

Pushing Boundaries: Two Amine Regenerator Case Studies

Jim Critchfield, Retired/Consultant

Joel Cantrell, Bryan Research & Engineering

Christina Daniels, INEOS GAS/SPEC Technology Group

**Presented at the 30th Annual Brimstone Sulfur Recovery Symposium
Vail Colorado September 11, 2023**

The refinery amine unit regenerator is a significant consumer of energy, and with so much emphasis currently placed on carbon intensity, there is an increased interest in the performance of regenerators. The primary function of the amine unit regenerator is to reverse the acid gas absorption reaction, and by doing so provide adequately-stripped lean solvent for the amine absorber(s). However other often-minor considerations can also influence regenerator energy requirements.

In this paper we focus on two case studies that highlight regenerator operation. By analyzing the conditions in the regenerators with the PROMAX^{®1} amine simulator, we examine performance and consequences when the systems operate near low-energy and turn-down boundaries.

Both Cases are set in refinery primary treating systems. The first case explores the consequences of a step change in operating conditions resulting from a refinery crude diet change. The second case reveals effects of a loss of lean/rich exchanger performance. “Rules of Thumb” for managing reboiler duty are considered for both cases.

Paper Scope

This paper focuses on refinery primary treating system regenerators. Two specific real-world cases are discussed.

Energy efficiency is a current emphasis, and in a typical refinery amine system, the regenerator is the major energy consumer...and therefore a target for conservation. In our carbon-focused reality, refinery energy advocates seek low-energy modes of operation. For amine regenerators, low energy operation may have unintended consequences, however. Forewarned is forearmed; we discuss some consequences of low-energy operation.

The Case studies are analyzed through process simulation to understand the operating conditions, and comparison is made to operating “rules of thumb” for the Cases. Simulation is also used to extend the operating conditions to demonstrate the effect of minor process considerations on system performance.

¹ ProMax is a registered trademark of Bryan Research and Engineering, LLC

Refiners Operate in a World of Change

Over the last four decades, US refiners have responded to shifting crude diets and moving environmental targets. Beginning in the mid-1980's, refiners experienced increasing sulfur content in their average crude diet; that trend leveled off after the year 2000. In particular, PADD 3 (NM, TX, AR, LA, AL, MS) has seen a decline in the concentration of sulfur in the crude diet, coinciding with the growth in USA NGL produced from hydraulic fracturing. The chilling effect of COVID on the economy – in particular, travel – hit refineries; creating a negative “COVID Bobble” that further decreased sulfur in the crude diet. Data of the US EIA illustrate these trends in [Figure 1](#).

Starting in the mid 1990's, US refiners responded to the increasing demand for sulfur capacity by investing in sulfur units (Amine, Claus, and Tail Gas systems). Environmental pressures also played a role in capacity build; in many cases acid-gas flaring-avoidance drove decisions to install spare capacity. Data from the US EIA in [Figure 2](#) demonstrate that sulfur charge rates to refineries trended with installed capacity until the early 2000's; after that time, sulfur capacity continued to grow while sulfur charge rates plateaued. Then the COVID Bobble hit; in 2020 charge rates dropped further, creating even lower capacity utilization.

[Figure 3](#) shows inferred capacity utilization of sulfur units in refineries, compared to EIA's Crude Utilization metric. The COVID Bobble hit the overall refinery industry hard, but it hit refinery sulfur process units harder. The bottom line is that sulfur processing units in refineries have had to deal with shifting load utilization, and for many, turndown operation has become a new normal.

Refinery operations produce hydrocarbon gases containing CO₂ as well as H₂S; a primary source of CO₂ in fuel gas is the FCC. FCC-generated CO₂ is at least partially absorbed in refinery amine units and contributes

Figure 1




Figure 2





Figure 3



to acid gas loading. EIA data for FCC charge and capacity imply that CO₂ production has been more stable over time than sulfur charge rates have been. As a result, not only has the total amount of acid gas in the refinery changed, but so has its average composition.

It's well-recognized that refinery amine systems become polluted with heat stable amine salts (HSAS); the strong effect of this pollution on system performance has been a prolific subject of publication. Conversion units such as FCC are a known source of HSAS pollution. Relatively stable FCC rates suggest HSAS continue to plague refinery amine units even as sulfur loads change.

The new normal is that refinery sulfur systems must successfully surf throughput *and* composition changes. This macro environment has implications when applying Rules of Thumb for energy targets; the effects of rate and composition changes must be understood and incorporated into practical targets for operating conditions in sulfur units.

The refinery in Case 1 experienced the effect of these trends on amine unit performance: the case study occurred during this period of new normal, amplified in the COVID Bobble. The sulfur load and the solvent composition changed during the case study, and effects are discussed later in this paper.

Regeneration Energy

The primary goal of the amine regenerator is to make properly-stripped lean amine so that the refinery's H₂S absorbers stay on specification. The amine regenerator consumes energy to achieve that goal. There are three major contributors to energy required in the regenerator:

1. Sensible heat is the heat required to bring the rich solvent (and reflux) up to temperature.
2. Heat of reaction is the heat required to keep the solvent hot as the endothermic evolution of acid gases occurs.
3. Heat of vaporization (of water) is the heat required to create dilution steam that "pushes" the acid gases away from the solvent, creating driving force, sending the acid gas up and out of the column.

The total heat requirement is provided by steam supplied to the reboiler, subject to two minor puts-and-takes:

4. Preheat of rich amine is the heat recovered from hot, lean amine as it preheats rich amine in the L/R exchanger. This offsets heat that otherwise would be added to the column.
5. Cooling of the regenerator overhead by heat removed in the condenser happens when cold reflux returns to the column.

Since the reboiler steam requirement is the largest single component of energy input into the amine system, it's the obvious target variable for energy optimization efforts.

