U.P.B. Sci. Bull., Series C, Vol. 78, Iss. 2, 2016 ISSN 2286-3540

SMART GRID ENERGY MANAGEMENT SYSTEMS: CONTROL MECHANISMS FOR ELECTRICITY DEMAND

Andrei HORHOIANU[1], Mircea EREMIA[2]

The paper deals with the mechanisms that can be used to optimize electricity consumption both within an oil processing plant and across multiple plants in a production site. Assuming that the production plants are using energy management controllers (EMCs) to control the operation of some of their equipment, we demonstrate, using simulation, that some solutions may in fact be more peaky than the "non-scheduled" solution, thereby negating some of the benefits (for the utility) of off-peak pricing models. In the end of the paper, a distributed scheduling mechanism to reduce peak demand within a whole production site is presented as an optimal solution.

Keywords: peak demand, energy management systems, load scheduling

1. Introduction

Two main approaches to development of energy management systems (EMS) for a smart grid, can be found in the literature, each having its own advantages and drawbacks. These two approaches are described and analyzed next.

A. Centralized energy management system (CEMS)

A few cases of CEMS have been discussed in the literature, taking into account varied smart grid characteristics and configurations. A general framework for the development of CEMS is proposed in [1], while a CEMS for a smart grid composed of hydrogen storage and wind power, utilizing a dynamic linear programming (LP) formulation is presented in [2]. An LP solution technique together with heuristics is proposed in [3] for the implementation of a CEMS for a PV storage smart grid, while [4] proposes a purely heuristic optimization approach.

A typical CEMS architecture is shown in Fig. 1, where a central agent collects all the relevant information from the different smart grid actors to perform an optimization and determine the inputs of the control system for the next period.

[1] PhD Student, Dept. of Electrical Power Systems, University POLITEHNICA of Bucharest, Romania, e-mail: andrei_horhoianu@yahoo.com
[2] Professor Emeritus, Dept. of Electrical Power Systems, University POLITEHNICA of Bucharest, Romania, e-mail:eremia1@yahoo.com

Depending on the particular resources present in the smart grid, the input variables of the CEMS can be [5]:

- Forecasted power output of the non-dispatchable generators for the following N consecutive periods.
- Forecasted local load for the following N consecutive periods
- State of charge of the energy storage systems (ESS).
- Operational limits of dispatchable generators and ESS.
- Security and reliability constraints of the smart grid.
- Interconnection status.
- Main grid energy price forecasting

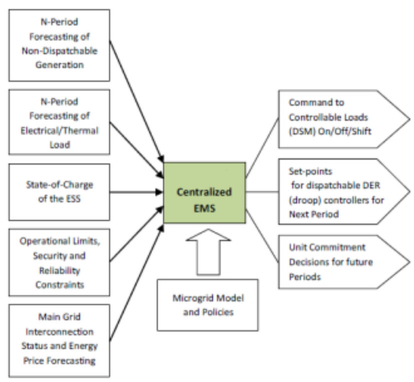

Fig. 1. Smart Grid Centralized Energy Management System [5]

Once all the input variables are gathered in the CEMS, a multi-stage optimization is performed in order to determine the optimal dispatch of units according to a defined cost function, over a pre-specified time frame. Output variables of the EMS are the reference values of the control system (e.g., output power and/or terminal voltage) for each dispatchable DER, together with binary decision variables for connecting or disconnecting loads for load shifting. An additional output variable is the unit commitment (UC) decision of the dispatchable generators (if required); however, this problem can be solved at a lower frequency than the dispatch, and separately. The main advantages of the centralized approach are: allowing for a broad observability of the smart grid, and suitability for application of optimization techniques. Some of its disadvantages are: reduced flexibility, as it needs to be modified to incorporate additional generators, and extensive computational requirements to perform the optimization [5].

