



Optimizing Chemical Processes

Dr. Patrick Bangert, algorithmica technologies GmbH

A chemical plant's efficiency and profitability can be optimized using mathematical modeling. The optimization tells plant operators which set-points should be changed to obtain the maximum profitability. This method requires no engineering changes to be made to the plant. We show that a profitability increase of approximately 6% was possible in a specific chemical plant producing silanes, with an overall yield increase of 5.1% and an increase of 2.9% for the most profitable end product.

Introduction

Like other industries, the chemical industry constantly aims to enhance the profitability of its plants by increasing the production yield. Such increase may be achieved either via engineering changes, which are quite cost intensive, or via operational changes. Operational changes can increase yield and profitability without actually changing the equipment and the processes of the plant. The question is: What operational changes are needed to increase yield, or better still, to optimize yield under the given circumstances? Here, we compute the action required to achieve optimal yield at any time using machine learning. This method develops a mathematical model of the process based purely on historical data and is therefore very fast and economical to employ.



The Problem

While some smaller processes are automated using various intricate technologies, the overall processes are most often controlled by human operators. With operators working in shifts, no single operator controls the plant over the long-term but usually only for the time of his or her shift. It can be observed that the efficiency of the plant oscillates in a rough eight hour pattern showing that human decision making has a significant influence on the efficiency of the plant. Not only are some operators better than others, but it is also quite difficult to transfer the know-how and experience of the best operators to those operators who are less experienced and knowledgeable. Even where regular knowledge transfer systems are in place, this transfer may work to a certain extent, but usually not to a maximum effect. Hence, there are good operators and less than good operators.

Furthermore, the plant usually outputs several thousand measurements at high cadence. An operator cannot possibly keep track of even the most important of these at all times. The degree of complexity is most often too great for the human mind to handle and as a consequence, suboptimal decisions are taken, even by the best operators.

The challenge, then, is to optimize the efficiency of the plant by systematizing the way the plant is operated, depending less on the intuition of the operators and more on hard evidence. How is this done?

The basis for the optimization model is the historical data from the data historian that keeps track of the numerical values of different variables. The knowledge and experience of the operators are thus plainly visible in the data. If the history is long and detailed enough, this information is effectively all one needs to know about the plant. A human being could not use this information to learn about the plant because of the sheer volume of data. Machine learning is designed to extract the underlying pattern in a large set of numerical data and produce a simple equation – one we may use to make optimal decisions.

The purpose of the chemical plant to be examined here is to input certain chemical compounds (such as silicon and hydrogen compounds) and to subject them to the Müller-Rochow Synthesis (see figure 1) in order to obtain Di methyl chloride silanes ((CH₃)SiCl₂) and Tri methyl chloride silanes ((CH₃)SiCl₃), hereafter referred to as Di and Tri.

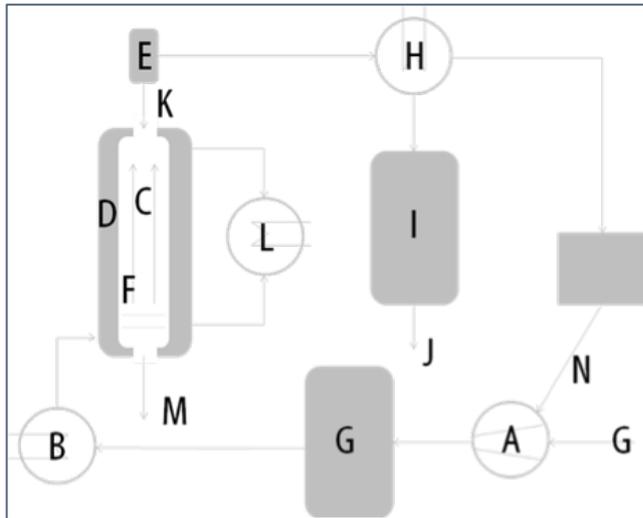


Figure 1: Sketch for the Müller-Rochow-Synthesis plant including: (A) Compressor, (B) Vaporizer, (C) Fluidized bed reactor, (D) Cooling jacket, (E) Cyclone, (F) Silicon / Copper (catalyst), (G) Methyl chloride, (H) Condenser, (I) Raw Silane, (J) To the distillation, (K) Silicon / Copper dust, (L) Heat exchanger, (M) Remainder, and (N) Back-flow methyl chloride.

