An energy market for trading electricity in smart grid neighbourhoods

Dejan Ilic, Per Goncalves Da Silva, Stamatis Karnouskos and Martin Griesemer
SAP Research, Karlsruhe, Germany
Email: {dejan.ilic,per.goncalves.da.silva,stamatis.karnouskos,martin.griesemer}@sap.com

***Abstract*—The smart grid vision relies on active interaction with all of its stakeholders. As consumers are acquiring energy generation capabilities, hence becoming prosumers (producers and consumers), a meaningful way to interact among them would be to trade over a marketplace. Market-driven interactions have been proposed as a promising potential interaction method due to the monetary incentives and other benefits involved for the participants [1]. In the Internet era an on-line marketplace is an thriving concept as it overcomes potential accessibility issues, however it is not clear how they should be structured, operated, what their limits and benefits might be. The design, implementation, modus-operandi as well as the assessment of such an energy market place for smart grid neighbourhoods is presented.**

I. Motivation

The current electricity grid follows with minor changes a 100 years old paradigm, in which electricity is mainly generated by a few far-off high-capacity generators and transmitted and distributed to end consumers. However, due to various reasons (including high cost, resource availability, sustainability, resilience, and environmental concerns) electrical energy is increasingly being produced by small-scale generators located closer to points-of-use. In conjunction with rapidly advancing information and communication technologies that empower the smart grid, this results in a highly sophisticated IT-based dynamic grid infrastructure [2].

As the paradigm changes from a centralized to a distributed production model, a significant impact is made on business in the electric energy domain. Assuming the continuation of the trend, in the near future, it will even be common for households to have their own production units, such as photo-voltaic panels, CHPs etc. These units will be used to meet internal demand, while surplus generation could possibly be used to meet local demand, in the neighbourhood. As such, considering that a significant percentage of the citizens have such generation capabilities, it makes sense to be able to trade an energy surplus locally. For instance, a surplus of energy generated from a rooftop photo-voltaic panel of one user, could be used to meet a portion of another user's demand. In this way, if energy is traded locally, a good portion of energy transmission costs can be reduced. Additionally, the dynamics of renewable resources, which is subject to localized conditions, such as the availability of sunlight or wind, could be better fitted if the surplus of other prosumers is available and can be communicated effectively.

Today, local generation is mostly bound to feed-in tariffs which however are slowly being reduced, and will finally be removed. Hence, a new approach is needed that will not rely on subsidization, but rather on free market rules. If such a market can be realized locally e.g. in a smart grid city, it will be possible to trade energy (both buying and selling) with various benefits for all stakeholders [1]. The existence of several local energy markets will bring new possibilities as well as challenges [3] to the traditional business relations and processes. However, the expectation is to offer a better way to manage highly volatile networks due to the large scale distribution of both production (e.g. solar panels, wind generation etc.) and consumption (e.g. electric cars etc.) by relying on economics. In this line of thought, a new kind of complex dynamic system will emerge that due to its cyber-physical nature [3], [4] has to be carefully designed and realized.

II. Neighbourhood Energy Trading

For energy trading to be realized, the necessary tools as well as timely information exchange between all stakeholders need to be provided. Part of it is also the actual smart metering i.e. the high granularity of metering data acquisition e.g. every 15 mins. Through better resolution of the production and consumption (prosumption) data and effective analysis, any market participant is able to monitor and even predict his energy behaviour. With smart meter technology deployments in large scale, as well as the necessary energy services [5], prosumers are able to offer and purchase electricity. As an example of these new capabilities, within the NOBEL project (www.ict-nobel.eu) a local energy market at smart neighbourhood/district level is realized and will be assessed with in the Spanish city of Alginet in 2012. The primary goal is to facilitate and manage the electricity trading between the citizens of a smart city. Additionally, the implied aim is also to use it for market-driven demand-response (DR).

