

# An Architectural View of Game Theoretic Control

Ragavendran  
Gopalakrishnan  
Dept. of Computer Science  
California Institute of Tech.

Jason R. Marden  
Dept. of Electrical, Computer  
and Energy Engineering  
Univ. of Colorado at Boulder

Adam Wierman  
Dept. of Computer Science  
California Institute of Tech.

## ABSTRACT

Game-theoretic control is a promising new approach for distributed resource allocation. In this paper, we describe how game-theoretic control can be viewed as having an intrinsic layered architecture, which provides a modularization that simplifies the control design. We illustrate this architectural view by presenting details about one particular instantiation using potential games as an interface. This example serves to highlight the strengths and limitations of the proposed architecture while also illustrating the relationship between game-theoretic control and other existing approaches to distributed resource allocation.

## 1. INTRODUCTION

Resource allocation is a fundamental problem that arises in nearly all computer systems, and increasingly it is a problem that needs to be solved in a distributed, decentralized manner, e.g., power control and frequency selection problems in wireless networks and coverage problems in sensor networks. Resultantly, there is a large and growing literature that focuses on developing distributed resource allocation protocols. This is an extremely diverse literature where protocols are designed using a wide variety of tools, e.g., distributed optimization [28, 44], distributed control [21, 32], physics-inspired control (e.g. Gibbs-sampler-based control) [17, 27], and game-theoretic control [2, 4, 42].

In this paper, we focus on game-theoretic control, which is a promising new approach for distributed resource allocation. The game-theoretic approach involves modeling the interactions of agents within a noncooperative game where the agents are ‘self-interested’. This is motivated by the fact that the underlying decision making architecture in economic systems is identical to the desired decision making architecture in distributed engineering systems, i.e., local decisions based on local information where the global behavior emerges from a compilation of these local decisions. This parallel makes it possible to utilize the broad set of economic/game-theoretic tools in distributed control. However, a key distinction between game theory for economic systems and game theory for engineering systems is that decision makers are inherited in economic systems while decisions makers are designed in engineering systems. This difference means that using game-theoretic tools for distributed control requires a new perspective on the economic literature.

Applying game-theoretic control requires specifying decision makers, their respective choices, their objective/utility functions, and the learning rules for the agents. In this pa-

per, we focus on two of these components: (i) the design of the agents’ utility functions, i.e., *utility design*, and (ii) the design of the distributed learning rules for the agents, i.e., *learning design*. The goal is to design the utility functions and the learning rules so that the emergent global behavior is desirable. There are wide-ranging advantages to the game-theoretic approach including robustness to failures and environmental disturbances, minimal communication requirements, and improved scalability.

Game-theoretic resource allocation designs are increasingly popular in a variety of wireless and sensor network applications, e.g., channel access control in wireless networks [5, 17], coverage problems in sensor networks [8], and power control in both [3, 7, 14]. A comprehensive survey of applications can be found in [4]. However, nearly all of these designs are highly application-specific, with both the utility and learning designs crafted carefully for the specific setting. There have been only a few papers that focus on general designs and even these papers tend to focus on only one aspect of the design – either the utility design, e.g., [25, 26], or the learning design, e.g., [22, 23, 37].

Our contribution in this paper is to present an ‘architectural’ view of game-theoretic control as a whole. We describe, at a high-level, our proposed architecture for game-theoretic control in Section 2. Then, to make the ideas more concrete, in Section 3 we present details about one particular instantiation of the architecture, based on potential games, which aligns with a number of current approaches. Additionally, we provide some new results highlighting how two other existing approaches for distributed system design (distributed constraint optimization and Gibbs-sampler-based control) can be viewed as instances of designs in this architecture. Finally, in Section 4 we discuss a number of important research directions outside of the potential games framework that are suggested by our architectural view of game-theoretic control.

## 2. A LAYERED ARCHITECTURE

Game-theoretic control has two key design tasks: utility design and learning design. Each of these tasks is complex and subject to application-specific constraints. Further, these two designs must be done in conjunction with one another in order to ensure desirable global behavior. Resultantly, many designers have carefully performed application-specific *co-design* of the utility function and the learning rule, e.g., [4, 33, 34, 36], which is typically a difficult task.