B. Distributed energy management system (DEMS)

A DEMS based on multi agent systems (MAS) for microgrids was first proposed in [6] as an alternative for coordinated operation of smart grids in a competitive market environment and with multiple generator owners. The relevant smart grid actors are grouped and represented by different agents that interact in a market environment in order to determine the operation of the smart grid. In this way, consumers, generators, ESS and the main grid participate in the market by buying and selling bids to the Central Microgrid Operator (CMO) based on their particular needs, availability, cost functions, technical limitations, expectations and forecasts. The CMO is responsible for the settlement of the smart grid market by matching buying and selling bids maximizing the social welfare, while ensuring the feasibility of the resulting operation plan. Additional agents assigned to different tasks such as load shifting and load curtailment to allow demand side management are proposed in [7] as well. The MAS-DEMS approach allows almost autonomous operation of the generating units in a smart grid, and reduces the need for manipulation of large amounts of data, thus reducing computation time. Another important advantage of DEMS is its flexibility, as it provides the plug-and-play feature, facilitating the installation and coordination of additional DER in the smart grid. On the other hand, DEMS based on MAS shows disadvantages compared to CEMS when applied to smart grids that require strong cooperation between the different DER in order to operate the system in a secure and reliable way. A typical DEMS model for a microgrid operating in grid-connected mode is shown in Fig. 2.

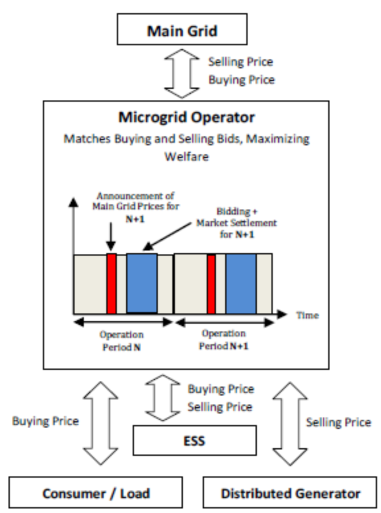

Fig. 2. Distributed Energy Management System [5]

In the case of isolated smart grids operating in stand-alone mode, the small number of generators in the gird and the uneven share of installed power, as well as the lack of a strong price signal from the main grid, make a DEMS more difficult to implement.

The emerging smart grid will provide industrial users flexibility in controlling their electricity costs. A primary driving force is the *smart meter*, which can deliver “real-time” electricity prices to all site facilities, potentially every fifteen minutes . The customer can make use of this information via an in-house *energy management controller* (EMC), which uses both prices and user preferences to control power usage across the plant. The EMC may be standalone or embedded either in the smart meter or in equipment. At the same time, customers may participate in Direct Load Control (DLC) that allows utilities to control some power consumption within an industrial plant during peak usage (thereby bypassing or even replacing the EMC) as a part of an energy savings subscription plan.

Such methods for controlling electricity consumption are part of *demand response*, which relies on varying the price of electricity throughout the day in order to reduce peak demand. Reduced peak demand lowers electricity bills and benefits utilities by reducing complexity of grid stability, equipment overload, brownouts, and blackouts. It also enables utilities to comply with government mandates to cut peak demand.

In this paper, we consider mechanisms to optimize electricity consumption both within an oil processing plant and across multiple plants in a production site. At the single oil processing plant level, we assume that equipment communicate with each other and with the EMC over an in-house area network (IHAN). Many equipment offer some degree of flexibility when they are operated. We present a simple optimization model that exploits this flexibility in order to determine the optimal timing of equipment operation, taking into account both electricity price fluctuations and user preferences. This optimization model represents the logic embedded in a simple EMC and aims to minimize customer’s cost, measured as a combination of the dollars paid to the utility and the inconvenience of delayed equipment use. Using simulation, we demonstrate that, if multiple plants each optimizes their equipment usage to take advantage of off-peak energy prices, the problem of demand peaks is not alleviated; rather, the demand peak simply shifts to the off-peak period, creating a new “rebound” peak [8] that is even more exaggerated. Thus, from the perspective of the utility, this optimal solution (for the customer) reduces the effectiveness of off-peak pricing models.

Next, we propose a distributed scheduling mechanism to reduce peak demand across a collection of local plants in an oil production site. Our scheme relies on a common control channel among a collection of plants that permits EMCs in each plant to communicate with each other. (The security and privacy concerns of exchanging power consumption data across plants is beyond the scope of this paper but is an important consideration in the practical implementation of this approach. Further, the control channel required for this scheme can be derived from the Automated Meter Reading Infrastructure.) The utility chooses a maximum allowable peak power consumption level for each plant, while also guaranteeing a minimum available level of power for each plant at all times. If a plant wishes to use more power than its minimum guaranteed level, then its EMC uses our proposed mechanism to contend with other plants' EMCs for the remaining available power. The proposed distributed scheduling scheme is designed to ensure that the total consumption across the production site does not exceed the target peak demand. Therefore, the utility can reduce the peak of its demand load (without simply rationing each plant's usage), no matter what pricing scheme is used.

