



Separating the inseparable

What reverse osmosis and cryogenic distillation teach each other

Sepehr Process Development Engineering Company · Process Engineering

ABSTRACT

Mixing happens on its own; separating never does. Thermodynamics fixes a minimum amount of work for pulling a mixture apart, and industry has converged on two very different ways of paying it: hold the mixture still and let one component pass a barrier, or change its phase and let the mixture sort itself. This note compares reverse osmosis and cryogenic air separation — their driving forces, their limits, and the three constraints they turn out to share.

1.06 kWh/m³

Thermodynamic minimum for seawater at 50 % recovery

27 bar


Osmotic pressure of seawater at 35 g/L TDS

12.8 °C


Boiling-point window between nitrogen and oxygen

FIGURE 1

Osmosis



Reverse osmosis




Two states of the same cell. Left: with no applied pressure, water crosses the membrane toward the saline side until the liquid head balances the osmotic pressure. Right: applied pressure above the osmotic pressure reverses the flow; water passes, ions do not.

1 Mixing is free. Separating is not.

Pour a spoonful of salt into a glass of water and it dissolves without help from anyone. Open a valve between a tank of nitrogen and a tank of oxygen and, given time, you get air. Nature performs both of these for nothing, because mixing increases entropy and the universe is generous with processes that do.

Running either one backwards costs work, and thermodynamics sets a floor on how little work will do. For seawater at 35 g/L, recovering half of it as fresh water requires at least about 1.06 kWh per cubic metre of product — an absolute minimum, at reversible conditions, with perfect equipment that does not exist. Real seawater RO plants land somewhere around three to four kWh/m³ once pretreatment, pumps and losses are counted. The gap between those two numbers is the entire engineering discipline.

Figure 2



The cost of the gap. Thermodynamic minimum against typical practice for seawater reverse osmosis. Everything above the first bar is irreversibility: friction, mixing, and the finite driving forces real equipment needs.

Industry has converged on two very different strategies for paying that bill. Understanding why each one exists — and where each one breaks down — is more useful than knowing either in isolation.

2 Strategy one: build a barrier

Reverse osmosis takes the direct approach. Place a semipermeable membrane between salt water and fresh water and, left alone, water flows *toward* the salt. It keeps flowing until the accumulated column of liquid pushes back hard enough to stop it. That equilibrium pressure is the osmotic pressure of the solution: roughly 27 bar for seawater, and it scales with the concentration of dissolved species.

Apply more than that on the saline side and the flow reverses. This is where the name comes from, and

why seawater plants operate at 55 to 70 bar — osmotic pressure is the toll, and only the pressure above it produces flux.

What happens inside the membrane is not filtration. The active layer of a polyamide composite membrane is on the order of 0.1 µm thick and has no pores in the ordinary sense. Water dissolves into the polymer matrix, diffuses across it under the pressure gradient, and desorbs on the other side. Salts, being charged and strongly hydrated, are far less soluble in that matrix and are left behind. This solution-diffusion mechanism is why RO removes ions that no mechanical filter could ever catch, and also why the membrane is so sensitive: anything that coats the surface interferes with the very first step.

FOUR CONSEQUENCES FOR PLANT DESIGN

Concentration polarization. Water leaves through the membrane; the salt it was carrying does not. Salts accumulate in a thin boundary layer at the surface, so the local osmotic pressure is always higher than the bulk feed analysis suggests. Crossflow velocity and feed-spacer geometry exist mainly to keep that layer thin.

Recovery has a chemical ceiling. As recovery rises the concentrate gets saltier, its osmotic pressure climbs, and sparingly soluble salts — calcium carbonate and sulfate, barium sulfate, silica — approach saturation in the last elements. Brackish systems commonly run at 70–85 % recovery; seawater systems at 35–50 %, because concentrate chemistry binds long before the equipment does.

Temperature is not a detail. Permeate flux changes on the order of 2–3 % per °C. A plant that comfortably meets design capacity in summer can fall short in winter on identical feed water.


Energy leaves with the reject. Concentrate exits the last element still at high pressure. Recovering it — with a pressure exchanger or turbocharger — is the single change that transformed seawater desalination economics.

3 Strategy two: change the phase

Air cannot be separated with a barrier at industrial scale. Nitrogen and oxygen molecules are nearly the same size, carry no charge, and interact with almost any material in nearly the same way. There is no practical membrane that says yes to one and no to the other with the selectivity a merchant gas plant needs.

