Exploring the Tradeoff between Waiting Time and Service Cost in Non-Asymptotic Operating Regimes

Bin Li, Ozgur Dalkilic and Atilla Eryilmaz

***Abstract*—Motivated by the problem of demand management in smart grids, we study the problem of minimizing a weighted-sum of the mean delay of user demands and the power generation cost, where the latter metric increases with both the mean and the variance of the service demand. The state-of-the-art algorithms for this problem are asymptotically optimal, i.e., they are optimal only when the mean delay of user demands increases to infinity or decreases to zero. Yet, these algorithms may perform poorly for moderate delay, which is the regime in which most applications operate. Hence, there is a pressing need for the design of algorithms that can operate efficiently in the moderate delay regime. We attack this challenging problem in a generic framework by first proposing two classes of parameterized algorithms, which include some existing policies as special instances. Then, we obtain the optimal designs by explicitly characterizing the mean delay and the power generation cost as a function of the algorithmic parameters. The proposed algorithms with the optimal parameters not only are asymptotically optimal but also outperform the existing algorithms uniformly for all cases.**

I. Introduction

Many resource allocation problems can be solved by stochastic network optimization techniques to obtain *asymptotically* optimal algorithms. Yet, deriving optimal control policies in non-asymptotic regimes is quite challenging, since it typically requires solving the optimization problem with evolving system state information. The dynamic programming technique (e.g. [1]) usually leads to an optimal solution, which has high complexity and thus is not applicable in practice. In the last decade, the powerful Lyapunov drift minimization technique (e.g., [2], [3]) has provided systematic means to develop efficient solutions that only depend on the current system state. More excitingly, a class of dynamic control algorithms has been proposed with asymptotic optimality guarantees. For example, for energy minimization problem in time-varying wireless networks, the existing algorithms guarantee performance optimality in the large-delay (e.g., [4], [5]) or in the small-delay (see [6]) regimes. However, they may perform poorly for moderate delay, which is the regime in which most applications operate.

This motivates us to design efficient algorithms that perform well in both asymptotic and non-asymptotic regimes. To that end, we consider the interesting and challenging demand management problem arising in the smart electrical grid, model it as a stochastic queue scheduling problem, and develop efficient electricity procurement algorithms. We note that the results and insights obtained in this work can be applied to more general stochastic network optimization problems.

In practice, low variability in service and low delay are desirable characteristics for many systems such as computer networks, cloud computing infrastructures, and smart electrical grid. In this investigation, we focus on the demand management problem in the electrical grid. The demand management problem involves optimization of costs and utilities of participants in the smart grid when demands have flexibilities such as being deferrable in time. In particular, our goal is to utilize the flexibility of deferrable demands to reduce variability of the electrical load while keeping the consumer delay at moderate levels.

The demand management problem recently attracted considerable interest in the literature (e.g., [7], [8], [9], [10]). In [7], authors employ dynamic programming to solve demand management problem under renewable generation uncertainty, but the optimal scheduling algorithm has high complexity. In [8], [9], demand management problem is modeled as a queue scheduling problem and low complexity algorithms based on Lyapunov drift minimization techniques are developed. However, the algorithms result in highly and abruptly varying load scheduling decisions, and they also suffer from the delay-cost tradeoff under non-asymptotic regimes. On the other hand, the authors of [10] propose a real-time algorithm that decreases the variability of the grid load, nevertheless delays experienced in demand can be quite large under the proposed algorithm.

In contrast to the previous works, we consider the mean delay experienced by consumers in the non-asymptotic regime and design algorithms that minimize both the cost and the variability of the electricity procurement. We propose asymptotically optimal load scheduling policies that outperform several key algorithms in the non-asymptotic regime. The following items list our main contributions along with the outline of the paper:

- In Section II, we introduce the system model for the demand management problem that accounts for delays experienced by users and the power generation cost in both long-term and short-term timescales.
- In Section III, we first review the existing solutions to our problem. Then, we propose two classes of efficient policies and obtain the optimal policies within these two classes. These optimal policies are shown to be asymptotically optimal.
- In Section IV, we perform simulations to show that our proposed algorithms outperform the existing algorithms uniformly for all cases.

