

Solving the Fair Electric Load Shedding Problem in Developing Countries

Olabambo Ifeoluwa Oluwasuji · Obaid
Malik · Jie Zhang · Sarvapali Dyanand
Ramchurn

Received: date / Accepted: date

Abstract Often because of limitations in generation capacity of power stations, many developing countries frequently resort to disconnecting large parts of the power grid from supply, a process termed load shedding. This leaves households in disconnected parts without electricity, causing them inconvenience and discomfort. Without fairness being taken into due consideration during load shedding, some households may suffer more than others. In this paper, we solve the fair load shedding problem (FLSP) by creating solutions which connect households to supply based on some fairness criteria (i.e., to fairly connect homes to supply in terms of duration, their electricity needs, and their demand), which we model as their utilities. First, we briefly describe some state-of-art household-level load shedding heuristics which meet the first criteria. Second, we model the FLSP as a resource allocation problem, which we formulate into two Mixed Integer Programming (MIP) problems based on the Multiple Knapsack Problem. In so doing, we use the utilitarian, egalitarian and envy-freeness social welfare metrics to develop objectives and constraints that ensure our FLSP solutions results in fair allocations that consider the utilities of agents. Then, we solve the FLSP and show that our MIP models maximize the groupwise and individual utilities of agents, and minimize the differences between their pairwise utilities under a number of experiments. When taken together, our endeavour establishes a set of benchmarks for fair load shedding schemes, and provide insights for designing fair allocation solutions for other scarce resources.

Keywords Load shedding · Fairness · Constrained optimization

O. I. Oluwasuji · O. Malik · J. Zhang · S. D. Ramchurn
School of Electronics and Computer Science
University of Southampton, Southampton, UK
E-mail: oio1a14,o.malik,jie.zhang,sdr1@soton.ac.uk

1 Introduction

Energy is the key driver for growth in developing countries. However, many such countries face significant challenges in providing enough energy to power their industries and communities [20]. To put this into perspective, the United Kingdom generates over 30 GW of electricity [49] for a population of around 65 million people. Comparatively, Nigeria, a developing country, generates under 8 MW for a population of over 170 million people [34]. Furthermore, the demand for electricity is steadily growing in developing countries, due to an increase in population, the modernization of poorer areas and an increase in the number of digital appliances and devices using electricity within these countries. In addition, while many of these countries do not have the means to channel the required investment into developing their energy sectors, it is also generally cumbersome and time consuming to increase the generation capacity of electric grids. As a result of these, the energy challenges in developing countries will be ongoing for the near future.

Although there is a wide gap between the demand on the grid and its supply capacity (or generation) in many developing countries, there is also the need to constantly maintain the alternating current frequency of the power system at its operational value (50 Hz or 60 Hz). The consequence of this in developing countries is that large parts of their electric networks are repeatedly disconnected from supply. This measure of disconnecting parts of a grid network from supply is termed load shedding. Load shedding reduces the strain on the grid and prevents it from collapsing. However, while load shedding results in the stability of grid operation, it also results in many homes within parts of the network being left without electricity. The effect of this is that the occupants of these homes are seriously discomforted and inconvenienced. Furthermore, some homes bear the brunt of these effects because standard load shedding techniques do not focus on fairly allocating electricity to homes, as much as they do on maintaining a demand-supply balance. As such, while load shedding will remain a common practice in developing countries for the near future, it is absolutely necessary to develop solutions which reduce unfairness in the system. Such solutions will improve the general availability of electricity, present a better platform for fighting poverty, increase the welfare of individuals and enhance societal development.

Against this background, this paper presents load shedding solutions that consider the heterogeneous electricity needs of households and uses these to fairly connect households to supply. The solutions build on existing research in two key areas, which we discuss in detail in Section 2. The first is in the area of resource allocation,¹ where solutions which fairly allocate resources within dynamic settings have been proposed. The second culminated in the design and development of cheap smart meter retrofits specifically for use in developing countries [3,21]. The retrofits provide functionalities upon which the load shedding problem can be modelled into a resource allocation problem that

¹ A resource allocation problem is a fair division problem whose solution involves finding an allocation of limited resources between a number of interested entities, subject to the availability of the resource and how interested the entities are in the resource [7].

implements load shedding at the household level, such that their electricity needs at different times can be taken into account.

On account of this, we present a novel approach to load shedding. Firstly, we execute load shedding at the household level, so that the electricity needs of households can be considered in designing fair load shedding solutions. In addition, household-level load shedding reduces waste and increases revenue by mitigating against overshedding that may arise when load shedding is implemented at network level.² We develop two sets of household-level solutions. The first considers the length of time individual households are connected to electricity, with the aim to keep these households connected to supply as equally as possible. The second uses the agent model of households to formulate the fair load shedding problem into Mixed Integer Programming (MIP) problems. The MIP models are designed to maximize different requirements, while minimizing pairwise differences between the duration agents are connected to supply, the level of satisfaction of their individual electricity needs, and the level of supply to them. It is noteworthy that all our solutions ensure demand is as close to supply as possible.

It should also be mentioned that this paper presents an approach that can be exploited when simulating datasets in future. Our approach involves analyzing the factors which affect the consumption of electricity in households, then using these to adapt a publicly available, verifiable dataset of household consumption (for households in a developed country) to our use case.³ We take this approach because we need a dataset to implement and evaluate our load shedding solutions, as no dataset of household electricity consumption for developing countries currently exists.

In addition, as will be seen in this paper, our approach presents a framework upon which other fairness problems involving constrained utility maximization (or resource allocation) may be generalized. Specifically, our approach dissects the general fairness problem in terms of modelling user utilities, preferences or comfort levels, and using these utilities within a constraint optimization solution that maximizes the utilities allocated to users (independently and collectively) and minimizes the differences between their individual allocations. In order to do so, we model some appropriate social welfare metrics into constraints which we use within mathematical programming models of the fairness and efficiency (allocation) problem.

The rest of the paper is organized as follows. In Section 2, we explore the bodies of work related to our research problem. We show how we create a relevant dataset and model households into agents in Section 3. We also describe our setting for the load shedding problem and discuss our key assumptions therein. In Section 4, we present and evaluate four household-level load shedding heuristics against some standard social welfare metrics. We also identify the shortcomings of these heuristics within the same section. To overcome these shortcomings, we model the fair load shedding problem into a

² Conventional load shedding approaches implement load shedding by disconnecting the supply to whole buses or substations on the power network. This is always likely to result in more load than is required being disconnected from supply.

³ We discuss this in detail in Section 3.

constrained optimization problem in Section 5. Using the same social welfare metrics (employed in Section 4), in Section 6, we evaluate and compare the results of all our solutions. We conclude in Section 7.

2 Related Work

A number of fair network-level load shedding techniques have been proposed [35]. The first was a simple technique that focused on reducing the same percentage of load on all buses of a power system (i.e., parts of a power system that are each made up of multiple households) at the event of an overload. As such, when there is a grid deficit, their solution executes load shedding by disconnecting equal amounts of load from all buses of the system in a simple and fast manner. However, it also results in too much electric load being disconnected from the grid. To mitigate against this, they proposed a more efficient technique that disconnects the same percentage of load on only a subset of buses on the grid. The set of affected buses includes all buses within a tree whose root is the bus of the initially failed line, and leaf is a generator, a peripheral bus or a sink bus. In comparison with their first approach, this results in more efficient load shedding, but reduces the size of the system being affected by the measure. However, the implication is that, firstly, only the bus on which an overload occurs, and some others around it, will always be affected by load shedding. As such, some buses which constitute small loads around the overloaded bus may be affected. Secondly, load is disconnected from overloaded buses at periods when they need electricity the most. The consequence of this is that the cost of load shedding to homes running critical activities within an overloaded bus (or buses around it) increases. As such, it is important to consider electricity needs when developing load shedding solutions.

In line with this, the electricity needs (or preferences) of buses on the power system were considered in [46]. In so doing, they designed a solution where load shedding depends on interactions between intelligent agents, each representing a bus on the system. The agents communicate with their neighbours and work out how much they can contribute to load shedding exercises, based on their electricity needs. They also determine the associated compensations for their contribution towards load shedding. In their solution, they minimize both the amount of load disconnected from supply and the aggregate costs of load shedding to the buses that participate in load shedding. Note that their solution is applicable to electric grids where power buses are able to communicate and interact with each other, as is the case in a smart grid. In addition, to be relevant to developing countries, there is a need to develop solutions with lower complexities, as are atypical of solutions which deal with user-determined (or dependent) preferences, incentives and non-compliance (as in [46]). This is because preferences may vary over time significantly, models of incentives are not always linear utility functions and consumers may not always behave rationally, thus increasing the complexity of solutions.

Nonetheless, each bus on the network is made up of lines, while each line is made up of multiple individual consumers (same in [35] above). As such,

Table 1 Summary of capabilities of a smart retrofitted meter [21].

Capabilities	
1	Two-way communication via GSM
2	Secure transmission of energy consumption data
3	Display of usage statistics
4	Remote connection and disconnection
5	Tamper proof mechanism
6	Power quality measurements (such as voltage and phase angle)
7	Setting of thresholds
8	Outage notifications and tariff updates

each bus is representative of a collection of individual consumers which all have heterogeneous electricity needs. A fair load shedding solution should take these heterogeneous needs into account in fairly implementing load shedding, as in a resource allocation problem.

Thereupon, we first propose that the needs of households within buses or substations should be considered when developing fair load shedding solutions. This is now possible because of recent advancements in the design of smart retrofits specifically for use within developing countries [3,21]. These retrofits employ Global System for Mobile Communications (GSM) technology for connecting individual meters and utilities, such that households can be remotely disconnected from and reconnected to supply. In addition, as a benefit of other capabilities of the retrofits, the consumption history of households can be used to analyze the needs of individual households, as well as to plan for load shedding ahead. Table 1 displays a summary of the capabilities of the designed retrofits, all of which present a foundation upon which electric grids in developing countries can evolve into smart grids.

Given this, in [33], we presented four heuristic algorithms an operator can employ to fairly disconnect households from supply during load shedding in developing countries. Our selection process was centred around ensuring that households are connected to electricity as equally as possible, in terms of number of hours. We employed the utilitarian, egalitarian and envy-freeness social welfare metrics in assessing the performance of these solutions. Together with our proposal for executing load shedding at the household level, our solutions also present an avenue for maximizing revenue and in effect, minimizing waste and utilizing electricity better. However, apart from providing solutions that connect households to electricity as equally as possible, we failed to take into account the heterogeneous needs of individual customers. As such, our solutions would, time and again, result in homes being disconnected from electricity when they need it the most (as in [35] above).

On account of this, the household-level Fair Load Shedding Problem (FLSP) can be mapped to a fair resource allocation problem (defined in Section 1) where the resource that needs to be divided among entities (i.e., households) is grid supply. Different resource allocation problems have been solved on electric grids using the multiagent system (MAS) approach. For example, in [5], the MAS approach was used to solve the resource allocation problem of trad-

ing electricity within the smart grid. They showed that this approach produces better results (in terms of fairness and efficiency) than other traditional (non-MAS-based) resource allocation methods. In addition, a resource allocation problem, where charging slots were allocated to electric vehicles (EVs) based on reports⁴ submitted by agents representing owners of EVs, was solved in [15]. This particular problem was solved in order to prevent overloading the network when charging the EVs. The problem was extended in [48] by considering EVs which have preferences that the system commits to fulfilling to an extent. Additionally, agent-based coordination algorithms were used to solve the optimal dispatch problem of allocating power outputs of intermittent power generators in [25], such that the electricity generated by these renewable sources is optimally utilized. Note that there is a necessity for self-interested participants to behave rationally and truthfully report their preferences in all of the studies. However, in our FLSP, we centrally generate these preferences in order to reduce the complexities and cost (in relation to those used in [15, 48, 25]) of eliciting the information from household agents.

To solve the FLSP, we use insights from the knapsack packing problem. A knapsack problem is one where a fixed-capacity knapsack is to be fitted with a set of items, each with its weight and value (or profit), in such a way that the knapsack holds the items with the highest values within its capacity [24]. This is often modelled as an optimization problem and solved using linear programming (or MIP), where a combination of items are selected to maximize the value of items packed into the knapsack, subject to the knapsack's capacity constraints. In this regard, the electric grid is analogous to a knapsack with a fixed supply capacity at a given time and agents with varying electricity needs and usage values.

