A Water-Filling Based Scheduling Algorithm for the Smart Grid

Merwais Shinwari, Amr Youssef, *Senior Member, IEEE*, and Walaa Hamouda, *Senior Member, IEEE*

***Abstract*—The processing and communication capabilities of the smart grid provide a solid foundation for enhancing its efficiency and reliability. These capabilities allow utility companies to adjust their offerings in a way that encourages consumers to reduce their peak hour consumption, resulting in a more efficient system. In this paper, we propose a method for scheduling a community's power consumption such that it becomes almost flat. Our methodology utilizes distributed schedulers that allocate time slots to soft loads probabilistically based on precalculated and predistributed demand forecast information. This approach requires no communication or coordination between scheduling nodes. Furthermore, the computation performed at each scheduling node is minimal. Obtaining a relatively constant consumption makes it possible to have a relatively constant billing rate and eliminates operational inefficiencies. We also analyze the fairness of our proposed approach, the effect of the possible errors in the demand forecast, and the participation incentives for consumers.**

***Index Terms*—Load balancing, probabilistic scheduling, smart grid, water-filling.**

I. Introduction

THE architecture of the power grid in use today follows the design proposed by Nikola Tesla over a century ago. At that time, electricity was considered a luxury that was primarily used for lighting. Today, the power grid is a critical infrastructure upon which most, if not all, other systems rely.

Historically, the electric grid was developed to deliver the power generated at remote power plants to the consumers. At that time, achieving this objective alone was a great success. Issues such as efficiency, greenness, and reliability were not a priority and addressing them was not feasible.

On the other hand, today's technologies allow much more. The smart grid, or the smarter grid, is about improving the efficiency of the existing power grid by employing state-of-the-art technologies [1]. Step by step, the grid will be enhanced to become an interconnected system of systems. This will allow real-time control of resources, better management and consequently, a more efficient system. The smart grid will enable participating parties to operate more efficiently which, in turn, will contribute to reducing both the financial and environmental costs of power generation and consumption [2], [3].

Manuscript received February 20, 2011; revised June 06, 2011; accepted November 07, 2011. Date of publication February 03, 2012; date of current version May 21, 2012. Paper no. TSG-00052-2011.

M. Shinwari and A. Youssef are with the Concordia Institute for Information Systems Engineering, Concordia University, Montreal, QC H3G 1M8, Canada (e-mail: m_shinwa@ciise.concordia.ca; youssef@ciise.concordia.ca).

W. Hamouda is with the Department of Electrical and Computer Engineering, Concordia University, Montreal, QC, H3G 1M8, Canada (e-mail: hamouda@ece.concordia.ca).

Digital Object Identifier 10.1109/TSG.2011.2177103

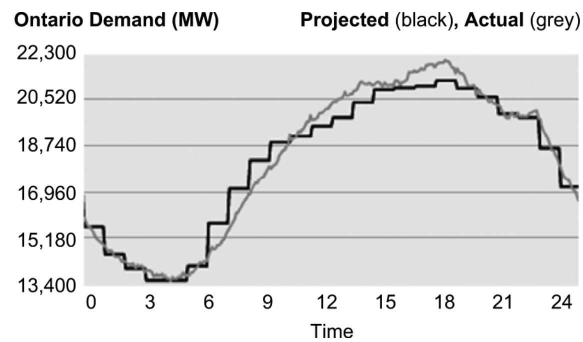

Fig. 1. Power demand during a typical summer weekday in Ontario, Canada [5].

The power grid can be made smarter in many ways, namely improving its efficiency and reliability, supporting distributed generation and storage, and facilitating consumption management for consumers [4]. The move towards a smarter power system is driven by many factors such as the high fuel prices, increasing energy demands, increasing environmental concerns, and the emergence of electric vehicles.

One area of inefficiency arises from the trend of the daily power consumption of a community. Fig. 1 shows the power consumption for the province of Ontario, Canada, on a typical summer weekday. As shown in the figure, the consumption ranges between 13 500 MW at about 4 A.M. and 21 000 MW at about 6 P.M. with the difference being about 65%. The consumption is lowest in the early hours of the morning, as most people are asleep. Then, it increases throughout the day peaking at about 6 P.M.. This is when people return home from work and add their home loads to the grid [5]. Furthermore, the figure shows that it is possible to forecast the power consumption with fair accuracy. In the figure, the black curve shows the forecasted consumption and the grey curve shows the actual (measured) consumption. In our work, we utilize the fact that it is possible to forecast the overall power consumption for a given community in order to develop a scheduler that doesn't require real-time data communication between the consumers' nodes.

With the current architecture of the grid, electric power must be consumed the instant it is generated. Because of this, power plants must have the capacity to support peak demand hours. This additional capacity becomes idle during low demand hours which results in operational inefficiency. Furthermore, at peak hours, high demand causes the utility companies to purchase power at higher rates. On the other hand, power is typically sold to consumers at a fixed rate. With such a setup, the time at which power is consumed is of no significance to the consumer. Therefore, consumers have no incentive to adjust their power consumption for a better bulk purchase rate.

Utilizing the capabilities of a smarter grid, utility companies are moving towards billing for power on a "time of use" basis. For example, *Hydro One*, which delivers electricity across the Canadian province of Ontario, is introducing a pricing scheme where each day is divided into on-peak hours, mid-peak hours and off-peak hours. In this scheme, power rates are most expensive during peak hours and less expensive during off-peak hours [5].

In this paper, we propose enhancing the power grid by enabling consumers to benefit from the capabilities of the smart grid. Specifically, we propose a method that allows consumers to shift part of their load to off-peak hours in a way that results in a relatively constant overall power demand profile. This results in a more efficient system, and a seemingly uniform pricing scheme.

The rest of this paper is organized as follows. In the next section, we briefly review some related works. In Section III, we discuss the idea of using instantaneous power rates, rather than fixed rates. In Section IV, we present an overview of water-filling, a concept on which our scheduler is based. Next, in Section V, we illustrate the logic of our scheduler. The proposed algorithm is presented in Section VI, and the supporting simulation results are presented in Section VII. In Section VIII, we analyze the fairness of our proposed algorithm, the effect of the possible error in demand forecast, and the participation incentives for consumers. Finally, we present our conclusion and recommendations for future work in Section IX.