Energy Targets

The correct amount of energy for the regenerator is that which provides the necessary composition of lean solvent. In refinery main systems, the sulfide that remains in the lean solvent limits the performance of the refinery's gas (and liquid) contactors. So there is incentive to push the lean loading down to very low values. Excessively deep stripping of lean solvent is energy intensive, and so "over stripping" is often considered to be a waste of energy.

A balancing act results and refiners usually have targets for both Low and High sulfide in solvent ... when they are able to successfully measure those targets. However, the measurement is usually performed offline, and delayed by the hours (or even days). Operators need practical, real-time targets to keep their plants running and so it's often "Rules of Thumb" that are actually applied to adjust energy consumption in regenerators.

The various Rules of Thumb reflect differences in available information, hardware and operating philosophy among refineries. Most commonly we encounter the following:

Rules of Thumb: Control Setpoints - The Big Three

The most commonly used Rules of Thumb are Steam-to-Feed, Reflux Ratio, and Overhead Temperature. These are readily available based on typically installed instrumentation, and are intuitive indications of heat reaching the top of the regenerator.

Steam-to-Feed (or Steam-to-Lean)

This is a simple ratio of heat input to solvent circulation. The flow rate of steam to the reboiler is divided by the circulation rate of solvent (typically, measured on the lean). Sensible heat is a portion of the net heat requirement: the warm lean solvent must be heated to regenerator temperature. So it's "sensible" that changed circulation would require different heat input.

The ratio predicts the required heat input only if other conditions in the system are constant. For instance, if the rich solvent contains more (or less) acid gas loading, the ratio alone doesn't compensate for that change. Also, if the regenerator is made more (or less) efficient by changes in the chemistry of solvent, the ratio misses that.

When rich solvent flow meter is used, steam-to-feed ratio is also sensitive to holdup during short-term level control events; for instance, an absorber foaming event that pauses rich solvent flow. In applications susceptible to such events, the selection of solvent flow meter location (lean, rich) may help to address that issue.

Reflux Ratio and Stripping Ratio

The hot acid gas at the top of the regenerator is saturated with water. Outside of the regenerator, the acid gas is cooled and condensed, and the cooled water returns to the top of the regenerator as reflux. The

amount of water in the hot gas is a measure of how much heat remains after warming up the solvent, reversing the reactions and producing acid gas.

There are two flavors of how to exploit that insight for energy control to the regenerator. Since the flowrate of reflux water is usually measured when it returns to the regenerator, and if the flowrate of acid gas is also measured, a simple “Reflux Ratio” is available. Because some water vapor is left in the acid gas downstream of the condenser, the Reflux Ratio changes somewhat depending on the condenser operation. An improved concept takes the condenser out of the picture, by calculating the water concentration of the acid gas in the regenerator overhead and correcting the Reflux Ratio for that value.

The strength of the Reflux Ratio or Stripping Ratio is that steam reaching the top of the regenerator suggests that most of the requirements of the heat-consuming reactions are met, and that the rich solvent has been preheated to regenerator temperature.

The weakness is that varying composition of the solvent, and the resulting changes in efficiency in the regenerator aren’t well considered. Another weakness is that acid gas flow meters are not usually considered “necessary to run” and may not receive attention: the orifice may be in poor condition. The instrument may not be properly pressure and temperature compensated, and the true meter factor may not be correct for varying composition of acid gas. Rarely do plants have analyzers on the acid gas stream to extract the actual values of H_2S and CO_2 molar flow. Usually, assumptions are made about that composition, grounded in past plant tests (or even, design compositions). The easiest assumption is that acid gas is all $\text{H}_2\text{S} + \text{CO}_2$, minus water: typically, a correction for the water vapor in acid gas is applied.

The reflux flowrate is a result of controlled level in the Reflux Accumulator Vessel. The success of that control may be challenged at low reflux rates: the measurement of flowrate may be the average of periods of low (or zero) flow rate when the controller struggles to maintain level.

Overhead Temperature

The Reflux and Stripping Ratios use measured flowrates (of liquid water and of acid gas) but not all amine units have those measurements (or trust them if they do). A simpler approach is to infer the water content of the regenerator overhead by measuring the temperature of the hot acid gas as it leaves the top of the regenerator. As with the Reflux and Stripping Ratios, the temperature is a measure of how much heat makes it to the top of the regenerator. A high Overhead Temperature is an indication that sufficient energy has been supplied to heat the rich solvent and reverse the reactions. Most plants have a temperature measurement, and many operators simply use this Rule of Thumb of “how hot is enough, how cold is too little”.

The strength of this approach is simplicity. One weakness is that the boiling point of the solvent varies with operating pressure. And at low overhead temperatures, the measurement loses sensitivity to process conditions, leaving the operator without a handle on the heat input. Just as with the other approaches above, the overhead temperature doesn’t fully consider changes in regenerator efficiency resulting from solvent composition changes.

Because they approximate the same thing – steam that reaches the regenerator overhead -- Overhead Temperature shares a weakness with Reflux and Stripping Ratios. When hot rich solvent flashes at the feed tray, that adds heat to the overhead. But that heat “upstairs” doesn’t help much with the dilution steam needed to strip H_2S deeply “downstairs”...at the bottom of the column. In a scenario where a lot of heat is supplied upstairs, the Overhead Temperature may misguide the steam controller to back off steam.

Rules of Thumb: Control Setpoints – Other Approaches

Temperature(s) Inside the Column

When very low heat rates are applied, the overhead temperature becomes insensitive to heat rate. An approach to improve control has been to measure temperature lower down in the column, where the relationship between heat rate and temperature is still sensitive. This approach has the weakness of losing sensitivity if the ‘sweet spot’ in the relationship temperature/heat input moves away from the point of measurement.

