INDUSTRIAL APPLICATIONS OF CHEMICAL PROCESS SYNTHESIS

Jeffrey J. Siirola

Eastman Chemical Company
Kingsport, Tennessee

Systematic approaches for the invention of conceptual chemical process designs have been proposed and discussed for more than twenty-five years. During that same time, the importance of front-end engineering, especially conceptual design, to product quality, health and safety, environmental impact, energy consumption, operability, capital and operating costs, and overall competitiveness has become *ever more apparent. A number of process synthesis frameworks, approaches, methods, and tools have now been developed to the point of industrial application. This chapter describes a framework for the industrial chemical plant innovation process, showing how process synthesis fits into that structure and how that framework has in turn influenced the development of systematic process synthesis methods. It also describes a number of industrial case studies in which process synthesis techniques have been successfully applied to the conceptual design of total process flowsheets, as well as to specific design subproblems including heat-integrated distillation trains, multiple-effect distillation, and the separation of azeotropic systems. Typical energy savings of 50% and net present cost reductions of 35% have been achieved in industrial practice using systematic process synthesis methodologies. Even greater benefits are expected to be realized as the next generation of approaches currently being developed is transferred to industry.*

I. Introduction

The manufacturing sector of the chemical processing industry is generally in the business of making materials rather than making artifacts. This is done in response to perceived needs and the belief that the materials offered will satisfy these needs in a valuable manner.

In contrast with artifact-making, material-making tends to involve more conversion and transformation than assembly, and is generally more capital- and energy-intensive than labor-intensive. Material-making sometimes involves substances that are toxic or otherwise hazardous to the environment. Furthermore, since material-making equipment costs are often sublinear functions of capacity, material-making facilities tend to be large and aggregated to take advantage of economies of scale, integrated material and energy flows, and centralized environmental mitigation. Many material-making facilities have very long operating lives, much longer than the life of many artifacts and most artifact-making machinery.

Chemical manufacturing plants come into existence through a series of actions sometimes called the *innovation process*. This process leads from the identification of a need to the operation of a material-making facility. The characteristics of material-making—relatively few but fairly large manufacturing facilities, very long operating lives, high initial capital costs, high continuing operating costs, and potential environmental impacts—place special importance on sound implementation of the innovation process and, in particular, on making good engineering design decisions within that process.

All existing chemical process designs were somehow invented. Some, especially lower volume products involving complex chemistry, have been implemented by a rather straightforward extrapolation of the laboratory procedure used to experimentally demonstrate the transformation of available raw materials into the desired product. Larger volume products, on the other hand, are more often implemented as continuous processes that bear little resemblance to either the procedure or the equipment used in the laboratory. Generally there is a combinatorially very large number of alternative pieces of equipment and interconnections among these pieces of equipment that will feasibly implement the desired chemistry. Identifying better process alternatives is a key activity within the innovation process.

Successful design engineers seem to build and evolve conceptual process flowsheets from a rich repertoire of past experiences and design heuristics. Such an experience base generally includes an extensive knowledge of available equipment, simple and complex unit operations, standard tricks or patterns (for example, strategies for breaking azeotropes), encyclopedias of complete flowsheets for existing chemicals, and some sense of hierarchy of which process design problems to tackle in what order and at what level of detail. This wealth of background information is copied directly or modified as necessary to fit the situation at hand. Generally time and resource constraints limit the number of conceptual process alternatives that may be generated and evaluated by the designer to a tiny fraction of the total number feasible. The key to discovering alternatives with superior economics is the judicious use of modern methods and tools together with a little good luck.

Occasionally, process designs are produced that are conceded by those skilled in the art as being *clever*. Perhaps most, if not all, world-beating designs exhibiting superior economics exploit something clever. What is technically feasible, what is competent, and what is clever process design? Is there a *best* design that cannot be beaten? Can the invention of chemical process designs be organized, systematized, or even automated? How can more or better alternatives be generated? These questions have been the focus of process synthesis research over the last twenty-five years. It is not the intention of this chapter to discuss the latest advances in that research. Rather, it is to illustrate that some of these results are beginning to have a real impact on industrial practice.

II. Industrial Chemical Process Innovation

To understand chemical process synthesis in an industrial environment, it may be useful to first discuss the chemical innovation process. *Invention* is discovery. A new material composition may be invented. The chemistry to transform raw materials into this new composition may be invented. The process flowsheet to implement this chemistry on an industrial scale may also be invented. But the invention of a flowsheet does not guarantee that the chemistry will be reduced to practice in an economical manner. That is accomplished by *innovation,* an organized multistage goal-directed process, which leads from the identification of a customer's need to the operation of a facility to produce a material believed to address that need.

The innovation process may be implemented in a number of different ways. The specific details and emphasis within the innovation process may differ, depending on whether the objective is to build a pioneering facility for a new chemical or an improved facility for an existing chemical. But in general, a product that addresses the need must be (or has already been) identified; a chemical route must be found from available raw materials that produces this desired product; and a facility that implements this chemistry must be conceived, designed, constructed, started, operated, and maintained (Fig. 1). Of special interest in the context of conceptual process design are the stages *basic chemistry* (in which the fundamental reaction chemistry is selected), *detailed chemistry* (in which supporting chemical details of catalysis, solvents, and reaction conditions

- **Need Identification**
- **Manufacturing Decision**
- **Basic Chemistry**
- **Detailed Chemistry**
- **Task Identification**
- **Unit Operations**
- **Basic Plant Engineering**
- **Detailed Engineering**
- **Vendor Specifications**
- **Component Acquisition**
- **Construction Plan and Schedule**
- **Plant Construction**
- **Operating Procedures**
- **Commissioning and Start-up**
- **Production Plan and Schedule**
- **Plant Operation and Maintenance**

FIG. 1. Innovation process sequence.

are defined), *task identification* (in which the physical operations to prepare raw materials for reaction and isolate reaction products for sale are identified), *unit operations* (in which the chemical and physical operations previously identified are associated with actual pieces of equipment), and *basic plant engineering* (where the supporting utilities and other facilities infrastructure are defined).

Each stage of the innovation process is in a sense implemented by all of the stages that follow. Consequently, there is a great deal of interaction among the stages. For example, the choices that can be made at one stage are clearly limited by selections made during previous stages. At the same time, the optimal choice to be made among alternatives identified at any stage may well depend on costs associated with subsequent stages that implement that choice but have yet to be addressed. Earlier stages are both less well-defined and less constrained than later stages. However, decisions made in these earlier stages typically prove to have a greater impact on the overall economic outcome of the entire venture. Because of the interacting nature of the stages, it is often necessary to revisit stages or iterate among the stages in order to converge to an acceptable solution.

A. Multiple Levels of Detail

Because of the interacting nature of the innovation process stages, it is possible that each visit to a stage may be approached with a slightly different objective or even conducted at different levels of detail. For example, each stage of the innovation process might be visited four times (Fig. 2).

1. The first pass through any stage of the innovation process is at the lowest level of detail, which might be called the *targeting* level. Its purpose might be to get a rough indication of what is to be accomplished at that stage of the innovation process as well as to see what is likely to be feasible given the choices that have been made previously in the process. Targeting is already familiar in the context of heat-integration networks where fairly simple analytical procedures can give much information about what is technically and perhaps economically feasible by bounding the expected performance even before the generation of any heat-exchanger network alternatives.
2. The second pass through each innovation stage might be called the *preliminary* or *conceptual* level. Here, a tentative, not too detailed solution to the given innovation problem is conceived. One general design paradigm to accomplish this essentially consists of a four-block procedure, as shown in Fig. 3. In the first *formulation* block, the goals for the particular stage of the innovation process are specified. This is followed by an iteration of three blocks consisting of *synthesis* (generation of a solution

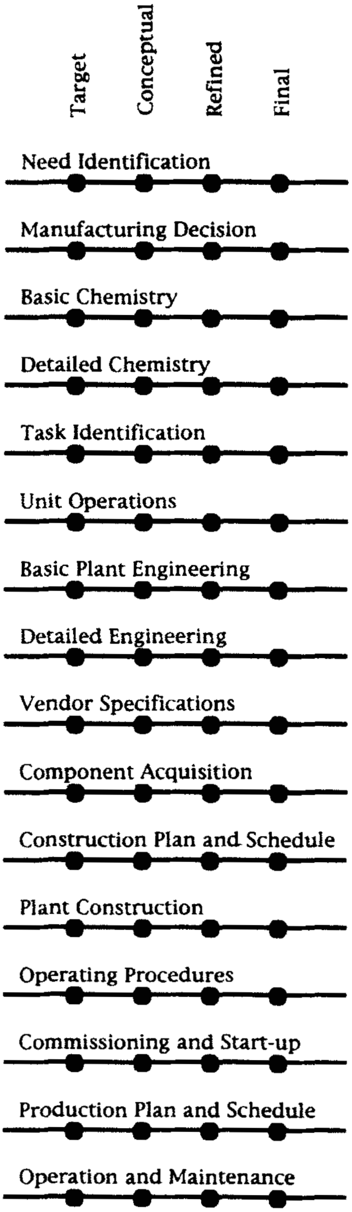

FIG. 2. Innovation process levels of detail.

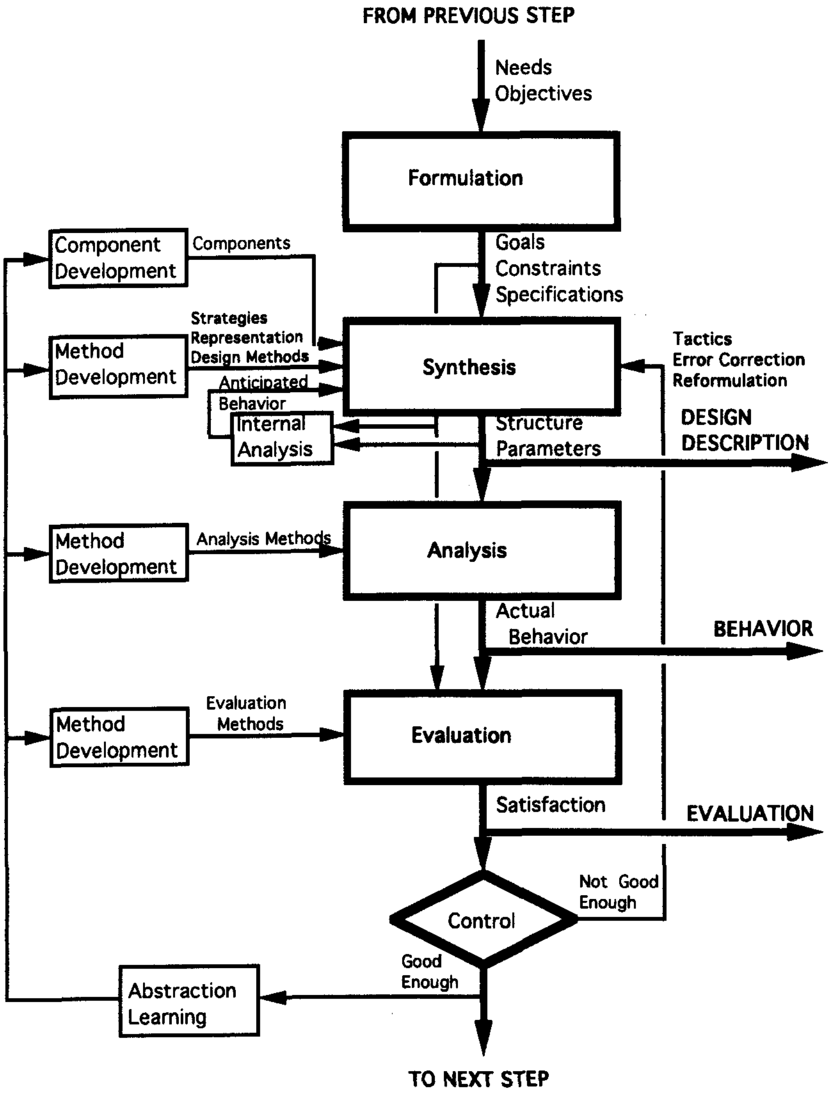

FIG. 3. General design paradigm.

alternative from available components), *analysis* (determination of the behavior of the alternative generated), and *evaluation* (comparison of the performance of the proposed alternative against the goals specified by the formulation block). If the performance of the proposed solution is judged to be satisfactory, the step is completed; otherwise, a new alternative must be generated, analyzed, and evaluated. In general, alternatives are not generated blindly, but are guided by some analysis internal to the synthesis step (sometimes encoded in the problem representation, sometimes encoded within synthesis rules of thumb, and sometimes performed by specialized analysis methods) so that alternatives are generated with anticipated behavior close to the desired behavior. In order to execute this general design strategy (which has variously been called *generate-and-test, propose–critique–modify,* and other names), components and formal or informal problem-specific alternative generation methods, analysis methods, and evaluation methods must be available. The development of such methods and tools for generating alternatives for the conceptual engineering steps of the innovation process is what much of process synthesis research is all about.