The central question that this paper explores is how to achieve a *modularization/decoupling* of utility design and

learning design. Such a decoupling would allow for the development of a rich set of utility designs and learning designs from which a specific utility design and a specific learning design can be chosen ‘off the shelf’ according to the requirements of the resource allocation problem being considered.

It turns out that game theoretic control can naturally be viewed within an ‘hourglass’ architecture, which provides the desired modularity. An hourglass architecture is a type of layered architecture where within the highest and lowest layers there is a large diversity of available designs, but near the middle, or waist, the design is highly constrained. The most famous example of such an ‘hourglass’ architecture is the IP network stack [20, 30, 45]; however, this architecture is quite common in computer systems and has also been observed in wide-ranging areas such as biology [12].

In the context of game-theoretic control, this architecture allows for a diverse set of possible utility and learning designs; but simplifies the design process by enforcing a particular structure on the games that can result from a utility design and on the games for which a learning design will ensure desirable global behavior. This enforced structure then serves as the interface between utility design and learning design; thus providing the desired modularity.

In Section 3, we discuss a concrete example of this constrained interface, the class of *potential games*. (See Figure 1 for an illustration.) This interface requires utility designs to guarantee that the resulting game is a potential game and requires learning rules to guarantee to provide desirable behavior when applied to potential games. Thus, requiring the structure of potential games enforces additional *constraints* on utility and learning design, but also *deconstrains* at a higher level by allowing modularization. Though this layered architecture was not explicitly used in prior work, the modularization provided by potential games underlies many successful examples of game-theoretic control [7, 16, 19, 25]. The change in perspective provided by this architectural view is not simply superficial; it highlights that the utility and learning designs in these papers can be ‘mixed and matched’ while still obtaining the same performance.

Though we focus on potential games in much of this paper, it is important to remember that they are not the only choice for the interface – we discuss moving beyond potential games in Section 4.

### 3. LAYERING VIA POTENTIAL GAMES

We now illustrate how using potential games as an interface provides modularization of utility and learning design. We focus on potential games because many recent applications of game-theoretic control have relied on potential games, e.g., [7, 16, 19, 25]. A key reason that potential games are a powerful choice for the interface is that they are a highly studied class of games in the economics literature, e.g., [13, 29, 35, 43, 47] and so there is a large literature of results that can be used in the context of game-theoretic control for both utility and learning design.

In this section, we highlight the variety of utility and learning designs that have been adapted from the economics literature and can be used interchangeably ‘off the shelf’, greatly simplifying the design process. However, we also illustrate that layering via potential games has some limitations, which highlights the need to consider other interfaces as well. Finally, we illustrate the relationship between layering with potential games and two other existing distributed

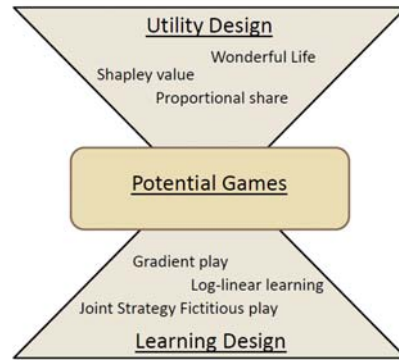


Figure 1: An illustration of the ‘hourglass’ architecture using potential games as the interface.

design approaches. In order to illustrate these issues formally, we first define a simple resource allocation model.

### 3.1 Preliminaries

#### 3.1.1 A model for resource allocation

Consider a set of distributed agents  $N = \{1, \dots, n\}$  and a set of resources  $R = \{r_1, \dots, r_m\}$  that are to be shared by the agents. Each agent  $i \in N$  is capable of selecting potentially multiple resources in  $R$ ; therefore, we say that agent  $i$  has action set  $\mathcal{A}_i \subseteq 2^R$ . An allocation, or an action profile, is represented by a tuple  $a = (a_1, a_2, \dots, a_n) \in \mathcal{A}$  where the set of possible allocations is denoted by  $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ . We will frequently denote an allocation  $a$  as  $(a_i, a_{-i})$  where  $a_{-i} \in \mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$  denotes the allocation of all agents except agent  $i$ .

