Energy study of a manufacturing plant

Chloé Desdouits
GIPSA-lab Grenoble Images Parole Signal Automatique
UMR 5216 CNRS – Grenoble INP
Université Joseph Fourier, Université Stendhal
11 rue des Mathématiques
Grenoble Campus – BP46
38402 Saint Martin D'Hères Cedex
France
and Université Montpellier 2
and Schneider Electric
chloe.desdouits@gipsa-lab.grenoble-inp.fr

Jean-Louis Bergerand
Schneider Electric
38TEC – 37 Quai Paul Louis Merlin
38050 Grenoble
France

Pierre-Alexis Berseneff
MINES ParisTech
60, boulevard Saint-Michel
75272 Paris Cedex 06
France

Claude Le Pape
Schneider Electric
38TEC – 37 Quai Paul Louis Merlin
38050 Grenoble
France

Dimitri Yanculovici
Orange Applications For Business
195 rue Lavoisier
38330 Montbonnot Saint Martin
France

Keywords

abatement, data collection, decision-making process, optimisation

Abstract

Nowadays, manufacturing plants usually compute production plans using a "no stock" strategy. This is efficient regarding storage and, to some extent, delivery tardiness costs, but not at all regarding energetic costs. Moreover, with the trend of energy consumption regulation laws and the emergence of the demand-response electricity market, plant managers have to imagine new ways to regulate their energy consumption. For that purpose, some sensors and a data acquisition system must be put in place.

Measuring energy data allows 1) the introduction of new Key Performance Indicators (KPIs) related to energy consumption while evaluating the performance of a plant, 2) to find out machines that consume more than the others and change production plans to respect limits or to minimize costs. In addition, production data can also be obtained from existing systems (like SCADA). Having energy and production data together bring the possibility to compute energy models of production activities and use them to build energy-aware production plans automatically. This way, the "no stock" or "just in time" strategy can be challenged by an energy-aware strategy.

During the past two years, we built such an energy data acquisition system in a Schneider Electric factory. It is non-intrusive, built from self-powered wireless sensors, and hence very easy to install. The combination of production data and energy data gathered in the same dashboard proved efficient and useful to understand energy consumption profiles. Indeed, we found out the baseline energy consumption of the plant and succeeded to underline days when consumption is abnormal. It opens a promising avenue to enrich the software suites for production management systems.

We also began to build automatically energy-aware production plans, based on energy models of production activities, to assist a plant manager. On benchmark data, energy savings of up to 15 % are sometimes achievable (depending on the plant's operating conditions and electricity tariffs) with no impact on tardiness costs.

Introduction

The work we are presenting in this paper has been carried out in the context of the Arrowhead European funded project. The study has been conducted based on a real life production line of a plant that manufactures electrical cabinets. The process is quite straightforward, the line in made of two parts, one dealing with the doors of the cabinets, the other one with the body. For both parts the decomposition is similar: a first step where the iron sheet is shaped by stamping and folding operations, a second step where the product parts are welded. Then there is a common step for painting the raw steel parts. The ending operation is a manual assembly of the door with its body to finish up the cabinet (see Figure 1).

The current scheduling strategy is led by a lowest stock possible and the activity is planned to produce the production order portfolio in due time, with the right quality level. No real attention is paid to the energy cost for the production line. Energy is considered as a utility that you consume to achieve quality production in due time. There is no reconciliation of the production log with the energy consumption over the corresponding period. There is very rarely any tracking of the energy cost per

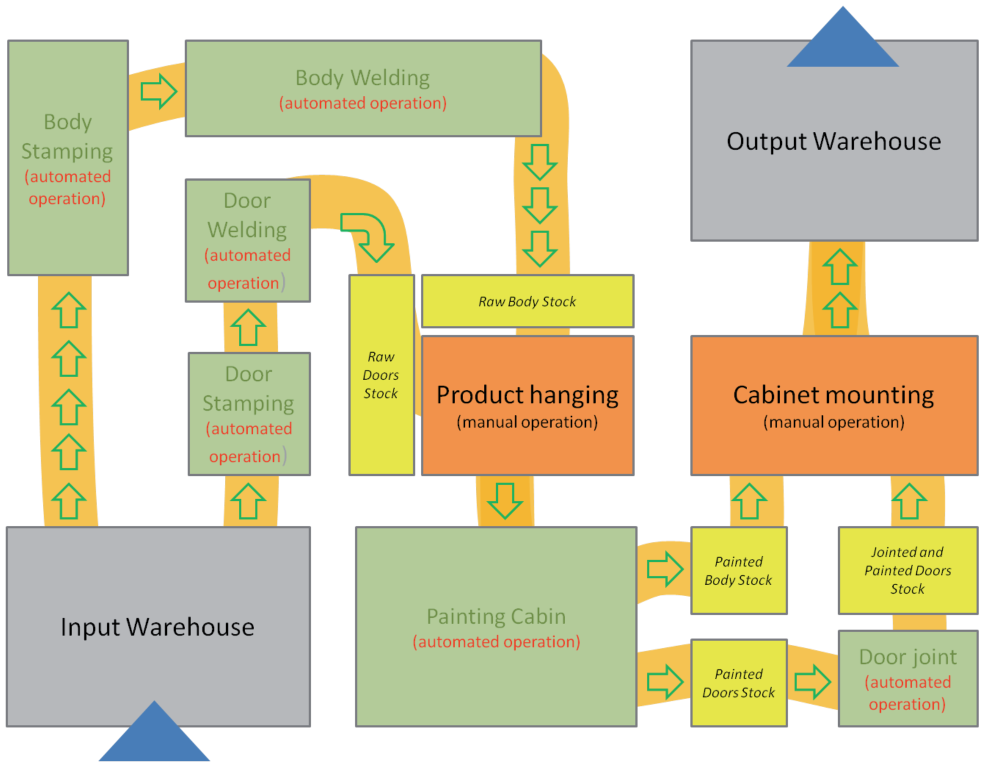

Figure 1. Manufacturing line description.

