A novel online energy management solution for energy plants

Mads Almassalkhi *(IEEE Member)*, Björn Simon, Archie Gupta

***Abstract*—Energy plants represent large energy consumers with a wide array of energy needs, assets (e.g. boilers, chillers, storage, on-site generation), and constraints on operations. An innovative energy management system for energy plants is presented in this paper. Through predictive optimization of plant assets, energy analytics, pricing signals, and historical and real-time data, the energy management system supplies energy plants with salient hourly, real-time recommendations and enables "what-if" analysis to achieve improved economic efficiency. Within a systems context, the paper draws upon ideas from power systems and highlights technical issues related to plant optimization. The paper also describes actual implementations of the energy management solution at two energy plants in the US, providing economic details and an analysis of the savings achieved.**

***Index Terms*—cogeneration, economic efficiency, energy plant, energy analytics, HVAC, multi-energy system, predictive optimization, tariffs.**

I. Introduction

UNIVERSITY campuses and industrial facilities with energy plants (EPs) represent large and complex multi-energy systems that both consume energy from external suppliers and produce their own energy needs. In an EP, the combination of direct-fired boilers, electric on-site power generation, heat-recovery mechanisms, pumps, compression and steam-absorption chillers, regulation on emissions, and a deregulated fuel markets with varying prices precludes real-time back-of-the-envelope calculations by an operator. Instead, the heavy technical lifting is usually performed by a multitude of external contractors who visit a customer plant every couple years to conduct capital planning studies and instruct operators on an ad-hoc set of rules (i.e. "rules-of-thumb") that permit economic efficiency and ensures reliability. However, the realities of maintenance, changes to equipment through repairs and upgrades, and the emphasis on reliability at the plant-level tends to undermine most efficiency measures and can rapidly make obsolete the work of the costly contractors. In addition, due to the multi-energy nature of EPs and since most contractors are specialized in one or two types of equipments or services, the recommendations from one contractor may compete against the rules of another contractor. Therefore, there is a need for contractors to be available more often (i.e. on-site) and to collaborate with each other on recommendations, which is cost-prohibitive. Instead, through multi-energy system models, predictive optimization algorithms, and energy analytics, we propose a holistic solution that provides real-time economically efficient recommendations, satisfies reliability requirements, and adapts to changes in the plant. We call the solution *Balance*. Balance empowers EP operators, engineers, and managers to conduct what-if studies themselves by providing actionable information at their fingertips and uncover and avoid hidden costs associated with complicated tariffs.

The idea of optimization and simulation is not new within power systems, where asset management is termed Unit Commitment (UC) and Economic Dispatch (ED) and the "assets" are generators and loads [1]. However, within the context of energy plants, online scheduling of assets is a novel and nontrivial application of optimization. In this work, Balance is not just scheduling assets as with UC nor directly setting the respective output/production levels as in a standard ED setting. The main differences between EPs and general electrical power systems lies at the multi-energy couplings and the type of actionable events. For EPs, for example, one does not set a chiller's cooling output, but rather manage a collection of set-points (e.g. temperatures), which determine specifically *how* energy is transferred through the system (i.e. via energy and mass balances). This significant difference means that Balance has to, for example, explicitly consider the physical interconnections between chillers, cooling towers, pumps, boilers, heat-recovery assets, etc, which naturally lends itself to multi-energy system analysis and modeling.

Multi-energy system analysis and simulation have been studied since at least the 1980's [2]. However, one of the most promising multi-energy methodologies have only been developed recently and is denoted the Energy Hub and is a modeling framework wherein conversion and storage processes are explicitly considered for the purposes of energy flows [3], [4]. Previous applications of energy hubs, however, have focused on high-level planning studies with simplified models and not online operations where high-fidelity energy models and awareness of real-time pricing signals are required. Therefore, this paper presents a state-of-the-art extension of energy hubs within the domain of / applied to energy plants. In addition to energy hubs, recent work has applied model-predictive control for thermal energy storage within an EP setting [5], but this work suffers from scalability issues as it was designed exclusively for the given customer and does not provide a meaningful interface nor an opportunity for energy analytics. Another suite of tools are called enterprise energy management systems (EEMs) [6]. EEMs provide data storage and enables *ex-post* analysis of billing data, but does not consider plant assets nor real-time operations.

As far as the authors are aware, no service currently integrates energy analytics, real-time pricing, and predictive optimization for the purposes of economic efficiency and capital planning (i.e. "what-if" analysis) within the EP domain, which is the main contribution of this paper and the product Balance.