To solve the FLSP over different allocations, we consider another variant of the knapsack problem, namely the multiple knapsack problem (MKP). A MKP differs from the classic knapsack problem in that it has m bins or knapsacks (where $m > 1$) [6]. The number of periods in which our FLSP will be solved is likened to the number of knapsacks within the MKP. However, other considerations are necessary before a suitable MKP representation of the fair load shedding problem can be developed. These considerations involve the introduction of constraints that ensure fairness when solving the FLSP as a MIP. Before we describe how we make these considerations, we first describe the methodology employed to generate a dataset applicable to the problem in the section that follows.

3 Modelling Household-level Load Shedding for Developing Countries

We focus on developing load shedding solutions for households. We do this because the residential sector constitutes a large percentage of the demand on

⁴ These reports indicate the periods within which the EVs will be available for charging. As such, they represent user preferences.

the grid. An example is in Nigeria,⁵ where the residential sector accounts for 51.3% of grid demand [30]. As such, effective household-level load shedding solutions will improve grid conditions and energy situations. We simulate a dataset that will be used to solve the household-level load shedding problem in the next section.

3.1 Simulating Developing Country Energy Consumption Data

Our solutions need to be simulated, implemented and evaluated using a relevant real-world household-level electricity consumption dataset of households in a developing country. However, to the best of our knowledge, no such dataset exists.⁶ Now, instead of creating or simulating an entirely artificial dataset, we consider collecting verifiable, authenticated, readily available household consumption data of homes in developed countries, then adapting the data to one representative of a developing country.⁷ This is because an adapted real-world dataset will preserve some of the consumption signatures and features typical to households.

A number of real-world household-level consumption datasets from *developed* countries exist. However, we consider those collected from a multiple of households. As such, in Table 2, we present a few of the datasets collected from over 20 households [27,51].

Our fair load shedding solutions should result in fair allocations over time. Therefore, a suitable dataset for implementing and evaluating our fair load shedding solutions should cover long periods. For this reason, the Tracebase [51,27] and PLAID [14] datasets are inadequate. Consequently, we are left with other datasets which cover a month and over. In further examining the remaining datasets, we consider a number of factors which affect the consumption of electricity in households. These include appliance usage, temperature and consumption habits. We discuss how these factors are used to determine our source of data and simulation approach in the sections that follow.

3.1.1 Appliance Usage

Homes in developing countries do not have as many appliances as those in developed countries. To illustrate this, in 2010, the average consumption of

⁵ We take Nigeria to be our representative developing country in this paper.

⁶ There is the iAWE dataset from India [51]. However, the dataset is collected from only one house. Only a dataset collected from multiple households is useful for simulating, implementing and evaluating our fair load shedding solutions.

⁷ We considered running a field trial to gather data of household consumption in Lagos, Nigeria. We had two options in mind. The first option was to deploy the retrofit smart meters we described in Section 2 within a network made up of a data collection centre and multiple households. Our data collection centre will then receive data from these households using a GSM receiver module, and store the data on a memory. The second option was to connect data collection kits to electric meters in multiple homes. These kits would collect the electricity consumption data for such households using non-invasive AC current sensors, and store these on micro-SD cards. We decided against the field trial majorly due to cost and time restrictions.

Table 2 Publicly available household electricity consumption datasets of 20 or more households [27,51]

	Dataset	Number of Houses	Data Point	Location	Data Length
1	Dataport	Over 1200	Appliance	USA	Over 4 years
2	HES	251	Submeter	United Kingdom	12 months (26 houses) & 1 month (225 houses)
3	Tracebase	158	Appliance	Germany	1 day
4	RBSA	101	Submeter	USA	27 months
5	PLAID	55	Appliance	USA	5 seconds
6	OCTES	33	Submeter	Finland, Iceland & Scotland	4 to 13 months
7	EEU data	23	Submeter	Canada	27 months
8	REFIT	20	Submeter	United Kingdom	24 months

an electrified household in Nigeria was $570kWh$ [55]. In contrast, $11,698kWh$ of electricity was consumed by an electrified home in the USA in the same year [55]. A reason for this is that, on average, homes in Nigeria are poorer than those in the USA, which directly impacts on the appliances used within a typical home in the country.

For this reason, any dataset adapted to the developing country context should be one from which the consumption data of appliances commonly used in developing countries can be extracted. Consequently, the HES [9], RBSA [51], OCTES [27,51], EEU (known as the Electrical-end-use data) [19] and REFIT [28] datasets are inadequate as they were collected at the submeter level. This leaves us with the dataset from Pecan Street Inc’s Dataport [37]. Thankfully, Dataport is the largest provider of disaggregated (i.e., appliance-level) customer energy data [36], as seen from Table 2. In Figure 1, we show the number of occurrences of individual appliances on Dataport [36]. Therefore, we collect the time-series appliance-level data from Dataport for the period of a year. We point out that the dataset on Dataport is collected from households in the USA.

Thereafter, we discover from multiple studies that the appliances typically available in a home in Nigeria include lighting, televisions, electric fans, DVD players, washing machines, electric irons, air conditioners, refrigerators, sewing machine and water pumps [31,45,11,26,32]. Hence, we extract the consumption data from the appliances which are common to both countries from our Dataport dataset. These include the consumption data of air conditioners, washing machines, lighting and refrigerators. In the next section, we discuss how we consider temperature in simulating our representative dataset.

3.1.2 Temperature

In simulating our representative dataset, in this section, we consider the effect which temperature has on the electricity consumed within a household. This consideration is one of the reasons for the wide contrast between the electricity




Fig. 1 Number of appliance occurrences in households on Dataport (reproduced from [36])

consumed in the USA and in Nigeria (as stated above). Now, the location of a household determines the external temperature the household becomes subjected to, while the external temperature influences the electricity consumed in the home. Consequently, the temperature in the USA is such that a typical home in the country spends energy on heating. For instance, in 2010, 41.5% of the average electricity consumed within a home in the USA was spent on heating, while 17.7% was spent on water heating. In turn, about 16% of the average electricity consumed within homes in Nigeria is spent on cooling [56]. To factor this in, we consider the similarities between the monthly temperature of Texas, USA and Lagos, Nigeria. We do this because many of the households whose appliance usage data is collected from Dataport are situated in Texas, while Lagos is largest city in Nigeria (and one of the largest in sub-Saharan Africa). Figure 2 shows how the average monthly temperature of Lagos and Texas [16, 17] are similar during summer in Texas. On this ground, from our Dataport dataset, we extract the data from appliances used within households in Texas during the 13 weeks of summer. By so doing, it is reasonable to assume that the consumption of appliances extracted from our Dataport dataset will resemble those typical of Lagos homes (particularly the air conditioning).

3.1.3 Consumption Patterns

Furthermore, we consider the consumption patterns within typical households in both countries by examining their typical load profiles. To derive the typical load profile of a home in the USA, we aggregate the data originally collected from Dataport and compute the average hourly consumption of a household. We derive the typical load profile of a home in Nigeria from [38]. Thereafter, we compute the load profile of our representative dataset (i.e., the Dataport dataset adapted so far) and present it in Figure 3. Following these, we discover that the average load profiles for the USA and Nigeria are similar to that in Figure 3. This depicts that occupants of households in both countries tend




Fig. 2 Average monthly temperature in Texas and Lagos




Fig. 3 The average load profile computed from our representative dataset. It shows that consumption reduces overnight, increases through the day and peaks in the evening.

to consume less during hours of the night (i.e., from 12AM to 6AM). It also suggests an increase in activities that require electricity in the mornings (i.e., from 7AM onward). In addition, it suggests that consumption peaks between 6PM and 8PM in both countries. Given this similarity, we aggregate the appliance consumption data of our representative dataset for each household. This makes up the overall household electricity consumption of the households. In so doing, we end up with a dataset that is representative of the household consumption of 367 typical Nigerian homes. Note that our approach to the development of this dataset is similar to that in [12], where demographics data was used for data transformation. Having described how we created a representative home consumption dataset for Nigeria, we will use it to formally model individual homes as agents, and to express a generic model of the FLSP. First, we formally model individual homes as agents in the next section.

3.2 Modelling Households as Agents

In this section, we model each household as an agent using the household consumption data developed above. We take the following steps in doing this. First, we collect the actual hourly consumption data for each household for up to four previous weeks, if the data is available. In this regard, we end up with a vector of 168 (i.e., from 24 hours \times 7 days) values for each week's worth

of data. We define this vector as $C_i^w = (c_i^{t=1}, \dots, c_i^{t=168})$, where $c_i \in \mathbb{R}_{>0}$ is the consumption of household i at the hour (t) during the week (w). Although no vector is available during the first week, C_i^1 becomes available after the first week (i.e., on the second week). In the same manner, (C_i^1, C_i^2) , (C_i^1, C_i^2, C_i^3) and $(C_i^1, C_i^2, C_i^3, C_i^4)$ become available on the third, fourth and fifth weeks respectively. From this point on, four vectors will provide a moving window over four weeks of data for each household. For example, the vectors $(C_i^4, C_i^5, C_i^6, C_i^7)$ will make up the consumption data over four weeks on the eighth week. Note that we consider weekly periods based on the assumption that the correlation between the electricity consumed on the same day of a week over weeks (e.g., between Sundays) is likely to be higher than those between different days of the week (e.g., between Sundays and Mondays). As such, the consumption patterns of a typical household will likely differ on different days of the week [52, 10], as their activities may differ and be particular to these days. Consequently, when computing the consumption profile of a household, it becomes necessary to consider its typical consumption pattern during each day of the week [52]. Figure 4(A) shows an example of the data collected for a household over all 168 hours on the fifth week.

Second, we model each household's consumption profile using the electricity it consumes in each hour of the week. We do this by computing the average hourly consumption and the standard deviation from this average over all 168 data points in the available vectors. Then, for each household i , the consumption profile is modelled as a vector, ζ_i , of values drawn from the normal distribution of the mean and variance as shown in Equation 1.⁸

$$\zeta_i = (\mathcal{N}(\mu_i^{t=1}, \sigma_i^{t=1}), \dots, \mathcal{N}(\mu_i^{t=168}, \sigma_i^{t=168})) \quad (1)$$

where ζ_i is the vector of the hourly consumption of i over the week, with the consumption for each hour drawn from a normal distribution of mean μ_i^t (i.e., $\mu_i^t = \sum_{w=1}^4 c_i^{w,t}$) and variance σ_i^t (i.e., $\sigma_i^t = (\sum_{w=1}^4 (c_i^{w,t} - \mu_i^t)^2)/4$). We show the consumption profile of our example household in Figure 4(B). Note that, as a result of replacing the oldest vector used in calculating the 168 hourly averages with the newest (when available), we are able to capture any changes in the consumption patterns of a household. In other words, our vector of averages take into account the effects of changes in season, appliance usage or habits.

Finally, we normalize the vector ζ_i , so that the consumption profiles of all households fall within the range $(\epsilon, 1)$, where $\epsilon \in \mathbb{R}_{>0}$ is a very small number. With this, we create a vector of *comfort* for each household, and finalize the modelling of households as agents. We define this comfort vector, Δ_i , for an agent i , as:

$$\Delta_i = \frac{\zeta_i}{\max_t \{\zeta_i\}} = (\delta_i^{t=1}, \dots, \delta_i^{t=168}) \quad (2)$$

⁸ By drawing from this normal distribution of our hourly consumption values over weekly periods, we may end up with a negative value of consumption. To avoid this (as it is unrealistic to have a negative consumption), we take an absolute value of any negative consumption value that is drawn from the distribution.