An approach that has been mentioned is to have multiple thermocouples throughout the column, sufficient to track changes in the temperature profile. And then to use that profile to control the heat rate. Application of this technique is uncommon, yet it potentially improves upon a single-point measurement which is vulnerable to conditions that shift the temperature profile away from the point of measurement.

Gas Analyzer – “Canary in the Coal Mine” Contactor

One practical approach has been to identify a contactor in the refinery that operates with a close-approach to equilibrium w.r.t. H_2S in the treated gas. Ideally, that contactor should operate at low pressure in order to have higher concentration of H_2S in the treated gas at equilibrium. To the extent that the equilibrium-approach remains in place with shifting conditions in the refinery, an analyzer on this gas stream will respond to changes in the lean solvent – especially depth of stripping. Gas analyzers are common technology in refineries, so there is comfort with the application and understanding of its reliability.

Analyzers require periodic calibration and maintenance; that may be difficult to sustain unless the analyzer is considered a “need to operate” item. One vulnerability is that the approach to the equilibrium condition – even the equilibrium condition itself -- in the “canary in the coal mine” contactor may change due to local condition changes (temperature of operation; pressure of operation; liquid/gas ratios; amount and composition of acid gas removed).

Gas analysis is usually rapid and it directly relates to the lean loading -- when the approach to equilibrium is close -- providing near real-time indication of over or under stripping. This technique works best with low-pressure contactors, as they provide more analytical leverage to see H_2S values in deeply stripped solvent. However, in a large, integrated refinery, the “Canary” contactor may be quite some distance from the regenerator. The resulting time-lag in that situation is a potential weakness when dealing with short term process changes.

Control Setpoint Validation

Online, Real Time Measurement of Sulfide in Lean Solvent

This is the holy grail of energy optimization, and over the years technology for this has been launched several times. Yet we haven't encountered widespread adoption. Rapid, direct measurement of the sulfide target makes logical sense from a control viewpoint.

Analyzers that employ optical measurement have the weakness of the dirty, potentially solids-laden and/or scale-building environment of an amine unit. Other types of techniques have been offered as well. Justifying this sort of expenditure based solely on steam savings was previously difficult; perhaps the new emphasis on carbon intensity will foster more widespread adoption in the future.

Offline, Time-Lagged Measurement of Sulfide in Lean Solvent

Usually this measurement involves discreet sample collection, transport from the location of the regenerator to a lab facility, time in the queue with a shared instrument that processes other samples potentially of different systems. Then reporting and time to react and incorporate the results into (possibly) adjusted targets for the hands-on control techniques.

The experience with measuring sulfides in lean amine in refinery environments is a checkered one. The collected amine samples are 'live' and subject to loss of sulfide over time – for instance from oxidation due to air contamination. Also the measurement itself has challenges; it can be a titration to an inflection endpoint yet multiple inflections may be detected. The instrument often is set to titrate both sulfide and mercaptide; for the same amount of titrant, these results differ by a factor of 2. Misidentification errors are common.

The most common analysis method (AgNO₃) is usually optimized for a certain mass of sulfide; deeply stripped and poorly stripped samples tend to have less practical accuracy. Human errors happen, including transcription errors.

Even when the samples and the instrument are accurate, results can be very noisy – for

good reason. This noise can be a real condition in the regeneration system itself; we will discuss this more later in the paper. It's understandable to encounter mistrust of the measurement among operators, who instead tend to rely on their Rules of Thumb (the real-time controls that they have at hand). Table 1 summarizes several different approaches to guide reboiler duty.

Technique	How to Get it
Overhead Temperature	Measured AG T before condenser
Reflux Ratio or Stripping Ratio	Reflux Return Flow Meter / Acid Gas Flow Meter or calculate with that: Moles Water in OVHD / Moles H ₂ S+CO ₂ in AG
Steam-to-Feed (or lean)	Steam Flow Meter / Lean Solvent Flow Meter
Temperature Elsewhere (vs OVHD)	Thermal wells; skin T's
Canary in Coal Mine	H ₂ S gas analysis – integrated Low P contactor (overcirculate; close approach to equilibrium)
Reset with Lean Sulfide – Measured in Lab	Amine sample → Lab: AgNO ₃ titration (and other); Opportunity for human and procedural errors
Reset with Lean Sulfide -- Online Analyzer	Various have been attempted; requires specialized online analyzers (some optical, some chemical)

Cases Defined

Case 1. An MEA-Era Refinery: Low Energy and Deep Turndown

Case 1 is a large PADD 3 refinery. The refinery amine system was originally designed to use MEA in its amine system, and the refinery installed parallel capacity during the “the golden age of sulfur investment”. Prior to the COVID Bobble, the refinery also converted the heritage-MEA unit to MDEA, a move which was intended to provide both capacity and lower energy use.

Heritage MEA systems often have some differences compared to modern MDEA designs. As many MEA systems operated installed thermal reclaimers, those systems tended to avoid swings in concentration of HSAS. Converting to MDEA meant eliminating the thermal reclaiming function as MDEA does not ‘fit’ MEA-designed thermal reclaimers. Instead of installing alternative reclaiming technology, the refiner of Case 1 chose to join the many refineries that opt for periodic merchant reclaiming. As a result, the MDEA solvent system now experiences a periodic high-low cycle in HSAS.

The refinery runs the two amine systems in parallel, with common rich amine feed. The heritage-MEA regenerator has a single wash stage. The newer regenerator has 5 wash stages.