The paper is organized as follows: Section II discusses the systems considerations, including tariffs and asset models. The

The authors are with Root3 Technologies Inc., Chicago, IL, and Ann Arbor, MI. {mads,bjoern,archie}@root3tech.com

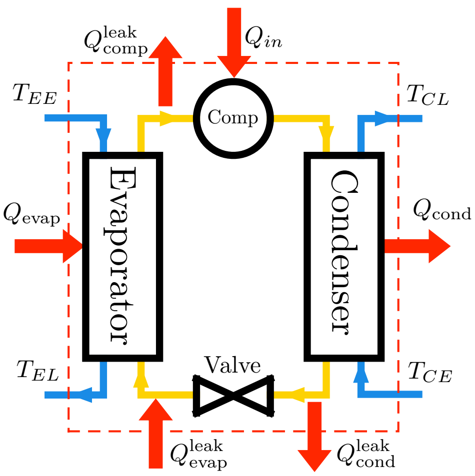

Fig. 1. Diagram of compression chiller with refrigerant (orange), water (blue), and energy (red) flows.

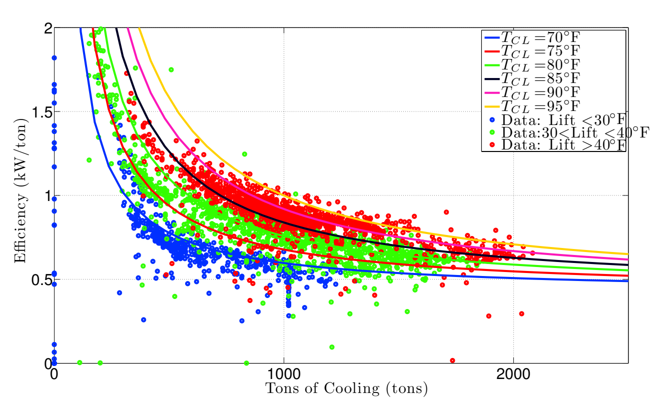

Fig. 2. Performance validation data of asset model for a VSD chiller highlighting the effect of part-load and lift on efficiency.

specific energy management services provided by Balance is described in Section III whereas Section IV discusses results from actual EP sites in the US. The paper is concluded in Section V with a description of future work.

II. System Considerations

An energy plant (EP) is a large energy consumer with a collection of energy consuming and producing assets that generally convert one form of energy into another and may even have ability to store energy. For example, boilers are utilized to convert natural gas into hot water for heating purposes. In addition to energy efficiency (i.e. minimal energy required to meet load) and the conversion and storage processes in an energy plant, consideration of energy tariffs is required to determine *economic efficiency* (i.e. minimal cost to supply load). A basic outline of general energy tariffs and relevant energy assets is presented below. For details on assets and tariffs, we refer the reader to [7], [8], [9].

A. Tariffs

While energy tariffs vary broadly among large energy consumers, a typical tariff (electricity or natural gas) consist of at least the following two cost items:

- **Energy charges**: is the cost per unit energy (e.g. \$0.08/kWh or \$5.00/mmBtu) and can be subject to variable time-of-use (TOU) rates (e.g. on/off-peak). Note that energy charges could also be supplied by real-time markets.
- **Demand-charges**: is the cost per unit power (e.g. \$10.00/kW) for a billing cycle and is usually only incurred after exceeding a minimum level of demand. For example, electric demand charges may only occur on demand that exceeds 8000 kW. The monthly or annual demand-charge is determined based on peak usage over a sampling period.

We will discuss application of tariffs within economic analysis of EPs in the next sections.

B. Chillers

Chillers take advantage of evaporative and condensing attributes of refrigerants to transfer heat and supply chilled water to cooling loads. A general chiller diagram is shown in Figure 1. The cooling provided by the chiller is based on the flow-rate of water ($\dot{m}_{evap}$) and the chilled water temperature change ($T_{EE} - T_{EL}$) across the evaporator. The heat transferred at the evaporator must be rejected at the condenser. The evaporator and condenser heat transfers are described as follows:

$$\dot{Q}_{evap}[k] = \dot{m}_{evap}[k]c_p(T_{EE}[k] - T_{EL}[k]), \tag{1}$$
$$\dot{Q}_{cond}[k] = \dot{m}_{cond}[k]c_p(T_{CL}[k] - T_{CE}[k]) \tag{2}$$

where c_p is the specific heat capacity of water (e.g. 0.0417 tons/gpm-F) and $T_{E/CE}$ and $T_{E/CL}$ are evaporator/condenser entering and leaving water temperatures, respectively. The energy balance for a chiller is then defined as:

$$\dot{Q}_{cond}[k] = \dot{Q}_{evap}[k] + \dot{Q}_{in}[k] + \dot{Q}_{loss}[k], \tag{3}$$

where $\dot{Q}_{in}$ represents the energy required to transfer the heat between evaporator and condenser through the refrigerant (e.g. compressor power, steam for absorption chiller). The term $\dot{Q}_{loss}$ represents losses within the chiller and include factors such as refrigerant leakage.

The efficiency of a chiller is often prescribed in terms of kW/ton. Generally, the efficiency of a chiller decreases at part-load and under high *lift* conditions, as shown in Figure 2. Lift is defined as the refrigerant pressure differential across evaporator and condenser and represents the work required from a compressor. A good indicator of lift is the temperature-differential: $T_{CL} - T_{EL}$, which can be managed through $\dot{m}_{cond}$ and will be discussed within the context of asset management later in the paper.

