HW1 - Mixers and Splitters

Course: CBE 253 - Aspen Plus Module
Sessions: 1-3 - Weeks 10-11 - Mar 24 - Mar 31
Assignment: HW1 - Methanol Stream Separation and Mixing
Due: March 31
Primary Mindset: Constraint
Connection forward: The IDEAL property method used here for methanol-water is a planted question. HW6 will generate T-xy diagrams for this same system using NRTL and Wilson - and the diagram will look nothing like what IDEAL predicts. The error you accepted here becomes the lesson there.

Simulation Context

A high-purity aqueous methanol feed (98.5 mol% methanol, 150 lbm/hr, 130 F, 25 psia) is split using an FSPLIT block: 40 wt% routes to Mixer M200 where it combines with 30 kg/hr water at 15 C and 1 atm to form Product A; the remaining 60 wt% routes to Mixer M100 where it combines with 15 kg/hr water at 70 F and 1 atm to form Product B. Both mixers operate at 1 atm. Property method: IDEAL.

The Entry

01 - Belief Update `CONSTRAINT`

I assumed that choosing IDEAL as a property method was a neutral starting point - a safe default that would give approximate answers without introducing meaningful error. After running the simulation and verifying the energy balance by hand, I no longer believe that. IDEAL assumes liquid activity coefficients of 1.0 for all components - it treats methanol and water as if they mix without any excess enthalpy, as if the molecules have no preference for their neighbors. Methanol and water are both strongly hydrogen-bonding molecules. They do not mix ideally. The product temperatures and compositions Aspen reported are approximations whose error I have not quantified, and the assignment did not ask me to.

> **Annotation:** Most students accept IDEAL without question because the simulation converges cleanly and the output looks reasonable. The student who records this belief update is asking the right question - not "did it converge?" but "what physical claim did I accept when I made that choice?" The payoff is deferred to HW6, where applying NRTL or Wilson to methanol-water produces a T-xy diagram that looks substantially different from IDEAL. A student who wrote this entry will recognize that difference immediately and be able to explain it. A student who did not will be surprised.

---

02 - Assumption + Failure Condition `CONSTRAINT`

I assumed the FSPLIT block correctly routes 40 wt% of the feed to M200 and 60 wt% to M100 based on the split fraction I specified. This assumption holds as long as the split fraction is specified on a mass basis and the feed composition is constant. It fails if the feed composition changes - FSPLIT applies the same split fraction regardless of composition, so a feed with a different methanol-to-water ratio would still split 40/60 by weight but would produce products with different compositions than expected. In a real plant where feed composition varies, an FSPLIT block would need to be paired with a controller or replaced with a more sophisticated splitting model.

> **Annotation:** The student has identified a real limitation of FSPLIT that the

> homework does not surface - it is a fixed-fraction splitter, not a composition-controlled splitter. The failure condition is specific and realistic: feed composition variation is common in real processes. This is the Constraint mindset applied correctly - the model works within its stated assumptions, and the student has named where those assumptions end.

---

03 - Math to Physical Link `CONSTRAINT`

The energy balance verification - comparing total inlet enthalpy to total outlet enthalpy for each mixer - connected the abstract conservation equation to a physical observable. In M100, the methanol-rich stream at 130 F mixes with cooler water at 70 F. The outlet temperature I calculated by hand from the enthalpy balance agreed with Aspen's output to within rounding. What this confirmed: Aspen is solving the same enthalpy balance I would solve by hand, just with property data pulled from the IDEAL model rather than steam tables or a handbook. The simulation is not doing something mysterious - it is doing arithmetic I can check.

> **Annotation:** The energy balance check is Question 3 in the homework, so most students complete it. Fewer connect it to what the check actually demonstrates - that the simulator is solving the same equations the student knows, with the same conservation laws, just faster and with more components. A student who understands this is positioned to diagnose convergence errors later in the course, because they know what Aspen is trying to do and can recognize when it fails to do it.

---

04 - Skeptical Challenge `CONSTRAINT`

A skeptical engineer would ask whether the product temperatures reported are physically meaningful given the IDEAL property method assumption. For a methanol-water system, mixing enthalpy is not zero - there is a measurable heat of mixing that IDEAL ignores. The outlet temperatures Aspen reported are therefore slightly wrong in a direction that depends on the actual excess enthalpy of mixing at these compositions. I cannot quantify the error without a better property model, but I know the error exists and that it affects the temperature more than the composition - because compositions are set by the mass balance, which is property-method-independent.

> **Annotation:** The student has made a precise and correct observation: the mass balance is property-method-independent (mass is conserved regardless of thermodynamic model), but the energy balance is not (enthalpy depends on the model). Naming the direction and type of error without knowing the magnitude is honest and technically correct. This is what engineering judgment under uncertainty looks like - not refusing to answer, but answering with explicitly stated limitations.

Primary Mindset: CONSTRAINT

The IDEAL property method is not a neutral choice - it is a constraint that bounds the validity of every temperature Aspen reported. The FSPLIT split fraction is not a physical law - it holds only when feed composition is constant. Naming the constraints you accepted, and the conditions under which they fail, is the intellectual work of this week.

