Proceedings of the Twenty-Eighth AAAI Conference on Innovative Applications (IAAI-16)

Optimizing Energy Costs in a Zinc and Lead Mine

Alan Kinsella
Boliden Tara Mines Ltd.
Navan, Co. Meath, Ireland

Alan F. Smeaton
Insight Centre for Data Analytics
Dublin City University, Ireland

Barry Hurley and **Barry O'Sullivan** and **Helmut Simonis**
Insight Centre for Data Analytics
University College Cork, Ireland

Abstract

Boliden Tara Mines Ltd. consumed 184.7 GWh of electricity in 2014, equating to over 1% of the national demand of Ireland or approximately 35,000 homes. Ireland's industrial electricity prices, at an average of 13 c/KWh in 2014, are amongst the most expensive in Europe. Cost effective electricity procurement is ever more pressing for businesses to remain competitive. In parallel, the proliferation of intelligent devices has led to the industrial Internet of Things paradigm becoming mainstream. As more and more devices become equipped with network connectivity, smart metering is fast becoming a means of giving energy users access to a rich array of consumption data. These modern sensor networks have facilitated the development of applications to process, analyse, and react to continuous data streams in real-time. Subsequently, future procurement and consumption decisions can be informed by a highly detailed evaluation of energy usage. With these considerations in mind, this paper uses variable energy prices from Ireland's Single Electricity Market, along with smart meter sensor data, to simulate the scheduling of an industrial-sized underground pump station in Tara Mines. The objective is to reduce the overall energy costs whilst still functioning within the system's operational constraints. An evaluation using real-world electricity prices and detailed sensor data for 2014 demonstrates significant savings of up to 10.72% over the year compared to the existing control systems.

Introduction

Today, the majority of electricity consumers pay a flat rate, with some suppliers offering a number of variations such as time-of-use schemes where the rate will change depending on the time of day, week, or month. However, with such schemes the price will only fluctuate marginally as compared with changes in wholesale price.

Many industries require continuity with their business activities and can only re-schedule electricity usage for short periods of time. Thus, in some cases these schemes do not incentivise reductions in consumption as the savings are not offset by the losses in profit.

Electricity suppliers purchase electricity from the wholesale market and sell it to the consumer at fixed rates. This rate absorbs the risk of fluctuations in wholesale price and incorporates it into the flat price. The risk factor contributes to an overall higher flat rate calculated and charged at the suppliers' discretion. In addition to this, the flat rate typically includes the raw cost of electricity, administration, supplier fuel hedges, and supplier profit margins.

(Kirschen 2003) highlights that when consumers pay variable electricity prices the levels of awareness increases to how time can influence electricity price and ultimately leads to developing initiatives that maximise the value of consumption. This paper proposes one such initiative by examining the following:

- a detail of the electricity usage at Tara Mines[1] specifically in the area of pumping;
- an analysis of the workings of Ireland's Electricity Markets and the present opportunity to achieve savings using real-time variable prices;
- the acquisition and fusion of sensor data such as flow meters, radar water level sensors, and smart energy meters to study the workings of *Pump Station #1* at Tara Mines;
- a formalisation of the pump station system constraints and energy-focused objective function; and
- a simulation of optimisation techniques to minimise energy costs using real-world sensor and market data, while comparing to a variety of baselines and existing control mechanisms.

Motivation

The Emergence of Smart Metering

One recent emerging opportunity of the *smart grid* is the development of schemes by energy suppliers to allow customers to directly respond to variable pricing. It has been identified that smart meters can enable variable pricing by tracking energy consumption at 15 minute intervals, giving consumers the possibility to understand and adjust usage to price curves in a timely manner. Policy makers in Europe

[1] http://www.boliden.com/Operations/Mines/Tara/

have focused heavily on the roll-out of smart meters in recent years, however, little attention has been given to ensure suppliers offer customers more advanced tariff schemes to coincide with these new levels of energy awareness. It is thought that if clear time-flexible pricing signals are communicated to customers then they can have the potential to save substantial amounts of money (Faruqui, Harris, and Hledik 2010).