By participating on the NOBEL market, participants can take optimal advantage of local conditions and consume electricity produced locally. Not only can one avoid the transportation costs and energy losses, but better planning around and management of local networks may be achieved. Eventually this concept could lead to interconnected micro-grid-like networks that may cooperate within a greater region and provide emergent behaviour to the electricity network. It has to be pointed out again that such a network is a system of systems [6], which will get even more complex as decentralization and integration of real-time information flows among prosumers increases.

The market idea has been in the heart of major roadmaps for the Smart Grid [7], [8]. Efforts already exist [9], [10] where a continuous double auction model is applied. In [9] an agent-based framework is described, and emphasis is given to trading agents and the architecture of the overall system, rather than the behaviour and outcomes of the system. In [10] the market is used to price the flow of electricity through inter-connectors, in order to manage network congestion. In the approach depicted in this paper, the focus is more on the market aspects and energy trading at neighbourhood/district level. Here, it is considered that the prices of energy trading include the necessary transmission costs, hence the transmission system costs are considered static or a varying percentage but are hidden from the end-user at this stage.

III. The NOBEL Market Design

The NOBEL market is based the stock exchange model, with the difference that the trading periods are discrete fixed-sized time slots throughout the day. To favour a decentralised approach, the order book and last transaction price (for all x) are made public. The timeslots are defined as $x \in X$, where X is the set of all timeslots where trading (buying or selling) can be realized. Each timeslot is considered to be of the same length δ and is limited by the start time t_d^x (this is also the delivery start time) and the trading end time $t_d^x + \delta$. The time interval from t_d^x to $t_d^x + \delta$ is the time frame energy is traded for. Every participant willing to trade can place a market order $o_i^x \in O^x$, where $i \in \mathbb{N}$ and O^x is the order book for timeslot x. Figure 1 depicts market slots where each slot's price develops as orders get matched, in this case to the slot x.

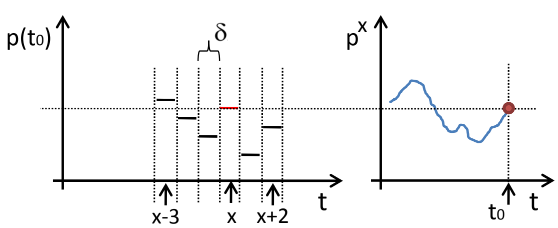

Figure 1. Market slots and their price development as orders (O^x) get matched

To facilitate meaningful interactions in the market, both consumers and producers should be capable of predicting their electricity demand/supply for a particular timeslot x. As timeslot x opens for trading at time $t_o^x = t_d^x - \tau$, where τ is fixed for every x before its delivery time. For $t_o^x \leq t \leq t_c^x$, where t_c^x is the closing time, the timeslot x is open for participants to place their orders. Orders can, of course, be adjusted. This is needed as no prediction tool can be completely accurate and therefore participants should be able to adjust their orders. Any prediction deviations e.g. caused by dynamically changing behaviours can possibly be considered by adjusting the orders, provided that the timeslot is still open for trading. For instance, if a household has overestimated the amount of energy it would use in a particular timeslot, it is able to sell the difference back to the market. Once the market order quantity is defined, the participant still has to put the limits for selling back the energy.

For each o, the type (buy or sell) is specified for the number of units $u \in \mathbb{N}^*$ and the price p (per unit) s/he is willing to trade for. In the simplest case, the participant can observe the top of each book "the best (highest) buy and the best (lowest) sell" also known as the "inside market". It then assumed that the market's consensus about the price lies somewhere between these two numbers. In Figure 2 the best buy price is 25.21, and the best sell price is 25.30 (the inside market is 25.21..., 25.30), and the "true price" of an energy block is in this interval. Now, the market participant can use the top of each order book as a reference point for positioning his initial quotes e.g. at 25.22..., 25.29 as depicted in Figure 2(b), and update his spread as the book evolves with new arrivals, transactions and cancellations.