One-Paragraph Contrast

A weak entry reads: *"I learned how to use the mixer and splitter blocks in Aspen Plus. I used the IDEAL property method as specified. The energy balance confirmed conservation of energy. A skeptical engineer might question whether my results are accurate."* That entry accepts IDEAL without examining what it assumes, reports the energy balance as a confirmation without connecting it to what was actually being confirmed, and produces a skeptical challenge so broad it applies to every simulation ever run. The blocks were placed correctly. The streams were connected. The simulation converged. The thinking is absent.

What This Entry Sets Up

The IDEAL property method for methanol-water is an open question that HW6 closes. When you generate the T-xy diagram for methanol-water using NRTL in HW6, the curve will show non-ideal behavior that IDEAL cannot predict. A student who wrote this entry will look at that diagram and immediately understand why it differs. A student who accepted IDEAL without question will have no framework for the difference and may assume the NRTL result is wrong rather than recognizing it as more accurate.

HW2 - Rankine Cycle

Course: CBE 253 - Aspen Plus Module
Sessions: 1-3 - Weeks 10-11 - Mar 24 - Mar 31
Assignment: HW2 - Steam Rankine Cycle
Due: March 31
Primary Mindset: Tradeoff
Connection forward: The cycle efficiency number calculated here is a permanent reference point. Every heat integration scheme, combined cycle configuration, or process improvement you encounter later in your engineering career will be evaluated against a baseline. You just calculated yours.

Simulation Context

4171 kg/hr of saturated liquid water at 0.025 MPa is pumped to 0.8 MPa (pump efficiency 0.8, driver efficiency 0.9), fed to a boiler producing superheated steam at 500 C, expanded through an isentropic turbine back to 0.025 MPa (isentropic efficiency 0.9, mechanical efficiency 0.9), then condensed to saturated liquid completing the cycle. Blocks: PUMP (P100), HEATER for boiler (H100), COMPR for turbine (C100), HEATER for condenser (H200). Property method: STEAMNBS.

The Entry

01 - Belief Update TRADEOFF

I assumed the pump and turbine were symmetric opposites - one puts work in, one takes work out, and they cancel in a straightforward way. After computing the cycle efficiency by hand, I no longer believe that. The pump work is small - roughly 1-2% of what the turbine produces at these conditions. The asymmetry exists because liquids are nearly incompressible: pumping a liquid to high pressure requires very little work compared to expanding a vapor through the same pressure drop. The turbine produces far more work than the pump consumes precisely because the working fluid changes phase between the two. That phase change is not incidental to the Rankine cycle - it is the entire reason the cycle is efficient. Removing it would collapse the work advantage.

> **Annotation:** The belief update identifies the physical mechanism behind a numerical result - the pump-turbine work asymmetry is not arbitrary, it follows directly from the incompressibility of liquids versus the expansibility of vapors. A student who understands this can immediately explain why gas turbine cycles, which compress a gas rather than pump a liquid, have lower theoretical efficiency than steam cycles for the same temperature ratio. That comparison is not in the homework but it is the natural next question for a student who has genuinely understood the Rankine cycle rather than just calculated its efficiency.

02 - Assumption + Failure Condition TRADEOFF

I assumed the specified pump and turbine efficiencies (0.8 pump isentropic, 0.9 turbine isentropic) represent realistic steady-state operating conditions. This assumption holds for a well-maintained plant operating near its design point at full load. It fails under partial load - when the turbine operates below capacity, isentropic efficiency drops, sometimes below 0.7 for large steam turbines at 50% load. The cycle efficiency I calculated is therefore an upper bound on what the plant achieves in real operation, where load varies continuously in response to grid demand. The condenser duty as a fraction of boiler input would be larger at partial load than my numbers suggest.

> **Annotation:** The partial load failure condition is not exotic - power plants cycle load constantly. A student who identifies this limitation is thinking about the simulation as a model of a real operating system with a load profile, not a textbook problem with a fixed operating point. The observation that condenser duty grows as a fraction of boiler input at partial load is thermodynamically correct and practically important - it affects the thermal discharge to whatever cooling medium the plant uses.

---

03 - Math to Physical Link `TRADEOFF`

The hand calculation for overall cycle efficiency - (Turbine Net Work - Pump Net Work) / Boiler Heat Duty - placed each Aspen output on the T-S diagram I already knew from thermodynamics. The boiler takes the stream from compressed liquid to superheated vapor: large enthalpy increase, large heat input, horizontal movement on the T-S diagram at constant pressure. The turbine extracts work by dropping enthalpy along a nearly isentropic path: vertical drop on the T-S diagram. The condenser rejects everything that was not converted to work - the remaining enthalpy leaves as heat at low temperature and low pressure. The numbers confirmed the diagram. Nothing Aspen reported was surprising once I placed each stream on the cycle.