1) Cooling Towers: The heat rejected by the condensed refrigerant into the water, $\dot{Q}_{cond}$, must be rejected to the atmosphere. One of the most common means of heat-rejection is cooling towers. A cooling tower is a heat-exchanger that uses ambient air to cool the water from T_{CL} back down to T_{CE} (i.e. the tower *range*). A limiting constraint is that

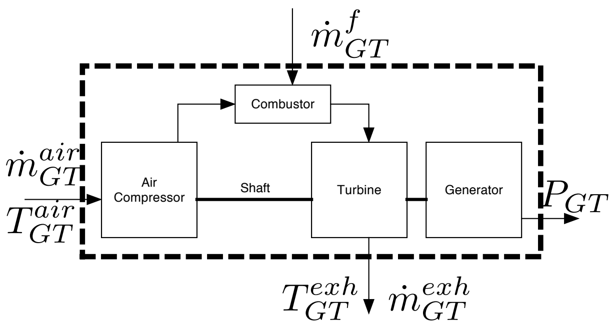

Fig. 3. Diagram of natural gas turbine generator with material flows.

$T_{CE} > T_{WB}$, where T_{WB} is the ambient wet-bulb temperature. The difference $T_{CE} - T_{WB}$ represents the tower *approach* and is generally greater than 2-4°F and represents one of the main drivers of cooling tower fan power.

2) Pre-cooling chilled water plant: By considering a lumped model of the chilled water in the pipes, it is possible to take advantage of the plant's inherent ability to store energy in the chilled water by initially producing very cold water (e.g. 38°F) and then coast (i.e. turn off chillers) to avoid an electric demand peak or expensive TOU rates. As the chilled water plant coasts, the chilled water temperature increases based on the load conditions and chillers can be brought back online in time to prevent the temperature from getting too high (e.g. 46°F).

C. Boilers

Boilers utilize natural gas through a combustion process to heat up feed-water and generate steam (or hot pressurized water). The efficiency of a boiler is nonlinear across its range and generally improves as load increases towards the design operating level after which efficiency decreases slightly. From a systems perspective, a boiler can be represented by a nonlinear one-input/one-output energy converter akin to an energy hub:

$$Q_{steam}[k] = f_B(Q_{fuel}[k]), \tag{4}$$

where Q_{fuel} is converted to Q_{steam} with boiler efficiency defined by $f_B(Q_{fuel}[k])$ and varying with the input level.

D. Natural Gas Turbine Generator

From a systems perspective an (air-breathing) natural gas turbine takes natural gas (and air) as an input and produces electricity and hot exhaust gasses through combustion and turbine processes. Figure 3 illustrates material flows through a gas turbine and the energy and mass balance of a gas turbine are given by the following:

$$0 = \dot{m}_{GT}^{exh}[k] - \dot{m}_{GT}^{air}[k] - \dot{m}_{GT}^{f}[k] \tag{5}$$

$$\begin{aligned} 0 = \dot{m}_{GT}^{air}[k] c_p^{air} T_{GT}^{air}[k] + \dot{m}_{GT}^{f}[k] LHV_f \\ - \dot{m}_{GT}^{exh}[k] c_p^{exh} T_{GT}^{exh}[k] - P_{GT}[k]. \end{aligned} \tag{6}$$

Note that $\dot{m}, T$ and P represents mass flow rate (exhaust gas, inlet air, and fuel), temperature, and power generated from the gas turbine, respectively. The constants c_p and LHV represent specific heat capacities of exhaust gas and air [Btu/lbm-F] and natural gas' lower heating-value [Btu/lbm], respectively.

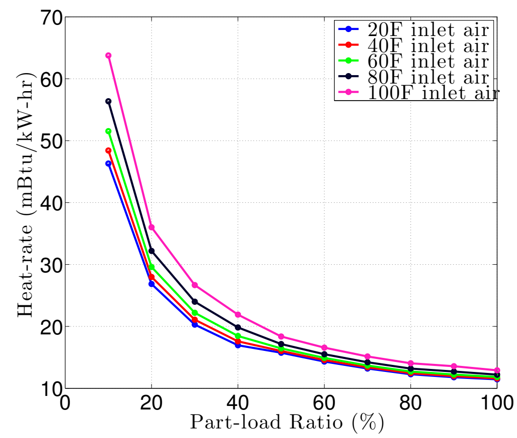

Fig. 4. Actual heat-rate of 8 MW turbine at 5000 feet altitude for different part-load and inlet air conditions.

