Unpicking Oversimplification: Relief and Flare Studies for Complex Plant using Dynamic Process Modelling

Dr Ben Firth, Senior Simulation Engineer, Flexible Process Consultants Ltd

"First-cut" and "bounding" case scenarios have their place early on in safety assessment for hazardous sites. They act to provide a screening tool to help understand where the real hazards may lie with a plant and provide a clear basis for subsequent safety analysis, studies and cases. However, application of simplistic approaches to assessment of process systems in fault conditions can lead to grossly over-engineered systems and worst of all, misunderstanding and dilution of attention and resources from unrevealed hazards and necessary risk reductions.

An occasion where this can happen is in complex, interconnected plant with site-wide relief and flare systems. Analysis of the various initiating events and combinations can be very complicated. Traditional approaches use a simplistic set of assumptions to ensure conservatism. The shortfall here is that these approaches can easily miss the interactive nature of emergency events, some of which can be counter-intuitive – those that are worse than expected, and others that are far less severe.

This paper presents the findings and selected safety cases from an assessment of two flare systems at a refinery using dynamic process modelling. The study included detailed modelling and rating of the Crude Overheads, Naphtha, Kerosene, Diesel, Isomerisation, Amine, Sour Water, and Fuel Gas units. The paper covers the process followed, which began with PFDs, a complete survey of plant equipment, building a steady state model to match plant conditions, and generation of a dynamic model running a variety of plant conditions from start-up to max throughput.

The modelling provided an understanding of both local and common mode failure cases (such as equipment failure, operator error, or utilities failure), and the resultant relief events. By modelling all the processes together, the interactions, knock-on effects, and cumulative flare loads could be studied in depth. This combination of rigorous calculation and narrative style allowed for a deeper understanding and clearer communication between engineers. This also brought to light specific vulnerabilities, with recommendations made to avoid it them.

The ultimate result of the study was that, by implementing a set of recommendations that had been tested in the model, two legacy flare systems could be kept in service, where traditional calculation methods would either have declared the headers undersized, or not been able to produce the appropriate actions to keep the flare in safe operation. This represented savings measured in the millions of pounds while advancing the design and operational knowledge of the Operator. The models provided substantiation to HAZOP and LOPA studies, as well as debottlenecking processes to achieve higher throughput.

We believe these tools represent a significant step forward for safety engineering analysis, where the increasing complexity of process plants means that traditional hazard assessment methods struggle to identify and properly quantify the risks, especially for multiple relief events. The reproducible nature of modelling removes the subjectivity inherent to HAZOP scenarios, and allows for future review to fully understand the reasoning and engineering behind past safety decisions.

Key words: Dynamic Simulation, Process Modelling, Relief Study, Flare Study, Refinery

Introduction

In the absence of inherently safe processes, pressure relief is an essential protection measure used across the process industries, and presumably one all process safety professionals are familiar with. Standards are available from multiple organisations for the development of relief cases, and sizing of devices to handle them. This exercise has become a fundamental part of most process design, undertaken at FEED, with the design engineer responsible for specifying the sizing cases and relief devices, using the design basis for the associated pressure equipment. To an extent, relief calculations are more a matter of following the steps prescribed in the standards, than undertaking an in-depth study of the hazards associated with the equipment, and this is probably for the best – these "first-cut" or "bounding" scenarios are over-simplified, but are likely to give the worst cases, and thus help make sure the relief devices are sized conservatively.

Along with the widespread use of relief devices, flaring has become ubiquitous across the oil and gas industry. The use of common headers to service large sections of a refinery is accepted practice, although sizing these systems is not as straightforward as with the individual devices. The standards become more subjective for these calculations, relying on experience, professional judgement, and rules of thumb to determine the sizing load. Multiple relief events with a common cause become the basis for sizing, with potentially complex unit interactions clouding the picture – standard calculation methods simply aren't equipped to deal with the interconnections of modern process plants. Because of the range of design pressures of the equipment sharing a single relief header, incorrect sizing could result in relief devices failing to operate, despite their own conservative basis. This risk will generally lead the designer to significantly overdesign the system to account for the extensive uncertainties.

Undertaking a relief study or device review at an existing facility is often essentially repeating the design process – though likely with the as-built equipment and process data – following the standards, and ensuring that all relief devices are appropriately sized, and their datasheets are up to date. This approach, again, is likely to be conservative, and perfectly acceptable in many circumstances.

Undertaking a flare study of an existing facility, however, is a much more complex situation. Unlike at the design stage, where over-conservatism may result in an acceptable overdesign, excessively conservative assumptions pose a serious business threat to operators of legacy systems – lengthy shutdowns and costly construction projects could easily result from the uncertainties of flare system loads. Worse still, making assumptions under pressure to avoid large capital expenditure could lead to the understanding of actual risks being lost, and the plant continuing to operate with an unclassified hazard.

Current Practices and Guidance

API is widely recognised as the leading standards body in pressure relief, with API 521 being most commonly used for process plants. Indeed, other standards bodies have adopted or adapted API standards (e.g. ISO 28300 is identical to API 2000). As mentioned above, these standards tend to take a prescriptive approach to relief scenarios (standardised fire calculations, maximum flow with blocked outlets, loss of cooling, etc.), and are familiar territory for most process safety professionals. However, API 521 states that "conventional methods for calculating relief loads are generally conservative and can lead to overly sized relief- and flare-system designs."

API 521 Section 5.3 discusses the determination of the design basis for disposal systems. The main sizing scenarios are assumed to be common cause, multiple relief events, and the design engineer is told to determine which relief events happen concurrently, based on the individual unit sizing cases. These are then combined, with credit taken for pressure control and blowdown valves reducing or eliminating the demand on the relief valves. The standard then goes on to state that this is still overly conservative – and may be refined through "Dynamic system load modelling" and "Load reduction credits", such as safety instrumented systems and operator intervention. The guidance for each of these is sparse – and for good reason; they are complex and difficult subjects, requiring expertise across a range of disciplines, and cannot be broken down into step by step instructions.

This required understanding of such complex scenarios has led organisations to develop rules of thumb for sizing, such as applying the full sizing load of a single relief device, accompanied by 50% of the loads from all other devices, and checking that the backpressure on the valve under consideration remains acceptable.