3. Because of the interacting nature of the innovation process stages, it is not efficient to make the attempt at generating a solution at the conceptual level particularly detailed. Rather, as additional information is developed in both previous and subsequent stages of the innovation process, a more *detailed* or *refined* solution can be attempted at a higher level of detail (see Fig. 2). The same basic four-block design procedure used at the conceptual level may be used in the refined level, but with different specific synthesis, analysis, and evaluation methods and tools capable of providing the required higher level of detail.
4. Finally, as yet more information becomes available, it may be necessary to repeat the refined level one *final* time using the same methods and tools that were used for the third pass.

B. A JOURNEY THROUGH THE INNOVATION PROCESS

One possible strategy might to attack each stage of the innovation process in a linear manner, going from target level to the conceptual level to the refined level to a repeat of the refined level for a single stage of the innovation process, and then moving on to the next stage. An alternative strategy that also accomplishes an iteration among the innovation process stages is a diagonal approach wherein the execution of the conceptual level at one stage just precedes the execution of the detailed level of the previous stage, and just follows the execution of the targeting level of the following stage, etc. (Fig. 4). This back-and-forth approach in different innovation stages as well as in different levels of detail enables a certain amount of look-ahead and allows preliminary downstream results to have an impact on the final solutions developed at previous stages of the innovation process. This format also defines the types and detail of information that may reasonably be expected to be known as each step of each stage of the innovation process is approached.

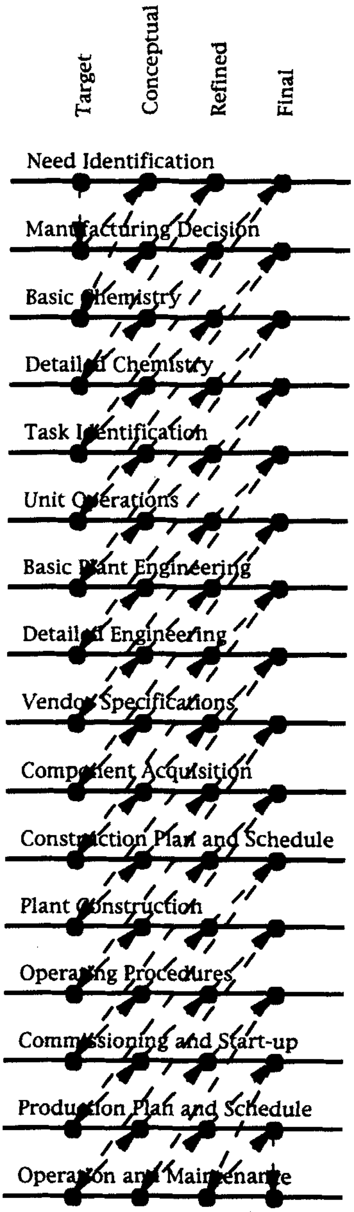

FIG. 4. Journey through the innovation sequence.

As an aside, this approach to the execution of the innovation process may also have an impact on the way industrial Research, Development, and Engineering functions might best be organized. A traditional organization might align specific innovation stages with particular company functions: for example, the need identification and manufacturing decision stages may be aligned with Business Development; basic and detailed chemistry stages with Research; task identification and unit operations stages with Development; plant and detailed engineering stages with engineering; and so forth (Fig. 5). Alternatively, the functional boundaries might be diagonal (Fig. 6), to more smoothly align with the iterative journey through the different levels of detail of the innovation process stages. This implies, for example, that in addition to the heavy chemistry emphasis usually associated with Research, it may also be appropriate for a research organization to maintain selected needs identification, business strategy, and conceptual process engineering functionalities at specific levels of detail.

III. Synthesis in Conceptual Process Engineering

Process synthesis is the invention of conceptual chemical process designs. Process synthesis is not just one step in the innovation process. Rather, it involves the generation of alternatives in all conceptual process engineering steps within the innovation process, including all the steps in the task identification stage and the less detailed steps of the unit operations stage. Because of the interacting nature of the innovation process stages, there is also some overlap with the preceding chemistry stages and the subsequent basic engineering stages.

The goal of conceptual process engineering is to develop a skeletal scheme for a material-making facility, as may be represented by a *flowsheet*, to implement the chemical conversion of available raw materials into desired product materials, fit for use, at the desired scale, safely, responsibly, economically, and on time. This flowsheet will be refined and optimized in subsequent basic and detailed engineering stages of the innovation process on the basis of additional experimental information, detailed calculations, standards, expert opinion, available equipment and construction capabilities, and other input. It is most desirable that this initial conceptual process scheme prove to be a better starting point for the remainder of the innovation process than other possible flowsheets. A flowsheet does not exist until it has been first synthesized.

Conceptual process engineering does not involve the invention of chemistry. However, it may very well help to select from among alternative candidate chemistries on the basis of what appears to be required to implement each chemical alternative in an industrial environment.

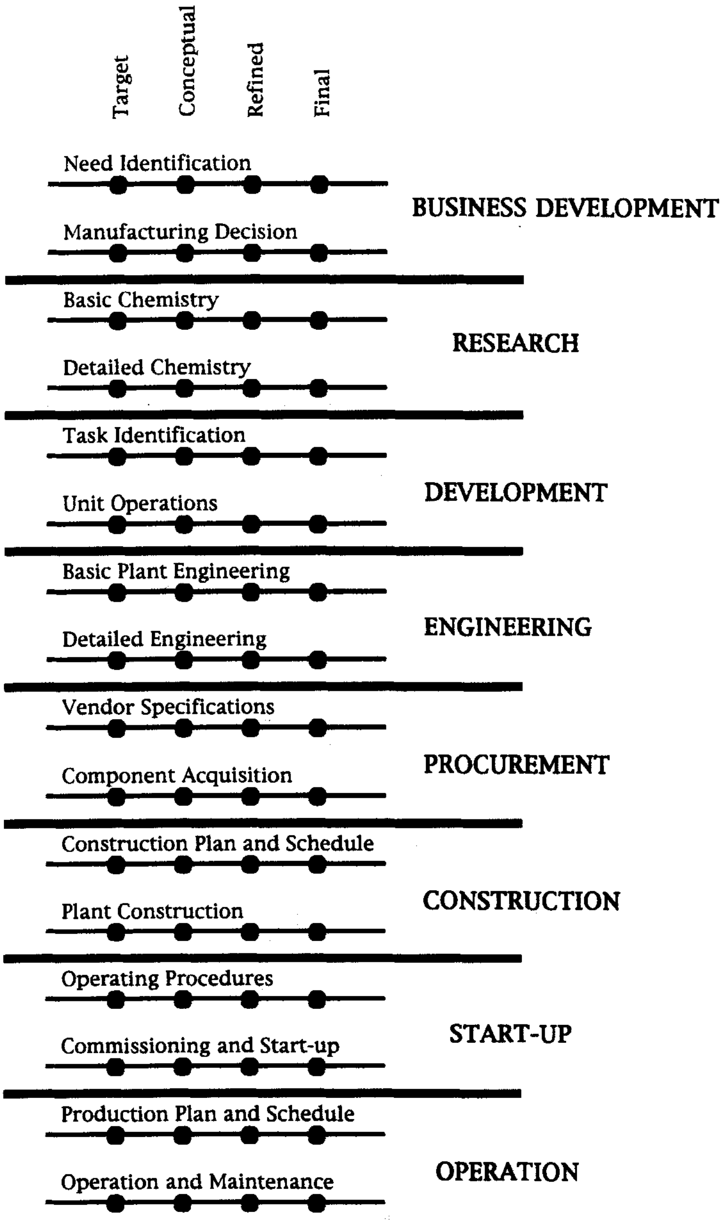

FIG. 5. Organization aligned with innovation stages.

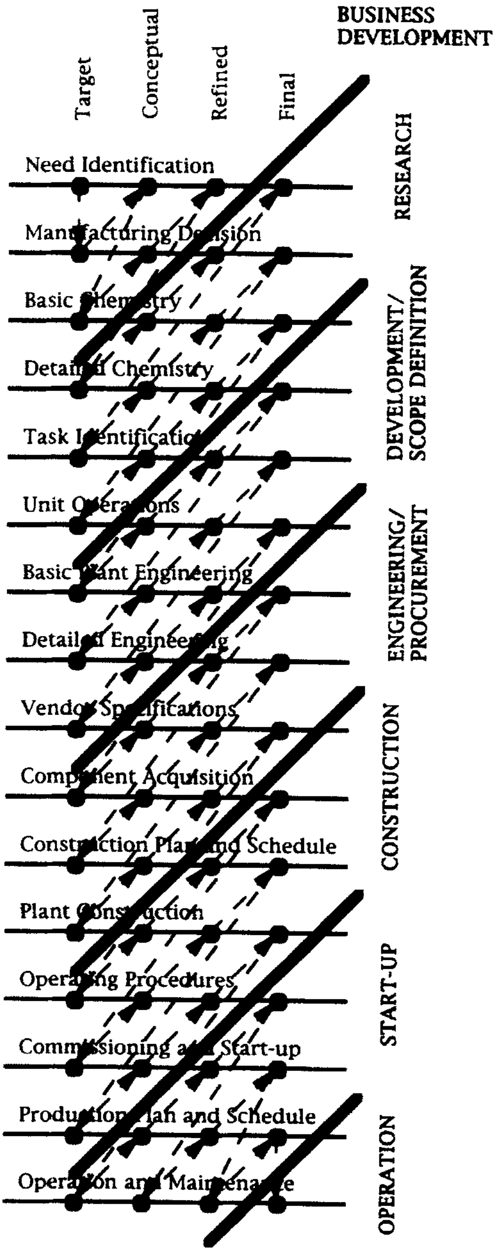

FIG. 6. Organization aligned with innovation journey.

There appears to be three fundamental approaches to the synthesis of chemical process flowsheets. The first, *systematic generation,* builds the flowsheet from smaller, more basic components strung together in such a way that raw materials eventually become transformed into the desired product. The second, *evolutionary modification,* starts with an existing flowsheet for the same or a similar product and then makes modifications as necessary to adopt the design to meet the objectives of the specific case at hand. The third, *superstructure optimization,* views synthesis as a mathematical optimization over structure; this approach starts with a larger superflowsheet that contains embedded within it many redundant alternatives and interconnections and then systematically strips the less desirable parts of the superstructure away.

Systematic generation from scratch sometimes leads to solutions that are already known. In the absence of the exhaustive generation of all alternatives, systematic generation cannot guarantee optimality. Likewise, the quality of solutions generated by evolutionary modification depends critically on the starting flowsheet as well as on the methods used to study and perturb it, thus rarely leading to clever designs. On the other hand, superstructure optimization offers the promise of simultaneous optimization of structural as well as other design parameters. However, it requires a starting superstructure from somewhere (which for some simple problems may be implicit in the formulation), as well as very extensive computational capability. Both *algorithmic* and *heuristic* methods have been proposed for the systematic generation and evolutionary modification approaches, whereas *mathematical programming,* especially mixed-integer nonlinear programming, is the obvious method for the superstructure optimization approach. Evolutionary modification is the approach traditionally used most frequently by conceptual design engineers. Systematic generation has been the basis of most of the academic process synthesis research results that so far have seen industrial application. Superstructure optimization, while not especially practical with current computing resources, offers tremendous potential for the future.

A. SYSTEMATIC GENERATION: MEANS–ENDS ANALYSIS

Although much (perhaps most) conceptual process design is in fact done by the evolutionary modification of whole or partial existing flowsheets, it is sometimes necessary to generate a flowsheet from scratch.

In dealing with the manufacture of materials, a number of physical and chemical properties of materials or streams of materials are of particular interest. These properties include the *molecular identity* of the material; the *amount* involved; the *composition* or purity of the material; its thermodynamic *phase;* its *temperature* and *pressure;* and possibly size, shape, and other physical *form* characteristics. Available raw materials, intermediates, and the desired product materials are all characterized by these same parameters of identity, amount, composition, phase, temperature and pressure, form, and so on.

The raw materials may be considered an *initial state.* The desired product may be considered a *goal state.* If the value of a particular property of a raw material is different from the desired value of the corresponding property in the desired product, a *property difference* is detected. One of the purposes of an industrial chemical process is to apply technologies in sequence such that these property differences are systematically eliminated and the raw materials thereby become transformed into the desired products. The systematic detection of state differences and the elimination of such differences through the application of appropriate corrective *operators* is the essence of a general goal-directed problem-solving paradigm known as *means–ends analysis* (Simon, 1969). The detection of property differences and the identification of methods to eliminate such differences is also the goal of the task identification stage of the innovation process. The specification of actual equipment to implement these difference-elimination methods is the goal of the unit operations stage of the innovation process.

