

CO₂ Unit Performance: When Actual Differs Radically from Design¹

Steven M. Fulk, Clayton E. Jones, Ralph H. Weiland
Optimized Gas Treating, Inc., Buda, TX, USA

In a case study of a Middle Eastern amine treating unit producing LNG quality gas with piperazine promoted MDEA, the ProTreat® mass and heat transfer rate based simulator is used to show that process responses to changing operating conditions can be counterintuitive, and occasionally our expectations can be quite wrong. Sometimes lower CO₂ in the raw gas does *not* make treating easier, the resulting very lightly loaded rich solvent may *not* lead to less heat load on the regenerator, and using more efficient heat integration can potentially ruin the operation altogether.

The most frequently used solvent for carbon dioxide removal in LNG production is aqueous *N*-methyl-diethanolamine (MDEA) promoted with a relatively small concentration of piperazine. MDEA has high capacity for CO₂ while, being nonreactive, needs considerably less regeneration energy than its carbamate-forming cousins. However, lack of reactivity makes absorption rates painfully slow. On the other hand, piperazine, a secondary diamine, is very highly reactive towards CO₂. By adding piperazine at modest concentrations to MDEA, the absorption rate is greatly increased without seriously burdening regeneration energy requirements.

In the traditional processing scheme, the absorber and solvent regenerator are heat integrated by direct coupling via a simple cross exchanger. The process includes sundry other equipment items such as pumps, valves, and fin-fan coolers for final lean solvent temperature control. Upwards of 80% or more of the operating cost of a CO₂ removal unit is the energy required for solvent regeneration; therefore, almost anything that can be done to improve heat integration and more efficiently use energy is worthwhile.

A more energy efficient alternative to the conventional process structure is to locate a low-pressure flash tower (LP-Flash) immediately downstream of the absorber (Figure 1). In this scheme, heat integration is achieved by heat exchange between the semi-lean and fully-lean amine streams. The actual regeneration column then is fed by a hot solvent stream already partially stripped of its carbon dioxide content. Thus, solvent regeneration is a two-step process: an LP-Flash Column followed by a regenerator.

To meet a 50 ppmv CO₂ specification with a high pressure absorber does not require the solvent to be stripped to a really low level. A loading value of 0.01 to 0.04 moles CO₂ per mole of total amine is usually adequate. Regardless of the flow configuration, this can be achieved with most of the trays or packing in the regenerator running relatively cold, roughly equal to the temperature of the solvent feed.

¹ LNG Industry, April, 2018

The lion's share of stripping takes place close to the bottom of the regenerator and in the reboiler. This, of course, requires stainless metallurgy, but the additional CAPEX is well worth the savings in OPEX.

