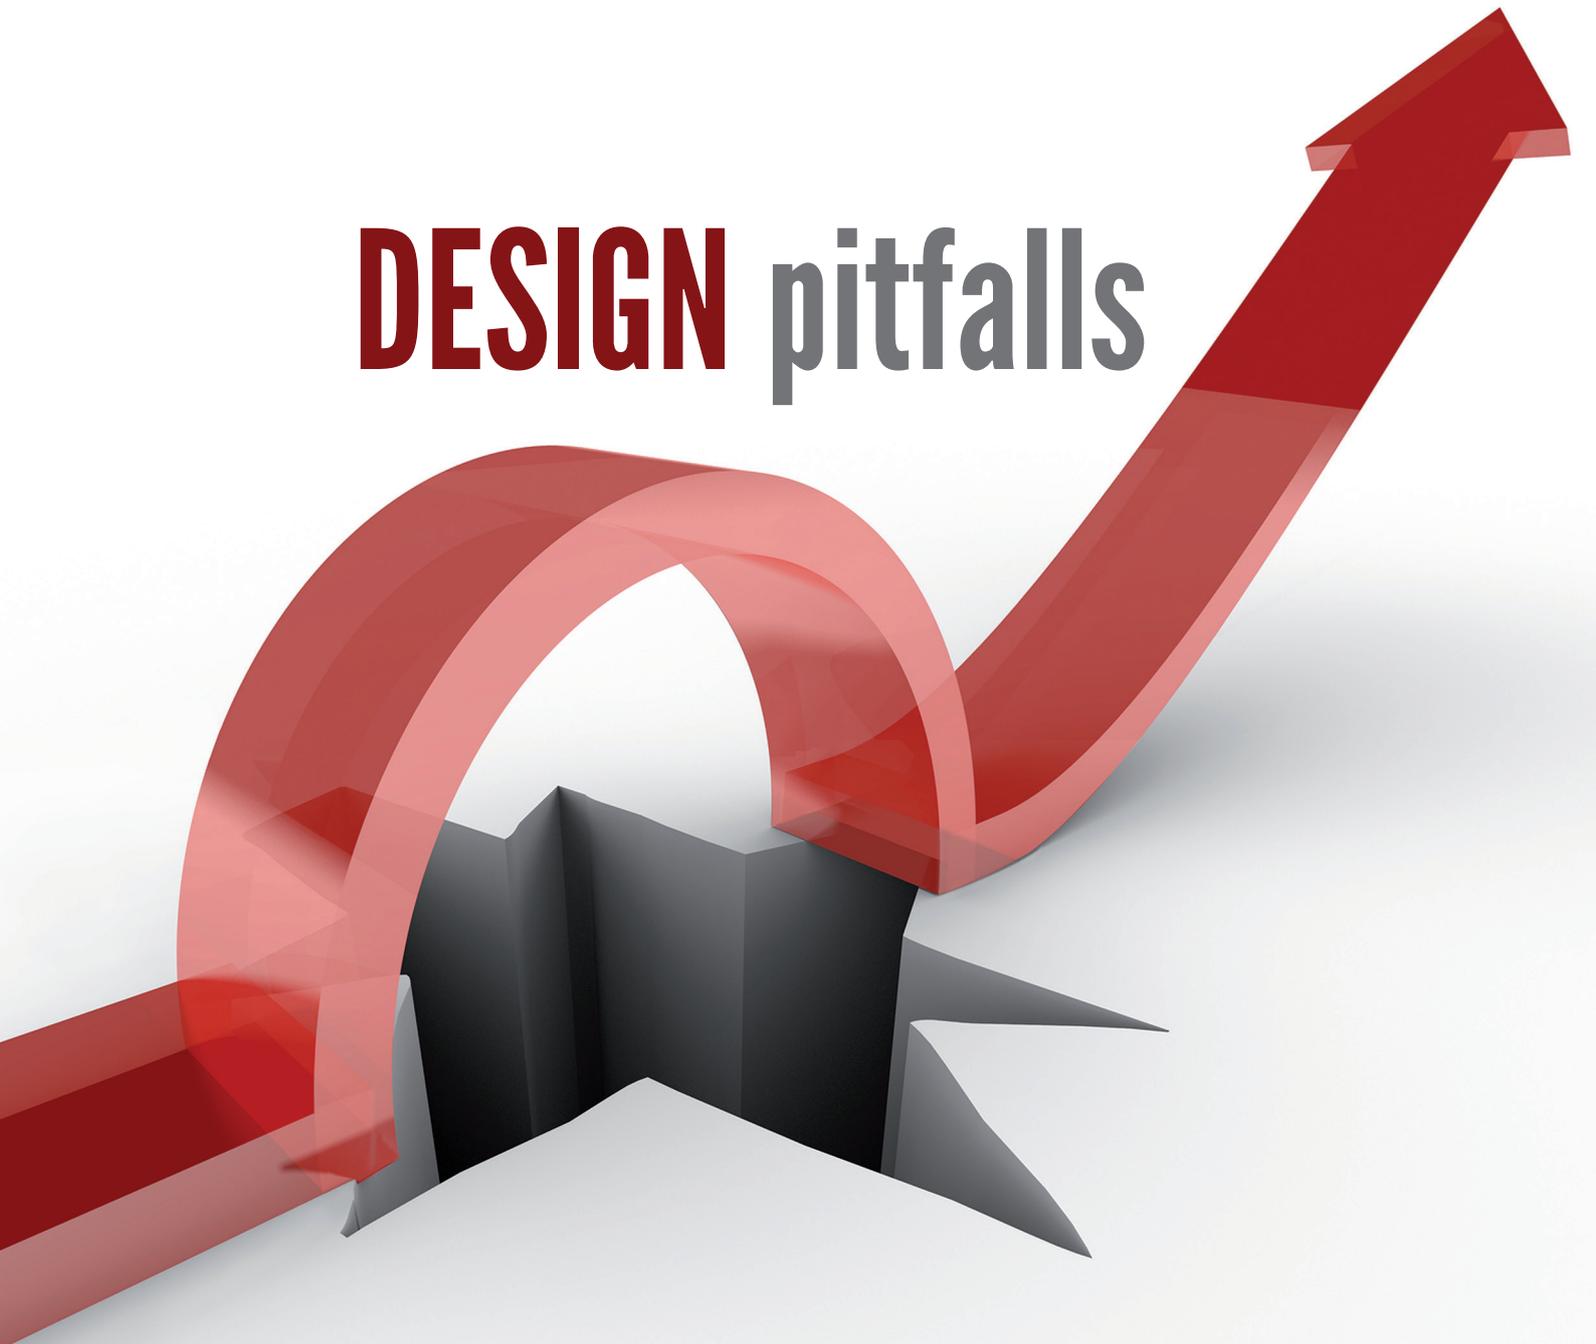


DESIGN pitfalls



Clayton E. Jones, Nathan A. Hatcher and Ralph H. Weiland, Optimized Gas Treating, Inc., USA, examine an operational pitfall involving a split-flow absorber in an LNG unit.

To prepare gas for liquefaction, its CO₂ content is typically reduced to 50 ppmv or below by absorption into an amine treating solvent. Most solvents are based on N-methyldiethanolamine (MDEA) promoted with lesser amounts of piperazine, although 2-(2-aminoethoxy) ethanol, known commercially as Huntsman's DIGLYCOLAMINE[®] agent (DGA[®]) and BASF's ADEG have found use in this application.

There are several pitfalls in designing the CO₂ removal system:

- Absorbers that use fast reacting solvents such as piperazine-promoted MDEA may be subject to instabilities should operating conditions deviate significantly from process licensor recommendations.
- Sometimes energy consumption can be greatly reduced and solvent usage optimised if the absorber is operated with fully-lean solvent flowing at a low rate to the column top, and a larger semi-lean solvent flow to a location part way down the tower. This is a split-flow arrangement that may allow the top part of the absorber to be made

considerably smaller in diameter than the bottom part, thereby saving shell and internals costs. However, there is a possible effect on mass transfer performance.