operation or per product reference. This reconciliation is not so easy to make, as the information, when existing, is stored in different tools, and managed and used by different actors of the plant: the production manager seldom is the energy manager in a plant of a reasonable size.

With the new energy challenge, the usage of energy has to come under scrutiny, on the one hand because of European regulation pressure on environment, on the other hand because of the increasing cost of energy. The first step and first need is to measure and understand energy consumption in a plant. The second step is linked with the prediction of an increasing variability or even volatility of the energy prices, due to the massive introduction of intermittent sources of energy production (mainly wind turbines and photovoltaic panels). This step consists of introducing some flexibility in the production plan to cope with the energy price variation.

Here comes the challenge of the "zero stock policy" inherited from kanban (Sugimori, et al. 1977), one of the current production management strong principles (if not a dogma). In a factory plant one of the only ways to get some flexibility without compromising the delivery date is to introduce some storage and to change the inventory policy. By introducing the variable price of energy, the time intervals when energy is spent are not equivalent. Some time intervals are more attractive than others to consume energy. These constraints can be taken into account in the scheduling of a production plan.

Some work on energy modelling of a manufacturing plant can be found in the literature. For example, (Le, et al. 2013) and (Dietmair and Verl 2009) address energy modelling of machines. (Dietmair and Verl 2009) have built a modelling framework for machine-tool energy consumption, based on state automation theory. (Le, et al. 2013) use a similar model and give an approach to reduce the number of required sensors in process tracking by identifying machines operational states, using a neural network. Following a different path, (Fysikopoulos, et al. 2014) propose a generalized approach to manufacturing energy efficiency by giving a definition of manufacturing energy efficiency at four different levels: Process level, Machine level, Production line level and Factory level. Interactions between these levels are studied and an example of scheduling is given at the production line level.

In this paper, work conducted in a Schneider Electric plant, with about 1M€ annual electricity bill, is introduced. Our goal is to understand the energy consumption of the manufacturing process and to manage it, in order to optimize the energy bill, as well as other Key Performance Indicators of the plant manager. It applies to any type of energy whose tariff varies with time of use. In the Schneider Electric plant under consideration, it concerns only electricity. Data acquisition from this plant is presented in the second section "Data acquisition system and visualization".

In practice, we have built a regression model to understand the electricity consumption of the plant with respect to the production portfolio. This is presented in the third section "Energy modelling of activities". This has proven to be a challenge. Indeed, this requires that (i) we get precise enough energy consumption data, (ii) we have historical production data recorded as precisely and accurately as possible and (iii) we are able to reconcile these two information sources into a synchronized table.

The regression model has been used to feed a scheduling optimisation tool that produces production plans. The fourth section "Energy aware scheduling" describes how we use the tool to evaluate the potential savings on the energy bill.

Data acquisition system and visualization

The energy data acquisition system (hardware and software components involved in the data collection process: sensors, data concentrator, transmission to a database) and the associated dashboards are first discussed. This acquisition system is described in Figure 2. Then, examples of production data are given.

ENERGY DATA

The focus was made on a selection of consuming machines in the production lines. Four machines have been selected; two per production lines: the door stamping, the door welding, the body stamping and the body welding. Those machines have been selected because their electricity consumption is significant enough and the assumption was made that the consumption profile would vary according to the type of items produced. The goal was to get an accurate enough consumption profile of a set of machines.

The system is built from a set of off-the-shelf components and bespoke components:

- The **energy sensors**: 2 prototypes of self-powered energy sensors have been used to monitor the machine electricity consumption: one type for mono-phased machines (E3 family Energy Estimator from Schneider Electric Innovation team) and one for three-phased machine (E4 family) ones. Both sensor types send the measured data through radio ZigBee Green Power (ZGP) protocol. The value sent is the cumulated value of energy measured since the beginning of their life; potential overflow is easily managed downstream. The information is not sent on a regular time basis, but on a regular quantum of energy.
- The **Sologate gateway**: this gateway has been used to take the advantage of both its support of the ZigBee protocol to receive sensors' data and its ability to expose the data through modbus protocol. The Sologate is capable of updating the energy values roughly every 100 ms.
- **The LINC gateway:** the role of this gateway is to be able to sample the energy value with a 1-second sampling rate. To achieve this, the LINC middleware provided by CEA (Louvel and Pacull 2014) has been deployed on a Raspberry Pi platform. It has been specialized in order to collect data and manage the Energy Operation database format. It queries the Sologate gateway through Modbus protocol and builds an aggregation of those data into a single file that contains time series made of the time stamped changes of energy value. Consequently, the timestamps are not regular, but the number of samples is smaller.
- The **Energy Operation** database and visualization tool: this is a proprietary Schneider-Electric tool offering a data storage capability and a web interface.