This refinery has a strong focus on energy conservation metrics – part of the reason for the conversion to an MDEA-based system. Sulfur complex operators here receive frequent guidance from energy managers who question (perceived) unnecessary energy consumption; a key variable watched is the sulfide concentration in the solvent. The refinery had the following target values for depth of stripping:

In the COVID Bobble, turndown operation was the norm; the regenerators were operating below 10% of system flood. During the data collection period, refinery H₂S production decreased in a step change. And the concentration of heat stable salts did as well. These changes are explored with the help of the ProMax[®] process simulator, and “Rules of Thumb” are tested against the results.

The refinery experienced high losses of amine during the COVID Bobble. The refinery conducted a study to look for sources of losses, and the Regenerator Reflux Purge was a usual suspect.

A minor consideration, sometimes overlooked in refinery amine units, is ammonia management. Operation at atypical condition influences ammonia management in amine units. Through simulation we examine how low energy operation affected ammonia management in Case 1’s regenerator.

From the standpoint of the regenerator, the following features are important for the Case:

Case 2. An MDEA Refinery with Heat Exchanger Performance Loss

Case 2 is another integrated refinery amine system that experienced changes in operating conditions. This refinery had operated for years with MDEA-based solvent, with accepted operating targets for energy input – their favorite was Overhead Temperature.

Over time, the lean/rich heat exchangers began to lose performance. Fouling of the heat transfer surface was suspected. A study was conducted to determine how much loss of heat transfer could be tolerated without losing system treating performance.

However, the refinery did not measure lean solvent loading, instead relying on offsite analysis by a service laboratory. As is often the case with offsite measurement, the values were not considered to be completely reliable due to the potential loss of amine.

Case 1

Rules of Thumb for Reboiler Steam


The data supporting Case 1 were measured during the COVID Bobble; sulfur rates in the refinery were already low prior to COVID and suddenly decreased further. Amine circulation decreased to <50% of design. As the regenerator was already at lower energy rates as a result of conversion to MDEA, turndown further stressed the operation: the column now operated at a very low percent of flood.

Within this turndown operating condition, the amine system confronted the composition changes w.r.t. HSAS. The COVID Bobble changed sulfur diet to the refinery as well. [Table 2](#) summarizes the swings in composition experienced during the period:

Table 2: Solvent Compositions in Case 1

		HSAS	Rich Loading
Composition 1	Start of period	Low	(Low) Normal
Composition 2	HSAS build	High	(Low) Normal
Composition 3	Low HSAS, sudden low S diet	Low	Very Low
Composition 4	Low S diet, HSAS build	High	Very Low

This refiner uses the Overhead Temperature “Rule of Thumb” to (via manual control) adjust steam supply to the reboiler. The operators then sample the lean solvent and the refinery laboratory measures the resulting concentration of H₂S. [Figure 4](#) shows simulation of that operating curve in the base case of operation before the COVID Bobble conditions. At the start of the period, the concentration of H₂S in the lean solvent responds proportionally to changes in the overhead temperature. This is the sort of operating line that makes life easier for plant operators: things change in an expected, proportional manner.



This plant has defined target values of H₂S concentration in the lean solvent. When very low values of H₂S in the lean solvent are measured, action is expected to decrease excess energy consumption in the amine

system. Not only are very low values of H₂S concentration perceived to be energy-wasteful, they are also perceived to be corrosion-risky. This refinery set a lower limit of about 50 ppmw H₂S in lean solvent.



The upper limit of H₂S in lean solvent results from achieving targeted depth of H₂S removal in the system's lower pressure contactors. That experience translates to about 500 ppmw H₂S in the lean solvent. The “goldilocks” zone for lean solvent H₂S concentration is therefore 50 – 500 ppmw. With the operating line shown in [Figure 5](#), this can be a happy situation. At the Start of Run condition, a Rule of Thumb of about 205°F delivered on-target H₂S.

However, the composition changes during the COVID Bobble shift this happy relationship between lean H₂S concentration and regenerator overhead temperature. [Figure 5](#) shows simulation results for the additional operating cases. New operating curves resulted for each composition. Sadly, one “Rule of Thumb” for Overhead Temperature does not apply here: for a given overhead temperature, widely different values of H₂S concentration in the lean solvent result. [Figure 5](#) shows that the Overhead Temperature

“Rule of Thumb” value required to achieve a targeted concentration of H₂S in the lean solvent varied substantially when the solvent conditions changed. To achieve a 300 ppmw target H₂S concentration, the Overhead Temperature varied from about 183°F to 223°F to achieve a 300 ppmw target H₂S concentration.


Different “Rules of Thumb” are often used for energy input to amine system regenerators. Each Rule has its proponents; and the merits and debits of each are well debated. One useful thing about having a process simulation tool (such as ProMax®) is that it's possible to test such merits (and debits) for one's own situation.

[Figures 6 and 7](#) show simulation results for Steam-to-Feed and Stripping Ratio Rules of Thumb (respectively). Unfortunately, as with Overhead Temperature neither of these two Rules of Thumb do a good job of handling the variability from composition changes in Case 1.




The key take-away in this: Rules of Thumb are useful for systems that do not experience solvent composition changes. For systems that undergo varying conditions, it's possible to use Rules of Thumb if one knows the operating line that one is currently on...but different targets for the Rule would be needed if conditions change.

The operating lines demonstrate yet another challenge for operators. The “HSAS Build” operating lines in the Figures demonstrate a real-world difficulty in managing H₂S concentration in lean solvent. As the operating line becomes vertical, small changes in the control variable have large responses.



An example of this challenge is shown in [Figure 8](#), which presents a hypothetical scenario based on the real-life experience of the unit engineer at the time of the COVID Bobble. HSAS had built to a high-normal concentration from the Start of Run condition. The refinery's lab measures (correct) low values of H₂S in the lean solvent. That fact came to the attention of the refinery's Energy Manager; who asks the sulfur unit staff: “why are you wasting energy?”. The operator takes action and cuts reboiler steam. The regenerator system walks down the operating line and new solvent samples are analyzed. But because the operating line is near-vertical at this point, varying values of H₂S in lean solvent are reported by the lab. It's human nature to question sudden, noisy information ... that's what happened here, too. Until small changes in process conditions also result in off-spec absorber performance.