Variable Electricity Price Opportunity

The Irish wholesale electricity prices vary at half hour intervals; in 2014, this ranged from a minimum of €4.69/MWh to a maximum of €955.38/MWh. Fluctuations in price can be considerable with a standard deviation of €37.3/MWh from mean prices of €56.9/MWh, even exceeding €600/MWh on many occasions. These variations offer the potential for large cost savings if energy intensive activities can be rescheduled to avoid the peaks and, where possible, increase workloads in the more cost effective valley periods.

Energy Cost-Aware Scheduling

(Ifrim, O'Sullivan, and Simonis 2012) developed models to accurately predict real-time energy price and used this knowledge to achieve significant energy cost savings in a scheduling context. They were able to make more accurate price predictions than those of the market operator's own forecast, SEMO (Single Electricity Market Operator).[2] (Grimes et al. 2014) studied the potential of using a methodology to construct energy cost-aware consumption schedules in an Irish industrial setting by developing price forecasts and scheduling usage from these forecasts. The schedules produced cost savings in contrast to cost-unaware schedules based on a flat tariff. As such, large industrial energy consumers can benefit from variable pricing structures even under usage constraint complexities and time-sensitive production processes.

(Aikema, Kiddle, and Simmonds 2011) and (Qureshi et al. 2009) also used scheduling techniques to make significant savings in the context of data-centre optimisation.

Mine Dewatering

The predominant energy consuming tasks in underground mining operations are crushing, conveying, ventilation, and pumping. Pumping is one of the most important tasks, using 10% of all energy. It has a significant cost of approximately €2.5M annually at the mine considered here. Water can enter the mine in three different ways: natural ground water, service water for the mining operation, and as a transport medium to push backfill cement down pipes to fill previously mined areas. Water flows in the mine by either free flowing in ground channels or by being pumped through pipelines. All the free flowing water from upper mine levels is collected at central sumps where a flocculant is added to accelerate deposition of fine material. To provide operational continuity and to avoid flooding of critical infrastructure, certain areas of the mine need to be kept water free and therefore water needs to be removed by pumping it to the surface. In the south-western section of the mine, water is diverted to horizontal sumps (mined out pits used to store water). The fine material settles within the sumps and a filtering system clarifies the overflow water prior to enter *Pump Station #5*. After a period of time, collected fine material at the bottom of the sumps is allowed to dry out, the barriers are removed and the fines are dug out. With an average flow of $260m^3/h$, *Pump Station #5*, transfers the cleaned water through a 300mm pipe over a distance of more than 2km and a vertical ascent of 500m to join the cleaned water from the upper levels of the mine, this water is collected in sumps before being pumped to surface by the main pump station, called *Pump Station #1*. *Pump Station #1*, to be examined in this paper, consumed 79,981 MWh of electricity in 2014 and is the largest pump station in the mine. At standard industrial retail prices of 13c/KWh, its total operating cost was roughly €1,039,757. One of the main benefits of energy cost-aware scheduling for mine water systems is the ability to retain water in the sumps, deferring pumping until cheaper prices are on the horizon. This fact, together with the possibility to run a combination of pumps concurrently to vary output flow speeds, gives a large degree of flexibility with which to optimise energy cost.

Pump Station #1: Current Practices, Challenges, and Opportunities

The operation of *Pump Station #1* is a complex task involving the running of four pumps to regulate water levels in sump storage to meet the dewatering requirement from all of the mine. The station is equipped with two twin capacity centrifugal pumps with an output flow capable of up to $950m^3/h$ on average. The station has residual capacity greater than its demand requirement in reserve as a precaution for peak water load situations. Figure 1 depicts the four pumps, sumps, pipe connections and civil structure that make up the pump station. Water is pumped from *Pump Station #5* to *Sump 1* located in *Pump Station #1*. An overflow from a *Lamella* which treats sediment water runs into *Sump 3* and then *Sump 2* and then subsequently to *Sump 1*.

Today, the PLC (Programmable Logic Controller) controlled station is programmed to use a *"duty-standby"* configuration whereby if water levels in *Sump 1* exceed a start set-point a pump will start to operate, corresponding to this if the water level in the sump drops to the stop level the pump will stop. Thus, the management of high water levels are dealt with by increasing the number of concurrently running pumps to adjust the output flow rate as appropriate until water levels in the sump are reduced. The sump level set-points have been refined over a number of years by experienced engineers and as a result it operates with a great deal of reliability. However, variable energy prices has not been considered in its development.