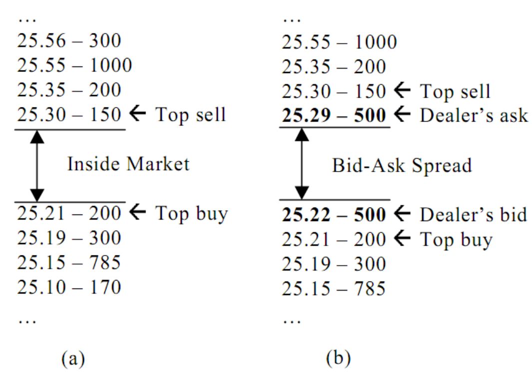

Figure 2. Order book example

Every order represents an acceptable price p for each unit of quantity in an order, of a participant. An order o_i^x may represent only a portion of the participant's electricity demand/supply to be met locally. If a buy order has the price p_{buy} and a sell order has price p_{sell}, a transaction (order matching) will occur only if $p_{buy} \geq p_{sell}$. The matching process will be repeated every time a new order is received. Therefore an order o_i^x may partially match with multiple orders from the slot's set of orders O^x. Since an order may not be fully executed, it possibly stays in the order book but only with the remaining (unmatched) quantity.

While new orders can be placed, updating p for existing ones is at the heart of our market. An order update within a slot x possibly affects the rest of the orders in the set O^x. Each order update first cancels the original order, and a new order is created. Update requests can be classified into two categories representing two different types of strategies. The first attempts to foresee the upcoming market movements (either from the order book imbalances or from short-term patterns), and adjust the spread according to these expectations. The second reasons solely on the information about the current inside market. These "non-predictive" strategies are inherently simpler and therefore better suited for our examination of the local electronic market making.

As the matching process is repeated every time an order is inserted or updated, it is important to understand key advantages/disadvantages of the proposed market. Similarly to the stock market model, matching is based on the First-Come, First-Served (FCFS) service policy. Every newly received order, update, or cancellation, will be sequentially executed. Therefore, if a market participant places a new order o_i^x at t_0, it has an advantage over participants placing orders on timeslot x at $t > t_0$. This policy also affects the price fluctuations. In order to fully execute an order, one side (buy or sell) has to overcome the price spread of the x. Consequently, timing in the updated and cancellation process becomes even more important. If a new order o_i^x is located in the execution queue before the o_j^x cancellation request, the two orders may be matched before o_j^x gets cancelled.

Every matched order in the trading period for a timeslot can be considered as the contracted good. This contract is made between participants of the matched orders executed by the matching algorithm. Unmatched orders will simply be aborted by the market execution cycle. These have no contractual value for any participant.

A timeslot closes for trading at time $t_c^x = t_d^x - \omega$, where ω is fixed time period before t_d^x and $\tau > \omega$. Once a timeslot is closed, historical analysis can be used for future decisions. As an example, a participant can use it to understand its final costs in comparison to average block price, or minimum price, of an acquired block in a particular slot, etc.

IV. Market Implementation

The proposed market has been designed and implemented; the following sections provide an insight to the architecture and configuration decisions we made.

A. Order Configuration

Given that this is a marketplace for trading energy, different order configurations should be made available to the participants. Using order configurations, one can express specific energy requirements, or usage patters. The order configurations are composed of two behaviours, as depicted in Table I. The first dimension, specifies whether units of an order can be partially matched, or if must be fully matched. “Fully match” indicated if a participant wants everything or nothing. The second dimension specifies if an order has to be matched immediately. If immediate match is required, possible matching is executed while the unmatched part of the order is automatically cancelled. Matching limitations of this dimension are the trading price and availability of the trading commodity. With these four order configurations, participants should be able to express their internal processes, or trading strategies. For instance, a fully matching order could be used for a process which requires the full amount of energy to be available for the entire duration. These same order configurations are also available in the EPEX intraday market (www.epexspot.com) as well as in other European and American markets.