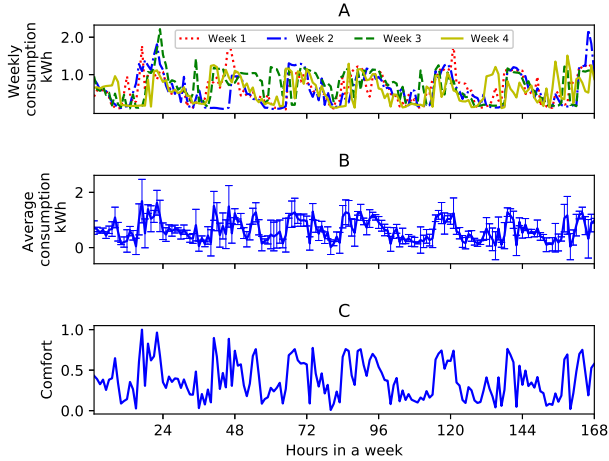


Fig. 4 Weekly consumption data of four prior weeks (A), hourly averages and errors of four-week data (B), and comfort (C).

where Δ_i is a vector of values $\delta_i^t \in \mathbb{R}_{>0}$. Consequently, Δ_i represents the preference agent i has for electricity during all hours in a week. A value close to 1 within the vector represents an hour during which an agent has a high need for electricity in the week. On the other hand, a value close to ϵ reflects the agent's low need for electricity at that hour of the week. The comfort profile of our example agent is shown in Figure 4(C).

Our formulation provides two benefits. (i) It represents a vector of utilities⁹ of each agent in terms of the electricity needs of agents. This vector is formulated in such a way that the utility of an agent is the highest during the hour it needs electricity the most, and the lowest during the hour it needs electricity the least. As such, an agent receives higher comfort utilities if it is connected to electricity at hours it needs more electricity. (ii) With this formulation, the electricity needs of agents can be uniquely quantified and interpersonally compared¹⁰ at the same time, without considering how much electricity the agents consume with respect to others. For example, values of 1 within the comfort vectors of two different agents represent the hours in a week that both agents need the most electricity. It is worthy of mention that we use the term “comfort” to qualify this vector because, the more an agent is connected to supply during hours with higher δ^t values, the more it benefits (or derives comfort) from electricity.

Note that our comfort formulation: (i) is an independent, linear utility function for each agent, (ii) is one that is relevant to developing countries

⁹ The utility of an agent is a numerical value that is used to represent the preference of the agent for the resource, such that agent i receives a utility of $u_i(P)$ from an allocation P [7]. This utility is a mapping from the consumption of agents.

¹⁰ Comparison between agents is possible because all agents' comfort values are on the same scale of ϵ to 1, not minding how much the electricity they demand (or consume) differ.

where, in reality, there is limited information (such as appliance-level consumption, internal and external temperature data, occupancy information from sensors etc.) that can be used to formulate comfort with more sophisticated approaches, and (iii) can be used to design our fair load shedding solutions to provide results similar to those of incentive compatible mechanisms.¹¹ In making our solutions incentive compatible, agents can receive their highest utilities by consuming electricity in their usual manner.¹² We express the FLSP in the next section.

3.3 The Fair Load Shedding Problem (FLSP)

Herein, we formally define the load shedding problem, based on the data and comfort model above. Our definition will be applicable to the rest of this paper.

We define I as a set of n agents, where each agent is denoted as i . Also, we derive the hourly estimated consumption (or demand) of each agent, \tilde{c}_i^t , at hour (t) from the representative data, as would be necessary when planning load shedding ahead. We do so by drawing from the normal distribution $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 0.05)$.¹³ The aggregated hourly estimated demand of the set of agents represents the hourly load on the system. We denote this hourly load as $l^t \in \mathbb{R}_{>0}$, where $l^t = \sum_{i=1}^n \tilde{c}_i^t$. Similarly, the hourly estimated supply capacity available for meeting the demand of agents in I is represented as $g^t \in \mathbb{R}_{>0}$.¹⁴ Now, in a developing country, it is often the case that l^t is greater than g^t . In this event, there is a deficit, d_t (i.e., $d_t = l^t - g^t$), on the system and the demand of all agents cannot be met. System operators then have to resort to load shedding in order to maintain a balance between demand (l^t) and supply (g^t) and keep the system in operation (as discussed in Chapter 1). In executing load shedding, we define a piece-wise variable A_i^t , which takes the value 1 if i is connected to electricity at t , and 0 otherwise. As such, an agent is either connected to supply or not. Now, we briefly discuss the assumptions upon which our pre-planned solutions are based.

3.4 Key Assumptions

In solving the fair load shedding problem, we use day-ahead hourly estimates of demand and supply to plan for load shedding a day ahead. Our solutions result in the selection of households to be disconnected from supply in a way that ensures fairness and grid stability. As such, we classify our solutions as

¹¹ Incentive compatible mechanisms are known to reward agents which act according to their preferences with the best outcomes for such actions [43].

¹² We design our fair load shedding solutions in this manner in Section 5.

¹³ We state later in Section 3.4 that we use accurate predictions of household consumption to solve the FLSP. For this reason, we specifically take σ as 0.05 to produce consumption estimates that are close to the actual consumption of homes. However, we evaluate our solutions under higher levels of uncertainty in Section 6.4.

¹⁴ For our purposes, we take the value of g^t for each day ahead as the average of the sum of hourly household consumption estimates for that day (i.e., $g^t = (\sum_{i=1}^{24} \sum_{i=1}^n \tilde{c}_i^t)/24$).

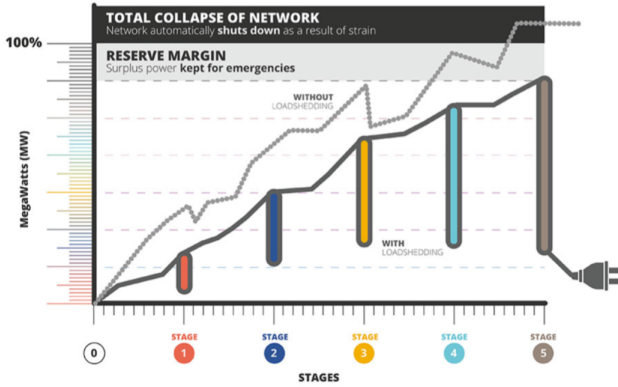


Fig. 5 Load shedding as implemented by the city of Johannesburg, indicating there is emergency reserve power [8]

planned, day-ahead, household-level load shedding options. We make the following assumptions in solving our problem in this manner:

1. **Household-level load control:** We assume that the retrofits which provide the means for household-level load control are generally available in homes within distribution networks in developing countries. This is necessary because our solutions warrant the execution of load shedding at the household level. We can make this assumption because household-level load control is possible with the smart meter retrofits discussed in Section 2.
2. **Household consumption estimates:** We assume that we receive (or compute) near-accurate estimates from homes (hence the level of uncertainty in deriving \hat{c}_i^t). We can make this assumption because it is possible to receive (or estimate) near-accurate household consumption estimates using elicitation approaches such as prediction-of-use games (see [54]) and scoring rule-based mechanism (see [42]), and with the tools (such as in [44]) and techniques (such as in [50]) for estimating household demand.
3. **Spinning reserve:** We assume that the sum of available supply and emergency power will always suffice for any uncertainties in demand estimates. We can make this assumption because generators ensure that they have a spinning reserve available, as they are unable to always predict consumption. This is the case in South Africa, another developing country, as depicted by the “reserve margin” shown in Figure 5. In addition, while some households may consume more than their estimates, others may consume less than their estimates. Such estimation errors may cancel themselves out. Furthermore, we assume that the spinning reserve will cater for any power flow concerns (e.g., transmission losses).
4. **Comfort vector:** We assume that comfort is independent of load shedding events. Although we design solutions which are able to factor in the effect of load shedding on consumption after supply is restored (i.e., the rebound

effect)¹⁵ in Section 5.2.2 and Section 5.2.3, it is necessary to assume that consumers may not know in advance when load shedding events will happen and so do not preemptively run some activities. As such, our comfort formulation in this paper embodies the preference of agents for electricity, even when they may be arbitrarily subjected to load shedding.

Based on these assumptions, in the next section, we first develop a number of heuristics that disconnect households from supply during load shedding. Our heuristics are designed with the objective to connect agents to supply for a similar number of hours, while ensuring that demand matches supply. We also assess the performance of these heuristics and analyze their results therein.

4 Household-Level Load Shedding Heuristics

In this section, we briefly summarize and evaluate the solutions presented in [33].¹⁶ In [33], four heuristics (i.e., **GA**, **CSA1**, **RSA** and **CSA2**) which focus on the single objective to minimize the maximum difference in the number of hours agents are connected to supply (i.e., to minimize envy in terms of connections as specified in Section 3.3) were presented. This objective is formulated as:

$$\min\{\max_{i,j}(|N_i - N_j|)\} \quad (3)$$

Specifically, **GA** tries to simulate the response of an operator to load shedding. It does so by creating a few groups of agents, such that the sum of the estimated demand of the agents in each group is enough to offset the deficit. Thereafter, it sums up the number of hours all agents in each group have been connected to supply, then disconnects the group of agents that have been connected to supply the most. The other heuristics (i.e., the **CSA1**, **RSA** and **CSA2**) use a round-robin scheme to select agents in rounds over all load shedding events. In this manner, when an agent is disconnected from supply, it will not be disconnected again until all agents have been disconnected in that round.¹⁷ The **CSA1** disconnects agents in each round based on their consumption, the **CSA2** disconnects agents in each round based on their comfort, and the **RSA** is designed to be agnostic of the consumption or comfort of agents, so that it randomly disconnects agents in each round.

In order to re-assess¹⁸ the performance of these heuristics, we point out that q is 2184 hours (as the representative dataset of hourly consumption

¹⁵ The rebound effect, illustrated in [39], is the reaction of consumers to load management measures by shifting their consumption to periods when electricity is more available.

¹⁶ Please see [33] for a more in-depth description of these solutions.

¹⁷ A round is completed when each agent has been disconnected once. Depending on the deficits, a round may either span over different load shedding implementations or terminate within the implementation it begins.

¹⁸ The performance of the heuristics were assessed in terms of connections and comfort in [33], as in a chore division setting. However, to conform with the results to be evaluated later in this paper, we reproduce the results in [33] in terms of connections and comfort, as in a cake-cutting setting.

Table 3 Utilitarian, egalitarian and envy-freeness results in terms of the number of hours agents are connected to supply

Heuristic	Utilitarian	Egalitarian	Envy-freeness
GA	631498	1607	697
CSA1	646971	1766	1
RSA	640233	1758	1
CSA2	640651	1745	1

developed in Section 3 is for a period of 13 weeks). In evaluating these heuristics, we use the utilitarian, egalitarian and envy-freeness metrics to evaluate these heuristics. The utilitarian criterion, as defined in [23], is the sum of the individual utilities of all agents in a system. We adopt this definition, and consequently adopt the utilitarian criterion as the total number of hours all n agents are connected to electricity within q hours (i.e., $\sum_{i=1}^n N_i$). In addition, [23] also defines the egalitarian criterion as the utility of the agent that is currently worst off. In our domain, we adopt the egalitarian criterion as the number of hours the agent connected to electricity the least within q hours was connected for (i.e., $\min_i\{N_i\}$). Finally, envy-freeness is a criterion of fair division that allocates resources to agents in such a way that no agent envies the allocation of another. However, we define this criterion differently. This is because, in our case, agents do not have information of the allocation of others, making it impossible for the standard definition of envy-freeness to be adopted. Instead, we define envy-freeness in terms of the highest difference between the utilities of all pairs of agents, such that if the agents were aware of their utilities, the overall envy within the agent-population will be minimal. As such, the envy-freeness criterion is defined as $\max_{i,j}\{|N_i - N_j|\}$, representing the maximum difference between the number of hours all pairs of agents are connected to supply for within q . We present the results based on these social welfare metrics in Table 3.¹⁹

Table 3 shows how the four heuristics achieve a somewhat similar performance in terms of how they connect agents to supply over the entire period that results from the data, as seen in the Utilitarian column. However, the performance of CSA1 is better than those of others. The same can be said of results shown in the Egalitarian column, where CSA1 connects the least connected agent to supply the highest number of times. However, in the Envy-freeness column,²⁰ CSA1, RSA and CSA2 all achieve a difference of one connection between the agents they connect the most and least to supply. This is due to the round-robin scheme they employ.