The social welfare function  $W(a)$  captures the global valuation of the agent allocation. In general, a resource allocation design seeks to find an allocation that optimizes the global welfare. In this work, we assume  $W(a)$  is *linearly separable* across resources, i.e.,  $W(a) = \sum_{r \in R} W_r(\{a\}_r)$  where  $\{a\}_r = \{i \in N : r \in a_i\}$  is the set of agents that are allocated to resource  $r$  in  $a$  and  $W_r : 2^N \rightarrow \mathbb{R}_+$  is the local welfare function for resource  $r$ . Hence, the welfare generated at a particular resource depends only on which agents are allocated to that resource. Further, we restrict our attention to *submodular* welfare functions, i.e., for each resource  $r \in R$  and any player sets  $X \subseteq Y \subseteq N$ ,

$$W_r(X) + W_r(Y) \geq W_r(X \cup Y) + W_r(X \cap Y).$$

A variety of resource allocation problems such as power control and coverage problems in sensor networks [7, 25], wireless access point assignment and frequency selection [17], and influence maximization [18] all have linearly separable, submodular welfare functions.

#### 3.1.2 Resource allocation games

Our goal is to utilize game theory to obtain distributed solutions to such resource allocation problems. This goal requires modeling the interactions of the agents in a noncooperative game theoretic environment where the agents act in a self-interested fashion. While we inherit the players  $N$ , the welfare function  $W$ , and the action sets  $\{\mathcal{A}_i\}_{i \in N}$ , we are left to design a utility function for each player of the form

$U_i : \mathcal{A} \rightarrow \mathbb{R}$ . A resource allocation game  $G$  is then defined by the tuple  $G = (N, \{\mathcal{A}_i\}, \{U_i\}, W)$ .

In general, a system designer has free reign in the design of utility functions; however, layering via potential games requires that the utility functions lead to a potential game. Formally, a game is called a *potential game*, if there exists a potential function  $\Phi : \mathcal{A} \rightarrow \mathbb{R}$  such that  $\forall i, \forall a_{-i} \in \mathcal{A}_{-i}$ , and  $\forall a_i, a'_i \in \mathcal{A}_i$ :

$$\Phi(a_i, a_{-i}) - \Phi(a'_i, a_{-i}) = U_i(a_i, a_{-i}) - U_i(a'_i, a_{-i}).$$

Potential games possess several nice properties that can be utilized in distributed control. One such property is the guaranteed existence of a *pure Nash equilibrium*. A (*pure Nash*) *equilibrium* is an action profile  $a^* \in \mathcal{A}$  such that for each player  $i$ ,  $U_i(a_i^*, a_{-i}^*) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^*)$ . In a distributed system, a pure Nash equilibrium takes on the role of a stable operating point.

To discuss the efficiency of games we use the notions of Price of Anarchy (PoA) and Price of Stability (PoS) [31], which compare the welfare of the set of equilibria to the globally optimal welfare. Let  $\mathcal{G}$  denote a set of games and  $S(G)$  denote the set of equilibria for a game  $G$ . Then

$$\begin{aligned} \text{PoA}(\mathcal{G}) &= \inf_{G \in \mathcal{G}} \left( \min_{a^* \in S(G)} \frac{W(a^*)}{\max_{a \in \mathcal{A}} W(a)} \right) \\ \text{PoS}(\mathcal{G}) &= \inf_{G \in \mathcal{G}} \left( \max_{a^* \in S(G)} \frac{W(a^*)}{\max_{a \in \mathcal{A}} W(a)} \right) \end{aligned}$$

So, the PoA measures the worst-case efficiency of any equilibrium while the PoS measures the worst-case efficiency of the best equilibrium across all games.

## 3.2 Utility Design

There are several desirable properties that a global planner must consider when designing utility functions. These properties include the following conditions on utility functions: (i) computable using only local information, (ii) guarantee the existence of an equilibrium, (iii) computable in polynomial time, (iv) guarantee PoS=1, (v) guarantee a PoA close to 1, and (vi) guarantee that utilities are budget-balanced, i.e., the sum of the utilities of all the players is equal to the social welfare for every action profile. In the context of the considered resource allocation games, we want to design utility functions that are *linearly separable*, i.e., that satisfy

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} f_r(i, \{a\}_r),$$

where  $f_r : N \times 2^N \rightarrow \mathbb{R}$ , i.e., a player's utility is only dependent on the resources the players choose and the other players that choose the same resources.

