

## Refinery Fuel System Control and Optimisation

Pavlos Ioakimidis, Henrik Terndrup, Andrew Trenchard

### Contact information (E-mail)

Pavlos Ioakimidis: [pioakimidis@hellenic-petroleum.gr](mailto:pioakimidis@hellenic-petroleum.gr)

Henrik Terndrup: [henrik.terndrup@tcatech.com](mailto:henrik.terndrup@tcatech.com)

Andrew Trenchard: [andy.trenchard@honeywell.com](mailto:andy.trenchard@honeywell.com)

### **Abstract**

This paper discusses an application of Honeywell's Robust Multivariable Predictive Control Technology (RMPCT) to the control and optimisation of a refinery fuel system. A plant-wide fuel system application was commissioned in August 2003 and has delivered significant financial and operational benefits.

The advanced control and optimisation application for the refinery fuel system was part of a larger APC program implemented by Honeywell Process Solutions and Intec at the Hellenic Petroleum Aspropyrgos refinery.

This paper briefly describes the fuel system of the Aspropyrgos refinery and the architecture of the fuel system application. The main focus, however, is on the practical experience gained in the period after the commissioning of the application.

The fuel system of the Aspropyrgos refinery provides fuel to two gas turbines as well as to the heaters of the process and utility areas of the refinery. The fuel supply consists of fuel oil from storage and fuel gas produced in the process units. One of the key objectives of the fuel system optimisation is to balance supply and demand in an optimum way in the presence of frequent disturbances primarily from the main fuel gas producers.

One of the challenges of the project was related to the fact that the application is distributed over several operator consoles located in two separate control centres. The

paper describes the application architecture, the operator interface and the infrastructure that supports the application along with other aspects of reliability and operator acceptance.

The application manipulates the base load fuel on a number of dual-fired furnaces. The requirements, in terms of burner configuration, heater instrumentation and regulatory control functionality, in order to support the fuel system application are discussed. Frequent and significant load disturbances from the fuel gas producers often require significant changes to the number of gas burners in service. The handling of this problem and the associated impact on the model gains are described.

The application has been designed to be flexible and open in view of future extensions, e.g. additional dual-fired furnaces or additional refinery fuels.

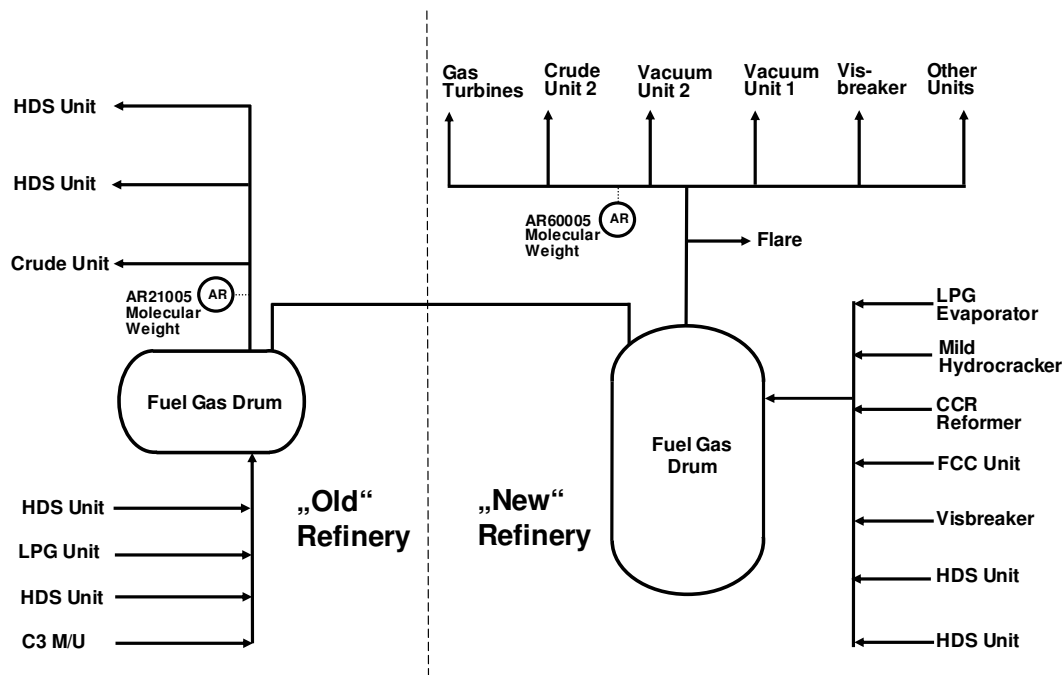
The implementation of the fuel system control and optimisation application has enabled Hellenic Petroleum to substantially reduce LPG losses to the fuel system and reduce flaring of fuel gas. The benefits are illustrated by comparison of pre-project base line data and post audit data.

### **Aspropyrgos Refinery Fuel System**

The Aspropyrgos refinery of Hellenic Petroleum is a quite complex 145,000 BPSD refinery, located near Athens, Greece.

The fuel gas system of the refinery consists of an extensive network of piping and vessels that connects a series of process units producing fuel gas with a number of fuel gas consumers (process unit furnaces, gas turbines, etc.) – a simplified sketch is provided in Figure 1.

