

Benchmarks and Models— Price-Responsive Devices with Forward Prices *Looking Forward and Looking Back*

Chellury Ram Sastry
Samsung
c.sastry@sta.samsung.com

William Cox
Cox Software Architects
wcox@coxsoftwarearchitects.com

Toby Considine
TC9 Incorporated
toby.considine@gmail.com

Keywords: Semantics, Price Communication, Collaborative Energy, OpenADR 2, OASIS Energy Interoperation, EMIX,

Abstract

We want to answer questions about device behavior in the context of indications or actual forward prices (for future time intervals). We demonstrate how existing market information can effectively address variability in forward prices.

Demand optimization will be crucial as we move forward. To this end, we need to enable suitable pricing price-responsive end nodes¹, along with pricing mechanisms and market structures to enable robust demand optimization.

In this paper we describe and model a device that takes available price information and makes choices about time of dispatch, and address a range of issues related to how the models respond to various prices. We validate the model and present results from simulation runs using one year of real time and day-ahead market clearing prices from one North American Independent System Operator (ISO-NE).

Regulation in the United States at the retail level is typically statewide, with each state having its own policies. The responsiveness of a device should be independent of such statewide discrepancies, and must respond only based on the configuration parameters set by the consumer. This ensures widespread portability and statewide interoperability.

¹ We use the terms end node, device, and facility interchangeably in this paper. The technical abstraction is the OASIS Energy Interoperation (OASIS, 2012) Virtual End Node or VEN.

Future work includes how to make decisions balancing the risk of up-to-the-minute purchases compared to premiums for forward purchases, validation of these techniques across other market environments, use of forward indicative prices, and time as a facet of the analysis.

1. INTRODUCTION

We build a framework for quantitative evaluation of forward price knowledge and device/node behavior to further the goal of understanding and building price-responsive devices so devices can engage in transactive operation using markets or using prices set by others,

In a previous paper (Sastry, Cox, & Considine, Price Normalization, 2011) we defined formal models and examined abstractions of price, which were focused on looking backward at recent price history. In this paper we build on our price-responsive node (or device) model and create quantitative models for device behavior with future price knowledge.

We already know a number of things about electricity and energy prices—they have a diurnal, often bimodal, and seasonal patterns ((Dept of Energy, 2008) *on peak* definition)—but what if we could know future prices? How would devices take advantage of the knowledge? How can we quantify decisions we might make? What’s the best we can do with that information? Can we do reasonably well with what we have, or are new price streams and/or markets needed?

To that end, our framework contributes to determining the economic value of

- Different kinds of forward price information
- Different time extent of future knowledge

- Application of that knowledge to future prices
- Algorithms to take advantage of forward price knowledge

2. PRICES, RESPONSE, AND ACTIONABLE INFORMATION

First we must talk about price, price streams, and response.

2.1.1. Semantics of Price

Another paper (Sastry, Cox, & Considine, Semantics of Price, 2012) includes a more complete discussion of the semantics of price. In this paper, we avoid consideration of the economic reality of prices to which a price-responsive node (or device) responds, except to use those prices in evaluating performance.

The key semantic terms we use are *transactable* and *non-transactable*, meaning respectively that a device may buy energy at the stated price (transactable) or may not.

The latter are sometimes called *indicative* prices, but without a connection to the economic reality of the node's environment, the actionability of indicative prices is questionable.

We will assume that there is an economic consequence for accepting or not accepting a price, and that a delay in execution that requires energy will take the price for the time of consumption.

2.2. Price-Responsive Nodes

The figure of merit for responsive devices is total price paid. This also allows comparison to a baseline set of prices. If the set is contiguous, and the time intervals are the same size, we call this a *price stream*.

We will assume that a price stream is available to the node, and that the node will make a choice:

- (a) Run now
- (b) Run later, and determine how much later

Following the model of (Sastry, Cox, & Considine, Price Normalization, 2011), we build a framework to model behavior of an energy-consuming device, a device that uses one unit of energy for one time interval in the stream. The results may then be applied to devices with varying energy use and run schedules.

2.2.1. End Nodes, Facilities, Devices, and Virtual End Nodes

Following the terminology of Energy Interoperation 1.0 (OASIS, 2012), we look at pairwise relationships. For demand response [OpenADR], we name the initiator of an

event the *Virtual Top Node* (VEN) and the recipients *Virtual End Nodes*.