The Solution

The optimum of profitability is achieved with a maximum of yield for a minimum of raw materials. The selectivity of particular end products is influenced by the amount of various catalyst and promoter materials added as well as diverse process variables such as temperatures and pressures. All these must be regulated to the best possible operation considering a number of features that the operator cannot control at all, such as the temperature of the environment or the quality of the raw materials.

To simplify matters, we shall view the entire plant as a black box. Raw materials go into the box and product comes out of the box. The box has some gauges with which we can sense what is going on inside the box, and it also has some dials with which we can control what goes on inside. Based on this information, we will want to determine the relationship of input to output given the restriction of the gauges (which we cannot control) and the dials (which we can control). The process of discovering this relationship is machine learning, which we will not treat here. The important thing to know is that machine learning is done automatically without the manual addition of human knowledge – it operates purely on the historical data of the plant.

The result is a set of formulae that describe what comes out of the plant for any particular input, dial setting and gauge measurement. This equation may then be turned around so that we can ask: For the current gauge measurement, which we must take as given since we cannot control it, the question is: What is the optimal dial setting and input of raw material? The term "optimal" is defined here as the highest profitability. The answer to that question is a concrete action that must be implemented by the operator to achieve that optimum.

Results

The results of the optimization were obtained in an experimental period lasting three months and encompassing three reactors. During the evaluation period, the operator implemented only the actions deemed useful by him or her. During the usage period, the operator implemented all actions computed by the model. During the reference period, the operator did not use the optimization model at all. The results for all periods are plotted in figure 2.

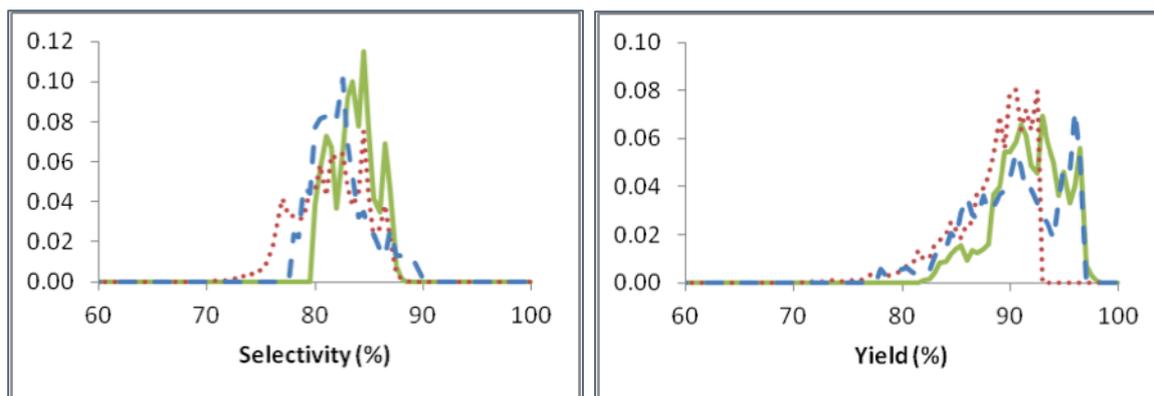


Figure 2: The probability distribution functions for selectivity and yield of Di for periods in which the optimization was not used (reference period: red dotted line), used whenever deemed useful (blue dashed line) and used fully (green solid line).

It is apparent, from these images alone, that we increase the selectivity and the yield with more use of the optimization and that we decrease the variance (i.e. the spreading out or distribution of the results) of both selectivity and yield as well. Decreasing the variance is desirable because it yields a more stable reaction over the long term and thus produces its output more uniformly over time. Numerically, the results are displayed in table 1 below.

	Selectivity (%)	Yield (%)
Reference	79.8 ± 3.6	86.6 ± 4.2
Evaluation	79.9 ± 2.5	89.7 ± 4.3
Usage	82.7 ± 1.9	91.7 ± 3.2

Table 1: For both selectivity and yield, we compute the mean ± the standard deviation for all three periods.

The results show that the selectivity can be increased by approximately 2.9% and the yield by approximately 5.1% absolute by comparing the usage with the reference period. Together these two factors yield an increase in profitability of approximately 6% in the plant.

It is to be emphasized that this profitability increase of 6% was made possible through a change of operator behavior only (as assisted by the computational optimization) and no capital expenditures were necessary.



