Efficient Computation of Shapley Values for Demand Response Programs

Gearóid O'Brien
Department of Electrical Engineering
Stanford University
Stanford, CA
gobrien@stanford.edu

Abbas El Gamal
Department of Electrical Engineering
Stanford University
Stanford, CA
abbas@ee.stanford.edu

Ram Rajagopal
Department of Civil and Environmental Engineering
Stanford University
Stanford, CA
ramr@stanford.edu

Categories and Subject Descriptors: G.3 [Probability and Statistics]: Probabilistic algorithms

General Terms: Algorithms, Economics, Theory

Keywords: Demand Response; Shapley Value; Game Theory; Cooperative Games

1. INTRODUCTION

Demand Response (DR) is currently a major area of research. We propose analyzing demand response schemes in a game theoretic setting and utilizing the Shapley Value for fairly compensating participants of such schemes. As exact computation of the Shapley Value is intractable in general, we propose a stratified sampling method that can dramatically reduce the variance of the Shapley Value estimate when compared to previous methods.

DR ([1] gives an overview of DR) can be achieved by a combination of three actions. Usage can be curtailed by reducing their consumption level, a process known as "load shedding." Local generation can be used to take loads "off-grid," and loads can be adjusted temporally so that the load consumes power at a later time [5].

We focus on the cases where penalties are submodular functions, and payments are supermodular functions of the set of participants. In this case, the Shapley Value concept provides a payment incentive that guarantees that every participant is better off joining a single grand coalition. Computing these optimal payments requires a combinatorial calculation. We propose a statistical approach to this challenge and show that the approach works well in practice. It generalizes an early idea of sampling coalitions [3], but furthers it by considering a careful statistical computation.

2. PROBLEM FORMULATION

A DR program consists of an operator and customers (agents) $\mathcal{X}$ who have agreed to participate. Each agent provides a load profile of power consumed which is not known in advance. The program manager requests a given level of DR at a certain time, and the agents must react as requested. If the coalition fails to meet the target, an ex-post penalty is charged which should be fairly shared by all participants in the program. As an example, agents receive a fixed payment c_i for participating and then subtracts a penalty. The total paid out to the entire coalition is then $P(\mathcal{X}) = \sum_i c_i - f(\mathcal{X})$, where $f(\mathcal{X})$ is the total penalty assigned to the coalition and is known as the *characteristic function*. In many practical DR programs, the payment c_i has already been completed and so it is only necessary to fairly divide the total penalty $f(\mathcal{X})$ amoung the agents. If the game is structured so that the characteristic function is *submodular* (or conversely, if the payment function is supermodular), then no subset of players has any incentive to leave the coalition.

Shapley proposed a solution that is both unique and fair for dividing $f(\mathcal{X})$ amoung the agents. The Shapley value for agent $x \in \mathcal{X}$ is a weighted average of the marginal contributions of an agent to all possible coalitions and can be defined in the equivalent equations (1) and (2).

$$\phi_x(f) = \sum_{\mathcal{A} \subseteq \mathcal{X} \setminus \{x\}} \frac{|\mathcal{A}|!(|\mathcal{X}| - |\mathcal{A}| - 1)!}{|\mathcal{X}|!} \left[\rho_x(\mathcal{A})\right] \quad (1)$$

$$= \frac{1}{|\mathcal{X}|!} \sum_R \left[\rho_x(\mathcal{A}_x^R)\right] \quad (2)$$

$\rho_x(\mathcal{Y}) = f(\mathcal{Y} \cup \{x\}) - f(\mathcal{Y})$ is the marginal contribution of agent x to coalition $\mathcal{Y}$. In (2), R is an ordering of the agents in $\mathcal{X}$ and $\mathcal{A}_x^R$ is the set of agents that precede x in a given ordering R.

Aside from directly calculating (1), generating functions have been previously examined [2]. This approach requires large arrays in order to reduce complexity to polynomial time, although it does produce an exact solution for the Shapley Value. The intractability of calculating the Shapley value has given rise to a number of studies which attempt to estimate it via approximation methods [3]. In [4], the authors use the Shapley value to allocate transmission service costs among network users in energy markets.

The main contribution of this paper is proposing a new stratified sampling method for estimating the Shapley Value, thus reducing the variance of the estimated value when compared to other randomized sampling methods.

3. ESTIMATING THE SHAPLEY VALUE USING STRATIFIED SAMPLING

Stratified sampling is a method of dividing the population into strata and then sampling from these strata. Here, the population is the marginal contributions of a given participant to every possible coalition. We divide this population into strata where each stratum contains the marginal contribution of a given participant to coalitions with an equal number of agents.

e-Energy'13, May 21–24, 2013, Berkeley, California, USA.
ACM 978-1-4503-2052-8/13/05.

It can be shown that the Shapley Value is equal to the mean of the average marginal contributon from each stratum. By sampling appropriately from the strata, the variance of the Shapley Value estimate can be reduced. We introduce the notation $\boldsymbol{\sigma}^2$ to represent a vector of the strata variances and $\boldsymbol{\mu}$ to represent a vector of the strata means, indexing both vectors by i. Further, N is the total number of samples taken and n is the number of strata.