So cryogenic air separation uses the only property that differs meaningfully: the temperature at which each component condenses. Nitrogen boils at -195.8°C , argon at -185.9°C , oxygen at -183.0°C . Cool air until it liquefies, then boil it back off carefully, and the nitrogen leaves first while the oxygen stays behind in the liquid. That 12.8°C window is the entire physical basis of the process.

Figure 3




The whole margin. Normal boiling points of the three main components of air. Argon sits between the other two, which is precisely why it is the hardest to recover.

A narrow window has three consequences that shape every air separation unit ever built.

It takes a great many stages. A small boiling-point difference means low relative volatility, and low relative volatility means the separation has to be repeated many times over. Modern columns use structured packing to provide that stage count at very low pressure drop — which matters enormously, because pressure drop in the low-pressure column translates directly into compressor power at the front of the plant.

Figure 4



One coupled system. The condenser at the top of the high-pressure column is simultaneously the reboiler of the low-pressure column. A single exchanger serves both duties.

The energy balance is closed and tight. Essentially all of the refrigeration is produced and recovered inside the process itself: incoming air is cooled against outgoing cold products in the main heat exchanger, and the double column is arranged so that the condenser of the high-pressure column is the reboiler of the low-pressure column. Every watt of heat that leaks into the cold box has to be paid for at the compressor, which is why the whole assembly sits inside an insulated, perlite-filled enclosure.

Anything that freezes must be removed first. Water, carbon dioxide and hydrocarbons are taken out upstream by molecular sieve adsorbers before the air ever reaches the cold end. This is not housekeeping. In the cold box these components would freeze and plug passages — and hydrocarbons would concentrate in the liquid oxygen at the bottom of the column, which is the classic hazard of the process. Front-end purification is a safety system that happens to look like a process step.

4 What the two have in common

The last increment costs the most. In reverse osmosis it is the final few percent of recovery, where osmotic pressure and scaling limits rise together. In an air separation unit it is the last decimal of purity, and argon, which needs its own column and the steadiest operation of the whole plant. In both cases the cost curve turns sharply upward near the end, and a specification written one decimal tighter than the process actually requires can be one of the most expensive sentences in a project.

Protecting the separation step is most of the engineering. The membrane and the column each perform an elegant physical trick, and each is defenceless without what sits in front of it. Media filtration, coagulation, cartridge filters and antiscalant exist so the membrane can do its job; molecular sieves and aftercoolers exist so the column can do its. In both plants the majority of the equipment count — and the majority of the operating problems — lives upstream of the step everyone names the plant after.

Measure the right thing, and normalize it. RO performance is only meaningful as normalized permeate flow and normalized salt passage, corrected for temperature, feed salinity and recovery; a raw gauge reading tells you the weather, not the condition of the membranes. An air separation unit is judged on specific power and purity together, because either one alone can be made to look excellent at the expense of the other.

5 Why this matters on paper

Both processes reward decisions made early and punish those deferred. Recovery, staging and pretreatment are chosen in basic design and are expensive to revisit after the concrete is poured. Column pressure drop, heat exchanger approach temperatures and the purification cycle are settled long before anyone orders steel.

That is the part of engineering we find most interesting: not the equipment itself, but the small number of decisions made early that determine what a plant will be capable of for the next twenty years.

ABOUT

Sepehr Process Development Engineering Company works across process engineering, electrical and instrumentation, and planning and ICT — for water and wastewater treatment, air separation units, cryogenic storage tanks, and oil, gas and petrochemical projects, from feasibility through commissioning and optimization.

6 Nomenclature

π — osmotic pressure: the pressure that exactly stops spontaneous flow across a semipermeable membrane.

Recovery — permeate flow divided by feed flow, as a percentage.

Flux (LMH) — permeate flow per unit membrane area, in litres per square metre per hour.

SDI — silt density index; a filtration test that judges how likely a feed water is to foul a membrane.

Relative volatility — the ratio measuring how easily two components separate by distillation. Values close to one demand many stages.

TSA — temperature swing adsorption; the molecular sieve cycle that removes water, CO₂ and hydrocarbons ahead of the cold box.

GOX / LOX / GAN / LIN / LAR — gaseous and liquid oxygen, gaseous and liquid nitrogen, liquid argon.

A QUESTION FOR READERS

If you work on either side of this — membranes or cryogenics — which constraint binds first in your projects: energy, feed chemistry, or the specification itself?