Figure 1  
Aspropyrgos Refinery Fuel Gas System



Synoptically, i.e. from a layout point of view, the Aspropyrgos fuel gas system can be considered in two parts:

1. The fuel gas network of the process units of the “old” refinery (i.e. before 1985). This system supplies a series of process units that formed a classical hydro-skimming refinery and that are connected via fuel gas production and distribution headers and a relatively small fuel gas drum.
2. The fuel gas network of the “new” refinery units (i.e. after 1985). This system supplies a series of process units, which, together with units of the “old” refinery, form a complex, conversion based, refinery (including FCC and Mild Hydrocracking units). The new process units are connected via separate fuel gas production and distribution headers and a relatively large fuel gas drum. Naturally, this network is far more extensive than the previous one. Also, the major fuel gas producers of the integrated refinery complex are all in the “new” refinery. The two networks communicate via piping that connects the two fuel gas drums.

Moreover, two gas turbines are connected to the fuel gas network. The main fuel supply to these gas turbines is fuel gas, but the possibility to use propane as fuel, as supplement to fuel gas, also exists in order to satisfy the required load.

The fuel gas producers are located in the various process units (e.g. HDS, FCC, CCR, Hydrocracker, Hydrogen Plant, Visbreaker). Quantitatively, the most important fuel gas producers are the FCC and the Hydrocracking units.

The main consumers are the two gas turbines and various furnaces within the process units (e.g. CDU's, VDU's, FCC, CCR, Hydrocracker, Hydrogen Plant, Visbreaker). Quantitatively, the most important consumers are the gas turbines and the large furnaces of the CDU's, the VDU's and the CCR unit. The utility boilers of the refinery are single fuel boilers firing only fuel oil.

### **Characteristics of the Aspropyrgos Fuel Gas Network**

As result of the production and consumption levels of the various producers and consumers, the fuel gas network of the Aspropyrgos refinery exhibits the following operational characteristics:

1. The “old” refinery network is usually in deficit (more consumption than production).
2. However, the “new” refinery is usually in surplus. Hence, there is normally a net flow of fuel gas from the “new” refinery to the “old” refinery.
3. If we consider only normal operating disturbances, i.e. we do not consider abnormal situations like an unforeseen unit shutdown, the most important disturbances in the total fuel gas production level are caused by the main producers, namely the FCC and Hydrocracking Units. The magnitude of these disturbances is usually around  $\pm 5-6$  % of total fuel gas production and it is characteristic for these disturbances that they are neither particularly fast, nor particularly slow, i.e. it usually takes 2-3 hours to decline from the higher production level to the lower and vice versa. We can typically trace the origin of these disturbances either to changes in the operational parameters, e.g. feed flow

or feedstock type, of the corresponding units, or to changes in ambient conditions, most importantly of course the ambient temperature. In general, however, it is not considered desirable to try to attenuate the disturbances by manipulation of parameters that affect the production, e.g. unit feed rates or severities of operation.

4. On the other hand, there are also disturbances in the total fuel gas consumption level. Again, we consider only normal operating disturbances and we do not consider abnormal situations, like an unforeseen gas turbine shutdown. The origin of these disturbances is usually related to changes to operational parameters (e.g. feed flow or feedstock type), which trigger variations in the heat duty of the corresponding furnaces. Their magnitudes are usually around  $\pm 1\%$  of total fuel gas consumption, i.e. much less than that of the production disturbances. Also here, one of the characteristics of these disturbances is that they are neither very fast, nor very slow. As for the production level disturbances, it usually takes 2-3 hours to decline from the higher consumption level to the lower and vice versa. And likewise, it is typically not considered desirable to try to attenuate the disturbances by manipulation of parameters that would affect the consumption level, e.g. heater coil outlet temperatures or unit feed rates.
5. Finally, a very important characteristic of the total network is its “capacitance”. This means its ability to store quantities of fuel gas, and, hence, its ability to attenuate and smooth out the effect of a change in the balance “total production versus total consumption” on the network pressure. In other words, this capacitance acts as a cushion and as a result, the larger this capacitance is, the longer it will take for a disturbance of a given magnitude to drive the network pressure from one value to another. In our case, this property has been estimated (based on step-test data) to have a value of around 15-18% of the total hourly fuel gas production. In other words, the amount of fuel gas stored in the network is sufficient to supply all consumers for 10 minutes. This figure is probably quite typical for a refinery of this size and complexity. Unfortunately, however, it allows only a very small amount of time for the operators to cope with network production/consumption imbalance. Typically, corrective action is required within 25 minutes for the normal and regularly occurring disturbances and within 6-7 minutes in case of more severe disturbances. Of course, additional fuel gas

storage capacity in the network would help, but this would be at a significant installation cost, which is unlikely to be economically justified.

### **The Fuel Gas Network Control Problem**

The primary control target of the fuel gas network is to keep the pressure of the network in control – nor too high nor too low – without necessarily targeting a specific pre-selected value.