II. Related Work

Caron and Kesidis [6] proposed a dynamic pricing scheme that encourages consumers to adjust their power use with the objective of getting a flat overall consumption. The authors showed that finding an optimum schedule is NP-hard, and then presented methodologies to study how close one can get to the ideal case. This was done based on the amount of information the consumers are willing to share with their utility company. Furthermore, the authors studied the outcomes based on several scheduling policies, namely, uniform, ALOHA I, ALOHA II, and Time/Slackness. In their work, the authors have also compared the performance of these scheduling policies.

Xiong *et al.* [14] proposed an approach based on communication protocols that reduce the power demand in an attempt to produce a uniform power consumption over time. Similar to other approaches, the authors divided the power consumption into real-time loads (non-schedulable) and schedulable loads. They also defined a "target" power level and modeled their algorithm to schedule the power use such that this target is not exceeded. The algorithm uses a specific structure that consists of three main phases, namely power update phase, power request phase, and power scheduling phase. Furthermore, the authors simulated their approach and showed that it is possible for a consumer to keep the power demand below the defined target.

Gatsis and Giannakis [15] presented a cooperative scheduling approach between the utility company and the consumers. In this approach, loads are classified into loads that must run, loads that must consume a given amount of power (e.g., a rechargeable battery), and loads that are adjustable in power consumption but the adjustment could cause consumer dissatisfaction (e.g., climate control). This was modeled into a convex optimization problem that was solved using the distributed subgradient method. The authors also presented simulation results that show that it is possible to meet the constraints above in an optimum way.

Chen *et al.* [7] proposed the use of a real-time pricing rate, and formed a Stackelberg game [12] between the utility company and energy management controllers that are to be deployed at each home. The game was setup such that the controllers play the role of the follower and the utility company plays the role of the leader. The authors simulated their proposed methodology and concluded that their approach saves money for the consumers and ensures that rebound peaks do not appear.

Mohsenian-Rad *et al.* [8] presented an autonomous incentive based algorithm for scheduling power consumption. In this scheme, loads are classified into soft (or schedulable) consumption and hard (or nonschedulable) consumption. Soft consumption represents usages that do not have strict time constraints, and hard consumption represents usages that have strict time constraints. The authors also proposed the use of energy consumption scheduling devices (ECS) as a component of the smart meters. In this model, an ECS communicates with other ECSs in its neighborhood sharing its scheduling information. Running their proposed distributed algorithm, each ECS computes and broadcasts its optimal schedule. The algorithm repeats until no ECS announces any change of schedule.

The approaches presented above do flatten the overall demand of a community. However, the need to continuously update other nodes with scheduling information and to solve optimization problems poses a great overhead in communications and processing. Furthermore, sharing detailed scheduling information with other consumers presents a major privacy concern [19]. In what follows, we propose a simple heuristic scheduling method that eliminates these requirements. In our proposed method, consumer nodes do not need to share detailed information with each other and they also do not perform any complex processing. Rather, we utilize statistical information on consumption trends that the utility companies could easily make available to their subscribers.

III. Instantaneous Power Rate

If consumers are billed at a constant rate, they will have no incentive to consume power with any pattern. Let $D(t)$ denote the power demand of a community throughout a time period T. For our purpose, we define T to be the duration of one day, i.e., 24 hours. With power consumed at will, $D(t)$ gets a shape as shown in Fig. 1. Let $Pr(t)$ denote the corresponding price rate as a function of time. Furthermore, let $Pr(t)$ be a function of the instantaneous demand. Therefore, we can define $Pr(t)$ as

$$Pr(t) = F(D(t)) \tag{1}$$

where F is a function that relates price to demand.

With the communication capabilities of the smart grid, distributing real-time pricing information can be easily achieved.

If power is priced as a real-time function of demand, it becomes possible for a community to influence the power price rate $Pr(t)$ by adjusting its consumption throughout the day. As shown in [6], [8], a coordinating community that shifts part of its consumption to off-peak hours can significantly reduce its power bill while consuming the same amount of power.

If power is priced as a function of consumption, and the consumption trend has some minimum value $\min$ that occurs at time $t\min$, $Pr(t)$ will also have a minimum value at time $t\min$. Therefore, if consumers are able to shift their loads, possibly by scheduling, it would appear that the most cost effective approach is to target $t\min$. However, as the number of consumers who follow this reasoning increases, the time $t\min$ becomes a high demand point and consequently a highly priced point.

Thus, a fair methodology is needed to allocate the low-rated hours to consumers without creating new points of high demand. A perfect situation would be one where the demand function becomes a constant which also leads to a constant price function. In the following sections, we propose a water-filling based method for scheduling power consumption such that the overall demand becomes relatively constant. We also present simulation examples to show the effectiveness of the proposed methodology.

IV. Water-Filling

The classical water-filling algorithm, traditionally used in communications theory, solves the problem of maximizing the mutual information between the input and the output of a communication channel that is composed of several subchannels and that is subject to a global power constraint. In this section, we first present the mathematical model for the channel power allocation problem, its solution (the water-filling result), and explain its implication. We then explain the analogy between the channel power allocation problem and the load scheduling problem illustrating how the water-filling result will be used to build a scheduler for our purpose.

Given a communication channel with multiple subchannels $(1, 2, 3 \ldots L)$ that are subject to noise levels $(\lambda_1^{-1}, \lambda_2^{-1}, \lambda_3^{-1} \ldots \lambda_L^{-1})$, it is desired to transmit a signal such that the channel capacity (i.e., the information rate) is maximized. The signal to be transmitted is subject to the global power constraint $\sum_{i=1}^{L} x_i \leq P_r$, where x_i is the power allocated to subchannel i, and P_T is the total power to be transmitted. This problem can be modeled as the following optimization problem

$$\max_{\{x_i\}} \sum_{i=1}^{L} \log(1 + x_i \lambda_i) \tag{2}$$

subject to

$$\sum_{i=1}^{L} x_i \leq P_t, \qquad x_i \geq 0, \text{ and } 1 \leq i \leq L. \tag{3}$$

Using the Lagrange Multipliers method, the solution of this problem is given by

$$x_i = (\mu - \lambda_i^{-1})^+ \tag{4}$$

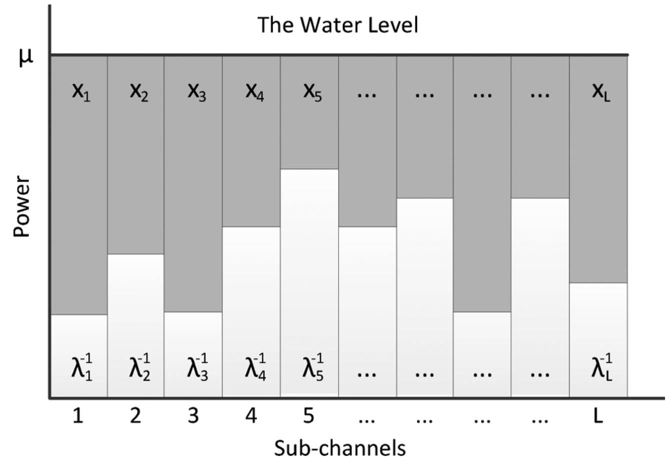

Fig. 2. Water-filling graphical illustration. More transmission power is allocated to channels with less noise.

where $(x)^+ \stackrel{\text{def}}{=} \max(0, x)$ and μ is a constant chosen to satisfy the power constraint in (3).

