Practical Implementation of System Safety Approaches and STAMP in the Oil and Gas Sector

Conor J Crowley C.Eng. F.I.Chem.E. Process Safety Team Lead, Atkins, Kirkgate House, Upperkirkgate, Aberdeen

With the increase in facility automation and capability of control systems, modern production facilities in the oil and gas sector are becoming significantly more complex, and the operators of the facilities more and more distant from the plants they are controlling. Plant modification, gradual degradation and facility life extension means in many cases that the facilities being controlled are unlikely to operate as per original design. In addition, while the circumstances in any major accident are not likely to be replicated exactly elsewhere, as Dame Judith Hackitt points out "There are no new accidents, just old accidents happening to new people".

It has been pointed out that one way to deal with the complexities of modern plants is to apply a "systems safety" approach. This has been notably developed by Prof Nancy Leveson at MIT, in her STAMP (System Theoretic Accident Model and Processes) model of accidents. The model treats systems as dynamic, with interacting systems of measurement, feedback and control, and design and operational constraints, combining to keep a plant within the safe operating region, and that system degradation combined with inadequate constraint control combine to give the conditions for accidents.

It's clear that the STAMP model has good application in design, operation, and accident investigation, but there is little in the literature to show how this would apply outside of the aviation industry examples, from where the STAMP model emerged.

This paper describes how the STAMP model could be applied within the context of some real-life Oil and Gas examples, covering well operations, production and late-life field management. It also explains how the author believes that it can be used to examine potential "black swan" accident events, and assist safety professionals in pre-emptive accident investigation, shining additional understanding on complex high risk industries.

Introduction

"All models are wrong, but some are useful". George Box & Norman Draper, 1987.

Modern processing facilities are complex machines, and with the ongoing march of computer speed increases, both the plants themselves and the systems we use to control and shutdown the facilities grow more comprehensive and complex in tandem with the systems they are attached to. In the offshore oil and gas production world, it is not unusual to be mixing or attempting to integrate equipment built in the 1970s or earlier with state of the art integrated control and shutdown systems. While pneumatic control systems were already fading out of use in the early years after Piper Alpha and are a distant memory to most engineers now, it can be argued that our design approaches have not always kept pace with the complexity and opportunity the march of progress has delivered. As a result, we are using systems significantly more complex and inter-related than ever before.

One way to deal with the complexities of modern plants is to apply a "systems safety" approach. This has been notably developed by Prof Nancy Leveson at MIT, in her STAMP (System Theoretic Accident Model and Processes) model of accidents. In her book "Engineering a Safer World - Systems Thinking Applied to Safety" (Leveson, 2011), she presents a number of limiting assumptions with other ways of examining accidents. These are shown in Table 1.

The STAMP model is based on systems theory and replaces some limiting assumptions with proposed corresponding revised assumptions.

Limiting Assumptions (Leveson)	Alternative Assumptions
Safety is always increased by increasing system or component reliability: if components do not fail, then accidents will not occur.	High reliability is neither necessary nor sufficient for safety • Systems can be reliable but unsafe - individual components may operate in line with their requirements, but the combination of components and external factors may still result in an accident. • Systems can be unreliable but safe - for instance, they may fail into a safe state, or operators may not follow exact procedures but adapt their approach to prevent a developing situation resulting in an incident.
Accidents are caused by chains of directly related events. We can understand accidents and assess risk by looking at	Accidents are complex processes involving the entire socio-technical system. Traditional event-chain

Limiting Assumptions (Leveson)	Alternative Assumptions
the chains of events leading to the loss.	models cannot describe this process adequately
Most accidents are caused by operator error. Rewarding safe behaviour and punishing unsafe behaviour will eliminate or reduce accidents significantly.	Operator behaviours is a product of the environment in which it occurs. To reduce operator "error" we must change the environment in which the operator works.
Highly reliable software is safe	Highly reliable software is not necessarily safe. Increasing software reliability or reducing implementation errors will have little impact on safety.
Major accidents occur from the chance simultaneous occurrence of random events.	Systems will tend to migrate towards states of higher risk. Such migration is predictable and can be prevented by appropriate system design or detected during operations using indicators of increasing risk.