Finally, we introduce a new EMC optimization model for a single plant optimization. This model accounts for the fact that, due to the proposed distributed scheduling scheme, a plant may not receive as much power as desired, and therefore the scheduling problem becomes much harder. It is therefore much more realistic than the first optimization model. On the other hand, it is more computationally intensive.

The problem of scheduling electricity consumption across multiple plants can be addressed by considering that the plants are analogous to "energy aggregators". This solution consists in a central hub to coordinate communications and energy control across these aggregators. In contrast, we propose here a decentralized demand response scheme for multiple plants located in an oil production site.

2. Single Plant Scheduling

A. Optimization Model

In this section, we consider the problem of optimizing equipment timing within a single plant. The planning horizon for the model consists of T discrete time periods. The plant has N equipment. Each equipment n, when on, consumes an amount c_n of power (measured in kW). (When off, the equipment is assumed to consume 0 kW, though this assumption can easily be relaxed.) User requests for equipment n are random, as is the length of time the equipment remains on. If equipment n is off in period t, then the probability that it is requested in period $t+1$ is given by λ_{nt}. Similarly, if equipment n is on in period t, then the probability that it completes its operation and turns off in period $t+1$ is given by μ_{nt}. Note that these probability parameters may vary over time and across equipment. A special case occurs when $\lambda_{nt} = \lambda_n$ and $\mu_{nt} = \mu_n$ for all t, in which case the duration of equipment *n*'s off and on times are geometrically distributed with probability parameters λ_n and μ_n, respectively.

When the user requests equipment *n*, the EMC may turn it on immediately, or it may choose to delay turning it on. The user specifies a maximum allowable delay of d_n time periods for equipment *n*, with $d_n \geq 0$. (If $d_n = 0$, then the equipment is "non-schedulable" and must be turned on immediately when the user requests it.) Each period of delay incurs a cost of $\psi_n^1 \geq 0$, which represents the inconvenience to the user introduced by the delay and is measured in \$/hr.

The cost of electricity in period *t* is denoted by π_t and measured in \$/kWh. We assume for simplicity that electricity prices are deterministic (but dynamic) throughout the planning horizon, but if the prices are in fact stochastic, then π_t can simply be replaced by its mean. If equipment *n* is requested in period *t*, the decision of when to turn it on is simple: we must simply find the *s* that solves:

$$\min_{t \leq s \leq t+d_n} (s-t)\,\psi_n^1 + \sum_{r=s}^{T} \left(\prod_{i=s}^{r-1} (1-\pi_{ni})\right) \pi_r c_n \quad (1)$$

The first term represents the delay cost (in dollars) incurred by waiting until period *s* to turn the equipment on. The second term (also measured in dollars) represents the expected energy cost while the equipment is on. The product term within the second term calculates the probability that the equipment is still on in period *r*. (We take the product $\prod_{i=s}^{r-1}$ to equal 1 if *r* = *s*). This minimization problem can be solved in $O(T^2)$ time since the product over *i* can be updated within the loop that calculates the sum over *r*, using one operation per iteration. Therefore, the problem of optimizing all *N* equipment can be solved in $O(NT^2)$ time. This algorithm works since each equipment can be optimized individually. This approach would not work if, for example, electricity prices were a non-linear function of the load or if (as in Section 3) there are constraints on the power used by a given plant.

B. Simulation

We simulated an oil production site consisting of 50 oil gathering plants, each containing 3 equipment that are scheduled by an EMC. The equipment parameters are summarized in Table 1. The request probability λ_{nt} was assumed to vary throughout a 24-hour period, with a minimum of "Min λ_{nt}" and a maximum, attained at 6:00 PM, of "Max λ_{nt}". On the other hand, the μ_{nt} values were assumed to be stationary. So, for example, the water pump has a peak usage level of λ =0.0704, corresponding to a mean inter-request time of 13.7 hours; its lowest usage level is λ =0.010, corresponding to a mean inter-request time of 100 hours; and the water pump operates for a mean of 3 hours. Power usage values are

based on the operational records available for typical oil gathering plant composed of 1 water pump, 1 oil pump and 1 gas compressor.
Electricity prices are assumed to take two levels, corresponding to peak and off-peak hours. During the peak period, from 10:00 AM to 10:00 PM, electricity costs \$0.21/kWh, and at all other times it costs \$0.014/kWh. (These are actual rates from Con Edison's time-of-use pricing model in New York City [9]).