Visualisation of Energy dashboards (the Energy Operation tool)

Energy Operation is a SaaS (Software as a Service) tool developed by Schneider Electric. Its goal is either to store data coming from meters installed in a building (or coming from third party services), or to provide a customizable user interface accessible from the Internet in order to make the data understandable by the user.

It can store data coming from various ecosystems such as comfort sensors, energy meters (heating, electricity, water, gas) and building management systems. In our case we also want to push data coming from the manufacturing execution system.

A set of dashboards has been configured in order to show data and support analysis. For example, Figure 3 shows the global evolution of the consumption of each machine over a month period, the total consumption of the four machines per day and per machine, the weight of each machine in the total

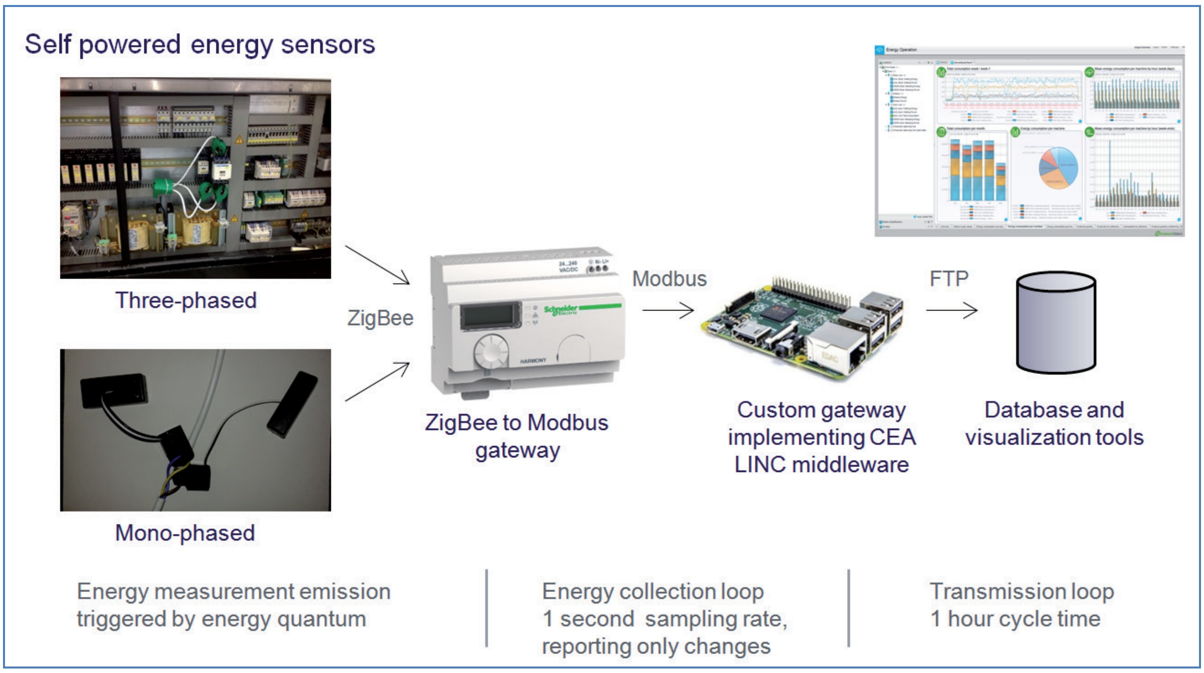

Figure 2. Energy Data acquisition system.