To explain a little bit better this point, the following analogy might be helpful:

The fuel gas network resembles in many respects a liquid storage tank. There are a number of streams feeding the tank (=network) and a number of streams that are fed by the tank. The target for the control of the tank is simply to manipulate some streams in order to keep the level between some specified safe limits (to avoid emptying or overflowing the tank). However, there may be a “soft” target, e.g. 50% level, just in order to have equal capacitance to attenuate disturbances in both directions. In any case, this level control can only be accomplished if the material balance is satisfied, i.e. total input flow equals total output flow, at steady state. In order to achieve this, the system must have available a sufficient number of input or output flows to manipulate as the remaining flows fluctuate freely (=wild flows).

In our case, the operational range for the fuel gas pressure is around  $\pm 8-9\%$  of the average “target” network pressure (which in turn lies around 3 bar<sub>g</sub>) in order to protect a number of burners, compressors and vessels from under-pressure and over-pressure respectively.

As already explained above, it is essential that a number of streams are available to external manipulation in order to control the pressure. In our case the available streams include a propane make-up stream (connected to the overhead system of a depropaniser column) and a vent stream to the flare system. The former is used as an automatic make-up that acts if the pressure drops below a certain low setpoint and the latter acts automatically if the pressure increases above a high setpoint. Two additional handles are available to cope with very big disturbances: a) a large LPG

evaporator for use in case of large deficit, and b) a bigger vent valve for situations of large surplus.

However, neither of these valves/streams is the economically optimum way to control the network. The make-up valve uses an expensive product to supply the fuel gas network, and the off-gas valve directs valuable fuel gas to flare. In both cases there is a significant downgrading of product (propane  $\rightarrow$  fuel gas & fuel gas  $\rightarrow$  flare respectively), so there is a strong incentive to minimise the use of these valves.

To accomplish this task it is necessary to find other streams to manipulate. The most obvious choice is the fuel oil consumption of various dual fired furnaces, which can be varied, at least to some extent, without significant impact on the operation of the process units.

### **Base Case Operation**

Flaring of fuel gas is obviously a very visible sign of suboptimal operation of the network. Hence, the primary target for the base case operation was to minimise the flaring of fuel gas. This was accomplished as follows:

- 1) Choose a high load furnace with a large number of both oil and gas burners (in our case a CDU1 furnace),
- 2) Switch on or off fuel oil and fuel gas (=exchange fuels) to these burners according to the fuel gas network pressure, i.e. in case the pressure of fuel gas falls, then switch on a number of fuel oil burners instead of fuel gas.
- 3) If above manipulation was not enough to stabilise the network pressure, then do the same thing with the burners of other furnaces (usual order: CDU2, VDU1, VBU, VDU2, CCR) for gross correction and continue fine-tuning with the CDU1 furnace.

Although this method could be described as optimal in principle, it was not very successful in actual practice for a number of reasons, e.g.

- As described above, the network behaves like a liquid level, i.e. as an integrator using control terminology. Hence, even a small discrepancy between production and consumption will eventually drive the pressure to one of its limits, i.e. to the setpoint of the propane make-up or the flaring controller;
- The network capacity is not particularly large, so it does not take very long to reach the limits;
- The manipulation of consumption by lighting or switching off burners is step-wise and not continuous;
- Due to the large number of variables involved, is not easy to predict the future trend of the pressure. Hence, the adjustments were typically initiated only after the actual opening of either of the two valves (and usually only after significant make-up or flaring had persisted for a while ...);
- Also, since the responsibility for the optimal control of the fuel system is shared between a number of control room operators located in two control room buildings, it was sometimes difficult to ensure that all furnaces acted in the same direction and prevent that one furnace (accidentally or due to specific needs) acted in the opposite direction, potentially negating the action of the rest;
- Finally, the continuous lighting and switching off of oil and gas burners distressed both the furnaces and operating personnel, which discouraged frequent adjustments to correct the fuel gas balance.

As a result, a quite significant amount of propane make-up and flaring took place in order to keep the fuel gas pressure under control.

The minimisation of the use of these valves was in fact the main incentive for installing a new advanced control and optimisation application for the fuel system.

## **Objectives of New Advanced Control Strategy**

The advanced control and optimisation application for the fuel system has been designed to improve the control of the fuel gas network and achieve the following benefits:

- Minimise flaring
- Minimise propane make-up
- Reduce frequency of switching burners on/off
- Maintain or if possible improve safety and performance of furnaces.

Further requirements for the application include:

- Be capable to incorporate/handle future process changes (e.g. new furnaces),
- Be capable of handling possible future fuel type changes (e.g. introduction of natural gas),
- Be capable of satisfying possible future secondary objectives and targets (e.g. maintain gas turbine load at maximum).

Specifically, the design has the following features:

- The application acts directly on the setpoint of the PIC controller of a base load fuel of each furnace that participate in the scheme;
- The operator has the ability to switch the action from one PIC to the other (from fuel oil to fuel gas and vice versa);
- The system will not allow both fuels to be made available; one fuel should execute the process Coil Outlet Temperature (COT) control task;
- The limits for the setpoint of the PIC's that are available to the controller are amenable to change by the operator that is responsible for the furnace;
- A number of safety features are included in the application to avoid adversely affecting the operation of the furnace (e.g. PIC maximum rate of change, protection of low/high values of the secondary PIC that performs the COT and duty control etc.);
- The existing make-up and flare valves remain as they were, but the controller has been set up to try to avoid reaching their pressure level of actuation.