Bin Li, Ozgur Dalkilic and Atilla Eryilmaz ({li.1318, dalkilic.1, eryilmaz.2}@osu.edu) are with the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, Ohio 43210 USA.

This work is supported by the NSF grants: CAREER-CNS-0953515 and CCF-0916664. Also, the work of A. Eryilmaz was in part supported by the QNRF grant number NPRP 09-1168-2-455.

II. System Model

We consider an optimal demand management problem in the smart grid with N customers with deferrable demands and a utility serving these demands as shown in Fig. 1. We assume

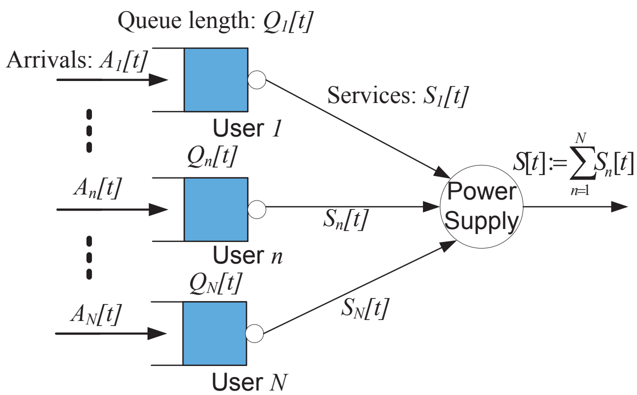

Fig. 1. Demand management in smart grids

that each customer maintains a queue storing his/her demands for using electricity. Let $Q_n[t]$ denote the queue length of the n^{th} customer's demands in time slot t. We use $A_n[t]$ and $S_n[t]$ to denote the amount of demand and service of the customer n in slot t, respectively. Thus, the evolution of queue length $Q_n[t]$ is described as follows:

$$Q_n[t+1] = \max\{0, Q_n[t] + A_n[t] - S_n[t]\}, \forall n = 1, 2, ..., N.$$

We say queue n is stable if

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[Q_n[t]] < \infty.$$

We say that the system is stable if all queues are stable.

Let $S[t] \triangleq \sum_{n=1}^{N} S_n[t]$ be the total amount of service that is provided in slot t. On one hand, each customer prefers to get service as soon as possible. Hence, the larger the queue length is, the larger the loss the customer suffers from. To model this effect, we use $Q_n[t]$ to denote the cost associated with the n^{th} customer. On the other hand, power generation is subject to costs in two timescales: long-term cost of sustaining a constant power and short-term ramp up/down cost of changing the power generation level. As such, we consider a cost function $C(S)$ that is a convex and increasing function of S. The design objective can be expressed as

$$\min_{\pi\in\Pi} \limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[\frac{1}{K}\sum_{n=1}^{N} Q_n[t] + C(S[t])\right],$$

where K is a positive constant characterizing the relative importance of the power generation cost over the cost associated with customers, and Π denotes the class of stationary and stable policies that govern the selection of customers to serve and the power generation rate, and stabilize the system. Stability condition implies that for each policy $\pi \in \Pi$, there exist random variables $\overline{Q}_n^{(\pi)}$, $\overline{S}_n^{(\pi)}$ and $\overline{S}^{(\pi)}$ such that $Q_n[t] \Rightarrow \overline{Q}_n^{(\pi)}$, $S_n[t] \Rightarrow \overline{S}_n^{(\pi)}$, $S[t] \Rightarrow \overline{S}^{(\pi)}$, where "$\Rightarrow$" means convergence in distribution. We assume that $\mathbb{E}[C(\overline{S}^{(\pi)})] < \infty$ and $\mathbb{E}[\overline{Q}_n^{(\pi)}] < \infty$, $\forall n$, under each policy $\pi \in \Pi$. Thus, our goal becomes

$$\min_{\pi\in\Pi} \mathbb{E}\left[\frac{1}{K}\sum_{n=1}^{N} \overline{Q}_n^{(\pi)} + C(\overline{S}^{(\pi)})\right]. \quad (1)$$