Table I
Order Configurations

Fully Match \ Immediate Match	YES	NO
YES	Standard	Immediate or Cancel (IOC)
NO	All Or None (AON)	Fill Or Kill

B. Architecture

The proposed and implemented market architecture can be seen in Figure 3. A market participant submits an order to the market through the market communication manager (who performs the necessary security checks). Firstly, the order validity check is made within the verification module. It checks if an order complies with all market rules, that is, price limitations, order configuration, timeslot validity etc. If an order fails the verification process, an appropriate response is returned to the market participant. Otherwise the verification module forwards it to the market kernel module and participant gets notified of a successful order submission.

The market kernel module is responsible for managing the life-cycle of the timeslots, their order books, and the order matching process. Once a new order comes in, it is passed to the matching algorithm, along with the order book for the order’s timeslot. If there is a matching executed, the state of an order changes, and the market output manager is notified of the transaction. Since an order might not match as soon as it arrives, the output of the market is handled asynchronously. Market participants and other tools (e.g. an external analytics service) can subscribe to this module and receive the produced market notifications. Of course, the information a subscriber has access to, is limited by his credentials. Thus, it is ensured that a market participant only has access to the outcome of his/her orders.

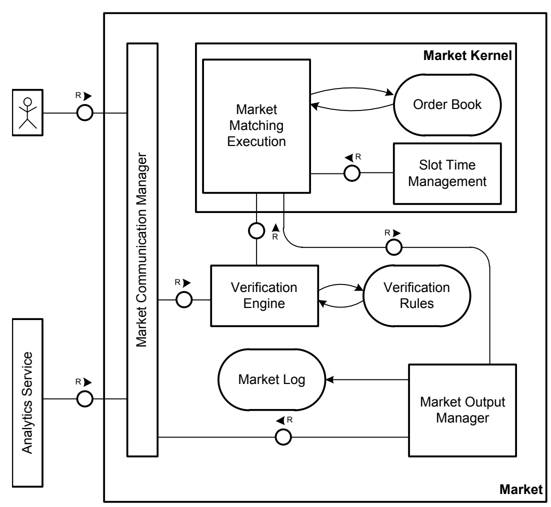

Figure 3. Overview of market modules

It must be made clear that the architecture of Figure 3 depicts the main market components, upon which other stakeholders and services can built their functionalities. For instance the analytics service has the credentials to gather all market outcomes (it stores them as they are received) in order to process the data with various algorithms and analyse the market behaviour.

Some of the parameters of the market can be changed at runtime by updating the verification rules module. All parameters can be changed without difficulties except for parameters related to timeslot management, e.g. opening or closing time. Due to the effect such a change can have on market processes, parameter changes only affect timeslots which are in a previous state than the one the parameter defines. That is:

- Changing the value of τ only affects the upcoming timeslots
- Changing the value of ω affects all timeslots except the already closed ones

The matching algorithm must cope with the different order configuration outlined in Table I. As such it must handle different cases when a particular order type is being matched with another order type. In the case where an order requires immediate matching, the trigger order passed to the process is used. If, after a run of the matching algorithm, this order is not fully executed, all the remaining quantity is cancelled. When one or two orders which need full matching match together, the availability of all quantity has to be checked.

To reduce CPU cycles of the matching process, all orders are ordered according to their price and their type (BUY or SELL). Their quantities are added until the required quantity is reached. If the quantity is achieved and the price limit is not breached, the entire quantity needed by the order is available at, or bellow, the required price; hence a trade takes place. However if the incoming order only required partial matching, and the top order required full matching, this process also needs to be run again for that particular order, as there can be previous orders with whom the total quantity can produce a match.

The matching algorithm may not let a full matching order block trading. For instance, if the cheapest sell order required fully matching, and there is no combination of buy orders which can fulfil the sell order, the market should not wait until the sell order can be processed before processing the other sell orders. Therefore, in this case, the top order is ignored, and matching is attempted with the remainder of the orders.

The proposed market has been implemented in Java SE version 6 and is available into two forms (i) as a simulator with accelerated time clock that may run locally and (ii) as an Internet application that runs on the *Glassfish 3.1 Application Server* making its functionalities available as *Java REST services*. The business data are stored in a*mySQL* DB (www.mysql.com). All external communication is done via an encrypted channel i.e. *HTTPS* and a security (with role-based authorization and authentication) framework is in place based on *Apache Shiro* (shiro.apache.org). Additionally for performance reasons, all services interact using *Google Protocol Buffers* (code.google.com/p/protobuf/) which offer a highly efficient binary format. The developed market forms an integral part of a platform providing energy services for the smart grid city [5].