When systems have such a ‘cliff’ in an operating line, it's usually best to avoid operating at the cliff's edge. In this example, only 5% more energy provides some distance from the cliff. When presented with such a challenge, simulation tools can help one demonstrate that Optimum and Minimum aren't always the same thing: a true optimum incorporates non-performance risks when we are able to quantify them.

Amine Loss from Reflux Purge

Another key difference between heritage MEA systems and modern MDEA designs is the premise for managing amine in the regenerator reflux. Heritage MEA systems often relied upon amine building significant concentration in the reflux loop, for corrosion avoidance. Refinery MEA systems, like any other refinery amine system, absorb ammonia when treating the refinery's ammonia-bearing hydrotreated


gases. The absorbed ammonia accumulates in the regenerator overhead, and if not removed from the regenerator system, ammonia may accumulate to problematic concentration.

In heritage MEA systems, prevention was the bulk of the ammonia strategy: up-front water washing of the hydrotreated gases removed the large majority of the ammonia, and the MEA systems dealt with only a small amount of remaining ammonia.

During the “golden age of sulfur investment”, sour water stripper capacity wasn’t always added in concert with amine and sulfur capacity. Conservation of sour water production was a common strategy to avoid building more SWS capacity – minimizing hydrotreated-gas water washing helped avoid the need. That trend meant more ammonia reached amine systems, increasing the importance of managing ammonia in the amine reflux. Case 1 experienced this ammonia trend; however the heritage-MEA system design wasn’t intended to manage such levels of ammonia. High purge rates in this design resulted in unintended amine losses; a fact that was tolerated ... until conversion of the amine system to MDEA.

MDEA is less volatile than MEA, and people expect to find low losses of amine in reflux. [Figure 9](#) shows simulation results for amine concentration in the reflux loop in the heritage-MEA design, compared to the same system operating with the MDEA-based solvent.

The regenerator has only 1 wash stage, and to demonstrate the potential for amine losses, simulation was done with no purge of reflux. MEA is found at expected concentration in the reflux, and less MDEA is present (even in this 1 wash stage system). Even though MDEA is indeed less volatile than MEA, when reflux is purged for ammonia control, the design of Case 1’s reflux loop creates potential for significant amine losses.



Reflux Loop Operation

Operation at extreme turndown during the COVID Bobble brought amine losses into focus. Amine replacement is a budget-line-item, and at times of lower cash flow, those items are natural targets for conservation. The usual potential loss mechanisms were pursued by the refinery – carryover events, filter changes, solubility losses, etc. But the heritage-MEA system operation presented an additional loss potential: excessive amine in the reflux loop.

Figure 10 is a conceptual diagram of the regenerator top + reflux loop, showing the single wash stage. The system originally operated with MEA, which required high reboiler steam rates and therefore made generous reflux flowrates. But with MDEA operation, turndown of the reflux loop required manual intervention by the operators. Steam condensate was added to baseload the flowrate of reflux through the reflux pumps, in order to satisfy their minimum-flow requirement. The flow of steam condensate was set with a manual valve and monitored with a local flow meter.

Water added into the reflux loop merges with the amine system and would ultimately fill the amine system if not managed; the operators manually purged water to the sour water stripper. The control of that purge was with a manual valve, monitored with a local FI.

The combined reflux loop flowrate was monitored at the discharge of the reflux pump with a flow transmitter. No direct measurement of *reflux return to the regenerator* exists in this design; during the period of data collection, reflux return was not known accurately. The reflux return flow was adjusted to control the water level in the reflux accumulator.

This reflux loop arrangement was a necessity to adapt to minimum flow operation, especially at turndown with MDEA as the solvent. However it is far from ideal for the stable operation of the regenerator. The reflux return flowrate naturally varies in a reflux loop, as the level control system attempts to manage the variability of the operation. When the reflux return flowrate is a large value, that variation has less effect on the regenerator. However at turndown, the reflux return variability may cause the level control valve to close at some points, starving the wash tray of liquid makeup.




Figure 10

Figures 11 and 12 depict the concept of flow variation when at low reflux return flowrate. Normal flow variation occurs, and when flowrates are high, that variation is a small percentage of the total ... and the system sees minimal effect.

When reflux return flowrates are very low, the variation in flowrate becomes significant. If a minor system upset should occur – for instance, holdup of solvent in the contactors due to a foaming event – the reflux loop sees a step change and that step change may starve the wash stage of liquid flow. In the examples of Figures 11 and 12, we show conceptually how that sort of upset condition might limit wash stage performance and result in higher, erratic concentration of amine in the reflux loop.




Figure 11

In application we have observed that type of impact on reflux return loop operation when the LCV operated at less than 15% open.

In the turndown condition of Case 1, simulation demonstrated that reflux return rates were less than 15% of design rates; therefore we expect that Case 1's regenerator overhead was vulnerable to this sort of upset condition.

Reflux Loop Simulation



Case 1 operating data encompass a wide range of conditions: data were collected for operation at high rich loading and low rich loading, both with significant concentration of HSAS. We focus on these data as they were collected when the refinery was searching for causes of amine loss.

Simulation of the wash tray in this Case comes with significant challenges.