In many cases there are fairly obvious technologies (in some cases tautologies) for reducing or eliminating property differences. Familiar examples include *chemical reaction* to change molecular identity, *mixing* and *splitting* (and *purchase*) to change amount, *separation* to change concentration and purity, *enthalpy modification* to change phase, temperature, and pressure, etc. In some cases such close relationship exists between a difference-elimination method and an obvious type of equipment to implement the method that many designers think directly in terms of *equipment* when developing conceptual process designs, thereby translating property differences directly into unit operation solutions.

Generally, property-changing methods or tasks (the operators in means–ends analysis terminology) are applied to existing streams to produce new streams closer in properties to the desired product material. Most often operators are applied and the conceptual design is developed in a *synthetic, forward,* or *opportunistic* direction, in the same direction as material flow, from the source of raw material through any reactions and toward final products, with recycles added as an afterthought. These operators are applied until no further property differences remain. This is in contrast to the generally *retrosynthetic* (or *backward* from the desired product) approach taken by most chemists in the development of complex organic reaction sequences. In conceptual design, both the starting materials and the desired product are known, while in reaction sequence development sometimes only the desired product is known. Opportunistic approaches are possibly somewhat more comfortable because they are in a sense "anchored." At any point, the design generated is a feasible consequence of the raw materials and the operators chosen. Given the remaining goals, what must I do now? Where can I go from here?

One complication is that often property-changing operators can only be applied to a stream when certain other properties of the stream are within specified values, which may not be true at the time. For example, a method to select only crystals greater than a given size can be applied only if a stream contains solids. Similarly, a separation method expected to exploit relative volatility differences can be applied only if enthalpy conditions permit simultaneous liquid and vapor phases. If the *preconditions* for the immediate application of an operator believed to be useful are not met, a new design *subproblem* may be formulated whose objective is to reduce property differences between the initial stream and the conditions necessary for the application of the operator. This *recursive strategy* is a common feature of the means–ends analysis paradigm.

Alternative solutions are generated when more than one operator is identified that can reduce or eliminate a property difference. The decision of which operator to choose might be made on the basis of some evaluation at the time the operators are being examined. Alternatively, each may be chosen separately, then the consequences followed separately (leading to alternative design solutions) and each final solution evaluated. Yet another possibility is that all feasible alternative operators are selected and applied in parallel, leading to a redundant design or superstructure. At the end of the design process, the superstructure is reevaluated in its entirety, and the less economical redundant portions eliminated.

Another problem is that an operator may not completely eliminate a property difference. In such a case, another operator for the same property difference may need to be specified. For example, a drying task might follow a filtration task to completely remove a liquid contaminant from a desired solid. Furthermore, the application of an operator may not exclusively change a single property difference. The side effects of a difference-elimination method may change other properties as well. For example, the application of a concentration-changing separation method may also change the pressure or the temperature or the phase. In the worst case, the paradigm may not converge, but rather oscillate repeatedly, eliminating one difference but creating a second difference as a side effect. Other problems include the possibility of not finding a method in the repertoire that can eliminate a property difference. If this occurs in a recursive subproblem (that is, while attempting to meet the preconditions for the application of another operator), it may not be fatal if an alternative operator for the original problem is available. The paradigm fails, however, if no operator can be found to correct a property difference at the top level.

B. HIERARCHICAL APPROACHES

Alternative designs may also result when more than one property difference exists between a stream and its goal and these differences are eliminated in different orders. However, a natural *hierarchy* among property differences seems to exist in the chemical process domain. The hierarchy is the same as the order in which the differences were previously mentioned: identity first, then amount, then concentration, then phase, then temperature and pressure, then form. The hierarchy arises because properties lower in the hierarchy are often more readily manipulated in order to satisfy the preconditions for the application of difference-elimination operators for properties higher in the hierarchy. Alternative solutions are still possible when multiple differences at the same level in the hierarchy exist: for example, when both temperature and pressure need to be changed or when a mixture of more than two components is to be separated and sent to different destinations.

Therefore, a natural approach to systematic generation of a conceptual design to a desired product given available raw materials using both the principles of hierarchical properties and the most obvious difference-elimination methods would be to attack identity differences first (resolved with reaction methods), then amount (resolved with mixing, splitting, or purchase methods), then composition (resolved with mixing or separation methods), then phase, temperature, and pressure (resolved with enthalpy-changing methods), then form properties (size, shape, etc.), as in Fig. 7. Preconditions are met by adjusting properties lower in the property hierarchy if necessary. The conceptual designer, thinking directly in terms of unit operations rather than tasks, would consider reactors first; then mixers, columns, decanters, filters, dryers, and the like; then heat exchangers, pumps, compressors, valves, etc., to satisfy goal properties for the desired product as well as the subgoals that arise as preconditions for the application of unit operations higher in the equipment hierarchy. As before, alternative designs result when alternative unit operations can be used to resolve the

Property Difference	Resolution Method
Molecular Identity	**Reaction**
Amount	**Mixing, Splitting, Purchase**
Composition	**Mixing, Separation**
Phase	**Enthalpy Change**
Temperature, Pressure	**Enthalpy Change**
Form	**Various**

FIG. 7. Property hierarchy and common difference-elimination methods.

same a property difference, or when multiple property differences at the same level in the property hierarchy are resolved in different orders.

The classical means–ends analysis paradigm, although recursive in nature, is nevertheless not iterative. However, for efficiency, many chemical processes exploit *recycle* of mass and energy. Recycle introduces a serious complication since even the existence, much less the properties, of the mass or energy to be recycled may not be known at the time property differences at the point of recycle were resolved. Sometimes the interactions are minimized by the way recycle tasks are implemented in equipment, as in indirect energy recovery with a heat exchanger. Other times the interactions are more severe, as in mass recycle by mixing. Recycles are generally considered at the end of one pass through the conceptual design. If the effects of recycle are severe, particularly with respect to structural validity, the design may need to be repeated from the beginning, this time assuming the existence of the potential recycle, until the design process converges.

C. THE ADAPTIVE INITIAL DESIGN SYNTHESIZER

A number of experimental chemical process synthesis procedures have been developed with a hierarchical architecture based on both the hierarchical structure of the chemical innovation process and the physical property hierarchy. Examples include PIP (Kirkwood *et al.*, 1988) and BALTAZAR (Mahalec and Motard, 1977). This hierarchical approach is also evident in the Pinch Technology Onion Diagram (Linnhoff and Ahmad, 1983). The hierarchical approach is also the foundation for the two early textbooks on process synthesis methodology (Rudd *et al.*, 1973; Douglas, 1988).

One such procedure, the Adaptive Initial DEsign Synthesizer (Siirola and Rudd, 1971), is instructive. The AIDES process synthesis hierarchy consists of *reaction path* (chemistry identification), *species allocation* (target mass flows among raw materials, reactions, products, and wastes including possible recycle of incompletely converted reactants), *task identification* (property difference detection and elimination method identification), *task integration* (association of tasks with actual unit operations, including combining adjacent tasks and also combining complementary tasks), *utilities system* (basic plant design and preliminary operating cost estimation), and *equipment design* (target level equipment design and preliminary capital cost estimation) (Fig. 8). With characteristics similar to the overall innovation process of which process synthesis is a part, the initial stages of the AIDES hierarchy are less well defined, but have a greater impact on the overall outcome of the design. Stages later in the procedure are more constrained and better defined, but have comparably less potential for impact. The stages clearly interact with each other. Decisions at one stage are

- **Reaction Path**
- **Species Allocation**
- **Task Identification**
- **Task Integration**
- **Utilities System**
- **Equipment Design**

FIG. 8. AIDES process synthesis hierarchy.

constrained by choices made at earlier stages, but optimal choices are influenced by how all the subsequent stages prove to be executed. Similar to the innovation process, passing one time through each stage sequentially does not generate exceptional results. Better are schemes to iterate among the stages possibly at different levels of detail. Better yet is the possibility of solving several stages simultaneously. The AIDES procedure employed the means–ends analysis paradigm. A key feature is that task identification, task integration, and equipment design are discrete and separate activities. In particular, and contrary to common conventional design practice, property differences are not resolved directly in terms of specific equipment.

The software version of the AIDES process synthesis procedure (Siirola, 1970) was one of the first applications of *artificial intelligence* in chemical engineering, exploiting dynamic storage allocation, linked data structures, and symbolic manipulation. More importantly, it employed a number of techniques at each stage of the synthesis hierarchy to anticipate certain downstream effects including a *linear programming* formulation for species allocation which sought to minimize the difficulty of the resulting separation problems that would arise because of the allocation of species in a source stream to different destinations. It also employed *multiobjective programming* to prioritize both the property differences to be addressed and the selection among competing property difference-elimination methods. It also used a *targeting* algorithm to maximize energy recovery and minimize utility requirements. Unlike the later PIP program and other expert systems approaches, however, it was not overtly rule-based. The automated means–ends analysis implementation was biased to propose obvious difference-elimination methods for each class of property differences, and preconditions were met, if necessary, by adjusting properties *lower* in the property hierarchy. For example, separations might be identified in order to get appropriately pure feed conditions for reaction tasks, and temperatures might be adjusted in order to get better conditions for separation tasks. However, AIDES would never consider changing the identity of a species to get better conditions for a separation, as that would have violated the property hierarchy. Multiple design alternatives were generated by applying different operators (for example, by exploiting different physical phenomena to effect a separation) and by attacking multiple differences at the same level of the property hierarchy in different orders. The AIDES program can interact as tightly or loosely with the conceptual designer as desired. It can accept advice or criticism, have its recommendations overruled, or formulate its own solution independently. The program is adaptive, and if feedback is provided it can *learn* from its performance (by adjusting internal multiobjective weighting parameters).

The AIDES program definitely has its limitations. The internal representation of solution thermodynamics is simplistic; it cannot, for example, handle azeotropes. Also, the species allocation, once proposed, is adhered to too rigidly and its task integration capability is somewhat weak. In industrial situations, its performance is judged to be good, but only slightly better than conceptual designers that resolve property differences directly in terms of common unit operations.

However sophomoric the AIDES computer implementation, the hierarchical process synthesis procedure itself offers three important lessons which have had a significant industrial impact—lessons that can be used even manually by conceptual process designers.

1. *Keep task identification distinct from task integration and equipment design.* Think carefully about what needs to be accomplished (tasks), then consider how it will be accomplished (equipment). In other words, remember the architect's maxim, "Form follows function."
2. *Do not adhere inflexibly to the property hierarchy.* The fact is, any property, not just those lower in the property hierarchy, might be worth adjusting as a precondition to the application of an operator, including properties at the same level or even at a higher level in the hierarchy. For example, identity change can sometimes be a route to resolving concentration differences, either because it turns an offending species into something else easier to separate—or better yet into something useful or at least innocuous so that it doesn't have to be separated at all.
3. *Be wary of convenience.* Much chemical laboratory practice is driven by the desire for high yields at ambient conditions. Thus there is a tendency toward irreversible procedures. For example, acid chlorides are used rather than acids for esterifications and amidations to avoid reaction equilibrium limitations. But such laboratory convenience comes at the price of more energetic and expensive reagents, and the production of very low-energy-state coproducts, such as volatile gases and salts, which are often of little economic value and large environmental impact. In a similar vein, it is more convenient in manufacturing operations to avoid the complications of material and energy recovery and recycle. Again, such luxury is often at the expense of greater raw material and energy requirements and costs, as well as greater thermal and material loading on the environment.

IV. A Total Flowsheet Example

Consider the following industrial example of a process for the production of methyl acetate, first synthesized in the conventional conceptual process design manner, then modified using evolutionary approaches, and finally resynthesized using the lessons from the hierarchical process synthesis procedure experience. For simplicity here, only identity, amount, and composition differences are resolved (that is, temperature and pressure changers will be ignored).

A. Conventional Conceptual Process Design

Methyl acetate may be produced by the elementary equilibrium-limited acid-catalyzed esterification reaction of methyl alcohol and acetic acid which also by-produces water. The reaction is nearly athermic, with an equilibrium constant on the order of unity.

$$\text{Acetic acid} + \text{methanol} \Leftrightarrow \text{Methyl acetate} + \text{water}$$

Acetates form azeotropes with both their corresponding alcohols and water, but the azeotrope with water is generally heterogeneous and lower-boiling. Therefore, acetates may be produced from the acetic acid and alcohol in a reactor fitted with a fractionating column which removes the acetate–water heterogeneous azeotrope overhead while forcing unreacted acetic acid and alcohol back to the reactor. The acetate and water are separated in a decanter and purified additionally as necessary (Witzeman and Agreda, 1993). However, this textbook flowsheet is not applicable to the production of methyl acetate because the methyl acetate–methanol azeotrope boils at a lower temperature than the methyl acetate–water azeotrope; furthermore, the latter azeotrope is unique among acetate–water azeotropes in that it is homogeneous. Therefore, a new conceptual process design must be generated.

