty but is it “fair”?
 - Suppose we have $Ag = \{1, 2\}$, with the following Characteristic Function:
 - $v(\{1\}) = 5$
 - $v(\{2\}) = 5$
 - $v(\{1,2\}) = 20$
 - The outcome $\langle 20, 0 \rangle$ (i.e., agent 1 gets everything) **will not be in the core**, since agent 2 can object; **by working on its own it can do better**, because $v(\{2\}) = 5$
 - However, outcome $\langle 14, 6 \rangle$ is in the core, as **agent 2 gets more than working on its own**, and thus has no objection.
- But **is it “fair”** on agent 2 to get only a payoff of 6, if agent 1 gets 14???

Sharing the 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.***
 - The Shapley value is one that satisfies the axioms opposite!

Symmetry

Agents that make the same contribution should get the same payoff. I.E. the amount an agent gets should only depend on their contribution.

Dummy Player

These are agents that never have any synergy with any coalition, and thus only get what they can earn on their own.

Additivity

If two games are combined, the value an agent gets should be the sum of the values it gets in the individual games.

Marginal Contribution

- The **Shapley value** for an agent is based on the marginal contribution of that agent to a coalition (for all permutations of coalitions)
- Let $\delta_i(C)$ be the amount that agent i adds by joining a coalition $C \subseteq Ag$
 - i.e. the **marginal contribution** of i to C is defined as $\delta_i(C) = \nu(C \cup \{i\}) - \nu(C)$
 - Note that if $\delta_i(C) = \nu(\{i\})$ then there is **no added value** from i joining C since the amount i adds is the same as if i would earn on its own
- The Shapley value for i , denoted φ_i , is the value that agent i in Ag is given in the game $\langle Ag, \nu \rangle$

Shapley Axioms: Symmetry



- *Agents that make the same contribution should get the same payoff*
 - The amount an agent gets should only depend on their contribution
 - Agents i and j are interchangeable if $\delta_i(C) = \delta_j(C)$ for every $C \subseteq Ag \setminus \{i, j\}$
- The symmetry axiom states:
 - *If i and j are interchangeable, then $\varphi_i = \varphi_j$*


Shapley Axioms: Dummy Player



- *Agents that never have any synergy with any coalition, and thus only get what they can earn on their own.*
- An agent is a dummy player if $\delta_i(C) = \nu(\{i\})$ for every $C \subseteq Ag \setminus \{i\}$
 - i.e. an agent only adds to a coalition what it could get on its own
- The dummy player axiom states:
 - *If i is a dummy player, then $\varphi_i = \nu(\{i\})$*

Shapley Axioms: Additivity

- *If two games are combined, the value an agent gets should be the sum of the values it gets in the individual games*
 - I.e. an agent doesn't gain or lose by playing more than once
 - Let $G^1 = \langle Ag, \nu^1 \rangle$ and $G^2 = \langle Ag, \nu^2 \rangle$ be games with the same agents
 - Let $i \in Ag$ be one of the agents
 - Let φ^1_i and φ^2_i be the value agent i receives in games G^1 and G^2 respectively
 - Let $G^{1+2} = \langle Ag, \nu^{1+2} \rangle$ be the game such that $\nu^{1+2}(C) = \nu^1(C) + \nu^2(C)$
- The additivity axiom states:
 - *The value φ^{1+2}_i of agent i in game G^{1+2} should be $\varphi^1_i + \varphi^2_i$*



Shapley value

- Recall that we stated:
 - The **Shapley value** for an agent is based on the marginal contribution of that agent to a coalition (**for all permutations of coalitions**)
 - The marginal contribution can be dependent on the **order** in which an agent joins a coalition
 - This is because an agent may have a larger contribution if it is the first to join, than if it is the last!
- For example, if $Ag = \{1,2,3\}$ then the set of all possible orderings, $\Pi(Ag)$ is given as
 - $\Pi(Ag) = \{(1,2,3), (1,3,2), (2,1,3), (2,3,1), (3,1,2), (3,2,1)\}$
- We have defined the **marginal contribution** of i to C as $\delta_i(C) = \nu(C \cup \{i\}) - \nu(C)$
- The **Shapley value** for i is defined as:
$$\varphi_i = \frac{\sum_{o \in \Pi(Ag)} \delta_i(C_i(o))}{|Ag|!}$$

Shapley Example

- Suppose we have $Ag = \{1, 2\}$, with the following characteristic function

$$\nu(\{1\}) = 5$$

$$\nu(\{2\}) = 10$$

$$\nu(\{1, 2\}) = 20$$

- We can now calculate the marginal contribution $\delta_i(C)$ of each agent $i \in C$, for each coalition $C \subseteq Ag$

$$\delta_1(\emptyset) = \nu(\emptyset \cup \{1\}) - \nu(\emptyset) = (5 - 0) = 5$$

$$\delta_1(\{2\}) = \nu(\{2\} \cup \{1\}) - \nu(\{2\}) = (20 - 10) = 10$$

$$\delta_2(\emptyset) = \nu(\emptyset \cup \{2\}) - \nu(\emptyset) = (10 - 0) = 10$$

$$\delta_2(\{1\}) = \nu(\{1\} \cup \{2\}) - \nu(\{1\}) = (20 - 5) = 15$$

- Finally, we can calculate the individual Shapley values for each i :

$$\varphi_1 = \frac{\delta_1(\emptyset) + \delta_1(\{2\})}{|Ag|!} = \frac{5 + 10}{2} = 7.5$$

$$\varphi_2 = \frac{\delta_2(\emptyset) + \delta_2(\{1\})}{|Ag|!} = \frac{10 + 15}{2} = 12.5$$

