Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

A Graphical Representation for Games in Partition Function Form

Oskar Skibski[1], Tomasz P. Michalak[2,3], Yuko Sakurai[1,4], Michael Wooldridge[2], Makoto Yokoo[1]
[1]Department of Informatics, Kyushu University, Japan
[2]Department of Computer Science, University of Oxford, UK
[3]Institute of Informatics, University of Warsaw, Poland
[4]JST PRESTO, Japan

Abstract

We propose a novel representation for coalitional games with externalities, called *Partition Decision Trees*. This representation is based on rooted directed trees, where non-leaf nodes are labelled with agents' names, leaf nodes are labelled with payoff vectors, and edges indicate membership of agents in coalitions. We show that this representation is fully expressive, and for certain classes of games significantly more concise than an extensive representation. Most importantly, Partition Decision Trees are the first formalism in the literature under which most of the direct extensions of the Shapley value to games with externalities can be computed in polynomial time.

Introduction

Many models of coalitional games considered in the computer science literature assume that when a coalition of agents forms, it does not impact other coalitions in the system (Chalkiadakis et al. 2011). This assumption does not hold in many important settings. For example, oligopolistic markets are a common economic setting in which coalitions have a non-negligible influence on one other. This influence, called *externalities*, is also commonplace in task-based multi-agent systems with overlapping agent preferences. If a coalition achieves a desired goal, then this goal may become unachievable by other coalitions (Rahwan et al. 2012; Yokoo et al. 2005; Bachrach and Rosenschein 2008).

Coalitional games with externalities raise considerable game-theoretic and computational challenges. Despite many efforts, there is currently no consensus in the literature on how to extend key game-theoretic solution concepts to this more complex setting. The concept that has attracted the most attention is the Shapley value—the unique payoff division scheme for coalitional games without externalities that satisfies certain intuitive axioms (Shapley 1953). Unfortunately, extending the Shapley value to games with externalities is non-trivial, since the original axioms are too weak to guarantee uniqueness. Among a number of proposals,[1] six extensions are especially appealing: *the externality-free value* (Do and Norde 2007; de Clippel and Serrano 2008), *the McQuillin value* (McQuillin 2009; Skibski 2011), *the Bolger value* (Bolger 1989), *Macho-Stadler et al. value* (Macho-Stadler et al. 2007), *the Hu-Yang value* (Hu and Yang 2010), and *the Myerson value* (Myerson 1977). Unlike other extensions, these preserve the spirit of the original Shapley axiomatization, and will be the focus of this paper.

The standard formalism for coalitional games with externalities, called *Partition Function Games* (Thrall and Lucas 1963), is computationally challenging. It requires the consideration of all partitions of agents, i.e., *coalition structures*, into coalitions and the specification of the values of all coalitions for every coalition structure in which they are embedded, the complexity of which is $\omega(n^{\frac{1}{2}n})$ and $O(n^n)$ (where n is the number of agents) (Sandholm et al. 1999).

Given that the Partition Function Games model has prohibitive time and space requirements, alternative representations of games with externalities were developed in the computer science literature (Michalak et al. 2009; 2010a; 2010b; Ichimura et al. 2011). These representations aim to balance expressiveness, compactness, and computational tractability. Unfortunately, the formalisms designed to facilitate efficient computation of the Shapley value extensions to games with externalities (see Related Work for more details), facilitate polynomial algorithms only for the externality-free value and the McQuillin value. In other words, no representation that allows for efficient computation of the remaining extensions has been discovered so far.

In this paper, we address this issue. We propose a novel representation for coalitional games with externalities, called *Partition Decision Trees*. This representation is based on rooted directed trees, in which non-leaf nodes are labelled with agents' names, leaf nodes are labelled with payoff vectors, and edges indicate membership of agents in coalitions. We show that this representation is fully expressive, i.e., it can represent any game with externalities. It can also be significantly more concise than the extensive representation. Most importantly, our representation is the first concise formalism in the literature that facilitates polynomial computation of almost all the aforementioned extensions of the Shapley value. In particular, the externality-free, the McQuillin, the Macho-Stadler et al. and the Myerson values can all be computed in time $O(n \times |\mathcal{T}|)$, whereas the Hu-Yang value can be computed in time $O(n^3 \times |\mathcal{T}|)$, where $|\mathcal{T}|$ is the size of the representation.

[1]See Skibski et al. (2013) for a recent overview.

Preliminaries

Let $N = \{a, b, c, \dots\}$ be a finite set of agents. A game $\hat{v}$ in *characteristic function form* (without externalities) is a function that assigns a real-valued payoff to every non-empty subset (coalition) of agents: $\hat{v} : 2^N \to \mathbb{R}$, with $\hat{v}(\emptyset) = 0$.

In games with externalities, the value of a coalition may depend on how other agents have organized themselves into coalitions. A *partition* of N is a set of disjoint coalitions that collectively cover N. A coalition S that is part of a partition P is called an *embedded coalition* in P, and is denoted by (S, P). The set of all partitions and the set of all embedded coalitions are denoted by $\mathcal{P}(N)$ and $EC(N)$, respectively. A game v in *partition function form* (with externalities) is a function that assigns a real-valued payoff to every embedded coalition: $v : EC(N) \to \mathbb{R}$, where $v(\emptyset, P) = 0$ for all P.