The efficiency of a gas turbine is often given in terms of the heat-rate (i.e. how much fuel is required to produce a kW) and is strongly coupled to the following items:

- **Part-load operation**: When the turbine is operated at part-load (also called off-design operations), fuel efficiency decreases rapidly below 60% capacity. At low capacities (e.g. $< 50\%$), emissions become a factor and may prevent the turbine from being run.
- **Inlet air conditions**: A large compressor is required to provide the combustion process with compressed air. The compressor is more efficient when the air is more dense, which occurs at lower temperatures. Since the compressor is powered by the turbine (via a shaft), having a more efficient compressor means that more kW can be generated by the turbine for the plant electric load. In fact, a rule-of-thumb states that for each 18°F rise in inlet air temperature, power output decreases by about 9%, which is the motivation behind most air inlet cooling systems. Therefore, it is important to consider weather conditions in the performance of a gas turbine. Note that relative humidity has a negligible effect on efficiency.
- **Turbine elevation**: Similarly to the inlet air condition, the higher the elevation of the turbine, the less dense is the air, which results in a less efficient compressor and reduces fuel efficiency. Generally, speaking, every 1000 ft increase in elevation decreases the power output by about 3.5%.

The heat-rate of an actual 8 MW natural gas turbine generator from the southwest (elevation of 5000 ft) is illustrated in Figure 4 and highlights the loss of fuel efficiency at part-load for different inlet air conditions. However, in addition to turbine fuel efficiency, economic efficiency of the gas turbine generator must be considered as it competes against the TOU rates from the local utility's tariff. For example, if TOU rates are very low, it may be uneconomical to utilize a natural gas turbine, regardless of how efficient it is. However, due

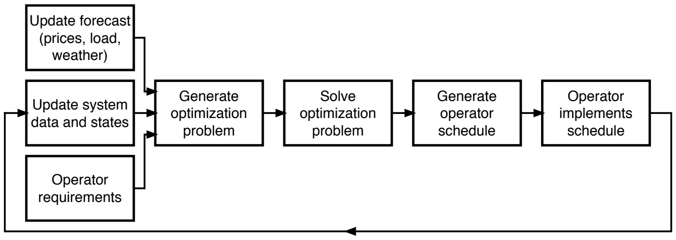

Fig. 5. Data flow within Balance.

to demand charges, even if TOU rates are low, it may still be valuable to operate the natural gas turbine. This is because the cost of generation is offset by possible demand-charge reduction and the value of heat-recovered steam. We will discuss gas turbine economic analysis further in Case Study 2.

1) Heat-Recovery Steam Generator: As briefly mentioned above, one can employ a heat-recovery process for the exhaust gas to heat up water and generate steam. This is often accomplished with a heat-recovery steam generator (HRSG), which is a heat exchanger that on one side has hot exhaust gas entering and on the other side has feedwater coming in. Through the exchange of heat (hence the name), the feedwater evaporates into usable steam while the exhaust gas cools down and is sent to the atmosphere. The HRSG performance is determined through its variable *effectiveness* rating, ϵ, which is the ratio of the actual heat-transfer to the maximum possible heat transfer and should not be confused with the efficiency. The relationship between gas turbine exhaust gas and generated steam is given by the following:

$$\begin{aligned}\epsilon[k]\dot{m}_{GT}^{exh}[k]c_p^{exh}\left(T_{GT}^{exh}[k]-T_{HSRG}^{steam}[k]\right) = \\ \dot{m}_{HRSG}^{steam}[k]\left(h_{HRSG}^{steam}[k]-h_{HSRG}^{feed}[k]\right)\end{aligned} \tag{7}$$

III. Online Energy Management with *Balance*

In the previous sections, we have described tariffs and system model considerations. Now, consider an actual EP consumer with time-varying energy price signals, a real-time data stream, and at least one year of historical plant data for asset model validation. This section will discuss how Balance improves economic efficiency and best manages energy costs. Figure 5 illustrates the data flow within Balance. The underlying engine of Balance is a multi-period optimization algorithm that minimizes energy costs, schedules equipments, and determines optimal set-points for plant assets and loads. However, before solving any optimization problem, an accurate multi-energy load and weather forecast is computed on the fly, which supplies a forecast of hourly heating (e.g. lbm/hr steam), cooling (e.g. tons of cooling), and electric loads and outside air and wet-bulb temperatures for the next 24 hours and is validated against historical data. Next, the real-time data stream is utilized to capture the current state of the system to consider startup costs, ramp-rates, and minimum up/down time constraints that impact scheduling of assets. Finally, based on the customer's utility tariff or, if applicable, forecasted real-time market prices, the future energy rates are sent to Balance and optimization can commence.