Dynamic Simulation

Dynamic simulation is real-time version of traditional process modelling, and allows time-dependent events to be understood in more depth. While largely originally used for unit-level design and operator training simulators, it is finding increasing use in larger scale site-wide models, for analysing controller responses and effects of disturbances, including pressure relief, flare studies, and HIPS design (Nezami, 2008; Sirven et al, 2011; Nair et al, 2010). It has also been used to reproduce accidents (Manca & Brambilla, 2012). Tools such as Dynsim, HYSYS, UniSim, and GProms are commercially available, each with particular strengths.

API 521 has two sections covering the use of dynamic simulation in relief and flare studies. Section 4.3.3 discusses the use for calculating required relief capacities of individual devices, and states that it "provides an alternative method to better define the relief load and improves the understanding of what happens during relief." The scenarios are developed in largely the same way, initially with no credit taken for control responses. The remaining guidance consists of a few warnings to avoid over-reliance on the simulation without full understanding of the system.

Section 5.3.4.2 comprises a paragraph on the use of "dynamic system load modelling" for flare system sizing. The use of such tools is desirable because the flare loadings from individual units during a relief event will not peak simultaneously, and so testing the overall sequence of events will lead to a more realistic sizing case.

The Study

Flex Process carried out a comprehensive relief and flare study for the Milford Haven Refinery, using dynamic simulation as a key tool. The refinery was originally built in 1973, with its capacity increasing significantly over the years, and plans were in the works to expand by a further 10%. However, the flare systems had not been rebuilt in conjunction with the successive expansions, and the HSE expressed concern over their ability to handle the increased design capacity.

The study covered two of the site's three flare headers, and incorporated most of the refinery units, including:

- Crude Tower Overheads
- LPG Train (Depropaniser, Debutaniser)
- Naphtha Reformer
- Isomerisation Unit
- Kerosene Hydrotreater
- Diesel Hydrotreater
- Ancillary systems (Fuel Gas, Sour Water, Amine)

The study was undertaken as both a relief and a flare study; the individual pressure systems and their devices were subjected to all identified relief scenarios, while the effect of all multiple unit events on the flare systems were examined in detail. This combined approach allowed the impact of heat integration, unit interconnections, and common mode failures to be simulated, without losing sight of the underlying causes and behaviour at the unit level. The study had two overall goals:

1. Ensure each relief device was correctly designed
2. Ensure the flare system was adequate for the plant's full capacity

Modelling Complex Processes

Due to the lack of strong guidance from API on how to undertake a dynamic modelling study, it is necessary to define best practice for such an undertaking. By ensuring a rigorous and auditable process is followed, a reliable model, with sufficiently conservative assumptions, can be built to reflect actual plant behaviour, and perform relief calculations at a level of detail that is simply not available with the traditional approaches. The general workflow of a simulation project is illustrated in Figure 1.

Figure 1. Simulation project workflow

Data Gathering

The level of data required for a dynamic model is significantly higher than for a conventional study. While equipment datasheets and isometrics are used to a similar extent, more in-depth equipment data is incorporated, including column tray design, heat exchanger layouts, valve trim characteristics, and actuator speeds.

Similarly, all available process data should be made available to the simulation engineers, as the goal is to reproduce actual plant performance in the model. For both steady state and dynamic simulations, this typically takes the form of a single point "snapshot" of the plant, where at an agreed upon point of steady operation, all the available operating data is downloaded. There are several vulnerabilities with this approach – unusual controller settings, in place to deal with a short-term issue with a unit; fouling, reducing throughput and heat transfer; erroneous instrumentation; catalyst at the beginning or end of its regeneration cycle; unusual ambient conditions; transient conditions not identified due to the complexity of the plant; plant settings unique to the operator at the time – all of which can create an unrepresentative snapshot of the overall process data.

To avoid these issues, process data should also be looked at in terms of distribution – how often does the data fall within certain bounds? This approach will also allow for the development of worst case plant conditions, which can be used to ensure conservatism is maintained while remaining within the realistic bounds of the plant's operational history. Because the data is examined in terms of distribution, there is no limit to the amount that can be examined. In this study, hourly data points over multiple years were used. The distribution of data also reveals differences between how the engineers may expect the plant to operate, and where it actually does. As can be seen from Figure 2, a single point on a plant can be incredibly variable, and being able to quantify the frequency that operations fall within the snapshot range is invaluable to creating a rational model. This approach thus allows the engineer to prioritise matching the well-controlled points, while achieving realistic results at the more variable.

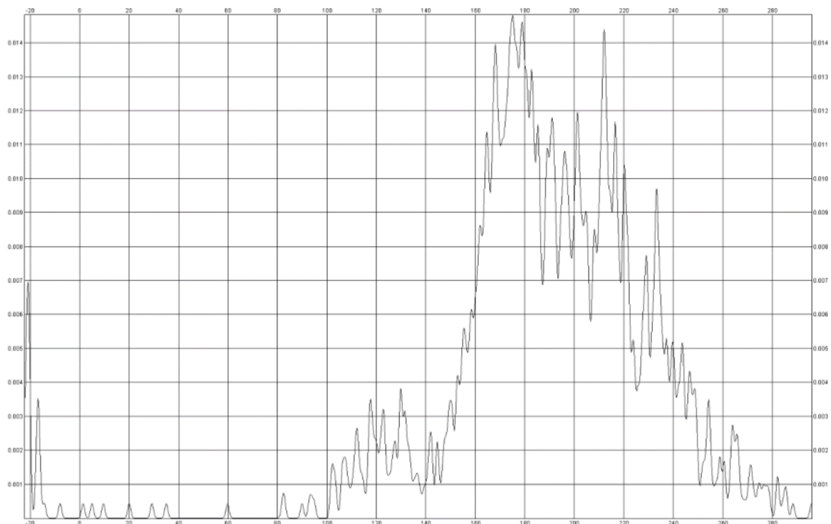

Figure 2. Probability distribution of reflux rate from a column (m^3/hr)

Model Building

Building a dynamic model to match plant data follows a general procedure:

1. Generate a project PFD at the required level of detail

2. Build a steady state model to develop the chemistry, heat transfer correlations, and ensure overall mass and heat balance
3. Build a dynamic model to match the steady state model
4. Generate multiple initial conditions for the dynamic model to match the various operating states of the plant
5. Reviewing the model conditions with plant engineers

The use of a steady state model is mentioned in API 521 Section 4.3.3. Presumably this is for design cases. In legacy plant, although it would be possible to proceed without the steady state model, it is best to generate one first, in order to reconcile plant data, and to gain a deeper understanding of complex process units, such as reactors. This work is much more efficient with a steady state model, where convergence is relatively instantaneous.