The practical setup of this optimization took approximately two days of time for the operating personnel. The computation time for the computer to construct the necessary functions was about one month. The computer interfaces for input and output of the data are standardized in the industry and can thus be applied readily without delay.

Thus, within about one month, the model can be fully operational without occupying the operators for much time. The approach is thus practicable in a real industrial plant.

References

- [1] Bishop, C.M.: Pattern Recognition and Machine Learning. Heidelberg: Springer 2006
- [2] Mandic, D., Chambers, J.: Recurrent Neural Networks for Prediction: Learning Algorithms, Architectures and Stability. Hoboken: Wiley 2001
- [3] Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes. Cambridge: Cambridge University Press 2010
- [4] Bangert, P.D., Czernitzky, J.-A. (2010): Increase of Overall Combined-Heat-and-Power (CHP) Efficiency via Mathematical Modeling. Proceedings of the VGB PowerTech e.V. Conference "Fachtagung Dampferzeuger, Industrie- und Heizkraftwerke 2010"
- [5] Bangert, P.D. (2010): Two-Day Advance Prediction of a Blade Tear on a Steam Turbine of a Coal Power Plant. 175--182. in Link, M. (2010): Schwingungsanalyse & Identifikation -- Proceedings of the Conference of 23./24.03.2010 in Leonberg, Germany. (VDI-Berichte No. 2093) ISBN: 978-3180920931.
- [6] Bangert, P.D. (2009): Braids and Knots. 1--73. in Berger, M.A. et. al. (2009): Lectures on Topological Fluid Mechanics. (Springer Verlag, Berlin) ISBN: 978-3642008368.
- [7] Bangert, P.D. (2005): Optimizing Simulated Annealing. Vol. 3. 198-202. Proceedings of SCI 2005 -- The 9th World Multi-Conference on Systemics, Cybernetics and Informatics from 10.-13.07.2005 in Orlando, FL, USA.



Mathematics of the Optimization

Any measurement made in the plant falls into one of three categories: The operator can directly modify the value by changing a set-point, cannot modify the value at all (because the value is set by the world outside the plant) or can indirectly modify it by changing some set-point (e.g. vibrations can be increased by increasing the pressure). We call those controllable, uncontrollable and semi-controllable measurements. This is similar to the distinction that mathematicians make between variables and parameters. Both are changing quantities but parameters are fixed by the world and variables can be changed by the mathematician.

The goal, in this case profit, is a function of both controllable and semi-controllable values as variables and the uncontrollable values as parameters, i.e. $g(c, s; u)$.

What we want to compute is the set of controllable values that yields the largest possible profit given that we cannot modify the uncontrollable values. Thus, we are looking for the point that has highest profit among all those points that have the same uncontrollable values as the present point. That is why, we use machine learning to get a model of the plant. This means a set of formulas that compute the semi-controllable values as a function of the controllable values as variables and the uncontrollable values as parameters, i.e. $s = f(c; u)$.

Machine learning [1] has the goal of determining functions from data by using particularly flexible functional templates as a basis and computing the best possible values for the parameters of this template such that the function arrived at fits the data in the best possible way. Particular emphasis must be paid to the fact that the process depends upon its own history. A special class of methods, known as recurrent neural networks, is particularly good at capturing this time-dependence and so they are used here [2]. It would be beyond the scope of this paper to describe how these are arrived at.

We may now substitute the model into the goal to get $g(c, f(c; u); u)$. The major change here is that we were able to get rid of all the measurements in the plant that are not directly accessible to us as actions or boundary conditions. As the uncontrollable must not change, we may now use optimization techniques to find the set of controllable values that maximize the goal, i.e. $g(c', f(c'; u); u)$. We use simulated annealing as this is a technique that is able to find the global optimum in very complex situations [6,7].

The difference between the optimal and current controllable values ($c' - c$) is then the action that the operator must implement in order to achieve the best possible profit. This action is a time-dependent action that may change from moment to moment. As such, it is frequently not possible to implement this action in a single step.



Figure 3 displays this problem graphically using real data taken from the current process. The two axes on the horizontal plane indicate two controllable variables and the vertical axis displays the goal function. We can easily see that the change in a controllable variable can produce a dramatic change in the goal. The two paths displayed represent the reactions to the current situation by a human operator (the upper path) and the computer program (the lower path). They initially begin on the left at the current operational point. Because of their differing operational philosophies, the paths deviate and eventually arrive at different final states. This is a practical example of the human operator making decisions that he believes are best but that are, in fact, not the best possible. The computed path is better than the human path by about 5% in the profit goal.