In the initial design, five furnaces are incorporated in the scheme (CDU1, CDU2, VDU1, VDU2, VBU). The number of furnaces to be included was decided from the following considerations: there should be enough furnaces available to handle the normal disturbances of the network. The more furnaces that are available the better, as less movement will be required on each individual furnace and also bigger disturbances can be confronted. On the other hand, more furnaces means that more operator consoles are involved and more furnace instrumentation is required, as described in the following sections.

Actually, the application is based on the same principle as the manual procedure. In other words, the basic idea is still to alter the fuel gas flow to the dual fuel furnaces according to the fuel gas balance. The difference is that it is now done automatically and continuously.

The application, as well as the actual experience with it, is further described later in this paper. Before entering this field, however, we should describe some of the prerequisites and infrastructure needed to support the application.

### **Improved Regulatory Control**


A series of items had to be addressed before embarking on the implementation of the actual APC application for the fuel system. A couple of the most important issues are discussed below.

#### Fuel Gas Molecular Weight Fluctuations

It had been known for a long time that there was a persistent and significant variation in the COT of the CDU1 furnace, which correlated with the opening of propane make-up valve. Considering that the CDU1 and the propane make-up valve are both part of the “old” network and that this network is usually short of fuel gas, this variation could be attributed to the resulting variation in fuel gas molecular weight and heating value. This hypothesis was also confirmed by spot lab analyses.

For this reason, it was decided to install two on-line density analysers in the main headers that are connected to the respective fuel gas vessels. The interesting results that we observed after installation of these analysers are clearly depicted in Figure 2. As easily seen, the molecular weight of the fuel gas in the two vessels is significantly different, indicating that each vessel collects gas from totally different types of units. Moreover, the molecular weight of the “old” vessel, AR21005, is strongly affected by the operation of the propane make-up valve, PIC2601A.MV, which in its turn affects the COT of CDU1. This demonstrates the necessity for incorporating MW correction in the new Advanced Furnace Control scheme.

Figure 2  
Fuel Gas Molecular Weight Fluctuations




### New Advanced Furnace Control Scheme

At an early stage in the project, it was demonstrated that changing the setpoint of a PIC of a base load fuel in a furnace (even at a very low rate) would adversely affect the process COT control. In all trials that took place, the COT standard deviation increased markedly during a base load fuel PIC ramp. This observation confirmed that

a more elaborate control scheme for the COT of the furnaces should be implemented before activating the fuel gas balancing application.

A typical example of the existing furnace control scheme is displayed in Figure 3. The operator assigns one of the fuels to execute the variable load control (=process COT control) by putting the PIC in CAS (cascade means that the output of the COT controller is used as the set-point of PIC controller) or in PRD (in this mode the output of the COT controller directly moves the output of the PIC controller). The other fuel is used as a base load (by putting the PIC of that fuel in AUT).


Figure 3  
Typical Furnace Control Scheme Existing Prior to the APC Project



With this control scheme, the only way to compensate for process side or fuel side disturbances is via the COT feedback. Such disturbances could be a) process CIT, b) process flow, c) fuel gas conditions (P, T, MW, LHV), d) base load fuel duty.

The new Advanced Furnace Control scheme (Figure 4), which is implemented at the DCS level just as the old one, has the capability to handle these disturbances in a feed-forward fashion. As described in the above-mentioned figure, the scheme considers CIT changes, process flow changes, disturbances in fuel gas conditions (P, MW, LHV) and base load fuel duty. Orifice temperature correction is not applied as it can be considered relatively constant and measurements were not generally available. The LHV is estimated from a correlation between the MW and the LHV. In order to facilitate the very important compensation for fuel oil base load duty, it was necessary to install new fuel oil meters (positive displacement type) to replace the original orifice meters that generally had poor accuracy and reliability.

Figure 4  
Simplified Sketch of New Advanced Furnace Control System



In all furnace where this control scheme has been installed it has achieved a reduction of the COT standard deviation by at least 50%. In the case of CDU1, the reduction was by a factor of more than three (not surprisingly...). Moreover, there is an even more important and marked improvement during the burner on/off operations.

A potential drawback of the new scheme is that it relies on additional instruments, like fuel gas and fuel oil flow meters, where the old scheme considers only temperatures and pressures. Hence, careful attention to reliability is required in connection with the installation and maintenance of these meters, especially the fuel oil flow meters. However, the new scheme incorporates a series of safety switches that protect the COT control from the effects of instrument failure (e.g. rate-of-change, freeze or out-of-limits signal).

### **Control System Architecture**

A simplified sketch of the control system architecture is provided in Figure 5. The process units and utility systems of the Aspropyrgos refinery are controlled by Yokogawa DCS equipment (Centum V, Centum CS, Centum CS 3000). The process units are controlled from two separate control centres located within the process area approximately 600m apart. Regulatory control, including the Advanced Furnace Control schemes described in the previous section, is achieved within Yokogawa Field Control Stations, predominantly Centum V Stations.