While market and transactive operations do not necessarily use the same interaction graph that would be used for demand response, calling the recipients of prices (whether they be the result of market operations, operations in which the node might or might not participate) *Nodes* or *Virtual End Nodes* is reasonably consistent with the Energy Interoperation standard terminology.²

A price and product definition has meaning in a specific Market Context as defined in Energy Interoperation and Energy Market Information Exchange (OASIS, 2012)

2.2.2. Actionable Information

To be useful, information must provide information to guide actions. We call this *actionable* information.

The actions in our model are as stated in Section 2.2. There is a measurable value to not acting when (e.g.) prices are high, and to acting when (e.g.) prices are low. We will not consider other costs such as the cost for delay, only the value of delaying to obtain better prices.

2.2.3. Node Details

Factoring out quantities and units, both of energy and time, we describe a price-responsive node as one that changes its behavior when prices differ.

Our framework could be applied to demand markets and prices, but in this paper we use only energy markets and prices.

As stated, device uses a single unit of energy for one time period; the time period is the basic interval size for the price stream.

2.3. Look Forward, Look Back

We will describe pseudocode (Wikipedia) for a price-responsive device. The device will analyze information, perhaps looking forward, back, or both, in a price stream (see Figure 1 below) relevant to its business, and apply its business information, requirements, and goals in deciding what action to take.

² The actors that interact with markets are called Parties and likewise the recipient of price information is also called a Party.




Figure 1. Price Stream Consisting of Historical, Real-Time, and Future prices

The retrospective analysis described in (Sastry, Cox, & Considine, Price Normalization, 2011) included abstractions of past prices into simple levels; those were designed used to determine whether a current price is high, low, or at some intermediate state.

Levels may be interpreted as nominal prices, though disconnected from actual costs, and that the same node algorithms could use an abstracted price stream.³

Past information is, in essence, not transactable, as the time and circumstances have passed. There are, however significant markets (called *forward markets*) that allow transactions or allow a viewer to understand future prices.

2.4. Discussion

Knowledge of the future makes the results of actions more predictable. As such, forward price knowledge is a big step toward efficiently operating to minimize energy cost.

We look only at energy costs, not at other costs that may be included in an electric bill. Since the energy costs are variable and a primary determinant of total cost, we do not need that refinement for this discussion.

3. TERMINOLOGY

In this section we summarize the terminology we have already introduced. Need the rest.

3.1. Abstractions and Nominal Prices

Common approaches to conveying prices involve their abstractions such as a set of levels, where each level corresponds to a range of prices. However, in [Sastry, Cox, Considine, Grid-Interop 2011], we argue that such abstractions lead to inconsistent device behavior, and is fraught with lack of interoperability. We further argue that actual (nominal) prices must be conveyed in a standardized and consistent manner.

³ For useful behavior, the actual price should be available; hence we use actual prices even though they may not be the prices that the end node actually pays.

3.2. Locational Marginal Price

A Locational Marginal Price (LMP) is a price for a particular time interval at a specific location. There may be hundreds of nodes in a market operator’s territory, and thousands addressed in multi-market forward exchanges.

3.3. Time Series—Streams

Price streams are a specific example of the streams formally defined elsewhere [Energy Interoperation] as extensions to iCalendar [iCalendar], which itself is the basis of nearly all personal calendars and scheduling.

Mathematically, a stream resembles a time series of data; there are many statistical tools that may be applied to time series [The R Project for Statistical Computing].

3.3.1. Substitution of Streams

For our discussion it is irrelevant to the node precisely what the source of a price stream is; the behavior is similar, and the application to forward or present-time prices is explicitly stated in an Energy Interoperation/EMIX stream [].

The figure below (Figure 2 from [Sastry, Cox, Considine, Grid-Interop 2012]) shows the separation between an end node and source of a price stream.




Figure 2. Price Stream to a VEN

In Figure 2, Node A (a VEN, facility, or device) receives a price stream from one or more sources (we show one source, B). The blue lines indicate interfaces to the VEN and price source. C labels interoperation between end nodes and price source.

3.3.2. Two Streams

We use two price streams, both from the New England Independent System Operator (ISO-NE).