This problem is notoriously difficult, since the objective involves the system state information. In the rest of the paper, we consider a single-queue system that maintains the key nature of the general objective and leads to a more tractable analysis. We assume that $\{A[t]\}_{t\geq 0}$ are independently and identically distributed over time with mean λ and variance σ^2. Also, we set $C(x) = x^2$, where $(\overline{S}^{(\pi)})^2$ conveniently characterizes both average service rate and service variation, since $\mathbb{E}[(\overline{S}^{(\pi)})^2] = (\mathbb{E}[\overline{S}^{(\pi)}])^2 + \text{Var}(\overline{S}^{(\pi)})$. Thus, in the rest of the paper, we consider the following optimization problem:

$$\min_{\pi\in\Pi} \quad \tfrac{1}{K}\mathbb{E}[\overline{Q}^{(\pi)}] + \mathbb{E}[(\overline{S}^{(\pi)})^2]. \quad (2)$$

III. Efficient Algorithm Design and Analysis

In this section, we first review the state-of-the-art algorithms that target to solve the optimization problem (2) and discuss their deficiencies in various cases. This, then, motivates us to develop efficient algorithms that not only achieve the desirable properties of the existing schemes but also outperform them uniformly for all cases.

A. Benchmark Algorithms

Several algorithms that are based on heuristic and existing solutions have been proposed to solve the difficult problem (2). Next, we present the following key algorithms and discuss their performance.

(i) **"Strawman" algorithm**: Let $S^{(\text{strawman})}[t] = A[t], \forall t \geq 1$. Then $Q^{(\text{strawman})}[t] \equiv 0, \forall t \geq 1$. Hence,

$$\frac{1}{K}\mathbb{E}\left[\overline{Q}^{(\text{strawman})}\right] + \mathbb{E}\left[\left(\overline{S}^{(\text{strawman})}\right)^2\right] = \mathbb{E}[A^2[t]] = \lambda^2 + \sigma^2.$$

Note that if the arrival has a high variance, then the "strawman" algorithm yields a cost far from the optimal value and thus performs poorly.

(ii) **Energy Efficient Control Algorithm (EECA)** (see [2]): Let $S^{(\text{EECA})}[t] = \frac{Q[t]}{2K}$, $\forall t \geq 1$. This algorithm is developed through the Lyapunov drift minimization technique (e.g., [2], [3]). It is shown that this algorithm can achieve the average generation cost within $O(\frac{1}{K})$ of the minimum generation cost required for network stability at the expense of network delay growing with $O(K)$. In [4], the authors show that the network delay should at least grow with $O(\sqrt{K})$ when the average generation cost is within $O(\frac{1}{K})$ of the minimum generation cost for maintaining network stability. This indicates that the EECA algorithm may not be optimal in the asymptotic regime as K goes to infinity.

(3) **Berry-Gallager (BG) algorithm** (see [4], [5]): Let service rate $S^{(\text{BG})}[t]$ be $\lambda + \frac{1}{K}$ if $Q[t] > K$ and be $\max\{\lambda - \frac{1}{K}, 0\}$,

otherwise. Note that the BG algorithm is shown to be asymptotically optimal. That is, it is optimal when K goes to infinity, i.e., when the network delay goes to infinity. However, for each $K > 0$, especially for small K, $S^{(\mathrm{BG})}[t]$ may change abruptly between consecutive time slots. Intuitively, it leads to large variation of service and thus the BG algorithm may not be a good choice for moderate K values. Based on this intuition, we propose the following efficient algorithms, in which the service changes smoothly.

B. Proposed Algorithms

In this subsection, we first propose the classes of Queue-Scaled Service (QSS) algorithms with parameter $0 < \theta \leq 1$ and Queue and Arrival Scaled Service (QASS) algorithms with parameters $0 \leq \alpha \leq 1$ and $0 < \beta \leq 1$, and then obtain their corresponding optimal designs.

QSS algorithm with parameter θ: Let $S^{(\mathrm{QSS})} = \theta Q[t]$, where $0 < \theta \leq 1$. When $0 < \frac{1}{2K} \leq 1$, the EECA algorithm can be seen as a special case of the QSS algorithm with $\theta = \frac{1}{2K}$.