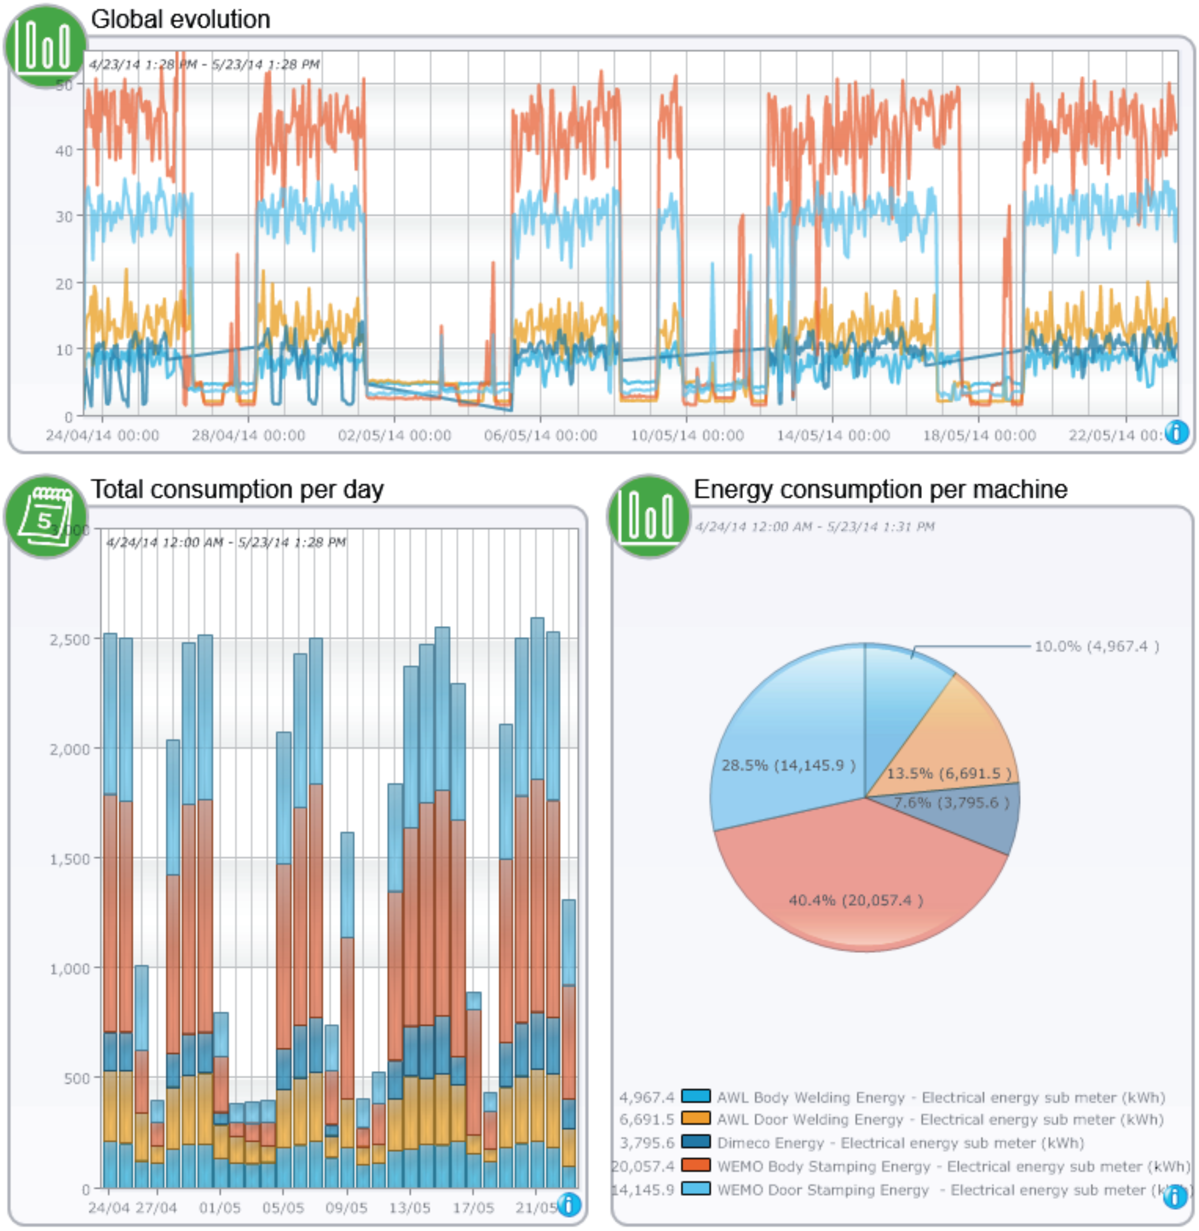

Figure 3. Energy Operation dashboard example.

consumption. Energy values are expressed in kWh in all the diagrams. We can see that stamping is the most consuming activity; almost 70 % of the global electricity consumption for the given month. With these dashboards, it is however difficult to determine whether the variation of electricity consumption from one reference to the other can be exploited to generate savings.

PRODUCTION DATA

Production data are extracted from the Manufacturing Execution Systems (MES) deployed into the plant. The extracted production data still need to be formatted and processed in order to retrieve the key information. It is to be noticed also that part of the production information is the result of a manual keying; therefore, the time stamping can suffer from some inaccuracy. In particular, start and end times of activities cannot be estimated accurately enough.

A FIRST MANUAL DATA MINING

Figure 4 shows the evolution of the power consumption of the body production line (in Watts), over a two-month period. We can see that the baseline power is about 7,500 W for the body part of the production lane. A similar visual analysis helped us to highlight that the baseline power consumption of the whole lane is about 10 kW when machines are idle. Moreover, machines are never switched off and their energy consumption during non-productive periods represents about 3 % of the total energy consumption.

Those observations are used in the following section, in order to check the consistency of the consumption model built. They also highlight flexibilities that can be exploited by the scheduler to reduce the electricity bill.

On the other hand, a first analysis showed that no production data are available on the Sundays, while the electricity consumption measured is higher than the baseload. The plant manager explained that this is due to Sunday maintenances, not logged in the production system. This observation confirmed that the energy acquisition system put in place is accurate enough to highlight anomalies.

Energy modelling of activities

From the gathered data, we wish to evaluate and model the energy consumption of the production activities executed in the plant. To enable this evaluation, we developed a procedure in two stages. Each of the two stages (further described below) is implemented as a web service on a Schneider Electric platform "in the cloud". The first service synchronizes the schedule and energy consumption data. The second service uses the result of this synchronization to characterize the energy consumption for each relevant "activity". The first service is optional: if the user already has synchronized data, the second service can be used directly.

The overall procedure shall be appropriate (sufficient) under the following theoretical assumptions:

- Sufficient production and energy data are available and these data are (globally) clean enough.
- If the energy consumption of the process depends on external factors, we are able to get the associated data and to take them into account in the modelling. For example, if the energy consumption varies with the external temperature, we consider that we can take into account the relevant weather data.
- If manufacturing activities can overlap in time and consume energy measured by the same sensor (e.g., either because the same machine works on several activities in pipeline or in parallel, or because a unique sensor is used for a group of machines), then the energy consumption of each activity shall not vary too much during the course of the activity. Otherwise, the activity would need to be decomposed in energetically homogeneous sub-activities.

