Journal of Energy Storage 7 (2016) 258–269

Contents lists available at ScienceDirect

Journal of Energy Storage

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

Multi-objective optimization of demand response in a datacenter with lithium-ion battery storage

A. Mamun[a], I. Narayanan[b], D. Wang[c], 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 & Engineering, The Pennsylvania State University, University Park, PA 16802, USA
[c] Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA

ARTICLE INFO

Article history:
Received 12 May 2016
Received in revised form 2 August 2016
Accepted 3 August 2016
Available online xxx

Keywords:
Lithium-ion battery
Battery aging
Datacenters
Demand response
Optimization
Power management

ABSTRACT

This article optimizes lithium-ion battery management in a datacenter to: (*i*) maximize the dollar savings attainable through peak shaving, while (*ii*) minimizing battery degradation. To the best of the authors' knowledge, such multi-objective optimal datacenter battery management remains relatively unexplored. We solve this optimization problem using a second-order model of battery charge dynamics, coupled with a physics-based model of battery aging via solid electrolyte interphase (SEI) growth. Our optimization study focuses on a classical feedforward-feedback energy management policy, where feedforward control is used for peak shaving, and feedback is used for tracking a desired battery state of charge (SOC). Three feedforward-feedback architectures are examined: a proportional (P) control architecture, a proportional-integral (PI) architecture, and a PI architecture with a deadband in its feedforward path. We optimize these architectures' parameters using differential evolution, for real datacenter power demand histories. Our results show a significant Pareto tradeoff between dollar savings and battery longevity for all architectures. The introduction of a deadband furnishes a more attractive Pareto front by allowing the feedforward controller to focus on shaving larger peaks. Moreover, the use of integral control improves the robustness of the feedback policy to demand uncertainties and battery pack sizing.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

This article examines the use of electrochemical batteries for datacenter demand response (DR). The article focuses on minimizing a Pareto combination of total electricity cost and battery aging by optimizing the control policy used for demand response. This work is motivated by the rapid growth of both the datacenter industry and its electricity needs. The cost of energy is a significant factor in datacenter operation. A large, energy efficient datacenter may, for instance, need as much as 10 MW of electricity and spend approximately $1 million on electricity bills per month, which is 30–50% of its monthly operating cost [1]. The high operating cost of purchasing electricity and the capital cost of the power infrastructure to draw electricity from the grid is affected by datacenter workload characteristics. Datacenter workloads often fluctuate significantly due to scheduled virus scans, media and cloud services, and flash crowd visitors [2,3]. Operating expenses (OpEx) are affected by these fluctuations since utility companies often charge separately for peak power. Demanding a large amount of power during peak hours imposes additional cost penalties due to spot pricing or time-of-day tariffs [4]. A large capital expenditure (CapEx) is also required to enable the delivery of peak power through the power infrastructure, even though the probabilities of requiring peak power are often very low [5]. Fig. 1 shows the normalized power demand for 7 days for one of the clusters in a Microsoft datacenter [6]. The figure shows a wide power demand range including a few large peaks and significant fluctuations in cluster-level power demand.

The economic penalty associated with fluctuating datacenter power demand can be reduced through demand response. The term "demand response" refers to any process that changes consumer electricity demand based on electricity price [1]. Researcher have considered different information theoretic (IT) [7–10] and cooling capacity knobs [11–13] to perform demand response in datacenters. The IT knobs are used to (i) "throttle" workload by reducing server speeds when electricity is expensive, (ii) "shift" workload temporally from peaks to adjacent valleys, or (iii) "transfer" the workload to other datacenters experiencing

* Corresponding author.
E-mail address: hkf2@engr.psu.edu (H.K. Fathy).

http://dx.doi.org/10.1016/j.est.2016.08.002
2352-152X/© 2016 Elsevier Ltd. All rights reserved.