Equation (4) is known as the water-filling result and it implies that for an optimum solution, more transmission power should be allocated to subchannels with less noise. The constant μ (called the water-level) is selected such that power constraint is satisfied. As shown in Fig. 2, this capacity achieving solution has the visual interpretation of pouring water over the curve given by the subchannel noise (i.e., the inverse of the subchannel gains), and hence the name water-filling or water pouring [9], [16], [17].

Applying the water-filling result for centralized applications where the power allocation is performed by a single entity is a straightforward application of (4) once the water-level is known [10].

The water-filling approach naturally fits the load balancing problem considered in this paper. Specifically, the hard loads are analogues to the noise inherent in the channel, the soft loads are analogues to the power that is to be transmitted through the channel, and the objective is to schedule the soft loads such that the overall energy consumption becomes flat which yields an efficient system. Thus, similar to the classical capacity maximizing problem above, where more transmission power is allocated to subchannels with less noise, soft loads can be allocated to the time slots with less hard loads. This reduces the peak-to-average ratio and, in the event that sufficient soft loads are available, flattens the overall energy consumption profile. As will be explained in Section V, the main difference between our proposed load balancing approach and the traditional water-filling solution above is that our schedulers allocate soft loads probabilistically following a distribution produced from the water-filling result. This enables our protocol to eliminate the need for continuous communication and synchronization between scheduling nodes. As depicted in Fig. 3, our proposed algorithm can be summarized as follows:

Demand Forecast: This is the process where the utility company produces a forecast for the community's power demand.

Water Level Computation: This process computes the correct water-level for use with the forecasted demand pro-

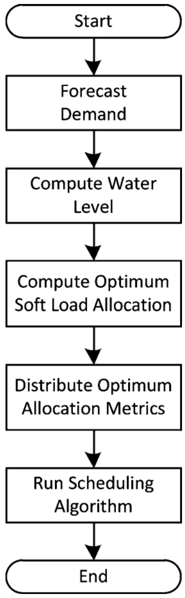

Fig. 3. Basic steps of our proposed algorithm.

file. In our case, as will be explained in the next section, the water-level is set to the lowest point where a constant demand can be achieved.

Optimal Allocation Computation: Once the water-level and the demand forecast are known, the optimal allocation of soft loads can be obtained as a direct application of (4).

Optimum Allocation Metrics Distribution: The optimal allocation produced in the previous process is applicable for a central scheduling entity. In our case, this process produces the probabilistic metrics that enable the distributed schedulers to attain the optimum distribution.

Scheduling Algorithm Execution: The scheduling algorithm, presented in detail in Section VI, uses the metrics distributed through the previous process to achieve an overall flat power consumption profile in a distributed way.

V. Scheduling Power Use

As mentioned previously, a method is needed to allocate consumers' soft consumption to the low priced time slots in a fair way. The allocation strategy has to satisfy the following conditions:

1 Consumers that follow the algorithm pay the same amount for a shifted unit of power regardless of the time slots to which their loads are allocated.

2 It is not possible for a consumer to have a more cost effective allocation.

Both of these conditions will be satisfied if it is possible to achieve and maintain a constant demand. To realize this, we propose allocating soft loads to time slots probabilistically in a way analogous to the water-filling approach described in the previous section. That is, more loads are probabilistically allocated to low demand hours, and less loads are probabilistically allocated to high demand hours. Once an overall constant is achieved, the flat consumption profile is maintained by allocating any additional soft loads uniformly.

Let $H(t)$ denote a function that models the hard power consumption for a community, and let its maximum value be $\max$. Furthermore, let $S(t)$ denote a function that models the soft power consumption of the community. Thus we have

$$D(t) = H(t) + S(t). \tag{5}$$

Our objective is to distributively schedule soft loads such that the overall power demand for the community becomes a constant. Therefore, after scheduling, the following equation must be satisfied:

$$H(t) + S(t) = D(t) = Constant. \tag{6}$$

We then evaluate a Scheduling Distribution (SD) that, when followed, complements the hard loads' consumption such that the total consumption becomes a constant. Similar to water-filling, we relate SD to the difference between the hard loads and the value $\max$. Let $C(t)$ denote the function that complements $H(t)$ to the value $\max$. Thus we have

$$C(t) = \max - H(t). \tag{7}$$

Having $C(t)$, we define SD as a probability distribution that produces $C(t)$ when followed randomly for a sufficiently large number of times. We can obtain SD by normalizing $C(t)$ such that the area under its curve is 1. Thus we have

$$SD(t) = \frac{C(t)}{\int_0^T C(t)dt}. \tag{8}$$

Equation (5) shows that the demand $D(t)$ is a sum of both, the hard and the soft power consumption $S(t)$ and $H(t)$, respectively. Therefore, strictly following (8), the shape of $D(t)$ will either be a constant, or inclined towards $H(t)$ or $S(t)$. This depends on the ratio of the total hard consumption to the total soft consumption. We summarize the possible outcomes as follows:

1. If the amount of soft consumption is too small compared to the hard consumption, scheduling following $S(D)$ will have a minimum effect since $H(t)$ will dominate. Therefore, $D(t)$ will have the shape of $H(t)$ rather than a constant. In this case, scheduling can reduce the peak-to-average ratio of the consumption profile, but cannot flatten it.
2. If the amount of soft consumption is too large with respect to hard consumption, scheduling following $S(D)$ will cause the scheduled loads to dominate, which results in $D(t)$ following the shape of $C(t)$ rather than a constant (i.e., causing an overshot).
3. If the ratio of the soft consumption to the hard consumption equals the ratio of the area under $C(t)$ to the area under $H(t)$, scheduling following the distribution stated will complement $H(t)$ causing $D(t)$ to become flat.