Let us note that the second service provides model quality indicators. Bad values of these indicators are often reported when one of the above assumptions does not apply. A more complex "visual" version of the procedure is currently under study, in order to deal with more complex environments, for which the automatic procedure is insufficient.

Services implementation and data analysis were performed with RapidMiner 5 (Hofmann and Klinkenberg 2013), although in a version customized by Schneider Electric.

FIRST STAGE: ENERGY AND PRODUCTION DATA SYNCHRONIZATION

The first stage consists of synchronizing the energy consumption data (one time series per sensor) with production data extracted from the manufacturing execution system (or another legacy system) monitoring the plant. The synchronization algorithm first divides the time line according to the energy time series and determines the portion of each production activity in each time interval. Roughly speaking, the synchronization algorithm consists in re-sampling the production data at the sampling rate of the energy data.

Moreover, the synchronization algorithm offers the option to aggregate or not contiguous time intervals with similar production information into a unique time interval. When aggregation is done, similar time intervals are replaced by a unique time interval, summing both the energy consumption and the production amount data. Such an aggregation presents two advantages: first, the resulting table is smaller and easier to read; second, minor differences during the execution of the same activities are automatically averaged. It presents two drawbacks: the time intervals in the resulting table have different durations and do not coincide with the

Figure 4. Power consumption of the body line over time.

sampling rate of the energy consumption sensor; second, energy differences during the execution of the same activities are no longer visible.

SECOND STAGE: CONSTRUCTION OF THE ENERGY MODEL OF THE PLANT

The second stage consists in constructing an energy model. It relies on multivariate regression analysis to identify a consumption baseline (i.e., power consumed in the plant, even when it is idle) and the energy required in order to execute a unit of each type of activity in the plant. This can be expressed in the following way:

$$\forall i \in \{1, \dots, \mathrm{H}\},$$
$$total\_measured\_energy_i = \sum_{r \in R} (mean\_power\_of\_reference_r \times batch\_size_{r,i} \times \tau_i) + baseload \times \tau_i + error_i$$

Where:

H is the number of time periods considered;

τ_i is the duration of the time period i;

$total\_measured\_energy_i$ is the amount of energy measured by the energy sensor during time period i;

R is the set of identifiers of product references;

$mean\_power\_of\_reference_r$ is the power consumed by producing one unit of reference *r*; this is the coefficient we want to compute for reference r;

$batch\_size_{r,i}$ is the quantity of reference r produced during period i;

baseload is the power consumed by the line whatever the production;

$error_i$ is a term that represents regression model errors.

For both the baseline and each activity, usual statistic quality indicators (absolute error, t-statistic, and p-value) are also provided, enabling to qualify the confidence in the model.

VALIDATION AND APPLICATION

The reader can easily imagine that the result is very sensitive to the precision and accuracy of historical data. We have first validated the algorithm on a set of 34 benchmarks in the ideal case (with perfect energy consumption and manufacturing production data). Then, we have introduced perturbations of these data in order to quantify the algorithm sensitivity. More precisely, we considered the following types of perturbations of the data:

i. a significant energy consumption baseline (i.e., energy consumed when no manufacturing activity is ongoing in the plant);

ii. random fluctuations of the energy consumption and/or random errors from the sensors;

iii. reduced frequency of sensor measurements;

iv. imprecise manufacturing execution system data (start and end times of production activities).

Roughly, the obtained results show an appropriate robustness to perturbations (i) and (ii), provided enough measurements are available. Perturbations of type (iii) are well handled, as long as the frequency of sensor measurements remains below the characteristic frequency of the manufacturing process. Perturbations of type (iv) are the most annoying. Indeed, the information available for making the analysis quickly degrades with the imprecision of activity start and end times.

For the Schneider Electric plant, we applied the two stages on the data of the body production line, with two different synchronization strategies:

- First, we kept periods of the same duration τ . We can see a part of the results obtained on Figure 5. There is a line per kind of reference r , and the "value" column corresponds to the $mean\_power\_of\_reference_r$ value, in Watts per product unit. The three other columns are quality indicators (for example, pValue should be as near zero as possible). We obtained a baseload power consumption of 7,299 W for the body production line, consistent with the value of 7,500 W suggested by the visual analysis (see Figure 2). However, coefficients $mean\_power\_of\_reference_r$ obtained have a very bad confidence level. This can be explained by a significant variability of the energy consumption during the production of a specific reference.
- The second synchronization strategy aggregates time periods with similar production, as explained above. The baseload obtained is higher (which is obviously wrong), but the activities consumptions are more realistic. The high baseload is due to imprecision in the production log. Indeed, when an activity begins several minutes before being logged, machines are considered idle while they are actually producing a reference and consuming energy. As consecutive non-productive periods are aggregated, this phenomenon virtually increases the computed baseload. On the other hand, activities consumptions are probably lower than the real consumptions (due to the too high baseload), but consistent between one another (see Figure 6). Even if this result is not perfect, this is much more satisfactory than the previous result, as it enables to identify the relative consumption of each activity, and hence the activities to perform when the price of electricity is low.