Water Flow and Sump Storage Capacities

Treated water flows in overflow form with a cascade effect as part of an incremental filtration process from *Sump 3* to *Sump 2* to *Sump 1* at a constant rate of $275m^3/h$ (See Figure 2). The flow to *Sump 1* remains at this constant rate once

[2] http://www.sem-o.com/

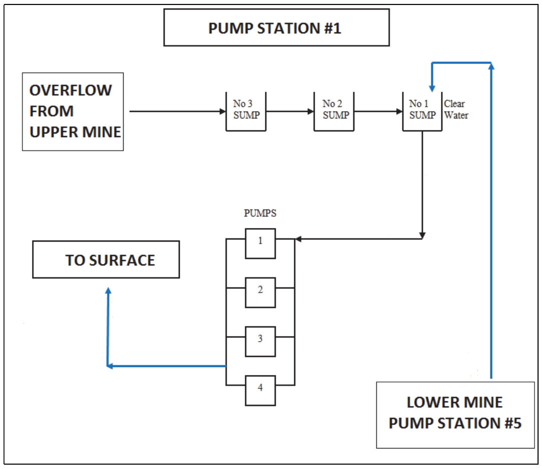

Figure 1: Schematic Pump Station #1

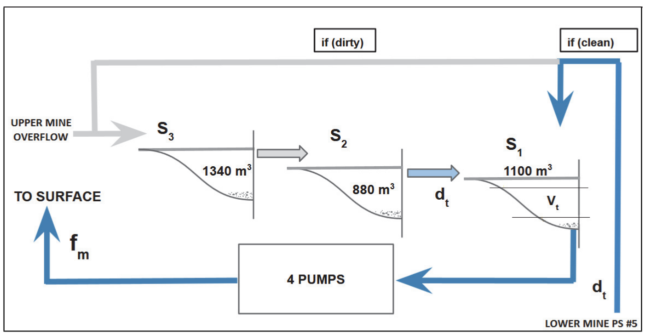

Figure 2: Water Flows

Sump 3 and *Sump 2* are at adequately high levels. Based on analysis of past data, these levels remain consistent between 60% and 80% to provide this flow. As sediments can cause pump life to shorten, the sump overflows are created by design to ensure the water that reaches *Sump 1* is as sediment-free as possible. This is done by allowing rock fine particles to settle at the bottom of the sumps and removed once there is adequate build up. *Sump 1* has a storage capacity of $1100m^3$. The dewatering demand flowing into *Sump 1* has two sources: a constant overflow from the upper mine levels, and a fluctuating flow through a pipeline from *Pump Station #5* in the lower mine. It is not recommended for *Sump 1* levels to go under 10% or over 90%. This buffer is required as the capacity will slowly adjust as sediments settle to the bottom of the sump.

Single Electricity Market Operator (SEMO)

Understanding the operations of SEMO, Ireland's electricity market, is essential to scheduling consumption with dynamic market prices. SEMO operates a mandatory pool market, where all electricity on the island of Ireland is bought and sold. Sale and purchase from the pool is completed with a code of practice by which generators sell to the pool at the marginal cost of producing a unit of electricity €/MWh known as the System Marginal Price ('SMP').

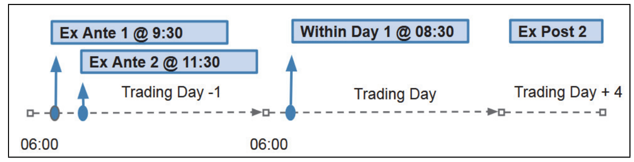

Figure 3: SEMO Market Publication Times

The market operates by stacking generator bids in order of price, cutting the stack at the point which makes up the demand to be met on the day of consumption. This process is completed a day in advance known as the Ex-Ante market (SEMO 2014). On the day Eirgrid, the system operator, communicates in real-time to all the generators to start up or turn off to balance demand. Of course, it is not possible to know all price influencing factors in advance and so to promote market efficiency the market executes several forecast price runs before electricity is consumed on the day, e.g., the Ex-Ante 1 ('EA1'), the day ahead forecast Ex-Ante 2 ('EA2'), the day ahead forecast Within Day 1 ('WD1'). This price data is released to market participants at different intervals throughout the trading day through a series of publications (See Figure 3). A settlement price, the Ex-Post 2 (EP2), is reached four days after consumption. Factoring in the actual happenings of the day, this is the actual price paid by market stakeholders.