Nonetheless, although three of our algorithms achieve their objective, some questions arise. What if agent i is mostly connected to supply when the occupants are home (and, as such, need electricity the more) but agent j is mostly connected to supply when they are away (and, as such, need electricity the less)? Would an agent prefer to be connected to supply when it needs electric-

¹⁹ Note that the results presented in Table 3 are obtained from solving the fair load shedding problem once.

²⁰ Note that the smaller the number is, the better the result in this column.

ity the more, even if it ends up being connected to supply for a fewer hours? We believe the fairness scenarios posed by these questions are not addressed by these heuristics. Therefore, to better achieve fairness, it is important that we model the FLSP as a resource allocation problem that considers the electricity needs of agents. We do this in the next section.

5 Optimizing Fair Load Shedding

With the notion of comfort, we now have a number of objectives to work with, vis-à-vis the number of hours agents are connected for, the comfort delivered to agents and the demand of agents. Upon this background, in this section, we use the consumption and comfort values of agents to model the FLSP as a constrained optimization problem.²¹ In so doing, we adapt the utilitarian, egalitarian and envy-freeness social welfare metrics in terms of comfort, demand and connections as objectives and constraints, and use these to design an MIP that allocates electricity to individual households over a period.

5.1 The Knapsack MIP Formulation

We first model the objective of the fair load shedding problem with insights gained from the knapsack problem for an hour (see Section 2). We take the values of the items in a knapsack problem to be the comfort values of the agents. In addition, we take the weights of these items to be the consumption (or demand) of each agent. As such, the capacity of the knapsack is taken as the supply capacity of the grid. With these, we formulate FLSP for an hour, t , as:

$$\begin{aligned} \max \quad & \sum_{i=1}^n \delta_i^t A_i^t, \\ \text{s.t.} \quad & \sum_{i=1}^n \tilde{c}_i^t A_i^t \leq g^t \end{aligned} \tag{4}$$

where δ_i^t is the comfort value of i at t and A_i^t connects i to supply at t (i.e., $A_i^t = 1$) or otherwise. This is an implementation of the utilitarian criterion in terms of comfort, in that the summed utility of all agents is maximized [29]. It is subjected to the constraint that limits the total electricity supplied to the available supply capacity, thereby implementing load shedding. With this formulation, the comfort of all agents is maximized at t while keeping the electricity supplied within the supply capacity. As such, the MIP ensures that

²¹ We believe the consumption and comfort values of agents provide the tools upon which our solutions can become an algorithmic framework and methodology which may be generalized to other problems involving constrained utility maximization, where the utilities can be suitably modelled (as we did with comfort) and where the constraints involve notions of fairness.

electricity is supplied to agents that need it more at t . However, this in no way considers fairness. Instead, fairness considerations can only be incorporated into this problem when it is solved over a number of hours [13]. For this reason, we extend this into a Multiple Knapsack Problem (MKP) formulation in the next section.

5.2 The MKP MIP Formulation

Herein, each hour is modelled as a knapsack in the MKP formulation. With each agent having different demands and different comfort values for different hours in a period made up of p hours, the formulation in Equation 4 becomes:

$$\begin{aligned} \max \quad & \sum_{t=1}^p \sum_{i=1}^n \delta_i^t A_i^t & (01), \\ \text{s.t.} \quad & \sum_{i=1}^n \tilde{c}_i^t A_i^t \leq g^t \quad \forall t \in \{1, \dots, p\} & (C1), \\ & \beta_2 \geq \sum_{t=1}^p A_i^t \geq \beta_1 \quad \forall i \in I & (C2), \quad (5) \\ & \sum_{t=1}^p \delta_i^t A_i^t \geq \beta_3 \quad \forall i \in I & (C3), \\ & \sum_{t=1}^p \tilde{c}_i^t A_i^t \geq \beta_4 \quad \forall i \in I & (C4) \end{aligned}$$

As such, the comfort of all agents is maximized over p^{22} hours. The objective is also an implementation of the utilitarian social welfare metric that is subjected to a number of constraints. Constraint C1 ensures that the supply capacity is not exceeded for every hourly solution to the MIP problem. C2 to C4 are constraints which we formulate to fulfill some fairness criteria using the egalitarian and envy-freeness metrics. We describe these criteria and constraints in the sections that follow.

5.2.1 Fairness Criterion Based on Number of Hours of Connection

We begin by discussing a fairness criterion based on the number of hours individual agents are connected to supply (one of the utilities of agents), as was the focus of our heuristics in Section 4. This criteria is formulated as a constraint, C2. The constraint is constructed using the egalitarian and envy-freeness metrics.

The egalitarian metric is described as the number of hours which the agent least connected to supply is connected for. As such, we constrain our FLSP

²² It is noteworthy that, because we look to plan for load shedding a day ahead, we take p to be 24 (hours).

Algorithm 1: Deriving an upper bound and a lower bound for the number of hours all agents should be connected to supply the day ahead

Data: The number of agents in I , n ; the hourly consumption of each agent, \tilde{c}_i^t ; the hourly supply capacity, g^t ; the hourly deficit, $d^t = l^t - g^t$

```

1  $F \subseteq t \in \{1, \dots, 24\} \forall d^t > 0$  // Hours load shedding is necessary,  $F$ 
2  $c_\mu^f = \left( \sum_{i=1}^n \tilde{c}_i^f \right) / n$  // Average hourly demand of agents in  $f \in F$ 
3  $\pi^f = g^t / c_\mu^f$  // Estimated number of households connected in  $F$ 
4  $\Pi = \sum_{f \in F} \pi^f$  // Total estimated households connected hourly in  $F$ 
5  $N_\mu^i = \Pi / n$  // Estimated average connections per agent in  $F$ 
6  $\beta_1 = \lfloor N_\mu^i \rfloor + 24 - |F|$  // Round down average day ahead connections per agent
7  $\beta_2 = \lceil N_\mu^i \rceil + 24 - |F|$  // Round up average day ahead connections per agent
```

to ensure that there is a lower bound, $\beta_1 \in \{0, \dots, 24\}$, to the number of hours every agent will be connected to supply within 24 hours. The egalitarian criterion can be satisfied in terms of connections in that a high value for β_1 will ensure all agents are connected to supply for a minimum number of hours within the day ahead.

In order to satisfy the envy-freeness criterion, we specify an upper bound, $\beta_2 \in \{0, \dots, 24\}$, on the number of hours every agent will be connected to supply within 24 hours. Thereupon, we are able to limit the differences between the number of hours all pairs of agents are connected to supply within 24 hours. Furthermore, we use the hourly supply capacity (g^t) and the demand of all agents ($\sum_{i=1}^n \tilde{c}_i^t$) for each hour within the day-ahead to derive the values for parameters β_1 and β_2 . We do this using the set of computations shown in Algorithm 1.

Algorithm 1 first finds the hours within the day ahead in which load shedding is necessary (Line 1). For each of these hours, it then computes the average hourly estimated demand of agents (Line 2). Thereafter, it divides the hourly supply capacity by the average hourly estimated demand of agents during load shedding hours. This provides what we say is an estimate of the total number of agents which can be connected to electricity in these hours (Line 3). Next, it sums up these hourly estimates (Line 4). Thereafter, it divides this total number of connections by the total number of agents, to obtain an average number of hours each agent can be connected to supply during load shedding hours (Line 5). To compute an exact number of whole hours agents can be connected to supply in the day ahead, it rounds down (Line 6) and up (Line 7) the number computed in Line 5 plus the number of hours no load shedding is necessary.

β_1 and β_2 are used in C2 to determine the number of hours which agents will be connected to supply the day ahead. It should be noted that the values of these parameters are absolutely dependent on the data used in solving the FLSP (i.e., the consumption of agents and the supply capacity from our representative dataset). However, the above steps have generated parameters that result in feasible solutions in this case and can be used in different scenarios

(i.e., for different datasets or system characteristics) as we show in Section 6.6. We discuss C3 and its associated fairness criterion in the next section.

5.2.2 Fairness Criteria Based on Comfort

Next, we consider the comfort of each agent over each day's period (another of the utilities of agents, as we have previously mentioned). We use this consideration in formulating another constraint, C3, based on the egalitarian and envy-freeness metrics. However, we do not have a comfort capacity for computing bounds (as we did in Section 5.2.1, where bounds which determine the number of hours that agents will be connected to supply daily are derived from supply capacity). In addition, we do not have a single yardstick that we can equally base the comfort of all agents on (as we did in Section 5.2.1, where all agents had the same 24 hours in which they can be connected to supply). Furthermore, the comfort of each agent for each hour is a function of their consumption over 168 hours of a week. By this, we mean that the comfort enjoyed by two agents which consume the same amount of electricity in a day is unlikely to be the same because we derive their comfort from their weekly historical consumption data (as described in a Section 3.2). For this reason, we do not have the adequate information to analytically determine lower and upper bounds for implementing the egalitarian and envy-freeness metrics in terms of comfort.

To remedy this, we define a factor (or parameter), β_3 , which represents the percentage of the summed comfort of each agent within a day as $\beta_3 = \alpha_3 \sum_{t=1}^p \delta_i^t \quad \forall \quad i \in I$. We set β_3 as a lower bound of the comfort that must be delivered to every agent daily. We do not select an upper bound because the constraint C2 already ensures the solution to the MIP does not result in any agent being connected to supply all hours of the day. As such, there is no day an agent enjoys all of its summed comfort.²³ To this end, our MIP produces solutions that satisfy the egalitarian and envy-freeness metrics as a result of the lower bound. With respect to the egalitarian metric, the lower bound ensures all agents are delivered a minimum level of their comfort daily. Whereas, the envy-freeness metric is satisfied with respect to the value of parameter α_3 which limits the difference between the comfort enjoyed by all agents.

In arriving at a value for α_3 , we provide a grid of values in the range $0 < \alpha_3 < 1$ to our solver.²⁴ Using a binary search algorithm, we find values which result in feasible solutions within the range. From these values, we then select a value, α_3 , which maximizes the solution to the MIP.

In addition, we attempt to factor in the rebound effect (described in Section 3.4) herein. It is necessary to do so because if i is disconnected from supply at t , it is likely to need electricity more (than it would have if it was not

²³ Although this is dependent on data, our representative dataset (in Section 3.1) and problem formulation (in Section 3.3) are such load shedding will be necessary during a number of hours daily and no agent is connected to supply all hours of the day.

²⁴ We use the standard mixed integer linear programming solver CPLEX to solve this constrained optimization problem.

disconnected) in the next hour, $t + 1$. This is because some activities which consumers may have been deprived of at t are likely to become more important at $t + 1$. In this regard, we take δ_i^{t+1} as a value randomly selected between the computed comfort value for that hour and the maximum comfort value of the week (i.e., $\delta_i^{t+1} = \{\delta_i^{t+1}, 1\}$). We discuss an additional fairness criterion, modelled as constraint **C4**, in the next section.

5.2.3 Fairness Criteria Based on Electricity Supply

We begin by pointing out that if an agent is delivered a certain level of its summed comfort over a particular day, there is no guarantee that the agent will be supplied the same level of its total demand that day. This is because comfort is derived from historical consumption while demand is estimated. In addition, as stated in Section 5.2.2, comfort is derived from weekly consumption. Therefore, it is necessary to also formulate another constraint, **C4**, that delivers a minimum percentage of each agent's daily demand (the third of the utilities of agents). We formulate this lower bound as a percentage because all agents have different demands over different hours such that no single amount of electricity supply will equally satisfy them. In addition, as in Section 5.2.2, we do not have enough information from which we can analytically determine this lower bound. Likewise, we do not compute an upper bound because our constraint in Section 5.2.1 already ensures that no agent is connected to supply all 24 hours of a day as an effect of load shedding.

A lower bound presents a basis for which our MIP formulation can conform to the egalitarian and envy-freeness metrics. With respect to the egalitarian metric, the lower bound ensures all agents are supplied a minimum level of their demand daily. It also determines the maximum difference between the pairwise percentages of the daily total demand that is supplied to the agents, such that the envy-freeness metric is satisfied to a degree determined by the lower bound. We define this lower bound as $\beta_4 = \alpha_4 \sum_{i=1}^p \tilde{c}_i^t \quad \forall i \in I$. On this ground, our FLSP is subjected to the constraint **C4** in Equation 5. As such, a minimum level of demand is supplied to each agent daily and the differences between the supply shares of all agents are reduced.