Recent work [25] has observed that the task of utility design is strongly related to the cost sharing literature in economics. Here, we present two promising utility designs that have thus emerged, though it is important to point out that there are a variety of other possible choices, e.g., [10, 25].

(i) **Wonderful life utility design (WLU)** [46]: This design defines the utility of each player  $i$  as their marginal contribution to the social welfare,

$$\begin{aligned} U_i(a_i, a_{-i}) &= W(a_i, a_{-i}) - W(\emptyset, a_{-i}) \\ &= \sum_{r \in a_i} W_r(\{a\}_r) - W_r(\{a\}_r \setminus i). \end{aligned}$$

Note that this utility design is linearly separable. Additionally, it has been shown in [25] that WLU leads to a potential game with  $\Phi = W$  [46] and that  $\text{PoS} = 1$  and  $\text{PoA} = 1/2$  for submodular resource allocation games. However, WLU is not budget-balanced.

(ii) **Shapley value utility design (SVU)** [40]: This design defines the utility of each player  $i$  as their Shapley value to the social welfare,

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} Sh_i(\{a\}_r; W_r)$$

where  $Sh_i(\{a\}_r; W_r)$ , the Shapley value of player  $i$  at resource  $r$ , is

$$\sum_{S \subseteq \{a\}_r: i \in S} \frac{(|\{a\}_r| - 2)! (|S| - 1)!}{|\{a\}_r|!} (W_r(S) - W_r(S \setminus \{i\})).$$

Notice that this utility design is also linearly separable. Furthermore, SVU leads to a potential game [43], is budget-balanced, and guarantees that  $\text{PoA} = \text{PoS} = 1/2$  for submodular resource allocation games [25]. However, in general, SVU is not polynomial-time computable [11]. Note that there is also a 'weighted' SVU [39] that has the same properties.

## 3.3 Learning design

As with utility design, there are several desirable properties that a global planner must consider when designing distributed learning rules. These properties include (i) asymptotic global behavior, (ii) equilibrium selection, (iii) informational dependencies and (iv) convergence rates.

A learning design takes on the form of a one-shot repeated game where at each time  $t \in \{0, 1, 2, \dots\}$ , each player  $i \in N$  simultaneously chooses an action  $a_i(t) \in \mathcal{A}_i$  according to probability distribution  $p_i(t)$  and receives a utility  $U_i(a(t))$  where  $a(t) = (a_1(t), \dots, a_n(t))$ . We refer to  $p_i(t)$  as the *strategy* of player  $i$  at time  $t$  and denote the probability that player  $i$  will play action  $a_i$  at time  $t$  by  $p_i^{a_i}(t)$ . A player's strategy at time  $t > 0$  can rely only on actions (and their corresponding utilities) from times  $\{0, 1, 2, \dots, t-1\}$ .

Here, we present two promising learning designs that have emerged, though it is important to point out that there are a variety of other possible choices, e.g., [15, 37, 38, 47].

(i) **Joint Strategy Fictitious Play (JSFP)** [23]: In JSFP, each agent  $i$  starts with an initial strategy  $p_i(0)$ , and then at every time step  $t > 0$ , maintains a hypothetical payoff for each action  $a_i$  of the form

$$V_i^{a_i}(t) = \sum_{\tau=0}^{t-1} \frac{1}{t} U_i(a_i, a_{-i}(\tau)).$$

Note that this hypothetical payoff can be computed recursively and that a player only needs access to the payoff for alternative actions at each time step. Using this hypothetical payoff, the strategy of player  $i$  at time  $t > 0$  is of the form

$$p_i^{a_i(t-1)}(t) = \epsilon, \quad p_i^{a_i^*}(t) = 1 - \epsilon$$

where  $a_i^* \in \arg \max_{a_i \in \mathcal{A}_i} V_i^{a_i}(t)$  and  $\epsilon > 0$  is referred to as 'inertia'. For any potential game, if all players adhere to this strategy, then the global behavior will converge almost surely to a pure Nash equilibrium [23].