A *value* of a game is a vector that distributes the payoff of the *grand coalition*, i.e., $v(N)$ or $v(N, \{N, \emptyset\})$, among the agents. While, in principle, any distribution is admissible, we seek those that meet certain desirable criteria (or axioms). In particular, let us consider the following axioms:[2]

- *Efficiency*—the entire available payoff is distributed among agents: $\sum_{i \in N} \varphi_i(v) = v(N, \{N, \emptyset\})$ for every v;
- *Symmetry*—payoffs do not depend on the agents' names: $\varphi_i(f(v)) = \varphi_{f(i)}(v)$ for every v and every bijection $f : N \to N$, where $f(v)(S, P) \overset{\text{def}}{=} v(f(S), \{f(T) \mid T \in P\})$ and $f(S) \overset{\text{def}}{=} \{f(i) \mid i \in S\}$;
- *Additivity*—the sum of payoffs in two separate games equals the payoff in the combined game: $\varphi(\beta_1 v_1 + \beta_2 v_2) = \beta_1 \varphi(v_1) + \beta_2 \varphi(v_2)$ for all v_1, v_2 and scalars $\beta_1, \beta_2 \in \mathbb{R}$, $(v_1 + v_2)(S, P) \overset{\text{def}}{=} v_1(S, P) + v_2(S, P)$ and $(\beta v)(S, P) \overset{\text{def}}{=} \beta \cdot v(S, P)$;[3]
- *Null-Player Axiom*—agents with no impact on the value of any coalition should get nothing: $\forall_{(S,P) \in EC(N), i \in S} \forall_{T \in P \cup \{\emptyset\}} v(S, P) - v(S_{-i}, \tau_i^T(P)) = 0 \Rightarrow \varphi_i(v) = 0$ for every v and $i \in N$, where $\tau_i^T(P) \overset{\text{def}}{=} P \setminus \{S, T\} \cup \{S \setminus \{i\}, T \cup \{i\}\}$.

Shapley (1953) famously proved that in games without externalities there exists a unique value that satisfies all four axioms. This value can be obtained through the following process. Assume that agents enter the grand coalition in a random order. As an agent i enters, he receives a payoff that equals his marginal contribution to the group of agents that he joins: $v(S \cup \{i\}) - v(S)$. Then, the Shapley value is the expected outcome of agents' contributions over all orders:

$$SV_i(\hat{v}) \overset{\text{def}}{=} \frac{1}{|N|!} \sum_{\pi \in \Omega(N)} \hat{v}(C_i^\pi \cup \{i\}) - \hat{v}(C_i^\pi), \quad (1)$$

where $\Omega(N)$ is the set of all orders (permutations of N), and C_i^π is the set of agents that appear in permutation π after i. Let us present the Shapley value in a more concise form. Let $\zeta_S^i \overset{\text{def}}{=} \frac{(|S|-1)!(|N|-|S|)!}{|N|!}$ if $i \in S$ and $\zeta_S^i \overset{\text{def}}{=} -\frac{|S|!(|N|-|S|-1)!}{|N|!}$, otherwise. We have:

$$SV_i(\hat{v}) \overset{\text{def}}{=} \sum_{S \subseteq N} \zeta_S^i \cdot \hat{v}(S). \quad (2)$$

In our proofs we will also need a class of elementary games $\langle e^{(S,P)} \rangle_{(S,P) \in EC(N)}$, where only one coalition embedded in one partition has a non-zero value. Formally, for every embedded coalition $(S, P) \in EC(N)$, we define:

$$e^{(S,P)}(\tilde{S}, \tilde{P}) \overset{\text{def}}{=} \begin{cases} 1 & \text{if } (\tilde{S}, \tilde{P}) = (S, P), \\ 0 & \text{otherwise,} \end{cases}$$

for every embedded coalition $(\tilde{S}, \tilde{P}) \in EC(N)$.

Extended Shapley values: Skibski et al. (2013) showed that only the six values from the literature that are listed in the introduction satisfy a straightforward translation of Shapley's axioms to games with externalities. The authors also showed that these values can be obtained using a process that is similar to the one presented by Shapley. Specifically, assume that the agents leave the grand coalition in a random order and divide themselves into coalitions outside. In each step, one agent departs and with some weight[4] enters an existing coalition outside, or forms a new coalition. As a result of leaving the (grand) coalition, the agent is granted his *elementary marginal contribution*, *i.e.*, the loss of the coalition that he left. Now, the agent's Shapley value extended to games with externalities is the expected marginal contribution of this agent, i.e. the weighted average of his elementary marginal contributions over all permutations.

The formulas for the six aforementioned values are:

- *the externality-free value* (Do and Norde 2007; de Clippel and Serrano 2008)—an agent leaving the coalition always forms a new one:
$$\varphi_i^{free}(v) \overset{\text{def}}{=} \sum_{S \subseteq N} \zeta_S^i \cdot v(S, \{S\} \cup \{\{j\} \mid j \in N \setminus S\}). \quad (3)$$
- *the McQuillin value* (McQuillin 2009; Skibski 2011)—an agent always chooses an existing coalition outside:
$$\varphi_i^{McQ}(v) \overset{\text{def}}{=} \sum_{S \subseteq N} \zeta_S^i \cdot v(S, \{S, N \setminus S\}). \quad (4)$$
- *the Bolger value* (Bolger 1989)—an agent chooses each option with the same probability:
$$\varphi_i^B(v) \overset{\text{def}}{=} \sum_{(S,P)} \beta^i(S, P) v(S, P), \quad (5)$$
where $\beta^i(S, P) = \frac{(|S|-1)!}{|N|!} \sum_{\omega \in \Omega(N \setminus S)} pr_\omega(S, P)$ if $i \in S$ and $\beta^i(S, P) = -\frac{|S|!}{|N|!} \sum_{\omega \in \Omega(N \setminus S \cup \{i\})} pr_{\omega||i}(S, P)$, otherwise. Here, $\omega||i$ is a concatenation of two permutations: ω and i, and $pr_\omega(S, P) = \prod_{j \in N \setminus S} \frac{1}{|\{T \setminus (S \cup C_j^\omega \cup \{j\}) | T \in P\}| + 1}$.

[2] As these axioms can be used for games both with and without externalities, we state them in the more general form.

[3] This translation of Additivity is consistent with Bolger (1989), Macho-Stadler et al. (2007), McQuillin (2009), and Skibski et al. (2013).

