Energy 154 (2018) 80–95

Contents lists available at ScienceDirect

Energy

journal homepage: www.elsevier.com/locate/energy

Collective learning of lithium-ion aging model parameters for battery health-conscious demand response in datacenters

A. Mamun [a], A. Sivasubramaniam [b], H.K. Fathy [a, *]

[a] Department of Mechanical and Nuclear Engineering, The Pennsylvania State University, University Park, PA 16802, USA
[b] Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA 16802, USA

A R T I C L E I N F O

Article history:
Received 31 May 2017
Received in revised form
17 March 2018
Accepted 12 April 2018
Available online 17 April 2018

Keywords:
Lithium-ion
Particle filter
Batteries
Health degradation
Collective learning
Demand response

A B S T R A C T

This paper examines the degree to which a large-scale datacenter employing lithium-ion (Li-ion) batteries for demand response can learn the physics-based aging and degradation dynamics of the underlying batteries by measuring their input/output current/voltage data. Battery degradation dynamics are chemistry dependent and change significantly for newer chemistries. Moreover, characterizing these degradation dynamics requires time-consuming and expensive laboratory testing. Together, these facts motivate the following question: is it possible to use battery current and voltage measurements to learn battery degradation behavior in a datacenter where numerous distributed batteries are being used for demand response? If so, what are the benefits and challenges associated with such learning? The goal of this paper is to provide preliminary answers to these questions building on earlier work by authors on health-conscious stochastic battery control. Specifically, we show that when datacenters exploit its demand management flexibility at the rack-level to control different batteries in accordance with different management policies, the resulting data is sufficiently rich to the point where (1) the learning of degradation behavior is possible within the span of approximately 1 year, (2) with a reasonable number of cells, (3) even when the batteries are used simultaneously for degradation learning and demand response.

1. Introduction

This paper examines the challenge of using datacenters as a cyber-physical resource for learning battery aging behavior when a large number of batteries is employed for demand response. The paper focuses on the development of a model-based technique for the collective learning of degradation dynamics. A preliminary study is performed to test the possibility of learning battery aging behavior using measurement data obtained from a large number of lithium-ion cells during demand response. The term demand response refers to any process that changes consumer power and energy demand based on electricity price [1]. In the context of a datacenter, demand response generally refers to the process of modifying the overall power demand profile to reduce amortized capital and operating expenses. Cost of power is a significant factor in datacenter operation and contributes around 30–40% of total monthly operating expenses [2]. A significant portion of this large power cost is a consequence of datacenter workload variations resulting from many factors, for example, scheduled virus scans, media services, and flash crowd visitor etc. [3]. The large power peaks can create disturbance in the system. The occasional power peaks increase the capital expenses since the datacenter power infrastructure need to be provisioned for the peak power. Monthly operating cost also increases since peak power draw is penalized as a consequence of higher peak load on the grid. In the year 2013, US datacenters overall used 91 billion kWh of electricity, at an estimated cost of $6.7b. By the year 2020, this consumption is projected to increase up to 140 billion kWh/year, which is equivalent to the output of 50 coal-fired power plants [4]. Therefore, even a small improvement in existing datacenter demand response policies will cut the electricity cost and significantly reduce carbon emission by lowering the peak load on the power grid. Power provisioning and demand response in datacenters are highly similar to the operation of microgrids in terms of distributed storage and fluctuating loads [5]. Several works in the microgrid area focus on implementing hybrid energy storage systems including fuel cells [6], Li-ion

* Corresponding author.
E-mail addresses: aamamun@stanford.edu (A. Mamun), anand@cse.psu.edu (A. Sivasubramaniam), hkf2@engr.psu.edu (H.K. Fathy).

https://doi.org/10.1016/j.energy.2018.04.070

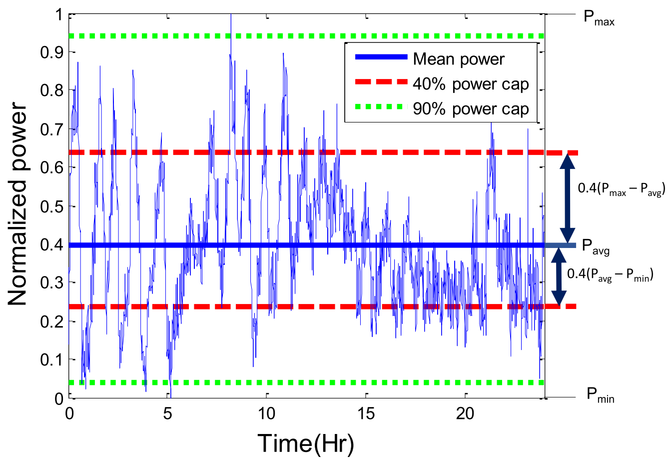

Fig. 1. Power cap is defined as a limit on maximum and minimum power drawn from the grid.

batteries and renewable generation [7] which are relevant to datacenter power management. Transient stability [8] of power systems and fault analysis [9,10] in microgrid also play important role in power management using energy storage devices.