(ii) *Log-linear learning* [6, 24]: In many settings, it is desirable to converge to a specific Nash equilibrium, rather than just any equilibrium. This is termed ‘equilibrium selection’ and is not provided by JSFP. Log-linear learning provides equilibrium selection, but it also requires more structure on the learning environment. Under log-linear learning, each agent  $i$  starts with an initial strategy  $p_i(0)$ , and then at every time step  $t > 0$ , one agent  $i \in N$  is randomly selected to update its action while all other agents are required to keep their actions fixed, i.e.,  $a_{-i}(t) = a_{-i}(t-1)$ . Player  $i$  plays a strategy at time step  $t > 0$  of the form

$$p_i^{a_i}(t) = \frac{e^{\frac{1}{T}U_i(a_i, a_{-i}(t-1))}}{\sum_{a'_i \in \mathcal{A}_i} e^{\frac{1}{T}U_i(a'_i, a_{-i}(t-1))}}$$

where  $T \geq 0$  is a temperature coefficient. For any potential game, if all players adhere to this mechanism, then the global behavior obeys an ergodic Markov chain with a unique stationary distribution given by:

$$\mu^a = \frac{e^{\frac{1}{T}\Phi(a)}}{\sum_{a' \in \mathcal{A}} e^{\frac{1}{T}\Phi(a')}}.$$

As the process cools (anneals), i.e.,  $T \rightarrow 0^+$ , all the weight of the stationary distribution falls on the action profiles that maximize the potential function, thus providing equilibrium selection [24]. Often, e.g., for WLU where  $\Phi = W$ , this guarantees that the global performance achieves the bound set by the PoS, highlighting the importance of PoS as a measure for efficiency.

### 3.4 Limitations

To this point we have highlighted that using potential games as an interface provides modularization and diverse options for utility and learning design. However, recent research [26] has identified that there are also fundamental limitations of layering via potential games. In particular, it is not possible for a utility design to achieve all of the properties one might desire.

For example, there are conflicts between maintaining a  $PoS = 1$  and being budget-balanced. Specifically, a limitation of layering via potential games is that *any linearly separable, budget-balanced utility design that guarantees the existence of an equilibrium for all resource allocation games has  $PoS \leq 1/2$*  [26]. WLU and SVU provide an illustration of this limitation: SVU is linearly separable and budget-balanced but has  $PoS = 1/2$  while WLU is linearly separable and has  $PoS = 1$  but is not budget-balanced.

Notice that this limitation focuses on budget-balanced utility design. While being local and having a PoS close to one are nearly always desirable, not all applications require that the utility design be budget-balanced. Thus, in many cases, this limitation may not be relevant. However, limitations of potential games have only begun to be studied, and there are likely more severe ones yet to be discovered.

### 3.5 Relationships to other approaches

In this section, we highlight some interesting connections between layering via potential games and other distributed system design tools. We show formally that two other tools (distributed constraint optimization and Gibbs-sampler-based control) can be viewed as instances of layering via potential

games (though they do not fall into our simple illustrative model of resource allocation games). This relationship highlights that in retrospect, one could have used the layered architecture presented in this paper to design these techniques modularly via ‘off the shelf’ utility and learning designs. Further, it highlights that other utility and learning designs can easily be swapped into these designs.

#### 3.5.1 Gibbs-sampler-based control

Gibbs-sampler-based control [17, 27] is a popular physics-inspired approach for wireless protocol design. To introduce this technique, we adapt the description of [17]. We represent a distributed system by an undirected graph  $G(N, E)$  with  $|N| = n$ . Each node stores a state variable from a finite state space  $S$ . The state of the graph is  $s = (s_1, \dots, s_n)$ . An energy function  $\mathcal{E} : S^n \rightarrow \mathbb{R}$  represents the global cost of the system as a function of its state. The objective is to find a state of minimum energy. The Gibbs sampler approach provides an efficient solution for this problem, if  $\mathcal{E}(s)$  is of the form

$$\mathcal{E}(s) = \sum_k \sum_{M \in C_k} V(M)$$

where  $C_k$  is the set of all cliques of order  $k$ , and  $V : 2^N \rightarrow \mathbb{R}_+$  is such that  $V(M)$  depends only on the states of nodes in  $M$ , and is zero if  $M$  is not a clique. We then define the *local energy* of a node  $i \in N$  to be the sum of those terms in  $\mathcal{E}(s)$  that involve  $s_i$ ,

$$\mathcal{E}_i(s_i, (s_j)_{j \neq i}) = \sum_k \sum_{M \in C_k: i \in M} V(M)$$

The Gibbs measure associated with an energy function  $\mathcal{E}$  and temperature  $T > 0$  is defined as the following probability distribution on the states of the graph,

$$\pi(s) = e^{-\frac{\mathcal{E}(s)}{T}} / \left( \sum_{s' \in S^n} e^{-\frac{\mathcal{E}(s')}{T}} \right) \quad (1)$$