The advanced control and optimisation application for the Aspropyrgos Refinery Fuel System is based on Honeywell's Robust Multivariable Predictive Control Technology (RMPCT) also known as the Profit Controller. The fuel system application has 16 manipulated variables, 2 disturbance variables and 35 controlled variables. The application is executed on an NT server together with 12 other Profit Controller applications handling the control and optimisation of other parts of the refinery and a variety of supporting calculations, e.g. inferred properties.

The Honeywell applications running on the NT server communicate with the DCS via a Yokogawa ACG gateway. At present, an alternative solution based on the OPC

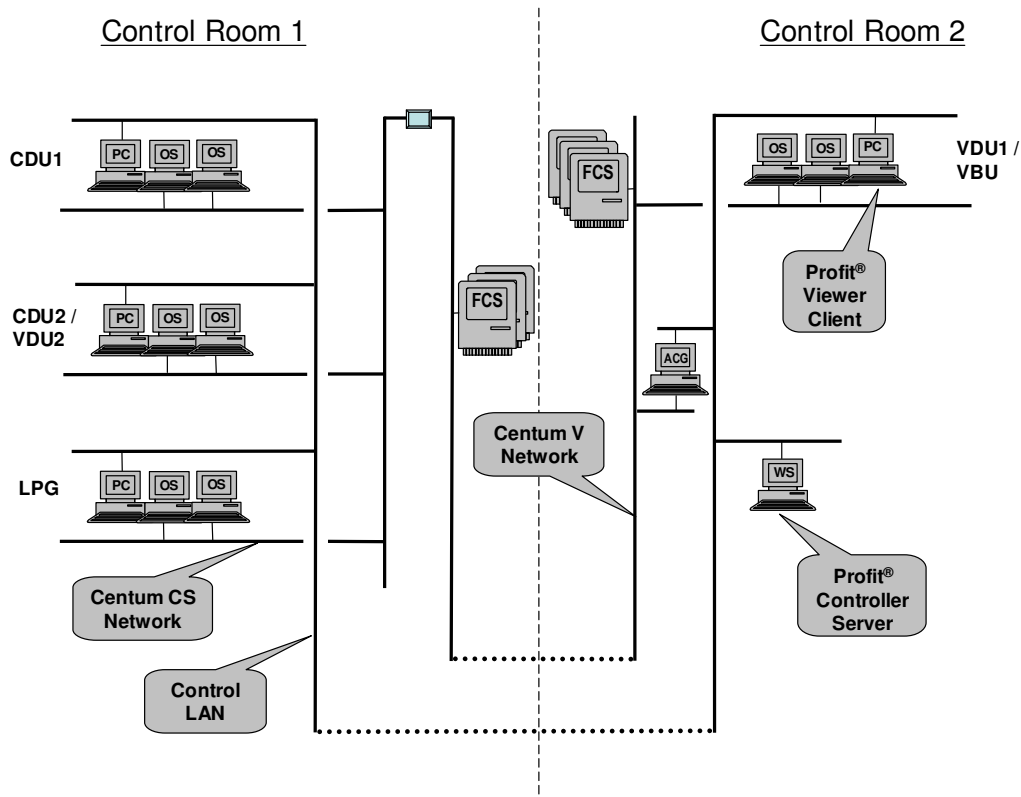
standard is also available, but the ACG connection provides adequate performance and reliability and therefore has been retained.

The fuel system application is currently designed to manipulate furnaces of five of the major process units of the refinery.

The application is shared by up to four control room operators who are working on separate consoles partly located in the “old” control room and partly in the “new” control room. The operators use Honeywell’s Profit Viewer client software to monitor and interact with the application via a dedicated control network, which connects the two control centres via fibre optic cables.

Watchdog functionality is provided at the DCS level to alert each operator in case of any system or communication failure.

Figure 5  
Simplified Sketch of System Architecture



## **Some Early Spin-Off Benefits of the Project**

It is noteworthy that significant benefits were achieved already during the implementation, i.e. prior to that actual commissioning of the fuel system application. These benefits are partly related to the improved regulatory controls described in the previous section, but there were also a couple of other significant spin-off effects.

During the design phase of the project, we configured some pseudo-tag calculations of “total production flow” and “total consumption flow” and made them available on-line to the operators. These pseudo-tags helped the operators a lot understanding the behaviour of the network pressure at any given point in time during normal operation. As explained before, there are two issues that used to preclude fast corrective action on the furnaces to stabilise the fuel gas network: a) the relative small capacitance of the network and b) the multitude of flows involved. To tackle the latter problem, we created two composite pseudo-flow indications: A “total production flow” calculation that incorporates all the measured (and some inferred) production flows, after P, T, MW correction and a “total consumption flow” that incorporates all corrected consumption flows. The use of these pseudo-tags is actually twofold: 1) one version of them is used as feed-forward input in the RMPCT application, 2) the actual values and their difference were configured also in the real time data base of the refinery (an OSI PI system) and made available to the operators. Although the absolute values are not totally accurate (due to the large number of flow meters involved and the various P-T-MW corrections), their fluctuation with time (=trend) provides very useful information helping the operators understand what is happening in the total network, i.e. if it is in deficit or in surplus, at any point in time. In practice this simple tool has proved very valuable to the every-day operation, helping operators to significantly reduce the use of the propane make-up and fuel gas flaring valves, even before the installation of the actual RMPCT application for the fuel system.