> **Annotation:** The T-S diagram connection is the most important link in this entry. Students who complete the efficiency calculation without placing streams on the diagram are doing arithmetic without physical grounding. A student who can place each Aspen stream result on the T-S diagram - compressed liquid, superheated steam, wet vapor after expansion, saturated liquid after condensation - has integrated the simulation output with thermodynamic knowledge they already possessed. That integration is what Field 03 is designed to capture.

---

04 - Skeptical Challenge `TRADEOFF`

A skeptical engineer would challenge the thermal efficiency calculation and ask what happens to the energy the cycle rejects. The boiler inputs a large heat duty. The turbine extracts work. The condenser rejects everything not converted - and that heat goes somewhere: a river, a cooling tower, or the atmosphere. The simulation reports condenser duty as a number but does not model where that heat goes or what constraint governs its rejection. In a real plant sited near a water body, condenser duty determines the thermal discharge to the water, which is regulated under environmental permits. The efficiency I calculated is thermodynamically correct but environmentally incomplete - the condenser duty is not just a residual, it is a design constraint with regulatory consequences.

> **Annotation:** The skeptical challenge goes outside the simulation boundary deliberately. Condenser heat rejection is a real siting and permitting constraint - thermal discharge regulations have forced redesigns and shutdowns of real plants. A student who notices that the condenser duty in their Aspen output represents a real environmental constraint is beginning to think like a practicing engineer. The Tradeoff is explicit: cycle efficiency was maximized at the cost of a condenser duty that has to go somewhere, and where it goes is not the simulator's problem - it is the engineer's.

Primary Mindset: TRADEOFF

Every joule the boiler puts in either leaves as turbine work or leaves as condenser heat

rejection. There is no third option. Increasing cycle efficiency means extracting more work from the same boiler input - which means reducing what the condenser rejects. The tradeoff between work output and heat rejection is not a design choice to be optimized away; it is a thermodynamic constraint. The Carnot limit exists for this reason.

One-Paragraph Contrast

A weak entry reads: *"I simulated the Rankine cycle in Aspen using a pump, boiler, turbine, and condenser. The cycle efficiency came out to about 30%. I assumed the efficiencies given were correct. A skeptical engineer might question the accuracy of my simulation."* That entry reports the efficiency number without connecting it to the T-S diagram, accepts the specified efficiencies without identifying where they fail, and produces a skeptical challenge that challenges nothing specific. The pump-turbine work asymmetry goes unnoticed. The condenser duty is a number, not a constraint. The simulation was completed. The cycle was not understood.

What This Entry Sets Up

The Rankine cycle efficiency number is a permanent reference point - roughly 30% for these conditions, meaning roughly 70% of the boiler input leaves as condenser heat rejection. Any process improvement, heat integration scheme, or combined cycle configuration encountered later in your engineering career will be evaluated against a baseline like this one. The condenser duty concern raised in Field 04 connects directly to HW6 - where property method selection and thermodynamic analysis tools will give you the means to characterize real fluid behavior more precisely than STEAMNBS alone. The cycle you simulated here is ideal in its block structure; real plants are more complex, but this calculation is where your intuition for steam cycle performance begins.

HW3 - Flash Separators

Course: CBE 253 - Aspen Plus Module
Sessions: 4 - Week 12 - Apr 2
Assignment: HW3 - Moist Air Separation using SEP, HEATER, and FLASH2
Due: April 7
Primary Mindset: Fragility
Connection forward: The property method concern raised here - UNIFAC at cryogenic conditions - reappears directly in HW6 when property method selection becomes the explicit subject of the course.

Simulation Context

Moist air (160 lbmol/hr, 0.77 N2 / 0.215 O2 / 0.015 H2O mol fraction) at 65 F and 14 psia enters a SEP block that removes all water. The dry O2/N2 stream is compressed to 180 psia and cooled to -280 F via a HEATER block, then enters two sequential FLASH2 units: first at 80 psia with vapor fraction 0.55, then at 18 psia with vapor fraction 0.70. Three product streams result - two vapors and one liquid. Property method: UNIFAC with STEAMNBS for free water.

The Entry

01 - Belief Update `APPROXIMATION`

I assumed going in that the SEP block and the FLASH2 block were doing similar things -

both separating a stream into two products. After running the simulation I no longer believe that. The SEP block enforces a split I specify regardless of thermodynamics; it does not care whether that split is physically achievable at those conditions. The FLASH2 block at 80 psia and -280 F is doing a real phase equilibrium calculation - the vapor fraction of 0.55 emerges from the VLE, it is not imposed. The difference matters because the SEP result can be physically unrealistic and I would never know from the output alone.

> **Annotation:** The belief update is specific and traceable to a real moment in the simulation. The student is not reporting that they learned about flash calculations - they are recording a precise conceptual shift: SEP imposes a split, FLASH2 calculates one. That distinction is the thermodynamic content of this session, stated in the student's own words with their own example. This understanding is prerequisite to using separators correctly in any future simulation.