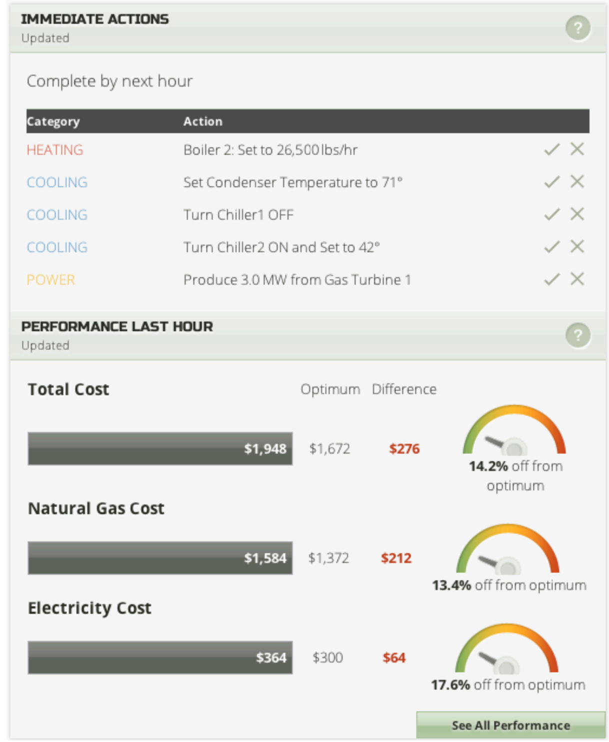

Fig. 6. Balance user interface with a list of recommendations.

The optimization formulation considers the following:

1) objective function: forecasted energy costs, including start-up, fixed, and shut-down costs.
2) the system configuration: energy flow paths through the plant to the forecasted loads (i.e. the multi-energy grid).
3) asset-specific models: defines asset efficiency and couples asset inputs and states (possibly across time-steps).
4) asset-specific constraints: minimum up/down times, ramp-rates, min/max operating levels, and bounds.
5) customer-specific constraints: specific couplings required by a customer, such as, maintenance events and other operational limitations not apparent from the physical system alone.

The resulting output is then processed to generate the recommended actions for the next hour, which are shown to the human operator in the user-interface and then implemented. The user-interface along with a compilation of recommendations is shown in Figure 6. After the recommendations have been issued, the current system state is measured against the optimized recommendations to evaluate how close the plant is operating to economic optimum.

Finally, the established real-time data stream enables operators to review recent economic-, efficiency-, and forecast-related performances. In addition, Balance tracks actual vs. optimal performance for each assets and allow operators and engineers to simulate the plant under different conditions, such as after installment of additional chiller capacity or a more efficient boiler (i.e. what-if analysis).

Remark III.1. The operators are severely time-constrained and often perform duties outside of the control room, which limited the effectiveness of employing Balance with human operators "in-the-loop." Therefore, efforts are currently underway to move Balance towards set-point automation and trade off the frequency of asset cycling against the economic benefits.

IV. Two Case-studies with *Balance*

The following section discusses two real case studies that highlights the capabilities of Balance. The first case-study represents a large university in the Midwestern US with 120 buildings served by two boiler/chiller EPs. The second case-study considers a large campus in Southwestern US of more than 100 buildings with natural gas and steam turbines, absorption and electric chillers, and boilers.

A. *Case study 1: Coordinating EPs and chiller lift*

In the midwest, two separate EPs (EP1 and EP2 with a total of 6 boilers, 4 chillers, and 10 cooling towers) coordinate energy management for an entire campus. Balance was able to reduce energy costs by \$200,000 over a year. The savings were \$100,000 and \$15,000 from more fuel efficient operation of natural gas boilers and electric chillers, respectively, while additional savings of \$90,000 were achieved through a reduction of natural gas demand charges. In addition, Balance provided the operators in the two plants with a portal through which they can communicate, which improved coordination further.

The boiler fuel savings were a direct function of the models developed and validated against real data. With proper models of boilers at part-load, Balance can dispatch them to operate the steam plant in a system-wide optimum. The natural gas demand-charge savings was a result of Balance's what-if capability and consideration of utility tariffs. The tariff stipulates that during the time-period $\mathbf{T}$ = December 1st - March 1st, daily (9AM-9AM) natural gas totals will be computed for each EP and the demand charge (DC) is then the sum of the individuals peaks over the 3-month period. That is,

$$DC = \sum_{i=1}^{2}\left(\max_{k\in\mathbf{T}}\left\{\sum_{\tau=1}^{24} Q_{fuel}^{\text{EP},i}[24(k-1)+\tau]\right\}\right). \quad (8)$$

The demand-charge DC set by (8) is then utilized to determine the cost of natural gas from March 1st to next year's March 1st, with the caveat of comparing demand charges to the last 12-month historical peaks and picking which ever is largest for each month. Balance's role in demand charge reduction was two-fold:

- Balance enhanced coordination between the two plants by providing operators with consistent forecasts and the ability to have the plants directly communicate with one another about impending maintenances and daily peak heating loads.
- Balance was able to ration fuel daily with a daily fuel limit constraint in place for EP1. The hard limit was necessary to ensure that EP2 had enough capacity remaining to handle possible contingencies issues without causing EP1 to set a large peak.