If methanol and acetic acid are available as raw materials and methyl acetate is the desired product, according to the property-difference hierarchy, an identity difference is first detected between the desired product and each of the raw materials. A known chemical reaction operator, namely the esterification reaction, can be applied to a mixture of the raw materials brought to the proper conditions to produce methyl acetate and eliminate the identity difference between the reaction effluent and the desired product. Thinking directly in terms of equipment, this operator may be immediately implemented, for example, as a stirred tank reactor.

The effluent of the reactor contains product methyl acetate, by-product water, and significant unreacted methanol and acetic acid. The species allocation is trivial; methyl acetate is directed toward the product, water to wastewater treatment, and both methanol and acetic acid are recycled back to the reactor in order to improve feedstock usage efficiency. Since the reactor effluent does not meet the composition specifications of any desired destination, separation tasks need to be identified to eliminate composition differences between the reactor effluent and the various destinations of its various components.

The species are all liquids at ambient conditions. Normal boiling points vary from 57 to 118°C. It might be possible to exploit volatility differences—that is, distillation—to effect the required separations. Water and methyl acetate form a low-boiling azeotrope at 95 wt% methyl acetate boiling about 56°C. Water and methyl acetate are immiscible between about 23 and 92 wt% methyl acetate. Thus, as previously stated, this binary system is unusual in that it exhibits both azeotropy and heterogeneity, but the azeotrope is homogeneous. Also, methanol and methyl acetate form an azeotrope at 81 wt% methyl acetate boiling even lower at 54°C. There are no other binary or ternary azeotropes in the system, although water and acetic acid almost form an azeotrope exhibiting a tangent pinch on the water end.

How should the required separation system be synthesized? One common sequencing practice is to perform the easiest split first. Considering the boiling point distribution of the components and azeotropes and thinking in terms of equipment, an opportunistic first split is a distillation column in which all of the methyl acetate, all of the methanol, and as much water as azeotropes with the methyl acetate are taken as overhead. All of the acetic acid and the remaining water are taken as bottoms (Fig. 9). At first, this may not seem like much

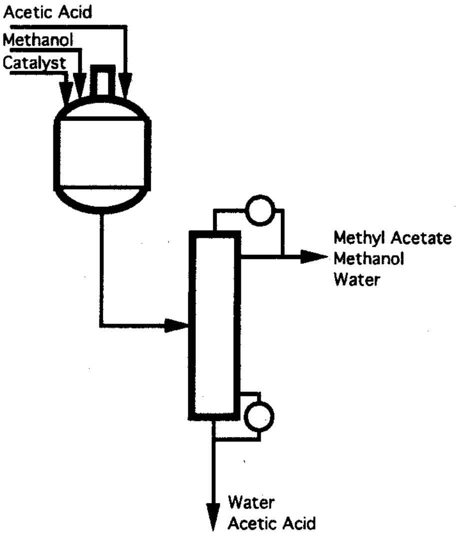

FIG. 9. Methyl acetate: Opportunistic first split.

progress since the composition of neither column effluent stream is directly suitable for any destination. However, the bottom stream is essentially a binary mixture of acetic acid and water, and the separation of that mixture is a problem that has been solved before.

The separation of acetic acid–water mixtures is a feature of a number of large-scale industrial processes including manufacture of cellulose acetate, terephthalic acid, and even acetic acid itself. Simple binary distillation is feasible but expensive because of the tangent pinch at the top of the column. The literature, however, contains a number of flowsheets for this separation. Most involve azeotropic distillation with an entrainer sometimes preceded by solvent extraction to circumvent the tangent pinch, but each is different because of the acetic acid feed concentration addressed and the specifics of the azeotropic entrainer or extraction solvent. One particular textbook flowsheet (Lodal, 1993) suitable for the acetic acid–water composition from the underflow of the splitting column in this example is shown in Fig. 10. The design consists of a solvent extractor, an azeo column and decanter, flash columns for solvent recovery from the extractor raffinate and the decanter aqueous layer, and a color column for purification of the azeo column underflow. All of the acetic acid appears in the color column distillate suitable for recycle back to the reactor. Water exits from the two flash columns suitable for discharge to the wastewater treatment system. Ethyl acetate in this case is used for both the extraction solvent and the azeo column entrainer, and is totally recycled within the flowsheet. Other flowsheets using different solvents or entrainers resulting in different configurations are also reported in the literature (Siirola, 1995).

This now leaves the overhead from the first splitting column. The homogeneous azeotropes in this mixture are a significant problem. However, a well-known pattern for breaking homogeneous azeotropes involves extractive distillation with a high-boiling solvent. Ethylene glycol is a suitable solvent in this case. A standard textbook configuration consisting of an extractive distillation column followed immediately by a solvent recovery column produces a pure methyl acetate product as distillate from the first column and a methanol–water mixture as distillate from the second. The methanol–water mixture may be separated in one final binary distillation column, producing methanol suitable for recycle back to the reactor and water suitable for discharge to the wastewater treatment system. Composition specifications of the product methyl acetate and the by-product water are met, and all unreacted methanol and acetic acid is recycled to the reactor. No additional identity, amount, or concentration differences exist and the conceptual design is now complete. The flowsheet includes one reactor, one extractor, one decanter, and eight distillation columns, and employs two different external mass separation agents (Fig. 11).

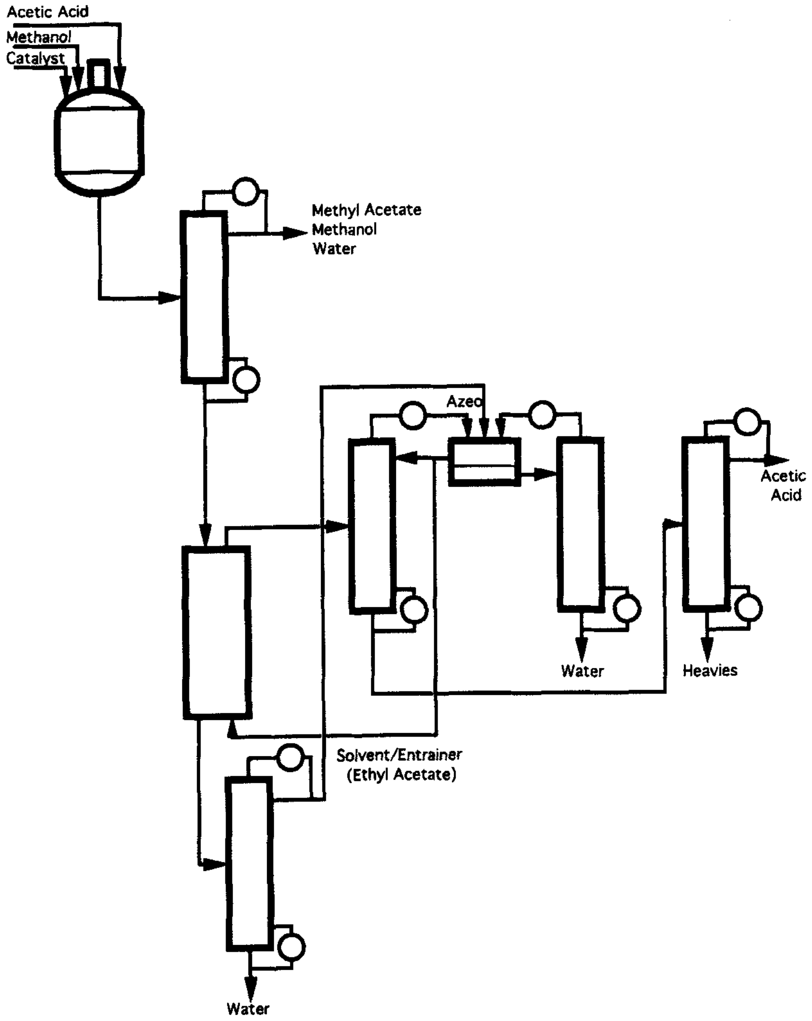

FIG. 10. Methyl acetate: Textbook acetic acid dehydration.

B. EVOLUTIONARY MODIFICATION

The classic flowsheet for the production of acetates was unsuitable for adoption for methyl acetate. However, the flowsheet generated by the conventional conceptual process design approach using literature schemes and standard patterns seems especially complicated for such simple chemistry involving so few components. This first flowsheet is a candidate for evolutionary modification.

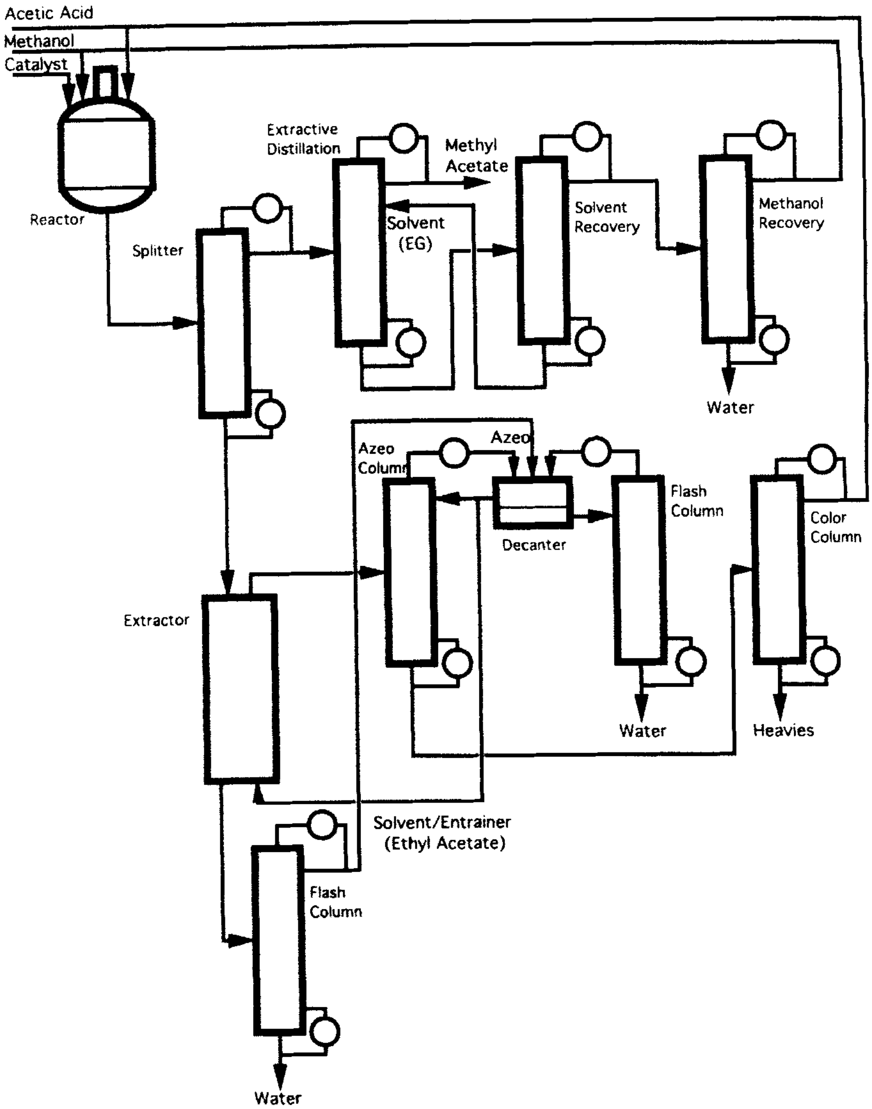

FIG. 11. Methyl acetate: Complete conventional flowsheet.

One evolutionary strategy is to search for thermodynamic inefficiencies to attack, for example, through exergy analysis. Another strategy is to search for structural inefficiencies including design redundancies. In our first flowsheet, both flash columns have similar feed compositions, both overheads go to the same decanter, and both underflows go to wastewater treatment. Therefore, these two columns can probably be combined. Second, the acetic acid color column, which was a feature of the textbook flowsheet copied for this part of the design, may not be required. Upon investigation, it is found that the original flowsheet was designed to handle dirty acetic acid–water mixtures containing minor unspecified high-boiling contaminants. In the present case, it is determined that the acetic acid composition in the bottoms of the azeo column is sufficiently pure to be recycled to the reactor directly, and thus another column is eliminated. Enthalpy-changing tasks were not considered in this example, but if they had been, a number of energy recovery opportunities might also be evaluated using heat integration techniques.

Another point of concern is the two external mass separating agents. Both were chosen from textbook patterns and flowsheets. Although both are readily recovered and recycled within their respective sections of the process, a common evolutionary heuristic suggests that if possible, mass separation agents should be chosen from among components already present within the system. First, consider the agent that must serve the dual role of extraction solvent and azeotropic entrainer. The only possible component in the original system that might work as an acetic acid extraction solvent is methyl acetate. However, as previously noted, the region of immiscibility is not large, and depending on the feed composition, it might not be a feasible extraction solvent at all. Moreover, its azeotrope with water is not heterogeneous, so the subsequent azeotropic distillate cannot be broken by decantation. Methyl acetate appears to be totally unsuitable as an internal mass separating agent for this problem.