Due to high energy density and reduced deployment costs, Li-ion battery storage is currently being considered for demand response in datacenters [11,12]. When performing demand response with Li-ion batteries, battery degradation behavior and lifetime play a significant role in the overall savings. Moreover, the deployment of new and more efficient battery chemistries might also improve the demand response savings. These insights are explained further in later sections. They act as primary motivations of this paper to study the feasibility of using datacenter as a cyber-physical laboratory to learn battery aging behavior in parallel to using them for demand response. If Li-ion batteries are used as uninterruptible power supply (UPS) systems in datacenters, the high capital cost of Li-ion batteries can be compensated over time by their long calendar life [13]. However, using Li-ion batteries for demand response requires frequent charging and discharging which will impact battery health in terms of power and energy capacity. If not optimized for health, a demand response control policy might cause premature failure and shorter end of life (EOL) of batteries and offset the economic benefit of demand response. Recent work by Liu et al. shows that battery aging-aware power management in a green datacenter can extend battery life up to 69% [14]. Preliminary analysis by the authors also show that battery health-conscious demand response policy can reduce battery degradation significantly compared to a policy that is not optimized for health (Fig. 2) [15]. In that analysis, the demand response policy is to limit grid power draw within a certain range called "power cap". "Power cap" is defined as a percentage of the difference between the maximum and minimum power compared to an average power demanded by the datacenter. Fig. 1 explains the definition with the examples of 40% and 90% power cap given a sample power demand. Smaller percentage value of power cap means that lower peak power will be drawn from the grid which will result in higher savings. This definition of a power cap is used throughout this paper and not necessarily congruent with the definitions adopted by the existing datacenter demand response literature. However, the intent is to use the power cap constraint for different datacenter demand response policies as an indicator of these policies aggressiveness.

Fig. 2.a demonstrates how battery health (solid-electrolyte interphase (SEI) layer-based capacity fade) is affected by different levels of aggressiveness of demand response policies (i,e,. power cap). It shows that a health-conscious stochastic control policy for demand response can reduce battery health degradation upto 88% compared to a stochastic control policy that is not optimized for health. Fig. 2.a also shows that a tight power cap (e.g., 40%) has a smaller room for improvement in long-term battery aging. The reason behind that is the requirement of deep charging and discharging to shave large peaks and the necessity of maintaining a high average state of charge (SOC) to allow such deep cycling. Fig. 2.b shows the combined capital expense (CapEx) and operating expense (OpEx) savings obtained from demand response with and without optimizing for battery health. For both cases, the savings and health degradation increase almost linearly with power cap. The slope of the two fitted lines shows combined CapEx and OpEx savings per unit capacity of the energy storage. When demand response is performed in a health-conscious manner, the combined CapEx and OpEx savings are \$62/Ahr of installed Lithium-ion battery capacity. The savings diminish to \$29/Ahr when battery health is not optimized during demand response. Therefore, health-conscious stochastic control can increase the savings more than 100% compared to a stochastic demand response policy that is not optimized for health. Given the highest estimated market price of Li-ion energy storage in 2015 (\$600/kWh [16]), the nominal cell voltage (3.3 V), and effective usable capacity (EOL at 20% capacity fade), the cost of usable Li-ion cell capacity becomes \$0.9/Ahr. Since this cost is much

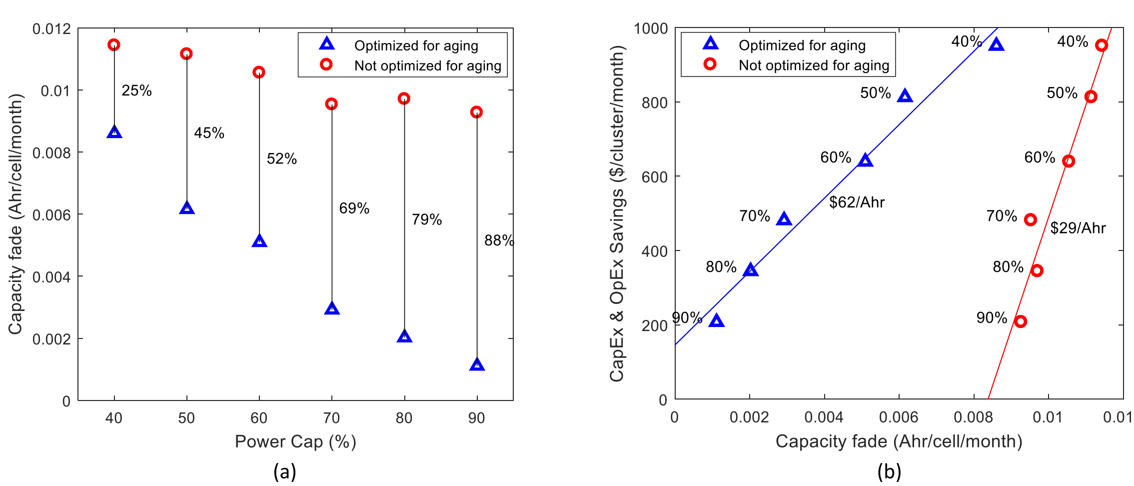

Fig. 2. Optimal (a) capacity fade per month and (b) dollar savings per month by implementing control policy obtained by solving Stochastic Dynamic Programming (SDP) problem when different power caps are imposed on grid power.