Table 1 Accident Model Assumptions (Leveson)

The STAMP Accident Model.

The model proposes that events leading to losses occur only because safety constraints were not successfully enforced. Three basic constructs underlie STAMP: safety constraints, hierarchical safety control structures and process models.

These constraints are not only enforced within individual systems and unit operations, but are also influenced by a hierarchy of philosophy, design and operational controls which seek to impose control on the lower levels. In all of this, the concept of a "process model", i.e. a representation of how individual and combined systems work, is important to consider.

The Constraint as a Basic Element of Safety

The most basic concept in STAMP is not an event, but rather a constraint. For an event to lead to a loss, a safety constraint has to be unsuccessfully enforced.

Constraints fall into a range of categories,

- Passive Constraints
    - These include inherent properties of the systems, such as design pressures and temperatures, physical interlocks, and also process safety aspects such as layout, physical barriers, bunding arrangements, etc.
- Active Constraints
    - These rely on action to provide protection, and generally involve detection of a developing hazardous event or condition, typically by measuring something, interpretation of the measurement and response using some final elements, which must be completed before the developing situation results in a loss.

As a simple visualisation of the accident space, Figure 1 is a map of a typical system.

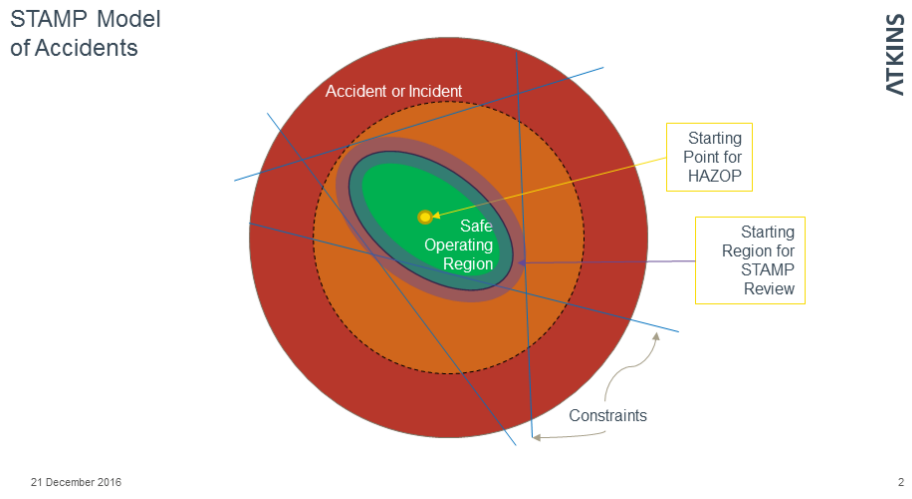

Figure 1. STAMP Accident Map

In the map, there is a safe operating region, where losses are not possible, and a loss zone, visualised in red. The constraints are intended to ensure that the accident line is not crossed, but the diagram is also drawn to show that the constraints on their own may not keep an accident from happening - a similar idea to the "holes" in the Swiss Cheese model.

It is worth noting that the starting point for traditional HAZOP studies is the premise that there is a notional steady state, and that the plant is safe in that state. This is the equivalent to starting from the middle of the safe operating region above, whereas we will consider areas closer to the edge of the safe operating region.

Dynamic Systems Under Control

Traditional models of incidents, including the "Domino Effect" and "Swiss Cheese Models", if applied literally, can lead to the interpretation that accidents are random events, and that it's only when you are unlucky enough for a barrier to fail that an incident results. In many cases, these barrier failures, and the particular action taken by an individual are focused on in incident investigations, but this is not always useful in improving safety.

Leveson provides an illustration at unit-operation level of a typical system under control, see Figure 2.

A controlled process in general will take inputs from connected systems, handle disturbances, and generate outputs, which in turn can become inputs to other connected systems. The state of the controlled process is detected using sensors, which often feed into an automated controller. These have either explicit or implicit models of how the controlled process should react, and then apply control algorithms to change the state of actuators (such as valve position, heater output, etc.) which will in turn change the state of the controlled process.