Shapley Value (reminder)

Marginal Contribution:

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

Shapley value:

$$\varphi_i = \frac{\sum_{o \in \Pi(Ag)} \delta_i(C_i(o))}{|Ag|!}$$

Representing Coalitional Games

- It is important for an agent to know if the core of a coalition is non-empty
 - Problem: a naive, obvious representation of a 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
 - An n -player game consists of $2^n - 1$ coalitions
 - e.g. a 100-player game would require 1.2×10^{30} lines

```
% Representation of a Simple  
% Characteristic Function Game
```

```
% List of Agents
```

```
1,2,3
```

```
% Characteristic Function
```

```
1 = 5
```

```
2 = 5
```

```
3 = 5
```

```
1,2 = 10
```

```
1,3 = 10
```

```
2,3 = 10
```

```
1,2,3 = 25
```

Representing Characteristic Functions?

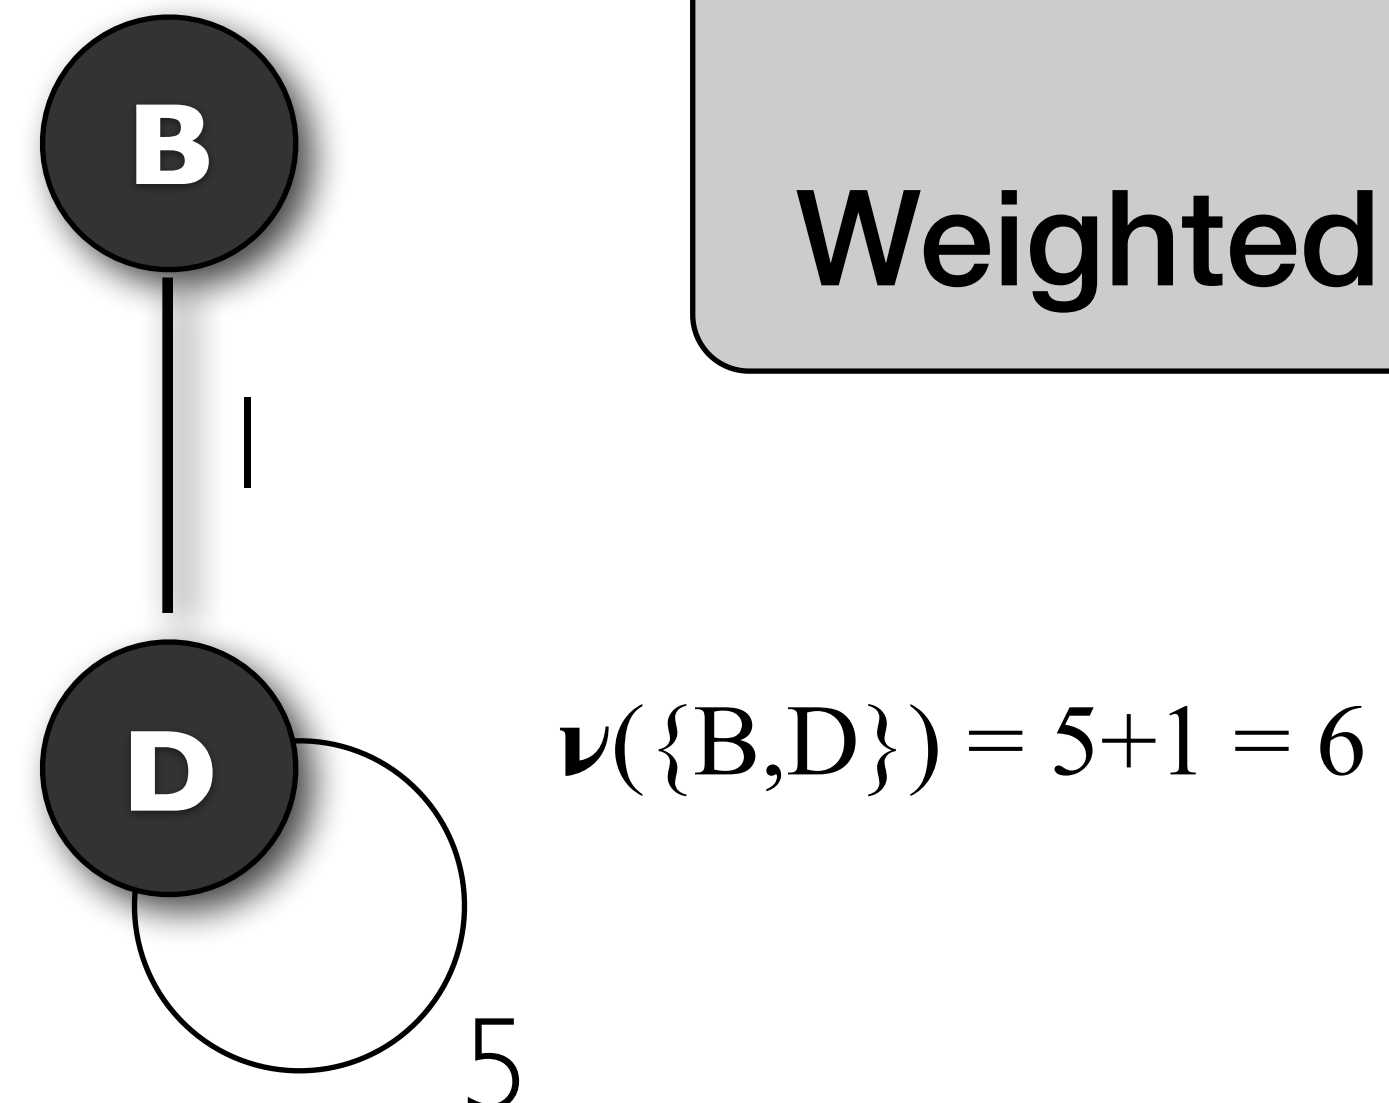
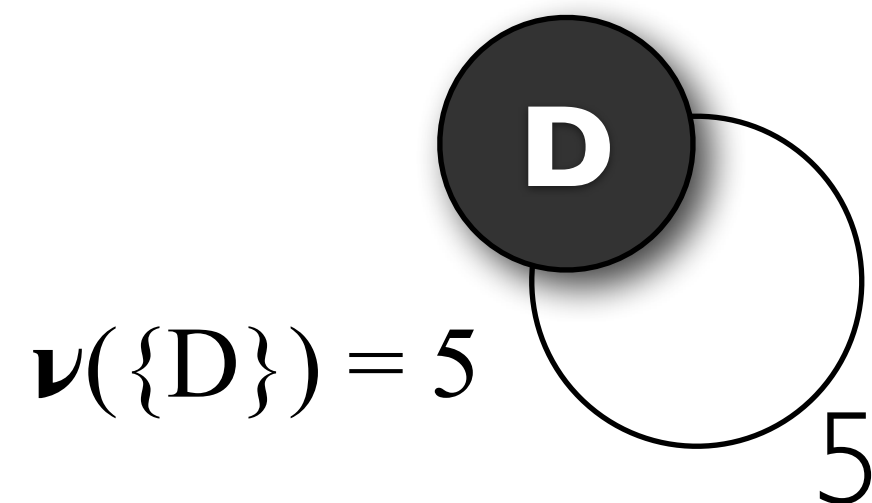
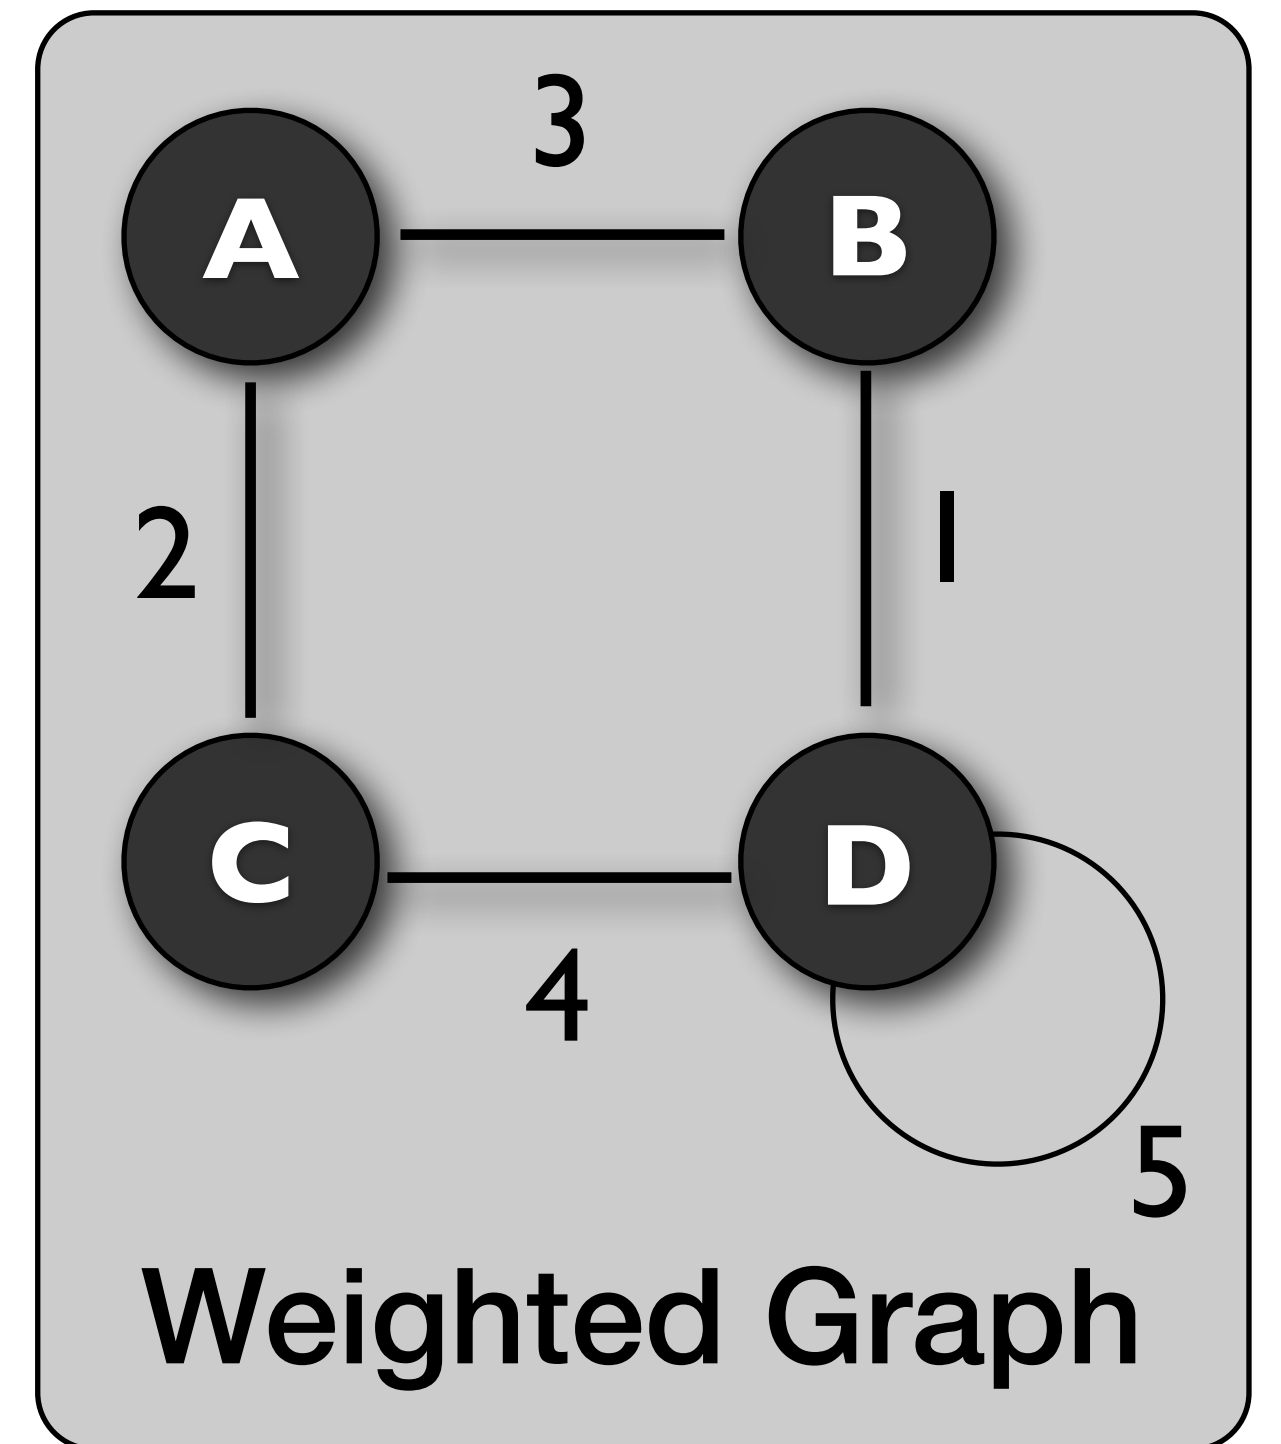