It is easy to show that the QSS algorithm stabilizes the system for any arrival process by studying the drift of the linear Lyapunov function. Hence, there exist random variables $\overline{Q}^{(\mathrm{QSS})}$ and $\overline{S}^{(\mathrm{QSS})}$ such that $Q^{(\mathrm{QSS})}[t] \Rightarrow \overline{Q}^{(\mathrm{QSS})}$ and $S^{(\mathrm{QSS})}[t] \Rightarrow \overline{S}^{(\mathrm{QSS})}$.

Next, we will obtain the optimal QSS algorithm. To achieve it, we perform a Lyapunov-drift-based analysis of the steady-state behavior of the stochastic network (see [11]). By choosing linear and quadratic Lyapunov functions and setting their drifts to zero respectively, we obtain

$$\mathbb{E}\left[\overline{Q}^{(\mathrm{QSS})}\right] = \frac{\lambda}{\theta}, \tag{3}$$

$$\mathbb{E}\left[\left(\overline{S}^{(\mathrm{QSS})}\right)^2\right] = \lambda^2 + \frac{\theta}{2-\theta}\sigma^2. \tag{4}$$

By substituting (3) and (4) into (2), our objective becomes

$$\min_{0<\theta\leq 1} \quad f_{\mathrm{QSS}}(\theta) \triangleq \lambda^2 + \frac{1}{K}\frac{\lambda}{\theta} + \frac{\theta}{2-\theta}\sigma^2. \tag{5}$$

Note that this function is convex. Also, consider

$$\begin{aligned} f'_{\mathrm{QSS}}(\theta) =& \frac{1}{K}\frac{-\lambda}{\theta^2} + \frac{2}{(\theta-2)^2}\sigma^2 \\ =& \frac{2\sigma^2}{\theta^2}\left(\frac{\theta}{2-\theta} + \frac{1}{\sqrt{2\gamma}}\right)\left(\frac{\theta}{2-\theta} - \frac{1}{\sqrt{2\gamma}}\right), \end{aligned} \tag{6}$$

where $\gamma \triangleq \frac{K\sigma^2}{\lambda}$ is proportional to weighting parameter K and normalized variance $\frac{\sigma^2}{\lambda}$.

It is easy to check that $f_{\mathrm{QSS}}(\theta)$ is decreasing in the interval $\left(0, \min\left\{1, \frac{2}{1+\sqrt{2\gamma}}\right\}\right)$ and increasing in the interval $\left[\min\left\{1, \frac{2}{1+\sqrt{2\gamma}}\right\}, 1\right]$. Hence, the optimal solution θ^* to the problem (5) is

(i) If $\gamma \leq \frac{1}{2}$, then

$$\theta^* = 1, \tag{7}$$

$$f^*_{\mathrm{QSS}} = \lambda^2 + \sigma^2\left(\frac{1}{\gamma} + 1\right); \tag{8}$$

(ii) If $\gamma > \frac{1}{2}$, then

$$\theta^* = \frac{2}{1+\sqrt{2\gamma}}, \tag{9}$$

$$f^*_{\mathrm{QSS}} = \lambda^2 + \sigma^2\left(\frac{1}{2\gamma} + \sqrt{\frac{2}{\gamma}}\right). \tag{10}$$

The optimal choice of θ^* as a function of γ is depicted in Fig. 2, which shows that as γ scales, θ^* decays at the rate of $\Theta(\frac{1}{\sqrt{K}})$ towards zero. This indicates that the scaling of $\Theta(\frac{1}{K})$ under the EECA algorithm may not be the best scaling for non-asymptotic regimes (also see Fig. 3).

Also, note that as K goes to infinity, f^*_{QSS} converges to λ^2. For any algorithm $\pi \in \Pi$, we have

$$\begin{aligned} &\frac{1}{K}\mathbb{E}[\overline{Q}^{(\pi)}] + \mathbb{E}[(\overline{S}^{(\pi)})^2] \\ \geq& \mathbb{E}[(\overline{S}^{(\pi)})^2] \overset{(a)}{\geq} \left(\mathbb{E}[\overline{S}^{(\pi)}]\right)^2 \overset{(b)}{\geq} \lambda^2, \end{aligned} \tag{11}$$

where the step (a) follows from Jensen's inequality, and step (b) uses the fact that $\mathbb{E}[\overline{S}^{(\pi)}] \geq \lambda$, since the algorithm π stabilizes the system. Hence, the optimal QSS algorithm is asymptotically optimal.