Similar to parameter α_3 , our solver arrives at a value for parameter α_4 within the range, $0 < \alpha_4 < 1$, such that the solution to the MIP is maximized. Also, α_4 determines the value of β_4 (see above), and depends on the demand and supply capacity of the system. We attempt to factor in the rebound effect herein also, because if i is disconnected from supply at t , it is likely to consume electricity more (than it would have if it was not disconnected) in the next hour, $t + 1$. This is due to appliances like refrigerators cycling on and home occupants possibly running activities they were deprived of at t . In this regard, we take \tilde{c}_i^{t+1} as a value randomly selected between the computed consumption value for that hour and the maximum consumption value in its consumption profile (i.e., $\tilde{c}_i^{t+1} = \{\tilde{c}_i^{t+1}, \max_t \{\zeta_i\}\}$).

When taken together, the solution of the FLSP MIP selects the agents to be connected to supply at each hour within a day, such that comfort is maximized over the day. It ensures that the electricity supplied to agents

(being a function of comfort) is as high as possible without exceeding the hourly supply capacity. It also ensures that the number of hours each agent is connected to supply within the day, the percentage of the daily total comfort of each agent and the percentage of the daily total consumption of each agent all satisfy the egalitarian and envy-freeness metrics.

5.3 Maximizing Supply: An Optional MIP Objective

Next, we consider the FLSP in the context of revenue maximization (instead of comfort maximization). Hence, we formulate an objective which maximizes supply below:

$$\max \sum_{t=1}^p \sum_{i=1}^n \tilde{c}_i^t A_i^t \quad (02) \quad (6)$$

The objective is likewise subjected to all constraints C1 to C4 to form another solution to the FLSP.

Consequently, we offer a pair of solutions to the FLSP. These solutions present options that can be utilized in different conditions and environments, depending on the desired objectives and requirements. In highlighting these objectives and requirements, we describe these solutions as:

1. **The Comfort Model (CM)**: We present Equation 5 as CM, being a solution that maximizes the comfort objective (or utility). In so doing, we think of an environment which places a premium on supplying electricity to households based on their needs. The key objective within this environment would be to maximize comfort. A consideration may be that, if a household has access to electricity when it needs it more, the household is likely to maintain its consumption patterns. As a consequence, the feasibility of day-ahead fair load shedding schemes is increased.
2. **The Supply Model (SM)**: Likewise, we present Equation 6 as the objective of SM, being a solution that maximizes the supply objective (or utility). In so doing, we think of an environment that places a premium on maximizing revenue. The objective within the environment would be to maximize the access of households to electricity. A consideration may be that a scheme such as this will result in the least waste and the highest revenue, as we highlighted above.

We highlight that since both comfort (as in CM) and demand (as in SM) are related, each of the models above maximizes the objective of the other to an extent. This is the sort of compromise that exists when solving multi-objective problems. In addition, while both of the models are individually less complex than a multi-objective model, they may produce results which Pareto dominated those of their multi-objective counterpart. We evaluate the results of these constrained optimization models in the section that follows.

Table 4 Results of MIP models and heuristics in terms of hours of connection to supply on the average, along with their standard deviations (SD) within parenthesis

Solution	Utilitarian (SD)	Egalitarian (SD)	Envy-freeness (SD)
CM	717031 (3950)	1920 (3.24)	123 (2.09)
SM	709676 (3878)	1922 (3.41)	71 (2.04)
GA	629534 (4178)	1609 (4.69)	695 (3.28)
CSA1	647439 (3063)	1764 (2.27)	1 (0.00)
RSA	643504 (4094)	1753 (4.33)	1 (0.00)
CSA2	641002 (3154)	1746 (2.38)	1 (0.00)

6 Evaluation of Results

We assess the performance of our solutions in this section. In the case of our MIP solutions, as previously stated, parameters α_3 and α_4 are derived using grid search within the ranges $0 \leq \alpha_3 \leq 1$ and $0 \leq \alpha_4 \leq 1$ respectively. This results in parameters the $\alpha_3 = 0.78$; $\alpha_4 = 0.75$ for **CM**, and $\alpha_3 = 0.8$; $\alpha_4 = 0.8$ for **SM**. On this ground, **CM** aims to supply 78% of agents' daily comfort and 75% of agents' daily demand to the agents daily. Furthermore, **SM** is designed to supply 80% of agents' daily comfort and 80% of agents' daily demand to the agents daily. In addition, β_1 and β_2 are 21 and 22 respectively within both models. We run nine independent simulations of our models to consider how they perform on the average. We present these average results under different experiments in the sections that follow.

6.1 Fairness and Efficiency in terms of Connections

In this section, we evaluate the performance of all our load shedding solutions with respect to the number of hours they connect agents to supply individually (i.e., the egalitarian and envy-freeness social welfare metrics) and collectively (i.e., the utilitarian social welfare metric) on the average herein over q (i.e., 2184) hours. We present the average results obtained by each solution after a number of independent implementations, along with their standard deviations (in parenthesis) in Table 4.

Within the Utilitarian column in Table 4, we see that **CM** connects the entire population of agents to supply more hours than the other solutions on the average. Its performance is 1.04% better than that of **SM**, which performs second best. The performances of the heuristics under this consideration are in line with our discussions in Section 4, but we highlight that the best performing heuristic under this consideration (i.e., **CSA1**) connects agents to electric supply 10.75% less than **CM** on the average. Whereas, **RSA**, **CSA2** and **GA** obtain results that are 11.43%, 11.86% and 13.90% worse than that of **CM** respectively

As seen in the Egalitarian column, **CM** and **SM** result in every individual agent being connected to supply for longer periods. As such, they connect the worst off agents to supply for longer than the heuristics do. This is so because of constraint, **C2**, which ensures every agent is connected to supply for at least 21 hours daily.

Furthermore, because of the round-robin scheme utilized by **CSA1**, **RSA** and **CSA2**, they connect agents to supply with pairwise differences that are lower than those of **CM** and **SM**, as shown within the envy-freeness column. However, the effect of this is that **CSA1**, **RSA** and **CSA2** connect the agents most connected to supply for 1765, 1754 and 1747 hours respectively (i.e., the egalitarian value added to the envy-freeness value). These are all much lower than the number of hours which the MIP models connect the worst off agents to supply (1920 for **CM** and 1922 for **SM**). On account of this, although the MIP models result in the difference between the utilities (in terms of hours of connection) obtained by agents being higher, they also result in all agents obtaining higher utilities all round (i.e., they Pareto dominate the heuristics under this consideration). Next, we evaluate the performance of our solutions in terms of comfort.

6.2 Fairness and Efficiency in terms of Comfort

We developed the notion of comfort in order to consider the preference (or electricity needs) of agents when allocating electricity. Now, we evaluate the results obtained by all our load shedding solutions using the utilitarian, egalitarian and envy-freeness social welfare metrics. On the one hand, the results obtained by our MIP model under the utilitarian metric can be directly presented in this evaluation, being an addition of the comfort delivered to all agents during q hours. On the other hand, their results under the egalitarian and envy-freeness metrics are dependent on the comfort delivered to individual agents in q hours. These results should not be presented without first developing a unifying scale for all agents, since the total comfort values of agents within q hours are different.²⁵ It was why we defined β_3 as a lower bound of percentages of each agent's summed comfort to be delivered to the agent within a period in Section 5.2.2, as opposed to specifying a single value of comfort to be delivered to all agents. Therefore, we provide a unifying scale for assessing the results of our solutions. In so doing, we define the *Agent Comfort Share* (ACS), δ_i^* , as the total comfort enjoyed by each agent in q hours, divided by its total comfort over q hours (i.e., $\delta_i^* = \sum_{t=1}^q \delta_i^t A_i^t / \sum_{t=1}^q \delta_i^t$). To this end, we present the results obtained by each solution after a number of independent implementations, along with their standard deviations (in parenthesis) in Table 5.

The Utilitarian column shows the sum of the comfort enjoyed by all agents within q hours. Of all fair load shedding solutions, **CM** delivers the maximum comfort to all agents under this metric. This is as expected, with the objective of the model being to maximize comfort. The next best performing solution in this category is **SM**, with **CSA1** and **GA** also performing well.

It is under the egalitarian and envy-freeness metrics that our MIP solutions are expected to obtain results which are in line with their purpose. We begin with the egalitarian metric, where the Egalitarian column displays the ACSs of the agents with the minimum ACSs. Table 5 shows how **SM** delivers the highest

²⁵ As such, any solution that delivers an equal amount of comfort to all agents does not necessarily distribute comfort evenly.

Table 5 Results of MIP models and heuristics in terms of comfort delivered on the average, along with their standard deviations (SD) within parenthesis

Solution	Utilitarian (SD)	Egalitarian (SD)	Envy-freeness (SD)
CM	303217 (3447)	0.81 (0.01)	0.13 (0.02)
SM	292135 (3802)	0.83 (0.01)	0.09 (0.02)
GA	291021 (5198)	0.38 (0.04)	0.56 (0.03)
CSA1	291909 (3201)	0.67 (0.02)	0.25 (0.02)
RSA	268564 (5106)	0.65 (0.04)	0.28 (0.03)
CSA2	270262 (3112)	0.64 (0.02)	0.28 (0.02)

ACS to the worst off agent (in terms of their ACSs) after all load shedding periods. In turn, the result obtained by CM is only 2.47% worse than that of SM. It is worse because its comfort constraint factor is lower than that of SM (α_3 is 0.78% for CM and 0.80% for SM). Nonetheless, CM and SM deliver 81% and 83% of the ACSs of the agents with the minimum respectively, so that they both satisfy constraint C3. However, CSA1, which performs better than other heuristics, only delivers 67% of its ACS to the worst off agent. These results show that our MIP solutions both consider and satisfy the individual electricity needs of agents by supplying electricity to them when they need it the most.

Next, we assess the results of the fair load shedding solutions under the envy-freeness metric. Notably, even though CSA1, RSA, and CSA2 heuristics connect agents to electric supply a very even number of hours, they still fail to distribute comfort as evenly as any of the MIP models. Instead, SM evenly distributes comfort the best, with the maximum difference between the ACR delivered to all agents only 0.09. In turn, the best performing heuristic (i.e., CSA1) obtains a result that is 177.78% worse than SM's. Contrariwise, the result obtained by CM, is 5.92% worse than SM's, owing to CM's weaker constraint.

In summary, CM and SM both outperform the four heuristics under the utilitarian metric, deliver more comfort to all agents than any of the heuristics with respect to the egalitarian metric, and distribute comfort more evenly than any of the heuristics, with respect to the envy-freeness metric. We show in the next section how our MIP models compare to the heuristics when considering the amount of electricity supplied to individual agents.

6.3 Fairness and Efficiency in terms of Supply

We assess the performance of all fair load shedding solutions based on the amount of electricity supplied in this section. In the Utilitarian column, we show the sum of electricity (in kWh) supplied to all agents within q hours. However, as with the egalitarian and envy-freeness results presented in Section 6.2, we formulate a basis upon which the electricity supplied to meet the heterogeneous demand of each agent can be compared. This is because the sum of electricity supplied to agents within q hours are different, being the reason for defining β_4 as a lower bound of percentages of each agent's

Table 6 Results of MIP models and heuristics in terms of electricity supplied on the average, along with their standard deviations (SD) within parenthesis

Solution	Utilitarian (SD)	Egalitarian (SD)	Envy-freeness (SD)
CM	1340015 (8299)	0.78 (0.01)	0.17 (0.02)
SM	1347801 (8304)	0.83 (0.01)	0.11 (0.02)
GA	1297020 (11264)	0.35 (0.04)	0.58 (0.03)
CSA1	1296939 (7564)	0.66 (0.02)	0.28 (0.02)
RSA	1344945 (11284)	0.68 (0.03)	0.25 (0.03)
CSA2	1345537 (7388)	0.63 (0.03)	0.30 (0.02)

summed demand. We formulate our basis for comparison as the *Agent Supply Rate* (ASR), c_i^* , which we define as the summed electric demand supplied to each agent in q hours, divided by its total electric demand over q hours (i.e., $c_i^* = \sum_{t=1}^q c_i^t A_i^t / \sum_{t=1}^q c_i^t$). We present the average results obtained by each solution following a number of independent simulations, along with their standard deviations (in parenthesis) in Table 6.