The *Gibbs sampler* is an iterative procedure where during each step, each node  $i$ , given the states of all other nodes, samples its new state from the following distribution on  $S$ ,

$$\mu(s_i) = e^{-\frac{\mathcal{E}_i(s_i, (s_j)_{j \neq i})}{T}} / \left( \sum_{s'_i \in S} e^{-\frac{\mathcal{E}_i(s'_i, (s_j)_{j \neq i})}{T}} \right), \quad s_i \in S$$

When  $T$  is fixed, the Gibbs sampler converges to a steady state that is distributed according to (1). Finally, for convergence to the global minimum of the energy function, we use the ‘annealed’ Gibbs sampler, which adds a small decrease of  $T$  to this algorithm at every step. When this decrease with time  $t > 0$  is proportional to  $1/\log(1+t)$ , the system converges to a set of states of minimal global energy.

The above description of Gibbs-sampler-based control already highlights the connection with potential games – it is equivalent to using WLU in combination with log-linear learning. It is immediate to see that the Gibbs sampler is the same as the log-linear learning algorithm described in Section 3.3. To show that the utility design is equivalent to WLU, we construct a game as follows. Consider  $N$  to be the set of players, the common state space to be their action sets, and the negative of the global energy function to be the social welfare function. Denote by  $C_k(H)$ , the set

of cliques of order  $k$  in graph  $H$ . The WLU design is:

$$\begin{aligned} U_i(s) &= W(s) - W(\emptyset, s_{-i}) \\ &= -\mathcal{E}(s) + \mathcal{E}(\emptyset, s_{-i}) \\ &= -\sum_k \sum_{M \in C_k(G)} V(M) + \sum_k \sum_{M \in C_k(G-\{i\})} V(M) \end{aligned}$$

Since  $M \in C_k(G - \{i\}) \iff M \in C_k(G)$  and  $i \notin M$ , all these terms cancel out, leaving only the terms in  $-\mathcal{E}(s)$  for which  $i \in M$ . Hence,

$$U_i(s) = -\sum_k \sum_{M \in C_k, i \in M} V(M) = -\mathcal{E}_i(s)$$

### 3.5.2 Distributed constraint optimization

A constraint optimization problem is specified by a set of variables  $N = \{1, \dots, n\}$ , each of which takes a value  $s_i$  from a finite state space  $S$ , a set of constraints  $C = \{c_1, c_2, \dots, c_m\}$ , and a global objective function  $W : S^n \rightarrow \mathbb{R}$ , that encodes the relative desirability of each possible state  $s \in S^n$  of the system. A constraint  $c = \langle N_c, R_c \rangle$  is specified by the set of variables  $N_c \subseteq N$  over which it is defined, and a relation  $R_c \subset S^{|N_c|}$  between those variables. A function  $U_c(s_c)$  specifies the reward for satisfying constraint  $c$ , where  $s_c$  is the configuration of the states of the variables in  $N_c$ . The global objective function is typically written as  $W(s) = \sum_{c \in C} U_c(s_c)$ . The problem is to find a global maximizer of  $W$ . Given this, a ‘distributed’ constraint optimization problem (DCOP) is produced when a set of autonomous agents each independently control the state of a variable.

In [9], the authors show that DCOPs can be viewed in the context of potential games, thus allowing any of the learning designs described earlier to be applied. Here, we highlight that this ‘DCOP game’ corresponds specifically to choosing WLU as the utility design. Consider the autonomous agents of the DCOP to be the players, the common state space to be their action sets, and the global objective function to be the social welfare function. Denote by  $C(M)$  the set of constraints involving any of the variables in the set  $M \subseteq N$ . The WLU design is:

$$\begin{aligned} U_i(s) &= W(s) - W(\emptyset, s_{-i}) \\ &= \sum_{c \in C(N)} U_c(s_c) - \sum_{c \in C(N \setminus \{i\})} U_c(s_c) \\ &= \sum_{c \in C(\{i\})} U_c(s_c) \end{aligned}$$

which is exactly the utility function suggested for the DCOP game in [9]. The last step follows by observing that when variable  $i$  is not part of the DCOP, all constraints not containing  $i$  in the original DCOP are not affected.