---

02 - Assumption + Failure Condition `FRAGILITY`

I assumed UNIFAC is appropriate for an N2/O2 mixture at cryogenic conditions. This assumption likely fails below approximately -280 F where both components are near or below their normal boiling points (N2 at -320 F, O2 at -297 F). UNIFAC is an activity coefficient model developed for liquid mixtures of organic compounds - applying it to a near-critical cryogenic system is outside its intended range. The simulation converged and gave plausible numbers, but I cannot verify that UNIFAC is producing accurate K-values at these conditions without comparing against an equation of state like Peng-Robinson, which is designed for this regime.

> **Annotation:** The student identified a specific technical concern - UNIFAC at cryogenic conditions - and located the failure condition precisely: near the normal boiling points of the components. The honest admission that the simulation converged and gave plausible numbers but cannot be verified is exactly the right epistemic position. This is the Fragility mindset made explicit. It also implicitly identifies a design check: run it again with Peng-Robinson and compare. That check is what HW6 will teach explicitly.

---

03 - Math to Physical Link `SYSTEMS`

The Rachford-Rice equation governs the vapor fraction in the FLASH2 block. At 80 psia and -280 F, the calculation produced VF = 0.55 as specified - meaning roughly 55% of the molar flow left as vapor, enriched in nitrogen (the more volatile component at these conditions, lower boiling point). That vapor enrichment is visible in the stream table: the vapor product from F100 shows approximately 82 mol% N2 versus 77% entering the flash. The equation predicted a composition shift that is visible physically as nitrogen preferentially partitioning to the vapor phase - which is the basis for cryogenic air separation in industrial practice.

> **Annotation:** The student connects a specific equation (Rachford-Rice), a specific numerical output (VF = 0.55), and a specific physical observable (nitrogen enrichment in vapor) in one paragraph. They then extend it to industrial context - cryogenic air separation - which demonstrates that the simulation is a model of a real process. The numbers are approximate and clearly labeled as such, which is honest and appropriate.

---

04 - Skeptical Challenge `TRADEOFF`

A skeptical engineer would challenge the energy accounting in this flowsheet. The HEATER block cools the stream to -280 F - that duty has to come from somewhere, and in a real plant it would come from a refrigeration cycle that consumes significant work. The simulation reports a cooling duty but does not model where that cooling comes from or what it costs in electricity. The process as simulated is thermodynamically incomplete: I know the separation works at these conditions but I do not know the true energy cost, which is the number that determines whether the process is economically viable.

> **Annotation:** The challenge is specific and consequential - not a generic observation about heat losses but a precise identification of what the model omits (the refrigeration cycle) and why that omission matters (energy cost determines economic viability). The student names what was traded and what was lost. This entry would immediately generate a productive conversation in class.

Primary Mindset: FRAGILITY

The UNIFAC concern is the most consequential judgment in this entry. The student identified the regime where the model stops being trustworthy before the model gave any visible signal of failure - the simulation converged cleanly despite operating outside UNIFAC's valid range. Finding where a model breaks before it breaks is the Fragility mindset precisely.

One-Paragraph Contrast

A weak entry reads: *"I learned more about how flash separations work. I assumed the property method was appropriate. The flash calculation determined the vapor and liquid compositions of the outlet streams. A skeptical engineer would question whether the simulation is accurate."* That entry answers all four fields in form but none in substance. No specific belief is named. The property method concern is circular - "appropriate" is not an assumption, it is a conclusion. The math-to-physical link describes what flash calculations do in general, not what this simulation showed specifically. The skeptical challenge challenges nothing. The simulation was run. The thinking is absent.

What This Entry Sets Up

The UNIFAC concern raised in Field 02 is not resolved here - it is left open deliberately. HW6 introduces the property method selection framework that explains why Peng-Robinson is the right choice for nonpolar gases at high pressure. A student who wrote this entry arrives at HW6 with a specific, remembered reason to care about property method selection - not as an abstract course topic but as an unresolved question from their own simulation.

HW4 - Distillation DSTWU

Course: CBE 253 - Aspen Plus Module
Sessions: 5-6 - Weeks 12-13 - Apr 7-9
Assignment: HW4 - Distillation of Styrene / Ethylbenzene / Toluene using DSTWU
Due: April 14
Primary Mindset: Approximation
Connection forward: The DSTWU outputs recorded here become the starting inputs for the RADFRAC simulation in HW5. Where this entry ends - with an estimate and its known limitations - HW5 begins.

Simulation Context

A reactor output mixture of 35 mol% styrene, 45 mol% ethylbenzene, and 20 mol% toluene is fed to a distillation column at 120 lbmol/hr, 310 F, 30 psia. The goal is to produce a high-purity styrene bottoms product (>=0.995 mol fraction) and an ethylbenzene-rich overhead (>=98.5% ethylbenzene recovery). DSTWU shortcut method, SRK property method. Total condenser, column top pressure 24 psia.