Before starting the project, a small study was conducted on the fuel gas network operation in order to establish economics, justification of the project and the basis for the initial design of the application. This study was based on the approximate historical price relations listed in Table-1.

Table 1  
Approximate Costs of Various Fuel Streams

STREAM NAME	VALUE (\$/GCAL)	BASIS
Flare Gas	0	
Fuel Oil	1.0 X	HFO Value
Gas Turbine Feed	1.6 X	Value of Electricity and Steam
Propane	2.0 X	LPG Value
Hydrogen	2.3 X	Steam Reforming of LPG

In other words: If we assign to the price of fuel oil a value equal to X \$/Gcal then the prices of the other streams are


- a) Flaring = 0 \$/Gcal
- b) Gas turbine Feed = 1.6X \$/Gcal.
- c) Propane = 2X \$/Gcal
- d) Hydrogen = 2.3X \$/Kcal

Obviously, using hydrogen-rich streams generated by steam reforming of LPG will never be an economically attractive way to overcome a fuel gas deficit. We shall return later to the observation that the Gas Turbine feed is in fact a better candidate for closing the fuel gas balance than the propane make-up. Another observation from this particular price set is that, from an economics point of view, it is about equally detrimental to use propane make-up and to use flaring to balance the network. This is probably no surprise to those familiar with petroleum economics. However, this observation does to some extent contradict the usual practice the operators had in the past. In order to avoid flaring, they preferred to keep the pressure of the network near the low limit, where the make-up valve was very often activated. Based on the economics, the operator should in fact aim to keep the fuel gas pressure approximately mid distance between the limits where the two valves actuate. Hence, they now try to stay mid distance from the setpoints of the two valves. Having also available the on-line pseudo-tag calculations, they managed to substantially decrease the average opening of the make-up valve (at the expense of a moderate increase in

the opening of flaring valve, which, however, by no means cancels the benefit achieved).

The benefits captured prior to the implementation of the application are illustrated in Figure 6. From October 2001, the combination of the effects described above reduced the average opening of the make-up valve (22.1% → 12.3%), but this was at the expense of slightly increased flaring (2.8% → 5.2%). The value of the net reduction of losses was estimated at around 730K\$/yr. This significant improvement before implementation of the actual application is not at all unusual. Other advanced control projects undertaken in the past in the refinery also confirmed that indirect benefits, which are often achieved prior to the implementation of the actual application, often account for a very significant part (as much as 50%, sometimes more) of the total benefits from an advanced control project. We believe this is a fairly well established fact within process control community.

Figure 6  
Propane Make-Up and Fuel Gas Flaring



## **Application Benefits and Service Factor**

The results of the implementation of the actual application are summarised below, together with a number of observations made during the post-commissioning period.

The on-line application finally went into operation in August 2003. Incidentally, the reasons for the very long elapsed time of the project were not related to the RMPCT application, which was designed and commissioned in a relatively short time frame. The long lead-time was caused by the time required to complete the instrumentation work required to support the application, in particular the Advanced Furnace Controls.

The commissioning of the on-line application further reduced the average opening of the make-up valve (12.3% → 5.5%) and at the same time decreased the flaring (5.2% → 4.3%). The total effect was a further reduction of losses, which has been estimated at 540K\$/yr. Thus, the total loss reduction is around 1270K\$/yr and the payout period of the total project (FO meters + Advanced Furnace Control + RMPCT application) was less than 6 months.

It is also worthwhile to note that only in rare circumstances did we actually have all five furnaces available to the application. In fact, from October 2003 onwards, the typical number was 3, sometimes just 2 furnaces. This is related to the fact that the refinery, and hence the furnaces, is now approaching the end of its 5-year turnaround cycle. Hence, various equipment problems make operators reluctant to accept base load fuel PIC movement in some of the furnaces. The result of this lack of manipulated variables was that the average opening of the make-up valve increased again (5.5% → 7.9%) as well as the flaring valve (4.3% → 4.5%), resulting in a 10% reduction of the benefits that had been achieved by September 2003.

In this context, it is also important to realise the difference between furnaces that are equipped with dual fuel burners compared to furnaces that are equipped with single fuel burners. A furnace is not only characterised by the use of one or two fuels in order to achieve its thermal load, but also by the way that each fuel is distributed over the firebox of the furnace. Specifically, there are two possible cases:

- a) The furnace may have fuel gas going to a number of dedicated gas burners, while the rest of the burners are dedicated oil burners.
- b) Alternatively, each burner utilises both fuels at the same time (dual fuel or combination burners).