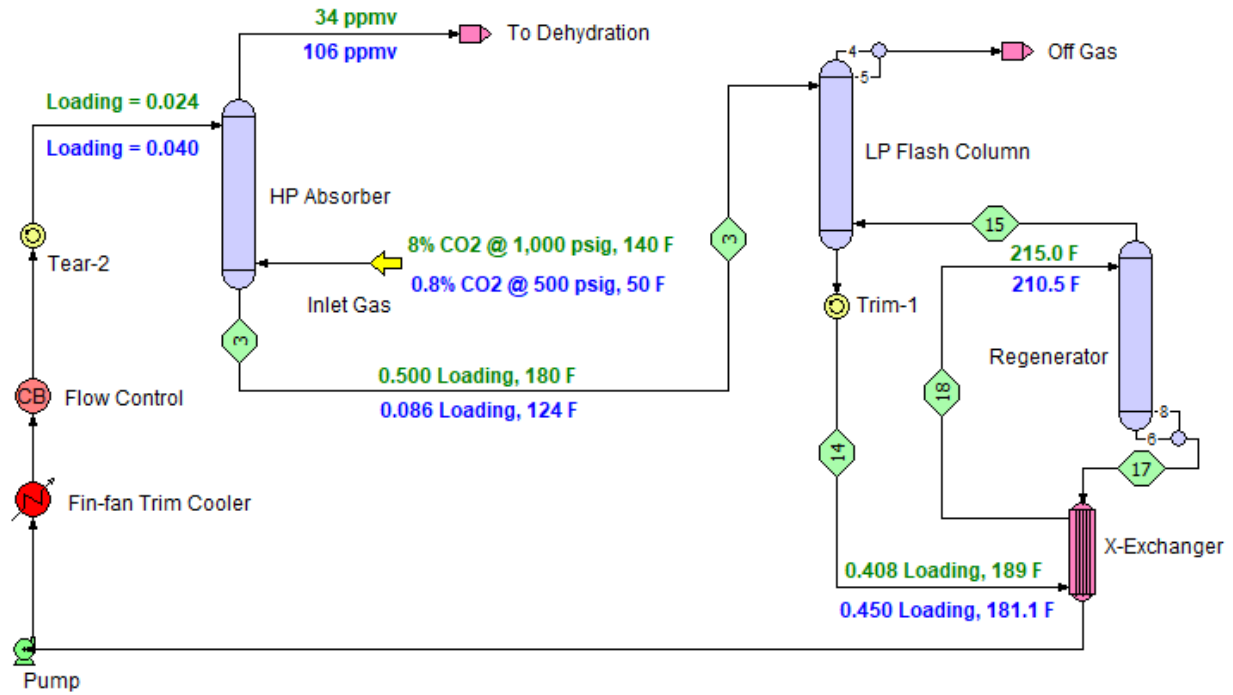


Figure 1 Process Scheme with LP-Flash Column and Regenerator. Annotations in Green are Design; Blue are Operating

Ignoring for the moment the solvent loading and temperature annotations in Figure 1, the energy (in the form of steam) in the vapour leaving the regenerator (Stream 15) has already been expended and it cannot be used to strip additional CO₂ from the solvent even if the solvent bypasses the LP-Flash Column and Stream 3 enters directly into the regenerator. However, potentially it can have its energy content increased via cross exchange with hot semi-lean feed (Stream 18). This increases the temperature of Stream 15 and, via partial vapourisation of solvent, replenishes its steam content for further use in the LP-Flash column. This rationale applies especially when the regenerator operates with a collapsed steam flow, i.e., when most of the regenerator is operating relatively cold, the vapour has already cooled off, and stripping is confined to the bottom part of the column and its reboiler.

The CO₂ Unit as Designed

The specific carbon dioxide removal system being discussed was originally designed and built to produce <50 ppmv CO₂ treated gas from an inlet gas expected to contain 8% CO₂ at a pressure of about 1,000 psig (69 barg) using a solvent with 35 wt% MDEA and 2 wt% piperazine. The rich amine is regenerated in two steps: first the LP-Flash Column and then the Regenerator. The annotations in green in Figure 1 show simulated parameter values for the design case. The absorber (30 feet of No. 1.5 Super

Raschig Rings®) produces gas having 34 ppmv CO₂, well below the maximum of 50 ppmv. The rich amine loading drops from 0.500 to 0.408 in the LP-Flash Column, a 15-ft (4.57-m) bed of No. 1.5 Raschig Super-Rings®. The loading drops to about 0.43 immediately upon entry into the column because the rich feed stream flashes from 90 psig to 14 psig as it enters. Further contact within the packing decreases the loading from 0.43 to 0.408 mol/mol. Finally, the rich loading drops from 0.408 to 0.024 in the regenerator, a 25-ft (7.62-m) bed of No. 1.5 Raschig Super-Rings® with kettle-type reboiler.

LP-Flash Column

The solid green line in Figure 2 shows the solvent CO₂ loading profile across the LP-Flash column for this 8% CO₂ design case. About 20% of the stripping is done in the LP-Flash Column. The loading profile is fairly flat across the column except for the sudden reduction (0.500 to 0.430) from flashing upon column entry due to the reduction in pressure from 90 psig to 14 psig (6 to 1 barg), and a further loading reduction from 0.43 to 0.408 close to the bottom of the column as the solvent contacts the entering stripping steam. Stripping at the bottom of the column is minimal owing to the low energy content and quantity of stripping steam.