Since it's impossible to determine exactly what will happen beforehand, the actual cost will differ to the forecast. Table 1 shows a €0.74 difference in mean and €4.67 difference in standard deviation between WD1 and EP2 prices.

Table 1: Forecast & Settlement Prices 2014 - SEMO

Price	Mean	Std. Dev.	Min	Max	Range
EA2	55.01	28.88	18.46	471.58	453.12
WD1	56.16	32.63	12.6	488.36	475.76
EP2	56.9	37.3	4.69	955.38	950.69

Electricity is unique compared to most other commodities as it cannot be stored economically with current technology. Consequently this gives rise to large fluctuations in wholesale price due to factors such as transmission congestion, generator maintenance, technical constraints, fuel prices, generator efficiency, generator start-up costs and demand fluctuations, etc. (Enright 2013). Ireland has witnessed a significant uptake of renewable wind generation in recent years which, given the intermittent nature of this energy form, further adds to the price volatility. Figure 4 illustrates a typical trend from a sample day (31/10/2014). The potential to use market forecasts is evident when one considers the prospect of moving energy usage within the day. The differences in the price profiles between the EA2, WD1, and EP2 are shown over a 24 hour period. Although prices differ quantitatively overall, a common feature of the EA2 and WD1 forecast fluctuations is their ability to identify movements in EP2 settlement prices.

In fact, if we examine the covariance matrix associated with each price on the diagonal axes in Figure 5, the red

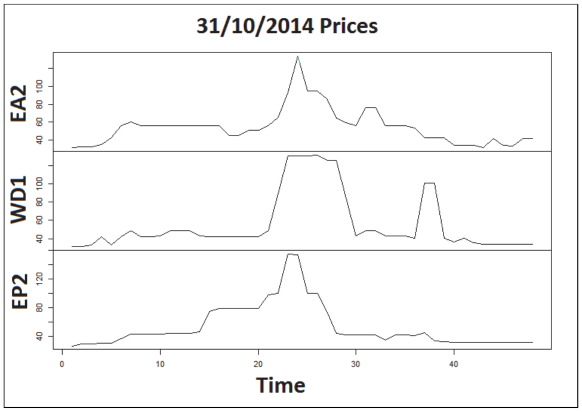

Figure 4: Sample EA2, WD1 & EP2 price runs

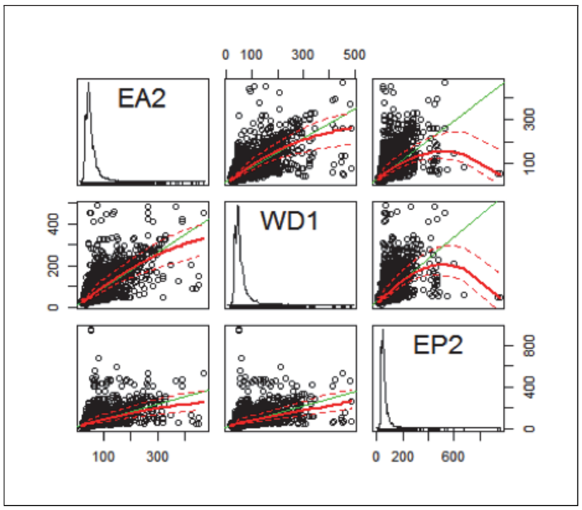

Figure 5: SEMO price runs covariance plot matrix

lines indicate linear patterns between all EA2, WD1, and EP2 prices over all of 2014.