With the increased demand for electric vehicles, which are a form of energy storage and can usually be scheduled, it is expected that the energy consumption of an average home will double [11]. This suggests that having a large amount of soft consumption is a reasonable assumption which eliminates the first outcome above.

Having a large amount of soft consumption, it becomes possible to adjust the second outcome to result in an overall consumption that is constant. This can be done by having each node schedule its soft loads partially following SD and partially following a uniform distribution. More precisely, SD should be followed to a point where it complements the hard consumption to the water-level. After this point, to maintain the overall constant demand, nodes should schedule any remaining soft loads uniformly. For each node, a threshold defines the point of switching from scheduling following SD to scheduling uniformly. This threshold is computed as a factor R of the node's hard consumption. The value of R is to be provided by the utility company and its computation is presented in the Appendix.

Therefore, for a given node k with a total hard consumption, H_k, and a total soft consumption, S_k, the threshold is given by W_k where

$$W_k = R \times H_k. \tag{9}$$

Specifically,

$$W_k = \frac{\int_0^T C(t)dt}{\int_0^T H(t)dt} \times H_k. \tag{10}$$

Therefore, each node will schedule an amount, W_k, of its soft loads following SD, and any remaining soft loads, $(S_k - W_k)$, uniformly. Scheduling loads as described above results in an overall flat consumption.

VI. Proposed Distributed Algorithm

In this section, we present an algorithm for scheduling the soft loads throughout the day to achieve the objectives described in the previous section. Our algorithm is based on a probabilistic model that uses forecasted approximations (or projections) of consumption trends rather than real-time computed values. In this section, we describe the environment, the information available to each consumer (scheduling node), and the proposed algorithm. In the following section, we present a simulation of this algorithm for a small residential community.

The main advantage in having participating nodes schedule probabilistically following a centrally computed distribution is to eliminate the need for continuous communication and synchronization between nodes. Furthermore, this also reduces the processing required at each node. This approach is possible provided that some basic information related to the overall usage trend can be made available to each scheduling node.

Obtaining perfectly accurate values for $C(t)$ would require gathering and processing the community's demand in real-time. This would cause a large overhead in communications and would have to be made available upfront. Fortunately, a community's overall consumption follows a trend, which, as shown in Fig. 1, is fairly predictable [20]. Using a projection may introduce a margin of error (see Section VIII), but it eliminates the need of gathering and distributing consumer's data in real-time. Furthermore, in practice, loads to be scheduled may not be equal in power demand or in the duration of operation. Therefore, scheduling loads to obtain a perfect complementing function $C(t)$ may not be possible as this would require loads to be infinitely small and have no operation time restrictions whatsoever. Despite these limitations, as shown by our simulations, scheduling following the proposed method significantly reduces the peak-to-average ratio of the overall demand. This is due to the fact that the power consumption of individual appliances are significantly small in comparison with the overall power demand of a community.

In what follows we assume that the following information can be made available to each scheduling node:

1. A distribution (SD) that, when followed, complements the trend of the hard consumption towards a constant.
2. A ratio R for computing the threshold that indicates how much of the soft load should be scheduled following SD.

The algorithm below reflects the methodology discussed in Section V. At each scheduling node, the initialization phase acquires the centrally computed scheduling information. At this point, an accumulator is defined to keep track of how much load had already been scheduled. If the total amount of scheduled loads is less than the threshold, the next load will be scheduled following and the accumulator will be incremented by the weight of this load. When the threshold is exceeded, the next item to be scheduled will be scheduled uniformly across all time slots.

Algorithm 1 Scheduling Soft Loads to Flatten the Overall Power Consumption Profile

```
Scheduling Algorithm

Initialization
  Acquire SD
  Acquire Threshold
  Reset Accumulator

For each Soft Load, weighted W
  If (Accumulator < Threshold)
    Schedule Soft Load following SD
    Accumulator = Accumulator + W
  Else
    Schedule Soft Load uniformly
    Accumulator = Accumulator + W
  End If
End For
```

VII. Simulation Results

To demonstrate the effectiveness of our proposed algorithm, we simulate the basic power consumption of a residential community. Initially, we simulate the normal use pattern where each home consumes power without scheduling any loads. These initial rounds resemble historic data used in demand forecasting by the utility company. Next, we use the simulated data to compute SD as described in Section V. Finally, we simulate the community's power consumption when following the proposed algorithm and present our results.

To do this, we first define a selection of appliances available in most homes and assign them their typical power ratings. Table I shows some of the basic attributes of these appliances. We approximate the operation time for each appliance in the second column. The third column shows whether this load can

TABLE I
SIMULATED HOUSEHOLD APPLIANCES

Appliance	*Operation Time*	*Soft Load*	*Continuous Use*
Clothes Dryer	1 hour	Yes	Yes
Electric Vehicle	5 hour	Yes	No
Clothes Washer	0.5 hour	Yes	Yes
Climate Control	5 hour	No	N/A
Water Heater	6 hour	No	N/A
Range (1st run)	1 hour	No	N/A
Range (2nd run)	1 hour	No	N/A
Electronics	5 hour	No	N/A
Lighting	6 hour	No	N/A
Fridge	24 hour	No	N/A
Kitchen App. (1st run)	1 hour	No	N/A
Kitchen App. (2nd run)	1 hour	No	N/A

TABLE II
APPLIANCE OPERATION TIME SELECTION METHODOLOGY

Appliance	*Starts Following*
Clothes Dryer	Normal dist. mean at 5PM
Electric Vehicle	Normal dist. mean at 6PM
Clothes Washer	Normal dist. mean at 5PM
Climate Control	Uniform dist. over 24 hours
Water Heater	Uniform dist. from 8AM to 12AM
Range (1st run)	Normal dist. mean at 1PM
Range (2nd run)	Normal dist. mean at 6PM
Electronics	Uniform dist. from 3PM to 1AM
Lighting	Uniform dist. from 8AM to 1AM
Fridge	Uniform dist. over 24 hours
Kitchen Appliances (1st run)	Normal dist. mean at 1PM
Kitchen Appliances (2nd run)	Normal dist. mean at 6PM

be shifted or not based on its nature of use, i.e., if it is a soft load. The last column shows if the time slots have to be allocated continuously. That is, if the nature of the appliance dictates that all the required slots must be scheduled back to back. This is only considered for schedulable appliances.

In our simulation, we divide the scheduling interval into 30-minute time slots and the appliances are turned on for the duration of a multiple of these slots. Each appliance is turned on probabilistically at a time slot either following a uniform or a normal distribution. The details of the distributions used for each appliance are shown in Table II.