For instance, if a separator level is increasing, this should be detected by a level instrument, and a control algorithm triggered to, say, open an outlet control valve by a certain amount to restore the level to the target. If the control system is not able to affect the level quickly enough, then a separate sensor is often triggered, which should, say, cause a unit trip.

For complex operating plants, the automated controllers are often embedded within an integrated control and shutdown system (ICSS), with many hundreds or thousands of input sensors, trips, and actuators. Typically in these cases, there will be an operator or team of operators monitoring the entire plant, and responding to alarms and system alerts. The automated controllers will present a display of the plant status, the inputs as measured by the sensors, and any current state of the actuators. The human controller can take manual or computer assisted actions on top of the automated control system, based on their procedures, and their wider understanding of the plant, such as environmental impacts. In all cases, they are using a mental model both of the way the process works, and also the way that the automation system should work.

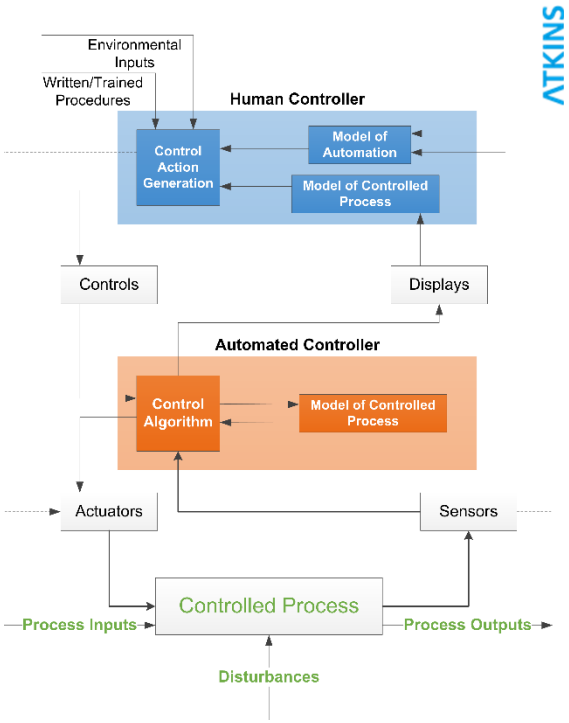

Figure 2 Unit Operation and Control

Without any control or constraints, a system will inevitably degrade and may drift towards an accident. A typical disturbance (see Figure 3) would be detected and control action taken to restore the plant into a safe state, albeit often a different state from the starting position, depending on the control algorithms applied. In complex systems, there will also be the issue that a control action which would be safe in many circumstances would in fact move the plant to an unsafe state and result in an incident.