- Split-flow plants can be especially sensitive to departures from normal operating conditions, to the extent that even a slight change in a crucial process parameter can lead to failure-to-treat by a very wide margin. In these cases, small changes do not lead to the expectedly small responses in performance.

These pitfalls are by no means unique to treating gas in LNG production. For example, ammonia production where CO₂ is removed to a few hundreds of ppm is subject to the same concerns. Hydrogen production and the manufacture of various synthesis gases are other examples. The commonality, however, is the removal of CO₂ to concentrations measured in tens and hundreds of ppm using solvents having fast reaction kinetics with CO₂.

The first item in this list, column instabilities, has been addressed for both absorbers and regenerators in other articles.^{1,2} The focus of this article is on an operational pitfall involving a split-flow absorber in an LNG unit. First, some of the basic principles on which mass transfer rate-based simulation relies are explained. Then, a detailed case study is used to expose a design and operational issue that caused an otherwise excellent design to fail.

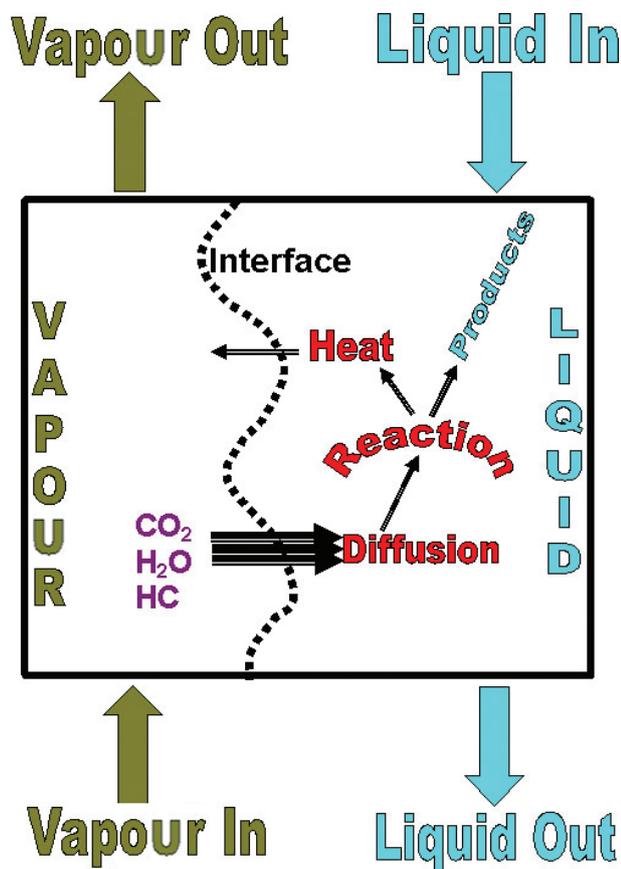


Figure 1. Schematic of mass transfer between vapour and liquid being contacted within the biphase of an absorber.

Mass transfer rate-based simulation of amine units

There are two fundamentally different methodologies for simulating any column used to separate mixtures of chemical species: ideal-stage and mass transfer rate-based. Calculations employing ideal stages have been used for about 80 years. The method is so old and well understood that it needs no exposition here. However, it suffers from a serious drawback – to connect ideal stages to real trays and tower packing, numerical values for tray efficiency and height equivalent to a theoretical plate (HETP) are required. Unfortunately, in amine treating, there is no reliable way to determine the right values, other than using good quality data from an already operating plant. Of course, for a grassroots design there simply are no data. Thus, while the ideal stage approach can be used to fit the ideal-stage model to data, it must be viewed strictly as a regression or fitting model, one that completely lacks any predictive power. Inability to reliably predict performance leads to uncertainty and oftentimes gross overdesign and this is a serious weakness.

One of the easiest ways to understand the principles of mass transfer rate-based simulation is by analogy with heat transfer. Heat exchangers have been designed this way for so long that after a first course in heat transfer, little thought is given to the fact that the methodology's basis is a heat transfer rate model. The simplest case is a shell and tube exchanger heating water with condensing steam on the shell side.