Data Input

The models' level of detail is defined by the project PFD. This tends to be more detailed than a simple process diagram, but can omit much of the instrumentation and certain lines from the P&IDs. For example, there is no need for the level gauge piping to be built into a model, but bypasses and start-up lines may be part of relief cases, and so are generally included. By defining this level of detail at the beginning of the project, continuous iteration of the models as small details are added can be avoided. It also aids in collating the correct equipment data.

A key process in developing a reliable model is the mark-up of all relevant data. Pipe lengths, equipment sizes, heat transfer performance, etc. should all be clearly marked for each piece of equipment. This can then be input into the model, prior to another engineer checking all the inputs against the marked-up data. This fills a similar role to having a competent engineer check calculation sheets – while not guaranteed to result in perfection, it provides a vital check on input data, and prevents obvious errors from making their way into the final results.

Data checking must go beyond this to ensure that the model can be relied on for a critical safety assessment. Each pump, control valve, and relief device should be visually checked to ensure that there are no discrepancies between the data in the plant's equipment library, and the actual installed equipment on-site.

The final data input is the controller settings. Matching the control schemes and tunings used by the site allow simulations to predict the complex effects of controllers operating throughout an emergency scenario. While many times controllers will reduce the severity of an event, there are instances where a controller will in fact worsen an event, or cause an unforeseen disruption on another part of the plant. Additionally, for flare studies, it is unreasonable to take no credit for controllers – much of the flare load during an event may come from pressure control valves opening, including on units that don't overpressure.

Interconnections

While in the past the computational intensity of large scale, detailed dynamic simulations would have been prohibitive (c.f. Spooner, 1994), modern high power multi-core processors have opened the door to modelling complex plants. A single PC can be powerful enough to run an entire refinery simulation faster than real-time. Because of this, models can be linked by the utilities and networks that connect the real plants: fuel gas, flare header, heat integration, etc. By implementing these interconnections within the models, the interactions and knock-on effects of process disruptions can be demonstrated, often with unexpected results. This is important because knock-on events and unit interactions are not necessarily obvious as having a common cause in the way that utility failures are. This may result in a relief event being missed from a flare scenario (which wasn't an issue for the individual unit's device sizing, where the sizing case was almost certainly worse), and the backpressure from concurrent relief streams not being considered.

Model Rating

Both the priority and the main challenge with a process simulation is that it matches the process it is modelling. Temperatures, pressures, product compositions, and valve positions should all be within the normal operating range of the historical data. Once the model is built and operating, the simulation engineer's primary job is to compare the rating data with the model, identify the discrepancies, and understand and resolve them.

Matching rating data must be done in collaboration with process plant engineers. They have the experience and knowledge of their own processes to help identify and check the probable causes of data discrepancies. A simple example is one that occurred multiple times during the project: a line was seeing more flow than would be expected from the position of the valve. The answer was quickly verified by the site – the control valve was undersized, and the bypass valve was open, and had been for some time. Similarly, pumps were identified that had a different impeller installed than was recorded. This level of detail is only possible when reproducing operating conditions in a model; any conventional calculation would have to assume the pumps' library data were correct, or that the bypass valves were closed.

Initial Conditions

Once a model is rated, it can be altered to create a range of initial conditions, including conditions that the plant has never undergone, such as running at maximum design capacity. These initial conditions represent the various combinations of operating states the plant may be in, such as minimum turndown, low or high liquid levels, different stages of start-up or shut-down. These initial conditions provide the basis for the model to be run, either as sensitivity studies for repeating cases, or as starting points for unusual events, such as operator error during start-up.

Scenario Development

Plant fault scenarios are developed using the following steps:

1. Identify initiating event
2. Identify relevant initial conditions
3. Identify and select a latent failure (if applicable)
4. Apply case to all relevant initial conditions

Initiating Events

Working within the guidance of cases prescribed by API 521, the fault conditions causing an event can be developed to reflect the scenarios to be examined. The initiating events used for the study are listed below in Table 1, along with examples.

Table 1. Initiating event classifications, with examples

Initiating Event	Example
External Fire	Pool Fire around multiple vessels in a single fire zone
Power Failure	Loss of power at motor, distribution board, or site level supply
Process Control Failure	Loss of level control, fails to open
Trips	Downstream unit trips
Air Failure	Loss of instrument air – either local or global
Mechanical Failure	Heat exchanger tube rupture
Heat Input	Exothermic reaction heat
Utility Failure	Loss of steam
Mass Input	Reaction product, or connected services
Abnormal Event	On upstream or downstream systems
High-Low Pressure	Valve opens leading to a liquid to gas blow-by case
Mass Out	Increased downstream suction
Heat Out	Planned depressurisation
Human Factors	Operator error on start-up

Initiating events were applied to every piece of equipment in the models – each pump, valve, controller, and electricity breaker was failed, and the scenario allowed to play out in a real-time simulation.

Latent Failures

In order to introduce an adequate level of conservatism to the calculations, it is necessary to consider latent failures. These may include a temporary bypass being open, check valve failure, control valve seizing, or upstream cooling not being online or not at full capacity.

The challenge when introducing latent failures is finding the correct balance between worsening the case and avoiding invoking double jeopardy – multiple unrelated instantaneous failures. A latent failure should be chosen such that the plant can continue to operate for a prolonged period of time. However, it is also unreasonable to introduce all possible latent failures to a unit simultaneously. For this reason, the study used one latent failure for each scenario wherever possible, unless more were justifiable.

Credible Cases

The goal of a relief and flare study such as this is to improve the safety of the plant without introducing so much conservatism that the justification for modifications is lost amidst the assumptions. By methodically building a scenario from a demonstrable starting point, with a reasonable initiating event and a justifiable latent failure, the study will produce a quantifiable result, which can be reviewed and interrogated, and be robust enough to stand the scrutiny of experienced engineers. Part of this balance is ensuring that the correct initial conditions have been used, for example, scenarios worsened by valves being in their minimum turndown position must be assessed using a minimum turndown initial condition.

Site-Wide Scenarios

For multiple relief events and common mode failures, the approach must be slightly different from single unit scenarios. Where controlled responses can reduce the demand on relief devices (e.g. a pressure controller will take some or all of the load from a relief valve), having all valves across an entire refinery under manual control is neither credible, nor conservative – the movement of heat and material in a realistic fashion is the point of such a study, and artificially restricting this will result in unrealistic scenarios, which are just as likely to reduce the severity of an event as increase it.