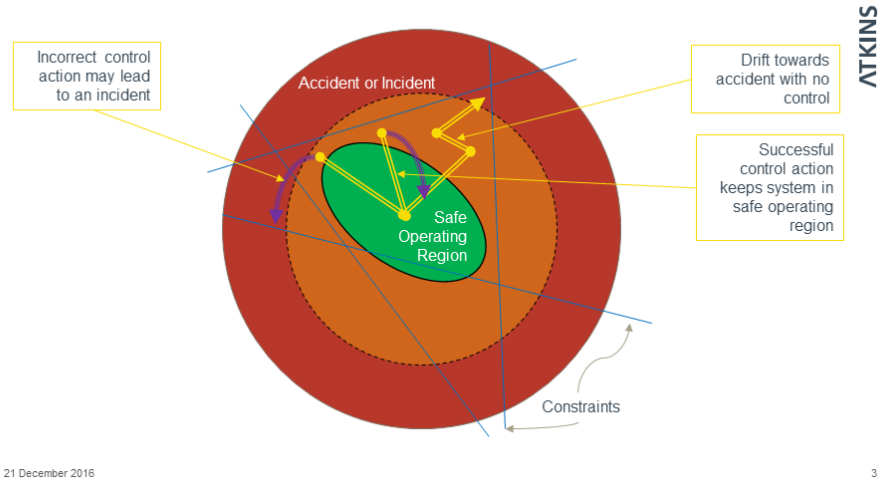

Figure 3. Control Actions on the Accident Map

There will be some disturbances that will activate an active or passive constraint. For instance, a sudden pressure rise may be higher than can be handled in the control system, and if the pressure continues to rise, a pressure safety valve (PSV) may lift. If the rate of pressure rise is not extreme, the operation of the PSV should prevent the vessel being compromised. However, in some scenarios, the rise in pressure may be so extreme that even with the PSV operating, the pressure limits for the system may be already compromised, and even with the control action taking the system back into less hazardous areas, the damage is already done (see Figure 4).

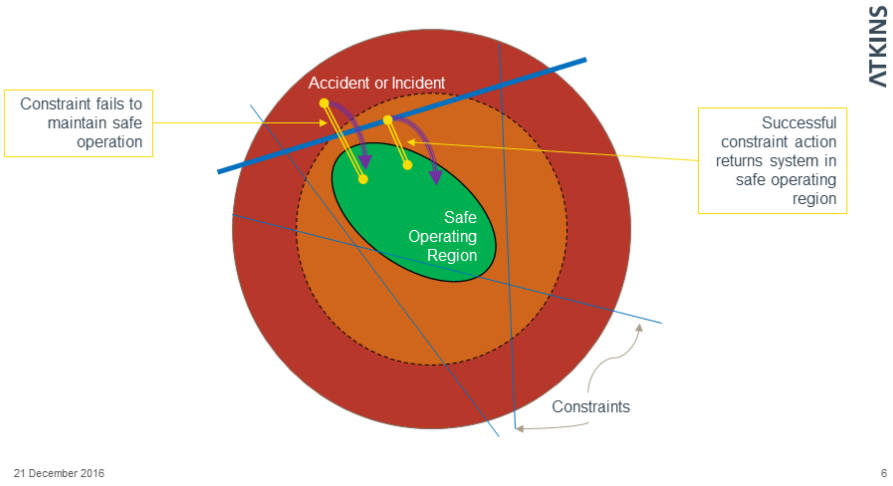

Figure 4. Constraints on Accident Map

Hierarchy of Control

In the management of high hazard industries, there are many layers of influence and control on the design and operation of a facility. At the highest level, the country legislature and standards bodies provide the over-arching framework in which risks are managed. Below that there are company and contractor approaches, standards and methodologies, which all result in plant that are operated in line with company procedures, and managed in line with a company management system. Third parties may also have an oversight role, for instance, with the independent verification of performance standards on offshore production facilities.

At each level of the hierarchy, there is a system of control in place, either as a management system or a physical/electronic control system. These systems will interact in complex ways, and it is not always clear or even calculable what the actual impact of the control on the sub-systems will be.

Process Models

The concept of a process model in the context of STAMP is far wider than process simulation: put simply, it's the way that inputs, processes and disturbances combine to provide outputs for the system. At the unit operation level in the hierarchy, the process model will link the physical and thermodynamic processes going on, e.g. how a primary separator allows fluids from a production well to flow in a vessel producing output streams of gas, oil and water of the required conditions and

compositions. Variation in the chosen operating parameters on the plant can produce different outcomes, and disturbances, such as slugging in upstream pipelines, can cause failures of interfaces to form, and result in contamination of outlet streams.

In upstream oil and gas, the knowledge of the exact fluids to be processed is not always well defined at the project phase, and reception facilities are designed to have a wide range of effective operation. This is informed by detailed analysis, such as flow assurance modelling of feed pipelines, but by their nature these may not reflect the reality of actual conditions. In operation, the process operator may gain much more experience of the dynamics of the inflow to the facilities, but this may lead to assumptions of how things work based on developing personal cause-effect relationships that are not valid across the operating range. Changes in well composition over time or subsequent field developments may well move the actual operation of any plant to a range not originally assessed or envisaged. Also, there are few good mechanisms to feed the actual operations back to the design personnel, even if there is a remaining involvement between the system designers and the people operating the plant. With the ultimate fruits of the design effort often being installed many miles offshore, and with the fact that it is rare for an engineer to maintain full contact with a facility from engineering through operations, there is plenty of potential for issues to slip through the gaps.