[4] While these weights can be any real number, monotonic values are obtained only for non-negative weights. In this case, it is more intuitive to think of a weight as the probability with which an agent enters a coalition outside. Thus, in what follows, we will refer to weights as probabilities.

- *the Macho-Stadler et al. value* (Macho-Stadler et al. 2007)—in the k-th step, an agent chooses a coalition of size b with probability b/k, and forms a new one with probability $1/k$:

$$\varphi_i^{MSt}(v) \stackrel{\text{def}}{=} \sum_{(S,P)} \zeta_S^i \frac{\Pi_{T\in P\setminus\{S\}}(|T|-1)!}{(|N|-|S|)!} v(S,P). \quad (6)$$

- *the Hu-Yang value* (Hu and Yang 2010)—the probabilities of forming all partitions should be the same:

$$\varphi_i^{HY}(v) \stackrel{\text{def}}{=} \sum_{(S,P)} \zeta_S^i \frac{\theta(S,P)}{|\mathcal{P}(N)|} v(S,P), \quad (7)$$

  where $\theta(S,P)$ is the number of partitions from $\mathcal{P}(N)$ that can be obtained from $P \setminus S$ by inserting agents from S.

- *the Myerson value* (Myerson 1977)—a non-monotonic value, obtained through the process with negative weights (with no interpretation in terms of probability):

$$\varphi_i^{M}(v) \stackrel{\text{def}}{=} \sum_{(S,P)} (-1)^{|P|-1}(|P|-1)! \cdot \mu^i_{(S,P)} v(S,P), \quad (8)$$

  where $\mu^i_{(S,P)} \stackrel{\text{def}}{=} \frac{1}{|N|} - \sum_{T\in P\setminus S, i\notin S} \frac{1}{(|P|-1)(|N|-|T|)}$.

Partition Decision Trees

In this section, we propose Partition Decision Trees (PDT)—our representation for coalitional games with externalities.

PDT rules: We represent the game by a set of *PDT rules*. A single PDT rule consists of a *rooted directed tree*, where non-leaf nodes are labelled with agents' names, leaf nodes are labelled with payoff vectors, and edges are labelled with numbers that correspond to coalitions. Formally, a PDT rule T is a tuple $T = \langle V, E, x, f_V, f_E \rangle$, where (V, E) is a directed tree with root x (i.e., V is a set of nodes, $E \subseteq V \times V$ is a set of directed edges, and for every $y \in V$ there exists exactly one path from x to y); $f_V : V \to N \cup \mathbb{R}^{\mathbb{N}}$ is a label function for nodes, and $f_E : E \to \{1, 2, \ldots, |N|\}$ is a label function for edges, with the assumption that $f_V(v) \in V$ for every non-leaf v and $f_V(v) \in \mathbb{R}^{\mathbb{N}}$ for every leaf v.

Given a rule T, let $\Pi(T)$ be the set of paths from the root to any leaf (we will only consider such paths), and let $last(\pi)$ denote the last, leaf node in any $\pi \in \Pi(T)$. Now, every path $\pi = (v_1, v_2, \ldots, v_k) \in \Pi(T)$ represents a partition of agents, where $f_E(v_i, v_{i+1})$ is the number of the coalition to which agent $f_V(v_i)$ belongs. Thus, for every such path:

- all non-leaf nodes on the path are labelled with different agents, i.e., $|\{f_V(v_1), f_V(v_2), \ldots, f_V(v_{k-1})\}| = k-1$;
- for every path from x to v_i, the set of labels on the edges is the set of consecutive natural numbers beginning with 1. Thus, $f_E(v_1, v_2) = 1$, and a label of an edge is not bigger than the maximal label used earlier on this path plus one: $f_E(v_i, v_{i+1}) \le max_{1\le j<i} f_E(v_j, v_{j+1}) + 1$ for $1 \le i < k$;
- the label of a leaf node has exactly the same size as the number of coalitions $|f_V(v_k)| = max_{1\le j<k} f_E(v_j, v_{j+1})$.

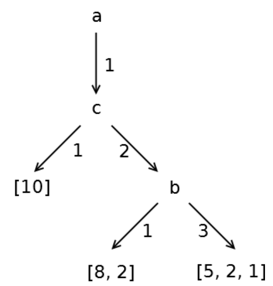

*"If agents **a** and **c** are together, then add* 10 *to the value of this coalition; if not, then:*

- *if **b** is in s coalition with **a**, then add* 8 *to the value of this coalition, and add* 2 *to the value of the coalition with **c**;*
- *if **b** is in a coalition with **c**, then do not add anything;*
- *if **b** is not in a coalition with **a** and **c**, then add* 5 *to the value of the coalition with **a**,* 2 *to the value of the coalition with **c**, and* 1 *to the value of the coalition with **b**."*

Figure 1: An example of a PDT rule.

Since we want to ensure that all partitions are different, we label all outgoing edges of a node with different numbers: if $f_E(v_i, v_j) = f_E(v_i, v_l)$ for some v_i, v_j, v_l, then $v_j = v_l$.

Satisfiability of the PDT rule: Partition $P \in \mathcal{P}$ satisfies $\pi = (v_1, v_2, \ldots, v_k) \in \Pi(T)$, denoted by $P \sim \pi$, if it covers the partition described by this path; thus, for any two members of the same coalition that appear on the path, the labels of outgoing edges are the same (i.e., for every $a, b \in S \in P$, if $f_V(v_i) = a$ and $f_V(v_j) = b$ for some $1 \le i, j < k$, then $f_E(v_i, v_{i+1}) = f_E(v_j, v_{j+1})$). Since paths describe different partitions, there exists no more than one path in one PDT rule satisfied by a partition P. If $P \sim \pi$, then there exists a mapping from the coalitions in partition P to the set of labels of edges (coalition's numbers) with zero: $g_\pi^P : P \to \{1, 2, \ldots, max_{1\le j<k} f_E(v_j, v_{j+1})\} \cup \{0\}$, where 0 is assigned to coalitions which agents does not appear on the path. For example, path π, i.e., a $\to^1 c \to^2 b \to^1 [8, 2]$, from Figure 1 is satisfied by $P = \{\{a, b, d\}, \{c\}, \{e\}, \{f\}\}$ with the mapping $g_\pi^P(\{a, b, d\}) = 1$, $g_\pi^P(\{c\}) = 2$, and $g_\pi^P(\{e\}) = g_\pi^P(\{f\}) = 0$.