- Two approaches to this problem:
 - try to find a **complete** representation that is succinct in “most” cases
 - try to find a representation that is **not complete but is always succinct**
- A common approach:
 - interpret characteristic function over a combinatorial structure.
- We look at two possible approaches:
 - Induced Subgraph and Marginal Contribution Networks

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 is 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



Induced Subgraph



$$\nu(\{A, B, C\}) = 3 + 2 = 5$$

$$\varphi_A = \frac{1}{2} \sum_{j \neq i} \mathcal{W}_{i,j} = \frac{3 + 2}{2} = 2.5$$

$$\varphi_B = \frac{1}{2} \sum_{j \neq i} \mathcal{W}_{i,j} = \frac{3}{2} = 1.5$$

$$\varphi_C = \frac{1}{2} \sum_{j \neq i} \mathcal{W}_{i,j} = \frac{2}{2} = 1$$

- Representation is **succinct**, but **not complete**
 - there are characteristic functions that cannot be captured using this representation
- Determining emptiness of the core is NP-complete
 - Checking whether a specific distribution is in the core is co-NP-complete
- Shapley value can be calculated in polynomial time

$$\varphi_i = \frac{1}{2} \sum_{j \neq i} \mathcal{W}_{i,j}$$

- i.e. an agent gets **half the income from the edges in the graph to which it is attached.**

Marginal Contribution Nets

- Characteristic function ν represented as rules:

pattern \longrightarrow 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.