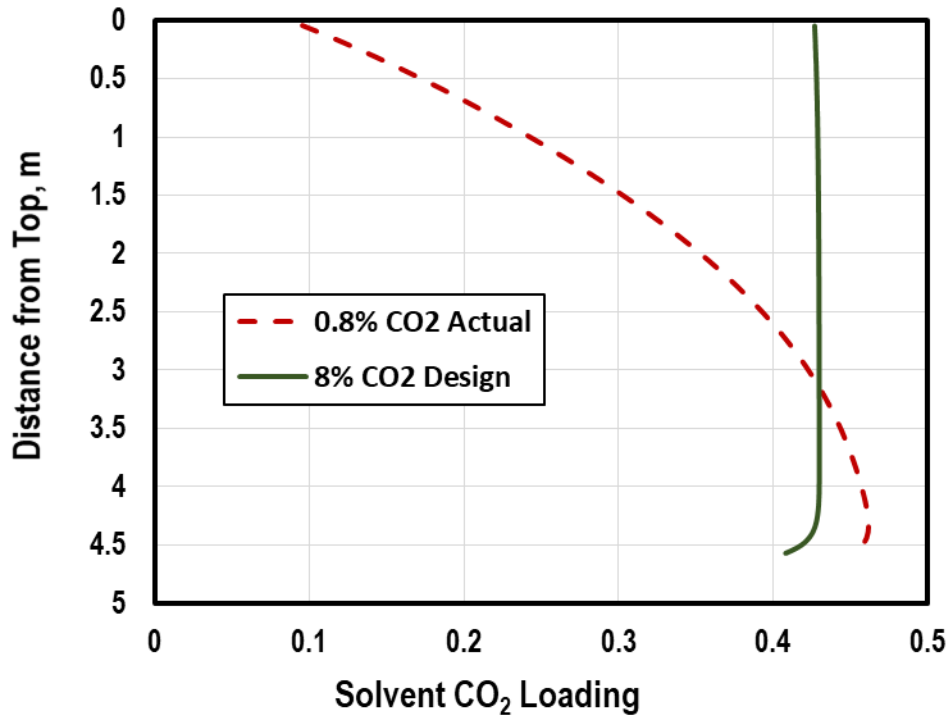


Figure 2 CO₂ Loading Profiles within the LP-Flash Column

Regenerator

A 73,000-ft² (6,780-m²) cross exchanger is needed to raise the semi-lean solvent (Stream 14) to 215°F (101.7°C) and this allows the regenerator to produce a lean solvent loaded to 0.024 mole CO₂ per mole total amine. The solid green line in Figure 3 shows the loading profile across the regenerator for the design case. The regenerator operates with collapsed temperature and steam flow profiles over the upper part of the column because the reboiler steam flow is being controlled to keep the lean solvent loading just

sufficient to achieve slightly below 50 ppmv CO₂ in the treated gas. This is reminiscent of the way the regenerator in a CO₂ capture plant would be operated where the goal is not to reach super low CO₂ levels in the treated gas, but rather to limit absorption to about 90% CO₂ removal (10% CO₂ slip)—in both cases, removal is limited by limiting reboiler duty. Treating for LNG production may not be quite as striking in this regard as it is in CO₂ capture, but in any event, treating is controlled to 50 ppmv by limiting the reboiler steam supply.

The process metrics and loading profiles shown and discussed here are typical of a well-designed CO₂ removal system for efficient operation in LNG production. But what happened in the actual operating case?

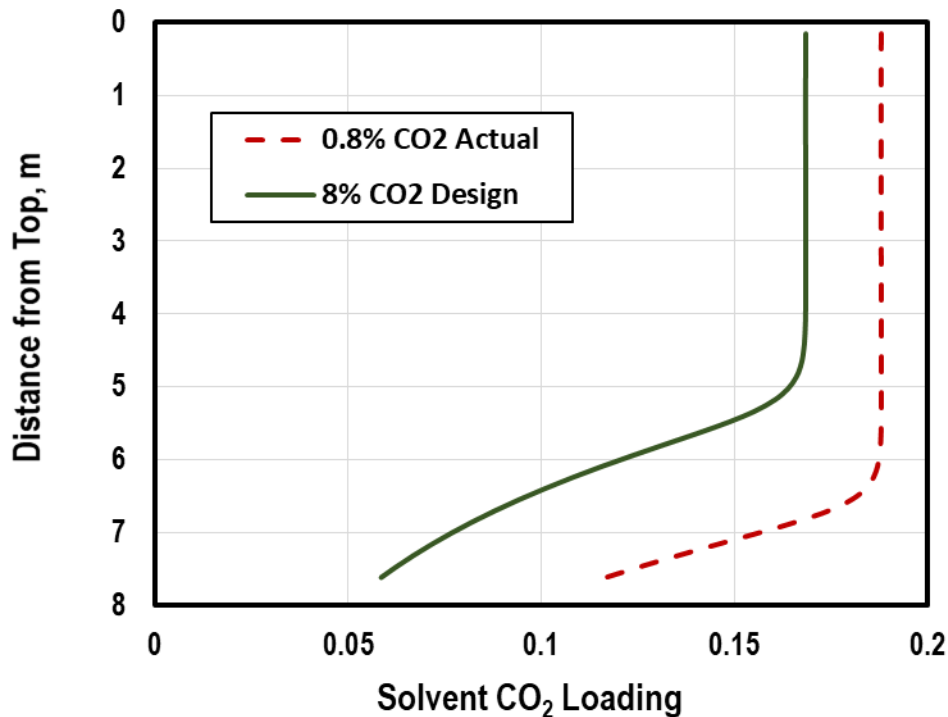


Figure 3 CO₂ Loading Profiles within the Regenerator

The CO₂ Unit as Operated

When a production unit is first started up it is often presented with inlet gas that is quite different from what it was designed to treat. That was certainly the case here where the first wells brought on stream were at half the expected pressure; the gas contained only a tenth the expected CO₂ (0.8% versus nominally 8%) and was considerably colder.