System Architecture

An extensive underground fibre optic network connects a PLC network to all fixed electrical equipment underground. Monitoring of this equipment is done using a Supervisory Control and Data Acquisition (SCADA) system. SCADA is a computer system for gathering and analysing real-time data used to monitor and control a large variety of industrial equipment. *Pump Station #1* is equipped with piezometer water level sensors, flow meters and smart energy meters providing complete coverage of all sump levels, electricity usage and flow rate data. The KWh readings on the four pumps are acquired from smart meters recorded at 15 minute intervals sourced from an Energy Management System.

Figure 6 depicts the high-level architecture of the scheduling system. Data is combined from the smart meters connected to the energy management system with water level and flow sensors in the SCADA together with market price data. This is fed into the optimisation model to produce operational schedule for the pump station.

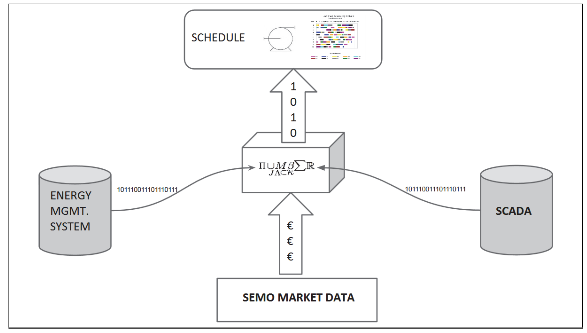

Figure 6: Operational Architecture

Operational Plan. An optimisation model is run every 15 minutes which considers a two hour planning horizon. The outcome decides the pumping configuration for the next time period while considering the current sump levels, a forecast of the mine dewatering demand, and the forecasted price vector. The model is re-run at the next interval with the latest sump levels, reflecting the actual demand which occurred.

Pump Station #1

A series of interviews with engineers elicited a number of practical concerns and constraints on the operation of the pump station.

The pump station is composed of four pumps, each varying in pumping capacity and energy consumption. In theory, 2^4 configurations of the four pumps are possible, however due to operational constraints with control systems we considered only five active configurations chosen by the engineers. These configurations, along with their respective pumping capacity and power consumption are listed in Table 2.

Table 2: Mode Configurations

P_1	P_2	P_3	P_4	Flow	Power
OFF	OFF	ON	OFF	$480m^3/h$	692KW
ON	OFF	ON	OFF	$628m^3/h$	1000KW
ON	ON	ON	OFF	$640m^3/h$	1204KW
ON	ON	OFF	ON	$776m^3/h$	1288KW
ON	ON	ON	ON	$910m^3/h$	1716KW

Note that the net flow when multiple pumps are running is not simply a linear combination of their individual flows due to the physical behaviour of varying pressures and pipe resistance. Additionally, in practice, the pumping flow and power consumption vary from design specification and degrades over time, therefore the values presented have been recorded empirically from localised sensor data and need to be updated periodically.

In order to prolong the overall lifespan of the pumps they should not be cycled on/off rapidly, and so, the motor duty is modelled at a minimum active time of 15-minute intervals, mainly to avoid overheating.

Forecasting Mine Dewatering Demand

Pump Station #1's fluctuating pumping demand will depend on the levels of ground water in the mine and whether or not pockets of underground water aquifers are hit during the mining processes. Removal of this water in both situations is done by discretionary pumping into the ground channels or directly to the sumps using portable submersible water pumps. Hence, the timing and occurrence of these events carry a very high degree of uncertainty, no periodicity or seasonality.

Data is acquired from flow meters at intermediate pump stations in the lower south west sections of the mine, this water will eventually flow to *Pump Station #5* and subsequently pumped to *Pump Station #1*. Thus, this data is used as an indicator of the future pumping demand after a certain time delay. The forecast demand is produced by monitoring the intermediate pump station flow output and adding ground water estimations issued from geological reports of the area to the flow.

A constant of $130m^3/h$ obtained from geological estimates of ground water in the area is added to the intermediate sump flow time series to make up the full flow that enters *Pump Station #5*. A cross correlation function is then computed to identify the lags of intermediary pump stations x_t and *Pump Station #1* demand y_t. A multiple regression is used where y_t is the linear function of lags of the intermediate sump variables, the model works well with all coefficients being statistically significant. The regression model resulted in a mean absolute percentage error of 7.36% and root mean squared error of 17.74. However, given that it is not possible to assume the forecast demand will be perfectly accurate, the sump levels are updated before each iteration using the actual levels recorded from sensor data.