The streams use one-hour intervals, and give a price that applies to each interval. The prices in the streams are clearing prices in two different markets maintained by ISO-NE and obtained from their historical data. They are

- (1) The *Day Ahead* market (DA)
- (2) The *Real Time* (or same-day) market (RT)

These are published and available going back several years. We selected July 1, 2010 through July 31, as we have additional data sets that cover that time with abstractions of price (Sastry, Cox, & Considine, Price Normalization, 2011) at different but compatible granularity for that time period (see Future Work)

For this discussion it suffices to know that the DA market clears for a given date at 5pm the day before; the real time markets clear close to the start time of the five-minute interval referenced.

4. THE DATA SETS AND CHANGES

The source is (ISO-NE, 2012). Much information irrelevant to our purpose is in those files. Note that the prices are not transactable, as they reflect market clearing, and as recorded are preliminary, not final settlement prices.

First, we restrict ourselves to only considering a single [ISO-NE] Zone, SEMASS for Southeast Massachusetts.

The prices are only indirectly related to end node prices, and are expressed in dollars per megawatt-hour.

To line up with another data set we used we restricted the dates from July 1, 2010 through July 31, 2011.

We used the *2010_smd_hourly* and *2011_smd_hourly* datasets, and performed the following adjustments and restrictions:

- Hourly was the only granularity real-time data available.
- Date was changed to a single date & time format rather than separate day/hour
- Only rows for Zone 4006 Z.SEMASS were used (restricting to one Zone LMP)
- We retained the Day-Ahead Locational Marginal Price and Real-Time LMP
- All sub-components of the LMP were discarded; the Energy Component is nearly all of the LMP, so we used the LMP.
- The data for hour 2 of March 13, 2011 (the omitted hour for daylight savings time) was replaced with the average of the values immediately before and immediately after.

The information was approximately 9,500 records in each data set plus extras retained at the end to simplify look-ahead computations.

Analysis was done with Microsoft Excel and the R Statistics Package (R Statistical Package, 2012).

5. THE BENCHMARK AND THE EXPERIMENTS

5.1. Knowledge of the Future

We want to establish economic, quantitative value for various forms of forward price information, as well as a framework in which to evaluate various node management algorithms.

We are using only forward prices, with varying degrees of consideration of future prices.

5.2. Horizon

We define the term *horizon* to mean how many price stream intervals in the future we will consider, as shown in Figure 3



Figure 3. Price Stream Intervals

The Horizon is how far forward prices will be considered, including the current time interval. Assume we are at the time interval 1. A horizon of 1 is just the current interval. A horizon of 4 considers the prices for intervals 1 through 4; a horizon of 8 considers intervals 1 through 8, and so forth.

5.3. Benchmark for a Horizon

We define a means of determining optimal behavior. This parallels the development of *Belady's Algorithm* (Belady, 1966) (ERCOT, 2012) for disk head scheduling, where complete knowledge of the incoming requests is available. For the benchmark we assume perfect knowledge of near-term forward prices.

Adapting Aho's description of Belady's algorithm, we assume we know the price stream

$$\text{Price Stream } P = p_{i_1}, p_{i_2}, \dots$$

We then formally define the horizon to be the maximum value of i that applies to time t .

We define the benchmark price B_t^h for a time t and a horizon h as follows:

$$B_t^h = \text{Minimum}(p_1, p_2, \dots, p_h)$$

In other words, the benchmark price is the lowest price within the horizon.

In the rest of this paper we choose the first such block if there are ties. This addresses the common case where known future prices are determined (e.g.) by time of use; this gives a bias to the earlier, and reduced computed delays in our models. In dynamic price streams we would expect that ties will be rare.

For prices in a horizon, we often use the offset:

$$Offset = first\ i\ \exists\ p_i = B_t^h$$

We will use this to select the correct prices in Section 6.3.

6. EXPERIMENTS

We carried out three experiments detailed in the following sections with results and discussion.

The three experiments are

1. RT—Real-time (hourly) prices are used. The benchmark price is computed for each time interval
2. DA—Day-ahead (hourly) prices are used. This forward market clears typically at 5pm the day before the time intervals for which prices are determined.
3. DAselRT—Day-ahead prices are used to determine the interval to use; the RT prices are then used to compute the price actually paid.

In the following sections we describe each experiment, and display and discuss the results.