According to the mapping g_π^P for $P \sim \pi$, for every coalition S embedded in partition P such that agents in S appear on the path π, there exists a unique value in a payoff node: $f_V(last(\pi))[g_\pi^P(S)]$. All other coalitions have zero value:

$$\omega_\pi^P(S) = \begin{cases} f_V(last(\pi))[g_\pi^P(S)] & \text{if } g_\pi^P(S) > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

for every $P \sim \pi$ and coalition $S \in P$.

Now, the value of coalition S embedded in P in game v^T described by the PDT rule T is the value from the path satisfied by partition P (if it exists):

$$v^T(S,P) \stackrel{\text{def}}{=} \sum_{\substack{\pi\in\Pi(T),\\ P\sim\pi}} \omega_\pi^P(S). \quad (10)$$

The sum symbol can be misleading here, as in every PDT rule there exists at most one path satisfied by a particular partition P. If $\mathcal{T} = \{T_1, T_2, \ldots\}$ is the set of PDT rules, then the game $v^{\mathcal{T}}$ described by the set of rules $\mathcal{T}$ is the sum of games described by each of the following rules:

$$v^{\mathcal{T}}(S,P) \stackrel{\text{def}}{=} \sum_{T \in \mathcal{T}} v^T(S,P) = \sum_{T \in \mathcal{T}} \sum_{\substack{\pi \in \Pi(T), \\ P \sim \pi}} \omega_\pi^P(S). \quad (11)$$

A sample rule is presented in Figure 1. For complexity results, we assume that the size of a PDT rule is defined as the sum of the sizes of labels, and thus equals $O($*number of nodes* + *number of values in the leaf nodes*$)$.

Properties of the PDT representation: We will first show that PDTs are fully expressive.

Lemma 1. *Partition Decision Trees are fully expressive, i.e., every game v can be represented using this representation.*

Proof. The class of elementary games $e^{(S,P)}$ (in which only one coalition in one partition has a non-zero value) forms a basis for games with externalities (Bolger 1989). Thus, it is enough to show that the game $e^{(S,P)}$ can be represented using PDT for every $(S,P) \in EC(N)$. To this end, let π be any order of agents N (i.e., $\pi : \{1, 2, \ldots, |N|\} \to N$) with $\pi(1) \in S$. Consider a PDT rule $\langle V, E, x, f_V, f_E\rangle$, such that $\langle V, E\rangle$ is a path $(v_1, v_2, \ldots, v_n, v_{n+1})$ of length $|N| + 1$, x is the first node on this path, $f_V(v_k) = \pi(k)$ for $k \le n$, $f_V(v_{n+1}) = [1, 0, \ldots, 0]$, and $f_E(v_k, v_{k+1}) =$ *coalition of agent* $\pi(k)$. Now, only partition P satisfies this path, and the only non-zero valued coalition is S. □

In the proof above, we used only paths, i.e., every non-leaf node had only one outgoing edge. However, using only paths may decrease conciseness by a factor of $|N|$. Let us denote by $|\Pi(T)|$ the sum of sizes of all paths in T. The following lemma shows that if a game is represented using the PDT rule T, then the corresponding representation formed by all paths from the tree T can be $|N|$ times bigger than the size of the rule. It is key for our computational results that it cannot be bigger than that.

Lemma 2. *For any PDT rule T, $|\Pi(T)| \le |N||T|$ holds. Moreover, there exists a T such that $|\Pi(T)| = \Theta(|N||T|)$.*

Proof. We defined the size of a PDT rule as the number of nodes plus the number of values in leaf nodes. Since each path contains one leaf node, the number of values in the leaf nodes does not change if we split the rule into a set of paths. The length of any path is not bigger than $|N|$, thus the number of nodes is less than or equal to $|N||$*number of leafs*$|$, which is less than $|N||T|$.

To see that a PDT rule can sometimes be $|N|$ times more concise than a set of paths, consider the following set of paths (assuming z is the last, n-th agent of N):

$$\begin{array}{rl} a \xrightarrow{1} b \xrightarrow{1} \ldots \xrightarrow{1} z \xrightarrow{2} & [x_{n-1}, y_{n-1}] \\ \ldots & \\ a \xrightarrow{1} b \xrightarrow{1} c \xrightarrow{2} & [x_2, y_2] \\ a \xrightarrow{1} b \xrightarrow{2} & [x_1, y_1]. \end{array}$$

This set of paths can be described using n non-leaf nodes (that form the path $(a, b, \ldots, z)$, each node (except for a and z) with two edges labelled 1 and 2), $n-1$ leaf nodes and $2n-2$ values in the leaf nodes, while the number of non-leaf nodes equals $2 + 3 + \ldots + (n+1) = \Theta(n^2)$. □

The above lemma is also useful for comparisons with different representations. For arbitrary value vectors in the leaf nodes, in order to represent the same game, Weighted MC-Nets (Michalak et al. 2010b) would have a rule for every path of the PDT tree (with size being equal to the size of the path), and Embedded MC-Nets would have a rule for every value in the leaf of every path. Thus, PDT representations can be more concise than these alternate representations.

Since the partition function is always exponential, it is easy to see that PDT can be exponentially more concise.