Optimisation Model

The following notation will be used in the optimisation model: m is a configuration from set of pumping configurations M, f_m and p_m, respectively, correspond to the pumping flow and power consumption for configuration m. t is the time index from a time horizon T, d_t is the forecasted demand into the sump at time t, and c_t is the forecasted electricity price at time t. Additionally, we have the inital sump level v_0, and inventory cost q to penalise the final sump level.

The variables consist of Boolean variables x_{mt} to say if configuration m is active at time t, and integer variables v_t corresponding to the volume of the sump at time t. For flooding concerns, the level of the sump s_1 should not go above 80% or below 20% of its capacity as it could be pumping some residual sediment. Only the working capacity of the 20-80% range will be modelled by the v_t variables.

The model consists of the following constraints. At most one configuration can be active at any one time point:

$$\forall t \in T: \quad \sum_{m \in M} x_{mt} \leq 1.$$

The sump volume is channelled between time points as the sum of the level at the previous time point, plus the new incoming demand, minus any pumping which was done:

$$\forall t \in T: \quad v_{t+1} = v_t + d_t - \sum_{m \in M} x_{mt} \cdot f_m.$$

The objective is to minimise the energy cost of all active pumping configurations across the time horizon but also to include a penalty for the final sump volume. This is to avoid the optimisation leaving the sump full at the final time point, but an alternative method could be to extend the time horizon considered:

$$\text{minimise} \sum_{m \in M} \sum_{t \in T} x_{mt} \cdot p_m \cdot c_t + v_{|T|} \cdot q.$$

Evaluation

All simulations were run over a period 363 days in 2014 from the 2$^{\text{nd}}$ of January 2014 to the 30$^{\text{th}}$ of December 2014, resulting in 4356 two hour sliding time horizon optimisations. The model uses two hours as the highly volatile dewatering demand can only be meaningfully forecasted over this horizon. In each case the optimisation problem was solved using Numberjack[3] and a Mixed Integer Programming solver, SCIP 3.1.0.[4]

To solve the optimisation model for a single 2 hour horizon required a maximum of 12.6 seconds, minimum of 0.03 seconds, average of 0.7 seconds and a standard deviation of 1.1 seconds on Mac OS Yosemite with a 2.26 GHz Intel Core 2 Duo processor and 8 GB of RAM. This short solving time enables the system to react very quickly to the situation, which is an important consideration in such a dynamic scheduling scenario.

Evaluation Baselines

Baseline A: Oracle. Here a schedule is realised with perfect knowledge of price and demand using EP2 and actual demand sensor data. This is the optimal schedule and provides a lower bound on the potential savings. Such an omniscient oracle schedule has a cost of €400,983, which was used as a lowerbound baseline for all other schedule simulations.

Baseline B: Existing control. The existing configuration was not created to consider variable price, however, to measure the cost advantages of optimising, running existing operation against market prices should realistically be more cost effective than existing control using EP2 prices. Indeed, we see that existing control costs €457,930 giving a 14.27% difference from the oracle solution.

Baseline C & D: EA2 and WD price forecasts. Due to the real-time nature of the schedule, adequate time is available to use both EA2 and WD1 forecasts and so a simulation is completed using both price vectors to determine which has the most cost effective outcome. EA2 prices cost €415,354,

[3] http://numberjack.ucc.ie/

[4] http://scip.zib.de/

Table 3: Results of Cost-Aware Schedules

Baseline	⟨Schedule, Demand, Tariff⟩	m^3	KWh	Cost	% Diff.	€/m^3
A. Perfect Knowledge	⟨EP2,Actual,EP2⟩	4,679,160	7,955,041	€400,983	0%	0.0856955
B. Existing Control	⟨Set-Points,Actual,EP2⟩	4,676,247	8,006,250	€457,930	14.27%	0.0979268
C. EA2 Forecast	⟨EA2,Actual,EP2⟩	4,679,541	7,945,439	€415,354	3.56%	0.0887596
D. WD1 Forecast	⟨WD1,Actual,EP2⟩	4,679,182	7,974,137	€416,284	3.82%	0.0889651
E. WD1& Dewater Forecast	⟨WD1,Forecast,EP2⟩	4,678,491	7,967,075	€415,156	3.55%	0.0887371