6.1. RT: Benchmark Prices for the RT Data

6.1.1. Description

We define device behavior with the following pseudo-code:

Loop

 Consider all prices from the interval just starting to the event horizon

 Select the lowest price

 Delay until the interval associated with the first occurrence of the lowest price

Repeat

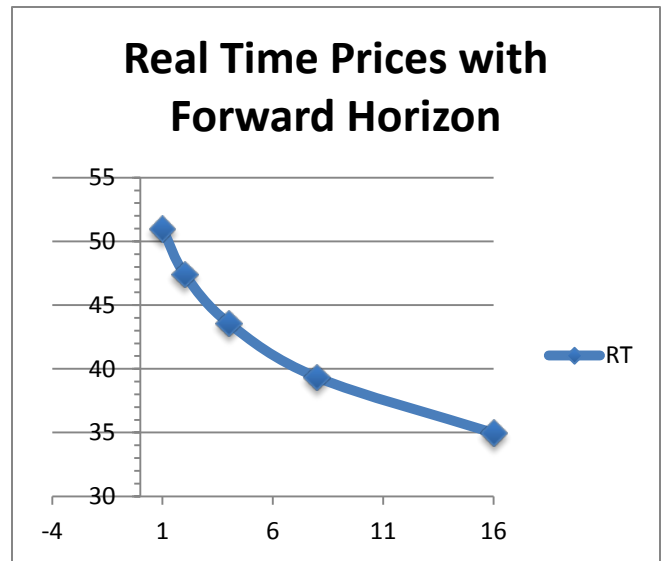
We look at the optimal choice for each interval in the data set, extending the data values two days past the end to allow for additional horizon values to be assessed.

A Benchmark takes two parameters: the data set with its associated product type, and the horizon.

Evaluate the mean, standard deviation, maximum, and minimum for each horizon in 1, 2, 4, 8, and 16.

6.1.2. Results

The results are presented in the graph and table below.



Horizon	RT Mean
1	51.0
2	47.4
4	43.6
8	39.3
16	35.0

6.1.3. Discussion

As one would expect, the mean price to be paid decreased with larger horizon values. The curve is very nearly a section of a parabola; when plotted with logarithmic X-axis the curve is nearly a straight line.

The mean for Horizon 1 is exactly that for simply taking the price in the interval which one considers first. The maximum benefit in this graph is at horizon 16, and gives around 31% benefit over taking the current price without delayed operation.

6.2. Experiment 2: Benchmark Prices for the DA Data

6.2.1. Description


We ran a similar experiment for the Day-Ahead market.

The semantics of the market provide that prices for today were determined in the market yesterday (typical closing time is 5pm the day before).

The information is therefore, with a 5pm closing, available for between 7 and 31 hours ahead depending on the hour of the day that we consider interval 1.

6.2.2. Results

The results are presented in the graph and table below. We include the RT series for comparison.



Horizon	DA Mean
1	50.7
2	48.8
4	45.8
8	41.6
16	36.8

6.2.3. Discussion

The DA curve is completely above the RT curve except at horizon 1.

The RT curve takes into account unforeseen (at least not foreseen the day before) contingencies, unplanned congestion, and so forth, all of which have the effect of making the RT stream more volatile (the sample standard deviation for RT at horizon 1 is 27.8; for DA it's 22.6).

However the contingencies are relatively infrequent, and the excursions above and below the mean are more extreme, so more advantage can be taken of the volatility.

6.3. Experiment 3: Benchmark Prices using DA to Determine RT Price Applied

6.3.1. Description

This is our first experiment to quantify the value of forward information from a less reliable source than the benchmarks. We are motivated in part by the recent decision of ERCOT to provide some actionable forward prices to DR providers through what they refer to as an "Hour Ahead Market" (HAM). The HAM is a new market which will be installed between the DAM and the RTM [(ERCOT, 2012)].

We combine information from the DA and RT datasets as follows:

For a given time T_i determine the offset from the current time to the interval with the lowest Day-Ahead price.