Rule set (rs) 2:

$$a \wedge b \rightarrow 5$$

$$b \rightarrow 2$$

$$c \rightarrow 4$$

$$b \wedge \neg c \rightarrow -2$$

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

- Example (rule set 1):

$$a \wedge b \rightarrow 5$$

$$b \rightarrow 2$$

- We have: $\nu_{rs1}(\{a\}) = 0$, $\nu_{rs1}(\{b\}) = 2$, and $\nu_{rs1}(\{a, b\}) = 5+2 = 7$.

$$\nu_{rs2}(\{a\}) = 0$$

no rules apply

$$\nu_{rs2}(\{b\}) = 2 + -2 = 0$$

2nd and 4th rules

$$\nu_{rs2}(\{c\}) = 4$$

3rd rule

$$\nu_{rs2}(\{a, b\}) = 5 + 2 + -2 = 5$$

1st, 2nd and 4th rules

$$\nu_{rs2}(\{a, c\}) = 4$$

3rd rule

$$\nu_{rs2}(\{b, c\}) = 2 + 4 = 6$$

2nd and 3rd rules

$$\nu_{rs2}(\{a, b, c\}) = 5 + 2 + 4 = 11$$

1st, 2nd and 3rd rules

- We can also allow negations in rules (i.e. for when an agent is not present).

Marginal Contribution Nets

- Calculating the **Shapley value for marginal contribution nets** is similar to that for induced subgraphs
 - Again, Shapley's symmetry axiom applies to each agent
 - The contributions from agents in the same rule is equal
 - The additivity property means that:
 - we calculate the Shapley value for each rule
 - sum over the rules to calculate the Shapley value for each agent
- Handling negative values requires a different method

Calculating the Shapley Value

$$\varphi_i = \sum_{r \in rs; i \text{ occurs in lhs of } r} \varphi_i^r$$

where

$$\varphi_i^{1 \wedge \dots \wedge l \rightarrow x} = \frac{x}{l}$$

$$\begin{aligned} a \wedge b &\rightarrow 5 \\ b &\rightarrow 2 \\ c &\rightarrow 4 \end{aligned}$$

$$\varphi_A = \sum_{r \in rs; A \text{ occurs in lhs of } r} \varphi_A^r = \frac{5}{2} = 2.5$$

$$\varphi_B = \sum_{r \in rs; B \text{ occurs in lhs of } r} \varphi_B^r = \frac{5}{2} + 2 = 4.5$$

$$\varphi_C = \sum_{r \in rs; C \text{ occurs in lhs of } r} \varphi_C^r = 4$$

Coalition Structure Generation

- In addition to representing the characteristic function, there is the challenge of calculating them!
 - Remember, for a set of n agents in Ag , there will be $2^n - 1$ distinct coalitions
- Shehory & Kraus (1998) proposed a method whereby agents ***distributed the calculation amongst themselves***
 - Resulted in a communication overhead, in coordinating which agent calculated the characteristic function value for which coalition
 - Rahwan & Jennings (2007) proposed the DVCD approach for allocating coalition value calculations to agents without the need for communication
 - However, agents could be incentivised to mis-represent the calculations for those coalitions in which they were not a member
 - This was resolved by Riley, Atkinson, Dunne & Payne (2015) through the use of ***(n,s)-sequences***

(n,s)-sequences

- Riley et.al. proposed a mechanism for calculating ***the coalition value calculation share*** for an agent, based on the agent's id
 - Given a set of agents in Ag where $n=|Ag|$, the agents are labelled $1 \dots n$
 - A coalition of size $1 \leq s < n$ can be generated given an (n,s)-sequence \underline{t} by first calculating the aggregate offset for each position in \underline{t} and given the agent x , determine its coalition value calculation share (for some $1 \leq x \leq n$):

$$x_i \equiv \begin{cases} x & \text{if } i = 1 \\ (x + \sum_{k=0}^{i-2} (t_k + 1)) \bmod n & \text{if } 2 \leq i \leq s \end{cases}$$

- Note that the result is an agent in the range $(1 \leq x \leq n)$
 - i.e. the result is congruent modulo n

Example

- The (n,s) -sequences for coalitions of size $s=3$, for a set of agents $Ag = \{1,2,3,4,5,6\}$ are $\langle 0,0,3 \rangle$, $\langle 0,1,2 \rangle$, $\langle 0,2,1 \rangle$, $\langle 1,1,1 \rangle$
- These are used to generate the coalition value calculation shares for each agent x
 - If each agent generates their share...
 - ... all of the coalitions of size s will be generated
- Duplications occur if there is a repeated periodic sub-sequence in the (n,s) -sequence (e.g. $\langle 1,1,1 \rangle$)
 - If $s=4$, then $\langle 0,1,0,1 \rangle$ has the repeating sub-sequence $\langle \dots 0,1 \dots \rangle$, but $\langle 0,0,1,1 \rangle$ has no repeating sequence
- By tracking which agent generates coalitions from the repeated sequence, duplications can be eliminated