Design heuristics applied:

- Actual reflux ratio = 1.2-1.5 x minimum reflux ratio
- Actual equilibrium stages ~= 2 x minimum stages

The Entry

01 - Belief Update `APPROXIMATION`

I assumed the heuristics - 1.2-1.5x minimum reflux ratio, twice the minimum stages - were rules with a physical basis that would reliably produce an optimal design. After working through the DSTWU iterations I no longer believe that. The heuristics encode *economic* experience from an era of specific energy costs and construction norms, not thermodynamic principles. At different energy prices or with different feed compositions the optimal point shifts. DSTWU showed me that when I set the reflux ratio to 1.2x minimum, the required stages increased significantly - and when I pushed toward 1.5x, the stage count dropped but energy cost went up. The heuristic range is a search window, not an answer. I picked a point inside that window based on the RR-versus-stages plot, not because any single point was provably optimal.

> **Annotation:** This is the belief update the assignment is designed to produce. The student arrived thinking heuristics are laws and left understanding they are compressed historical judgment. The specific observation - that 1.2x and 1.5x bound a tradeoff rather than prescribe a solution - shows the student actually ran the sensitivity and watched the curve. That is the Approximation mindset made explicit: every heuristic encodes assumptions about context that may not hold.

---

02 - Assumption + Failure Condition `APPROXIMATION`

I assumed SRK is appropriate for this styrene/ethylbenzene/toluene system. This assumption is reasonable - all three are nonpolar hydrocarbons and SRK handles nonpolar systems well - but it fails if small amounts of water or oxygenated species are present in the feed, as can happen in real dehydrogenation reactor effluents. SRK does not handle polar-nonpolar interactions accurately. In this simulation the feed is specified as pure hydrocarbon, so the assumption holds. In a real plant the feed would need to be characterized more carefully before accepting SRK as the property method.

> **Annotation:** The student correctly validated the property method choice for the specified system, then located precisely where it would fail in a real application. The connection to real reactor effluents is not a digression - it is the signal that the student is thinking about the simulation as a model of something physical, not as an exercise in itself.

---

03 - Math to Physical Link `SYSTEMS`

The Fenske equation underlies the minimum stages calculation in DSTWU - it assumes total reflux and constant relative volatility. At total reflux, every molecule that goes up comes back down, so the separation per stage is maximized. The minimum number of stages DSTWU reported (approximately 11 for this system) is therefore a lower bound that can never be achieved in practice, because real operation requires a finite distillate flow. The actual stages (approximately 22, following the 2x heuristic) account for the real reflux ratio. What connected for me: the temperature difference between the distillate (~290 F) and the bottoms (~340 F) maps directly to the volatility difference between ethylbenzene and styrene - styrene's higher boiling point is what drives the separation and sets the reboiler temperature.

> **Annotation:** Two connections in one field, both legitimate. The Fenske equation link is conceptual - minimum stages as a limit, not an operating point. The temperature-to-volatility link is observable - the student read the stream table and traced the number back to a physical property. The column operating between those temperatures is not arbitrary - it is set by the thermodynamics of the system.

---

04 - Skeptical Challenge `TRADEOFF`

A skeptical engineer would challenge whether 0.995 mol fraction styrene in the bottoms is actually achievable at the reflux ratio and stage count I selected, and would ask to see the stream table, not just the block report. DSTWU is a shortcut - it estimates the operating conditions that *should* achieve the spec, but it does not verify that they do. The block report gives minimum stages and reflux ratio under idealized assumptions. Whether those translate to the specified purity in a rigorous tray-by-tray calculation is not confirmed until RADFRAC is run. I am reporting a design estimate, not a verified result.

> **Annotation:** The student has correctly identified the epistemological status of DSTWU output: it is an estimate whose validity is contingent on the rigorous calculation agreeing with it. DSTWU is a tool for getting to RADFRAC, not a substitute for it. A student who understands this is ready for HW5. A student who treats DSTWU output as the final answer will be confused when RADFRAC gives different results.

Primary Mindset: APPROXIMATION

Every number DSTWU reports is conditional on assumptions that RADFRAC will test. The minimum stages, the optimal reflux ratio, the distillate-to-feed ratio - all are estimates derived from simplified thermodynamic assumptions (constant relative volatility, constant molar overflow). Recording which assumptions were made and where they are likely to hold is the intellectual work of this assignment.

One-Paragraph Contrast

A weak entry reads: *"I used DSTWU to find the number of stages and reflux ratio. I set the reflux ratio to 1.3 times the minimum and the stages to twice the minimum as the heuristic says. The block report showed the required values and the design criteria were met."* That entry reports procedure without recording judgment. It does not name what the heuristic encodes, does not question whether SRK is appropriate, does not connect the temperature profile to the volatility difference, and does not flag that DSTWU results are estimates awaiting RADFRAC confirmation. The simulation was run correctly. The thinking is invisible.

What This Entry Sets Up

The DSTWU outputs from this entry - approximately 22 actual stages, reflux ratio in the 1.3-1.4x minimum range, distillate-to-feed ratio near 0.55, feed stage near the middle of the column - become the starting inputs for RADFRAC in HW5. The skeptical challenge in Field 04 is not rhetorical: it is the literal setup for the next assignment. When RADFRAC disagrees with DSTWU, the student who wrote this entry will have already predicted that disagreement and will be able to explain why it happened rather than treating it as an error.