Appliances that are modeled by a uniform distribution are constrained with a time frame. That is, they can only operate during that time frame. Furthermore, appliances that are modeled by a normal distribution are constrained with a mean value. Therefore, most of the time, these appliances will be turned on at the time that corresponds to the mean. If an appliance is to be simulated for two operating periods, we list it twice. For example, we assume that kitchen appliances are used twice a day and, consequently, they have two entries with two different mean values.

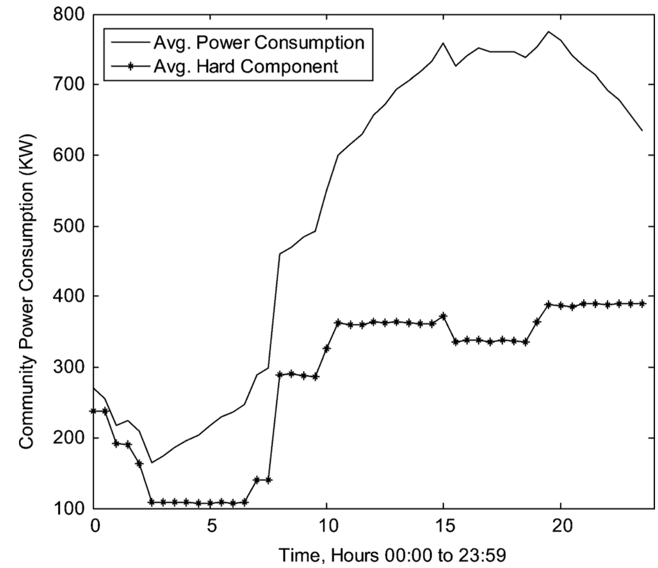

Fig. 4. Average power consumption for the simulated community over 50 rounds without scheduling (mean = 524, standard deviation = 226.8, peak-to-average ratio = 1.47).

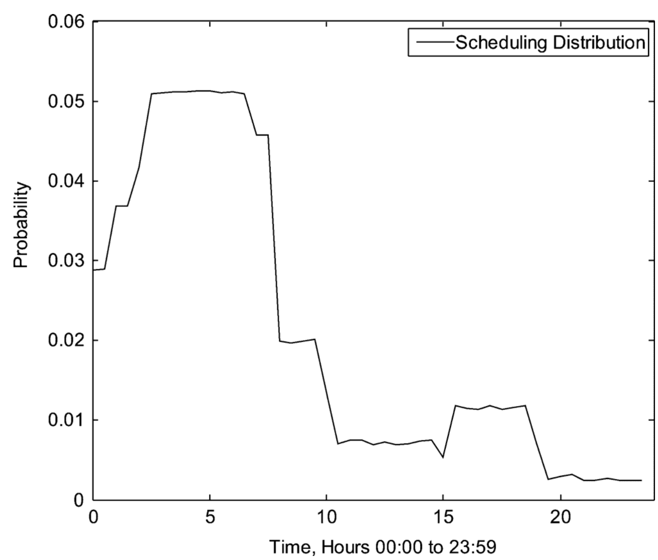

Fig. 5. The Scheduling Distribution (SD).

Following Tables I and II, simulating the power consumption for a community of a thousand homes (nodes) results in the overall power demand shown in Fig. 4. In this figure, the upper curve shows the total consumption, i.e., the summation of the soft consumption and the hard consumption. The figure also shows the hard consumption component of the total use.

The hard consumption component is used to compute SD as described in (7) and (8). SD, shown in Fig. 5, is computed once based on an average case and is made available to all nodes. As the figure shows, a node scheduling loads following this distribution has a higher probability of scheduling at 5 A.M. in comparison to 6 P.M.. Moreover scheduling loads is not limited to off-peak hours, rather, it extends throughout the day with varying probabilities.

The results corresponding to the case when following our scheduling algorithm are shown in Fig. 6. The figure shows the overall hard loads of the community being low in the early hours of the morning and increasing throughout the day. It also

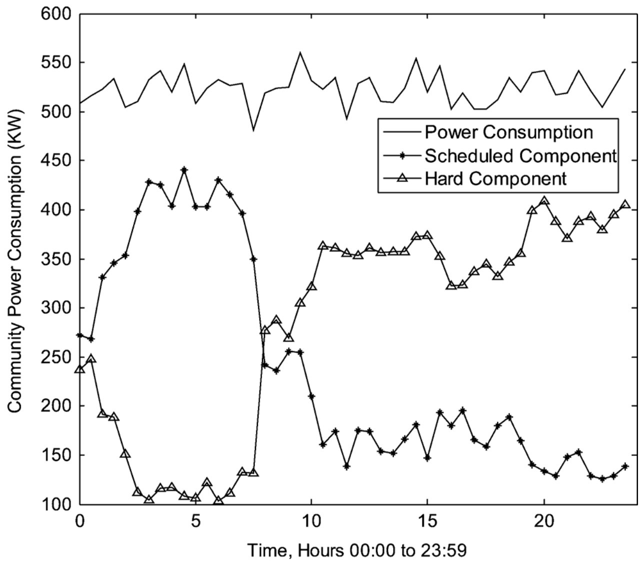

Fig. 6. A single round power consumption for the simulated community with scheduling (mean = 523, standard deviation = 15.8, peak-to-average ratio = 1.069).

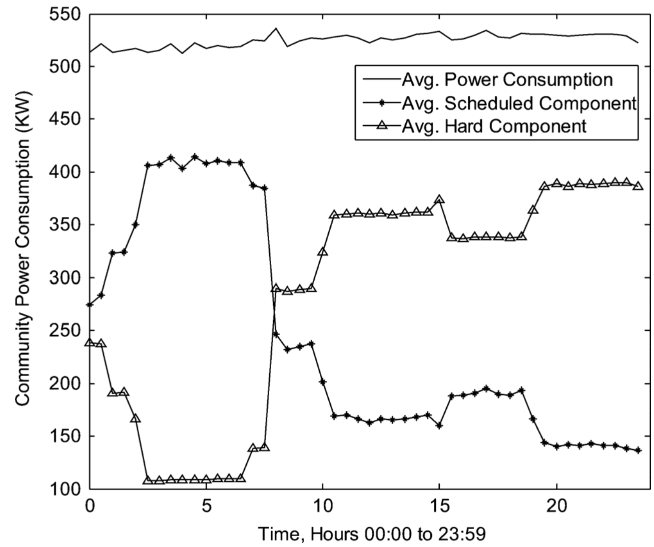

Fig. 7. Average power consumption for the simulated community over 50 runs with scheduling (mean = 524, standard deviation = 6.1, peak-to-average ratio = 1.021).