Latent failures for a site-wide event must therefore be approached in a more nuanced way. The failure of valves to operate must now be considered a latent failure, as opposed to being intrinsic to the scenario. For example, a column's pressure controller will now be allowed to respond and send incondensables to flare, and a level controller will be left to send liquid to the next unit, but the controller on the reboiler will be left in manual mode, preventing heat input from being reduced. The number of latent failures was set at one per major operating unit, unless more were justifiable.

Running the Models

With the models built and rated, the initial conditions generated, and the scenarios and latent failures defined, the models must be run by an experienced simulation engineer. Understanding both the software and the process is an indispensable check on the quality of the results, as well as vital to explaining the underlying phenomena that drive a scenario from initiating event to relief load.

Implementing Scenarios

Scenarios are implemented via scripted events, which are auditable and automated, making the initiating event, latent failure, and all other assumptions repeatable and recordable. For individual unit events, the script also changes all controllers to manual mode. As operator action is a mitigating factor which should not be considered, the simulation engineer runs the model for at least 30 minutes of simulation time, to ensure that all events are captured within a reasonable window of inaction, unless the model has reached a steady state.

Example 1: VNS Power Failure

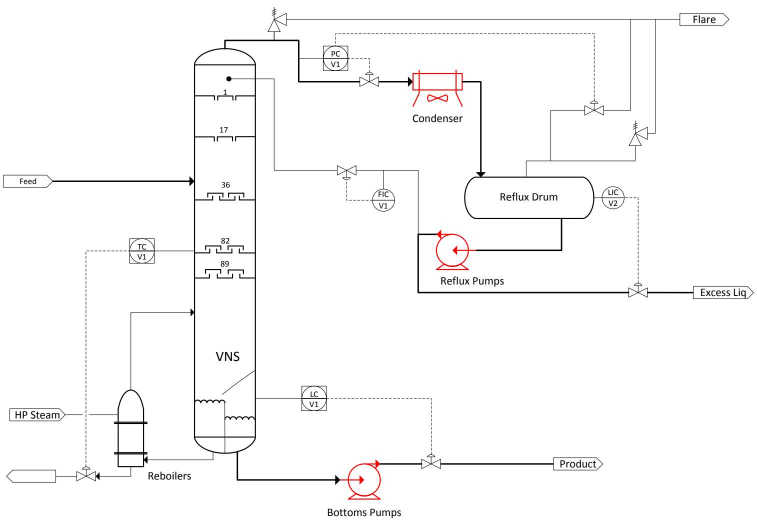

Figure 3. Simplified PFD of VNS

Scenario: Loss of power to Virgin Naphtha Splitter (VNS). This scenario was of specific interest to the study, because the site had experienced a power failure event which resulted in liquid reaching the flare. The VNS was suspected to be the culprit, but it had not been proven. On failure of a specific board, the VNS would lose its bottoms and reflux pumps and the condenser fans, as shown in Figure 3, resulting in a build-up of liquid in the column and reflux drum, as the upstream debutaniser emptied under pressure and the reboiler continued to operate. Once the reflux drum overfilled, the condenser would flood and the column lose the sink for the reboiler heat. Following API 521, this would essentially be considered a blocked outlet, and the main column relief valve would be sized to handle the event. This had been identified and carried out in previous studies. However, the effect of static pressure from the top of the column to the lower elevation reflux drum would mean that there would be substantial liquid relief via the reflux drum pressure controller and relief valve. The column would also relieve once the condenser flooded, resulting in two-phase flow in the sub-header. This interaction could easily be missed following conventional methods, where sizing is the priority over accurate event diagnosis.

Recording Results

In order to capture the results of a simulation in a meaningful manner, and maximise its use beyond passing or failing equipment, the results of each scenario must be explained both as a sequence of events and as ultimate numbers. By producing both static answers (e.g. peak relief rate) and time-based graphs, accompanied with a narrative describing the phenomena as they occurred in the model, the simulation engineer can generate a scenario report that is both comprehensive and engaging. The narrative style also provides the opportunity to identify other contributing factors which may not have been obvious beforehand, and develop sensitivity studies for deeper analysis.

At the end of a simulation run, the resulting state should be saved, along with the data historian, so that the data logged throughout the scenario can be easily accessed. Additionally, the simulation engineer can re-simulate the case as many times as needed until they have gained a clear understanding of the sequence and underlying phenomena, so that they can explain both the process and the thermodynamic contributions to the end result.

Sensitivity Studies

Sensitivity studies involve searching for the worst credible scenario, by altering the initial conditions, controller modes, or latent failures of a safety case. Selection of different initial conditions may cause a case to become significantly worse, or change the character of the event completely. This happened in a case illustrated by Example 2.

Another aspect of sensitivity studies is to introduce onerous assumptions, to ensure that conservatism is completely met. While there may be scenarios where credit could be taken for outlets to remain in their minimum normal position (see e.g. API 521, Section 4.4.8.2), where the relief capacity margin is narrow, or there is any uncertainty over the scenario, these outlets may be assumed closed to direct all of the required relieving rate to pass through the relief device.

Example 2: High Pressure Blow-by

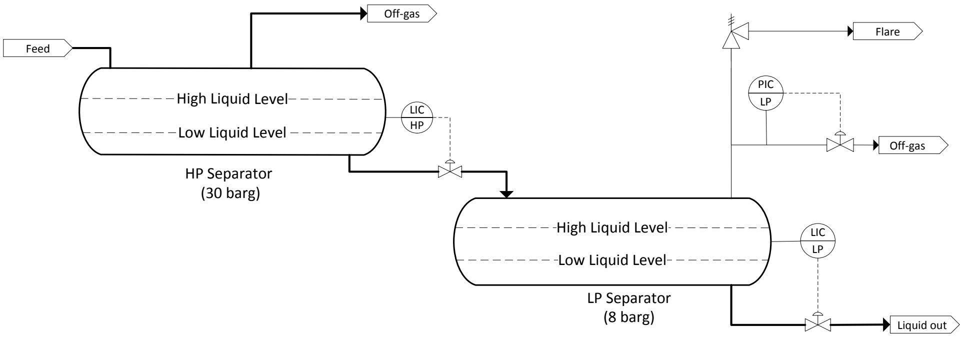

Figure 4. Simplified PFD of high pressure and low pressure separator vessels

Scenario: Loss of level control LIC-HP. On loss of level control, the HP Separator will run dry and vapour with blow through into the LP Separator, causing a vapour relief case. However, if the vessels both have high liquid levels, the LP Separator is full of liquid before breakthrough, leading to an immediate two-phase relief load, which is significantly larger than the simple high-pressure vapour breakthrough case.