Nomenclature

A_n	Sheet area of the negative electrode (m^2)
C_{cap}	Capital cost of power provisioning ($\$\,W^{-1}(12\ years)^{-1}$)
C_e	Concentration of lithium-ion in the electrolyte ($mol\,m^{-3}$)
C_{energy}	Unit cost of energy ($\$\,(kWh)^{-1}$)
C_{peak}	Unit cost of peak power ($\$\,kW^{-1}\,month^{-1}$)
C_r	Charge capacity of the capacitor in the RC pair (F)
$C_{s,n,max}$	Maximum concentration of lithium-ion in the negative electrode ($mol\,m^{-3}$)
$C_{s,n,surf}$	Surface concentration of lithium-ion in the negative electrode ($mol\,m^{-3}$)
E_{charge}	Energy stored in the battery (Wh)
$E_{discharge}$	Energy released from the battery (Wh)
F	Faraday's constant ($C\,mol^{-1}$)
I	Input current (A)
J_{main}	Volumetric current density of the main reaction ($A\,m^{-3}$)
J_{side}	Volumetric current density of the side reaction ($A\,m^{-3}$)
J_{tot}	Total intercalation current ($A\,m^{-3}$)
L_n	Thickness of the negative electrode (m)
M_p	Molar mass of SEI layer material ($kg\,mol^{-1}$)
P_{batt}	Battery input power (W)
$P_{batt,dis}$	Battery discharge power (W)
$P_{cap,des}$	Desired power cap for "ideal" load leveling (W)
$P_{cap,max}$	Peak power demanded from utilities with demand response (W)
P_d	Datacenter power demand (W)
$P_{d,avg}$	Average datacenter power demand (W)
$P_{d,max}$	Maximum datacenter power demand (W)
$P_{d,min}$	Minimum datacenter power demand (W)
Q	Cell charge capacity (C)
Q_{loss}	Capacity lost due to side reaction (Ah)
R	Lumped resistance of the cell and the connector (Ω)
R_{film}	Total film resistance ($\Omega\,m^{-2}$)
R_r	Resistance of the resistor in the RC pair (m^{-2})
R_{SEI}	Initial resistance of SEI layer ($\Omega\,m^{-2}$)
R_u	Universal gas constant ($J\,mol^{-1}$)
SOC_{des}	Desired SOC of the cell
T	Cell temperature (K)
$U_{n,ref}$	Main reaction equilibrium potential of the negative electrode (V)
$U_{s,ref}$	Side reaction equilibrium potential of (V)
V	Terminal voltage (V)
V_{oc}	Open circuit voltage (V)
X_1	Cell state of charge
X_2	Charge in the capacitor in the RC pair (C)
Z	Optimization objective
a_n	Specific surface area of the negative electrode (m^{-1})
$i_{o,n}$	Main reaction exchange current density in the negative electrode ($A\,m^{-2}$)
$i_{o,s}$	Side reaction exchange current density ($A\,m^{-2}$)
k_n	Reaction rate constant at the negative electrode ($A\,m^{-2}\,(mol\,m^{-3})^{1.5}$)
$x_{n,surf}$	Surface state of charge of the negative electrode
Δj	Potential difference between solid and electrolyte (V)
α, β	Proportional gain
α_a, α_c	Apparent charge transfer coefficients for anode and cathode
γ	Deadband width gain
δ	Integral gain
δ_{film}	Time varying SEI film thickness (m)
η_n	Main reaction overpotential at the negative electrode (V)
η_s	Side reaction overpotential (V)
κ_P	Conductivity of the SEI film
ρ_p	Density of side reaction product ($kg\,m^{-3}$)
τ	Time constant (s)

either smaller workloads or cheaper energy availability (or both). Since these knobs involve tradeoffs between energy cost and performance degradation, energy storage has come into play as an additional knob. The idea is to store grid electricity when it is either abundant or inexpensive (or both) and use the stored energy during high workload or peak hours.

The literature already examines the problem of utilizing existing storage in datacenter uninterruptible power supply (UPS) systems or using a separate battery pack for demand response [1,4,5,14]. The primary research focus of demand response in datacenters with UPS is the minimization of total cost of ownership (TCO) [15,16]. This is the sum of amortized capital and operating costs over a time horizon [17]. Queuing theory-based Lyapunov optimization techniques are also used to find near-optimal solution to minimize monthly electricity bills [1,18]. These optimization analyses typically focus on lead-acid battery storage, although lithium-ion batteries are also considered by some researchers. For demand response optimization, lithium-ion batteries are generally modeled purely as charge integrators, and battery lifetime is modeled using depth of discharge (DOD) based lifetime charts and/or charge processed models [16,19–21]. These models are quite limited in their ability to capture the fundamental physical phenomenon affecting lithium-ion battery behavior. For example, an elementary charge integration model does not capture internal battery diffusion dynamics, and therefore fails in capturing the dynamic constraints imposed on battery charging/discharging by diffusion. One goal in this article is to extend the existing literature on datacenter demand response through the use of a lithium-ion battery model that captures both charge integration and voltage relaxation dynamics.

In addition to using a suitable battery model, a control scheme is also necessary to optimally utilize batteries for demand response. Optimal battery utilization in this case means using the batteries to minimize the electricity cost as much as possible with minimum battery degradation. Load leveling using batteries will reduce the discounted monthly capital and operating cost associated with datacenter electricity demand. However, deep charge and discharge to shave large peaks might accelerate battery aging. Therefore using battery charge capacity for load leveling during demand response becomes a multi-objective optimization problem, where minimizing battery health degradation and electricity costs are the two competing objectives. To the best of our knowledge, the extensive study of control strategies for model-based, health-aware battery control during demand response remains relatively unexplored in the datacenter power management literature. Previous work by authors addresses this issue by designing an optimal proportional controller for Li-ion batteries in large-scale datacenters [22]. The work assumes the availability of an optimally sized, power-efficient storage system (i.e., flywheels or ultracapacitors) for emergency power plus a separate Lithium-ion battery pack for demand response since such a hybrid storage solution is often more cost-effective than a battery-only solution [23]. In this previous work we (i) build a second-order model of a Lithium-ion battery that captures both ohmic and diffusion

Download English Version:

https://daneshyari.com/en/article/1133082

Download Persian Version:

https://daneshyari.com/article/1133082

Daneshyari.com