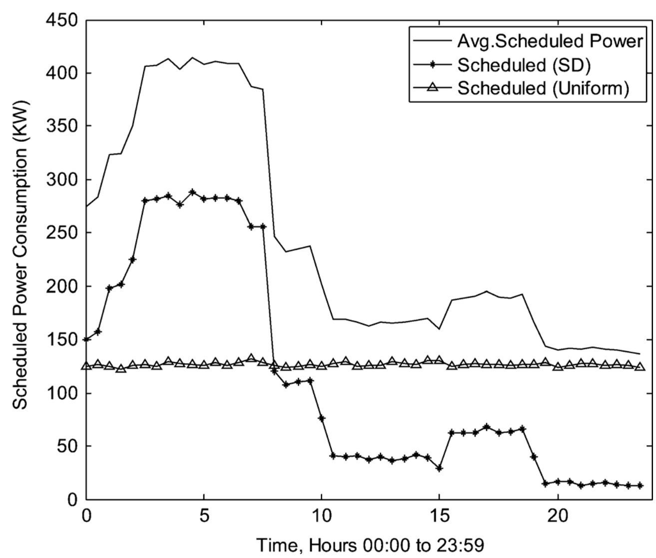

Fig. 8. Average scheduled power consumption for the simulated community over 50 runs.

shows the overall scheduled consumption and the total use. As depicted in the figure, the distributed schedulers clearly counter the hard consumption by scheduling most loads towards the early morning hours. Furthermore, the schedulers prevent a peak-hour from appearing at low demand hours and maintain an overall relatively flat consumption. Our approach significantly reduces the peak-to-average ratio (to 1.069 from 1.47) and the standard deviation (to 15.8 from 226) with a relatively small processing and communication overhead. Furthermore, it ensures that consumers pay relatively the same amount (per unit of power) regardless of the time slot their loads were allocated.

Fig. 7 shows the average power consumption profile obtained by repeating the simulation for 50 runs. As shown in the figure, the average power consumption is much closer to a constant with minimal peak-to-average ratio of 1.02, and a standard deviation of 6.1. Fig. 8 shows the breakdown of how the community's loads were scheduled in the average case. As shown, a portion of the soft loads are scheduled following SD and any remaining loads are scheduled uniformly. Furthermore, the figure shows that scheduling nodes can be enhanced to accommodate further operation time constraints. For example, in the portion scheduled uniformly, a scheduler could locally swap loads to accommodate additional time constraints.

VIII. Algorithm Analysis

In this section, we present our analysis of the proposed approach from various perspectives. Specifically, we analyze the fairness of the algorithm, the effect of errors in demand forecasting, and the participation incentives for consumers from a game theoretic perspective.

A. Fairness

The fairness of the proposed approach is an important factor that would determine the willingness of consumers to follow the scheduling algorithm. To assess the fairness, we use the difference between the starting time for soft load devices without any scheduling and after following the scheduling algorithm as a measure of inconvenience amongst all users. A scheduling algorithm would be considered unfair if some users would have a significantly larger difference compared to others. On the other hand, if all users have relatively the same average change in time, then the algorithm would be considered fair.

More precisely, let Δ_i denote the average absolute value of change in the starting time of the scheduled soft loads belonging to the $\mathrm{i^{th}}$ user (reference to the case without scheduling). In other words, Δ_i is a measure of the inconvenience caused to the $\mathrm{i^{th}}$ user for opting to follow the scheduling algorithm. Let Δ denote the average of Δ_i's for all users. Fig. 9 below shows the average deviation of Δ_i from Δ obtained by simulating our scheduling algorithm with a thousand users for a thousand runs.

As shown in the figure, on average, the discrepancy among users is very low, i.e., the soft loads of most users are shifted by almost the same amount of time which reflects the fairness of the proposed approach.

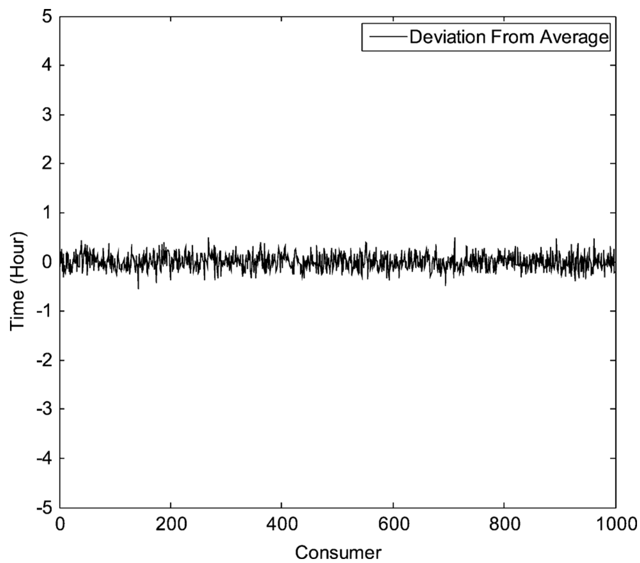

Fig. 9. Fairness of the proposed algorithm. A low deviation from average amongst all users indicates a fair algorithm.

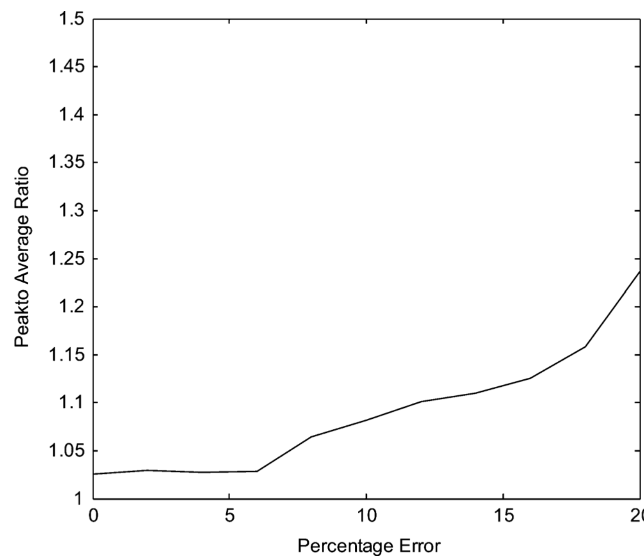

Fig. 10. The effect of error in consumption forecast on the peak-to-average ratio after scheduling. The peak-to-average ratio remains within 1.1 for up to about $\pm 10\%$ in error.