The shell side temperature is constant at the saturation temperature of the steam, but on the tube side the water temperature gradually rises along the length of the tubes. If this heat exchanger were modelled as a single ideal stage, the outlet water temperature would be in equilibrium with the steam so it would be at the steam temperature regardless of the physical size of the exchanger itself. The exchanger's efficiency in this model would be the value needed to achieve a 'match' with the real measured outlet temperature. The idealised model had to be corrected empirically, after the fact. This certainly cannot be called predictive – it is just a fit of a poor model to measured data. Could this model be used to design another exchanger? Yes, but only if the conditions were the same. Could it be used to 'predict' performance at other operating conditions (different water flow rate, different steam temperature)? Possibly, but with decreasing reliability as conditions move away from those of the original, matching case. This is patently not the best way to design heat exchangers. Yet, it is still the most commonly used method for designing much more expensive and complex mass transfer equipment.

Modern heat exchanger design is done using sophisticated software. However, the underlying principles have remained largely unchanged since long before the publication of Kern's classic work nearly 65 years ago.³ Whatever the method used for carrying out the calculations, be it hand calculation or computer programs, heat exchangers are designed using heat transfer rate models. Correlations are used that relate tube and shell side film coefficients to flows and fluid properties. Then the heat transfer rate is calculated using these coefficients along with temperature difference driving forces and heat transfer areas. The engineer provides the exchanger geometry and specifies the design conditions

(fluid flow rates and properties), and these are sufficient for the design to be done completely without engineer-supplied guesses. The engineer is never asked for an efficiency of any kind, or for artificial parameters needed to patch up model inadequacies. This is a 100% predictive model – but only in the last 25 years has mass transfer equipment design and analysis been put on the same solid basis. Ideal-stage calculations still dominate, despite the availability of vastly superior methods. The penalty is often over design to compensate for the inherent uncertainty, and sometimes outright failure.

At its core, a mass transfer rate model focuses on the rate of exchange of chemical components between two phases as driven by the extent to which temperatures and species concentrations between phases are not at equilibrium. This is the very antithesis of the ideal-stage model with its assumption of perfect equilibrium, not disequilibrium. Figure 1 shows schematically two phases being contacted in some portion of the biphase on a tray or in a segment of packing within a column. The main premise is that the mass transfer process is characterised by interfacial areas and mass transfer coefficients in the vapour and liquid phases on either side of the interface. In no sense are these parameters adjustable any more than are the film coefficients used in heat transfer. The flow of high and low viscosity fluids across trays and over packing agitates the fluids and results in specific values for mass transfer coefficients.

Engineers are not free to pick and choose coefficient values to obtain the preferred result or some preconceived notion of what the separation should be. On the contrary, values of the coefficients are a function of fluid flows, fluid properties, and the mechanical design of the specific tower internals as expressed in well-established correlations. The parallel with heat transfer is extremely close. Tower performance is predetermined and predictable in exactly the same way and to exactly the same extent as heat exchanger performance. Every tray in an operating plant is a real tray. The continuous nature of contact in a packed column is handled by dividing the height into a large number of short segments and numerically integrating across the column. Each short segment has its own unique characteristic mass transfer parameters that depend on the particular packing (vendor, brand, material, size) in each segment and on the fluids and their flow rates. Both the simulated and real separation a column achieves depends directly on what specific internals are in the column. Mass transfer rate-based simulation has the authenticity that simply cannot be achieved with theoretical stages regardless of embellishments. This is the basis for the ProTreat® simulator, the original mass transfer rate-based tool used for the simulation of a wide variety of gas treating processes.

Case study: split flow plant with limited solvent cooling

Case studies are one of the most effective ways of understanding some of the pitfalls in design. This particular situation involved the CO₂ removal section of an LNG plant using piperazine-promoted MDEA to remove CO₂ from predominantly methane to below 50 ppmv. Focus is on the absorption side of the process. Although it could have been built as a single tower with a mid-tower solvent flow of

semi-lean (partially stripped) amine, the absorption section was actually built as two towers in series, as shown in Figure 2.

