What about reactor APC?

By
Y. Zak Friedman, PhD
Principal Consultant

June 2007

What about reactor APC?

The premise of APC is that there is money to be made by pushing a unit against its limits, while keeping the product qualities at target. This concept is easy to understand but difficult to implement, because while some qualities and unit constraints are measured, many are not. Expressing unmeasured process parameters in a numerical format suitable for our multi-variable control tools is the challenge of our industry. Applications that can do it mine gold, and those that cannot – join the high pile of APC garbage. Miles of paper have already been devoted to inferential modeling of product qualities, and hence, in this editorial I would like to concentrate mostly on other unmeasured constraints that profoundly affect unit performance. You can structure an APC application to push the unit to extremes, only to find out months later that a reactor has fouled up and production must now be cut to meet the desired run length. Ignore the effect of unit conditions on reactor fouling to become first a hero, then a pariah.

Whatever a unit makes, the products are separated usually by distillation, and many product quality inferential methods rely, either via first principles or regression, on distillation thermodynamics. Column temperature profile, pressure and internal reflux conform to thermodynamic laws, and hence provide information about the products. But the luxury of well known thermodynamic principles does not exist when dealing with chemical reactions. How should we handle this problem? Ignoring the effect of unit conditions on by-products and catalyst deterioration is not an option, as that would render the APC counterproductive. As a minimum we should attempt to model deposition of coke on catalyst or other surfaces, reactors, this being the most widespread run time constraint.

Consider the influence of ethylene cracking furnaces run length on advanced control. Ethylene plants have a battery of cracking furnaces, processing different feeds at different conditions. These furnaces gradually coke up, and once a month or so come down for decoking, which takes several days. Over-cracking shortens the run length, resulting in throughput penalty, whereas under-cracking results in yield penalty. Decades ago ethylene cracking furnace APC was based on expensive, high maintenance, furnace transfer line analyzers for determination of the right cracking severity for each individual furnace. Such approach provided reasonable severity control but no sexy optimization. Today, people are reluctant to make use of transfer line analyzers, and yet ethylene plant APC has evolved into a sophisticated closed loop optimization application. On line ethylene plant optimization APC has been claimed to be more successful than other industrial closed loop optimizers because the prices of feed and products are known, and kinetic models exist to predict cracking furnace yield patterns.

Not being convinced that ethylene furnace yield models are verifiable for partially fouled furnaces without transfer line analyzers, I still accept the notion that when feed economics, composition and product prices are known, optimization works better. The optimizer determines furnace COT (coil outlet temperature) and throughput for each of the furnaces, to optimize total production. One should pause here for a minute and reflect about the danger of making the wrong severity decision. If the object is to maximize ethylene production, the penalty for premature decoking is quite severe. Would you risk your reputation without seeing evidence that the kinetic model can indeed estimate furnace yield and run length? To my knowledge comparisons between predicted versus actual run lengths are yet to be presented in the open literature. Suppose one succeeds in modeling transfer line composition but not furnace run length, could that be a basis for plant wide optimization? Would it not be better to use simple old style severity control, and remove the furnaces out of the optimizer? A much reduced scope of optimizing the product separation section would still be feasible.

Moving from ethylene crackers to refinery units, reformers come to mind as units whose performance is sensitive to the rate of catalyst coking. Reformers converting naphtha aliphatic molecules into aromatic ones, and coke is a by-product of the reaction. Modern CCRs (continuous regeneration reformers) are not as sensitive to the rate of catalyst fouling as older semi-regenerative units, but even so, catalyst travels through the reactor for a week or more, before moving into the regenerator, and should the catalyst accumulate too much coke, that coke would overload the regenerator and force a throughput cut. Reformer APC optimizers must consider the trade-off between catalyst fouling versus higher severity or throughput, only that is not easily done. The reforming reactions are a function of feed composition whereas naphtha PNA (paraffin, naphthene, aromatic) content is not measured frequently. Without PNA knowledge it is neither possible to estimate extent of aromatization, nor coke deposition rate. We know of only one commercial model capable of inferring PNA, extent of aromatization and rate of fouling from unit conditions. Some units use on-stream octane or aromatics analyzers as APC input to help control reactor severity, and that is OK as long as the APC does not attempt to maximize feed. If an APC is configured for feed maximization, it better be able to determine the rate of catalyst coking or it would drive the unit right into a regenerator constraint.

Delayed coking is another important refinery unit, perhaps the most profitable one, converting vacuum pitch into coke and distillates. Implementing APC on a coker is one of the most difficult control applications, because cokers are semi-continuous, with coke drums as thermal reactors being filled up with coke, then cooled, opened, emptied, closed, reheated with process streams and reconnected as reactors again. Coping with the huge disturbances of drum switching is quite difficult, and in addition the coker APC has a dense control matrix, meaning – each manipulated variable affects many constraints. Controlling dense matrix units smoothly without violating constraints requires very good dynamic control models. Of relevance to our discussion today is the unmeasured but important constraints of level of coke in the coke drum. The level is measured only by a level switch, activated once when the level reaches about 70% and a second time when it is at 90% and an immediate drum switch is called for. This level is often a throughput limit, and even if it is not a limit, advanced knowledge of drum fill rate is of value because drum emptying is a manual labor intensive operation that must be organized hours in advance. A coker APC I was involved in recently was not constrained by drum fill rate, but still we had implemented drum fill model as an indicator.

Hydrotreaters, from desulfurization to hydrocracking to selective desulfurization, are all about bringing the unit feed in contact with hydrogen and catalyst. While these units have, as a result of ever tightening specifications, become significant refinery bottlenecks, their APC sophistication has lagged. APC applications tackle separation downstream of the hydrotreater, but as far as the reactor, APC by and large follow the licensor recommendations with respect to treat ratio and temperatures. I am not aware of attempts to trade-off severity, throughput and run length in real time. Imagine how much money APC could make if we had a model of run length versus reactor conditions. We could increase throughput on some feeds by 10% or so, opening up a total refinery bottleneck. I am not aware that anyone has tried a run length model in closed loop, but the model must exist in some form. Otherwise, how do people know what size reactor is needed?

Lastly we note that coking is many times a constraint even on units that do not have reactors. Vacuum units typically operate at 400°C, and at those temperatures one has to pay attention to the possibility of coking transfer line or trays. Fear of coking dictate operational limits on COT and wetting of packing sections, and those limits are conservative, in that they take into account the worst feed. If we are able to more accurately assess the relations between COT, and column liquid loading, versus run length the vacuum unit yields would much improve.

I have said in a previous editorial that inferential product property modelling is the Achilles heel of our industry. It would be appropriate expand the Achilles heel to include reactor fouling rate. I would assess that if we could infer reactor fouling rate in real time we would be able to increase the throughput to that unit by 10% or so without any added investment.