B. *Effect of Errors in the Demand Forecast*

The proposed algorithm relies on historical power consumption data to forecast the consumption for the upcoming day. This information is used to compute SD that guides scheduling nodes in distributing their loads throughout the day. Therefore, the effectiveness of the scheduling algorithm depends on the accuracy of the forecasted demand.

As Fig. 1 shows, utility companies have the ability to forecast the power demand of a community with good accuracy. Furthermore, typical demand forecasting performance indicators that are available in the public domain (e.g., see [20]) show that when the projection is made a day ahead, the estimation error is usually within $\pm 10\%$.

To evaluate the robustness of our proposed approach against errors in the demand forecast, we introduced a uniformly distributed error between e and $-e$ to the forecasted data; where e is determined as a percentage of the instantaneous consumption at a given time slot. This can be seen as an extreme case as it introduces random errors at all time slots. Fig. 10 shows the effect of error on the peak-to-average ratio of the community's demand after scheduling.

As shown in the figure, the peak-to-average ratio of the demand profile after scheduling remains below 1.1 when the error is bounded to $\pm 10\%$ which reflects the effectiveness and robustness of using forecasted data in scheduling soft loads. Furthermore, the figure shows that the peak-to-average ratio of the scheduling community is not affected by errors up to about 5%. This can be attributed to the inherent difference between runs used to produce SD and runs used to simulate the community's consumption, i.e., due to the probabilistic nature of our proposed approach and simulations.

C. *Participation Incentives*

To further examine the consumer's participation in the proposed scheme, we analyze the incentives of participation from a game theoretic perspective. We take two main factors into account, namely the action of the community and the inconvenience introduced by scheduling.

A consumer's utility from scheduling consumption is directly related to the community's trend of power usage. If the community does not shift its consumption, a consumer would have a greater incentive in shifting loads towards lower demand hours. On the other hand, if sufficient consumers shift their consumption, a consumer would have less incentive in shifting his/her loads.

Allocating utilities for savings in power costs and possible inconvenience of the scheduling process is a great challenge. In what follows, we focus on the outcome for an individual in the proposed setup. The consumer's incentive would be in reducing his/her electric bill, taking the convenience of this reduction as a main factor. That is, a consumer could decide to schedule or not to schedule soft loads in an environment where the community schedules or does not schedule. In our analysis, we use the following notation. If the community does not shift loads, the savings introduced by the consumer's action are represented by S, and savings missed are represented by $-S$. Furthermore, if the community does shift its loads, the introduced savings are represented by O, and savings missed by $-O$. We represent the inconvenience of scheduling by $-I$. Thus, the possible combinations, summarized in Table III, are:

1. The consumer follows the scheduling algorithm in an environment that follows the algorithm. In this case, the consumer's soft loads draw power at their allocated time. This results in the near optimum overall consumption, which offers the community the best possible rate. We represent the incentive of this outcome by $O - I$.
2. The consumer does not follow the scheduling algorithm in an environment that follows the algorithm. In this case, the consumer misses the slot allocated by the algorithm, and consumes power at a different time resulting in increased demand at that time. Because of the loss, we represent this outcome by $-O$.

TABLE III
SCHEDULING OUTCOME FOR THE CONSUMER

	Consumer Schedules	*Consumer Does Not Schedule*
Community Schedules	*O-I*	*-O*
Community Does Not Schedule	*S-I*	*-S*

3. The consumer follows the scheduling algorithm in an environment that does not follow the algorithm. In this case, the consumer's soft loads are mostly moved toward low demand hours. This allows the consumer to introduce savings in power cost. We represent the incentive of this outcome by $S - I$.
4. The consumer does not follow the scheduling algorithm in an environment that does not follow the algorithm. In this case, neither the consumer, nor the community benefit from the opportunity. Due to the loss, we represent the outcome by $-S$.

With $S > O$, analyzing the previous table shows that if $O > I/2$, scheduling loads is a dominant strategy. That is, the left column would always be greater than the right column. Therefore, given the constraint above, the choice "to schedule" would be a strongly dominant strategy and consequently, a Nash equilibrium [12], [13].

IX. CONCLUSION AND FUTURE WORK

The communication and computation capabilities of the smart grid make it possible to have a dynamic pricing scheme where the rate for power is a direct function of the instantaneous demand. In particular, it becomes possible for a community to adjust its power consumption controlling the price rate and eliminating operational inefficiencies in addition to benefiting from lower power costs. Within this context, we proposed a methodology that allows consumers to shift part of their soft loads to off-peak hours in a probabilistic way that results in a relatively constant overall power consumption profile. Our results confirm the efficiency and fairness of the proposed scheduling algorithm.

It should be noted that throughout our analysis, we have considered appliances to be either completely shiftable (soft loads), or completely nonshiftable (hard loads). In reality, some appliances can only accommodate small tolerance in their operation time, and therefore, cannot be considered completely shiftable. One possible area of research on this front would be to model the constraints associated with such appliances and study their effect on the peak-to-average ratio of the scheduling process.

APPENDIX

In this appendix, we describe the method used in computing the ratio needed to calculate the threshold in (9). Assuming we have a community with n nodes, let $S_k(t)$ represent the soft power consumption of the k^{th} node as a function of time. Similarly, let $H_k(t)$ represent the hard power consumption for that node. Therefore, the overall soft and the overall hard consumption for all members of the community can be respectively represented by

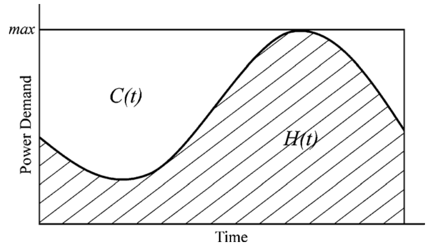

Fig. 11. Complementing the hard consumption $H(t)$ to the water-level $\max$ using $C(t)$.