The system was simulated for 5 days (120 hours), with the first and last day omitted from the results as warm-up and warm-down intervals. Although Table 1 reports parameters in terms of hours, the simulation and scheduling algorithm used 10-minute time periods, and all parameters were adjusted accordingly. We simulated the system twice, once assuming that an EMC schedules the equipment using the model discussed in Section 2 A and once assuming that no scheduling is performed and equipment are turned on as soon as they are requested. The two simulations used the same sample path of equipment request times and "on" durations.

Table 1

Equipment Parameters

Name	c_n	d_n	ψ_n^1	ψ_n^2	Min λ_{nt}	Max λ_{nt}	μ_{nt}
Water Pump	1.8	6	0.10	2.5	0.0100	0.0704	0.283
Oil Pump	3.4	4	0.25	2.5	0.0392	0.1193	0.632
Gas Compressor	5.0	2	0.40	5.0	0.0952	0.2078	0.865

The average cost per plant in the non-scheduled system over days 2–4 is \$7.29, compared to an average cost in the scheduled system of \$6.24 for energy costs and \$0.68 in delay costs. Therefore, the EMC scheduler saves the average plant 14.5% in energy costs, or 5.1% when delay costs are also factored in. The EMC scheduling model is therefore effective in reducing consumers' costs.

Unfortunately, it also defeats the purpose of the utility's off-peak pricing scheme. Fig. 3 plots the electricity consumption, in kW, across all 50 plants, for days 2–4 of the simulation. The shaded bands represent peak pricing periods. Although the scheduled system reduces the peak 6:00 PM demand to a small extent, it also creates a new, even larger "rebound" peak, immediately after the off-peak prices begin. The maximum load is larger in the scheduled system than in the non-scheduled system (83.2 vs. 68.0 kW), as is the standard deviation of the load across periods (13.2 vs. 11.3 kW). We note that this rebound peak occurs even though the request intensity during off-peak hours is low. Therefore, this simulation demonstrates that, at least under certain assumptions about consumer usage patterns and electricity prices, the off-peak pricing model fails to achieve its goal of reducing load peak, and may even worsen the problem. In the next section, we propose a mechanism that the utility can use to ensure a more level load throughout the day.

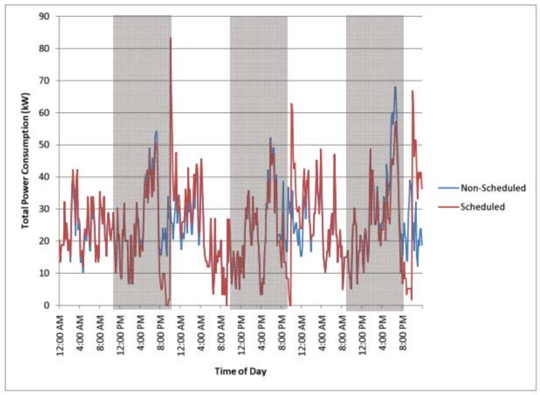

Fig. 3. Simulation results: Total power usage by period

3. Site Level Scheduling

We next consider a decentralized approach to support site-level load scheduling. To do so, we assume that there exists a communication network between the EMCs in each plant located within the analyzed site. To simplify discussion here, we assume that all EMCs are "one hop" away and can transmit/receive each other's signal sent over a common control channel. This channel may be supported by an underlying smart metering infrastructure or some other local communication network. Access to the control channel is granted to the EMCs by following the protocols of the supporting network. We assume here that the control channel has high capacity and supports the exchange of EMC scheduling packets with high priority and high reliability.

We describe below a scheduling-level exchange of packets over this channel that enables each EMCs to compete for power while collectively maintaining a relatively low and level peak demand. In contrast to purely pricing-based resource allocation, our proposed scheme does not require local EMCs to compete in a retail-level capacity market (where power is auctioned and bought). Instead, it extends access control methods typically used in communication networks to randomize an EMC's access to the local power capacity for a site. Specifically, we assume that time is divided into *scheduling slots* and a peak total power demand for the site is set (by the utility) as $P_{max,t}$ for slot t[2]. The scheme described here aims to meet $P_{max,t}$, which the utility can fix to be below the typical peak demand for time *t*.