Done properly, both steps can be automated such that the computer continuously keeps track of the optimum point and alerts the user to necessary actions in order to keep the physical process at the computed optimum. The methods of machine learning work in such a way as to update the model with every new measurement. Therefore, the model validates itself over time as it always checks its performance against experimental verification and alters itself if necessary. Even if changes are made to the process, the model will learn them autonomously after some time has passed.

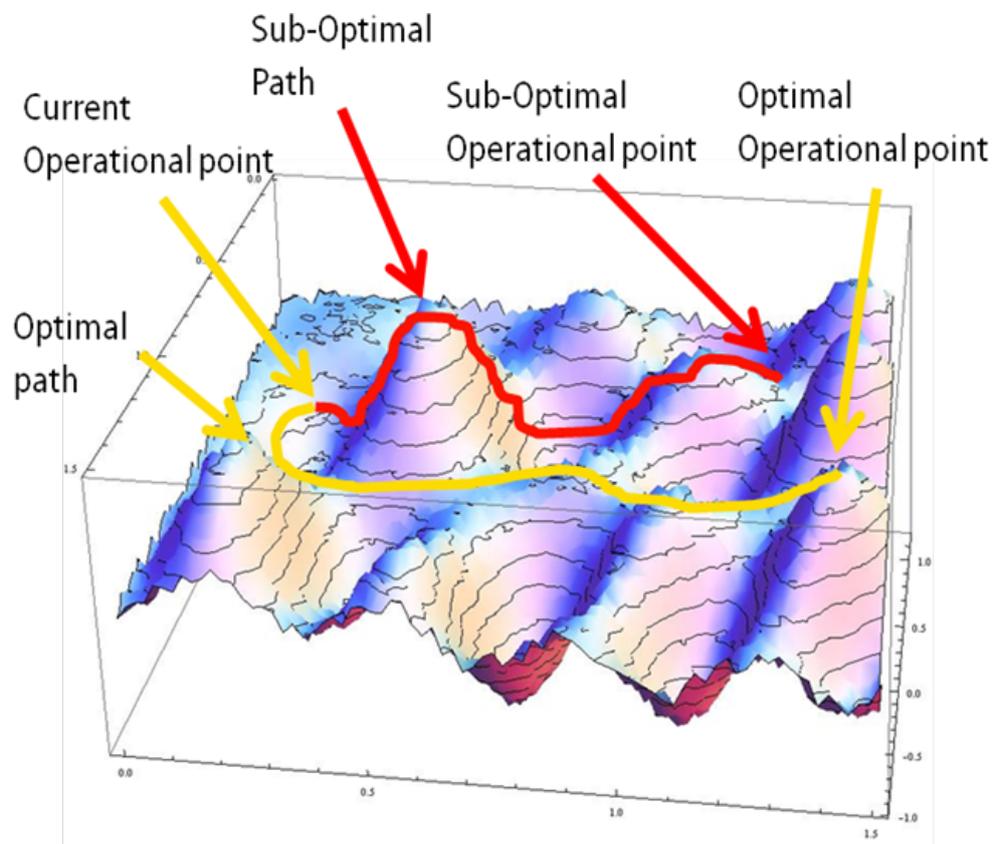


Figure 3: The dependency of the goal on two controllable variables. The upper path displays the reaction of a human operator and the lower path displays the reaction of the optimization system.

Specific Design

For the specific current application, the molecules are produced in three separate reactors and then brought together for shipment. We are to optimize the global performance of the plant but are able to make changes for each reactor separately. We now outline the variables that were under our influence. These should be understood as being per reactor.

In this case, the controllable variables are: Temperature of the reactor, amount of raw material to the jet mill, steam pressure to the jet mill, amount of Methylene Chloride (MeCl) to the reactor, pressure of the reactor and others relating to the processes before the synthesis itself.

The uncontrollable variables are: X-ray fluorescence spectroscopy measurements on 17 different elements. The semi-controllable variables are the other variables that are measured in the system. In total, there were almost 1000 variables measured at different cadences.

The goal function is the financial gain of the reaction. We compute the input raw materials and the output end products. Each amount is multiplied with the currently relevant financial cost or revenue.