In general, furnaces with dual fuel burners have proven much more effective in the fuel system application than furnaces with single fuel burners. The dual fuel burner configuration has the general advantage that it allows a much wider operational window, i.e. distance between the high and the low limits, for each individual PIC. The problem with the single fuel burner configuration is that with this configuration the gas-to-oil ratio of the furnace can only be altered by significantly altering the distribution of heat duty between the burners. However, this also requires altering the distribution of air supply to the burners, which in most cases can only be done locally by the field operator. Hence, in order to prevent possible local air deficiency at some of the burners, which could damage the furnace when CO subsequently burns to CO<sub>2</sub> when mixed with the excess air from other burners, the operational PIC range has to be relatively narrow with the single fuel burner configuration. In the case of dual fuel burners an increase in the base load fuel will be offset by a corresponding decrease in the other fuel (in order to keep the COT constant) keeping the heat duty distribution between the burners more or less constant. Unfortunately, the dual fuel burner configuration is not possible everywhere, e.g. because

- a) Many burners have not been designed with this capability in mind and exhibit poor firing characteristics in such cases.
- b) Dual fuel burners often tend to result in longer flames, which may have undesirable affects on the furnace itself (e.g. skin temperatures).

In any case, irrespective of single or dual firing, is necessary to establish some safe operating window (high – low limits) for each PIC that participates in the fuel gas balancing application. These limits should consider not only the oxygen issue that has been described above, but also the operating limitations of each burner type. This is a task that is done best in the field by trial and error and subsequent visual inspection of the flame appearance and monitoring of skin temperature indications of each furnace.

The application runs smoothly and consequently the service factor is very high (>99%). This is also due to the decentralised structure of the application, which means that a problem in one of the furnaces that participate in the application will only result in the operator temporarily switching off that particular furnace, not the entire application. This will not really affect the operation of the rest of the application, except for its ability to reject disturbances, as there are now fewer variables available.

Furthermore, the application has demonstrated its ability to provide additional protection of the individual furnaces that participate in the scheme. Specifically, it uses the base load fuel PIC of each furnace not only to balance the fuel gas network, but also to prevent the corresponding PIC that performs the COT control from exceeding its high or low limits. During the evaluation period, it has been estimated that for every furnace that participated in the scheme, the number of PIC excursions has been reduced by about 85% and the magnitude of the excursions by about 65%.

Moreover, it has been observed that the need to light or switch off burners has been substantially reduced (in the order of 50%) after the implementation of the application. Together with the improved protection of PIC limits this contributes to a smooth and safe operation of the furnace. However, with the currently available furnaces and the control range of each furnace, burners on/off operations are still required.

The application writes directly to the setpoints of the PIC's of the base load fuels. This fits well into the advanced regulatory control structure described in a previous section and enables the application to very effectively manage fuel pressure or valve position constraints as mentioned above. One drawback, however, is that the effect of changing a base load PIC setpoint on the network pressure, i.e. the steady-state gain of the dynamic model between base load fuel pressure and the time derivative of the network pressure, varies significantly over time, depending on

- the number of base load fuel burners currently in service
- the type and condition of the burners,

- the possible throttling of local fuel valves of the burners, etc.

The most significant gain changes are probably caused by the variation in the number of oil and gas burners in service. In principle, this could be dealt with by requiring the operators to enter information about the number of burners in service into the DCS. However, this would increase the workload on the console operators and does not address the other potential reasons for variations in the process gain. A better approach, which was successfully applied within this application, is to make use of the heat release curves of the burners, which allows the process gain to be estimated on-line. This substantially improves the model prediction accuracy, and hence the performance of the application.

During the design phase of the project, the need for setting targets and priorities for the oil-to-gas ratios of the various furnaces participating in the scheme was discussed. This would effectively include additional optimisation targets to ensure that the optimum solution is always fully specified with no additional degrees of freedom. However, considering the frequency and the magnitude of the disturbances compared to the expected operating range of the available furnaces, we decided that such functionality would not be necessary. In practice, this turned out to be correct. The minimum-move solution inherent to the RMPCT algorithm does a good job stabilising the network pressure at minimum control effort in the presence of significant and frequent disturbances. Attempting to enforce a desired firing pattern using the very limited (often zero) available degrees of freedom would just introduce additional control movement without adding any value.

### **Future Enhancements**

As discussed in the previous section, it is still necessary to switch on and off burners on a regular basis in order to cope with larger disturbances in the fuel gas production level. In fact, the remaining losses of propane and fuel oil due to opening of the make-up and/or flaring valves are almost exclusively caused by lack of control flexibility leading to an infeasible problem for the fuel system application.

The only way to improve this situation is to make more control handles available to increase the flexibility of the application in terms of control range.

As clearly observed, the more furnaces participate in the application, the better is the disturbance rejection capability of the scheme and the less is the required movement of each furnace. Having additional furnaces available would not only almost eliminate finally the use of the two valves (there is still a 250k\$/yr loss to be captured), but at the same time it would

- a) provide an important back-up in case some furnaces are not available to the application (like at present), and
- b) further reduce the need for burner on/off operations.

The implementation of the infrastructure required to include a TAME unit furnace is already in progress and probably some CCR or Isomerisation unit furnaces will be added in the future.