a 3.56% difference from the oracle solution, WD1 prices at €416,284 show a 3.82% difference and 0.74% to EA2 forecasts. This is an interesting result as although WD1 contains more market data than EA2, it does not exhibit any notable advantage over the EA2 schedule. Both schedules are close to the oracle schedule (c.3% difference).

Baseline E - Simulating with WD1 & Dewatering Demand Forecasts. Simulating a real-life scenario using the dewatering forecast and WD1 price forecasts, a cost of €415,156 is observed. Results show a 10.72% cost reduction from *Baseline E* to the existing control of *Baseline B*, representing a significant advantage. Marginal differences between scheduling EA2 and WD1 forecasts is noted with EA2 exhibiting 0.26% more accuracy.

Conclusions

We studied an intelligent system for real-time optimisation of an underground mine pumping operation. The system uses a number of AI techniques to predict a highly volatile pumping demand and schedule pumping operations, while minimizing the overall electricity cost. Simulation results show significant saving opportunities using a real-time energy price, displaying a saving of 10.72% over the existing control, equating to roughly €40,000 on the wholesale cost annually. Such annual savings demonstrate the effectiveness of using AI to make real-time consumption decisions.

The deployment of such a system is dependent on the procurement of the market's real-time energy price. At the moment, the considered operation is billed on a two-rate time of day tariff. Switching to a real-time price is not a decision to be made without due consideration to the overall energy consumption, not just a single pumping operation. This paper serves as a proof of concept for applying intelligent optimisation to a large, dynamic real-world system.

In future, we plan to apply similar techniques to additional components in the mine operations such as the autogenous mill which consumes the majority of the overall electricity. We hope such an extension will enable the adoption of a real-time energy price to be realised in the near future.

Disclaimer

Any views or opinions expressed in this paper are those of the author(s) and not of Boliden Tara Mines Ltd. Tariff rates used are industry standard for 2014 (Statista 2014).

Acknowledgements

This paper is based on research partially conducted with the financial support of Boliden Tara Mines Ltd. and Science Foundation Ireland under grant number SFI/12/RC/2289.

References

Aikema, D.; Kiddle, C.; and Simmonds, R. 2011. Energy-cost-aware scheduling of HPC workloads. In *IEEE WOWMOM 2011*, 1–7.

Enright, M. M. 2013. Pricing Irish power futures. a comparative analysis of forward curves in European energy markets. In *International Conference on the European Energy Market (EEM)*. IEEE.

Faruqui, A.; Harris, D.; and Hledik, R. 2010. Unlocking the €53 billion savings from smart meters in the EU: How increasing the adoption of dynamic tariffs could make or break the EU's smart grid investment. *Energy Policy* 38(10):6222–6231.

Grimes, D.; Ifrim, G.; O'Sullivan, B.; and Simonis, H. 2014. Analyzing the Impact of Electricity Price Forecasting on Energy Cost-Aware Scheduling. *Sustainable Computing: Informatics and Systems* 4(4):276–291.

Ifrim, G.; O'Sullivan, B.; and Simonis, H. 2012. Properties of energy-price forecasts for scheduling. In *Principles and Practice of Constraint Programming, CP 2012*, 957–972.

Kirschen, D. S. 2003. Demand-side view of electricity markets. *IEEE Transactions on* 18(2):520–527.

Qureshi, A.; Weber, R.; Balakrishnan, H.; Guttag, J. V.; and Maggs, B. M. 2009. Cutting the electric bill for internet-scale systems. In *ACM SIGCOMM 2009*, 123–134.

SEMO. 2014. The single electricity market: Market update (April-June SEM-14-090 October).

Statista. 2014. Industrial sector electricity prices: Europe 2014. http://www.statista.com/statistics/263262/industrial-sector-electricity-prices-in-selected-european-countries/.