V. Evaluation

In order to evaluate the realized NOBEL market, a simulation platform was built in which trading agents could be defined and trade in the market. All output pertaining to the market, i.e. bids, offers and transactions was collected and analysed. In the following sections, the trading agent, scenario, and key measures are described.

A. Trading via Agents

To evaluate the market, agents that interact with the market on behalf of their stakeholders and buy and sell energy were implemented. As a very initial investigation, the focus has been on very simple approaches i.e. experimented with "Zero Intelligence" (ZI) agents [11]. A ZI has no memory, and no guiding trading strategy, and simply bids to the market using random (within a user-specified limit) pricing. Since the behaviour of ZI can help reach a high level of order matching, the equilibria of the market can be found. Thus, the aim is not to apply a trading strategy e.g. to maximize the profit, but to try to understand what the outcome of the market model is given its rules.

In our case, a ZI agent is defined by (i) the minimum price ℓ_{min}, (ii) the maximum price ℓ_{max} and (iii) the maximum sleep time t_{sleep}. The maximum and minimum price parameters define the range of price range the agent uses to randomly bid. In our implementation, each agent is a thread, and the maximum sleep time defines the duration the agent waits between running cycles.

In the implemented scenario, each agent has access to its future consumption and/or production behaviour (which of course represents the end-user's behaviour). Each agent trades only once per timeslot (after t_o^x). The order price is chosen from the range $[\ell_{min}, \ell_{max}] \in \mathbb{R}$ using uniform distribution. The quantity is defined as the forecast energy behaviour in the timeslot. As always the surplus is traded, if an agent has more production, a buy bid is submitted, and similarly if the agent has more consumption, a sell bid is submitted. Once the agent bids, it sleeps for a random time $t_s : 0 < t_s \leq t_{sleep}$. At each run of the agent, a new price and sleep time are generated.

B. Simulation Scenario

In the evaluation scenario, the timeslot duration $\delta = 15min$, as this is the typical smart metering measurement frequency. Also, t_o^x and t_c^x are set through preconfigured $\tau = 2days$ and $\omega = 2hours$. There are two types of agents in this scenario i.e. 50 producers and 50 consumers. Their price parameters are fixed to $\ell_{min} = 12.00$ and $\ell_{max} = 20.00$ (cents/kWh) with the double digit precision. The t_{sleep} parameter is set to a value smaller than δ, implying execution of (at least) one trading cycle for every slot opening.

For the market consumers we have used the German representative load profile for households $H0$. Figure 4 depicts this load curve. The are two types of producers, 80% are fit with photovoltaic (PV) panels. Their output is produced from the radiation data (as reported in [12]) with an assumed efficiency of 19%. The remaining 20% of production agents are equipped with wind turbines of $4kW$ power of 30% average efficiency.