$$S(t) = \sum_{k=1}^{n} S_k(t) \tag{11}$$

and

$$H(t) = \sum_{k=1}^{n} H_k(t). \tag{12}$$

For the hard demand curve $H(t)$, its maximum value $\max$ is the lowest value that can be used as a water-level that produces a constant demand. This is illustrated in Fig. 11. Therefore, based on water-filling, we define $C(t)$ as the function that complements $H(t)$ to $\max$, i.e.,

$$C(t) = \max - \mathrm{H(t)}. \tag{13}$$

We define the factor R as the ratio of the area under $C(t)$ to the area under $H(t)$. Thus we have

$$R = \frac{\int_0^T C(t)dt}{\int_0^T H(t)dt}. \tag{14}$$

From (12), we see that $H(t)$ was formed by summing each user's hard component. We propose that $C(t)$ be formed in a similar way. Assuming that consumers have a sufficient amount of soft loads, each will dedicate a portion to produce $C(t)$ That is

$$H(t) = \sum_{k=1}^{n} H_k(t) \tag{15}$$

$$R\int_0^T H(t)\,dt = R\sum_{k=1}^{n}\int_0^T H_k(t)\,dt \tag{16}$$

$$\frac{\int_0^T C(t)dt}{\int_0^T H(t)\,dt} \times \int_0^T H(t)\,dt = \int_0^T C(t)\,dt \tag{17}$$

$$\int_0^T C(t)\,dt = \sum_{k=1}^{n} R\int_0^T H(t)\,dt. \tag{18}$$

Therefore, the amount of soft loads needed to form $C(t)$ can be obtained if each consumer contributes with an amount of soft loads equivalent to a factor R of their hard consumption. Scheduling these soft loads probabilistically following SD, as defined in (8) produces $C(t)$, which in turn complements $H(t)$ to the constant $\max$.

Acknowledgment

The authors would like to thank the anonymous reviewers for their constructive comments that greatly helped improve this paper.

References

[1] "The smart grid: An introduction," U.S. Dept. Energy, 2009.
[2] E. Santacana, G. Rackliffe, L. Tang, and X. Feng, "Getting smart," *IEEE Power Energy Mag.*, vol. 8, no. 2, pp. 41–48, Mar./Apr. 2010.
[3] S. E. Collier, "Ten steps to a smarter grid," in *Proc. IEEE Rural Electr. Power Conf. (REPC)*, Apr. 2009, pp. B2–B7.
[4] R. Krishnan, "Meters of tomorrow," *IEEE Power Energy Mag.*, vol. 6, pp. 92–94, Mar. 2008.
[5] "Ontario demand and market prices," Independent Electricity System Operator, Sep. 2011 [Online]. Available: http://www.ieso.ca
[6] S. Caron and G. Kesidis, "Incentive-based energy consumption scheduling algorithms for the smart grid," in *Proc. 1st IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, 2010, pp. 391–396.
[7] C. Chen, S. Kishore, and L. V. Snyder, "An innovative RTP-based residential power scheduling scheme for smart grids," *Proc. ICASSP*, 2011.
[8] Mohsenian-Rad, V. Wong, J. Jatskevich, and R. Schober, "Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid," in *IEEE PES Conf. Innov. Smart Grid Technol.*, 2010.
[9] C. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, pp. 10–21, Jan. 1949.
[10] D. P. Palomar and J. R. Fonollosa, "Practical algorithms for a family of waterfilling solutions," *IEEE Trans. Signal Process.*, vol. 53, no. 2, pp. 686–695, 2005.
[11] Ipakchi and Albuyeh, "Grid of the future," *IEEE Power Energy Mag.*, vol. 7, no. 2, pp. 52–62, Mar.–Apr. 2009.
[12] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA: MIT, 1994.
[13] P. D. Straffin, *Game Theory and Strategy*. Cambridge, U.K.: Cambridge Univ. Press, 1996.
[14] G. Xiong, C. Chen, S. Kishore, and A. Yener, "Smart (in-home) power scheduling for demand response on the smart grid," *Proc. Innov. Smart Grid Technol. (ISGT)*, 2011.
[15] N. Gatsis and G. Giannakis, "Cooperative multi-residence demand response scheduling," in *Proc. 45th Annu. Conf. Inf. Sci. Syst. (CISS)*, 2011.
[16] T. M. Cover and A. T. Joy, *Elements of Information Theory*. Hoboken, NJ: Wiley-Interscience, 2006.
[17] Proakis and G. John, *Digital Communications*. Boston, MA: McGraw-Hill, 2001.
[18] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena Scientific, 1999.
[19] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. IEEE Smart Grid Commun. Conf*, Gaithersburg, MD, 2010.
[20] "IESO demand forecasting performance indicators," Demand Forecast Deviations Working Group, the Independent Electricity System Operator (IESO), Sep. 2011 [Online]. Available: http://www.ieso.ca/imoweb/marketData/marketData.asp

Merwais Shinwari received the B.Sc. degree from the College of Computer Science and Engineering (CCSE), Computer Engineering Department, King Fahd University of Petroleum and Minerals in 2007. He is currently working toward the M.S. degree at the Concordia Institute for Information Systems Engineering (CIISE), Concordia University, Montreal, QC, Canada.

His main research interests are in the area of the smart grid.

Amr Youssef (SM'06) received the B.Sc. and M.Sc. degrees from Cairo University, Cairo, Egypt, in 1990 and 1993, respectively, and the Ph.D. degree from Queens University, Kingston, ON., Canada, in 1997.

He is currently a Professor at the Concordia Institute for Information Systems Engineering (CIISE) at Concordia University, Montreal, QC, Canada. Before joining CIISE, he worked for Nortel Networks, the Center for Applied Cryptographic Research at the University of Waterloo, IBM, and Cairo University. His research interests include cryptology, network security, and wireless communications.

Walaa Hamouda (SM'06) received the M.A.Sc. and Ph.D. degrees from Queen's University, Kingston, ON, Canada, in 1998 and 2002, respectively, both in electrical and computer engineering.

In July 2002, he joined the Department of Electrical and Computer Engineering, Concordia University, Montreal, QC, Canada, where he is now an Associate Professor. From June 2006, he has been appointed as a Concordia University Research Chair in Communications and Networks. His current research interests are in wireless networks, MIMO space-time processing, multiuser communications, cross-layer design, source and channel coding.

Dr. Hamouda served as the technical co-chair of the Wireless Networks Symposium, IEEE Globecom 2012, the Ad-hoc-Sensor and Mesh Networking Symposium of the ICC 2010, the 25th Queen's Biennial Symposium on Communications, track co-chair of the Transmission Techniques, VTC-Fall 2012, and track chair of the Radio Access Techniques of the IEEE VTC-Fall 2006. From September 2005 to November 2008, he was the Chair of the IEEE Montreal chapter in Communications and Information Theory. He received many awards including the best paper award [WNS] of the ICC 2009, and the IEEE Canada Certificate of Appreciation in 2007, 2008. He serves as Associate Editor for the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY and the IEEE COMMUNICATIONS LETTERS, and the IET *Wireless Sensor Systems*.