Yet, for a sufficiently small γ, the optimal QSS algorithm performs worse than the strawman algorithm, which motivates us to incorporate both the queue-length and arrival information into the algorithm design.

QASS algorithm with parameters α **and** β: Let $S^{(\mathrm{QASS})}[t] = \alpha A[t] + \beta Q[t]$, where $0 \leq \alpha \leq 1, 0 < \beta \leq 1$. When $\alpha = 1, \beta = 0$, the QASS algorithm reduces to the strawman algorithm. The QSS algorithm can be seen as a special case of the QASS algorithm with $\alpha = 0$ and $\beta = \theta$.

Similar to the QSS algorithm, we can show that there exist random variables $\overline{Q}^{(\mathrm{QASS})}$ and $\overline{S}^{(\mathrm{QASS})}$ such that $Q^{(\mathrm{QASS})}[t] \Rightarrow \overline{Q}^{(\mathrm{QASS})}$ and $S^{(\mathrm{QASS})}[t] \Rightarrow \overline{S}^{(\mathrm{QASS})}$.

Next, we will obtain the optimal QASS policy. By choosing linear and quadratic Lyapunov functions and setting their drifts to zero respectively, we have

$$\mathbb{E}\left[\overline{Q}^{(\mathrm{QASS})}\right] = \frac{(1-\alpha)\lambda}{\beta}, \tag{12}$$

$$\mathbb{E}\left[\left(\overline{S}^{(\mathrm{QASS})}\right)^2\right] = \lambda^2 + \left(\alpha^2 + (1-\alpha)^2\frac{\beta}{2-\beta}\right)\sigma^2. \tag{13}$$

By substituting (12) and (13) into (2), the objective becomes

$$f_{\mathrm{QASS}}(\alpha,\beta) \triangleq \lambda^2 + \frac{1}{K}\frac{1-\alpha}{\beta}\lambda + \left(\alpha^2 + (1-\alpha)^2\frac{\beta}{2-\beta}\right)\sigma^2.$$

Our goal is

$$\min_{\alpha,\beta} \quad f_{\text{QASS}}(\alpha,\beta) \tag{14}$$

$$\text{subject to} \quad 0 \le \alpha \le 1, \tag{15}$$

$$0 < \beta \le 1. \tag{16}$$

We use α^* and β^* to denote the optimal solution to problem (14). Let $f^*_{\text{QASS}} \triangleq f_{\text{QASS}}(\alpha^*,\beta^*)$.

Note that this function is not jointly convex in both α and β. Hence, we solve problem (14) by minimizing over α and β sequentially. The detailed derivation of the optimal solution is available in Appendix A. It turns out that the optimal solution to the problem (14) also depends on the key parameter γ:

(i) If $\gamma \le \frac{1}{2}$, then

$$\alpha^* = 1, \tag{17}$$

$$f^*_{\text{QASS}} = \lambda^2 + \sigma^2.$$

If the arrival has a small variance, the optimal QASS algorithm reduces to the "strawman" policy. In this case, the generated cost is small.

(ii) If $\frac{1}{2} \le \gamma \le \frac{3}{2}$, then

$$\alpha^* = \frac{1}{2} + \frac{1}{4\gamma}, \tag{18}$$

$$\beta^* = 1, \tag{19}$$

$$f^*_{\text{QASS}} = \lambda^2 + \sigma^2 \left(-\frac{1}{8}(\frac{1}{\gamma} - 2)^2 + 1 \right).$$

If the arrival has a medium variance, the optimal QASS algorithm becomes $S^{(\text{QASS})}[t] = \alpha^* A[t] + Q[t]$.