Unexpected Events

Due to the vast complexity of many process plants, it is unreasonable to expect an engineer to manually identify every consequence from every event they examine. This is the very reason for the prescriptive nature of relief sizing standards – ensure a conservative sizing case, and leave everything else to the HAZOP. But even a HAZOP may fail to identify an unintuitive sequence of events, and no one is performing complex thermodynamic calculations in the meeting room to inform the team of the exact process conditions on every line of every P&ID simultaneously. As can be seen in Example 3, non-relief events can be picked up just as easily as relief events.

Example 3: Kerosene product line over-temperature

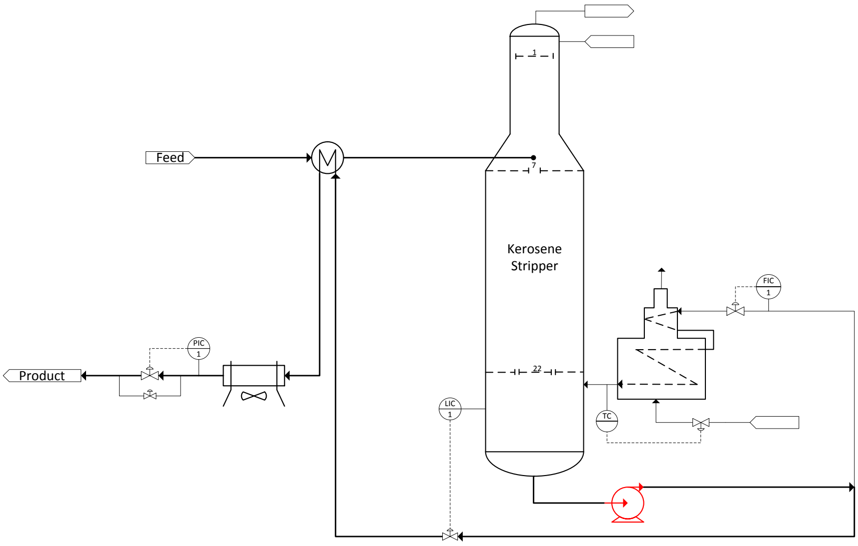

Figure 5. Simplified PFD of Kerosene Stripper

Scenario: Loss of power to the stripper bottoms pumps during maximum capacity operation. Additionally, the product line bypass is permanently open due to an undersized valve. Following the loss of bottoms pumps, the column sump begins fill with liquid, until eventually this reaches the reboiler return. This liquid then flows back through the reboiler line, where the lack of flow has fully opened the reboiler flow controller FIC1 has fully opened the valve, and the temperature controller TC1 has fully opened the fuel gas to the fired reboiler. This overheated fluid then goes to the Product/Feed Interchangers, well above the design temperature of the pipework.

Only one of the trips on the unit would have been called upon: low flow at the reboiler outlet. However, even taking this trip into account, only the fuel gas is shut off, and the pipework design temperature is still significantly exceeded due to slow flow through a hot furnace.

The event can be avoided by ensuring that the product valve is sufficiently sized to leave the bypass closed, and a non-return valve fitted to the reboiler circuit. This then affords two layers of protection to the system: the non-return valve and the pressure controller PIC1, which would shut the product valve in this scenario due to loss of pumps. The residual risk can then be categorised via risk assessment or LOPA.

Reporting

As with any engineering exercise, proper recording and reporting of outputs is essential to a successful study. In order to effectively communicate the complexity of events, as well as the engineering parameters required to size the relief devices, a reporting template was developed. The recommendations were also tested within this framework, allowing for both relief study and recommendations to be integrated.

Condensing Information

The challenge with reporting simulation results is the opposite of most engineering studies: too much data. Every section of every line, every column tray, vessel, valve, and pump can provide a data point for each model time-step, resulting in an overwhelming amount of data. Reducing these to the key parameters is the role of the simulation engineer, to allow others to understand the scenario and follow the narrative without having seen the simulation running. Additionally, the key engineering outputs must be recorded, so that they can be added to the relevant equipment datasheets. For this reason, a template was developed, requiring:

- Scenario description
- Initial conditions
- Initiating event
- Relevant assumptions
- Relief load
- Narrative description of events, with graphs where necessary

Where updated datasheets were required, these could be filled in directly with model data, with the relevant scenario referenced. Sensitivity studies were taken as sub-scenarios of the original case, but were treated identically with full reporting, as the scenario could be completely changed with new assumptions or latent failures.

Making Recommendations

With the simulation built and rated, the scenario defined, and the results known, recommendations can be rapidly developed and checked-out, meaning that it is not necessary to commission a new engineering project to investigate solutions for the findings of a dynamic relief and flare study. A suggested solution to an issue can be instantly built or instituted, and the scenario re-run. If the measure fails, it can be adjusted or replaced, and multiple concepts presented to the process safety engineer for consideration, with confidence that they are effective. This includes testing the effects of a resized relief device, a proposed trip, or the pipework hydraulics around a relief valve.

Example 4: Partial Power Failure

Scenario: Loss of half of power distribution. The power came into the site and was diversified across two main distribution systems: A and B. Loss of power side B led to a significant series of events, with 9 relief devices opening across 4 units, as well as many pressure controllers sending significant flow to the flare header.

The power failure affected equipment throughout the model, including:

1. Debutaniser overhead pumps, causing the drum to overfill (Figure 6)
2. VNS bottoms and overhead pumps (Figure 3, see Example 1)
3. Catalytic Naphtha Splitter (CNS) bottoms and overhead pumps (Figure 6), causing similar event to Example 1
4. Loss of overhead pumps and cooling on the naphtha reformer stripper and debutaniser, as well as the ultraformer Effluent Separator (Figure 7)
5. Loss of cooling and pumps throughout the kerosene unit, but not the feed pumps (Figure 8)

While the configuration of the failures on the naphtha reformer unit's columns was similar to the VNS and CNS, where overhead drum overfill and relief would be expected, the reboilers were heated using recycle gas from the reactor circuit as opposed to steam. This difference was key – the loss of pump-out for the Effluent Separator meant that would be the first item to overfill, which would trip (or disable) the recycle compressors. This loss of flow in fact causes the entire naphtha unit to wind down before it can contribute to the relief load.