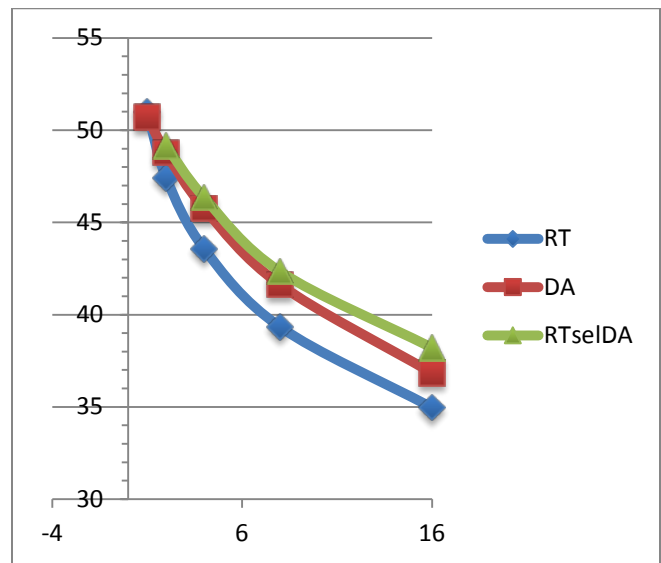
Then for that interval, we utilize the the current Real-Time price for that interval as the actionable information.

This mimics using readily available Day-Ahead information to determine how much delay is appropriate. Forward real-time price information may be much more difficult to determine and less reliable even if it were; the Day-Ahead clearing prices are easy to use.

The question we want to answer is "how well can I predict the future using yesterday's day-ahead markets?"

6.3.2. Results

The results are presented in the graph and table below.



Horizon	RTselDA Mean
1	51.0
2	49.2
4	46.4
8	42.3
16	38.2

6.3.3. Discussion

The value for horizon 1 is clearly the same as that for RT.

One interesting fact is that the mean for RT with horizon=4 is only slightly above the mean for either DA or RTselDA at horizon=8.

This tells us that in terms of average price paid, the mean cost for horizon=8 for DA and RTselDA are both greater than the mean cost for horizon=4 for RT.


This suggests that if forward real-time market information is not available, or is not very reliable (see Future Work below), then a reasonable substitute are the Day-Ahead prices.

Recall that the Day-Ahead prices have a 7 to 31 hour advance timeframe, so in effect the mean amount paid at h=7 is higher than the mean at higher horizons.

We can conclude that Day-Ahead clearing prices at h=8 or greater provide a better mean amount paid than accurate h=4 for Real-Time prices.

6.4. Consider Expected Delay and Second Moments

We did not directly show the standard deviations for the means in the previous sections. The graph below shows the differences in the standard deviations with the same color code as in the other graphs. (In the final version of the paper, we present a more detailed analysis including graphs for standard deviations).



As we mentioned above, the variation in the RT set is higher than in the DA for horizons up to 2. The variation for the third set (RTselDA) is higher than either of the others.

The increased variation is somewhat troublesome, as it increases the risk.

7. FUTURE WORK

In this section we briefly mention possible future work based on the ideas and implementations of the ideas in this paper.

7.1. Facets—Time of Day

Questions to consider include

Does the range of times affect the value? For example, appliances may primarily delay during the time from going to work until returning home, and from going to bed to arising. What are the implications from time-related patterns and facets?

7.2. Different Stream Intervals

It may be that running the same tools and simulated algorithms on five-minute data will produce different results, possibly due to the included time being shorter.

While there are efforts to prove the self-similarity of energy use (see e.g. [Leland, IEEE Comm Soc] for self-similarity of communication traffic), this is an open area that we may investigate at stream intervals that are currently provided.

7.3. Different Markets and Conditions

Applying the same analytic tools to price streams in other markets and at different times would be straightforward. Validate across a larger set of markets, times, and weather conditions.

7.4. Indicative Prices

Information on forward prices may be reduced from full knowledge. One way people have described this (Centolella, 2012) is to obtain *indicative* prices.

These might have sources such as interim simulated attempts at market clearing in advance of deadlines, or estimation of future direction.

How do the model and its outputs respond to partial information? Directional information is clearly useful with horizon=2, as the choice to run now or delay an hour is guided clearly.

But how accurate would the directional indications be? The model and framework would allow testing indications just as well as testing already-know (but “future” to the point analyzed) prices.

It would be interesting to compare the cost of creating and distributing indicative prices with the cost of using the day-ahead information as in RTselDA.