(iii) If $\gamma > \frac{3}{2}$, then

$$\beta^* = \frac{8}{1+\sqrt{1+32\gamma}}, \tag{20}$$

$$\alpha^* = \frac{\beta^*}{2} + \frac{2-\beta^*}{4\beta^*\gamma}, \tag{21}$$

$$f^*_{\text{QASS}} = \lambda^2 + \sigma^2 \left(\frac{1}{128\gamma^2} - \frac{5}{8\gamma} + \frac{1}{4}\sqrt{\frac{1}{\gamma^2} + \frac{32}{\gamma}} + \frac{1}{128\gamma}\sqrt{\frac{1}{\gamma^2} + \frac{32}{\gamma}} \right).$$

If the arrival has a large variance, the optimal QASS algorithm becomes $S^{(\text{QASS})}[t] = \alpha^* A[t] + \beta^* Q[t]$.

Note that f^*_{QASS} converges to λ^2 as K goes to infinity. Combining this with the lower bound (11), we see that the optimal QASS algorithm is also asymptotically optimal.

For sufficiently large K, both θ^* and β^* scale with $\frac{1}{\sqrt{K}}$ (see Fig. 2) while the EECA algorithm scales with $\frac{1}{K}$. This is the reason why both the QSS and QASS algorithms outperforms the EECA algorithm.

IV. Simulations

In order to show the promise of the proposed QSS and QASS policies, we compare their performances to those of the Energy Efficient Control Algorithm (EECA) in [2] and Berry-Gallager (BG) policy in [4] via simulations. The Tradeoff Optimal Control Algorithm (TOCA) algorithm in [5] has almost the same performance as the BG policy in a single-user case, and thus is omitted for brevity. In particular, we numerically study the performance of the mentioned policies w.r.t. the design parameter K and the system parameter $\frac{\sigma^2}{\lambda}$ which together constitute the γ parameter in QSS and QASS policies.

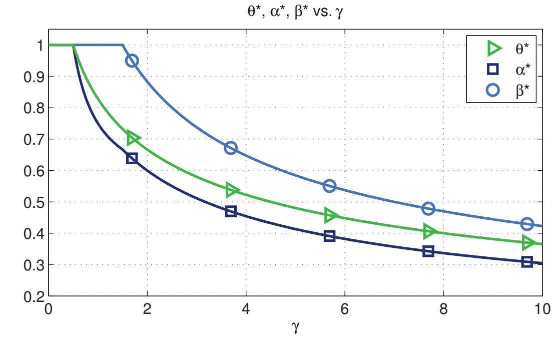

Fig. 2. The behavior of the optimal QSS and QASS policies

We demonstrate the effect of K on the cost achieved in Fig. 3. In the corresponding simulations, the arrival process is Poisson distributed with $\lambda = 10$, hence $\frac{\sigma^2}{\lambda} = 1$ and $\gamma = K$. The minimum achievable cost is $\lambda^2 = 100$, which is a trivial lower bound for problem (2). In Fig. 3, we observe that QASS policy performs best among all policies in all regimes, including the moderate delay regime, and it asymptotically achieves the optimal objective value. Although the BG policy also achieves optimal cost asymptotically, it is outperformed by QASS in the whole delay range and by QSS in the larger delay regime. As expected, QSS performs worse than the strawman policy when γ is small.

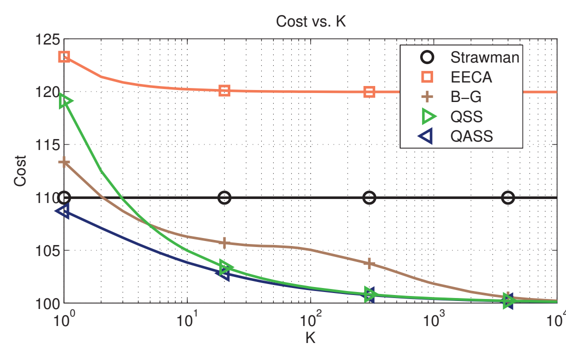

Fig. 3. Cost comparison of the proposed and the existing policies w.r.t. K. Arrivals are Poisson distributed, hence $\frac{\sigma^2}{\lambda} = 1$.