The Utilitarian column in Table 6 shows how SM supplies the most electricity to all agents over q hours. This is expected as the objective of the model is to maximize electric supply. In fact, both MIP models outperform all heuristics under the utilitarian metric. However, the results of the heuristics are between 3.63% (for CSA2) and 3.92% (for CSA1) worse than that of SM on the average.

It is under the other metrics that our MIP solutions are expected to obtain results which are in line with their purpose. We begin with the egalitarian metric, where the Egalitarian column displays the ASSs of the worst off agents (in terms of their ASSs). The column shows that SM delivers the most ASS of all worst off agents, due to its higher supply constraint factor (i.e., α_4). RSA, being the best performing heuristic, performs 22.01% worse than SM. In turn, the result obtained by CM is only 6.41% worse than that of SM. CM performs this worse because its supply constraint factor is much lower than that of SM (α_4 is 0.75% for CM and 0.80% for SM). Nonetheless, CM and SM deliver 78% and 83% of the ASSs of the worst off agents respectively, so that they both satisfy constraint C4. In turn, RSA only delivers 68% ASS to the worst off agent. This shows that our MIP solutions both consider and satisfy the individual demand of agents by supplying electricity to them based on their demand.

Finally, we assess the results of the fair load shedding solutions under the envy-freeness metric. It is noteworthy that although the round-robin heuristics connect agents to electric supply a very even number of hours, they still fail to meet the demand of agents as evenly as any of the MIP models. Instead, SM evenly distributes comfort the best, with the maximum difference between the ASR delivered to all agents only 0.11. In turn, the best performing heuristic, RSA, obtains a result that is 127.27% worse than SM's. Contrariwise, the result obtained by CM is 54.55% worse than SM's, owing to CM's weaker constraint.

In summary, the MIP models outperform the four heuristics under the utilitarian, egalitarian and envy-freeness metrics in terms of electricity supply. These MIP models specifically fulfill their purpose by meeting the demand of

Table 7 MIP constraint factors under different levels of uncertainty in prediction of consumption, and constant supply capacity

		$\pm 3\sigma$	$\pm 2.5\sigma$	$\pm 2\sigma$	$\pm 1.5\sigma$	$\pm \sigma$	$\pm 0.5\sigma$	$\sigma = \pm 0.5$
CM	β_1	21	21	21	21	21	21	21
	β_2	22	22	22	22	22	22	22
	α_3	0.78	0.78	0.78	0.78	0.78	0.78	0.78
	α_4	0.67	0.69	0.70	0.71	0.74	0.75	0.75
SM	β_1	21	21	21	21	21	21	21
	β_2	22	22	22	22	22	22	22
	α_3	0.80	0.80	0.80	0.80	0.80	0.80	0.80
	α_4	0.71	0.73	0.74	0.77	0.79	0.79	0.80

each agent more than any of the heuristics, with respect to the egalitarian metric. They also satisfy the demand of all agents more evenly than any of the heuristics, with respect to the envy-freeness metric. To show how our models react to poorer estimates of demand, we solve the FLSP using the data of consumption with different levels of uncertainty and evaluate the results in the section that follows.

6.4 Implementation with Different Levels of Uncertainty

We have so far evaluated our solutions based on accurate estimates of consumption (where $\sigma = \pm 0.5$). Although we assume that these accurate estimates are available, we yet evaluate their performance under different levels of uncertainties in our estimates of consumption. For this reason, we run simulations using six additional levels of uncertainty, by drawing from each of the normal distributions²⁶ $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 3\sigma)$,²⁷ $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 2.5\sigma)$, $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 2\sigma)$, $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 1.5\sigma)$, $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, \sigma)$ and $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 0.5\sigma)$ for each level of uncertainty. We also keep the supply capacity constant (as in Section 3.3) through all simulations. Note that for all these levels of uncertainty, we compute different constraint factors (i.e., different values for β_1 , β_2 , α_3 and α_4) for our MIPs, which we show in Table 7. Now, we briefly evaluate the performance of our solutions under these different levels of uncertainty with respect to how they connect agents to supply, the comfort they deliver, and the electricity they supply to agents against their actual demand.

6.4.1 Connections to Supply under Different Levels of Uncertainty

We display the average number of hours our solutions connect agents to supply individually (with respect to the egalitarian and envy-freeness metrics) and collectively (with respect to the utilitarian metric) under different levels of uncertainty in Figure 6.

²⁶ Note that whenever we draw a negative value of consumption from the distribution, we take an absolute value of that value.

²⁷ This is our highest level of uncertainty, being that we expect 99.73% of consumption values to lie within three standard deviations from c_i^t .

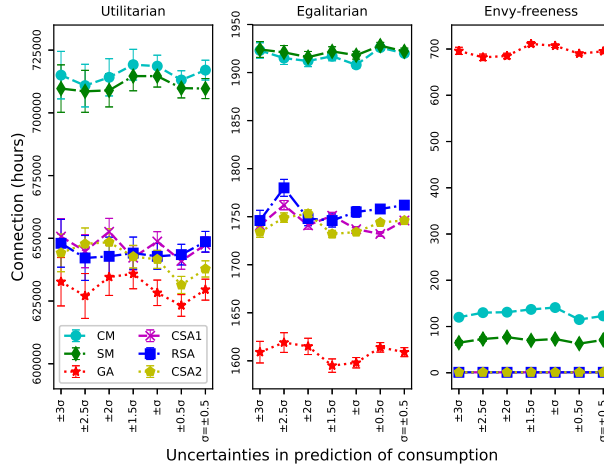


Fig. 6 Average number of connections to supply under different levels of uncertainty

Figure 6 shows how our solutions produce results that are increasingly inconsistent as more uncertainties are introduced into consumption. However, the number of agents they connect to supply does not show any consistent pattern. This is because the errors in consumption (predicted during our simulations) seem to cancel themselves out, such that the average number of agents our solutions connect to supply do not change in any pattern over all levels of uncertainty. These hold under the utilitarian, egalitarian and envy-freeness considerations.

6.4.2 Comfort Delivered under Different Levels of Uncertainty

We display the average comfort our solutions deliver to agents individually (with respect to the egalitarian and envy-freeness metrics) and collectively (with respect to the utilitarian metric) under different levels of uncertainty in Figure 7.

As in Section 6.4.1, Figure 7 shows how our solutions produce results that are increasingly inconsistent as more uncertainties are introduced into consumption, due to more varied consumption data being generated under higher levels of uncertainty. This holds under the utilitarian, egalitarian and envy-freeness considerations. Furthermore, the comfort our solutions deliver to agents collectively (under the utilitarian considerations) does not show any particular trend. This is due to comfort being formulated from historical (actual) consumption and not predicted consumption. However, our MIP models deliver slightly higher levels of comfort to the worst of agents as uncertainties reduce. This is so because the supply constraint factors they compute (see Table 7) are lower under higher levels of uncertainty, such that they impact on their egalitarian-comfort results. Contrariwise, our heuristics do not

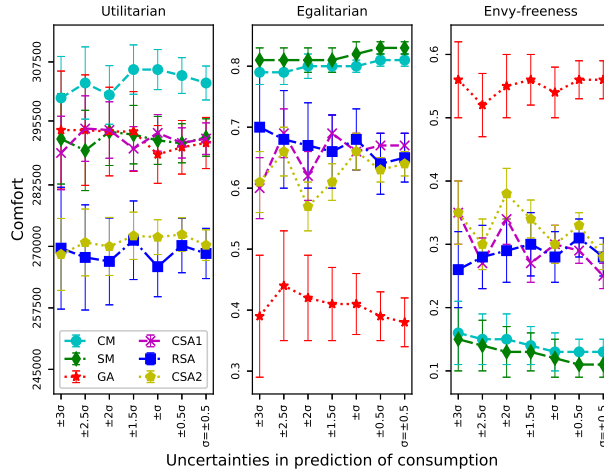


Fig. 7 Average comfort delivered under different levels of uncertainty

show any particular pattern under the egalitarian consideration. While the differences between the comfort delivered to the worst and best off agents by our MIP solutions slightly reduce with uncertainties (under the envy-freeness consideration), the results produced by our heuristics show no such patterns.

6.4.3 Electricity Supplied under Different Levels of Uncertainty

We display the average amount of electricity our solutions supply to agents individually (with respect to the egalitarian and envy-freeness metrics) and collectively (with respect to the utilitarian metric) under different levels of uncertainty in Figure 8.

Herein also, Figure 8 shows how our solutions produce results that are increasingly inconsistent as more uncertainties are introduced into consumption, due to more varied consumption data being generated under higher levels of uncertainty. This holds under the utilitarian, egalitarian and envy-freeness considerations. Furthermore, the electricity our solutions supply to agents collectively (under the utilitarian considerations) does not show any particular trend. This is due to the errors in predicted consumption which seem to cancel themselves out, such that the average sum of electricity supplied to agents fail to display any particular pattern over all levels of uncertainty. However, our MIP models supply higher amounts of electricity to the worst of agents as uncertainties reduce. This is mainly due to the supply constraint factors they compute as in Table 7. Contrariwise, our heuristics do not show any particular pattern under the egalitarian consideration. While the differences between the electricity supplied to the worst and best off agents by our MIP solutions consistently reduce with uncertainties (under the envy-freeness consideration), the results produced by our heuristics show no such patterns. To show that

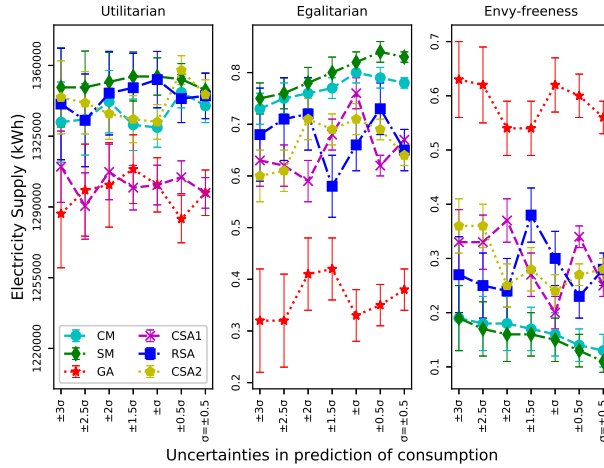


Fig. 8 Average electricity supplied under different levels of uncertainty

our models can be used within different settings, we show the results obtained by solving the FLSP on other datasets in the next section.

6.5 Implementation with other Datasets

To show further that our solutions generalize within the load shedding setting (and may generalize within other constrained utility maximization settings), we now show results of solving the FLSP with two other datasets: (i) the original dataset of homes in the USA (see Section 3) in Section 6.5.1, and (ii) a simulated dataset of 1000 homes based on the dataset of homes in Nigeria (see Section 3) in Section 6.5.2. Note that we return to computing predicted consumption from the normal distribution $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 0.05)$, as in all our evaluations.

6.5.1 Dataset of USA Homes

In this section, we show the results of solving the FLSP with the original dataset of household consumption of homes in the USA (from which we simulated the dataset in Section 3). This Pecan Street dataset is for 414 homes over 13 weeks (i.e., Q hours). We model all homes as agents by formulating their comfort profiles as we did in Section 3.2. Thereafter, we find parameters $\beta_1, \beta_2, \alpha_3$ and α_4 to be 21, 22, 0.78 and 0.72 respectively for CM. We also find parameters $\beta_1, \beta_2, \alpha_3$ and α_4 to be 21, 22, 0.79 and 0.78 respectively for SM. We solve the FLSP and show the results in Table 8.