Of course, there are many ways in which the industry has attempted to manage the issues of incomplete process models. For instance, both company standards and international standards attempt to provide guidance for design across a range of operations, and provision is made in the emerging standards for known significant issues. However, these cannot be ahead of the technological developments in all cases, and significant items tend to be covered in detail only after a major incident has occurred.

Applying STAMP in Upstream Oil and Gas Production

The STAMP Model and the systems theory underlying the approach clearly brings fresh insight to the consideration of accidents and how they develop, and Prof. Leveson and her team at MIT have put a number of publications, presentations and examples into the public domain. It is not the intention of this paper to summarise those, but rather to extend the range of examples used into the oil and gas context.

Degradation and Model Assumptions

The integrated control and shutdown system on a North Sea Production complex was built up of a group of interconnected nodes. The facility comprised four separate jacket structures, one for wellheads, two for production and one for utilities and accommodation. The jackets were connected by bridges, but segregated sufficiently to make escalation of an event between jackets unlikely. The control systems were housed in safe areas on each platform, and in line with common design approaches, detection of smoke or gas on the intakes to each of the safe areas on each platform would result in a very high level shutdown (Level 1), which shut down and depressurised all process facilities, and removed all electrical power apart from emergency power across the complex.

The ICSS system was again set up in a conventional way, with multiple redundant input, logical and output components, with advanced diagnostics and controls. The individual nodes were configured so that if there were significant faults in the system, the nodes would “fail safe” and initiate the highest level shutdown that was contained within the node.

Main power generation was provided on the utilities and quarters platform, along with emergency generators. With time, the emergency generators were proving relatively unreliable and efforts were ongoing to improve the system. However, an incident occurred on one of the platform jackets where there were multiple component failures in a single ICSS node. This node “failed safe” and initiated a Level 1 shutdown. This removed all normal power generation across the complex, but as the emergency generators did not function, had the effect of removing all the power. As the shutdown system was operating as if there were gas in a safe area, it would not allow the main power generators to restart, and it took close to two days to get power re-established to the facility. In the meantime personnel were exposed in an accommodation block with very limited life support and other facilities, not in response to a developing hazard but rather in response to a “fail safe” action. In the event, there were no injuries.

Applying a STAMP mindset to this incident, the following observations are noted.

There was a disparity between the process model of the designers and the reality of the separation of the different components of the system. While unexplained gas or fire in a safe area is not to be ignored, the fact that the full shutdown was taken across the complex regardless of the fact that the accommodation platform was 150m distant from the nearest facility, and up to 400m away from the wellheads area, was not taken into account in the design of the shutdown. In fact, the control system could have been configured to reflect this, treating each platform independently, but for simplicity, the control system was in fact configured as if it was entirely co-located in the same safe area.

Also, the assumption that the “fail-safe” option of shutting down and electrically isolating all non-emergency power was a safe control action was not true in all cases. In fact, the fail-safe approach (which occurred without any actual hazard being present) resulted in the accommodation being put into almost complete darkness (apart from battery powered electrical lighting).

Constraints not being Enforced

There have been a number of papers in previous Hazards conferences around the issues with the cargo tank vent on the Global Producer III, a tanker converted to an FPSO. The first incident on GPIII occurred on a day with very low wind-speed. While such days are rare in most offshore locations in the UKCS, it is equally rare that a tanker would be stationary in normal operations, and hence the vent would be sufficiently well located for normal tanker operations. It was only when

the geo-stationary FPSO encountered unusually still conditions while venting heavier than air hydrocarbon that the gas reached the production deck and ignited.

In this case, STAMP would consider the prevailing wind flow as one of the design controls on the vent, and explicitly considering this may have revealed further design considerations for low-wind conditions.