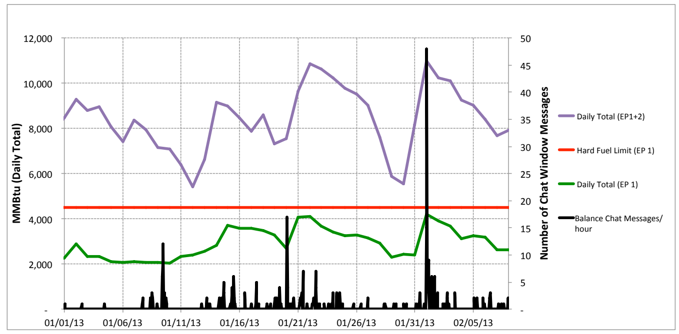

Fig. 7. Illustration that Balance was actively used by operators to mitigate natural gas demand charges.

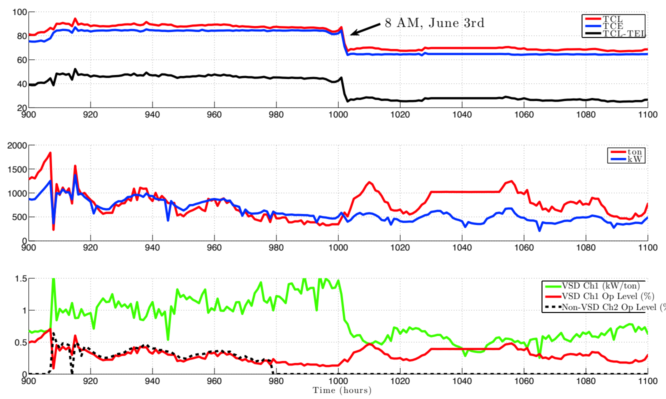

Fig. 8. Example of lift mismanagement and the impact on chiller efficiency.

This strategy reduced demand charges by 11.5% from the previous year. The role of Balance is well-illustrated in Figure 7 where it is apparent that communication via Balance peaked during peak heating load on February 1st. Note that Balance operates EP1 close to its maximum daily limit during the main peaks which highlights the effectiveness of the daily rationing method.

For chillers, it is also important to consider proper staging, however, even with proper staging, if set points are mismanaged the chilled water plant can be run inefficiently. To highlight the importance of lift management of chillers, the chiller shown Figure 2 was operated at part-load but under high lift conditions, which led to highly inefficient operations. As shown in Figure 8, at 8:00 AM on June 3rd, the set point T_{CE} was lowered which reduced lift ($:= T_{CL} - T_{EL}$) and improved efficiency by 50% as evidenced by the decrease in kW/ton in the lower subplot of Figure 8.

Lift management in Balance is achieved by regulating T_{CE} from cooling towers and $\dot{m}_{cond}$ from condenser water pumps. In (2), it is apparent that for a given $\dot{Q}_{cond}$ and T_{CE}, if $\dot{m}_{cond}$ is increased by speeding up pumps, then the condenser range $T_{CL} - T_{CE}$ decreases, which represents a decrease in T_{CL} and, consequently, reduces lift. However, speeding up condenser water pumps leads to increased power draw by the pumps, which implies that there is a trade-off between chiller lift reduction and condenser pump power. The relationship between lift, condenser power, and T_{CE} is one of the most dynamic relationships for chilled water plant set-points when it comes to energy efficiency, as different load and weather conditions beget different energy trade-off. Another method for reducing lift involves increasing T_{EL}, but it is generally not feasible due to hard constraints on the chilled water supply temperature. Nonetheless, with the chiller, pump, and tower asset models and for given load and weather conditions, Balance can simply compute the set points that maximize economic efficiency, which for the chilled water plant in this case-study led to savings of $15,000 for the year.

B. Case study 2: Cogeneration

A large southwestern EP has 1 steam absorption and 5 electric chillers, 2 steam boilers, 1 natural gas turbine with attached HRSG (i.e. cogeneration) and 1 steam turbine to supply a campus with cooling, heating, and electric needs. Unlike in Case study 1, the steam absorption chiller in this plant allows conversion of steam energy to provide cooling while the steam turbine can convert steam energy to electricity. This cross-coupling of energy flows and loads enables fuel flexibility, which can be utilized as a demand response mechanism or to reduce demand charges or avoid high TOU rates. For example, the absorption chiller can reduce the electric peak by up to 500 kW by displacing electric chillers while the steam turbine can bring down the utility's peak by another 500 kW. However, the focus of this case-study is on the interaction of the natural gas turbine and attached HRSG with the electric tariffs. Specifically, this customer can utilize natural gas to produce up to 7.2 MW (in a cold winter) from the gas turbine while the HRSG can output as much as 30,000 lbm/hr of steam. Their tariff structure is as follows:

1) Energy rates are split into on- and off-peak: 0.0821025 and 0.0327765 $/kWh, respectively. On-peak is for Monday-Friday 8AM to 8PM, while all other hours of the week are subject to off-peak rates.
2) Demand-charges are incurred for demand above 8 MW at the rate of $9.56/kW. For example, if the campus consumes 20 MW during on-peak in July, they are charged 12,000 kW $\times$ $9.56/kW $= \$114,720$ in demand-charges alone, which makes up about 20-30% of their monthly electric bill. Note that if they set a campus peak of 50 MW during a Saturday or Monday night (i.e. off-peak) this peak would *not* incur demand-charges.