Lemma 3. *Partition Decision Trees can be exponentially more concise than the partition function.*

Proof. It is enough to consider any polynomial-size PDT rule. For example, a PDT rule $a \to^1 1$ assigns value 1 to every coalition with agent a. The partition function which describes this game is exponential in $|N|$. □

In the game described by the rule above, all other agents are null-players. The fact that the size of PDT rules does not depend on the null-players will be crucial for our computational results in the next section (see Theorem 2).

Computing the Extended Shapley Values

In this section, we prove that, given the PDT representation, five out of the six extended Shapley values can be calculated in polynomial time. We start by observing that all the six values satisfy Additivity—one of Shapley's standard axioms. Thus, a value for a game described by a set of rules equals the sum of values for games described by a single rule. This comes from the additive definition of PDT rules (see formula (11)). In the following lemma, we argue that paths are also additive; hence, for additive values we can focus on calculating a value of a game defined by one path.

Lemma 4. *Let $\mathcal{T}$ be the set of PDT rules and $\Pi(\mathcal{T}) = \sum_{T \in \mathcal{T}} \Pi(T)$ the set of all paths. Assume that φ satisfies Additivity. Then, $\varphi(v^{\mathcal{T}}) = \sum_{\pi \in \Pi(\mathcal{T})} \varphi(v^\pi)$, i.e., a value for the game described by the set of PDT rules equals the sum of values of games described by all paths in the rules.*

Proof. Based on the additivity of φ and additivity of PDT rules, we have: $\varphi(v^{\mathcal{T}}) = \sum_{T \in \mathcal{T}} \varphi(v^T)$. Thus, it is enough to argue that $v^T = \sum_{\pi \in \Pi(T)} v^\pi$. Consider an embedded coalition (S,P). Based on the definition of PDT rules, in rule T there exists exactly one path, say $\tilde{\pi}$, satisfied by P. Thus, $v^T(S,P) = \omega_{\tilde{\pi}}^P(S)$, as all other paths of T will not be satisfied by P and contribute nothing to $v(S,P)$ (see formula (10)). In the same manner, $v^\pi(S,P) = 0$ for $\pi \in T$ such that $\pi \neq \tilde{\pi}$. Thus, we can state that:

$$v^T(S,P) = \sum_{\pi \in T} v^\pi(S,P),$$

which concludes the proof. □

To provide our main computational result, we need to strengthen the standard Null-Player Axiom. Based on the standard definition, an agent who does not have any impact on the value of any coalition (i.e., the null-player) should get nothing. This, however, does not imply that the agent has *no impact* on the payoffs of other agents. The stronger requirement of the payoff scheme is called the *Strong Null-Player Axiom*.

- *Strong Null-Player Axiom* (deleting a null-player from the game does not affect agents' payoffs): if i is a null-player in v, then $\varphi(v) = \varphi(v_{-i})$, where v_{-i} is a game defined by removing null-player i from game v, for every game v.[5]

Let us introduce the class of games that plays an important role in our computational analysis of PDT:

$$\Psi(N) \stackrel{\text{def}}{=} \{\sum_{S\in P} c_S e^{(S,P)} \mid P \in \mathcal{P}(N), \langle c_S\rangle_{S\in P} \in \mathbb{R}^{|P|}\}.$$

Thus, $\Psi(N)$ is the set of all games in which only coalitions in a single partition have non-zero values. Note that calculating the payoff of only one agent in an elementary game $e^{(S,P)}$ already requires $O(|N|)$ steps for five extended Shapley values. Interestingly, we show that for games from the class $\Psi(N)$, which are linear combinations of elementary games, the payoffs of *all* agents for four extended Shapley values can also be calculated in time $O(|N|)$.

Theorem 1. *The externality-free, the McQuillin, the Macho-Stadler et al., and the Myerson values can all be calculated in $O(|N|)$ time for $v \in \Psi(N)$.*

Proof. We will show that all four values can be calculated in linear time for every game $v \in \Psi(N)$. Let $v = \sum_{S\in P} c_S e^{(S,P)}$ for some $P\in\mathcal{P}(N)$ and $C = \sum_{S\in P} c_S$.