7.5. Balancing Risk and Costs of Purchases

Estimates of risk can be computed from the sampling and statistical techniques we have used. Consider how to apply those estimates to the buy ahead (usually at a price premium) versus buy at time of use.

8. SUMMARY AND CONCLUSIONS

We have demonstrated that there is value to be gained from forward price information, and have quantified that value.

We have graphically shown how to exploit forward price information and the value of that exploitation.

We have discussed how to extend and apply the work to address related problems in price-response.

By exercising our framework and the relevant horizon determination, we have shown a simple way to determine the first and second moments of the means for two different price streams and two different algorithms.

The framework in its current form easily computes the output values for any time series.

In the final version of the paper, we will show additional graphs with error envelopes. We are also working on an R function to do the computation direction, parameterized by horizon values.

This framework for quantitative analysis allows testing of algorithms and forward price information prior to large-scale deployment, and should serve as a price, product, and algorithm design tool.

8.1. Implications

We have demonstrated that transactable prices delivered using existing standards provide an actionable basis for optimizing energy purchases. Assumptions about the nature and volatility of the local power market can be reduced to reasonably simple and reasonably useful predictive algorithms. The particular assumptions are inherently local, and are outside the scope of this paper.

Different Virtual End Nodes will place different values on aspects of the market, and the value of price and availability arbitrage against that market. Because of the cleanliness of the model, it should be easy to develop applications, even on resource-constrained embedded devices.

This approach removes objections to participation in transactive markets by devices, even very small devices. Such participation could provide a path from today's large, ordered markets to radically different power markets, even microgrids, in the future.

In such markets, distributed power sources, even intermittent power sources such as local photovoltaic (PV) or Wind, can bring their product to bear in device oriented local markets.

9. REFERENCES

- Aho, A., Denning, P., & Ullman, J. (1971). Principles of Optimal Page Replacement. *Journal of the Association for Computing Machinery*, 18 (1).
- Belady, L. A. (1966). A Study of Replacement Algorithms for Virtual Storage Computers. *IBM Systems Journal*, 5 (2), 78-101.
- Centolella, P. A. (2012, October). *A Pricing Strategy for a Lean and Agile Electric Power Industry*. Retrieved from Electricity Policy: <http://ElectricityPolicy.com>
- Considine, T., & Cox, W. (2009). Smart Loads and Smart Grids--Creating the Smart Grid Business Case. *Grid-Interop 2009*. GridWise Architecture Council.
- Considine, T., & Cox, W. (2012). Understanding Microgrids as the Essential Smart Energy Architecture. *Grid-Interop (Accepted Paper)*.
- Cox, W., & Considine, T. (2011). Energy, Micromarkets, and Microgrids. *Grid-Interop 2011*. GridWise Architecture Council.
- Cox, W., & Considine, T. (2010). Structured Energy: A Topology of Microgrids (Presentation Only). *Grid-Interop*. GridWise Architecture Council.

Dept of Energy. (2008). *The Smart Grid: An Introduction*. Department of Energy.

ISO-NE. (n.d.). Retrieved October 15, 2012, from ISO New England: <http://www.iso-ne.com/>

ISO-NE. (2012). *Historical Data*. Retrieved October 15, 2012, from <http://www.iso-ne.com/markets/hstdata/index.html>

OASIS. (2012). *Energy Interoperation 1.0*.

OASIS. (2012). *Energy Market Information Exchange Version 1.0*.

Piette, M. A., Ghatikar, G., Kiliccote, S., Koch, E., Hennage, D., Palensky, P., et al. (2009). *Open Automated Demand Response Communications Specification (Version 1.0)*. California Energy Commission, PIER Program.

Sastry, C. R., Cox, W., & Considine, T. (2011). Price Normalization for Price-Responsive Devices--Algorithms and Issues. *Grid-Interop*. Phoenix: GridWise Architecture Council.

Sastry, C. R., Cox, W., & Considine, T. (2012). The Semantics of Price. *Grid-Interop*. GridWise Architecture Council.