Finally, we built an approximation of the activities energy consumption by keeping the baseload computed with non-aggregated data, and increasing the energy consumption of activities (computed with aggregated data), according to the difference between the two baseloads.

We found out that the power consumption of activities of the body production line varies between 13 kW and 23 kW, with a baseload of 7.3 kW. Due to the great variability of these consumptions, an efficient scheduling algorithm should be able to achieve savings on the electricity.

Indeed, this estimation of the energy consumption of each activity is used during scheduling to consider the cost of electricity consumption as one of the optimization criteria to minimize. This incites the scheduler to schedule the activities for which the electricity consumption is high at times at which the tariff is low. This way, a trade-off between storage costs and energy cost can be done, depending on the electricity tariff.

Row No.	name	value	error	tStatistic	pValue
421	Avg_productionAW16020116MCRA00002M0	0	1253.721598	0	1
422	Avg_productionAW15100223MCRA00003M0	0	1247.522016	0	1
423	Avg_productionAW12100676GHLA00062M0	994.936094	212.803318	4.675379	0.000003
424	Avg_productionAW13120363SRTA00044M0	427.319816	45.684442	9.353727	0
425	Avg_productionAENNA08361M0	0	243.265026	0	1
426	Avg_productionAW11120132TGFA40018M0	422.468515	78.641276	5.372096	0.000000
427	Avg_productionAW13110135TGFD10006M0	0	159.630895	0	1
428	Avg_productionAW16020486JCRA00003M0	6210.594032	1262.222747	4.920363	0.000001
429	Avg_productionAW14100675VMNA00006M0	0	1274.262075	0	1
430	Avg_productionA30348228_1M0	605.446678	101.797690	5.947548	0
431	Avg_productionAW15010097VMNB00002M0	0	1246.041148	0	1
432	Avg_productionAENNA08340M0	0	533.136979	0	1
433	Avg_productionA30349672_1M0	898.398615	219.565991	4.091702	0.000043
434	Avg_productionAW10040268TGFA70015M0	557.770394	99.196414	5.622889	0.000000
435	Avg_productionA30348784_1M0	0	88.273102	0	1
436	Avg_productionA18911119_1M0	480.680003	82.490633	5.827086	0
437	Avg_productionA18927440_1M0	511.016226	54.070091	9.450996	0
438	Avg_productionA18926924_1M0	440.193762	44.871857	9.810019	0
439	Constant	7298.940000	215.130005	33.928043	0

Figure 5. Sample of the regression results, on non-aggregated data.

Row No.	name	value	error	tStatistic	pValue
399	Avg_productionAS3D06913B20077M0	2953.481888	421.398350	7.008765	0
400	Avg_productionAW15100223MCRA00003M0	1077.810216	847.588638	1.271619	0.203533
401	Avg_productionAW12100676GHLA00062M0	771.717322	144.584323	5.337490	0.000000
402	Avg_productionAW13120363SRTA00044M0	390.414726	31.041020	12.577381	0
403	Avg_productionAENNA08361M0	605.522986	165.284325	3.663523	0.000250
404	Avg_productionAW11120132TGFA40018M0	384.822108	53.431408	7.202170	0
405	Avg_productionAW13110135TGFD10006M0	386.323274	108.456892	3.561998	0.000369
406	Avg_productionAW16020486JCRA00003M0	4708.357835	857.598353	5.490167	0.000000
407	Avg_productionA30348228_1M0	542.735261	69.164922	7.846973	0
408	Avg_productionAW15010097VMNB00002M0	382.635910	846.659515	0.451936	0.651323
409	Avg_productionAENNA08340M0	1316.873950	362.226967	3.635494	0.000279
410	Avg_productionA30349672_1M0	710.383284	149.242043	4.759941	0.000002
411	Avg_productionAW10040268TGFA70015M0	428.793810	73.671419	5.820355	0
412	Avg_productionA30348784_1M0	984.897664	499.641908	1.971207	0.048723
413	Avg_productionA18911119_1M0	415.807485	56.050032	7.418506	0
414	Avg_productionA18927440_1M0	432.071421	36.740501	11.760085	0
415	Avg_productionA18926924_1M0	360.716762	34.892763	10.337867	0
416	Constant	12819.176718	151.402511	84.669512	0

Figure 6. Sample of the regression results, on aggregated data.

Energy aware scheduling

The evolution of electricity tariffs, their variability in time, and the emerging possibility to save on energy costs by limiting or "shedding" electrical consumption during specific time periods, lead manufacturing plant managers to investigate whether they could schedule production differently to reduce their electricity bill while delivering their customers on time. Different concepts have been proposed on this subject, e.g., (Artigues, Lopez and Haït 2009), (Castro, Harjunkoski and Grossmann 2011), (Haït and Artigues 2011), (Pecero, et al. 2012). When no hard constraint applies on energy resources, the underlying optimisation problem takes as input:

i. a set of customer demands, for given quantities of given products, and for given due-dates;

ii. a set of recipes, describing the different activities required to manufacture the different products, their relationships (precedence constraints, possibly with minimal and maximal delays), as well as the resources (machines, energy) they require;

iii. a set of production orders (instantiating the recipes) aimed at satisfying the customer demands, together with information on their current status in the factory;

iv. time-varying parameters such as the electricity tariff to be applied over time.