To describe our scheme, we will first assume (in Section 3 A) that the plants must compete for *all* of the available power. Presumably, such a scheme would be unacceptable to consumers, who would want a guarantee that at least some of their power needs can be met at all times. Therefore, in Section 3 B, we modify the scheme to guarantee each plant a certain level of power at all times; each plant may then compete for additional power.

A. No Guaranteed Minimum

We assume that whenever a new load request is generated within a plant, the EMC will seek to meet the load requirement by coordinating with other neighborhood plants over the common control channel. During scheduling slot *t*, all loads that are currently being supported are termed as *active*. All EMCs that have an active load will continuously monitor the common control channel. When a new load is requested within a plant, its EMC will first require information about current active loads for the neighborhood. To do so, it will transmit a *probe packet* over the control channel in the next scheduling slot. To support transmissions from multiple EMCs in a given scheduling slot, we assume the slot is itself divided into *M* "minislots." We assume each EMC selects a random minislot (uniformly distributed in the range [*1,M*]) to send out a probe message. With sufficiently large *M* and slot durations, an EMC's probe transmission can be transmitted with negligible probability of collision with another EMC's transmission.

Once the first EMC's (say EMC *i*'s) probe transmission goes through successfully, all other contending EMCs cancel their probe transmission and wait another random number of minislots. This time the number of minislots chosen for the back off is uniformly distributed in the range from *m* to *M*, where *m* minislots is a duration long enough to obtain responses from the EMCs currently supporting active loads and one additional transmission from EMC *i*. This is because the successful probe message will require all EMCs supporting active loads to respond with a short *response packet* containing the power levels of their supported loads. We assume that in a site of *K* plants, the *m* minislots are long enough to support *K* such packets. EMCs from the K plants are assigned a transmission order during initial network formation and EMCs with active loads respond to the probe packet in this order over the *m* minislots.

Based on the response from the other EMCs, the EMC requesting the new load then computes the current total power usage and determines if its desired power demand can be supported within the total allowable site load of $P_{max,t}$. If so, the load joins the active set and sends a short *admission packet* containing its power consumption over the control channel; otherwise, the EMC enters random back off *at the scheduling layer* (for this load request) and re-attempts the inquiry procedure after its back off timer expires. The scheduling layer back off mechanism recognizes that this EMC's load cannot be supported in this current scheduling slot and waits for a random number of scheduling slots before reattempting a probe message and admission to the active set. We assume that at the scheduling layer, an EMC entering back off for a given load will select a random number of scheduling slots uniformly distributed in the range $[1,w]$ for some integer w.

Other EMCs attempting a new load admission in the current time slot (that had entered a second back off after the initial probe message) continue to monitor the control channel over these m minislots. Based on the probe response packets and the presence/absence of an admission packet, these EMCs know the current total power usage for the site at the end of the m minislots. If their requested power can be supported, they continue to monitor the channel until their bac koff timer (for the minislots) expires and then simply join the active set (if the total usage permits their admission within the constraint of $P_{max,t}$) and send out an admission packet indicating the power level of this new active load.

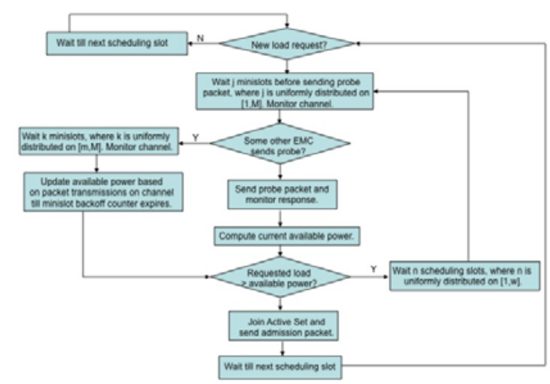

Fig. 4. Distributed Scheduling Flow Chart

The steps of this distributed scheduling algorithm are presented in Fig. 4, which shows the procedure that must be followed by each EMC in the site. At the same time, those EMCs that have active loads must monitor the channel to respond to any probe messages sent out in each scheduling slot. As noted above, the transmission times for these response packets are prescribed when the site EMCs are initially connected on the control channel.