## 4. MOVING BEYOND POTENTIAL GAMES

To this point, we have discussed a concrete example of the game-theoretic control architecture – using potential games as an interface between utility and learning design. However, potential games are only one, very restrictive class of games, and research has begun to uncover limitations of layering via potential games (see Section 3.4).

Thus, as we move forward, it is important to consider other options for the interface. Recent research is beginning to consider a variety of other classes of games as the basis for game-theoretic control. For example, [26] suggests ‘state-based potential games’, which are a limited form of Markov games, as a way to overcome the limitation of potential games described in Section 3.4. Other examples include [41], which proposes using conjectural equilibria in the context of multi-user power control, and [1] which proposes using oblivious equilibria in the context of large stochastic games. However, as yet, there is little understanding of the strengths and limitations of designs using these new classes of games. For example, which classes of games provide modularity when used as an interface? What is gained by broadening the interface from potential games to other classes? Is there a penalty for broadening the interface, e.g. slower convergence rates for learning?

Our hope is that the identification of an architectural view of game-theoretic control in this paper can help formalize and motivate these important directions for the field. We propose that a better understanding of the strengths and weaknesses of differing interfaces will provide useful insight into how to choose the appropriate interface for a given class of applications. For example, for some applications, the limitation described in Section 3.4 may not be important, while for others, the limitation may cause design outside of potential games to be preferable.

## 5. REFERENCES

- [1] S. Adlakha, R. Johari, G. Weintraub, and A. Goldsmith. Oblivious equilibrium: an approximation to large population dynamic games with concave utility. In *Game Theory for Networks*, pages 68–69, 2009.
- [2] T. Alpcan, L. Pavel, and N. Stefanovic. A control theoretic approach to noncooperative game design. In *IEEE CDC*, pages 8575–8580, 2009.
- [3] E. Altman and Z. Altman. S-modular games and power control in wireless networks. *Transactions on Automatic Control*, 48(5):839–842, 2003.
- [4] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter. A survey on networking games in telecommunications. *Computers and Operations Research*, 33(2):286–311, 2006.
- [5] E. Altman, R. El-Azouzi, and T. Jiménez. Slotted ALOHA as a game with partial information. *Computer Networks*, 45(6):701–713, 2004.
- [6] L. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [7] E. Campos-Náñez, E. Garcia, and C. Li. A game-theoretic approach to efficient power management in sensor networks. *Operations Research*, 56(3):552–561, 2008.
- [8] C. G. Cassandras and W. Li. Sensor networks and cooperative control. *European Journal of Control*, 11(4-5):436–463, 2005.
- [9] A. Chapman, A. Rogers, and N. R. Jennings. Benchmarking hybrid algorithms for distributed constraint optimization games. In *ACM OptMAS*, pages 1–11, 2008.