- The tray operates at severe turndown. Mass transfer and tray hydrodynamic correlations faithfully produce results in such regions, but one must consider results with a grain of salt. Directionally, it's believable that tray effectiveness decreases at extreme turndown, but absolute prediction is tricky at low percent flood.
- Tray maloperation is **not** simulated, yet this may be a reality in the plant. Weeping or dumping conditions might exist and those are not simulated.
- Simulation is steady state. Yet the reflux return to the regenerator in this Case is dependent on the level controller; even normal variation in flow due to level control may be significant when low reflux rates are experienced.
- Process conditions are not precisely defined, due to the use of linked local flow indication for condensate makeup and reflux purge.

With those caveats we proceeded to simulate the plant data.

Figure 13 summarizes the refinery's amine-in-reflux analyses. In the figure, it appears that a trend exists: higher overhead temperature results in more amine in the reflux. It also appears that the two data sets lay on the same trend: perhaps one operating line might correlate both sets of data?



Simulation explains more about these data (Figure 13). First we simulated the reflux loop with a worst-case mindset, starting with the higher-rich loading data: what would the reflux look like, if there was no reflux purging and no wash stage? Figure 14 shows that the resulting concentration of amine in the reflux was much higher than measured values. Next we added a single wash stage to the simulation, to match the design of the regenerator. The wash stage appeared to help remove amine from the reflux at higher overhead temperature, but at the coldest operating conditions the wash stage did not decrease the amine in the reflux loop by much.

Next, we included condensate addition and reflux purge in the simulation, using values based on verbal reports from the operators about their practices with the manual control of that system. We simulated the system with and without the single wash stage. The results show that the wash stage provides some, but not much, benefit at moderate operating temperature and that benefit decreases as the temperature cools.

Simulation answered the question: do the two sets of plant data lay on the same trend? Although the eye might say yes, the conditions were different enough that the data sets do not in fact connect. Clearly, solvent conditions matter when it comes to predicting amine concentration in the reflux loop.

We were curious why the wash tray loses effectiveness at lower overhead temperature. These temperatures correspond to minimum energy input into the system, and the refinery's focus was on operating near these conditions.

The answer is perhaps easier to understand if we think of the regenerator column as actually two separate devices: from the amine feed tray downwards, the regenerator is a stripping column. But the wash stage(s) are in fact an "amine absorber". Figure 15 shows this concept. We often use the terminology "amine absorber" to describe contactors that circulate amine to scrub H_2S from gases. However, that contactor is in fact an H_2S absorber. The wash section of a regenerator column is an absorber that uses water to scrub amine vapor from the gas produced in the column below. So, the wash section is an "amine absorber", and the things that things we know about absorbers apply:

- They work best when the stages operate in the proper percent of flood.
- They work best when enough liquid is circulated to scrub effectively.

Figure 14

Plant Data and Simulation - Amine in Reflux


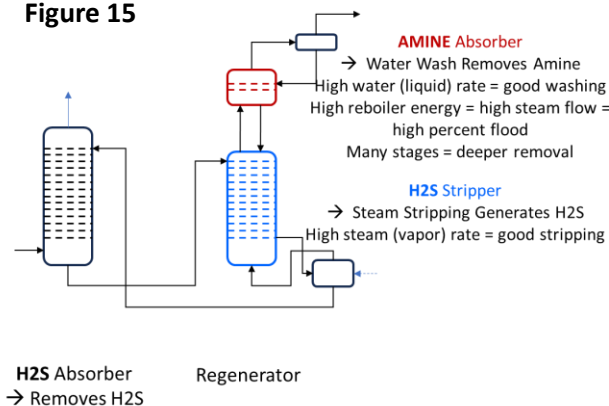



Figure 15



- And they work best when enough stages are in place to do the job.

At cold conditions, the tray operated at low percent of flood. Cold conditions also means little reflux was available to feed the wash stage. So when the overhead temperature is low, it make sense that the wash section has lower effectiveness. And with only one wash stage, limited performance is expected.

Figure 16 shows simulation results for the wash section if we were able to vary the number of the wash stages. Even at cold conditions, more wash stages results in the sort of lower amine concentration in reflux. In more modern MDEA-based unit designs, we typically encounter 2-4 wash stages. Even in the case of low liquid rates and low percent of flood at cold conditions, the “amine absorber” columns with the higher number of wash stages perform better.




This trend in number of wash stages was confirmed later, after the Case study. Even at turn down; the refinery measured low values of amine in the reflux from the new regenerator that has 5 wash stages.

Quantifying Amine Losses in the Reflux

The baseload of flow in the reflux loop through condensate addition made the system operable. But what was added, was removed as purge. We included the purge loop operation in the simulations and found that yes, the annual losses of amine in the reflux were substantial.

Figure 17 shows that these losses (for the single wash stage, operated in range of 185 – 205°F) were equivalent to the purchase of an addition 1-2 trucks per year of MDEA solvent. The shape of the curves in Figures 16 and 17 changes because the same baseload condensate flowrate was added to all the cases – consistent with how the plant operated. Dilution is more powerful at low reflux rates.



Losses, What Could Be Done Better?

First, the operability of the reflux loop could be improved. There are three functions of the reflux flowrate. The minimum flow of the reflux pumps at approximately 20% of the design flowrate determines the condensate addition rate. There are other approaches to turndown on a centrifugal pump; ideas such as a spillback loop, impeller change, variable drive controller all could be considered. With better turndown on the pumps, less condensate could be added (and therefore, less is purged). The pump impeller change

might be simplest, yet that comes with a loss of operability at the high range in the pump. Each of these interventions require physical change in the process and so are more problematic than the makeup/purge mode of operation.

The constant value of condensate added to the reflux loop often exceeded the minimum required to operate the reflux pumps. At higher overhead temperature, it would be possible to cut back on the condensate addition and minimize losses.