- the externality-free value: based on formula (3), we can make the following remarks. If P consists of more than one non-singleton coalition, then all payoffs equal zero. If P has one non-singleton coalition S, then $\varphi^{free}(v) = c_S\varphi^{free}(e^{(S,P)})$ and $\varphi_i^{free}(v) = c_S\zeta_S^i$. Now, ζ_S^i is equal for every agent in S, and it is equal for every agent not in S, and both values can be easily calculated in linear time. Finally, if P is the set of singleton coalitions, i.e., $P = \{\{i\} \mid i \in N\}$, then $\varphi_i^{free}(v) = c_{\{i\}}\varphi_i^{free}(e^{(\{i\},P)}) + \sum_{j\in N, j\neq i} c_{\{j\}}\varphi_i^{free}(e^{(\{j\},P)})$. This simplifies to: $\varphi_i^{free}(v) = (c_{\{i\}} - \frac{C-c_{\{i\}}}{|N|-1})/|N| = \frac{|N|c_{\{i\}}-C}{(|N|-1)|N|}$.
- the McQuillin value: based on formula (4), if $P = \{N\}$ then all payoffs equal $\frac{c_N}{|N|}$. If $|P| > 2$, then all payoffs equal zero. Otherwise, $P = \{S, N\setminus S\}$ and $\varphi^{McQ}(v) = c_S\varphi^{McQ}(e^{(S,P)}) + c_{N\setminus S}\varphi^{McQ}(e^{(N\setminus S,P)})$. For every agent i, this equals $c_S\zeta_S^i + c_{N\setminus S}\zeta_{N\setminus S}^i$ which takes only two different values (for agents from S and not from S), and both can be calculated in linear time.
- the Macho-Stadler et al. value: based on formula (6), for $i \in S$, $\varphi_i^{MSt}(e^{(S,P)}) = \frac{\prod_{T\in P}(|T|-1)!}{|N|!}$. Thus, the payoff $\alpha_P = \frac{\prod_{T\in P}(|T|-1)!}{|N|!}$ does not depend on i and S and can be calculated in linear time. Also, if $i \notin S$, then $\varphi_i^{MSt}(e^{(S,P)}) = -\frac{|S|}{|N|-|S|}\alpha_P$. From both equations and for $v = \sum_{S\in P} c_S e^{(S,P)}$, it follows that $\varphi_i^{MSt}(v) = \alpha_P(c_S - \sum_{T\in P, T\neq S}\frac{|T|}{|N|-|T|}c_T)$. Now, if we calculate $\sum_{T\in P}\frac{|T|}{|N|-|T|}c_T$ (denoted X) and α_P first, then $\varphi_i^{MSt}(v) = \alpha_P(c_S - X + \frac{|S|}{|N|-|S|}c_S) = \alpha_P(c_S\frac{|N|}{|N|-|S|} - X)$, which can be calculated in constant time for every agent $i \in N$.
- the Myerson value: based on formula (8), i's payoff in game $e^{(S,P)}$ equals $(-1)^{|P|-1}(|P|-1)!(\frac{1}{|N|} - \sum_{T\in P\setminus S, i\notin S}\frac{1}{(|P|-1)(|N|-|T|)})$. To simplify, denote $\beta_P = (-1)^{|P|-1}(|P|-1)!$, $\gamma_T = \frac{1}{(|P|-1)(|N|-|T|)}$ for every $T \in P$ and $\Gamma = \sum_{T\in P}\gamma_T$. Thus, for $i \in S$, $\varphi_i^M(e^{(S,P)}) = \beta_P(\frac{1}{|N|} - \Gamma + \gamma_S)$ and $\varphi_i^M(e^{(T,P)}) = \beta_P(\frac{1}{|N|} - \Gamma + \gamma_S + \gamma_T)$ for $T \neq S$. Assume that $v = \sum_{S\in P} c_S e^{(S,P)}$. Let us calculate $X = \sum_{T\in P} c_T(\frac{1}{|N|} - \Gamma + \gamma_T)$ first. Now, for $i \in S$, $\varphi_i^M(v) = \beta_P(X + \sum_{T\in P, T\neq S} c_T\gamma_S) = \beta_P(X + \gamma_S(C - c_S))$, which can be calculated in constant time for every agent.

□

Our main result is as follows:

Theorem 2. *Assume that φ satisfies Additivity and the Strong Null-Player Axiom. If $\varphi(v)$ can be calculated in time $O(|N|)$ for every game $v \in \Psi(N)$, then φ can be calculated in time $O(|N|\times|\mathcal{T}|)$ for every set of PDT rules $\mathcal{T}$. Furthermore, the externality-free, the McQuillin, the Macho-Stadler et al., and the Myerson values can all be calculated in time $O(|N|\times|\mathcal{T}|)$ for every set of PDT rules $\mathcal{T}$.*

Proof. We will show that a game represented by a PDT rule consisting of one path π is in $\Psi(M)$ for $|M| = O(|\pi|)$. This fact, combined with Lemma 2 and Lemma 4, shows that φ can be calculated in time $O(|N| \times |\mathcal{T}|)$ for every set of PDT rules $\mathcal{T}$. Moreover, this result combined with Theorem 1 yields the linear complexity for all four values.

Let N be the set of all agents and M be the set of all agents whose labels appear in the path π. Let $P = \{S_1, S_2, \ldots, S_k\}$ be the partition of M that corresponds to π (thus, S_1 gathers all nodes that have outgoing edges labelled with number 1, S_2 those labelled with number 2, etc.). We can argue that all agents from $N \setminus M$ are null-players, as their position in the partition does not affect the payoff of any coalition. This is due to the definition of path satisfiability, which does not concern agents who do not appear in the path. Thus, based on the Strong Null-Player Axiom, payoffs in the game of N agents equal the payoffs in the game of M non-null agents. Only the partition P of M satisfies the path. Thus, only coalitions embedded in this partition have non-zero value and the game $v^\pi \in \Psi(M)$. □

[5]Note that if i is a null-player in v, then creating a game without i is straightforward, as values of all embedded coalitions obtained by inserting agent i into $(S,P)\in EC(N\setminus\{i\})$ are equal. For example, game v_{-i} can be defined as follows: $v_{-i}(S,P) \stackrel{\text{def}}{=} v(S, P\cup\{i\})$ for every $(S,P)\in EC(N\setminus\{i\})$.

The Hu-Yang value does not satisfy the Strong Null-Player Axiom and we cannot limit our analysis to the agents on the path. Nevertheless, it can be calculated in polynomial time. To show this, we consider how adding the null-player changes the values of the game and provide a formula for the Hu-Yang value for a game with several null-players. Consequently, we calculate this value in time $O(|N|^3 \times |\mathcal{T}|)$.

Theorem 3. *The Hu-Yang value can be calculated in time $O(|N|^3 \times |\mathcal{T}|)$ for every set of PDT rules $\mathcal{T}$.*

Proof. We will show that $\varphi^{HY}(e^{(S,P)})$ can be calculated in time $O(|N|)$. Thus, for every game v in the form $v = \sum_{S \in P} c_S e^{(S,P)}$, $\varphi^{HY}(v)$ can be calculated in time $O(|N|^2)$. Using the analogous analysis as in Theorem 2, the Hu-Yang value can be calculated in time $O(|N|^3 \times |\mathcal{T}|)$.