B. Guaranteed Minimum

The mechanism described in the previous section would be unpalatable to consumers since it may result in them receiving little or no power at certain times. We now modify the mechanism to guarantee each plant a minimum level of power at all times, while also allowing a plant to compete for additional power. Specifically, we assume now that each plant from the site is allotted a base level of electrical power. For ease of exposition, we assume that each plant is allotted the same base power level P_t, although our mechanism can easily be adapted if this level differs from plant to plant. The process for determining the base power level is beyond the scope of this paper, but we note that it may be set in a number of ways. For example, it may simply be imposed by the utility, or the utility may offer consumers a menu of options, with larger P_t values incurring higher prices.

At any given time, a plant may require more or less than its base power level. If the requested load for a plant is less than P_t, then its EMC does not need to engage in distributed scheduling with the other neighborhood EMCs since its load is "low." However, if the requested load of a plant exceeds P_t (i.e., its load is "high"), then its EMC must coordinate and compete with other plants so that they collectively meet requirement. In our proposed mechanism, in each scheduling slot, an EMC for a high-load plants will first use as much as possible of its base power level P_b and then compete with other EMCs to obtain additional power to meet the rest of the load. If it is unable to obtain enough power to meet the entire demand, the EMC will prioritize equipment based on their delay penalties ψ_n^1.

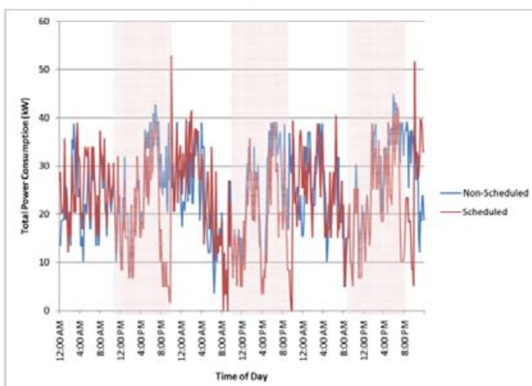

Fig. 5. Total power usage by period under distributed scheduling mechanism with $P_{max,t} = 40$ and $P_b = 4.0$

Under this mechanism, it is impossible to *guarantee* that all equipment is turned on with their delay d_n, since there is no guarantee that sufficient power is available at all times. Therefore, we introduce a second delay penalty ψ_n^2 that is incurred in each period of delay after d_n has elapsed.

That is, if equipment *n* is requested in period *t* and turned on in period $t + \Delta$, $\Delta > d_n$, then it incurs a delay cost of $\psi_n^1 d_n + \psi_n^2 (\Delta - d_n)$. The mechanism could also be constructed so that additional electricity may be purchased on demand, at a premium; if this premium is smaller than the additional delay cost, then the transaction is worthwhile.

C. Simulation

Using the same data as in the simulation in Section 2 B, we simulated the 50 oil gathering plants located in the same production site, now operating under the distributed scheduling mechanism described in Sections 3 A and 3 B. The ψ_n^2 values (not used in the previous simulation) are given in Table 1. In our simulation, we assume that the scheduling slots and the optimization time periods are equal (10 minutes) and coincide with each other. For the sake of simplicity, we also assume that *w =1*; that is, EMCs do not enter a random back off when their desired load cannot be met, but rather, they simply try again in the next scheduling slot. In addition, we do not simulate the transmission of probe packets, since the communication channel is resolved at a time scale that is much shorter than that of the scheduling slots/optimization periods. Fig. 5 plots energy usage over time assuming $P_{max,t} = 40$ for all t and $P_b = 4.0$. Note that the peak is greatly reduced and the load is significantly more level than in Fig. 3. We also simulated the system under a range of $P_{max,t}$ and P_b values. Fig. 6 plots the average total cost per plant for the scheduled system, over days 2–4 of the simulation, for $P_{max,t} \in\{20, 30,...,80, \infty\}$ and $P_b \in\{2, 4,...,10\}$. ($P_{max,t} = \infty$ represents the case in which the distributed scheduling mechanism is not used and all plants may use as much power as desired.) Note that the consumer cost increases only slightly or not at all as $P_{max,t}$ decreases from ∞. Therefore, significant reductions in peak demand are possible with minimal inconvenience/cost to the consumer. (Of course, in some cases, the peak demand exceeds $P_{max,t}$, as demonstrated in Fig. 5).