The R Project for Statistical Computing. (2012). Retrieved 2012 from <http://www.r-project.org/>

Wikipedia. (n.d.). *Pseudocode*. Retrieved October 15, 2012, from <http://en.wikipedia.org/wiki/Pseudocode>

ERCOT. ERCOT Considering an Hour Ahead Market. Retrieved October 21, 2012 from <http://beta.demandresponseinfo.org/library/posting/4798/>

iCalendar. Retrieved October 21, 2012 from <http://en.wikipedia.org/wiki/ICalendar>

Energy Interoperation. Retrieved on 10/22/12 from http://docs.oasis-open.org/energyinterop/ei/v1.0/cs02/energyinterop-v1.0-cs02.html#_Toc319920796

BIOGRAPHY

Chellury Ram Sastry

Dr. Chellury Ram Sastry is currently a Senior Manager at Samsung Telecommunications America LLC, focusing on various advanced R&D thrust areas including machine to machine communication, smart/connected home, and smart energy technologies including interoperability and standards.

Prior to joining Samsung, Ram was a Smart Grid Program Director with the Energy, Environment, and Material Sciences division at Battelle Memorial Institute, Columbus, OH. He was also with the Electricity Infrastructure Group in the Energy and Environment Directorate at Pacific Northwest National Laboratory (PNNL), Richland, WA managed and operated by Battelle. He was responsible for providing R&D, business development, and technical marketing leadership in various thrust areas including advanced smart grid enabled demand management to provide value-add services to residential and small commercial building customers, transmission/distribution modeling & simulation, and smart grid data analytics.

Prior to joining Battelle/PNNL, Ram was a Project Manager and Senior Research Scientist with Siemens Corporate Research (SCR), Princeton, NJ. One of the highlights of his tenure at SCR was an R&D program he was responsible for to enhance the product portfolio of various Siemens businesses (smart homes, remote health care, industrial automation etc) based on radio frequency identification (RFID), wireless sensor networks, and embedded machine-to-machine technologies.

He has published several papers in refereed journal and conference proceedings, and has been a plenary speaker at well-known conferences including Connectivity Week, Grid-Interop etc. He also has several patents against his name, and a number of provisional patent and patent applications under consideration.

Ram has a B.S. degree in electrical engineering from Indian Institute of Technology, Chennai, India M.S. /Ph.D. degrees in electrical engineering and an M.A. degree in Mathematics from University of Pittsburgh.

William Cox

William Cox is a leader in commercial and open source software definition, specification, design, and development.

He is active in the NIST Smart Grid Interoperability Panel and related activities. He contributed to the NIST conceptual model, architectural guidelines, and the NIST Framework 1.0.

Bill is co-chair of the OASIS Energy Interoperation and Energy Market Information Exchange Technical Committees, past Chair of the OASIS Technical Advisory Board, a member of the Smart Grid Architecture Committee, and of the WS-Calendar Technical Committee.

Bill has developed enterprise product architectures for Bell Labs, Unix System Labs, Novell, and BEA, and has done related standards work in OASIS, ebXML, the Java Community Process, Object Management Group, and the IEEE, typically working the boundaries between technology and business requirements.

He earned a Ph.D. and M.S. in Computer Sciences from the University of Wisconsin-Madison.

Toby Considine

Toby Considine is a recognized thought leader in applying IT to energy, physical security, and emergency response. He is a frequent conference speaker and provides advice to companies and consortia on new business models and integration strategies.

Toby has been integrating building systems and business processes for longer than he cares to confess. He has supported and managed interfaces to and between buildings, cogeneration plants, substations, chilled water plants, and steam and electrical distribution. This work led to Toby's focus on standards-based enterprise interaction with the engineered systems in buildings, and to his work in the Organization for the Advancement of Structured Information Standards (OASIS).

Toby is chair of the OASIS oBIX Technical Committee, a web services standard for interface between building systems and e-business, and of the OASIS WS-Calendar Technical Committee. He is editor of the OASIS Energy Interoperation and Energy Market Information Exchange (EMIX) Technical Committees and a former co-Chair of the OASIS Technical Advisory Board.

Toby has been leading national smart grid activities since delivering the plenary report on business and policy at the DOE GridWise Constitutional Convention in 2005. He is a member of the SGIP Smart Grid Architecture Committee, and is active in several of the NIST Smart Grid Domain Expert Workgroups.

Before coming to UNC, Mr. Considine developed enterprise systems for technology companies, apparel companies, manufacturing plants, architectural firms, and media companies old and new. Before that, Toby worked in pharmaceutical research following undergraduate work in developmental neuropharmacology at UNC.