On the other hand, considering the second mass separation agent, acetic acid is a possible candidate as an extractive distillation solvent to break the methyl acetate–water azeotrope. If acetic acid were substituted for the ethylene glycol, a methyl acetate–methanol mixture would be produced as distillate, and an acetic acid–water mixture would be produced in the underflow. The methyl acetate–methanol mixture (which also forms a homogeneous azeotrope) may be separated into methanol and sufficiently pure methyl acetate to meet the desired product specification either by binary distillation at reduced pressure (where the azeotrope disappears) or by pervaporation. The acetic acid–water mixture is sent directly to the previous part of the flowsheet with the underflow of the first splitting column, thus saving the solvent recovery column. Furthermore, the splitting column and the extractive distillation column might also be combined since one is fed from the other and both underflow to the same place. The final evolved design, Fig. 12, thus contains four fewer columns than the original flowsheet.

The relative performances of the original and evolved designs can be determined only by analysis. The evolved design has fewer columns, but has the added complication of either refrigeration or membranes, and the amount of acetic acid which must be separated from water is much increased. The results

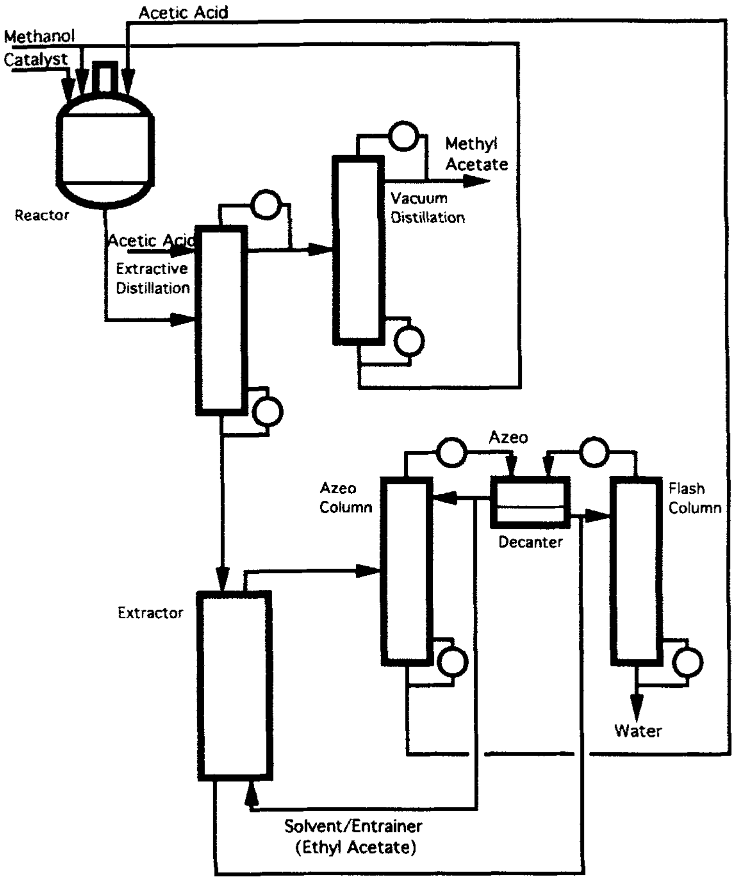

FIG. 12. Methyl acetate: Evolutionary modification.

of detailed simulations indicate in fact that combination of the flash columns and elimination of the color column is advantageous, but that using acetic acid as an extractive distillation mass separation agent in this case unfortunately is not.

C. THE HIERARCHICAL PROCESS SYNTHESIS PROCEDURE

If the hierarchical means–ends analysis synthesis procedure is applied to the methyl acetate problem, the task identification, task integration, and equipment design stages are kept completely separate. Following the property-difference hierarchy, an identity-changing reaction task (Task A) is identified first, as before. When examining the differences between the result of this reaction task application and the product methyl acetate and by-product water destinations, two sets of composition differences are detected. For the methyl acetate destination, acetic acid, water, and methanol must all be essentially completely removed. Similarly for the water destination, methyl acetate, methanol, and acetic acid must all be essentially completely removed. One way to attempt to resolve these concentration differences is separation exploiting relative volatility. Considering the methyl acetate destination first, the separation of acetic acid should be relatively easy because of the large boiling point difference. However, removal of the water and methanol may be difficult because of homogeneous azeotropes. Turning to the water destination, removal of both methyl acetate and methanol should be easy because of the large boiling point differences (the methyl acetate–water azeotrope is not an issue where the goal is to produce pure water composition, although it can affect meeting amount goals), but removal of the acetic acid may be difficult because of the tangent pinch.

Relative volatility is not the only property that might be exploited to address composition differences. For example, if the property hierarchy is ignored, then an identity-changing esterification reaction task with acetic acid might be chosen to eliminate methanol from the methyl acetate stream (Task B). Similarly, an esterification reaction task with methanol might be chosen to eliminate acetic acid from the water stream (Task C). Furthermore, as in the previous example, solvent-enhanced distillative separation using acetic acid might be used to break any methyl acetate–water azeotrope (Task D), and conventional distillative separation might be used to separate both methyl acetate (Task E) and methanol (Task F) from water, and acetic acid from methyl acetate (Task G). What results then are three reaction tasks (one to produce the desired product and two to remove offending species from destination streams), one solvent-enhanced distillative separation task (using acetic acid), and three conventional distillative separation tasks (Fig. 13).

There are a number of ways in which these tasks may be integrated into actual processing equipment. It turns out in this case that considering hydraulics, material flows, and energy requirements, all seven tasks may be integrated into one single piece of equipment (Fig. 14). All three reaction tasks are integrated into the center of a column, with the main reaction Task A in the middle, the methanol-removing reaction Task B above, and the acetic acid-removing reaction Task C below. Extractive distillation Task D sits above the reaction tasks with extraction solvent knockdown Task G above that. Low-boiler separation from water Tasks E and F are combined and sit below the reaction zones. Proper operation requires that the extractive zone must be above and distinct from the reactive zones (otherwise the azeotrope of water and methyl acetate formed at the top of the reaction zone cannot be broken), which requires that the catalyst facilitating the reaction be either introduced between zones D and B (if homogeneous) or installed only below zone D (if heterogeneous). Although this single piece of equipment is admittedly complex, the resulting capital cost and operating cost of the single-column reactive–extractive distillation design for methyl

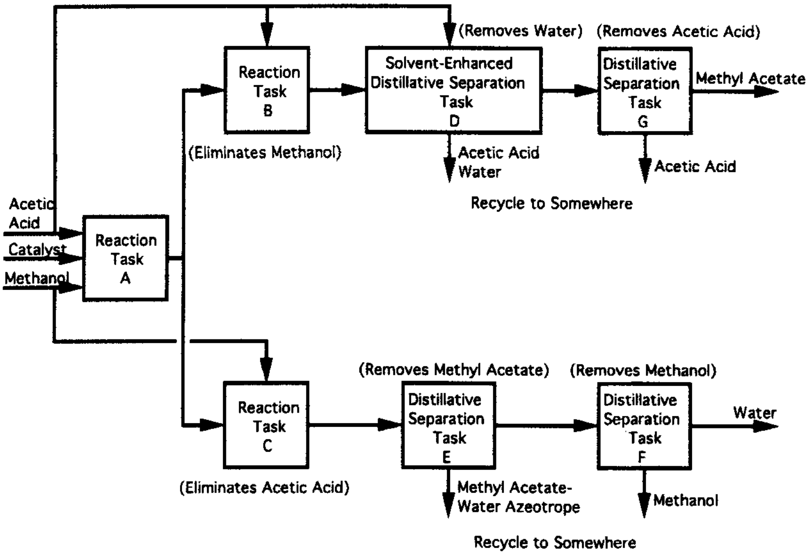

FIG. 13. Methyl acetate: Task identification.

acetate production are both just one-fifth of that for the optimized conventional flowsheet (Agreda and Partin, 1984). It is not clear how such a design might have been conceived without the explicit separation of task identification, task integration, and equipment design. The design is now in successful industrial operation in several plants.

To reiterate, most approaches to conceptual process design employ some kind of hierarchy to prioritize and order property-difference resolution. Sometimes the systematic generation of solutions from scratch is superior to copying existing patterns, tricks, and flowsheets. There may be some advantage to thinking of all properties as candidates for change to effect property-difference resolution, and selecting properties for change outside of the natural hierarchy sometimes leads to designs that are judged to be clever. Also useful is thinking primarily in terms of tasks rather than unit operations or equipment, being open to unique ways in which multiple tasks may be integrated into actual equipment, and being especially suspicious of laboratory convenience.

V. Energy Efficient Distillation Synthesis

A. HEAT-INTEGRATED DISTILLATION TRAINS

Heat integration, or the synthesis of heat-exchanger networks, is a classic problem of integrating complementary tasks at the task integration stage of the

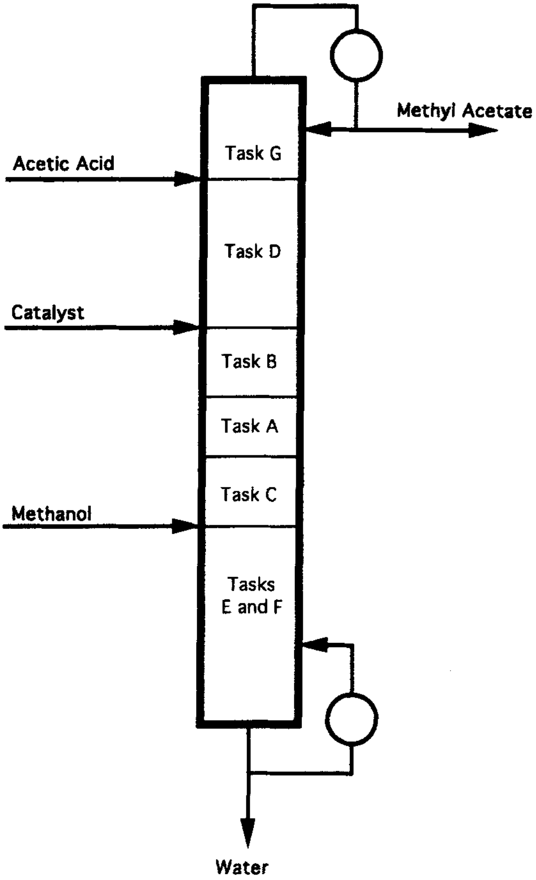

FIG. 14. Methyl acetate: Reactive extractive distillation process.

process synthesis hierarchy. The tradeoffs between capital and operating costs are well understood. Excellent procedures exist for targeting the expected amount of energy recovery, the remaining utilities requirements, and the expected capital requirements under a number of constraints even before any heat-exchanger networks have been actually designed. The *pinch* concept (Linnhoff and Hindmarsh, 1983) provides further guidance on restrictions on the placement of exchangers, the use of heat pumps, and determining situations where stream splitting and parallel heat-exchange designs are especially advantageous. It also introduced important visual representations of the relevant thermodynamics including *composite curves* and *grand composite curves.*

Procedures for the actual synthesis of heat-exchanger networks are somewhat less well developed; however, all of the standard process synthesis approaches—including systematic generation, evolutionary modification, and superstructure optimization using all the standard methods and tools including heuristics, expert systems, deterministic algorithms, and mathematical programming—have been discussed in the literature. Aspects of network flexibility, resiliency, and controllability have also been investigated. Available academic and commercial heat-integration software includes HEXTRAN, ADVENT, SUPERTARGET, MAGNETS, SYNHEAT, HEATNET, and others, as well as several proprietary codes.

Heat integration, which has been the most studied problem in the process synthesis literature, has also been the most widely applied synthesis technique in industry. The problem is well defined and highly constrained, being near the bottom of the synthesis hierarchy. Heat-integration technology has been applied to large industrial problems, including sections of individual processes, whole process plants, and even whole plant sites. The straightforward application of heat-integration technology to a typical chemical process design generally reduces the net present cost (NPC) by on the order of 10%, due largely to sensible heat recovery.

It is also well appreciated that greater savings may be possible if energy recovery is considered during, not just after, the task identification stage while the magnitudes and temperatures of the various heat sources and sinks are actually being determined. The following example considers the effect of closely coupling the conceptual detail level of the task identification stage of a distillation-based multicomponent separation problem with the target detail level of the subsequent enthalpy task integration stage.

The components involved in this example are proprietary, but the results are general (Siirola, 1981). During the species allocation stage of the process synthesis procedure, it was determined that each species of a particular four-component stream was required to be relatively pure at four different destinations. The components are liquids at ambient temperatures, have about equal relative volatility differences, and form no azeotropes. Distillative separation methods were selected to resolve all composition property differences. The feed stream composition was dominated (about 70%) by the heaviest component (D).

From exhaustive application of alternative simple distillation operators to all possible separations for this four-component system, or from application of ranked-list–based separations synthesis methods, it is easily shown that there are five different separation train structures for this four-component problem. Each can be generated systematically; or since this pattern of solutions is already well known, each can be written down immediately or design heuristics can be used to generate one or more of the structures expected to be most suitable. After each structure is synthesized, its performance can be analyzed and evaluated with a flowsheet simulator.