Let us recall that $\theta(S, P)$ is the number of partitions from $\mathcal{P}(N)$ that can be obtained from $P \setminus S$ by inserting agents from S. Thus, we can argue that, for every $(S, P), i \in S$:

$$\theta(S, P) = \sum_{T \in P \setminus S \cup \{\emptyset\}} \theta(S \setminus \{i\}, \tau_i^T(P)). \qquad (12)$$

Let us consider game $v_j^{\emptyset}$ obtained by adding a null-player j to the game v. Value $v(S, P)$ appears several times in the formula—as $v_j^{\emptyset}(S \cup \{j\}, \tau_{j+}^S(P))$, and also $\forall T \in P \setminus S \cup \{\emptyset\}$ as $v_j^{\emptyset}(S, \tau_{j+}^T(P))$, where $\tau_{j+}^T(P) \stackrel{\text{def}}{=} P \setminus \{T\} \cup \{T \cup \{j\}\}$.

$$\varphi_i^{HY}(v_j^{\emptyset}) = \sum_{(S,P)} v(S, P) \frac{|S|!(|N| - |S|)!}{(|N| + 1)!} \frac{\theta(S \cup \{j\}, \tau_{j+}^S(P))}{|\mathcal{P}(N \cup \{j\})|}$$
$$+ \sum_{T \in P \setminus S \cup \{\emptyset\}} v(S, P) \frac{(|S| - 1)!(|N| - |S| + 1)!}{(|N| + 1)!} \frac{\theta(S, \tau_{j+}^T(P))}{|\mathcal{P}(N \cup \{j\})|}.$$

Using formula (12) we get:

$$\varphi_i^{HY}(v_j^{\emptyset}) = \sum_{(S,P)} v(S, P) \frac{(|S| - 1)!(|N| - |S|)!}{|N|!} \frac{\theta(S \cup \{j\}, \tau_{j+}^S(P))}{|\mathcal{P}(N \cup \{j\})|}$$
$$= \sum_{(S,P)} \zeta_S^i \cdot \frac{\theta(S \cup \{j\}, \tau_{j+}^S(P))}{|\mathcal{P}(N \cup \{j\})|} v(S, P).$$

This transformation applied several times yields i's value in the game obtained by adding a set K of null-players:

$$\varphi_i^{HY}(v_K^{\emptyset}) = \sum_{(S,P)} \zeta_S^i \frac{\theta(S \cup K, P \setminus S \cup \{S \cup K\})}{|\mathcal{P}(N \cup K)|} v(S, P).$$

Now, it is enough to calculate $\theta(S \cup K, P \setminus S \cup \{S \cup K\})$. In general, if P is a partition of the set N, then we have:

$$\theta(S, P) = \sum_{i=0}^{|S|} \binom{|S|}{i} B_i (|P| - 1)^{|S|-i},$$

where B_i is i-th Bell number. Thus, $\theta(S \cup K, P \setminus S \cup \{S \cup K\})$ can be calculated in time $O(|S| + |K|) = O(|N|)$, which is enough to calculate all payoffs in time $O(|N|)$. □

The Bolger value is the only extension for which we did not provide a polynomial algorithm. Calculating this value is difficult even for an elementary game $e^{(S,P)}$, in which there exists only one coalition in one partition with a non-zero value (Bolger 1989 provided a recursive formula to calculate the coefficients). However, whether a polynomial algorithm exists under any concise representation is an open problem.

Related Work

The literature on concise representations of coalitional games can be divided into two broad categories (Wooldridge and Dunne 2006). The first category gives the characteristic function a specific interpretation in terms of combinatorial structures such as graphs, e.g., Deng and Papdimitriou (1994), Greco et al. (2009), Wooldridge and Dunne (2006), and the representations discussed in Aziz and de Keijzer (2011). Such representations are guaranteed to be succinct, however they can express only certain games. Our paper fits into the second category of representations, where the emphasis is placed on full expressivity, often at the expense of succinctness: MC-Nets (Ieong and Shoham 2006), its read-once extension (Elkind et al. 2009), Synergy Coalition Groups (Conitzer and Sandholm 2006), the Decision-Diagrams-based representations (Aadithya et al. 2011, Sakurai et al. 2011), and the vector-based representation (Tran-Thanh et al. 2013). While all the above models concern games with no externalities, representations for games with externalities include Embedded MC-Nets (Michalak et al. 2010a; Ichimura et al. 2011) and Weighted MC-Nets (Michalak et al. 2010b). To compare with Partition Decision Trees, if we consider conciseness, then these representations for games with externalities can be lined up in the following way:

partition-function form $\ll$ ***Partition Decision Trees*** $\ll$ *Embedded MC-Nets* $\ll$ *Weighted MC-Nets,*

where $A \ll B$ denotes that B can be exponentially more concise than A, and B is at most polynomially less concise than A.[6] On the other hand, if we consider time complexity of computing various direct extensions of the Shapley value:

partition-function form $=$ ***Partition Decision Trees*** $\gg$ *Embedded MC-Nets* $\gg$ *Weighted MC-Nets,*

where $A \gg B$ denotes that the set of extensions computable in polynomial time based on A is a superset of the corresponding set for B, and $=$ denotes that both sets are equal. Thus, we believe that the Partition Decision Trees representation is closer to the golden mean between the conciseness of the representation and its computational properties.

The underlying formalism behind our representation is that of decision trees. We note that the extensive form of a game in game theory is also a decision tree—the difference is in that we can be often more concise since we do not

[6]Although the PDT representation cannot be exponentially more concise than both MC-Nets representations, we can argue that there exist subclasses of Embedded and Weighted MC-Nets that we can concisely represent: for Weighted MC-Nets, these are the rules that in every block contain only one logical rule made of positive literals; for Embedded MC-Nets, these are the rules where only one logical rule is after the “bar”, and no negative literals appear.

model all moves. Thus, our representation can be directly used to model sequential games of coalition formation.

Finally, we mention the work by Michalak et al. (2009), with an overview of alternative ways to understand externalities. A comprehensive discussion on representation formalisms for various classes of games can be found in (Chalkiadakis, Elkind, and Wooldridge 2011).

Conclusions

In this paper, we presented and analyzed Partition Decision Trees—a new representation for coalitional games with externalities. Two directions for future research appear especially interesting. Firstly, one can think of allowing nodes to merge, which can further improve the conciseness of the representation, but may be more difficult to handle algorithmically. Secondly, it would be interesting to extend PDT in the similar direction in which MC-nets were extended by read-once MC-Nets (Elkind et al. 2009).