The structure of the application is quite versatile, so possible future changes in its structure will be straightforward to implement. Possible candidates that have been discussed include:

- a) introduction of natural gas as an additional fuel for the gas turbines,
- b) use the LPG vaporiser in addition to or in place of the propane make-up valve
- c) make the Gas Turbine feed available for manipulation by the fuel system application

Whereas the purpose of the addition of more furnaces to the scheme is to increase the flexibility in terms of *control range*, i.e. enable the application to do what it is currently doing, and what the operators used to do manually, in a more effective way, the three potential handles listed above are examples of variables that could be added to increase the *optimisation flexibility*, i.e. the ability to respond in different ways under different conditions.

As an example, the simple price scenario described in a previous section implies that the load of the Gas Turbines should be reduced in case of severe fuel gas deficits where the propane make-up valve remains open even at maximum oil firing. At the moment, it is not possible to manipulate the Gas Turbines from the fuel system application as they are not connected to the DCS, but there may be sufficient incentive to establish such a connection.

Another example would be the use of the LPG evaporator. This could have two advantages:

- a) Less disturbance to the process units
- b) Flexibility to support the fuel gas network with butane instead of propane if more economical

The latter could be a significant enhancement of the optimisation flexibility. Most of the time, fuel oil will probably remain the most attractive marginal fuel supply, but there will be situations, where butane is a more attractive supplementary fuel than propane or gas turbine feed. Probably, there will even be periods where butane is a cheaper marginal fuel than fuel oil.

In more general terms, there are in fact a number of further candidates that could be manipulated by the fuel system application, e.g.

- Fuel gas producers
- Gas fired furnaces

As described in the introduction, there is a large number of fuel gas producers located throughout the refinery. Moreover, the fuel gas production of each unit depends on a variety of operating parameters, e.g. operating severities. So far, in order to avoid having to extend the scope of the analysis to cover the operation of the process units of the refinery, we have conveniently distinguished between what you could call the “natural” fuel gas production and “artificial” fuel gas production. The “natural” fuel gas production is the quantity of fuel gas produced at the optimum set of refinery

operating parameters without taking into account the constraints of the fuel gas balance (e.g. the “natural” flow of FCC off-gas is a direct consequence of the FCC feed rate and severity). The “artificial” fuel gas production is additional to the “natural” fuel gas production and uses handles that can be manipulated without altering the key operating parameters.

Examples of “artificial” fuel gas production handles include streams that play an ancillary role in the operation of the unit like purge streams from HDS units, the automatic makeup to the fuel gas system from a depropaniser overhead system as described above, the LPG evaporator and to some extent the pressures and temperatures of flash drums and absorbers that route off gas streams to the fuel gas network.

So far, we have implicitly made the assumption that the “natural” fuel gas production should not be altered by the fuel system application. This assumption significantly simplifies both the analysis and the application by limiting the candidates for manipulated variables to only the “artificial” fuel gas production handles. However, one could argue that the distinction is somewhat arbitrary and that there is actually no guarantee that the above assumption is correct.

Analogous to the above discussion about “natural” versus “artificial” fuel gas *production*, we have also implicitly made the assumption that the key operating parameters that define what we could call the “natural” fuel *consumption* of the process units should not be altered by the fuel system application. Again, one could probably argue against the validity of this assumption, e.g. a situation where refinery fuel is at zero cost due to a prolonged period of fuel gas flaring could very well lead to a different optimum also for the key operating parameters of the process units.

These considerations could lead to a very large and extremely complex optimisation problem. Hence, we believe that our implicit assumptions, which effectively decouple the fuel system optimisation problem from the process unit optimisation problem, are necessary and reasonable as a first step.

On the other hand, the structure of the fuel system application is totally open, and there is no fundamental problem with solving the plant-wide fuel system optimisation problem together with the multi-unit optimisation problem of a cluster of process units.

## **Summary and Conclusions**

A plant wide application for the control and optimisation of the Aspropyrgos fuel system has been successfully implemented.

The overall result of the project is a reduction of the fuel costs of in the order of \$ 1.2 million per year. Around 60% of the benefits can be attributed to improvements of infrastructure, including regulatory controls, and improved operating strategies. The remaining 40% are direct benefits of the control and optimisation performed by the application.

Operator acceptance has been excellent and service factors very high (>99%). However, furnaces equipped with combination burners are much more effective in the scheme than furnaces with dedicated oil and gas burners.

The ability to update the models gains on-line and the inherent characteristics of the minimum-move algorithm of the RMPCT controller have contributed to the stability and performance of the application.

The remaining inefficiencies in the fuel system are primarily due to lack of control range flexibility due to the limited number of furnaces participating in the scheme and due to the limited flexibility of furnaces that are equipped with single fuel burners. This will be addressed in the near future by adding more furnaces to the application.

A further advantage of using a multivariable predictive control package as software platform for this type of application is that this leads to a very flexible solution, which can easily be amended and expanded. One of the future enhancements currently under consideration is to increase also the optimisation flexibility by adding other fuel

supplies and consumers to the application, e.g. natural gas, LPG evaporator and gas turbines.

As a longer-term vision, the application could even be combined with future multi-unit optimisers for the process units of the refinery.