Rich solvent flows through a cross-exchanger to the regenerator where it is contacted with steam stripping vapour generated in the regenerator's reboiler. For simplicity, the regeneration section is not shown. The largest part of the solvent is withdrawn after flowing only part way down the regenerator. It is used in the absorber labelled 'bulk removal'. This solvent has undergone only limited stripping so its CO₂ content is still fairly high. However, really well stripped solvent is unnecessary for bulk CO₂ removal where the objective is to remove most of the CO₂ without achieving anything like the <50 ppmv specification. The 50 ppmv specification is met by the polishing absorber downstream where a relatively small flow of well stripped solvent is all that is needed to remove the remaining CO₂ and meet the final CO₂ specification.

The 4 m dia. polishing column in this case contained two 5.5 m beds of 25 - 5 HIFLOW rings while two 1-pass bubble cap trays with 50 mm weirs served as wash trays to remove entrained and vaporised amine from the final treated gas. Wash water usage was 4055 kg/hr at a temperature of 50°C.

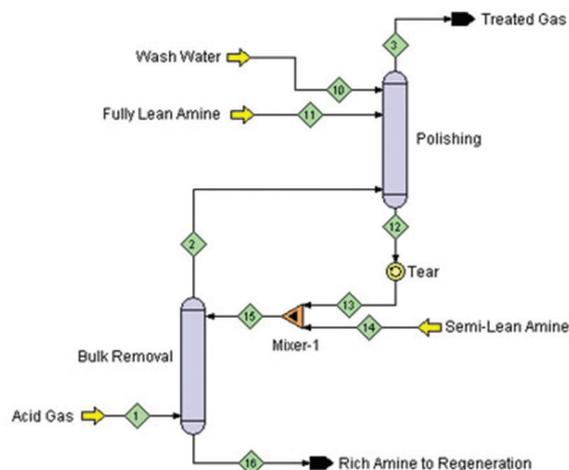


Figure 2. Process flow diagram of two CO₂ absorbers connected in a split-flow configuration.

Table 1. Inlet stream flows and compositions				
	Raw gas	Wash water	Fully lean	Semi-lean
Temperature (°C)	78.2	50	48	Variable
Pressure (bar(g))	45	45	45	45
Flow (kg/h)	328 000	4055	768 000	4 477 000
Composition				
Water	Saturated	100%	-	-
CO ₂ (mol% or mol/mol)	17.5%	0	0.021 m/m	0.388 m/m
MDEA (wt%)	0	0	37	37
Piperazine (wt%)	0	0	3	3
Methane (mol%)	82.5	0	0	0

The bulk removal absorber was 6 m dia. with two 5.5 m deep beds of 50 - 5 HIFLOW rings. Table 1 shows the flows, temperatures, pressures and compositions of the four inlet streams. CO₂ loading is in moles of CO₂ per mole of total amine (MDEA + piperazine).

The plant to which these data refer was unable to meet the < 50 ppmv CO₂ specification by several thousands of ppm. An undersized semi-lean cooler combined with very hot ambient air (used as the coolant) resulted in a minimum semi-lean temperature of about 80°C, some 10°C higher than the process licensors' recommended maximum of 70°C. A simulation study of this plant was carried out using the ProTreat[®] mass transfer rate-based simulator. As the simulation results of Figure 3 show, there is a critical temperature of about 76.8°C above which treating is predicted to fail precipitously. At that temperature the bulk removal absorber becomes thoroughly rich-end pinched to the extent that it is fully saturated with dissolved CO₂. The approach to equilibrium at the rich end of the absorber reaches virtually 100%. Any further increase in semi-lean solvent temperature reduces the solvent's capacity and the

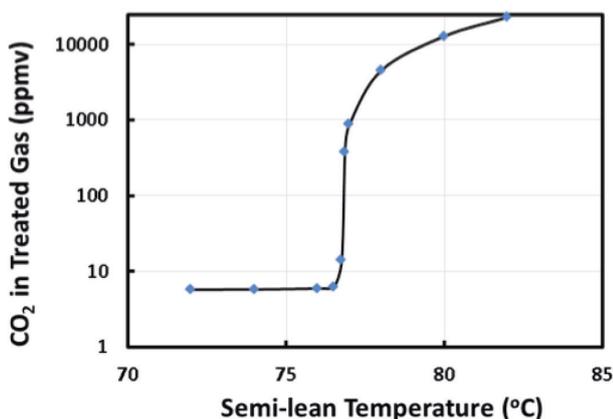


Figure 3. Effect of semi-lean temperature on treating in a split-flow CO₂ absorber system.