Applying STAMP in Drilling Operations

Control of constraints in drilling and well operations is significantly different to the production facilities. While production heavily relies on automated systems to monitor and control, and a small operations crew monitors a large number of key variables, in drilling there is often little or no automatic control, and management of the hazard relies on close monitoring of well variables by a group of personnel, and taking appropriate action at the correct time.

For instance, during drilling operations, the ultimate protective device known as the "Blow Out Preventer" is almost always operated manually only, and operates by either sealing around the drill string/tubulars in the well or by cutting the drill string/tubulars. As the last line of defence, this operation is not taken lightly, but the Macondo well incident is one well-known example of where the decision to close the BOP was taken too late.

For the specific case of drilling operations, the control cycle in Figure 2 can be redrawn with typical issues as shown in Figure 5. If an electronic system exists for the flows around the facility, this generally is just providing displays from sensor information and allowing remote operation of manual controls. The errors and potential action failures shown around the loop share many of the prompts that would traditionally be used in a batch HAZOP or human factors review.

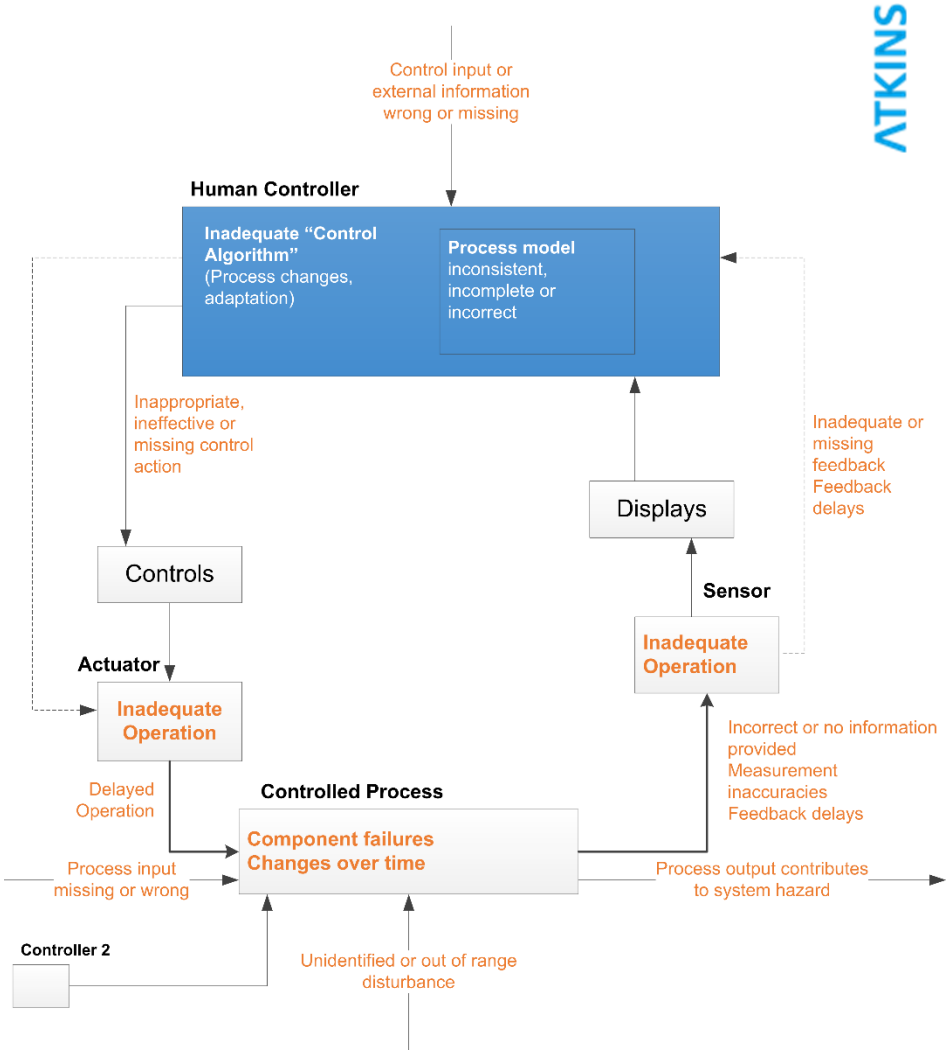

Figure 5. Human Controller and Potential Issues without ICSS

The STAMP Model of accidents does point out that accidents occur when the process model used by the controller (automated or human) does not match the process, and there are four main categories of control failures:

1. Incorrect or unsafe control commands are given.	Incorrect control actions would include routing fluids to the wrong place, breaking containment without sufficient barriers to flow, etc.
2. Required control actions (for safety) are not provided.	Failure to provide required control actions would include the use of a shear ram to cut a drill string that was not capable of doing so across a pipe joint
3. Potentially correct control actions are provided at the wrong time (too early or too late) or	In the case of Macondo, the decision to close the BOP was taken too late to prevent the loss of well control.
4. Control is stopped too soon or applied too long.	To safely remove entrained hydrocarbon from well drilling fluids, they are circulated into topsides gas removal facilities to safely remove the gas. If this is not carried out for long enough, then there may still be uncontained hydrocarbon in the well-bore that could result in a loss of primary containment in later procedural steps.

Table 2. Control Failure Categories

Consideration of these failure types in line with the plant design and the proposed drilling operations can provide a rich context to examine and test proposed combinations of well activities.

Pre-emptive Accident Investigation using STAMP Mindset

If there are no new accidents, just old accidents happening to new people, it could be argued that it should be easy to identify and therefore prevent the next big accidents. However, there are a number of challenges to this, including:

- The relative rarity of major accidents can induce complacency and the optimism bias of "it couldn't happen here".
- Degradation of systems can be a very gradual process, and the point at which a system moves to a point where it's one step away from a major incident may be hard to see.
- While there is rightful focus on performance standards and integrity management programmes, these may reinforce the view that enough is already being done to prevent any incident, whereas the majority of these programmes in fact attempt to manage to a given level of probability of failure, rather than to prevent failure in all cases.
- There are often invisible gaps in integrity management approaches, in that multiple people may well be focusing on the mechanical equipment, e.g. electrical, control, mechanical review of the operation of an emergency shutdown valve, but relatively little focus on the pipework connected to that valve.
- Accident investigations will often point to known degradations of the system, but only the benefit of hindsight automatically links these to the incident.

As illustrated in Figure 1, the valid starting point for any review based on STAMP is not necessarily the steady-state position that a traditional HAZOP would consider, but rather a zone around the transition point from "safe" towards "accident".

The STAMP literature does provide a methodology, referred to as STPA (the System Theoretic Process Analysis) technique, which is a largely paper-based design approach to apply the STAMP causality model. An alternative workshop-based implementation is described below.

The process starts with a preparation phase, in which each major unit operation or connected process is identified, the constraints to ensure the levels of safety required are considered and documented, the relevant sensors identified and their range of operation determined. Any persistent or likely disturbances that are inherent to the system would be considered, and the control mechanisms in place to address each constraint, be they active or passive, are considered. These would be documented across the process along with their associated actuators where relevant. The cross-connections between each systems and the direction of fluid and energy and other property transfer between the units would be identified. These would be documented in a matrix similar to Figure 6.

