Lowering Peak Demands with Batteries

Matthew P. Johnson Amotz Bar-Noy Yi Feng Ou Liu

Department of Computer Science
The Graduate Center of the City University of New York

There is increasing interest in saving fuel costs by use of renewable energy sources such as wind and solar power. Although such sources are highly desirable, and the power they provide is in a sense free, the typical disadvantage is unreliability: availability depends e.g. on weather conditions (it is not "dispatchable" on demand). Many companies seek to build efficient systems to gather such energy when available and store it, perhaps in modified form, for future use [5].

On the other hand, power companies charge some high-consumption clients not just for the total amount of power consumed, but also for how quickly they consume it. Within the billing period (typically a month), the client is charged for the amount of energy used (*usage charge*, in kWh) and for the maximum amount requested over time (*peak charge*, in kW). If demands are given as a sequence $(d_1, d_2, \ldots, d_n)$, then the total bill is of the form $c_1 \sum_i d_i + c_2 \max_i \{d_i\}$ (for some constants $c_1, c_2 > 0$), i.e., a weighted sum of the total usage and the maximum usage. (In practice, the discrete timeslots may be 30-minute averages [2].) This means that a client who powers a 100kW piece of machinery for one hour and then uses no more energy for the rest of the month would be charged more than a client who uses a total of 100kWh spread evenly over the course of the month. Since the per-unit cost for peak charges may be on the order of 100 times the per-unit cost for total usage [3], this difference can be significant.

This suggests a second use for the battery: to store *purchased* energy for future use. Indeed, at least one start-up company [1] is currently marketing such a battery-based system intended to reduce peak energy charges. In such a system, a battery is placed between the power company and a high-consumption client site, in order to smooth power requests and shave the peak. The client site will charge to the battery when demand is low and discharge when demand is high. Spikes in the demand curve can thus be rendered consistent with a relatively flat level of supplied power. The result is a lower cost for the client and a more manageable request curve for the provider.

It is interesting to note that a battery system may actually *raise* energy usage, since there may be energy loss due to inefficiency in AC/DC conversion. Serving peak requests during periods of high demand is a difficult and expensive task for the power company, however, and the event of a black-out inflicts high societal costs. While a battery system may involve higher total energy requests, it may benefit the system as a whole by easing the strain of peak demands. Combined with alternative energy sources such as solar panels, the system could even lower the net commercial power usage. Alternative energy sources are typically low-cost but unreliable, since they depend on external events such as the weather. With a battery, this energy can be stored until needed.

We may generalize this problem of minimaxing the request to any resource which is *tenable* in the sense that it may be obtained early and stored until needed. For example, companies frequently face shortages of popular products: "Plentiful supply [of Xboxes] would be possible only if Microsoft made millions of consoles in advance and stored them without releasing them, or if it built vast production lines that only ran for a few weeks–both economically unwise strategies," a recent news story asserted [4]. A producer could smooth the product production curve by increasing production and warehousing supply until future sales. But when should the producer "charge" and "discharge"? (In some domains, there may also be an unpredictable level of volunteer help.) A third application is the scheduling of jobs composed of generic work-units that may be done in advance. Although the problem is very general, we will use the language of energy and batteries for concreteness. Many features of this production problem, including uncertainty in future demand, a bounded warehouse size and a cost for storage in the warehouse have analogs in the battery problem.

In the online version of our problem, the essential choice faced at each timeslot is whether (and by how much) to invest in the future or to cash in a prior investment. The investment in our setting is a request for more energy than is needed at the time. If the algorithm only asks for the minimum required, then it is vulnerable to spikes in demand; if it asks for much more energy than it needs, then the greater request could itself introduce a new, higher peak. The strictness of the problem lies in the fact that the cost is not cumulative: we want *every* request to be low.

Acknowledgments. This work was supported by grants from the NSF (grant number 0332596) and the New York State Office of Science, Technology and Academic Research.

REFERENCES

[1] Gaia Power Technologies. gaiapowertech.com.
[2] ConEd electricity rates document. www.coned.com/documents/elec/043-059h.pdf.
[3] Orlando Utilities Commission website. www.ouc.com/account/rates/electric-comm.htm.
[4] T. Harford. The great Xbox shortage of 2005. *Slate*, Dec. 15, 2005 www.slate.com/id/2132071/.
[5] M. L. Wald. Storing sunshine. *The New York Times*, July 16, 2007 www.nytimes.com/2007/07/16/business/16storage.html.