Probably the most common industrial solution to a problem of this type is a train of three columns, all operated at essentially atmospheric pressure, and each

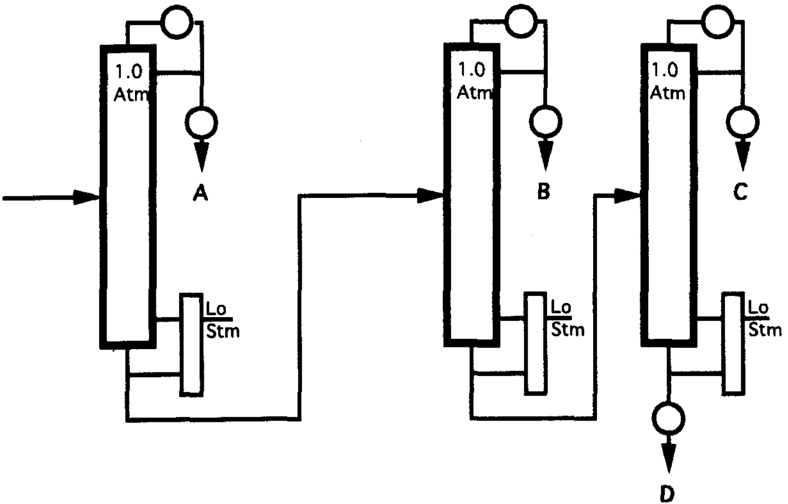

Relative NPC = 1.00

FIG. 15. Conventional sequence: Atmospheric operation.

operated to remove component relatively pure overhead one at a time (Fig. 15). In this example, this three-column train is the reference design to which we assign a relative NPC (including both capital and operating costs) of unity.

It turns out that, when all five sequences are analyzed, a different scheme has about 5% lower NPC than the reference case (Fig. 16). In this case, the

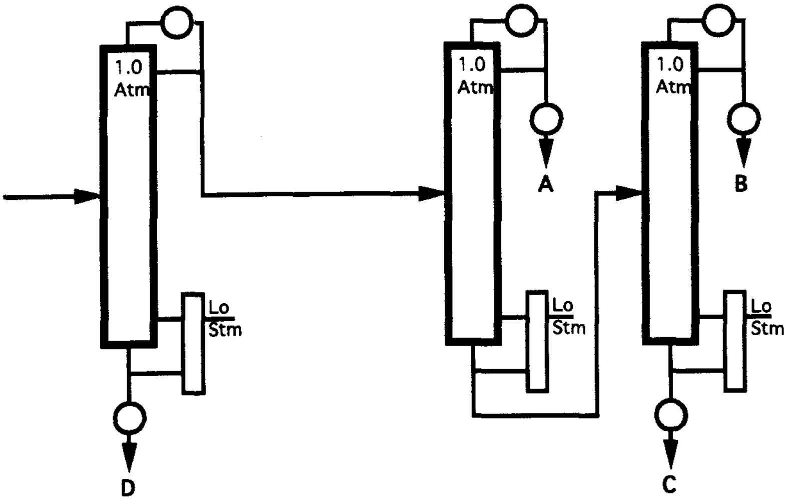

Relative NPC = 0.94

FIG. 16. Most plentiful first: Atmospheric operation.

dominant heavy species D is removed as the underflow in the first column, and then the remaining components are removed one at a time overhead. This is also the scheme that would have been generated by application of the most common distillation train sequencing heuristics (Seader and Westerberg, 1977).

Each of the distillative separation tasks is implemented as a separate standard single-feed distillation column with reboil and reflux. The thermal quality of each column feed might be adjusted externally to saturated liquid (although the size of the reboiler or condenser may be altered instead). There are also four product cooling (for storage) tasks. Altogether, therefore, there are potentially 13 enthalpy-changing tasks associated with this separations system (3 reboilers, 3 condensers, 3 feed conditioners, and 4 product coolers). Some of these tasks may be resolved with utilities and some integrated into a heat-exchanger network (see Fig. 17 for the separation sequence of Fig. 16). (In this example, there were no other significant enthalpy-changing tasks outside of the separation train to be simultaneously considered.) The NPC improvement from heat integration of the best separation scheme is on the order of 10% (15% reduction total from the reference case), more or less as expected. It should be noted that all of the heat-integration opportunities involve sensible heat.

Now one might return to the task identification stage, but at a higher level of detail to consider the possibility of separator designs not necessarily constrained to atmospheric pressure. In this example, the solution thermodynamics are such that the A/B split is favored by higher pressures and the C/D split is favored by lower pressures. In the best unintegrated case, the A/BC column is increased in pressure until just below where a more expensive reboiling utility would be required, while the ABC/D column is decreased in pressure until just

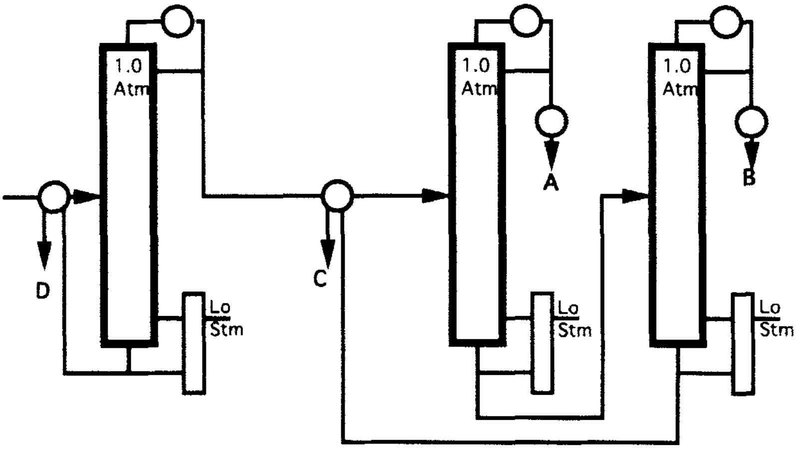

Relative NPC = 0.85

FIG. 17. Most plentiful first: Atmospheric operation, heat-integrated.

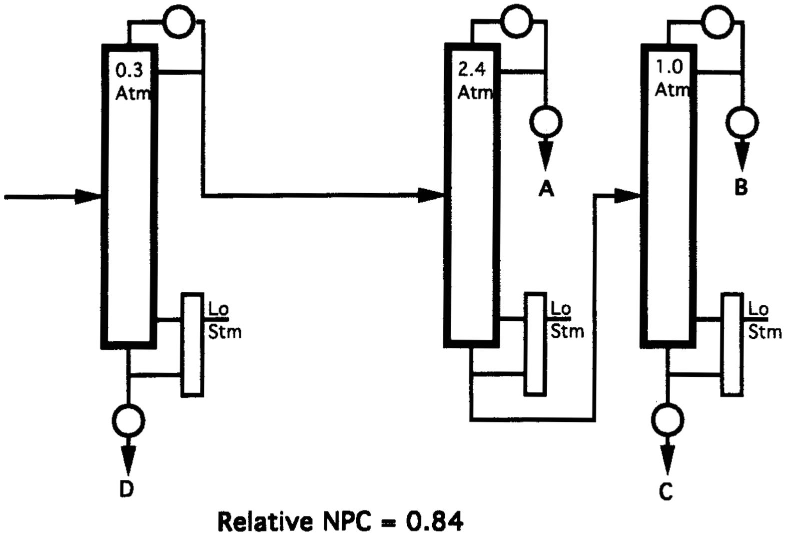

FIG. 18. Most plentiful first: Pressure-optimized.

above where a more expensive condensing utility would be required (Fig. 18). This pressure-optimized design has a 10% NPC improvement over the reference case due to flowsheet structure and column size minimization.

Once again the task integration stage may be revisited and a heat-exchanger network designed for the 13 enthalpy-changing tasks of the pressure-optimized design (Fig. 19). The result is a cost improvement similar to that for heat integrating all-atmospheric designs, and again represents only sensible heat recovery.

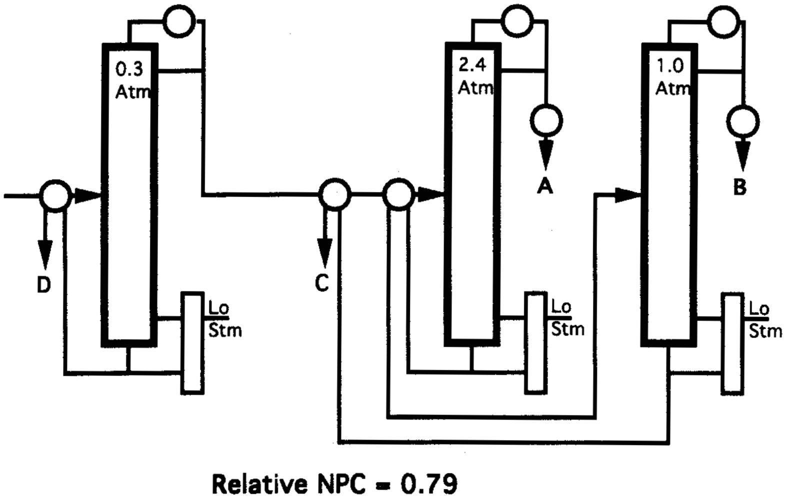

FIG. 19. Most plentiful first: Pressure-optimized, heat-integrated.

So much for the results of iterating between the task identification and task integration stages of the process synthesis procedure. What if the column pressure selection in the task identification stage were more tightly coupled with the heat-integration network generation of the task integration stage? A conceptual design level heat-integration procedure (Siirola, 1974) was made fast enough to be included within the objective function of the column-pressure optimization algorithm enabling such a coupling. Each objective function evaluation involved the design and costing of three columns; the determination of 13 enthalpy-changing tasks; the design and sizing of an associated heat-integrated energy recovery network including all reboilers, condensers, and other heat exchangers; and determination of the remaining utility requirements and costs. The result for this example of such a simultaneous column-pressure optimization with heat-exchanger network synthesis is shown in Fig. 20. The pressure of the ABC/D column is still low. However, the pressures of the remaining columns are both raised to a point where their reboilers require a higher cost utility and just below the point where they would require an even higher cost utility. However, latent heat from both of their condensers is recovered in reboilers for the first column. Operation of the two columns at these pressures does not make sense without the simultaneous latent heat integration. The NPC savings for the resultant design is a significant 35% reduction from the original reference case. The utility requirement was reduced by 50%.

The integration of enthalpy-changing tasks among different parts of a process can introduce interactions and control difficulties similar to other recycle

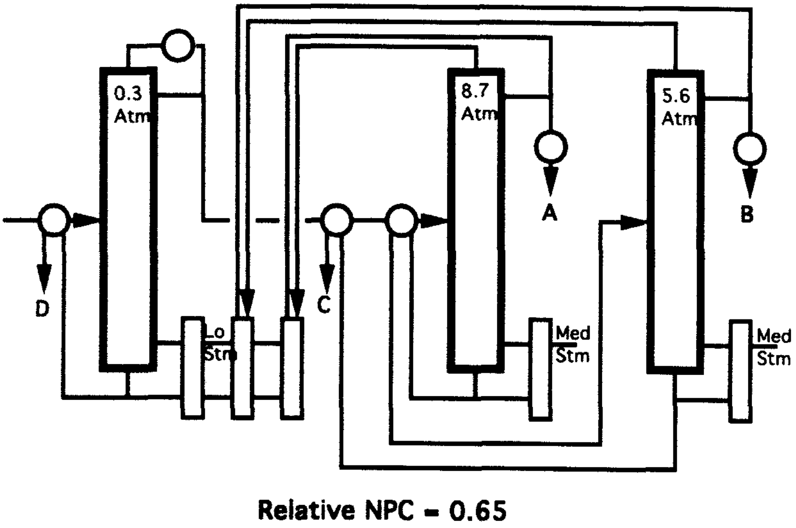

FIG. 20. Most plentiful first: Simultaneous pressure optimization and heat integration.

schemes. Sometimes these interactions may be decoupled through the use of an intermediate heat-transfer fluid between tasks, such as an intermediate-pressure steam, an imbalance of which could be made up from a mixture of standard utilities. Similarly, in a large plant, it may also be possible to effectively do heat integration mostly with the utility system, rather than within the process. This is especially true when there exists a large utility network with energy sources and sinks available at many levels of temperature. In such cases, it may be worthwhile to consider steam generation as a cooling utility. In the present example, simultaneously optimizing column pressures and heat integration while allowing the option of steam generation for credit produced the design in Fig. 21. The best separation train in this case turned out to be the more conventional single-species-at-a-time overhead arrangement. Higher pressure steam was used in all reboilers, and steam was generated in the condensers of two, for return to the utility network. There was only minimal sensible heat integration (for the fairly high-temperature first column). The NPC was again 35% below the original reference case. Because of perceived ease of operation, this was the design that was actually implemented industrially.