The objective is to assign resources, start times and end times to each activity of each production order, in a manner that both satisfies the relevant precedence and resource constraints and establishes a compromise between multiple performance objectives. More precisely, the goal is to minimize tardiness penalties corresponding to late customer deliveries, energy consumption costs, other manufacturing costs (e.g., when different machines are able to perform the same activity without requiring the same amount of manpower), and storage costs.

This problem is a generalization of classical scheduling problems and is thus strongly NP-Hard. Nevertheless, several techniques can be used to find good solutions to such scheduling problems, e.g., mixed-integer linear programming (MILP), constraint programming (CP), meta-heuristics, etc. All have their advantages and drawbacks. The current version of our scheduler uses CP to generate an initial schedule (CHOCO3), which is then improved using our own large neighbourhood search based on a MILP model (CPLEX12.6).

MODELLING

In the context of the study reported here, each activity has a unique execution mode requiring a specific machine (of capacity 1) and a given amount of electrical power throughout the execution of the activity. The duration of each activity is known. A simple form of energy tariff is also assumed: time intervals (buckets) are defined and, for each bucket, a price per kWh is provided.

In this context, the main variables of the problem are the start times of production order activities. ST(i, p, j) denotes the start time of the j^{th} activity of production order p, instantiating recipe i. ET(i, p, j) denotes the end time of the same activity and is equal to ST(i, p, j) + pt(i, p, j) where pt(i, p, j) is the known duration of the activity. Precedence constraints between two activities j_1 and j_2 of the same recipe translate into inequalities of the form ST(i, p, j_2) – ET(i, p, j_1) ≥ $\delta_{1,2}$ where $\delta_{1,2}$ is a constant positive or negative delay between the activities. Machine sharing constraints, stating that two activities requiring the same machine cannot overlap in time, take a disjunctive form [ST(i_1, p_1, j_1) – ET(i_2, p_2, j_2) ≥ 0 OR ST(i_2, p_2, j_2) – ET(i_1, p_1, j_1) ≥ 0]. In a MILP model, a Boolean variable PRECEDES(i_1, p_1, j_1, i_2, p_2, j_2) with value 1 if activity (i_1, p_1, j_1) precedes activity (i_2, p_2, j_2) and value 0 otherwise is introduced to replace the disjunction by two inequalities (as in (Applegate and Cook 1991)). Storage costs of intermediate products are expressed as a linear combination of delays between the end times of the activities that produce the intermediate product and the start times of the activities that consume it. Tardiness costs are expressed as a linear combination of the delays between the due dates of the demands and the end times of the activities that produce final products for these demands: TCOST(i, p, j, d) = w_d * $f_{i,p,j,d}$ * max(0, ET(i, p, j) – dt_d) where w_d is a weight associated to demand d, $f_{i,p,j,d}$ is the quantity of final product produced by activity (i, p, j) to serve demand d, and dt_d is the time at which demand d is due. In a MILP model, this equality is replaced by two inequalities TCOST(i, p, j, d) ≥ w_d * $f_{i,p,j,d}$ * (ET(i, p, j) – dt_d) and TCOST(i, p, j, d) ≥ 0.

Energy cost is the most difficult to represent as the cost incurred for each activity depends in a complex manner on when it executes. Following (Erschler, Lopez and Thuriot 1991), the overlap between an activity (i, p, j) and a time bucket [st_k et_k) is given as max(0, min(et_k – st_k, et_k – ST(i, p, j), ET(i, p, j) – st_k, pt(i, p, j))). This temporal overlap OVERLAP(i, p, j, k) can then be multiplied by the power required by activity (i, p, j) and the price per kWh over time bucket [st_k et_k) to get the energy cost. This looks simple at first glance, but in CP the expression of the overlap induces a lot of hidden variables (except if you develop a special global constraint, which we have not done yet), and in MILP the max(0, min(…)) structure also causes difficulties. After various trials, we converged on a model with 2 intermediate variables, one continuous OMAX(i, p, j, k) and one Boolean OBOOL(i, p, j, k), with the following constraints:

- OMAX(i, p, j, k) ≤ min(et_k – st_k, et_k – ST(i, p, j), ET(i, p, j) – st_k, pt(i, p, j)), which easily translates into 4 linear inequalities.
- 0 ≤ OVERLAP(i, p, j, k) ≤ (et_k – st_k) * OBOOL(i, p, j, k)
- Σ_k OVERLAP(i, p, j, k) = pt(i, p, j)
- OVERLAP(i, p, j, k) ≤ OMAX(i, p, j, k) + bigM * (1 – OBOOL(i, p, j, k)) where "bigM" is a large constant, e.g., the time horizon of the scheduling problem.

Moreover, we are currently working on handling activities preemption and piecewise-linear energy cost functions.

EXPERIMENTAL RESULTS

To assess the quality of various algorithmic combinations, we have used 38 benchmark instances from the literature (with recommended computing time limits) (Le Pape and Robert 2007), initially including no energy costs, to which we added energy consumption data. The reader is referred to (German, Desdouits and Le Pape 2015) for a description of the methods we tried and results. Depending on the instance and tariff, we observed that between 0 % and 15 % of the electricity bill could be saved without affecting tardiness costs.