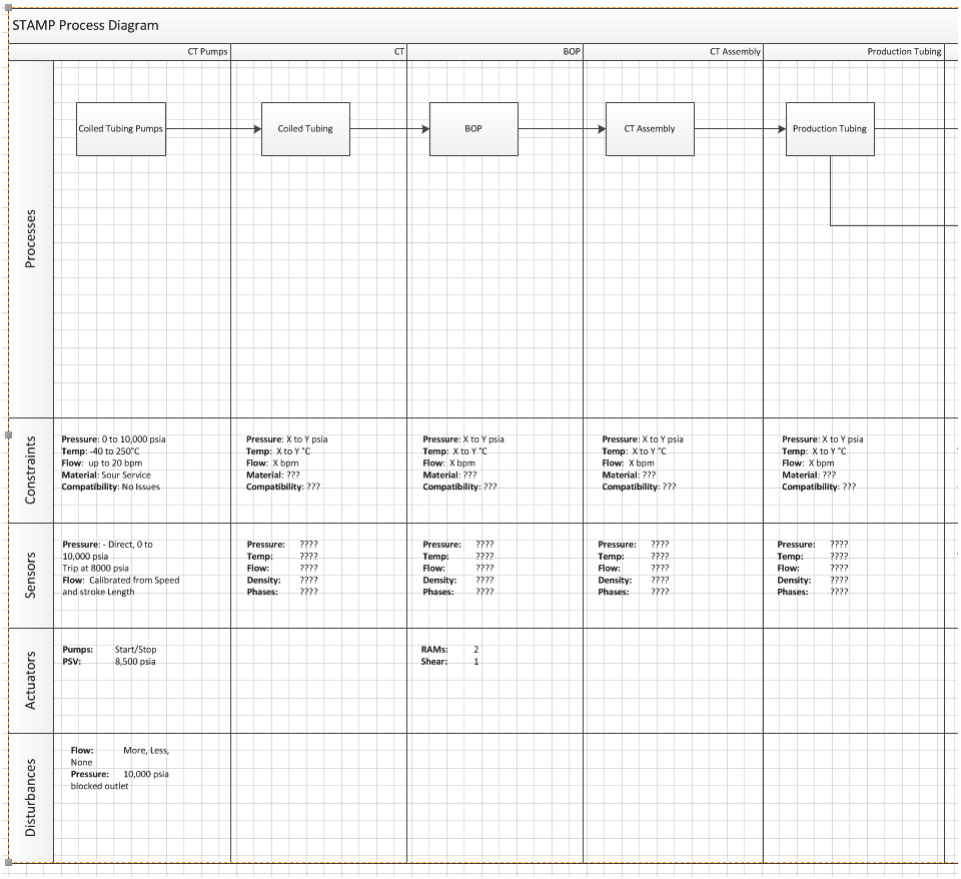

Figure 6. Sample STAMP Process Diagram

This diagram provides the basic template for the key operations and the potential control actions. These would then be used in a collaborative meeting to review the proposed activities or design, and identify key degradations that could happen, the additional factors that may turn such degradation into an incident, and the ability of the automatic or human controllers to interpret and react to these factors.

For instance, it is common in many drilling activities to carry out a test of pressure containment prior to moving to an additional step. A positive test outcome is of course not 100% definitive, so the STAMP Review would consider what would be the worst starting point to continue with the operation, and then look to define the possible ways in which this could lead to an uncontrolled flow of hydrocarbon, and, if the barriers are not successfully implemented, a potential well control incident.

Additionally, for drilling operations, there are often multiple personnel involved in the control of fluids into and out of the wellbore. For example, there may be personnel controlling the speed and rate of flow from high pressure pumps, others monitoring the return flow from the well, with the drilling supervisor in overall charge. There is limited information about the conditions at the drill-bit and along the drill string, so the process relies both on the design of the operation, in which the wells and fluid flows are modelled, and also the mental model of the personnel involved to ensure that the operation continues as originally intended. Carrying out the review in this format allows consideration of the drilling engineer's mental model in more detail, and ensure that it is clear that this understanding is shared by all, along with the possible signs that the reality of the well may be differing from the models, and how this can be handled and controlled in the field.

The approach is currently being trialled, but results are not available at the time of writing to demonstrate this in more detail. Having said that, the hindsight and optimism biases described above are explicitly handled in this approach, in that rather than trying to defend an existing approach to management of major hazards, the team are instead trying to work out any possible circumstances that could conceivably be present on their operation which would lead to a major hazard in 1-3 steps, going from "it couldn't happen here" to "let's find the ways it could happen here".

Conclusion

There are many different ways to describe an accident, and they provide their own insight into the prevention of accidents. We believe that STAMP can be another weapon/tool in the armoury of process safety professionals, and will be further developing this approach to extend its application from its heartland of aviation systems into the complex infrastructures and systems we deal with in the Energy Sector or Oil & Gas industry.