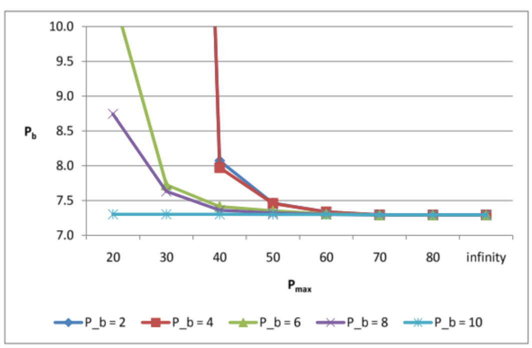

Fig. 6. Average total cost $P_b = f(P_{max,t})$ *for different values of power Pb*

As $P_{max,t}$ continues to increase, the cost begins to increase sharply as some equipment become delayed beyond d_n and therefore begin to incur the much higher delay cost of ψ_n^2. This occurs because the optimization model employed by the EMC implicitly assumes that the plant can use as much power as desired. If the EMC delays operation of an equipment, only to find that insufficient power is available when the equipment should be turned on, significant increases in cost may occur. To avoid this, the EMC optimization model should account for potential constraints on the available power. The algorithm described in the next section is designed to do this.

4. Dynamic Programming Algorithm

We now introduce a dynamic programming (DP) algorithm to optimize equipment operation times, subject to a capacity constraint on the total power available. This model assumes a single, fixed capacity level and ensures that the total scheduled power consumption never exceeds the capacity. Of course, under the distributed scheduling scheme from Section 3, the capacity is stochastic and changes over time. Our DP could be adapted to handle the stochastic capacity, but this would require knowing the probability distribution of the capacity in each period, and these distributions depend on the EMCs' decisions, which in turn depend on the distributions, and so on. The resulting equilibrium decisions (if any) would be optimal for the single plant scheduling problem, but determining them is significantly more difficult and is outside the scope of this paper. Of course, even though the EMC optimizes based on a fixed capacity level, it could adjust its decisions in real time as the stochastic capacity is realized.

We assume that the maximum available capacity in period *t* is given by b_t. For example, we might set $b_t = \alpha P_b$, where $\alpha \geq 1$. Our DP ensures that the load used by the plant in each period never exceeds b_t. It also ensures that each equipment is turned on before its maximum delay d_n has elapsed; therefore, the parameter ψ_n^2 is not relevant in this section. Let S_{tn} be the state vector of equipment *n* in period *t*. The possible values of S_{tn} are {-1, 0, 1,...,d}, where state -1 represents the equipment being on, state 0 represents it being off and not requested, and state i, $1 \leq i \leq d_n$, represents it having been requested for *i* periods. Let $S_t = (S_{tn})_{n=1}^{N}$ be the state vector for period *t*. Note that this state space may be quite large. For example, if there are *N* equipment and each has $d_n = d$, then there are $(d+2)^N$ possible state vectors S_t . Therefore, this DP approach is practical only for small-sized problems.

Let $f_t(S_t)$ be the optimal cost for periods *t*, *t +1,...,T* if the system begins period t in state S_t. In a given period, let O be the set of equipment that are currently on, let *R* be the set of equipment that have been requested but not yet turned on, and let $\overline{R} \subseteq R$ be the set of requested equipment that have reached the maximum delay d_n. We must choose which set $T \subseteq R$ of requested equipment to turn on. Let $f_{T+1}(S_t) \equiv 0$ for all S_t. Then $f_t(S_t)$ can be expressed recursively as

$$f_t(S_t) = \min_{T\subseteq R}\left\{\sum_{n\in O\cup T}\pi_t c_n + \sum_{n\in R\setminus T}\psi_n^1 + E[f_{t+1}(S_{t+1})] \mid \bar{R} \subseteq T, \sum_{n\in O\cup T} c_n \leq b_t\right\} \quad (2)$$

The minimization is taken over all possible subsets *T* of *R*. The first two terms inside the braces calculate the current period cost— the energy cost for equipment that already have been or are about to be turned on, and the delay cost for the equipment that have been requested but not yet turned on. The third term calculates the expected cost in all future periods; the expectation is taken over all possible states S_{t+1}, after accounting both for the decisions made in the current period (which equipment to turn on) and the random state transitions (equipment turning off or being requested) that occur at the beginning of the next period. The constraints on the last line ensure that all equipment that have reached their maximum delay are turned on and that the total load, among all equipment that are already on or are selected to be turned on, does not exceed the capacity.