Some amount of purging is required for ammonia management. Simulation of the Case 1 conditions sets that minimum purge at 10% of the pump design flow. In order to maintain water balance, condensate must be added to replace water lost to the purge for ammonia rejection. As this is a heritage-MEA design, one might try to argue for improved water wash upstream of the amine unit (at the hydrotreaters which produce ammonia). To the degree that one succeeds in decreasing the ammonia that reaches the amine unit, the purge can be decreased ... decreasing the amine loss.


The above do not address the issue of variability in level control at low net flowrate through the reflux loop, and the corresponding impact on the instantaneous reflux return flow rate. Modification to the control parameters and/or scheme could deal with variability. Alternatives to decrease the risk of starving wash trays might include switching the level control to the purge valve instead of the reflux return. However in doing so, attention would still be required to the sufficiency of purge for ammonia control.

Ammonia Breakthrough at Low Energy Input

Ammonia is a headache for amine units, but one that is often ignored. The amine system serves a purpose, protecting the acid gas from buildup of ammonia. By removing ammonia from the regenerator overhead with the reflux purge, ammonia is prevented from building up to problematic concentration.

At extremely low energy rates in the regenerator, ammonia distributes down into the regenerator column, and eventually 'breaks through' into the lean amine. Simulation demonstrates that at the conditions of Case 1, ammonia breakthrough actually occurs under conditions that still produce adequately H₂S-stripped lean solvent ([Figure 18](#)).

The consequences of ammonia in the lean solvent depend on the applications in the refinery. In other refineries we have experienced distribution of ammonia into treated products, creating off-specification products at high penalty for the refiner. Case 1's refiner wasn't aware of the potential of ammonia to move around the refinery with lean amine. In general there seems to be low awareness of this risk of low energy operation.



Key Take-Aways

Considering the Big Three “Rules of Thumb”: within a single solvent condition, any Rule of Thumb applied. However, none were robust enough to guide the Case 1 refinery through the composition changes experienced.

Despite the high boiling point of MDEA that might lead one to think “there will be no MDEA in reflux”, this plant design demonstrates that it can happen. Having more wash trays is fortunately more forgiving.

When operating at deep turndown and low energy ... things are not normal. In this heritage-MEA design, deep turndown leads to amine losses due to the need to baseload flow in the reflux loop. A more modern design, with 2 or more wash trays, would perform better but still be subject to the same trends.

At extremely low reflux rates, reflux return may at times “starve” the column due to reflux drum level control variability.

At very low energy input to the regenerator, ammonia redistribution into lean amine may occur, even before H₂S stripping is compromised.

Knowledge is power. Demonstrating risks of operation at low energy can inform the conversation around energy conservation targets, and (hopefully) result in operation away from cliffs and instead at true (process-reliability-inclusive) optimum conditions.

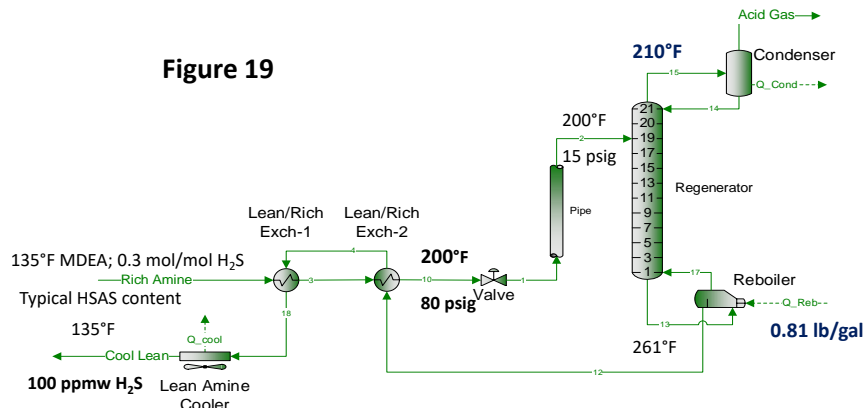
Case 2

The refinery in Case 2 experienced a loss of performance in the lean/rich heat exchanger, resulting in cooling of the rich amine. The technical staff sought answers as to how the cooling was affecting the operation of the regenerator. A study was undertaken to assess the system impact of declining heat exchange on the system, and to justify the need to clean the exchangers.

When the study was launched, heat exchanger fouling was already progressing, so the “start of run” condition represented an existing degree of decline in performance. Simulation was used to forecast the process impact at the end-point temperature expected when one of the two in-series exchangers was taken out of service for cleaning. The study also included the impact of cleaning both exchangers, resulting in significantly hotter rich amine.

The Plant

Figure 19 depicts the major elements of the regenerator system. Key to this study is that two exchangers exist in series, and that the pressure in the exchangers, and even downstream of the control valve, is sufficient to limit vapor breakout at most conditions.



The regenerator consists of 21 stages, two of which are wash stages. The typical operating conditions are shown in the figure. Note that there is a long run of pipe to feed the rich amine to the regenerator; the pressure in the pipe drops as the elevation increases, and hits 15 psig at the point of entry to the feed tray. The lower pressure near the top of the pipe is not sufficient to prevent vapor breakout at some conditions, so consideration was given to that risk in the study.

The simulation steps are summarized in Table 3. These steps are labeled to match the annotations on Figures 20A and 20B.

Table 3. Simulation Cases in Case 2


Condition	Rich Temperature	Reboiler Duty Manipulated?	Lean Solvent H2S
Base Case	200 F	Baseline Value	At Target
Step 1	170 F	Baseline Value	Allowed to float
Step 2	170 F	Changed as needed	At Target
Step 3	230 F	Baseline Value	Allow to float
Step 4	230 F	Changed as needed	At Target

Results of the Study

Figures 20A and 20B each show three operating curves for the regenerator; note that the Y-axis is a log scale. The targeted concentration of H₂S in the lean solvent was 100 ppmw; the two figures differ in the premise of reboiler energy control. Figure 20A demonstrates how the operating curves appear if the Overhead Temperature guides performance. Figure 20B does the same, except for Steam-to-Feed as the controlled variable.