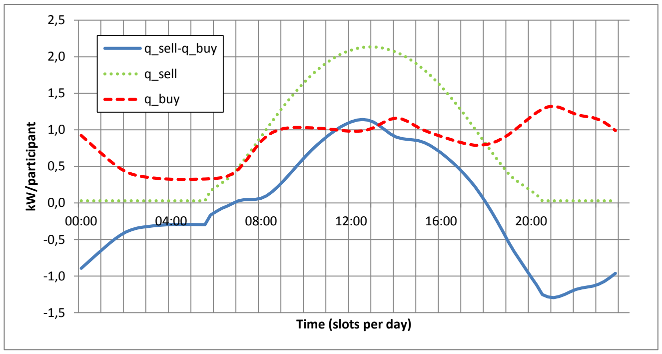

Figure 4. Aggregated production and consumption profile per agent used

Randomness is introduced for both production patterns. The PV area is randomized over an average area of $100m^2$ per agent with static radiation data (of May 2011). The wind turbine W outputs are uniformly distributed over range $[0 \ldots, 1.2]kW$. The later is calculated as $4kW * 30\%$ efficiency and hence the upper limit is $1.2kW * \frac{\delta}{60}h = 0,3kWh$ per agent. The two production types are aggregated on Figure 4. The consumption and production profiles of each agent are converted to energy based on the preconfigured slot size δ. As Figure 4 depicts, at some time the energy supply exceeds the energy demand. Although this scenario might only be considered realistic in the future (where large-scale distributed generation is in place), we use it to analyse the market behaviour.

C. Measures

The price statistics for each timeslot, as well as the allocative efficiency measure [11] were recorded for each timeslot over 365 days of simulation. Here, we measure the allocative efficiency as the ratio of traded energy to the maximum amount of energy that could have been traded in a timeslot. All measures where averaged over the same timeslot for each day of simulation. That is, for $\delta = 15min$ one day has $s = 96$ different trading timeslots. Daily data is then averaged based on the step size of $s \in \mathbb{N}^*$ as $X_j \subset X : X_j = \{x_{ds+j}\}$ where $j \in [0, s)$ represents the timeslot shift and $d \in \mathbb{N}^*$ total number of days. All the trades in the slot have trading price p, with average trade price $\bar{p}$. If expressed over X_j, we have the average price $\bar{p}^{X_j} = \langle \bar{p}^x \rangle, \forall x \in X_j$. For the average maximum price, as for the minimum, we have $\bar{p}^{X_j}_{max} = \langle \max(p^x) \rangle, \forall x \in X_j$.

In case of the allocative efficiency ε the trading quantity is considered. For each timeslot x the ε is calculated as ratio of total traded quantity $q_t otal$ and total tradeable quantity. In equation we have $\varepsilon = \frac{q_{total}}{\min(q_{buy}, q_{sell})}$ where q_{buy} and q_{sell} represent total available buy and sell quantity respectively. If expressed over X_j the average efficiency is calculated as $\varepsilon^{X_j} = \langle \varepsilon^x \rangle \forall x \in X_j$. Since it is expected that ε is affected by ratio between q_{buy} and q_{sell}, the absolute ratio r indicator is introduced. This indicator is calculated as $r = \frac{\max(q_{buy}, q_{sell})}{\min(q_{buy}, q_{sell})}$ having always $r \geq 1$. Similarly to efficiency, r can be expressed over set of timeslots as $r^{X_j} = \langle r^x \rangle, \forall x \in X_j$. The absolute ratio, r, quantifies the relationship between supply and demand. As such it is used as a guiding measure for understanding the behaviour of the price and the allocative efficiency measures.

VI. Results

A sanity test for any market model is whether the market outcome follows the rules of supply and demand. When there is more supply than demand, prices should decrease, and similarly when there is more demand than supply, prices should increase. As shown on Figure 5, this is clearly the case in our simulation.