Acknowledgements

Oskar Skibski and Makoto Yokoo were supported by JSPS KAKENHI Grant (24220003) and Yuko Sakurai was supported by JST PRESTO program. Tomasz Michalak and Michael Wooldridge were supported by the European Research Council under Advanced Grant 291528 ("RACE"). Tomasz Michalak was supported the Polish National Science grant DEC-2013/09/D/ST6/03920.

References

Aadithya, K. V.; Michalak, T. P.; and Jennings, N. R. 2011. Representation of coalitional games with algebraic decision diagrams. In *AAMAS*, 1121–1122.

Aziz, H., and de Keijzer, B. 2011. Complexity of coalition structure generation. In *AAMAS*, 191–198.

Bachrach, Y., and Rosenschein, J. S. 2008. Coalitional skill games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, 1023–1030.

Bolger, E. 1989. A set of axioms for a value for partition function games. *International Journal of Game Theory* 18(1):37–44.

Chalkiadakis, G.; Elkind, E.; and Wooldridge, M. 2011. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5(6):1–168.

Conitzer, V., and Sandholm, T. 2006. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence* 170(6):607–619.

de Clippel, G., and Serrano, R. 2008. Marginal contributions and externalities in the value. *Econometrica* 76(6):1413–1436.

Deng, X., and Papadimitriou, C. H. 1994. On the complexity of cooperative solution concepts. *Mathematics of Operations Research* 19(2):257–266.

Do, K. P., and Norde, H. 2007. The Shapley value for partition function form games. *International Game Theory Review* 9(02):353–360.

Elkind, E.; Goldberg, L. A.; Goldberg, P. W.; and Wooldridge, M. 2009. A tractable and expressive class of marginal contribution nets and its applications. *Mathematical Logic Quarterly* 55(4):362 – 376.

Greco, G.; Malizia, E.; Palopoli, L.; and Scarcello, F. 2009. On the complexity of compact coalitional games. In *IJCAI*, 147–152.

Hu, C.-C., and Yang, Y.-Y. 2010. An axiomatic characterization of a value for games in partition function form. *SERIEs* 1(4):475–487.

Ichimura, R.; Hasegawa, T.; Ueda, S.; Iwasaki, A.; and Yokoo, M. 2011. Extension of mc-net-based coalition structure generation: handling negative rules and externalities. In *AAMAS*, 1173–1174.

Ieong, S., and Shoham, Y. 2006. Marginal contribution nets: A compact representation scheme for coalitional games. In *ACM-EC*, 170–179.

Macho-Stadler, I.; Perez-Castrillo, D.; and Wettstein, D. 2007. Sharing the surplus: An extension of the Shapley value for environments with externalities. *Journal of Economic Theory* 135(1):339–356.

McQuillin, B. 2009. The extended and generalized Shapley value: Simultaneous consideration of coalitional externalities and coalitional structure. *Journal of Economic Theory* 144(2):696–721.

Michalak, T. P.; Rahwan, T.; Sroka, J.; Dowell, A.; Wooldridge, M. J.; McBurney, P. J.; and Jennings, N. R. 2009. On representing coalitional games with externalities. In Chuang, J.; Fortnow, L.; and Pu, P., eds., *ACM-EC*, 11–20.

Michalak, T.; Marciniak, D.; Szamotulski, M.; Rahwan, T.; Wooldridge, M.; McBurney, P.; and Jennings, N. 2010a. A logic-based representation for coalitional games with externalities. In *AAMAS*, 125–132.

Michalak, T.; Rahwan, T.; Marciniak, D.; Szamotulski, M.; and Jennings, N. 2010b. Computational aspects of extending the Shapley value to coalitional games with externalities. In *ECAI*, 197–202.

Myerson, R. 1977. Values of games in partition function form. *International Journal of Game Theory* 6:23–31.

Rahwan, T.; Michalak, T.; Wooldridge, M.; and Jennings, N. 2012. Anytime coaliton structure generation in multi-agent systems with positive or negative externalities. *Artificial Intelligence* 186:95–122.

Sakurai, Y.; Ueda, S.; Iwasaki, A.; Minato, S.-I.; and Yokoo, M. 2011. A compact representation scheme of coalitional games based on multi-terminal zero-suppressed binary decision diagrams. In *Agents in Principle, Agents in Practice*. Springer. 4–18.

Sandholm, T.; Larson, K.; Andersson, M.; Shehory, O.; and Tohmé, F. 1999. Coalition structure generation with worst case guarantees. *Artificial Intelligence* 111(1–2):209–238.

Shapley, L. 1953. A value for n-person games. In Kuhn, H., and Tucker, A., eds., *Contributions to the Theory of Games, volume II*. Princeton University Press. 307–317.

Skibski, O.; Michalak, T.; and Wooldridge, M. 2013. Marginality approach to Shapley value in games with externalities. technical report.

Skibski, O. 2011. Steady marginality: A uniform approach to Shapley value for games with externalities. In *SAGT*, volume 6982 of *LNCS*, 130–142. Springer.

Thrall, R., and Lucas, W. 1963. N-person games in partition function form. *Nav. Res. Logist. Quart.* 10:281–298.

Tran-Thanh, L.; Nguyen, T.-D.; Rahwan, T.; Rogers, A.; and Jennings, N. R. 2013. An efficient vector-based representation for coalitional games. In *IJCAI*, 383–389.

Wooldridge, M., and Dunne, P. E. 2006. On the computational complexity of coalitional resource games. *Artificial Intelligence* 170(10):835–871.

Yokoo, M.; Conitzer, V.; Sandholm, T.; Ohta, N.; and Iwasaki, A. 2005. Coalitional games in open anonymous environments. In *AAAI*, volume 5, 509–514.