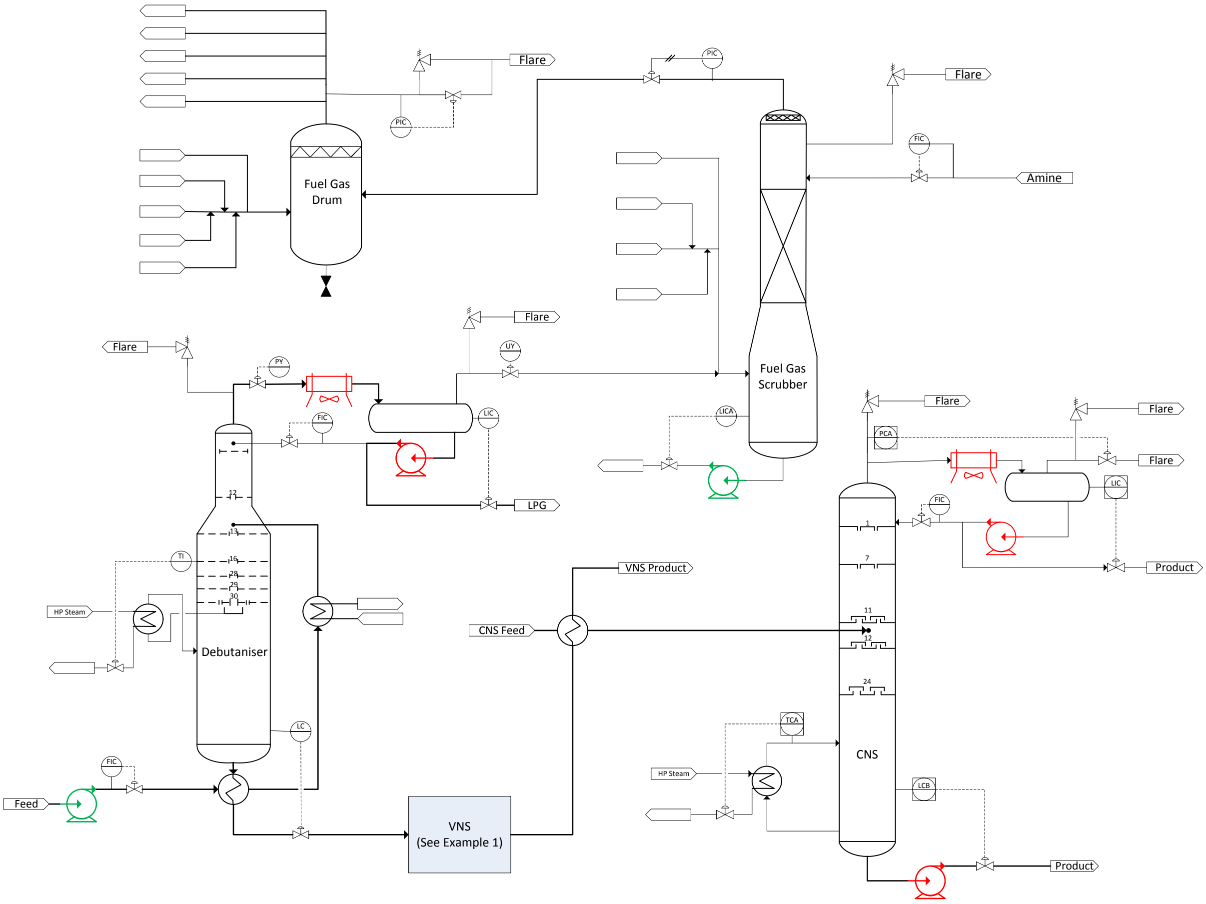

Figure 6. Simplified PFD of section of LPG and Fuel Gas Units

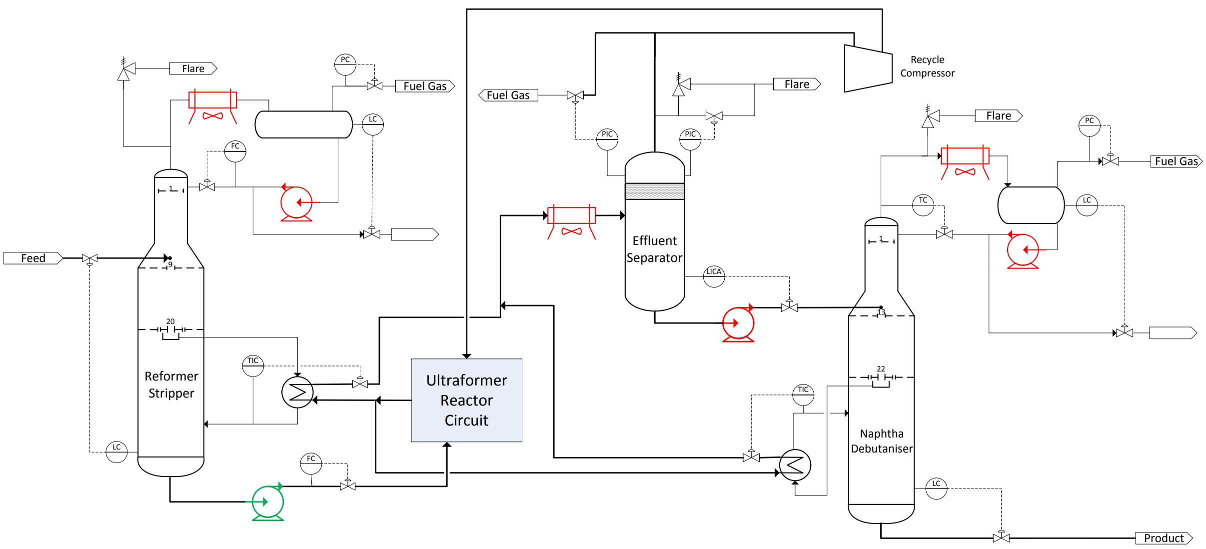

Figure 7. Simplified PFD of Naphtha Unit

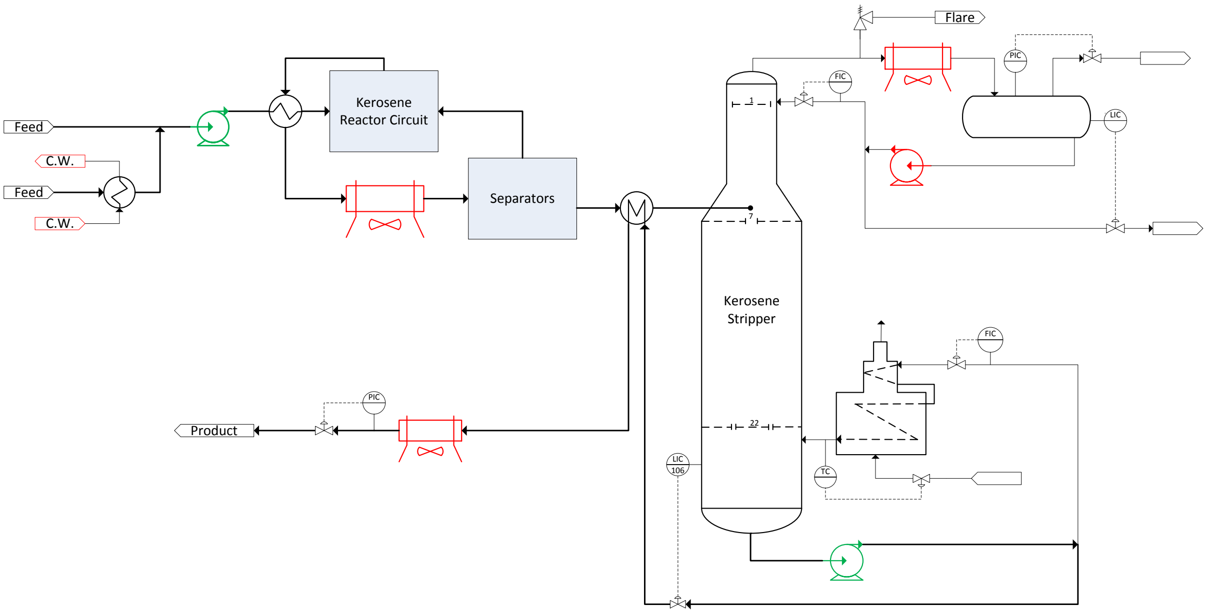

Figure 8. Simplified PFD of Kerosene Hydrotreater Unit

The loss of power to many fin-fan coolers significantly increased the supply to Fuel Gas, while furnaces winding down or tripping out reduced Fuel Gas demand, causing the pressure controller to send excess to flare, and the fuel gas to relieve for a short period. It should be noted that the Fuel Gas system itself was not directly affected by the loss of power –its connections to all of the other refinery units made it vulnerable to this event, and this would be easy to miss in a conventional study. This is followed by the kerosene stripper relieving as the overhead drum overfilled and backed pressure into the column.

Once the debutaniser overhead drum overfilled, flashing LPG was sent to the fuel gas system via the debutaniser pressure controller, causing a much larger relief event from the Fuel Gas Drum relief valve, and exposing the Fuel Gas Scrubber to very cold temperatures. After a build-up of pressure, the debutaniser overhead drum relieved more flashing LPG to flare, before being joined by vapour from the column relief valve. Finally, the VNS and CNS relieved, combined with liquid from their overhead drums, resulting in slug flow in the flare sub-header as well as adding to the already high backpressure in the main header. The cumulative backpressure reached a level beyond what the Kerosene Stripper and Fuel Gas Drum relief valves could safely operate against. Figure 9 shows the cumulative flows from major sources through the scenario, as well as the pressure in the flare header.

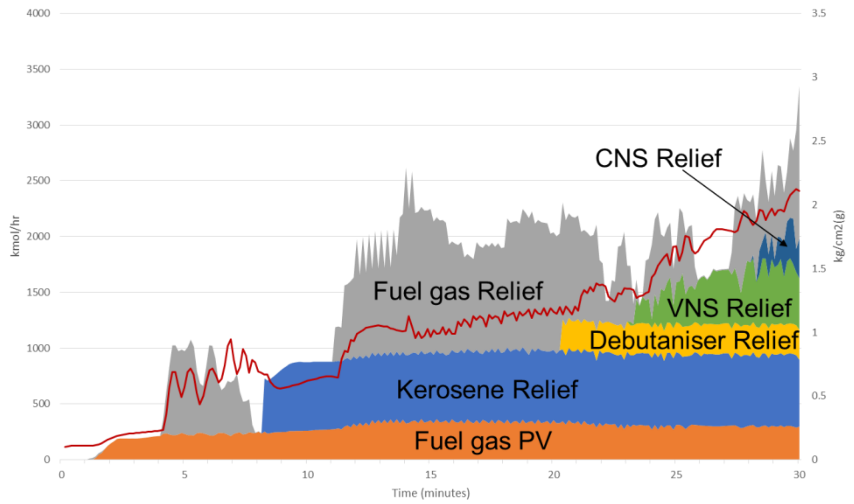

Figure 9. Cumulative flare flows (solid) and pressure (red line) due to loss of power side B

Although this case is severe, and does indeed lead to unacceptable pressurisation of the flare header, viewing the scenario in terms of time and using realistic flows allowed the engineers to develop modifications and avoid recommending excessive capital expenditure. These recommendations included:

1. Installing a high level trip on the debutaniser to prevent LPG going to fuel gas