This DP presents three main computational challenges, all due to the large state space: (1) the number of states S_t for which we must compute $f_t(S_t)$ is large; (2) the number of subsets $T\subseteq R$ in the minimization may be large, depending on $|R|$; and (3) the expectation $E_{S_{t+1}}$ [·] is difficult to calculate due to the large number of possible states to transition to. For small values of *N* and d_n, the DP may be solved exactly. However, for larger values, approximate dynamic programming (ADP) [10] approaches may be employed. For example, in our implementation we use sampling to calculate $E_{S_{t+1}}$ [·] and to determine which states S_t we must compute $f_t(\cdot)$ for. We also reduce the possible subsets *T* under consideration in the minimization by considering only $T = \bar{R}, T = \bar{R} \cup \{n\}$ for each $n \in R$ and *T* = *R*. Techniques such as these greatly improve the algorithm's execution speed, though of course the resulting algorithm no longer guarantees the optimal solution.

If the user does not specify a maximum allowable demand d_n, but only specifies the delay penalty ψ_n^1 then the DP can be simplified considerably. In this case, we can reduce the state space by collapsing the states {1, 2,...} into a single state and ignoring the constraint $R\subseteq T$.

5. Conclusion

In this paper, we propose a power scheduling protocol in a smart grid system, as well as two optimization methods for choosing the timing of equipment operation within an oil gathering plant in order to take advantage of lower off-peak energy prices. Our distributed scheduling mechanism guarantees plants a base power level while allowing them to compete for the remaining available power. Simulation results demonstrate that off-peak pricing models may exacerbate, rather than alleviate, the problem of demand peak, and that our distributed scheduling protocol can overcome this problem

REFERENCES

[1] *N. Hatziargyriou, G. Contaxis, M. Matos, J. A. P. Lopes, G. Kariniotakis, D. Mayer, J. Halliday, G. Dutton, P. Dokopoulos, A. Bakirtzis, J. Stefanakis, A. Gigantidou, P. O'Donnell, D. McCoy, M. J. Fernandes, J. M. S. Cotrim, and A. P. Figueira* "Energy management and control of island power systems with increased penetration from renewable sources" (IEEE-PES Winter Meeting, vol. 1, January. 2002, pp. 335-339).

[2] *M. Korpas and A. T. Holen* "Operation planning of hydrogen storage connected to wind power operating in a power market" (IEEE Transactions on Energy Conversion, vol. 21, no. 3, Sept. 2006, pp. 742-749).

[3] *S. Chakraborty and M. G. Simoes* "PV-smart grid operational cost minimization by neural forecasting and heuristic optimization" (Proc. IEEE Industry Applications Society Annual Meeting IAS '08, October. 2008, pp. 1-8).

[4] *E. Alvarez, A. C. Lopez, J. Gómez-Aleixandre, and N. de Abajo* "Online minimization of running costs, greenhouse gas emissions and the impact of distributed generation using smart grids on the electrical system" (Proc. IEEE PES/IAS Conference on Sustainable Alternative Energy (SAE), September. 2009, pp. 1-10).

[5] *David E. Olivares, Claudio A. Canizares, Mehard Kazerani* "A Centralized Optimal Energy Management System for Microgrods" (IEEE PES General Meeting, Minneapolis Minnesota, USA, 25-29 July 2010).

[6] *N. D. Hatziargyriou, A. Dimeas, A. G. Tsikalakis, J. A. P. Lopes, G. Karniotakis, and J. Oyarzabal*, "Management of microgrids in market environment" (Proc. International Conference on Future Power Systems, November. 2005, pp 7).

[7] *J. Oyarzabal, J. Jimeno, J. Ruela, A. Engler, and C. Hardt* "Agent based micro grid management system" (Proc. International Conference on Future Power Systems, November. 2005, pp 6).

[8] *M. LeMay, R. Nelli, G. Gross, and C.A. Gunter*. An integrated architecture for demand response communications and control. Proc. Of Hawaii International Conference on System Sciences, pages 174–183, 2008.

[9] *K. Belson*. Rewarding those who wait to flip the switch. *New York Times*, 2008. July 21

[10] *W. B. Powell*. Approximate Dynamic Programming: Solving the Curses of Dimensionality. Wiley-Interscience, New York, 2007.