HW5 - Reactor and Distillation

Course: CBE 253 - Aspen Plus Module
Sessions: 7-8 - Week 13 - Apr 14-16
Assignment: HW5 - Ethylbenzene Dehydrogenation with RSTOIC and RADFRAC
Due: April 21
Primary Mindset: Systems
Connection forward: The single-pass conversion calculated here - and the unreacted ethylbenzene leaving in the overhead - is the natural entry point for thinking about recycle streams and process economics. In a real styrene plant, that overhead is not waste, it is a recycle. HW6 begins to give you the property analysis tools to evaluate what separating and recycling it would cost.

Simulation Context

Pure ethylbenzene feed (100 lbmol/hr, 370 F, 55 psia) enters an RSTOIC reactor at 1100 F and 50 psia. Two reactions in series: dehydrogenation of ethylbenzene to styrene (0.60 fractional conversion), followed by further dehydrogenation of styrene to toluene and methane (0.34 fractional conversion of styrene). Reactor effluent is cooled to 185 F at 37 psia in a HEATER block, then fed to an adiabatic FLASH2 unit at 32 psia that removes most H2 and CH4 as vapor. The liquid product feeds a RADFRAC column with partial vapor condenser, 25 psia condenser pressure, ~0.1 psi/stage pressure drop. RADFRAC inputs seeded from HW4 DSTWU results. Target: >=99.8 mol% styrene in bottoms. Property method: SRK.

The Entry

01 - Belief Update `SYSTEMS`

I assumed the two reactions in the reactor were independent - that I could think about the ethylbenzene conversion and the styrene conversion separately and add the results. After setting up the RStoic block with "reactions in series" checked and watching how the second reaction consumed the styrene produced by the first, I no longer believe that. The second reaction cannot proceed until the first has produced styrene - the output of reaction one is the input to reaction two. The 0.34 fractional conversion of styrene applies to the styrene present after reaction one, not to the original feed. A student who treats the reactions as independent will calculate the wrong styrene yield and the wrong feed composition to the distillation column, which propagates error through the entire flowsheet.

> **Annotation:** The belief update identifies a sequencing error that is easy to make and consequential when made. The "reactions in series" checkbox in RStoic is not a formality - it changes the calculation. A student who understood why it changes the calculation has understood that process flowsheets are not collections of independent units but sequences where upstream outputs become downstream inputs. That is the Systems mindset precisely: the flowsheet as a connected whole, not a set of blocks.

---

02 - Assumption + Failure Condition `SYSTEMS`

I assumed that the HEATER block cooling the reactor effluent to 185 F at 37 psia would condense enough of the mixture that the subsequent FLASH2 unit would cleanly separate H2 and CH4 from the heavier hydrocarbons. This assumption holds when H2 and CH4 are the only light components - both have boiling points far below 185 F and will remain in the vapor phase. It fails if significant amounts of lighter cracking products form in the reactor at 1100 F - benzene, for instance, is more volatile than toluene and could partially report to the vapor stream from the flash unit along with H2 and CH4, reducing the liquid feed to the distillation column and changing its composition. The RStoic block does not model side reactions unless explicitly specified, so any cracking at 1100 F is invisible in this simulation.

> **Annotation:** The student has identified a real limitation of RStoic - it only models the reactions you tell it about. At 1100 F, ethylbenzene dehydrogenation is accompanied by thermal cracking reactions in real industrial reactors. The simulation is clean because unspecified reactions do not exist in RStoic. A student who recognizes that reactor temperature creates the possibility of side reactions the model ignores is thinking about the gap between the simulation and the physical system - which is the correct posture for every reactor simulation.

---

03 - Math to Physical Link `SYSTEMS`

The single-pass conversion of ethylbenzene - moles reacted divided by moles fed -

connects the reactor output directly to the distillation column design. At 0.60 fractional conversion of ethylbenzene in reaction one, 40 lbmol/hr of unreacted ethylbenzene leaves the reactor alongside the styrene product. That unreacted ethylbenzene reports predominantly to the RADFRAC overhead because ethylbenzene is more volatile than styrene - it is the light key in the distillation. The overhead stream is therefore not pure product; it is mostly ethylbenzene with some toluene and trace methane. The bottoms is the styrene product. The column is not purifying a product - it is separating an unreacted reactant from a product, which is the most common distillation task in continuous chemical manufacturing.

> **Annotation:** The student has connected the reactor conversion number to the distillation column function. This is the Systems connection the assignment is designed to produce - the reactor and the distillation column are not two separate exercises, they are two stages of one process. The observation that the overhead is mostly unreacted ethylbenzene rather than a byproduct is physically correct and sets up the natural next question: where does that ethylbenzene go in a real plant? The answer is recycle, which is the topic that connects this course to reactor design and process economics.