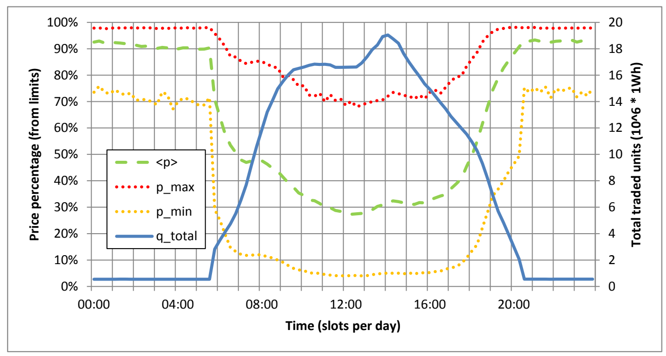

Figure 5. Market price behaviour

Figure 5 shows percentage curves of $\bar{p}^{X_j}_{min}, \bar{p}^{X_j}_{max}, \bar{p}^{X_j}, \forall j$ in respect to the ℓ_{min} and ℓ_{max} parameters. Additionally, it also depicts the difference between supply and demand on the market. As it can be seen, the price measures decrease in the region where supply outweighs the demand. In this case, the price follows the amount of supply available due to solar radiation, which peaks around each midday. The figure presents how averaged q_{total} over all X_j impacts prices on the market. Its variance highly impacts the percentage of the average price in its price window $\langle \ell_{min}, \ell_{max} \rangle$. As we observe, for $q_{total} \ll \max(q_{buy}, q_{sell})$ price averages around $\bar{p}^{X_j} \approx 90\%$, still $\bar{p}^{X_j}_{min} \approx 70\%$. In our case, when q_{total} rise to more than 10^7 traded market units (Wh), minimal averaged price $\bar{p}^{X_j}_{min} \approx 10\%$ for most of X_j. From Figure 4 one can see how the market price is correctly reflecting the high PV panel production during the day. An interesting point is when the mean price value $\bar{p}$ approximates $\langle \ell_{min}, \ell_{max} \rangle$. This is can be observed through the absolute ratio r, as will be shown that $r \approx 1$.

Although the market price fluctuations are high, the ε does not behave accordingly. The expectation is to have $\lim_{r \to \infty} \varepsilon = 1$, as shown on Figure 6. Since, in this case, there will be more than enough supply to meet the demand, or vice-versa. As depicted, the market efficiency drops down where supply meets demand ($q_{buy} \approx q_{sell}$) and stabilises back where supply exceeds demand ($q_{buy} \ll q_{sell}$). Interestingly these two drops show a strong effect on the efficiency curve, as it drops down to $\sim 77\%$ (but not lower). Thus, $\lim_{r \to 1} \varepsilon = \min(\varepsilon^{X_j}) \forall j$.

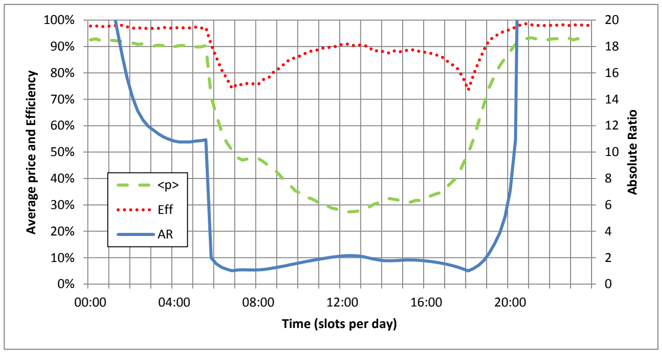

Figure 6. Market efficiency

It can be also observed, that the efficiency grows rapidly for low values of r. Figure 6 shows a high market efficiency where $e \approx 90\%$ only for $r \approx 2$. In other words demand is twice as much as the offered supply (or the other way around). This is an important point as is realistic to expect $q_{buy} \gg q_{sell}$ for current real-world deployments within the next 5–10 years. If $\varepsilon \approx 90\%$ (where demand is double wrt. to the offered supply), a seller should have a high monetary interest to be involved in local energy trading. Also an important point depicted in Figure 6 is the price behaviour, where we can see that its fluctuations impact only halfway the ε fluctuations.

The results show that the market prices follow the supply and demand. If there is more supply than demand, prices decrease, and on the other hand, if there is more demand than supply, prices increase. Even though the price fluctuations are significant, the allocative efficiency stays high. This aspect is interesting as matching is high even for low values of r, bottoming out at around 75%. This is an important outcome since the goal of the market is to facilitate trading between smart city stakeholders e.g. neighbours, therefore a high allocative efficiency would be needed in order for the market to work.

VII. Conclusion

If the trend of the distributed generation continues, local electricity trading can be technically realized in future smart grid cities. Using the discussed trading mechanisms, also market-driven demand response programs may be of interest. In local marketplaces consumers and producers can engage into energy trading for their neighbourhood, while several stakeholders may enjoy the generated benefits [1]. The efforts to design and implement such a local marketplace and investigate its impact have been analysed. The initial results show that such a market is a viable approach. It was observed how the market efficiency and the absolute ratio react to the energy surplus of its participants.