2. Re-rating the debutaniser overhead drum to avoid LPG relief on overfill
3. Removing the kerosene power to side A, where there was not a similar event

4. Installing high level trips and re-rating the overhead drums to prevent liquid relief from VNS and CNS
5. Undertaking a mechanical integrity assessment of the smaller diameter sections of the flare header near the towers to ensure they could withstand two-phase flow

These recommendations removed the kerosene stripper relief demand, and thus the impact of the high backpressure. Additionally the fuel gas system now only underwent the first portion of its relief event, ending before any other loads were added to the flare header, and removing the high backpressure issue. The introduction of trips increases the VNS flare load, but without other low pressure systems relieving concurrently this is acceptable. Finally, the slug flow data from the model was used directly in a mechanical model, which resulted in recommended additional pipe supports to reduce the risk of loss of containment in the event that the trips on the VNS or CNS failed.

Project Outcomes

By the end of the project, all identified relief and flare scenarios had been run, and all pressure systems and their protective devices sized, including inlet and outlet pipework. A full report on each relief device, as well as of each safety case, was presented to the client, along with checked-out recommendations to mitigate the identified hazards.

Key Recommendations

There were over 100 recommendations from the project overall, including:

- A list of pressure devices to re-size
- Re-design of relief valve pipework to meet required inlet and outlet pressure drop limits
- Reconfiguration of the electrical supply to minimise the hazards of various power failure scenarios
- Vessel re-ratings
- Installation of trips and other safety instrumented systems
- Actions to carry out SIL assessments

By implementing these recommendations, the study concluded that replacement of the existing flare header could be avoided, which would result in significant capital savings. Overall, 6% of relief devices were deemed undersized, 3% were oversized, and 22% had deficiencies in their pipework design. This is in line with industry figures that 8.3% of valves are undersized, and 20.5% have unacceptable inlet and outlet pressure drops (Berwanger et al, 2000).