Table 8 shows how our solutions perform in line with our evaluations in Section 6.1, Section 6.2 and Section 6.3, albeit under a different setting. This

Table 8 Results of MIP models and heuristics in terms of the average number of hours of connection to supply, average comfort delivered and average electricity supplied (following a number of implementations with the dataset of USA homes), along with their standard deviations (SD) within parenthesis

	CM (SD)	SM (SD)	GA (SD)	CSA1 (SD)	RSA (SD)	CSA2 (SD)
Connections	Utilitarian	810355 (4112)	712859 (4400)	725414 (2979)	722264 (4428)	719683 (3319)
	Egalitarian	1923 (3.27)	1926 (3.44)	1752 (2.30)	1744 (4.26)	1738 (2.30)
	Envy-freeness	145 (2.37)	97 (2.85)	723 (3.34)	1 (0)	1 (0)
Comfort	Utilitarian	339622 (3987)	326125 (4036)	326961 (3084)	300856 (4986)	302704 (3492)
	Egalitarian	0.81 (0.01)	0.82 (0.01)	0.68 (0.03)	0.67 (0.05)	0.65 (0.02)
	Envy-freeness	0.12 (0.02)	0.10 (0.03)	0.57 (0.03)	0.30 (0.04)	0.31 (0.03)
Supply	Utilitarian	3353701 (16447)	3261206 (20587)	3260997 (14197)	3268866 (20547)	3270249 (14228)
	Egalitarian	0.75 (0.01)	0.81 (0.02)	0.42 (0.04)	0.69 (0.03)	0.65 (0.03)
	Envy-freeness	0.17 (0.02)	0.11 (0.03)	0.55 (0.04)	0.28 (0.02)	0.30 (0.03)

Table 9 Results of MIP models and heuristics in terms of the average number of hours of connection to supply, average comfort delivered and average electricity supplied (following a number of implementations with the dataset of 1000 homes in developing countries), along with their standard deviations (SD) within parenthesis

	CM (SD)	SM (SD)	GA (SD)	CSA1 (SD)	RSA (SD)	CSA2 (SD)
Connections	Utilitarian	1953676 (7318)	1933673 (7330)	1766103 (6483)	1755480 (8296)	1746608 (6609)
	Egalitarian	1922 (3.31)	3.31 (3.46)	1766 (2.19)	1755 (4.42)	1746 (2.36)
	Envy-freeness	108 (1.86)	57 (1.54)	681 (3.11)	1 (0)	1 (0)
Comfort	Utilitarian	826212 (7769)	798271 (7837)	793047 (10949)	731820 (10895)	736455 (7025)
	Egalitarian	0.83 (0.01)	0.85 (0.02)	0.41 (0.05)	0.69 (0.02)	0.67 (0.03)
	Envy-freeness	0.11 (0.02)	0.8 (0.02)	0.53 (0.02)	0.24 (0.02)	0.26 (0.02)
Supply	Utilitarian	3656510 (17952)	3671976 (18076)	3557644 (23799)	3563821 (24252)	3567188 (15477)
	Egalitarian	0.79 (0.02)	0.85 (0.01)	0.42 (0.04)	0.70 (0.02)	0.66 (0.03)
	Envy-freeness	0.14 (0.02)	0.09 (0.03)	0.52 (0.04)	0.25 (0.02)	0.27 (0.03)

suggests that our solutions are applicable within other applicable settings with unique characteristics.

6.5.2 Dataset of Multiple Homes in Developing Countries

Herein, we show the results of solving the FLSP with the dataset of 1000 homes in developing countries. In Section 3, we generated the hourly consumption data of 367 homes in developing countries over three months. In order to see how our solutions scale and generalize, we generate the data of 1000 homes from the dataset for the same duration.²⁸ Thereafter, we model all homes as agents by formulating their comfort profiles as we did in Section 3.2. Then, we find parameters β_1 , β_2 , α_3 and α_4 to be 21, 22, 0.80 and 0.76 respectively for CM. We also find parameters β_1 , β_2 , α_3 and α_4 to be 21, 22, 0.82 and 0.82 respectively for SM. We solve the FLSP and show the results in Table 9.

As in Section 6.5.1, Table 9 also shows how our solutions perform in line with our evaluations in Section 6.1, Section 6.2 and Section 6.3 under a different setting. This suggests that our solutions are applicable within other applicable settings with unique characteristics. We conclude this section by discussing the time complexities of all load shedding solutions in the next section.

6.6 Computation Complexities of Load Shedding Solutions

Our heuristics are polynomial time algorithms that mainly depend on the number of agents (i.e., $O(n)$ for GA, $O(n \log n)$ for CSA1, $O(n)$ for RSA and $O(n \log n)$ for CSA2, where n is the number of agents). In turn, our MIP solutions are built upon a MKP formulation that has a non-polynomial time complexity. However, we solve our MIP using the CPLEX optimization package within the Python environment. CPLEX solves integer programming problems using efficient algorithms and relaxations.²⁹ We illustrate this by solving the FLSP for different populations of agents drawn from our main evaluation dataset,³⁰ and with other test datasets (i.e., datasets used in Section 6.5.1 and Section 6.5.2). We show the time taken to arrive at all solutions. It should be noted that, for all these datasets (with $\tilde{c}_i^t \sim \mathcal{N}(c_i^t, 0.05)$), our solver arrives at different factors (i.e., different values for β_1 , β_2 , α_3 and α_4) for both MIPs. This further suggests that our approach is applicable within many cases (i.e., different dataset or environments) with different system characteristics.³¹ We implement all load shedding solutions with these datasets and show their average execution times in Figure 9.

²⁸ We do this by drawing from a normal distribution with the consumption of homes in the dataset of Section 3 as the mean.

²⁹ Refer to the CPLEX Optimization Studio User's Manual [18].

³⁰ These populations (or datasets) are taken from the original developing country dataset simulated in Section 3.

³¹ By this we mean that, for each dataset, we take the steps in Section 3.3 to compute the hourly supply capacities and estimates of agent demands. As such, we develop a number of applicable settings, each with its unique characteristics.

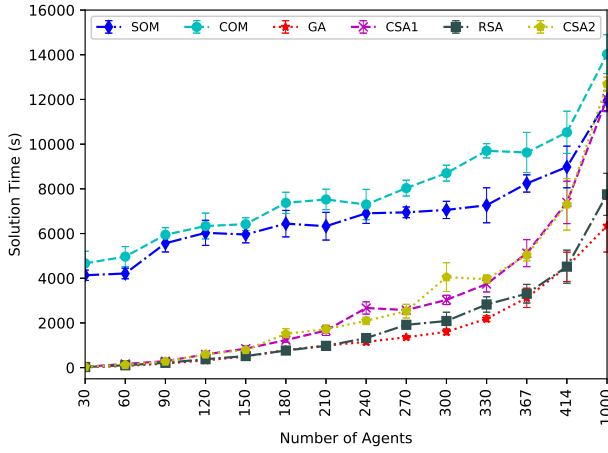


Fig. 9 The average runtimes of all load shedding solutions after nine executions

Table 10 MIP constraint factors under different settings (i.e., different population of agents)

		30	60	90	120	150	180	210	240	270	300	330	367	414	1000
CM	β_1	20	20	21	21	21	21	21	21	21	21	21	21	21	21
	β_2	21	21	22	22	22	22	22	22	22	22	22	22	22	22
	α_3	0.70	0.71	0.72	0.74	0.75	0.75	0.76	0.75	0.76	0.79	0.80	0.78	0.78	0.80
	α_4	0.73	0.73	0.74	0.74	0.74	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.72	0.76
SM	β_1	20	20	21	21	21	21	21	21	21	21	21	21	21	21
	β_2	21	21	22	22	22	22	22	22	22	22	22	22	22	22
	α_3	0.72	0.72	0.74	0.75	0.75	0.76	0.77	0.78	0.79	0.80	0.80	0.79	0.82	0.82
	α_4	0.77	0.78	0.79	0.80	0.79	0.81	0.80	0.80	0.80	0.79	0.80	0.80	0.78	0.82

As can be seen in Figure 9, the execution time of the MIPs grow polynomially and does not entirely depend on the population of agents. This is depicted by results which show that some MIPs with higher agent population are solved in less time than others (e.g., CM with 240 and 367 agents that have lower runtimes than the MIPs with 210 and 330 respectively, and SM with 150 and 210 agents that have lower runtimes than the MIPs with 120 and 180 respectively). To understand why this may have occurred, we present the constraint factors used to solve our MIPs in Table 10.

Table 10 suggests that the periods taken to execute our MIPs were also affected by their constraints (notably because the constraint factors used within the settings of 240 and 367 agents are lower than those used within the settings of 210 and 330 agents respectively by CM, and the constraint factor used within the settings of 150 and 210 agents are lower than those used within the settings of 120 and 180 agents respectively by SM). As such, the runtimes of our MIP solutions depend on their constraints. Additionally, their complexities will increase with the number of hours we optimize over (which we maintain as 24 hours in solving the FLSP a day ahead in this paper), and the number of variables in our formulation.

7 Conclusion

We summarize this paper in this section and discuss the open challenges that still exist. Specifically, in Section 7.1, we summarize the main body and results of this report while, in Section 7.2, we discuss the future lines of work that build on our solutions.

7.1 Summary of Work

In this paper, we introduced, formulated and provided a number of solutions to the load shedding problem. To implement our proposal, we simulated data that was representative of households in developing countries. We did this from publicly available datasets of electricity consumed by households in developed countries. Specifically, we collected appliance-level consumption data of households in the USA from Pecan Street Inc’s Dataport. Then, based on findings about appliances commonly used in both Nigeria (a typical developing country) and the USA [31,45,11,26,32], the temperature in cities of both countries and their consumption patterns [16,17,38], we developed the appliance-level data into household consumption data for households in a developing country. Thereafter, we modelled households into agents, each with its preference (or need) for consuming electricity. In so doing, we created a notion of comfort that resulted in some computed values which embody the electricity needs of households.

Following this, we developed four heuristic algorithms that considered different criteria for disconnecting agents from supply. All heuristics had the objective to connect agents to supply as evenly as possible, in terms of number of hours. We evaluated their results using the utilitarian, egalitarian and envy-freeness social welfare metrics. Thereafter, we highlighted the shortcoming of the heuristics, in that they fail to consider the agents’ preference (or needs) for electricity.

On this ground, we modelled the fair load shedding problem (FLSP) as two mixed integer programming (MIP) problems built upon a multiple knapsack problem formulation. The first had the objective to maximize the overall comfort of agents and the other had the objective to maximize overall supply to agents (both with respect to the utilitarian metric). Using the egalitarian and envy-freeness metrics, we developed a number of constraints which ensure that the number of hours individual agents are connected to supply, and the comfort delivered and electricity supplied to these individual agents, are as high and as equal as possible. We also included another constraint that executed load shedding when necessary, thus resulting in a balance between supply and demand.

Following this, we evaluated the performance of all six load shedding solutions (i.e., the four heuristics and the two MIPs). We performed the evaluations with respect to the number of hours agents were connected to electricity, the comfort delivered to agents individually and collectively, and the electricity supplied to agents individually and collectively. We showed how the MIP solutions outperformed the heuristics under these three key evaluations.

When taken altogether, our approach presents a framework upon which other fairness problems involving constrained utility maximization (or resource allocation) may be generalized. Our solutions also serves as benchmarks for designing future load shedding solutions, as well as solutions for allocating other scarce resources (e.g., water allocation problems addressed by [41,47]).

7.2 Future Work

We made some assumptions in developing our household-level fair load shedding solutions. One of these is that we receive near-accurate estimates from homes, based on which we took the estimates of household consumption as values drawn from a normal distribution of the electricity they actually consume. However, it will be useful to design solutions which incorporate the computation of these estimates in the future.

In addition to this, cost effective elicitation mechanisms applicable to developing countries may be designed to improve the accuracy of our comfort formulation. Such an elicitation mechanism may consider the needs of individual consumers in homes. An example is an occupant who needs an electric respiratory aid in the home for managing and treating sleep apnea or another who needs electricity to prepare for an upcoming examination. To be cost effective and applicable, such may elicit information (maybe through text messages) from homes in a bother-free manner (as has been studied in [53, 22]). this case, it is necessary to appropriately incentivize the true reporting of the information.