The impact of the variability of the arrival process is demonstrated in Fig. 4. The arrivals follow Gaussian distribution with values projected onto the non-negative real-line. The mean is set to be fixed at $\lambda = 10$, and standard deviation is varied between the simulations to obtain a range of different $\frac{\sigma^2}{\lambda}$ values. In Fig. 4, we observe that both QASS and QSS policies are fairly robust against the variations in $\frac{\sigma^2}{\lambda}$ ratio, while the cost values achieved by EECA and BG policies increase with

$\frac{\sigma^2}{\lambda}$. Furthermore, QASS performs consistently better than the other policies over the whole range of $\frac{\sigma^2}{\lambda}$ values. Note that, even though the cost of the Strawman policy is $\lambda^2+\sigma^2$ and decreases down to the lower bound λ^2 with decreasing $\frac{\sigma^2}{\lambda}$, QASS outperforms the Strawman policy since it is a special case of our QASS policy with $\beta = 0$.

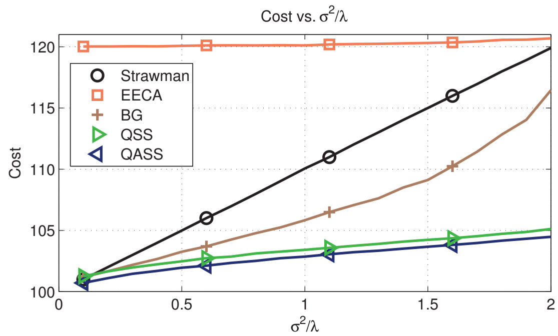

Fig. 4. Cost comparison of the proposed and the existing policies w.r.t. $\frac{\sigma^2}{\lambda}$. $K = 20$ and arrivals are Gaussian distributed with mean 10.

The above comparisons suggest gains in the mean cost of service achieved by our policies as measured in (2). Another metric of interest for service cost is *volatility* that measures temporal fluctuations of service, which will be part of our future investigations.

V. Conclusion

In this paper, we explore the tradeoff between the mean delay and the average service cost. Early works focus on the efficient algorithm design in the asymptotic regimes, i.e., when the delay goes to infinity or zero in the energy minimization problem in time-varying networks. These algorithms may not perform well in most applications that require moderate or low delay, hence this motivates us to design efficient algorithms for the non-asymptotic regimes. To that end, we propose two classes of parameterized algorithms based on queue length and arrival information, and obtain the optimal algorithms within these two classes. Both of the resulting algorithms with the optimized parameters not only are asymptotically optimal but also exhibit improved performance over the traditional schemes in all cases. For a quadratic service cost function, our investigations reveal the optimality of an alternative scaling of $\Theta(\frac{1}{\sqrt{K}})$ for the design parameter as a function of the delay weight K instead of the existing scaling of $\Theta(\frac{1}{K})$. This suggests the significance of refined parameter selection to better manage the non-asymptotic operation. In the future, we will seek a tighter lower bound and determine whether the proposed algorithms have a constant efficient ratio independent of the arrival statistics. Also, we will extend our results to the multi-user case.

References

[1] D. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice Hall, Englewood Cliffs, NJ, 1987.

[2] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, 52(7):2915–2934, 2006.

[3] M. J. Neely. *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.

[4] R. Berry and R. Gallager. Communication over fading channels with delay constraints. *IEEE Transactions on Information Theory*, 48(5):1135–1149, 2002.

[5] M. J. Neely. Optimal energy and delay tradeoffs for multiuser wireless downlinks. *IEEE Transactions on Information Theory*, 53(9):3095–3113, 2007.

[6] R. Berry. Optimal power-delay trade-offs in fading channels - small delay asymptotics. *IEEE Transactions on Information Theory*, 52(7):3939–3952, 2013.

[7] A. Papavasiliou and S. S Oren. Supplying renewable energy to deferrable loads: Algorithms and economic analysis. In *Proc. IEEE Power and Energy Society General Meeting*, July 2010.

[8] S. Chen, P. Sinha, and N. B Shroff. Scheduling heterogeneous delay tolerant tasks in smart grid with renewable energy. In *Proc. IEEE Conference on Decision and Control (CDC)*, Maui, Hawaii, December 2012.