Utility requirement savings in the neighborhood of 50% and NPC savings in the neighborhood of 35% for distillation trains are typical of what can be expected from the tight interaction of the process synthesis task identification and task integration stages involving simultaneous distillation sequence synthesis, operating-parameter optimization, and heat-integration network synthesis.

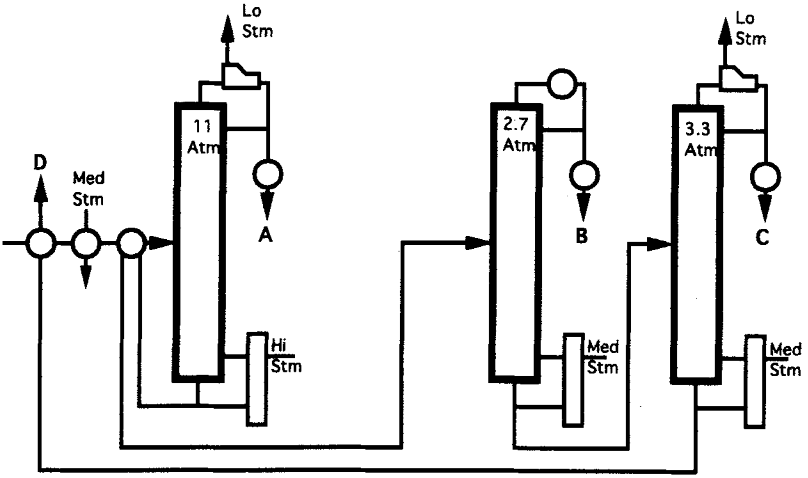

Relative NPC = 0.65

FIG. 21. Simultaneous pressure optimization and steam generation.

B. COMPLEX COLUMNS FOR BINARY SEPARATIONS

The previous section indicated that there can be a significant benefit to heat-integrated distillation sequences, with energy recovery either within the sequence or with the utility system. For processes that contain a single or a particularly dominant distillation, significant energy recovery may not be possible. Pinch analysis would suggest the possible use of a heat pump or vapor recompression in such cases. However, another possibility is to split the implementation of the distillative separation task into two columns in such a way that energy rejected from one might be recycled to the other, approximately halving the net utility requirement. The question is, given a separation that can be accomplished perfectly feasibly in a single simple distillation column, under what conditions is it more advantageous to implement the separation in two interlinked columns?

The motive for multiple-effect distillation for a single separation task is net energy reduction, generally at the expense of greater capital cost and operating complexity. For this to be accomplished, the condenser of one column must be at a higher temperature than the reboiler of the other. This may be done by operating the columns at different pressures, possibly assisted by the degree of separation (and hence temperature difference from bottom to top) designed for each column.

In one parametric study, six different designs (Fig. 22) were compared (Blakely, 1984). Design 1 was the single-column reference case. Design 2 consisted of a high-pressure and a low-pressure column. The feed was split and fed to both columns in parallel. Both columns performed the same sharp separation, producing pure products. The condenser of the high-pressure column was the reboiler of the low-pressure column.

In Design 3, the feed was sent to the high-pressure column, which produced a pure high-boiler bottom product but a mixed overhead product (a sloppy separation decreasing the temperature difference across the high-pressure column). The mixed distillate was then completely separated in the low-pressure column. Again, the condenser of the high-pressure column was the reboiler of the low-pressure column (heat integration in the same direction as flow).

In Design 4, the feed was sent to the low-pressure column, which produced a pure low-boiler distillate but a mixed underflow (again, a sloppy separation decreasing the temperature difference across the low-pressure column). The mixed bottoms was then completely separated in the high-pressure column. The condenser of the high-pressure column was the reboiler of the low-pressure column (heat integration in the opposite direction as flow).

In Design 5, both columns were operated at the same pressure. However, the first column produced a pure low-boiler distillate but a mixed underflow, thus allowing the use of a less expensive utility. The mixed bottoms was then com-

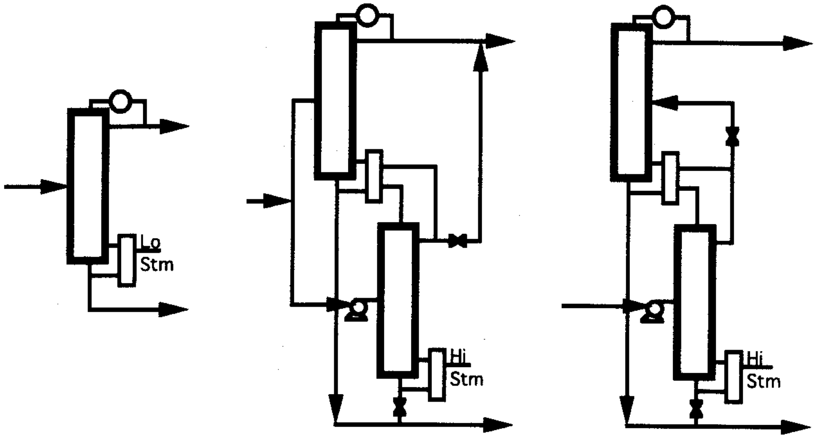

One Column
1

Split Feed
2

Sloppy First Top Split
3

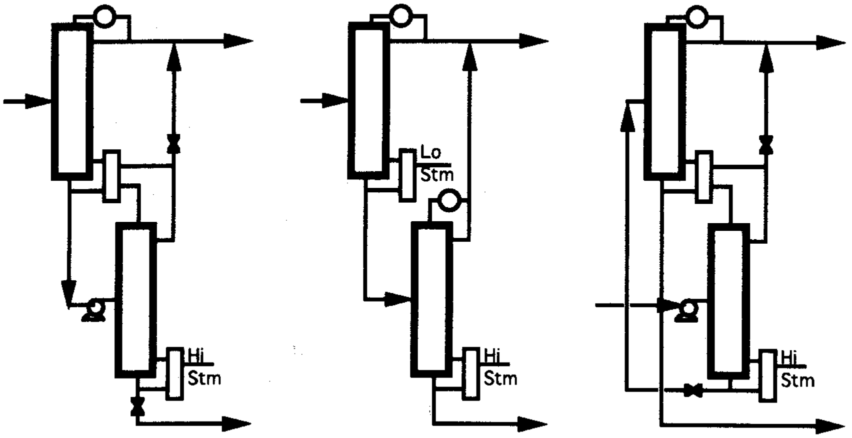

Sloppy First Bottom Split
Second Column Pressurized
4

Sloppy First Bottom Split
Two Steam Pressures
5

Sloppy First Bottom Split
First Column Pressurized
6

FIG. 22. Multiple-effect distillation configurations.

pletely separated in the second column, producing pure top and pure bottom products using a more expensive utility. No heat integration was involved.

In Design 6, the feed was sent to the high-pressure column, which produced a pure low-boiler distillate but a mixed underflow (again, a sloppy separation). The mixed bottoms was then completely separated in the low-pressure column.

The condenser of the high-pressure column was the reboiler of the low-pressure column (heat integration in the same direction as flow).

(The last obvious combination, with the feed sent to the low-pressure column producing a pure high-boiler bottoms product, with a mixed overhead completely separated in the high-pressure column, was not considered because it always requires a higher temperature utility than the other schemes.)

In the study, the pressure in the low-pressure column and the boiling point of the more volatile component were fixed. The composition of the mixed stream was chosen so that the heat duty of the high-pressure condenser was equal to that of the low-pressure reboiler. A parametric study on feed flowrate, feed composition, and relative volatility between the binary pair was conducted with a detailed process simulator. For each set of conditions, the NPC of each design was compared to that of the single-column reference case.

The specific results depend on the costs and temperatures of the utilities actually used in the study as well as on the capital cost correlations, and are not given here. However, the general trends are instructive. At very low flowrates, the economics of the single-column reference design was generally superior to all of the multiple-effect designs largely because at those scales the capital costs of the columns dominate the energy costs. However, at higher flowrates, there were conditions of feed composition and relative volatility for which the multiple-effect designs (but not the multiple-utility design) exhibited lower NPC than the single column. The shaded areas in Fig. 23 indicate the regions of feed composition and relative volatility in which the costs of each of the designs were at least 10% lower than the corresponding cost of the single-column design. There are still regions of high relative volatility and heavier feed composition where the single column is superior. But as feed flowrate increases, there are surprisingly more and more combinations of relative volatility and feed composition for which one or another of the multiple-effect schemes (each reducing the energy requirement by 50%) is worthy of examination. For some cases, the NPCs of the complex schemes were as much as 35% lower than the single-column reference design, savings similar in magnitude to typical savings in other multiple-column heat-integrated distillation designs.

In the production of methanol from synthesis gas, some water is by-produced from which the methanol product must be separated. This is the only major distillative separation task in the process, and because methanol has a high latent heat, it is very energy-intensive. For the feed conditions for this separation, the charts from this study indicate that Designs 2, 4, and 6 should be better than a single column, with Design 6 somewhat superior to the others. Steady-state simulation confirmed this prediction, and dynamic simulation confirmed the controllability of the tightly coupled column pair. Design 6 has been implemented industrially for this service with neither trim condensers nor reboilers (although multiple exchanger shells are required because of the low approach temperatures).

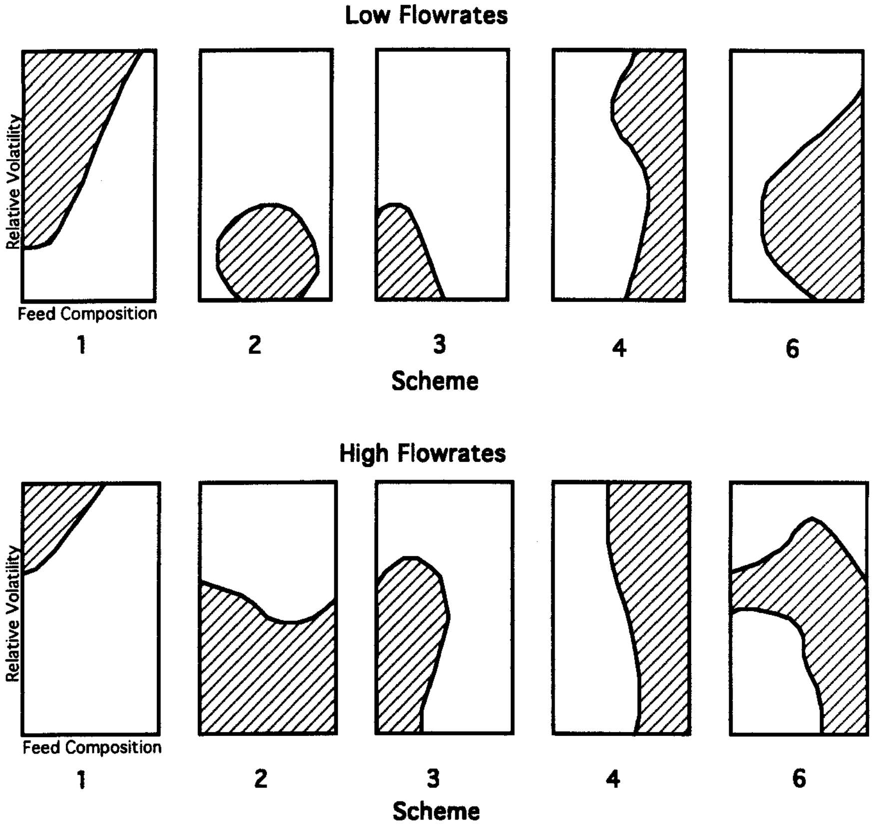

FIG. 23. Complex distillation: Regions of superior cost.

VI. Residue Curve Maps and the Separation of Azeotropic Mixtures

The synthesis tools required to propose property-difference resolution methods at the task identification stage of the process synthesis hierarchy depend in part on the nature of the system. In the case of distillative separations for resolving compositional differences, if the solution thermodynamics of the system are relatively ideal, a simple ranked-list volatility ordering may be sufficient to represent the thermodynamics for the purpose of separations flowsheet synthesis. Alternative separation trains can be synthesized at the targeting level of detail using list-processing techniques. However, if the system exhibits more nonideal thermodynamics involving azeotropes and regions of immiscibility, a more detailed representation of the thermodynamics may be necessary for the synthesis of separation schemes even at the targeting level.