Troubleshooting and Debottlenecking

The use of dynamic modelling for the relief and flare study brought with it additional benefits to the site, realised throughout the course of the project. As has been mentioned, rating the models involve matching plant data as accurately as possible. Discrepancies between model results and plant data often pointed to unrecorded operational issues, such as fouling or wrongly sized equipment. The very act of pushing the models to the design capacity of the site was in itself a debottlenecking exercise, and allowed simple recommendations to be made to increase the throughput of the plant, and explore the actual operational limits of the equipment.

Example 5: Diesel Hydrotreater Debottlenecking

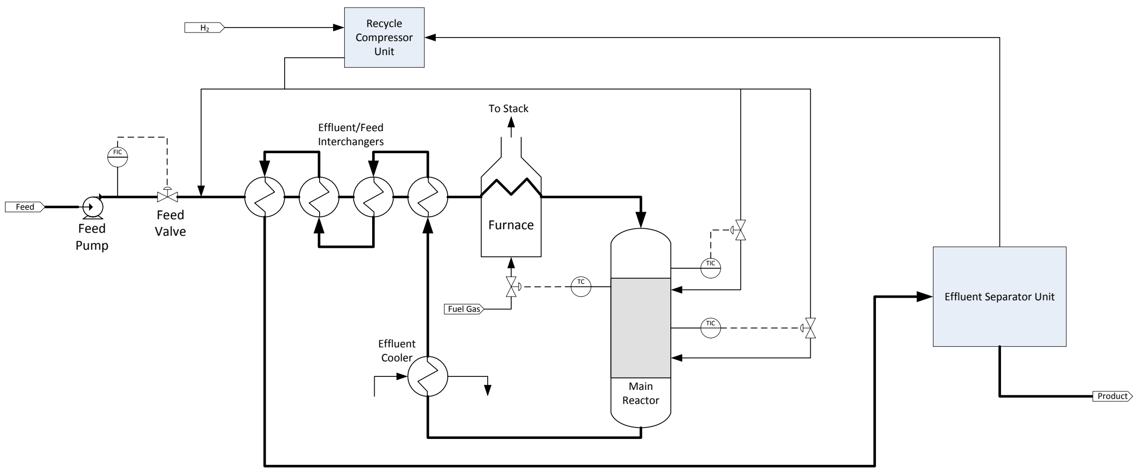

Figure 10. Simplified PFD of Diesel Hydrotreater

During the study, the possibility of debottlenecking the diesel hydrotreater unit was explored. The feed control valve had been identified as a limiting factor, as it was generally fully open. However, the site needed to fully understand the impact of increasing the throughput, particularly on the margin between operating and design pressure. The configuration of the heat exchangers, furnace and reactor resulted in a large pressure drop across the unit, which brought some equipment –

particularly the furnace – near to its design pressure. By building the required modifications into the model, the engineers were quickly able to:

- Recommend a re-rated pressure for the furnace tubes to maintain a safe operating margin
- Identify operating limits on gas recycle rate, to avoid unacceptable pressure drops
- Size the new feed valve against the pump curve and plant hydraulics
- Calculate the new power consumption of the feed pump, which would require a larger motor
- Ensure the relief capacity of the system would be sufficient, particularly for the two-phase case of blocked outlets
- Test for any new relief scenarios because of the modifications, specifically for settle-out; there were none

Additional Benefits

Additionally, the models were used to answer HAZOP actions for non-relief scenarios, as well as providing support for LOPA and SIL assessments. These were for transient events, where there isn't a simple method outside of dynamic simulation to obtain accurate and reliable answers. This included, for example, optimising the set point for high-pressure trips to reduce the likelihood of spurious activation, while avoiding relief events on unplanned shutdowns.

By generating a narrative of how each event unfolded, the simulation engineers were also able to determine the correct operator actions to prevent the scenario from worsening – or to point out situations where the operator is essentially powerless to prevent it.

Summary

Conventional pressure relief techniques are intentionally over-conservative, but this may result in over-designed solutions not practical for legacy plant. Dynamic simulation offers a method of analysing the relief and flare events of complex plant in a more accurate manner, which can reduce the onerousness of relief design while maintaining an adequate level of conservatism. Although API 521 is an excellent standard for relief and flare design, its guidance for using dynamic simulation is minimal, and it is desirable that best practice is formally developed.

Using dynamic simulation, most of a refinery's units and their flare systems were modelled and matched to plant data. A comprehensive study was then carried out, checking the sizing of both the relief devices and their pipework, as well as the capacity of the flare headers to handle the simulated loads. An auditable procedure was developed to guide the engineers through the process of building, rating, executing and reporting on the models. This included developing a philosophy for implementing models in a way that reflects API's guidance on relief load calculations, while allowing for realistic events to play out.

By modelling these complex processes, difficult, interconnected relief scenarios could be analysed and understood, with sensitivity studies undertaken to ensure a safe level of conservatism. Over 100 recommendations were made, but the need for a replacement of the existing flare network was avoided through accurate understanding of relief scenarios. Additionally, non-relief safety cases were highlighted, and plant bottlenecks identified. The models also found use supporting HAZOPs and LOPAs, as simulation removes the subjectivity inherent to scenario assessment.

References

American Petroleum Institute, API Standard 521:2014, Pressure-relieving and Depressuring Systems 6th Edition

Berwanger, P. C., Kreder, R. A. and Lee, W.-S., 2000, Analysis identifies deficiencies in existing pressure relief systems. *Proc. Safety Prog*., 19: 166–172.

Manca, D., and Brambilla, S., 2012, Dynamic simulation of the BP Texas City refinery accident." *J. Loss Prev. in the Proc. Ind*., 25(6): 950-957.

Nezami, P.L., 2008, Distillation Column Relief Loads – Part 1, *Hydrocarb. Proc.,* Sept 2008

Nair, A., Willetts, I., and Wade, A., 2010, Dynamic Simulation for Emissions Regulation, *Control Eng.,* Jan 2010: 43-45

Sirven, A., Grosclaude, J., Fenol, G. and Saada, J., 2011. Optimising safety relief and flare systems. Petroleum technology quarterly, 16(3).

Spooner, J.M., 1994, A Review of Computer Process Simulation in Industrial Pollution Prevention, *EPA/600/R-94/128*, US Environmental Protection Agency