---

04 - Skeptical Challenge `TRADEOFF`

A skeptical engineer would challenge the 99.8 mol% styrene specification and ask what it costs in reboiler and condenser duty to achieve it compared to, say, 99.0 mol%. The RADFRAC simulation converged to the specification, but getting there required adjusting the reflux ratio upward from the DSTWU estimate. Every increment of purity beyond 99.0% requires disproportionately more energy - the separation becomes harder as the light key concentration in the bottoms approaches zero. The condenser and reboiler duties at 99.8% purity are significantly higher than at 99.0%, and the specification of 99.8% was given without justification. A skeptical engineer would want to know what downstream process requires 99.8% rather than 99.0%, because the energy cost difference is real and the specification may be conservative.

> **Annotation:** The skeptical challenge questions the specification itself, not the simulation that achieved it. This is the correct target - specifications in engineering are not arbitrary, they exist because a downstream user or process requires a particular purity. A student who asks "why 99.8% and not 99.0%?" is thinking about the process in its industrial context. The observation that purity specifications become disproportionately expensive near 100% is thermodynamically correct and practically important - it is one of the reasons industrial distillation columns are designed to specification rather than to maximum achievable purity.

Primary Mindset: SYSTEMS

The reactor conversion sets the distillation feed composition. The distillation feed composition sets the reflux ratio and stage count needed to meet the purity specification. The purity specification determines the condenser and reboiler duties. Each unit in the flowsheet is a constraint on every unit downstream. Changing the reactor conversion by 10% changes the distillation column design. The flowsheet is a system - pulling one variable changes everything connected to it.

One-Paragraph Contrast

A weak entry reads: *"I simulated the reactor and distillation column in Aspen. The RStoic block modeled two reactions and the RADFRAC column separated the products. The bottoms product met the 99.8 mol% styrene specification. A skeptical engineer might question whether my column design is optimal."* That entry describes what blocks were used without recording what the student understood about how they interact. The reactions-in-series dependency goes unnoticed. The HEATER block assumption is accepted without examination. The single-pass conversion is reported as a number without connecting it to the distillation column feed composition. The purity specification is

accepted without questioning why it was set there. The simulation was completed. The system was not understood.

What This Entry Sets Up

The unreacted ethylbenzene in the overhead stream is a planted question about recycle. In a real styrene plant, that stream is recycled to the reactor feed - the process operates with recycle to improve overall yield. HW6 gives you the property analysis tools to begin characterizing the thermodynamic behavior of these component pairs. A student who understands that the overhead stream is mostly unreacted ethylbenzene will look at the ternary diagram in HW6 with a specific question: how does styrene, ethylbenzene, and toluene separate, and what does that mean for recycle design?

HW6 - Properties and Ternary Analysis

Course: CBE 253 - Aspen Plus Module
Sessions: 9-12 - Weeks 14-15 - Apr 21 - Apr 30
Assignment: HW6 - Property Analysis: n-Pentane/Methanol/Acetone System
Due: May 5
Primary Mindset: Decision
Connection forward: This is the final entry. The property method question planted in HW1 - what does IDEAL cost you for a hydrogen-bonding system? - has a quantitative answer now. The T-xy diagrams you generated this week are the evidence. Every simulation in this course used a property method someone chose. This week you learned what that choice means.

Simulation Context

Problem 1 - Binary analysis. n-Pentane/methanol system at 100 psia. T-xy diagrams and Gibbs free energy of mixing plots generated using three property methods: van Laar, Peng-Robinson, and UNIQUAC. Analysis of azeotrope prediction and two-liquid-phase behavior from Delta Gmix plots at 77 F.

Problem 2 - Ternary analysis. n-Pentane/methanol/acetone system at 100 psia using UNIQUAC. Ternary phase diagram generated to assess miscibility for each component pair, azeotrope locations, and phase behavior at overall composition 0.4 pentane / 0.4 methanol / 0.2 acetone.

The Entry

01 - Belief Update `DECISION`

I assumed that different property methods would give similar T-xy diagrams for the same system - that the choice of method was a refinement, not a fundamental change in what the simulation predicts. After generating T-xy diagrams for n-pentane/methanol with van Laar, Peng-Robinson, and UNIQUAC, I no longer believe that. The three diagrams are qualitatively different. Peng-Robinson, an equation of state designed for nonpolar systems, does not predict the azeotrope that van Laar and UNIQUAC show for this polar-nonpolar pair. The property method is not a refinement - it is a modeling assumption that determines what phenomena the simulation can see. A simulation using Peng-Robinson for pentane-methanol is blind to the azeotrope. It will converge, produce output, and report nothing unusual - because the model cannot represent the non-ideal liquid behavior that creates the azeotrope.

> **Annotation:** This is the belief update the entire course has been building toward. From HW1, where IDEAL was accepted without question, to this week, where

> three methods produce qualitatively different results for the same system. The student has arrived at the central lesson of property method selection: the choice determines what physics the model can represent, not just how accurately it represents it. A model that cannot represent an azeotrope will not predict one regardless of how carefully the simulation is set up. This understanding is prerequisite to any serious process simulation work.