As wells are brought in and out of production, the combined feed gas from the operating wells might (and usually will) vary substantially. In this case however, 8% CO₂ was the expected worst case, i.e. the highest anticipated CO₂ content. A reasonable person would expect that reducing a gas stream's CO₂ content by 90% would make it easier to treat. However, using the same solvent flowing at the same rate as for the 8% gas design basis, this plant in fact failed to produce LNG quality gas at all! This was a real

shock—the treated gas was only 106 ppmv CO₂. To meet the < 50 ppmv specification the reboiler steam flow had to be somewhat increased; whereas, operations expected to be able to *decrease* it very considerably!

In both cases, the absorber was lean-end pinched. This means that the treating level was completely controlled only by *lean solvent loading* and *temperature*. But in this case, the temperature in the upper part of the column was the same in both cases. Thus *solvent lean loading* was left as the *sole* determinant of absorber performance. In terms of overall carbon dioxide removal plant performance, this is the single overriding parameter, albeit one that is affected by flowsheet configuration and how other equipment items behave.

LP-Flash Column

An initially puzzling observation was that the LP-Flash Column was in effect operating as an excellent CO₂ *absorber*, taking a very lightly loaded (0.086 mol/mol) rich solvent from the absorber itself and producing a “semi-stripped” solvent that was loaded all the way up to 0.500 mol/mol — in fact, a *much higher loading than the design case* which was 0.408! The flash column turns into an absorber because the solvent feeding it is 56°F colder than expected; a combined effect of the colder inlet gas and less total heat of absorption from CO₂. The energy in the entering vapour (Stream 15) is being used to heat this cold solvent rather than to strip CO₂. In addition, the cold solvent is exposed to a high CO₂ content vapour stream. Thus, absorption ensues, and the dashed red line in Figure 2 shows the solvent loading *increasing* as it flows down the LP-Flash Column.

Regenerator

Before entering the regenerator, rich solvent is heated by cross exchange against the hot lean solvent (Stream 17) produced by the regenerator. However, the loaded solvent entering the cross-exchanger is some 8°F (4.5°C) colder than design so the limited area of the equipment is able to preheat the feed to only 210.5°F (99°C) rather than the design value of 215°F (101.7°C). The consequence for the regenerator is an even lower temperature that persists through even more of the regenerator. It results in the loading profile shown in dashed red in Figure 3. The regenerator feed is not only colder but it's also richer, and the column simply cannot cope with these adverse conditions—stripping rates are reduced, producing a final lean loading of only 0.040 mol/mol versus 0.024 in the design case. Lean solvent loading completely controls absorber performance and the now-higher lean loading causes the absorber to miss the treating specification.

What Was Expected and What Actually Happened

Rules of thumb garnered from experience suggest that removing only 1/10th of the CO₂ should have been easy to do even at 50% lower pressure with a colder lean solvent. In fact, it should have been able to be done at reduced steam flow into the reboiler, too. But even at full steam flow, the 0.8% CO₂ gas was not reduced even to the target maximum of 50 ppmv. Everything ran completely counter to all expectations. Although some details have been altered to conceal the identity of the plant and its owner, the case is not contrived; these observations are very close to what was actually experienced in a Middle Eastern LNG plant. And when it was seen, it was met with total disbelief. The ProTreat® mass transfer rate column model, however, removes the bias of expectations and separates what feels right from what actually is.

To achieve the same level of treating as the 8% CO₂ case, the reboiler duty has to be slightly increased from 105 MMBTU/h to 110 MMBTU/h. However, except for the CO₂ loading of the lean solvent,

the state of all process streams remains virtually unchanged as does the performance of the LP-Flash and Regenerator columns. The LP-Flash continues to operate as an absorber rather than a stripper and the Regenerator still shows a highly collapsed temperature profile, similar to the solid green line in Figure 3.

This unit's seemingly weird behaviour stems from the very low temperature of the rich solvent produced by the absorber, caused in part by the cold temperature of the inlet gas, and in part by the minor heating from the small amount of CO₂ absorbed. The effect of low rich solvent temperature cascades through to the LP-Flash causing it to absorb rather than strip, then it propagates through to the Regenerator as a more highly loaded but cooler solvent and causes the regenerator temperature to collapse even further. The combined LP-Flash and Regenerator columns show a huge internal recycle of carbon dioxide moving forward with the solvent and backward with the vapor. Processes that are highly heat-integrated and which depend on heat of reaction as part of their total heat input are susceptible to this kind of problem.

The primary purpose of this communication is not to suggest ways to regain the desired treating level; rather, it is to expose a case of seemingly abnormal and counterintuitive behaviour and to show how conventional thinking fails to explain it. In this particular example, the ProTreat® simulator's strict mass and heat transfer rate basis is almost essential to achieve a solid quantitative understanding of such behaviours and dependencies. It bridges the chasm between what we expect, and what engineering science says is true.