The variance of the Shapley Value estimate calculated by stratified sampling is minimized by sampling from all strata in proportion to its standard deviation. The variance of this estimate is

$$\sigma_{\text{SD}}^2 = \frac{1}{N}\left[\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{\sigma}_i\right]^2. \tag{3}$$

Comparing this variance to that of random sampling, where we do not distinguish between strata the variance is

$$\sigma_{\text{RS}}^2 = \frac{1}{N}\left[\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{\sigma}_i^2 + \text{var}(\boldsymbol{\mu})\right]. \tag{4}$$

As the strata standard deviations will not be available a priori in general, it is natural to ask what the *value* of this knowledge is. Applying the principle of indifference, we sample equally from each stratum. Here, the variance is

$$\sigma_{\text{EB}}^2 = \frac{1}{N}\left[\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{\sigma}_i^2\right]. \tag{5}$$

It is possible to therefore show:

THEOREM 3.1. *It is always better to sample equally from each stratum than to sample randomly, and if the standard deviations of the strata are known, it is better again to sample in proportion to those, since:*

$$\sigma_{\text{SD}}^2 \leq \sigma_{\text{EB}}^2 \leq \sigma_{\text{RS}}^2.$$

Next, we propose a reinforcement learning algorithm to approximate standard deviation sampling.

3.1 Implementing SD Weighted Sampling

This is a typical reinforcement learning problem where we seek to *exploit* the information we know (regarding the standard deviations) in order to sample correctly, but must also *explore* in order to calculate the standard deviations with accuracy. Defining $\pi_i(t)$ as the probability of sampling from stratum i at sample number t, and $\hat{\sigma}_i$ as the estimated standard deviation of stratum i:

$$\pi_i(t) = \epsilon_t\frac{1}{n} + (1-\epsilon_t)\frac{\hat{\sigma}_i}{\sum_j \hat{\sigma}_j},$$

A suitable function for ϵ_t would be similar to the sigmoid function $\epsilon_t = [1 + \exp(t - N)]^{-1}$, ensuring that the probablility of sampling from each stratum starts equal and ends proportional to the standard deviation. For each sample, the reinforcement learning algorithm chooses a stratum at random, weighted according to $\pi_i(t)$. After sampling, the weights are updated according to ϵ_t and we sample again until N samples have been taken.

4. EXAMPLE

A simple example of a supermodular payment structure is the Load Shedding DR program where a limit is placed on the aggregate load. For every time step during which the aggregate load exceeds this limit, loads that consumed power will be penalized in proportion with their Shapley value. The characteristic function representing such a DR program is $f(\mathcal{X}) = \left[\sum_{x\in\mathcal{X}} x - M\right]_+$. To evaluate this DR program, we take a small sample set $\mathcal{X}$ containing 20 loads. We can calculate an exact Shapley value for a representative load and in the process compare how the three different sampling techniques would have fared.

Figure 1 shows the mean and standard deviation for each stratum when using the above characteristic function. Figure 2 shows the variance of the Shapley value estimate for the sampling techniques. SD weighted sampling significantly reduced variance of the Shapley Value.

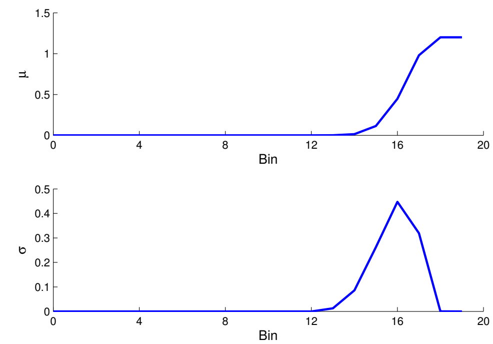

Figure 1: Strata Mean and Standard Deviation.

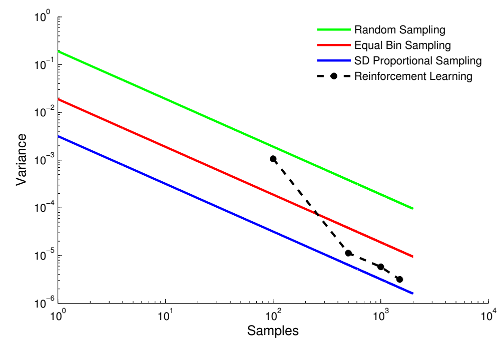

Figure 2: Change in variance with sample size.

5. REFERENCES

[1] M. Albadi and E. El-Saadany. Demand response in electricity markets: An overview. In *Power Engineering Society General Meeting, 2007. IEEE.*

[2] E. Algaba, J. Bilbao, J. F. Garcia, and J. Lopez. Computing power indices in weighted multiple majority games. *Mathematical Social Sciences*, 46(1):63 – 80, 2003.

[3] Y. Bachrach, E. Markakis, E. Resnick, A. D. Procaccia, J. S. Rosenschein, and A. Saberi. Approximating power indices: theoretical and empirical analysis. *Autonomous Agents and Multi-Agent Systems*, 20(2):105–122, Mar. 2010.

[4] M. Junqueira, L. Da Costa, L. Barroso, G. Oliveira, L. Thomé, and M. Pereira. An aumann-shapley approach to allocate transmission service cost among network users in electricity markets. *Power Systems, IEEE Transactions on*, 22(4):1532–1546, 2007.

[5] G. O'Brien and R. Rajagopal. A Method for Automatically Scheduling Notified Deferrable Loads. In *American Control Conference (ACC), 2013.* To appear.