---

02 - Assumption + Failure Condition `DECISION`

I assumed UNIQUAC is the most accurate of the three methods for this system because it accounts for molecular size and surface area differences - pentane and methanol are very different in both. This assumption is reasonable and supported by the decision tree in the course notes: for polar non-electrolyte systems at low to moderate pressure where liquid-liquid equilibrium may occur, activity coefficient models like UNIQUAC are preferred over equations of state. It fails if binary interaction parameters for pentane-methanol are not available in the Aspen database or were regressed from data outside the temperature and pressure range of this simulation. UNIQUAC with poorly fitted parameters can perform worse than a simpler model with well-fitted ones. I accepted the Aspen database parameters without verifying their source or regression range.

> **Annotation:** The student made a specific property method recommendation - UNIQUAC - and immediately identified its failure condition: parameter quality. This is the Decision mindset made operational. Recommending a property method is not just choosing from a decision tree; it is accepting responsibility for the parameters that implement the model. The honest admission that database parameters were accepted without verification is not a weakness in the entry - it is the correct statement of what was done and what was not. In professional practice, parameter verification is a required step before a simulation is used for design decisions.

---

03 - Math to Physical Link `DECISION`

The Gibbs free energy of mixing plot for pentane-methanol at 77 F connected thermodynamic theory to a visible prediction. A negative Delta Gmix across the full composition range means the components mix spontaneously - one liquid phase. A Delta Gmix curve with a positive region or an inflection point indicating local concavity signals the possibility of two liquid phases. For pentane-methanol at 77 F, UNIQUAC showed a Delta Gmix curve with the characteristic double-well shape - two minima separated by a local maximum - indicating that compositions in the middle range would spontaneously split into two liquid phases of different compositions. Van Laar predicted a similar shape. Peng-Robinson did not - its Delta Gmix curve was negative and concave throughout, missing the liquid-liquid behavior entirely. The mathematics of the mixing curve directly predicted a physical phenomenon visible in the ternary diagram.

> **Annotation:** The connection between Delta Gmix curve shape and liquid-liquid phase splitting is the most important mathematical link in this assignment. A student who can look at a Delta Gmix plot and read two liquid phases from the double-well shape has internalized the thermodynamic criterion for phase splitting - not as a formula to apply but as a visual pattern with a physical meaning. The observation that Peng-Robinson missed this behavior entirely closes the loop from Field 01: the method that cannot represent the non-ideal liquid behavior also cannot produce the Delta Gmix shape that signals it. The two observations confirm each other.

---

04 - Skeptical Challenge `DECISION`

A skeptical engineer would challenge my recommendation of UNIQUAC by asking for experimental data to validate it. I generated three T-xy diagrams that disagree with each other and selected UNIQUAC as most accurate based on theoretical reasoning - molecular size differences, suitability for polar systems, decision tree guidance. But I did not compare any of the three predictions against experimental VLE data for pentane-methanol. The NIST database contains experimental binary VLE data for this system. A recommendation of UNIQUAC that has not been checked against experiment is a theoretical preference, not

a validated choice. In a design context - sizing a distillation column, predicting azeotrope composition - the difference between a validated and unvalidated property method choice can be the difference between a column that meets specification and one that does not.

> **Annotation:** The skeptical challenge names the missing step that converts a reasonable choice into a defensible one: experimental validation. This is the highest-level engineering judgment in the course - not selecting from a decision tree, but knowing what it would take to trust the selection. The NIST database is specifically mentioned because it is available in Aspen and was introduced in the lecture material. A student who knows it exists and knows it was not consulted has correctly identified the gap between what they did and what professional practice requires.

Primary Mindset: DECISION

Every simulation in this course used a property method. In the early weeks that choice was given - IDEAL, UNIFAC, SRK, STEAMNBS. This week the choice was yours to make and defend. The decision is not which method appears highest on a ranking - it is which method best represents the physics of the system you are modeling, with parameters you can verify, for conditions within the range where the model was validated. That is a different kind of question than anything Aspen can answer for you.

One-Paragraph Contrast

A weak entry reads: *"I generated T-xy diagrams using three property methods and they looked different. I chose UNIQUAC because it is more accurate for polar systems. The ternary diagram showed some two-phase regions. A skeptical engineer might question whether I chose the right property method."* That entry reports that the diagrams differed without naming how or why. It recommends UNIQUAC without connecting the recommendation to molecular reasoning or the Delta Gmix evidence. It notes two-phase regions on the ternary diagram without connecting them to the Delta Gmix analysis from Problem 1. And it produces a skeptical challenge that is indistinguishable from a question the student already answered - without naming experimental validation as the missing step. The plots were generated. The decision was not made.

What This Entry Sets Up

This is the last entry. Look back at HW1 - the IDEAL property method for methanol-water, accepted without question because the assignment specified it. You now know what that acceptance cost: IDEAL cannot predict the non-ideal mixing behavior of a hydrogen-bonding system, would miss an azeotrope if one existed, and produces outlet temperatures that are wrong in a direction you could not quantify at the time. You can quantify it now. The six entries in this notebook trace a single arc: from accepting model assumptions silently to naming them, testing them, and defending the choices you make under uncertainty. That arc is what engineered judgment looks like in practice.