[9] M. J Neely, A. S Tehrani, and A. G Dimakis. Efficient algorithms for renewable energy allocation to delay tolerant consumers. In *Proc. IEEE Conference on Smart Grid Communications (SmartGridComm)*, Gaithersburg, Maryland, October 2010.

[10] O. Dalkilic, A. Eryilmaz, and X. Lin. Stable real-time pricing and scheduling for serving opportunistic users with deferrable loads. In *Proc. Allerton Conference on Communication, Control, and Computing (Allerton)*, Monticello, Illinois, October 2013.

[11] A. Eryilmaz and R. Srikant. Asymptotically tight steady-state queue-length bounds implied by drift conditions. *Queueing Systems Theory and Applications (QUESTA)*, 72(3-4):311–359, 2012.

Appendix A
Derivation of Optimal Solution to Problem (14)

$$h(\beta) = \min_{0\le\alpha\le1} f_{\text{QASS}}(\alpha,\beta) = \begin{cases} \lambda^2+\sigma^2 & \text{if } \beta \le \frac{1}{2\gamma}; \\ \lambda^2+\sigma^2\left(\frac{1}{\gamma}\frac{2-\beta}{2\beta} - \frac{1}{\gamma^2}\frac{2-\beta}{8\beta^2} + \frac{\beta}{2}\right) & \text{if } \beta > \frac{1}{2\gamma}. \end{cases}$$

$$\alpha^* = \operatorname*{argmin}_{0\le\alpha\le1} f_{\text{QASS}}(\alpha,\beta) = \begin{cases} 1 & \text{if } \beta \le \frac{1}{2\gamma}; \\ \frac{\beta}{2} + \frac{1}{\gamma}\frac{2-\beta}{4\beta} & \text{if } \beta > \frac{1}{2\gamma}. \end{cases}$$

$$h'(\beta) = \begin{cases} 0 & \text{if } \beta \le \frac{1}{2\gamma}; \\ \sigma^2\left(\frac{1}{\gamma^2}\frac{1}{2\beta^3} - \left(\frac{1}{\gamma}+\frac{1}{8\gamma^2}\right)\frac{1}{\beta^2} + \frac{1}{2}\right) & \text{if } \beta > \frac{1}{2\gamma}, \end{cases} = \begin{cases} 0 & \text{if } \beta \le \frac{1}{2\gamma}; \\ \frac{\sigma^2}{2\beta^3}\left(\beta - \frac{1}{2\gamma}\right)(\beta-\beta_1)(\beta-\beta_2) & \text{if } \beta > \frac{1}{2\gamma}, \end{cases}$$

where $\beta_1 \triangleq \sqrt{\frac{1}{16\gamma^2}+\frac{2}{\gamma}} - \frac{1}{4\gamma}$ and $\beta_2 \triangleq -\sqrt{\frac{1}{16\gamma^2}+\frac{2}{\gamma}} - \frac{1}{4\gamma}$. Note that $\beta_1 > \frac{1}{2\gamma}$ when $1 \ge \beta > \frac{1}{2\gamma}$, so $h(\beta)$ is decreasing in the interval $(0, \min\{\beta_1, 1\})$ and increasing in the interval $[\min\{\beta_1,1\},1]$. Hence, if $\beta_1 < 1$, i.e., $\gamma > \frac{3}{2}$, then

$$\beta^* = \beta_1 = \sqrt{\frac{1}{16\gamma^2}+\frac{2}{\gamma}} - \frac{1}{4\gamma} = \frac{8}{1+\sqrt{1+32\gamma}}, \tag{22}$$

$$\alpha^* = \frac{\beta^*}{2} + \frac{1}{\gamma}\frac{2-\beta^*}{4\beta^*}. \tag{23}$$

If $\beta_1 \ge 1$, i.e., $\gamma < \frac{3}{2}$, then, $\beta^* = 1$. In this case,

$$\alpha^* = \operatorname*{argmin}_{0\le\alpha\le1} f_{\text{QASS}}(\alpha,\beta) = \begin{cases} 1 & \text{if } \gamma \le \frac{1}{2}; \\ \frac{1}{2} + \frac{1}{4\gamma} & \text{if } \gamma > \frac{1}{2}. \end{cases} \tag{24}$$