Generating Coalition Value Calculation Shares

The following table lists the coalitions generated for the (n,s) -sequences where $n=6$ and $s=3$

	$\langle 0,0,3 \rangle$	$\langle 0,1,2 \rangle$	$\langle 0,2,1 \rangle$	$\langle 1,1,1 \rangle$
CV^3_1	1,2,3	1,2,4	1,2,5	
CV^3_2	2,3,4	2,3,5	2,3,6	
CV^3_3	3,4,5	3,4,6	3,4,1	
CV^3_4	4,5,6	4,5,1	4,5,2	4,6,2
CV^3_5	5,6,1	5,6,2	5,6,3	5,1,3
CV^3_6	6,1,2	6,1,3	6,1,4	

$$x_i \equiv \begin{cases} x & \text{if } i = 1 \\ (x + \sum_{k=0}^{i-2} (t_k + 1)) \bmod n & \text{if } 2 \leq i \leq s \end{cases}$$

Example

- The (n,s) -sequences for coalitions of size $s=3$, for a set of agents $Ag = \{1,2,3,4,5,6\}$ are $\langle 0,0,3 \rangle$, $\langle 0,1,2 \rangle$, $\langle 0,2,1 \rangle$, $\langle 1,1,1 \rangle$
- Therefore, if we consider agent 5:
 - $C(5, \langle 0,0,3 \rangle) \equiv \{5, 6, 1\}$
 - i.e. $\{5, (5+0+1) \bmod 6, ((5+0+1) + 0 + 1) \bmod 6\} \equiv \{5, 6, 1\}$
 - $C(5, \langle 0,1,2 \rangle) \equiv \{5, 6, 2\}$
 - i.e. $\{5, (5+0+1) \bmod 6, ((5+0+1) + 1 + 1) \bmod 6\} \equiv \{5, 6, 2\}$
 - $C(5, \langle 0,2,1 \rangle) \equiv \{5, 6, 3\}$
 - i.e. $\{5, (5+0+1) \bmod 6, ((5+0+1) + 2 + 1) \bmod 6\} \equiv \{5, 6, 3\}$
 - $C(5, \langle 1,1,1 \rangle) \equiv \{5, 1, 3\}$
 - i.e. $\{5, (5+1+1) \bmod 6, ((5+1+1) + 1 + 1) \bmod 6\} \equiv \{5, 1, 3\}$

Generating Coalition Value Calculation Shares

The following table lists the coalitions generated for the (n,s) -sequences where $n=6$ and $s=3$

	$\langle 0,0,3 \rangle$	$\langle 0,1,2 \rangle$	$\langle 0,2,1 \rangle$	$\langle 1,1,1 \rangle$
CV^3_1	1,2,3	1,2,4	1,2,5	
CV^3_2	2,3,4	2,3,5	2,3,6	
CV^3_3	3,4,5	3,4,6	3,4,1	
CV^3_4	4,5,6	4,5,1	4,5,2	4,6,2
CV^3_5	5,6,1	5,6,2	5,6,3	5,1,3
CV^3_6	6,1,2	6,1,3	6,1,4	