The Point B on Figures 20A and 20B is the Base Case for the simulation. The operating curve of the base case demonstrates that this plant operates near a “cliff”, a point where performance significantly degrades with a lower energy input.

Point 1 represents the location on the operating curve for 170 F rich solvent, if one heat exchanger was taken out of service but no adjustment was made to the reboiler duty. Note that stripping of H₂S is quite poor.



Point 2 shows the location on the operating curve at 170 F if reboiler duty is adjusted to meet the target 100 ppmw H₂S in the lean solvent. Point 3 is on the operating curve that represents both exchangers clean and at start of run conditions, but strips more deeply than the target. Point 4 is the location that corresponds to meeting the 100 ppmw H₂S target.

The verticality of the operating lines in [Figure 20A](#) represent a difficult-to-operate condition, made worse if one were to premise high lean solvent H₂S concentration. [Figure 20B](#) also shows near-vertical lines where it would be difficult to operate.

A horizontal slice of the two figures at 100 ppmw H₂S tells us about the change in reboiler duty required to meet target. [Figure 20B](#) is more intuitive at constant circulation because the x-axis is proportional to reboiler duty. The horizontal slice shows that cooling the rich solvent requires adding more reboiler duty to reach the stripping target, but heating the solvent does not yield much benefit in decreased reboiler duty.

Operating with cold rich amine requires an increase of ~8 MMBTU/hr in reboiler duty. Fortunately for this plant, that duty is within the design of the reboiler. At low heat rates in the reboiler, when H₂S stripping is limited — the conditions create a nearly vertical operating line. Small changes in reboiler duty result in poorly-stripped solvent: Points 1 and 2 in the figure highlight the difficulty of using Overhead Temperature to guide operation at such a condition.

Once the steam required in the stripping section of the column is satisfied enough that stripping target is reached, more heat at the top of the column is of little help. Comparing Points B and 4 in [Figure 20B](#)

demonstrates that lack of impact. It may seem counterintuitive that supplying very hot rich amine to the regenerator does not result in significant decrease in required reboiler duty. The application of heat to the top of the regenerator has the most impact on condenser duty; heat applied at that point results in creation of steam above the stripping stages, where H₂S is removed from the solvent.

Figure 21 shows the breakdown of contribution to net energy input into the regenerator for the 200°F, 170°F, and 230°F simulation cases (all of which strip to the 100 ppmw target H₂S concentration). The loss of rich amine heat at the colder condition is compensated by less heat removed in the condenser and more heat supplied in the reboiler. At the hotter condition, the addition of more rich amine heat is balanced by the removal of most of that heat in the condenser.

It is important to note that a *different* horizontal slice of Figure 20B at *high lean loading*, such as with CO₂ in Carbon Capture or very lax-performing H₂S systems, would give a very different result. Those systems do not operate stripping stages in pinched conditions and so hotter rich amine would be more beneficial for reboiler duty than Case 2's refinery main treating system.

Process Considerations: Flow Regime – Rich Amine

Operating at high rich amine temperature comes with increasing risk of two-phase flow in the rich pipe. Figure 22 shows that above about 215°F, vapor breakout causes the velocity to exceed about 5-6 ft/sec. The corrosion implication of high-velocity flow in carbon steel pipe is a consideration. Two-phase flow can also create instability; Figure 23 shows that at 230°F, the flow regime is well into the slug flow regime. The vibration / mechanical considerations can be serious and would require expert evaluation.


Process Considerations: Reflux Pump

As with Case 1, in this Case the cold conditions in the regenerator overhead result in low reflux flow rates. The minimum operating condition in the reflux flow loop must be addressed.

Figure 21 Energy Distribution – 100 ppmw H₂S Cases



Figure 22



Process Considerations: Lean Cooler Impact

Although this study focused on the impact of rich amine temperature on the regenerator, note that the lean amine cooler is directly affected by any change in the lean/rich exchanger. Colder rich amine means hotter lean amine, and more duty is required in the lean amine cooler.

Key Take-Aways

Colder rich amine requires more reboiler duty to compensate. The magnitude of the impact depends on the regenerator operating conditions.

Specific to the conditions of Case 2:

- Heating up the rich amine above the base case doesn't result in much benefit for decreased reboiler duty. Less deeply stripped cases may look different.
- The operating lines for this plant demonstrate "cliffs" and those are wisely avoided.
- Operating with very high rich temperature comes with risk of vapor breakout, and increased pipe velocity with troublesome/riskier (slugging) flow regime.

Final Thoughts

This paper demonstrates that operating at extremely low energy input comes with complications that are worth considering in advance. With the emphasis on saving energy for carbon footprint and cost control, the absolute minimum energy input may look appealing until operability considerations are included. Demonstrating risks of operation at low energy can inform the conversation around energy targets, and (hopefully) result in operation away from cliffs and instead at more process-reliability-inclusive conditions.

And when operating at conditions far away from the norm, the common "Rules of Thumb" for managing reboiler energy input may not adequately guide operations. Such guidelines are a starting point to understand what is fitting for the conditions that each refinery experiences.

References

Duns, H., Jr. and Ros, N. C. J.: Vertical Flow of Gas and Liquid Mixtures in Wells, Proc. Sixth World Pet. Congress, Frankfurt (Jun. 19-26, 1963) Section II, Page 22-PD6

Addington, L. and Ness, C.: An Evaluation of General "Rules of Thumb" in Amine Sweetening Unit Design and Operation, Gas Processors Association Convention, 2010

Figure 23

Duns and Ros (Upward) at 90 °