From the tariff information alone it is apparent the important role of the natural gas turbine in reducing demand charges. In fact, prior to installation of Balance, this customer spent about $240,000 extra last year due to poorly timed maintenance of the gas turbine causing electric utility procurement to spike and set peaks for six months out of the year. One particular instance was rather unfortunate with maintenance taking place over the weekend but the gas turbine was not ready to be turned on until Monday 8:30 AM, which set the peak for that month. From these experiences, it is apparent the need for considering economics when it comes to maintenance. In fact, with real-time and forecasted pricing signals, Balance can recommend maintenance and maintain economic efficiency.

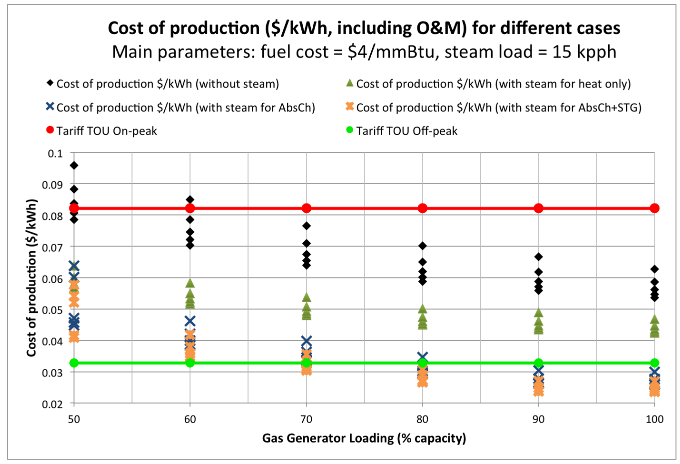

Fig. 9. Cogen economic analysis of cost of production with natural gas price at $4.0/mmBtu. Multiple identical symbols represent different inlet air conditions (20, 40, 60, 80, 100)°F.

The economics of natural gas turbine depend not only on demand charges, but also on the cost of production, which must consider the value of steam. The boilers in this southwestern EP can produce about 0.90 kpph of steam per mmBtu/hr of natural gas, which means that if the HRSG produces 30 kpph of steam from the gas turbine's exhaust gas *and* that steam is completely utilized, the value of the steam is tied to the cost of natural gas (from heating), the cost of cooling (from absorption chiller), and the cost of electricity (from steam-turbine). The cost of production, CoP [$/kWh], is then defined as:

$$\begin{aligned} CoP = (&\text{cost of gas turbine fuel and O\&M} \\ &- \text{value of steam for heating} \\ &- \text{value of steam for cooling} \\ &- \text{value of steam for electricity})\,/\text{kWh produced.} \end{aligned} \tag{9}$$

Using system models and energy prices, Balance finds the conditions when it is most valuable to operate the natural gas turbine for maximum economic efficiency. As shown, in Figures 9-10, the CoP is very sensitive to natural gas prices (which should be expected). As illustrated in Figure 9 with orange and blue crosses, when the natural gas costs are below $4.0/mmBtu and all of the steam from the HRSG can be utilized (recall, it produces 30 kpph at 100% GTG capacity), Balance informs the operators to continuously run the cogen unit at 100% capacity as the CoP is below even off-peak TOU rates with demand-charge reduction an additional economical bonus. However, if the steam cannot be utilized (i.e. absorption chiller and steam turbine are in maintenance and heating load is below 30 kpph), the cost of production increases to $0.045/kWh (as seen by the green triangles), which is much greater than off-peak TOU rate. Therefore, Balance will recommend that the cogen unit be turned off over the off-peak weekend (since daily cycling of cogen is often undesired). Note that even if no steam can be utilized (i.e. the black diamonds), the on-peak TOU rate is larger than CoP for all air inlet conditions, which means that the cogen should always run during on-peak.

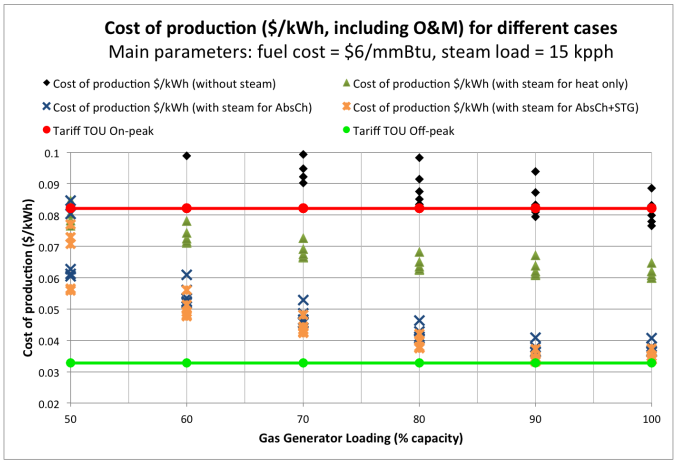

Fig. 10. Cogen economic analysis of cost of production with natural gas price at $6.0/mmBtu. Multiple identical symbols represent different inlet air conditions (20, 40, 60, 80, 100)°F.