In order to compute production schedules for the studied plant, we had to feed the above model with real plant data:

- energy consumption of production activities are computed using the method described in the previous section;
- processing times of activities are computed from production data;
- the electricity tariff is a typical peak/off-peak French tariff;
- due dates of customer demands are extracted from production data;
- tardiness costs depend on the production order: the tardiness penalty w_d for a customer order is set to 10 times the penalty for a stock order.

First results show that production schedules can be computed for our manufacturing plant, but for very short time horizon, because of the number of activities. The obtained schedules seem very close to the ones currently applied into the plant, which is encouraging for future work. In the future, our scheduler will have to handle at least one-week data to be usable in the real plant. Performance improvements, as well as large-scale experimentations on real data, are currently ongoing.

In the future, the goal is to build a simple decision-aid tool that could be tested in the plant. This will be achieved by:

1. computing a small set of nearly optimal production plans over a week period;
2. visualize these schedules and highlight the associated tardiness, energy and storage costs;

3. sending the result to the plant manager, who could use it to decide when to product which reference.

The plant manager would be asked which production plan he chooses and why. That would allow us to improve the decision-aid tool, in an iterative process.

Conclusion

In this paper, a methodology to enable automatic computation of optimal production plans is presented, as well as the application of this method to a real manufacturing plant.

In order to compute production plans, we had to know how much energy is consumed in producing each kind of reference. Thus, we measured energy consumed by each machine, over a long time period. Then we synchronized energy measures with production log, and used a regression algorithm to compute approximated energy consumption of each activity, as well as the baseload.

The approach has been validated on clean datasets from benchmark files. In addition, we confirmed for the studied plant, that energy consumption varies, depending on the kind of reference produced. This potentially enables savings on the electricity bill, by shedding energy-consuming activities when the electricity is cheap.

However, contrarily to what we expected, we did not find out any correlation between the size of the cabinet produced and the energy consumption. We will have to perform other experiments to understand what the energy consumption really depends on.

Moreover, a scheduling model of the plant is proposed, that still needs to be refined but that already allows us to handle small instances (e.g. instances with a few production orders). Future work will allow us to quantify possible savings of the electricity bill, depending on the electricity cost function. Indeed, when energy consumption is different from reference to reference, including energy cost in the optimization cost function can be particularly profitable. However, we still have to handle some disruptions in the data and to evaluate the economical interest of our approach for the real plant. The attractiveness of this approach is linked to the inherent flexibility of the process and the variation amplitude of the energy cost for a given energy tariff, compared to storage costs.

References

Applegate, David, and William Cook. "A computational study of the job-shop scheduling problem." *ORSA Journal on computing*, 1991: 149–156.

Artigues, Christian, Pierre Lopez, and Alain Haït. "Scheduling under energy constraints." *IESM.* Montreal, Canada, 2009.

Castro, Pedro M, Iiro Harjunkoski, and Ignacio E Grossmann. "Optimal scheduling of continuous plants with energy constraints." *Computers & chemical engineering*, 2011: 372–387.

Dietmair, A, and A Verl. "Energy consumption forecasting and optimisation for tool machines." *Energy*, 2009: 63.

Erschler, Jacques, Pierre Lopez, and Catherine Thuriot. "Raisonnement temporel sous contraintes de ressource et problèmes d'ordonnancement." *Revue d'intelligence artificielle*, 1991: 7–32.

Fysikopoulos, Apostolos, Georgios Pastras, Theocharis Alexopoulos, and George Chryssolouris. "On a generalized approach to manufacturing energy efficiency." *The International Journal of Advanced Manufacturing Technology*, 2014: 1437–1452.

German, Grigori, Chloé Desdouits, and Claude Le Pape. "Energy optimization in a manufacturing plant." *ROADEF.* 2015.

Haït, Alain, and Christian Artigues. "Un modèle PLNE à temps continu pour l'ordonnancement d'une aciérie avec régulation de consommation d'énergie." 2011.

Hofmann, Markus, and Ralf Klinkenberg. *RapidMiner: Data Mining Use Cases and Business Analytics Applications.* Chapman & Hall/CRC, 2013.

Le Pape, Claude, and A Robert. "Jeux de données pour l'évaluation d'algorithmes de planification et ordonnancement." *Conférence conjointe FRANCORO V/ROADEF.* 2007.

Le, Cao Vinh, et al. "Classification of energy consumption patterns for energy audit and machine scheduling in industrial manufacturing systems." *Transactions of the Institute of Measurement and Control*, 2013: 583–592.

Louvel, Maxime, and François Pacull. "LINC: A Compact Yet Powerful Coordination Environment." *Coordination Models and Languages: 16th IFIP WG 6.1.* Berlin, Germany: Springer Berlin Heidelberg, 2014. 83–98.

Pecero, Johnatan E., Héctor Joaquin Fraire Huacuja, Pascal Bouvry, Alejandro Santiago Pineda, Mario César Lòpez Locés, and Juan Javier Gonzàlez Barbosa. "On the energy optimization for precedence constrained applications using local search algorithms." *HPCS.* 2012.

Sugimori, Y., K. Kusunoki, F. Cho, and S. Uchikawa. "Toyota production system and Kanban system Materialization of just-in-time and respect-for-human system." *International Journal of Production Research* 15, no. 6 (1977): 553–564.

Acknowledgements

The work presented here has been partially funded by the Artemis Arrowhead project under grant agreement number 332987 (Artemis/Ecsel Joint Undertaking, supported by the European Commission and French Public Authorities).