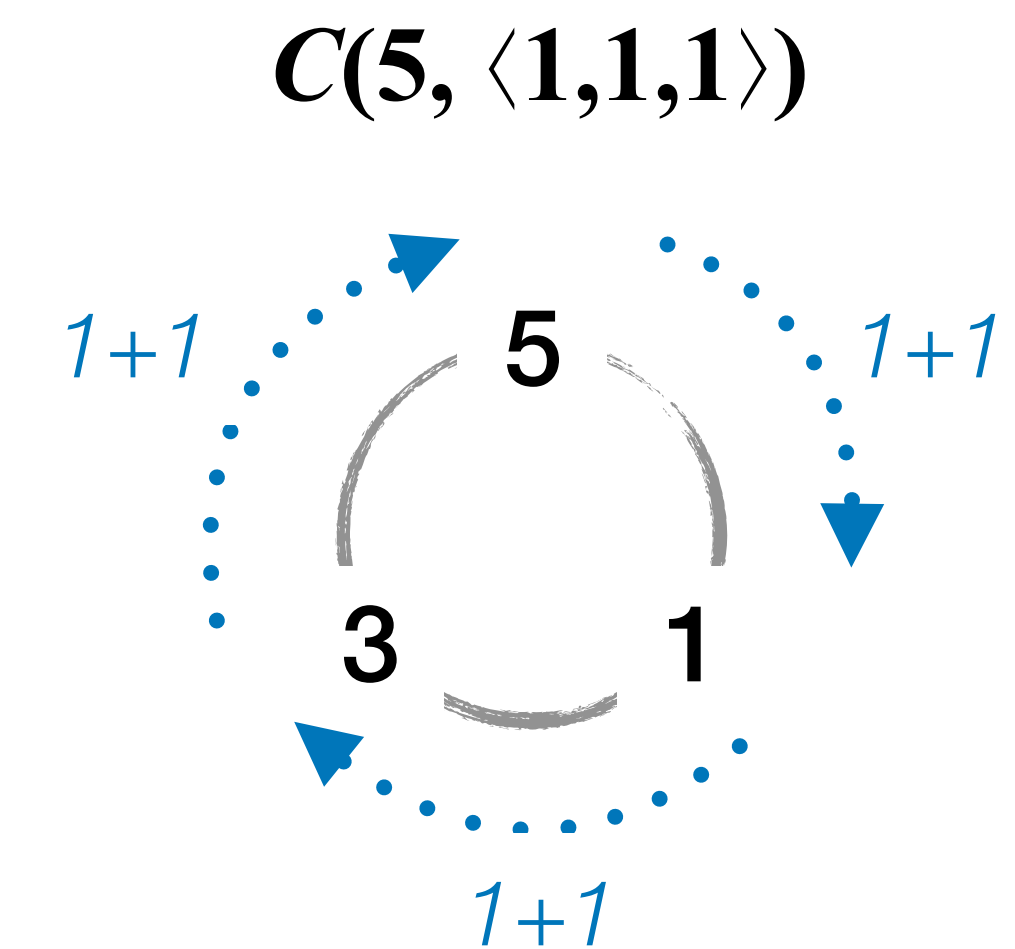

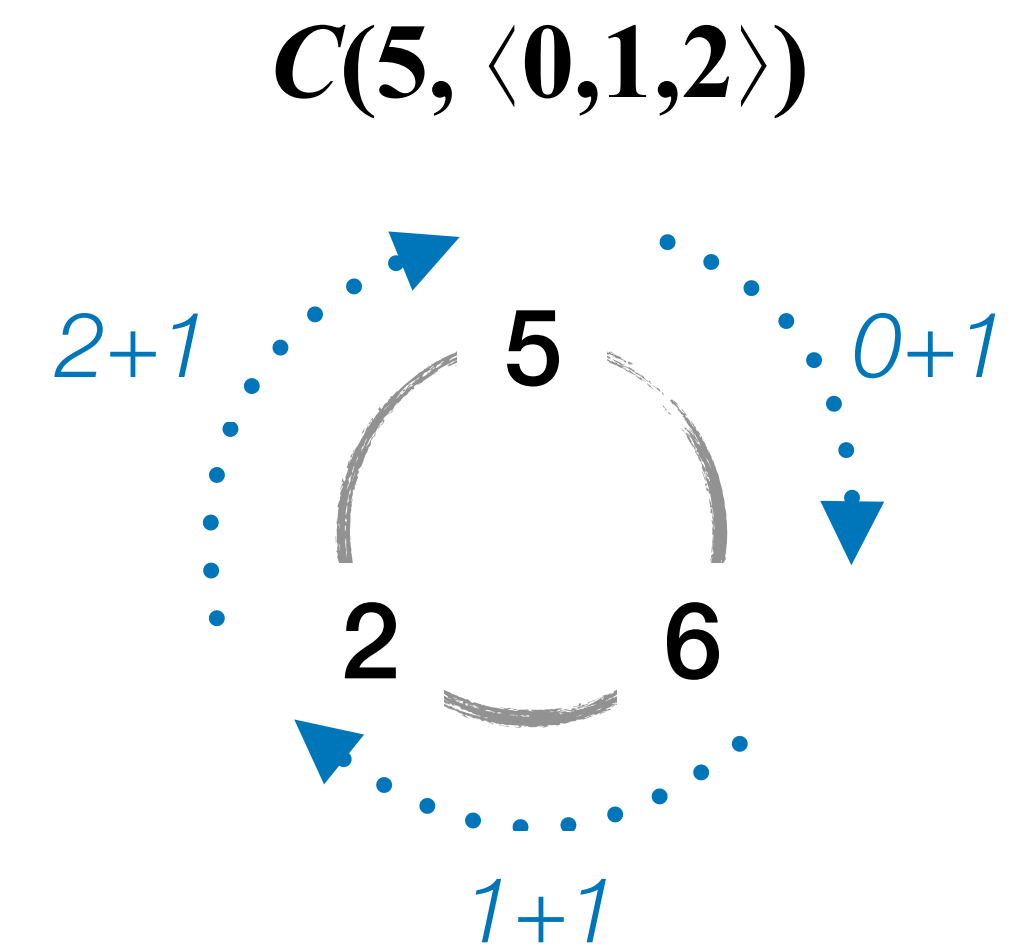
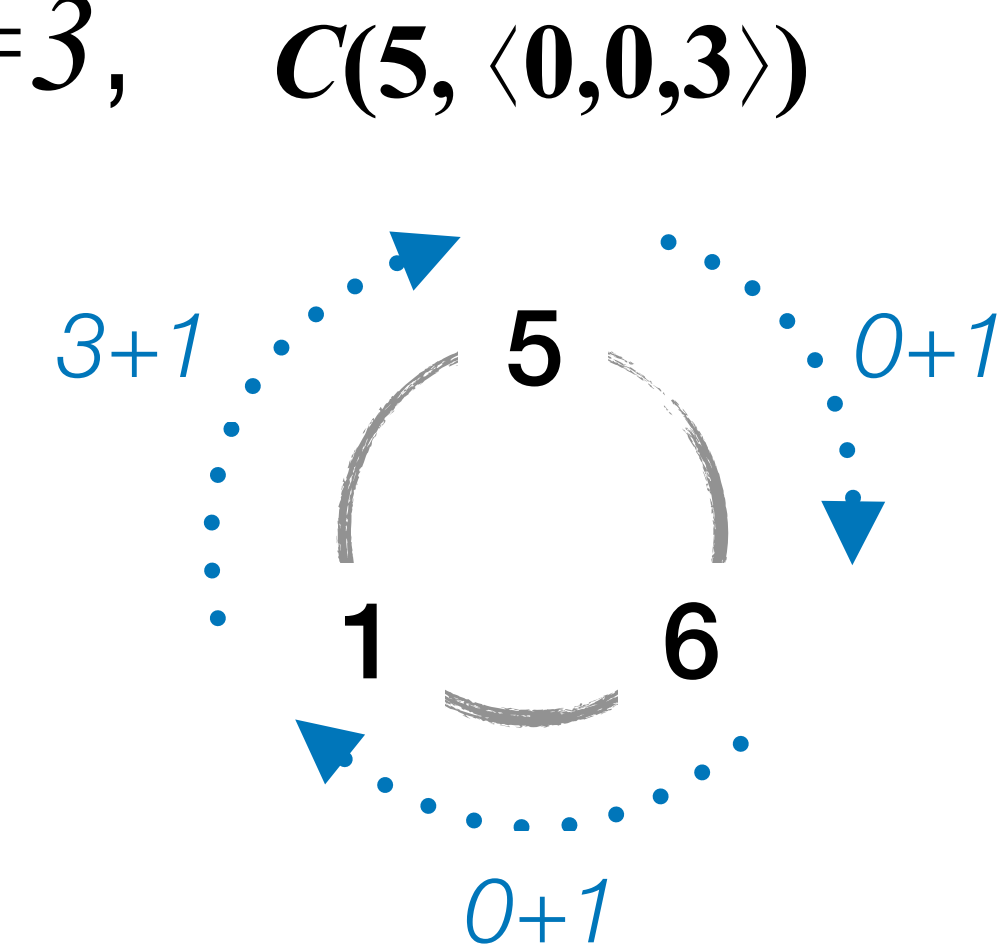
$$x_i \equiv \begin{cases} x & \text{if } i = 1 \\ (x + \sum_{k=0}^{i-2} (t_k + 1)) \bmod n & \text{if } 2 \leq i \leq s \end{cases}$$