However, as displayed in Figure 10, when the cost of natural gas increases to $6/mmBtu (as it was less than five years ago), the conclusions change. Namely, even with steam fully utilized (see blue and orange crosses), the CoP is still greater than off-peak TOU rate, which causes Balance to recommend bringing the cogen unit off-line during off-peak. In fact, when the steam cannot be utilized (black diamonds), it may seem to be uneconomical from the perspective of CoP to operate the cogen at all during on-peak, however, demand charges still incentivize Balance to recommend cogen operations during on-peak periods.

It is clear from the above analysis that changes to natural gas prices, heating loads, or availability of absorption chiller and steam turbine can alter the strategy for economic efficiency of the gas turbine. This suggests the need for recommendations that adapt to changing conditions, which is exactly what Balance does.

V. Conclusion

Through two case-studies, this paper highlights the value of utilizing historical and real-time data and pricing signals along with high-fidelity system models to enable economic optimization and energy analytics for large multi-energy consumers, such as energy plants. The two studies were made possible by applying the state-of-the-art product *Balance*, which provided online mitigation of natural gas and electric demand charges, optimization of chilled water and steam plants, and investigation of what-if scenarios for cogeneration assets and tariffs. It is important to point out that by using Balance, an energy plant is able to replace brittle ad-hoc rules of operations with robust dynamic plant recommendations that ensure economic efficiency and maintain reliability. That is, the dynamic nature of energy plants and pricing signals should no longer be subject to static ad-hoc operating rules that, when outdated, can become costly errors in implementation.

While Balance already offers a multitude of operational enhancements and insights into customer plants, future work will center on the application of machine learning to further speed up deployment of Balance and fully automated verification of incoming data. We are also interested in adapting the asset models online based on real-time data collection to allow Balance to automatically detect when asset performance deviates from expected behavior. Finally, we are interested in replacing time-based scheduling of maintenance with online condition-based scheduling while considering economic opportunities and efficiency.

VI. Acknowledgements

The authors would like to thank Morgan Leahy, Josh Keeler, Geoff Hancock, Colin Davy, Kathryn Schumacher, and Logan Niehaus for fruitful and interesting discussions on energy analytics, data preparation, and implementation.

References

[1] A. Wood and B. Wollenberg, *Power Generation, Operation, and Control*. Wiley-Interscience, second ed., 1996.
[2] W. van Gool, A. Hoogendoorn, R. Kümmel, and H. Groscurth, "Energy analysis and optimization of energy systems," *Intersociety Energy Conversion Engineering Conference*, pp. 1923 – 1928 vol.4, 1989.
[3] M. Geidl and G. Andersson, "Optimal power flow of multiple energy carriers," *IEEE Transactions on Power Systems*, vol. 22, no. 1, pp. 145 – 155, 2007.
[4] M. Almassalkhi and I. A. Hiskens, "Optimization framework for the analysis of large-scale networks of energy hubs," *Power Systems Computation Conference*, Aug 2011.
[5] Y. Ma, A. Kelman, A. Daly, and F. Borrelli, "Predictive Control for Energy Efficient Buildings with Thermal Storage: Modeling, Stimulation, and Experiments," *Control Systems, IEEE*, vol. 32, no. 1, pp. 44–64, 2012.
[6] Schneider Electric, Inc., "PowerLogic ION EEM 4.0," 2013.
[7] ASHRAE, *ASHRAE Handbook: Heating, Ventilating, and Air-Conditioning Systems and Equipment: Inch-Pound Edition*. ASHRAE, 2012.
[8] F. C. Knopf, *Modeling, Analysis and Optimization of Process and Energy Systems*. John Wiley & Sons, Dec. 2011.
[9] N. Petchers, *Combined Heating, Cooling and Power Handbook: Technologies and Applications, Second Edition*. Fairmont Press, Incorporated, 2012.

Mads R. Almassalkhi (S'06, M'14) obtained his Ph.D. degree in 2013 in electrical engineering: systems from the University of Michigan, Ann Arbor, USA. He is the Lead Systems Engineer at Root3 Technologies Inc.

Bjöwn Simon has 12 years of experience as a software developer and is the Chief Architect at Root3 Technologies Inc. He holds an M.S. degree in Civil Engineering from the Ruhr-University Bochum, Germany.

Archie Gupta is the CEO of Root3 Technologies Inc. and completed his MBA from University of Chicago - Booth School of Business in 2011 and holds an M.S. in electrical engineering from Clemson University, South Carolina, USA.