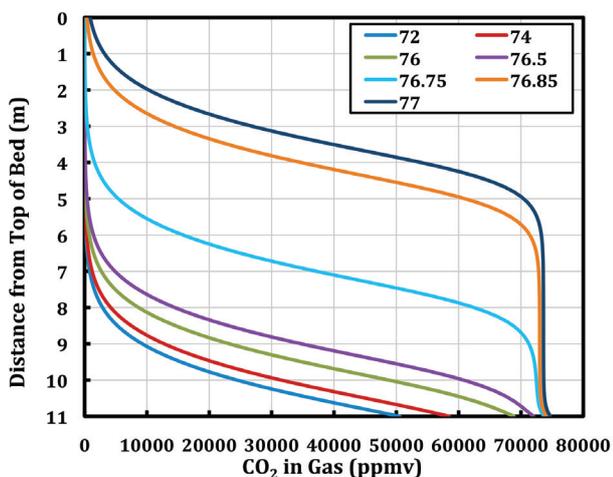


Figure 4. Profiles of CO₂ concentration in the polishing column as a function of semi-lean amine temperature.

excess CO₂ that cannot be absorbed spills over into the polishing column. Unfortunately, the spill-over is of such an amount that the polishing column cannot cope with the extra CO₂, and the treated gas quickly fails to meet specifications.

The reason for the failure is presented in Figure 4. It is evident that as the semi-lean temperature rises, the CO₂ profile in the polishing column goes from a typical lean-end pinched shape (72, 74, 76°C), where most absorption occurs close to the column base, to a typical rich-end pinched form at 76.85 and 77°C. At 76.75°C, the CO₂ profile is changing position in the column extremely rapidly. In fact, there were times in the plant when the absorption system rapidly went in and out of compliance with the treating specification, caused by favourable fluctuations in other process parameters such as gas composition, flow rates, etc.

In this case, the process operators initially felt there had to be something seriously wrong with the column internals for the plant to be unable to produce a gas any better than several thousands of ppmv CO₂ compared with the 50 ppmv design target. From a process standpoint, the cause turned out to be a lot more subtle. One does not normally expect a temperature change of a few degrees to cause a huge performance loss. Indeed, when the semi-lean temperature is low enough, changes of a few degrees have a vanishingly small effect on overall performance. This is reaffirmed by the flatness of the curve to the left of 76°C in Figure 3 and by the very small response of the CO₂ profile to temperature changes below 76°C. However, there is an extremely rapid change from lean-end to rich-end pinch conditions resulting from a seemingly small temperature difference near the cut-off temperature. This is revealed by the mass transfer rate-based simulations.

Lean and semi-lean amine temperatures were already as low as achievable in the plant, and solvent circulation pumps were already operating at capacity so that, short of a larger semi-lean exchanger, treating could be brought back on specification only by reducing the gas flow (production).

Conclusion

This case demonstrates the power of mass transfer rate-based simulation in diagnosing plant operating problems. What was initially attributed to malfunctioning tower internals turned out to be caused by inherent process sensitivity to an (unachievable) operating temperature. A simulation based on mass transfer rates can predict unanticipated phenomena such as this, but a simulation based on efficiency or HETP could not reveal this critical aspect of the process. To create confident and successful designs and to troubleshoot existing processes efficiently, it is vital to know which parameters control performance in different parts of the process. **LNG**

References

1. Hatcher, N. A., and Weiland, R. H., 'Amine Regenerator Instabilities', *Hydrocarbon Engineering*, October 2012 pp. 27 - 30.
2. Weiland, R. H., and Hatcher, N. A., 'Foundations of Failure', *Hydrocarbon Engineering*, December 2011, pp. 57 - 39.
3. Kern, D. Q., 'Process Heat Transfer', *McGraw-Hill Book Co.*: New York, 1950.