Example

- The (n,s) -sequences for coalitions of size $s=3$, for a set of agents $Ag = \{1,2,3,4,5,6\}$ are $\langle 0,0,3 \rangle$, $\langle 0,1,2 \rangle$, $\langle 0,2,1 \rangle$, $\langle 1,1,1 \rangle$

- Therefore, if we consider agent 5:

- $C(5, \langle 0,0,3 \rangle) \equiv \{5, 6, 1\}$
 - i.e. $\{5, (5+0+1) \bmod 6, ((5+0+1) + 0 + 1) \bmod 6\} \equiv \{5, 6, 1\}$
- $C(5, \langle 0,1,2 \rangle) \equiv \{5, 6, 2\}$
 - i.e. $\{5, (5+0+1) \bmod 6, ((5+0+1) + 1 + 1) \bmod 6\} \equiv \{5, 6, 2\}$
- $C(5, \langle 0,2,1 \rangle) \equiv \{5, 6, 3\}$
 - i.e. $\{5, (5+0+1) \bmod 6, ((5+0+1) + 2 + 1) \bmod 6\} \equiv \{5, 6, 3\}$
- $C(5, \langle 1,1,1 \rangle) \equiv \{5, 1, 3\}$
 - i.e. $\{5, (5+1+1) \bmod 6, ((5+1+1) + 1 + 1) \bmod 6\} \equiv \{5, 1, 3\}$



Summary

- In this lecture we have looked at mechanisms for identifying coalitions.
 - The notion of a stable coalition game was presented, through the idea of a **Core**.
 - The **Shapley Value** was then introduced, to determine the contribution that different agents may have on a coalition.
- The problem of representing coalitional games and characteristic functions was then discussed, including:
 - Induced Subgraphs
 - Marginal Contribution Nets.
- We finally talked about Coalition Structure Generation
- This is again an active research area, especially from a game-theoretic and computational complexity perspective.

Class Reading (Chapter 13):

“Marginal contribution nets: A compact representation scheme for coalition games”, S. Leong and Y. Shoham. Proceedings of the Sixth ACM Conference on Electronic Commerce (EC’05), Vancouver, Canada, 2005.

This is a technical article (but a very nice one), introducing the marginal contribution nets scheme.