- [10] H.-L. Chen, T. Roughgarden, and G. Valiant. Designing networks with good equilibria. In *ACM-SIAM SODA*, pages 854–863, 2008.
- [11] V. Conitzer and T. Sandholm. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *AAAI*, 2004.
- [12] M. Csete and J. C. Doyle. Bow ties, metabolism and disease. *Trends in Biotechnology*, 22(9):446–450, 2004.
- [13] P. Dubey, O. Haimanko, and A. Zapechelnyuk. Strategic complements and substitutes, and potential games. *Games and Economic Behavior*, 54(1):77–94, 2006.
- [14] D. Falomari, N. Mandayam, D. Goodman, and V. Shah. A new framework for power control in wireless data networks: games, utility, and pricing. In *Wireless Multimedia Network Technologies*, pages 289–310. 1999.
- [15] D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, MA, 1998.
- [16] A. Garcia, D. Reaume, and R. L. Smith. Fictitious play for finding system optimal routings in dynamic traffic networks. *Transportation Research Part B: Methodological*, 34(2):147–156, 2000.
- [17] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-based self organization of interfering 802.11 wireless access networks. In *IEEE INFOCOM*, pages 1451–1459, 2007.
- [18] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *ACM KDD*, pages 137–146, 2003.
- [19] R. S. Komali and A. B. MacKenzie. Distributed topology control in ad-hoc networks: a game theoretic perspective. In *IEEE CCNC*, pages 563–568, 2006.
- [20] J. F. Kurose and K. W. Ross. *Computer networking: a top-down approach*. Addison-Wesley, 2009.
- [21] L. E. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer. A cone-based distributed topology-control algorithm for wireless multi-hop networks. *Transactions on Networking*, 13(1):147–159, 2005.
- [22] J. R. Marden, G. Arslan, and J. S. Shamma. Regret based dynamics: convergence in weakly acyclic games. In *ACM AAMAS*, pages 1–8, 2007.
- [23] J. R. Marden, G. Arslan, and J. S. Shamma. Joint strategy fictitious play with inertia for potential games. *Transactions on Automatic Control*, 54(2):208–220, Feb 2009.
- [24] J. R. Marden and J. S. Shamma. Revisiting log-linear learning: Asynchrony, completeness and a payoff-based implementation,. In *Under submission*.
- [25] J. R. Marden and A. Wierman. Distributed welfare games. *Under submission*.
- [26] J. R. Marden and A. Wierman. Overcoming limitations of game-theoretic distributed control. In *IEEE CDC*, pages 6466–6471, 2009.
- [27] V. Mhatre, K. Papagiannaki, and F. Baccelli. Interference mitigation through power control in high density 802.11 wlans. In *IEEE INFOCOM*, pages 535–543, 2007.
- [28] A. Mishra, V. Shrivastava, D. Agrawal, S. Banerjee, and S. Ganguly. Distributed channel management in uncoordinated wireless environments. In *MOBICOM*, pages 170–181, 2006.
- [29] D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Economic Theory*, 68(1):258–265, 1996.
- [30] Nat.Research Council Comm. on the Internet in the Evolving Info. Infrastructure. The Internet’s coming of age, 2001.
- [31] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.
- [32] P. Parag, S. Shakkottai, and J.-F. Chamberland. Value-aware resource allocation for service guarantees in networks. In *IEEE INFOCOM*, 2010.
- [33] L. Pavel. An extension of duality to a game-theoretic framework. *Automatica*, 43:226237, 2007.
- [34] A. Rantzer. Distributed control using decompositions and games. In *IEEE CDC*, 2008.
- [35] W. H. Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97(1):81–108, 2001.
- [36] G. Scutari, D. P. Palomar, and J. Pang. Flexible design of cognitive radio wireless systems: from game theory to variational inequality theory. *IEEE Signal Processing Magazine*, 26(5):107–123, September 2009.
- [37] D. Shah and J. Shin. Dynamics in congestion games. In *ACM SIGMETRICS/Performance*, 2010.
- [38] J. Shamma and G. Arslan. Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria. *IEEE Trans. on Automatic Control*, 50(3):312–327, 2005.
- [39] L. S. Shapley. *Additive and non-additive set functions*. PhD thesis, Department of Mathematics, Princeton University, 1953.
- [40] L. S. Shapley. A value for n-person games. In *Contributions to the theory of games – II*. Princeton University Press, 1953.
- [41] Y. Su and M. van der Schaar. Conjectural equilibrium in multiuser power control games. *Transactions on Signal Processing*, 57(9):3638–3650, 2009.
- [42] Y. Su and M. van der Schaar. A new perspective on multi-user power control games in interference channels. *Transactions on Wireless Communications*, 8(6):2910–2919, 2009.
- [43] T. Ui. Shapley value representation of potential games. *Games and Economic Behavior*, 31(1):121–135, 2000.
- [44] E. G. Villegas, R. V. Ferré, and J. Josep Paradells. Frequency assignments in IEEE 802.11 WLANs with efficient spectrum sharing. *Wireless Communications and Mobile Computing*, 9(8):1125–1140, 2008.
- [45] W. Willinger and J. C. Doyle. Robustness and the internet: design and evolution. In *Robust design: a repertoire of biological, ecological, and engineering case studies*. Oxford University Press, 2005.
- [46] D. H. Wolpert and K. Tumer. An introduction to collective intelligence. In *Handbook of Agent technology*. AAAI, 1999.
- [47] H. P. Young. *Strategic Learning and its Limits*. Oxford University Press, 2005.