For simplicity, prediction errors were not considered in this evaluation, however, they may play an important role in evaluating economical perspectives of the market within the smart city. Since there are heavy interactions from the end-user side with the market, it is absolutely mandatory that the delegation of the user interactions to intelligent agents that can act on his behalf is realized. Hence, better strategies for those agents capturing real-world scenarios in order to get more realistic market evaluations need to be investigated. In future work it is planed to concentrate on further evaluating the market through diverse scenarios. The local conditions, the idiosyncrasy of the participants as well as the variety of acting strategies will need to be taken into account . Additionally a better understanding of how the introduced market parameters impact the overall trading efficiency for previously discussed dynamics of participants should be achieved. .

Acknowledgment

The authors would like to thank the partners of European Commission co-funded project NOBEL (www.ict-nobel.eu) for the fruitful discussions.

References

[1] S. Karnouskos, "Demand side management via prosumer interactions in a smart city energy marketplace," in *IEEE International Conference on Innovative Smart Grid Technologies (ISGT 2011), Manchester, UK*, Dec. 5–7 2011.

[2] R. Katz, D. Culler, S. Sanders, S. Alspaugh, Y. Chen, S. Dawson-Haggerty, P. Dutta, M. He, X. Jiang, L. Keys, A. Krioukov, K. Lutz, J. Ortiz, P. Mohan, E. Reutzel, J. Taneja, J. Hsu, and S. Shankar, "An information-centric energy infrastructure: The berkeley view," *Sustainable Computing: Informatics and Systems*, 2011.

[3] X. Yu, C. Cecati, T. Dillon, and M. Simões, "The new frontier of smart grids," *Industrial Electronics Magazine, IEEE*, vol. 5, no. 3, pp. 49 –63, sept. 2011.

[4] S. Karnouskos, "Cyber-Physical Systems in the SmartGrid," in *IEEE 9th International Conference on Industrial Informatics (INDIN), Lisbon, Portugal*, Jul. 26–29 2011.

[5] S. Karnouskos, P. Goncalves Da Silva, and D. Ilic, "Energy services for the smart grid city," in *6th IEEE International Conference on Digital Ecosystem Technologies – Complex Environment Engineering (IEEE DEST-CEE), Campione d'Italia, Italy*, Jun. 2012.

[6] M. Jamshidi, Ed., *Systems of Systems Engineering: Principles and Applications*. CRC Press, Nov. 2008.

[7] "SmartGrids SRA 2035 – Strategic Research Agenda," European Technology Platform SmartGrids, European Commission, Tech. Rep., Mar. 2012. [Online]. Available: http://www.smartgrids.eu/documents/20120308_sra2012.pdf

[8] Federation of German Industries (BDI), "Internet of Energy: ICT for energy markets of the future," BDI publication No. 439, February 2010. [Online]. Available: http://www.bdi.eu/BDI_english/download_content/ForschungTechnikUndInnovation/BDI_initiative_IoE_us-IdE-Broschure.pdf

[9] S. Lamparter, S. Becher, and J.-G. Fischer, "An agent-based market platform for smart grids," in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Industry track*, ser. AAMAS '10. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 1689–1696. [Online]. Available: http://dl.acm.org/citation.cfm?id=1838194.1838197

[10] P. Vytelingum, S. D. Ramchurn, T. Voice, A. Rogers, and N. R. Jennings, "Trading agents for the smart electricity grid," in *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada*, 10-14 May 2010, pp. 897–904.

[11] D. K. Gode and S. Sunder, "Allocative efficiency of markets with zero intelligence (ZI) traders: market as a partial substitute for individual rationality," *Journal of Political Economy*, vol. 101, no. 1, pp. 119–137, Feb. 1993. [Online]. Available: http://www.jstor.org/stable/2138676

[12] Solar radiation curve. [Online]. Available: http://re.jrc.ec.europa.eu/pvgis/apps4/pvest.php