Furthermore, approaches like “soft load shedding” (see [2]) are being developed to mitigate against the rebound effect (described in Section 3.4), and may be incorporated into our solutions in the future. In addition, it is important to consider that households which are arbitrarily subjected to load shedding may advance their consumption to periods when electricity is available. As such, future fair load shedding solutions should also consider the effect load shedding has on prior consumption. Although we made the attempt to consider the rebound effect in this work, our comfort formulation may be further improved to ensure it still embodies the preference of agents for electricity, even after load shedding influences their consumption.

We produced hourly solutions which are applicable because electricity billing is on an hourly basis (i.e., in kWh). Additionally, many energy, resource allocation and scheduling problems are solved in hourly or half-hourly time slots (e.g., in [1]). Likewise, RTPs are hourly (or half-hourly) pricing schemes that contribute to grid stability. However, in the future, our solution can be implemented using per minute data. This will be more computationally intensive, as suggested in Section 6.6. Nonetheless, to ensure the computational cost of using per minute data remains tractable, future solutions may solve the FLSP incrementally (e.g., one hour at a time), and employ approximation and relaxation techniques. Note that the computation complexity of such FLSP solutions is not a deterring factor, as they are not designed to be implemented at real-time. In addition, solutions that use per minute data should produce

more accurate results (from using more accurate comfort values and estimates of consumption).

Finally, our work can be designed to fit alongside other load management measures such as demand side management (DSM) (read more about DSM in [39]). As such, it is necessary to design incentives that encourage consumers to manage their consumption appropriately, and to have an understanding of how they will react to such incentives. In addition, with the right coordination mechanisms, our work can cope with the introduction of distributed renewable sources and storage on the grid [4]. However, it will be necessary to accurately predict demand (and supply, especially in the case of distributed renewable sources and storage) before our solutions can be effective when incorporated with these smart grid technologies [40].

References

1. Alam, M., Gerding, E.H., Rogers, A., Ramchurn, S.D.: A scalable, decentralised multi-issue negotiation protocol for energy exchange. In: International Joint Conference on Artificial Intelligence (IJCAI) (2015)
2. Aslam, T., Arshad, N.: Soft load shedding: An efficient approach to manage electricity demand in a renewable rich distribution system. pp. 101–107 (2018). DOI 10.5220/0006774001010107
3. Azasoo, J., Boateng, K.: A retrofit design science methodology for smart metering design in developing countries. In: 15th International Conference on Computational Science and Its Applications (2015)
4. Chalkiadakis, G., Robu, V., Kota, R., Rogers, A., Jennings, N.R.: Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants. Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems Innovative Applications Track (AAMAS 2011) (IFAAMAS), 787–794 (2011)
5. Chao, H.L., Hsiung, P.A.: A fair energy resource allocation strategy for micro grid. *Microprocessors and Microsystems* **42**, 235–244 (2016)
6. Chekuri, C., Khanna, S.: A ptas for the multiple knapsack problem. In: Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '00, pp. 213–222. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2000)
7. Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J.A., Sousa, P.: Issues in multiagent resource allocation. *Applied Soft Computing* (2006)
8. City Power Johannesburg: The Five Stages of Load Shedding (2014). URL <https://www.citypower.co.za/customers/Pages/Infographic.aspx>
9. Department for Environment, Food and Rural Affairs: Household Electricity Study (2011)
10. Do, L., Lin, K.H., Molnr, P.: Electricity consumption modelling: A case of germany. In: *Economic Modelling* (2016)
11. Emodi, N., Emodi, C., Murthy, G., Procaccia, A., Emodi, A.: Energy policy for low carbon development in nigeria: A leap model application. In: *Renewable and Sustainable Energy Reviews* (2017)
12. Fabini, D.H., Baridó, D.P.d.L., Omu, A., Taneja, J.: Mapping induced residential demand for electricity in kenya. In: Proceedings of the Fifth ACM Symposium on Computing for Development, ACM DEV-5 '14, pp. 43–52. ACM, New York, NY, USA (2014)
13. Freeman R., Z.S., Conitzer, V.: Fair and efficient social choice in dynamic settings (2017)
14. Gao, J., Giri, S., Kara, E.C., Bergés, M.: The Plug Load Appliance Identification Dataset (PLAID), a Public Dataset of High Resolution for Load Identification Research (2014). URL <http://plaidplug.com>
15. Gerding, E.H., Robu, V., Stein, S., Parkes, D.C., Roggers, A., Jennings, N.R.: Online mechanism design for electric vehicle charging. 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011) (IFAAMAS), 811–818 (2011)

16. Holiday Weather: Austin: Annual Weather Averages (2018). URL <https://www.holiday-weather.com/austin/averages/>
17. Holiday Weather: Lagos: Annual Weather Averages (2018). URL <https://www.holiday-weather.com/lagos/averages/>
18. IBM: IBM ILOG CPLEX Optimization Studio CPLEX User's Manual, Version 12 Release 6 (2015). URL https://www.ibm.com/support/knowledgecenter/SSSA5P_12.6.2/ilog.odms.studio.help/pdf/usrcplex.pdf. <https://goo.gl/1ZWRJM>
19. Johnson, G., Beausoleil-Morrison, I.: Electrical-end-use data from 23 houses sampled each minute for simulating micro-generation systems. *Applied Thermal Engineering* **114**, 1449 – 1456 (2017)
20. Kaygusuz, K.: Energy for sustainable development: A case of developing countries. *Renewable and Sustainable Energy Reviews* **16**(2), 1116 – 1126 (2012)
21. Keelson, E., Boateng, K., Ghansah, I.: A smart retrofitted meter for developing countries. In: *International Journal of Computer Applications* (2014)
22. Le, T., Tabakhi, A.M., Tran-Thanh, L., Yeoh, W., Son, T.C.: Preference elicitation with interdependency and user bother cost. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, pp. 1459–1467. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2018). URL <http://dl.acm.org/citation.cfm?id=3237383.3237918>
23. Leite, J., Alferes, J., Mito, B.: Resource allocation with answer-set programming. In: *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems* (2009)
24. Lust, T., Teghem, J.: The multiobjective multidimensional knapsack problem: a survey and a new approach. *CoRR* **abs/1007.4063** (2010)
25. Miller, S., Ramchurn, S.D., Rogers, A.: Optimal decentralised dispatch of embedded generation in the smart grid. In: *AAMAS 12 International Conference on Autonomous Agents and Multiagent Systems*. ACM (2012)
26. Monyei, C., Adewumi, A., Obolo, M., Sajou, S.: Nigeria's energy poverty: Insights and implications for smart policies and framework towards a smart nigeria electricity network. In: *Renewable and Sustainable Energy Reviews* (2018)
27. Murray, D., Liao, J., Stankovic, L., Stankovic, V., Hauxwell-Baldwin, R., Wilson, C., Coleman, M., Kane, T., Firth, S.: A data management platform for personalised real-time energy feedback. In: *Proceedings of the 8th International Conference on Energy Efficiency in Domestic Appliances and Lighting*. IET (2015)
28. Murray, D., Stankovic, L.: REFIT: Electrical Load Measurements (2015)
29. Myerson, R.: Utilitarianism, egalitarianism, and the timing effect in social choice problems. In: *Econometrica* (1981)
30. Nwachukwu, M., Ezedinma, N., Jiburum, U.: Comparative analysis of electricity consumption among residential, commercial and industrial sectors of the nigerias economy. In: *Journal of Energy Technologies and Policy* (2014)
31. Oji, J., Idusuyi, N., Aliu, T., Petinrin, M., Odejobi, O., Adetunji, A.: Utilization of solar energy for power generation in nigeria. In: *International Journal of Energy Engineering* (2012)
32. Oluwadamilola, O., Harrison, G., Weijde, H.: Modelling electricity and cooling load profiles for domestic customers in Nigeria (2015). URL http://www.wemcouncil.org/ICEMs/ICEM2017_PRES/ICEM2017_Poster_62_Oluwadamilola.pdf
33. Oluwasuji, O., Obaid, M., Zhang, J., Ramchurn, S.: Algorithms for fair load shedding in developing countries. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence* (2018)
34. Oyedepo, S.: Energy and sustainable development in nigeria: the way forward (2012)
35. Pahwa, S., Scoglio, C., Das, S., Schulz, N.: Load-shedding strategies for preventing cascading failures in power grid. *Electric Power Components and Systems* **41**(9), 879–895 (2013)
36. Parson, O., Fisher, G., Hersey, A., Batra, N., Kelly, J., Singh, A., Knottenbelt, W., Rogers, A.: Dataport and nilmtk: A building data set designed for non-intrusive load monitoring. In: *IEEE Global Conference on Signal and Information Processing (GlobalSIP)* (2015)
37. Pecan Street Inc.: Dataport: More than the Largest Source of Energy Data and Water Data (2018). URL <https://dataport.cloud>

38. Prinsloo, G., Dobson, R., Brent, A.: Scoping exercise to determine load profile archetype reference shapes for solar co-generation models in isolated off-grid rural african villages **28** (2016)
39. Ramchurn, S.D., Vytelingum, P., Rogers, A., Jennings, N.: Agent-based control for decentralised demand side management in the smart grid. In: The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS '11, pp. 5–12. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2011)
40. Ramchurn, S.D., Vytelingum, P., Rogers, A., Jennings, N.R.: Putting the ‘smarts’ into the smart grid: A grand challenge for artificial intelligence. *Commun. ACM* **55**(4), 86–97 (2012)
41. Read, L., Madani, K., Inanloo, B.: Optimality versus stability in water resource allocation. In: *Journal of Environmental Management* (2014)
42. Robu, V., Kota, R., Chalkiadakis, G., Rogers, A., Jennings, N.R.: Cooperative virtual power plant formation using scoring rules. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 3, AAMAS '12, pp. 1165–1166. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2012)
43. Rose, H., Rogers, A., Gerding, E.H.: A scoring rule-based mechanism for aggregate demand prediction in the smart grid. *AMMAS Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems* pp. 661–668 (2012)
44. Sakah, M., de la Rue du Can, S., Diawuo, F.A., Sedzro, M.D., Kuhn, C.: A study of appliance ownership and electricity consumption determinants in urban ghanaiian households. *Sustainable Cities and Society* (2018)
45. Salmon, C., Tanguy, J.: Rural electrification and household labor supply: Evidence from nigeria. In: *World Development* (2016)
46. Shi, B., Liu, J.: Decentralized control and fair load-shedding compensations to prevent cascading failures in a smart grid. In: *International Journal of Electrical Power and Energy Systems* (2015)
47. Song, W., Yuan, Y., Jiang, Y., Lei, X., Shu, D.: Rule-based water resource allocation in the central guizhou province, china. In: *Ecological Engineering* (2016)
48. Stein, S., Gerding, E., Robu, V., Jennings, N.: A model-based online mechanism with pre-commitment and its application to electric vehicle charging. Online (June), 4–8 (2012)
49. Stolworthy, M.: UK Electricity National Grid Demand and Output per Production Type (2018). URL <http://gridwatch.co.uk>
50. Suganthi, L., Samuel, A.A.: Energy models for demand forecasting a review. *Renewable and Sustainable Energy Reviews* **16**(2), 1223 – 1240 (2012)
51. The NILM wiki: NILM datasets (2014). URL <http://wiki.nilm.eu/datasets.html>
52. Truong, N., Tran-Thanh, L., Costanza, E., Ramchurn, S.: Towards appliance usage prediction for home energy management. In: Proceedings of the Fourth International Conference on Future Energy Systems (2013)
53. Truong, N.C., Baarslag, T., Ramchurn, G., Tran-Thanh, L.: Interactive scheduling of appliance usage in the home (2016). URL <https://eprints.soton.ac.uk/396670/>
54. Vinyals, M., Robu, V., Rogers, A., Jennings, N.R.: Prediction-of-use games: A cooperative game theory approach to sustainable energy tariffs. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14, pp. 829–836. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2014)
55. World Energy Council: Average household electricity use around the world (2010). URL <http://shrinkthatfootprint.com/average-household-electricity-consumption>
56. Yohanna, I., Iortyer, H., Ierve, A.: Assessment of energy use pattern in residential buildings of kano and kaduna northern nigeria. In: *American Journal of Engineering Research* (2013)