Residue curve maps (RCMs) are phase equilibrium representations for systems involving azeotropes (Foucher *et al.*, 1991). RCMs are derived from an analysis of a single-stage batch still, but they may be applied to the understanding of the thermodynamic behavior of solutions and the behavior of continuous distillative separation operations on those systems. For a three-component system, the analysis can be plotted on familiar triangular composition diagrams. However, the analysis may be performed on systems having any number of components. Residue curves trace the composition of the stillpot of a single-stage flash in time as the vapor is slowly removed. Curves may be generated from any starting composition. As time proceeds, the curves move toward the composition of the last drop remaining in the stillpot. For some mixtures, this composition is the highest-boiling pure component in the system. For other systems, this composition may be a high-boiling azeotrope. For yet other systems, the final composition varies depending on the initial composition in the still, sometimes ending at one composition, sometimes ending at another. Residue curves may also be mathematically extrapolated backward in time to give the most volatile composition whose residue curve would pass through the specified feed composition. When extrapolated both forward and backward in time, all residue curves that originate and terminate at the same two compositions are said to define a *region*. The demarcation between regions in which adjacent residue curves diverge and terminate at different compositions is called a *separatrix*. Separatrices are related to the existence of azeotropes in the system.

Residue curve maps are being employed in an industrially useful method for the synthesis of separation schemes for azeotropic systems. To a first approximation, the compositions of the overhead and underflow of a single-feed continuous distillation column lie on the same residue curve; therefore, it is not possible for the composition profile of a continuous distillation column to cross into different regions. The boundary that a distillation composition profile can never cross is called a *distillation boundary*. The precise location of distillation boundaries is a function of reflux ratio, but distillation boundaries closely approximate the RCM separatrices. Distillation boundaries connect azeotropic and pure component compositions just as separatrices do. If an RCM separatrix exists, a corresponding distillation boundary will also exist.

Distillation boundaries together with the composition diagram boundaries define *distillation regions* which closely approximate residue curve regions. All pure components and azeotropes in a system lie on the boundaries of such regions. The most volatile composition on the boundary of a distillation region (either a pure component or a minimum-boiling azeotrope) is called the *low-boiling node* (the origin of all residue curves within the region). The least volatile composition on the boundary of a region (either a pure component or a maximum-boiling azeotrope) is called the *high-boiling node* (the terminus of all residue curves in the region). All other pure components and azeotropes on the region boundary are called *saddles* and are never exactly reached by any residue curve. Saddle ternary azeotropes are particularly interesting because they are less obvious to determine experimentally (being neither minimum-boiling nor maximum-boiling), and are not generally recorded in the literature; however, their existence implies distillation boundaries that may have an important impact on separation system design.

The thermodynamic nature of nonideal systems places limitations on the exploitation of volatility to effect separation tasks, and on what may be accomplished with simple distillation. However, if the resolution of a property difference requires that a composition be changed from one distillation region to another, a number of possibilities still exist. In some cases, some new stream may be mixed with the initial stream to produce a composition on the other side of the distillation boundary (although care must be taken to determine where this new stream came from and how it might be continuously regenerated). In other cases, the system thermodynamics and hence the distillation boundary are pressure-sensitive. If the system pressure is changed, a stream initially on one side of a boundary may find itself on the other side of the shifted boundary and in a different distillation region. Finally, other phenomena, unrelated to volatility, may be exploited to change the composition of a stream. One common phenomenon often used in this context is liquid–liquid immiscibility. Others are based on solid–liquid equilibria, such as crystallization and adsorption, and various kinetic phenomena, such as diffusion through membranes.

A. DEHYDRATION OF ETHANOL

The traditional textbook flowsheet for the dehydration of ethanol involves three distillation columns. However, some industrial implementations for exactly the same separation use only two columns with a corresponding savings in capital cost (Doherty and Knapp, 1993). How can these two alternative designs be systematically understood?

The ethanol–water system has a well known minimum-boiling azeotrope at about 96 wt% ethanol. As ethanol and water are completely miscible in all proportions, the ethanol–water azeotrope in homogeneous. If pure ethanol is desired from a feed richer in water than the azeotropic composition, the azeotrope must be broken.

There is a standard flowsheet pattern for breaking heterogeneous binary azeotropes, but that, of course, is not applicable here. It is possible to break homogeneous azeotropes by pressure shifting the azeotrope or by extractive distillation with a suitable high-boiling solvent, in this case for example, ethylene glycol. Here we wish to consider a third alternative, the use of an entrainer and the exploitation of a heterogeneous ternary azeotrope. A number of specific entrainers will work, including benzene, cyclohexane, and hexane among many others.

Consider a typical entrainer that forms binary minimum-boiling azeotropes with both ethanol and water as well as a minimum-boiling ternary azeotrope. The RCM for such a system indicates that all residue curves originate at the ternary azeotrope (the lowest-boiling composition in the system). However, some residue curves terminate at pure water, some at pure ethanol, and some at pure entrainer, thus forming three distillation regions, with distillation boundaries connecting each of the binary azeotropes with the ternary azeotrope (Fig. 24). The binary feed and the pure water product are in one region, and the pure ethanol product is in another region. Also, as shown in Fig. 25, the ternary system contains a region of liquid–liquid immiscibility that includes the entrainer–water azeotrope and the ternary azeotrope.

The distillation boundary separating the desired ethanol pure product from the feed and the desired water pure product may be called a thermodynamic *critical feature*. In all solutions to this problem, this boundary must be crossed somehow. Overcoming this critical feature is so important that a hierarchical

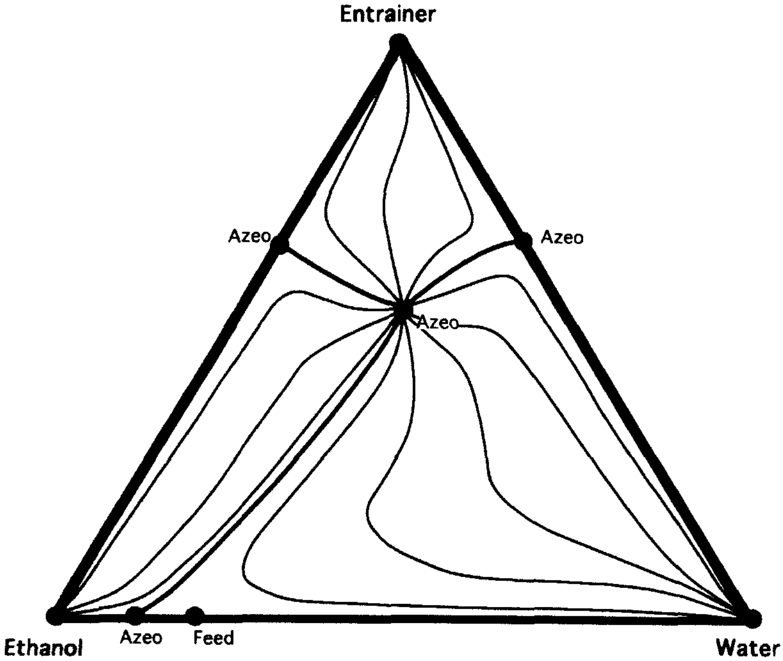

FIG. 24. Ethanol dehydration: Residue curves and distillation boundaries.

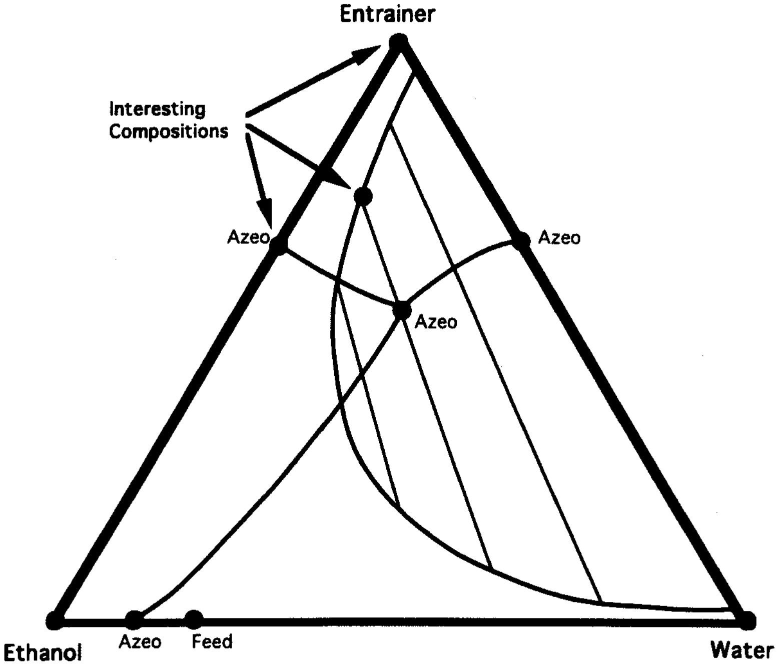

FIG. 25. Ethanol dehydration: Distillation regions and heterogeneous liquid–liquid region.

separations synthesis approach might be followed to concentrate on resolving this strategic problem first, rather than a straightforward opportunistic approach starting from the original feed stream. In other examples, avoiding a tangent pinch might be another type of critical feature that should be addressed first.

Assuming the distillation boundary does not shift significantly with pressure, at least two boundary-crossing methods are possible. The first is exploitation of that portion of the liquid–liquid immiscibility region having tie lines crossing the distillation boundary. The second is mixing with some new stream such that the resulting composition has crossed the boundary. The first method requires that somehow the feed composition be shifted onto one of the desired tie lines. It is not obvious, however, how that might be accomplished.

The second method requires that some composition be chosen so that when mixed with the original feed, the resulting mixture has crossed into the other distillation region. Criteria for selecting such a composition for mixing include the requirement that the composition be easily regenerated, for example, by simple distillation or decantation of a composition that may be generated by distillation. A material recycle is anticipated. From Fig. 25, it is seen that possible candidate compositions for this recycle include the pure entrainer, the ethanol–entrainer azeotrope, and the organic-rich end of the liquid–liquid tie line passing through the ternary azeotrope. Each choice will lead to a different design, which will be affected in part by how difficult the regeneration of the selected composition proves to be.

Let us choose for this example the organic-rich end of the tie line through the ternary azeotrope (as shown later, the ternary azeotrope, being the low-boiling node in all three distillation regions, is readily reached by distillation from any ternary composition in the system). This composition (which we will now call the *azeotropic entrainer*) may easily be mixed with the feed to produce a new composition in the distillation region which includes pure ethanol (Fig. 26). In this region pure ethanol in the highest-boiling composition (high-boiling node), while the ternary azeotrope is the lowest-boiling composition (low-boiling node). In fact, the amount of azeotropic entrainer to be mixed with the feed might just be selected such that the net composition is exactly on the straight line between pure ethanol and the ternary azeotrope.

A distillative separation task with this net feed may then be specified to simultaneously produce a pure ethanol underflow and the ternary azeotrope overhead. This is feasible because the net feed, distillate, and bottoms compositions are collinear by mass balance, and because the distillate and bottoms are on the same residue curve (since all the curves in that region originate at the ternary

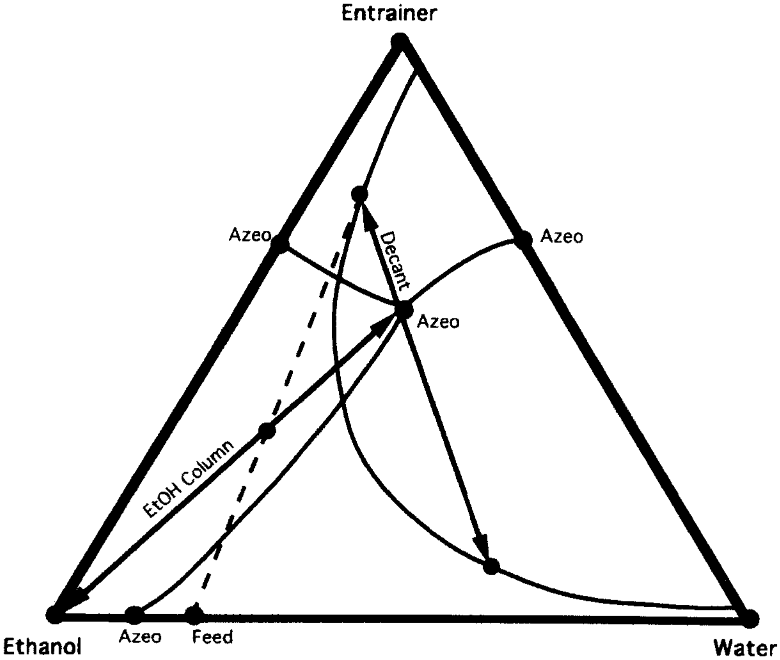

FIG. 26. Ethanol dehydration: Mixing to cross distillation boundary.

azeotrope and terminate at pure ethanol). In the integration of these mixing and separation tasks into actual equipment, the azeotropic entrainer may be directly mixed with the ethanol–water feed and introduced together into the middle of the column as a single feed, which would then be refluxed with condensed distillate in the conventional manner. Alternatively, since the azeotropic entrainer in this case has a lower boiling point than the ethanol–water feed, it may better be introduced separately as a second feed in the upper part of the column. Yet again, depending upon the staging and reflux requirements, the amount of azeotropic entrainer may enough to provide most or even all of the required column reflux at the very top of the column. In this last case, the staging may be adjusted so that the amount of azeotropic entrainer required by mass balance considerations (so that the net feed composition lies directly between pure ethanol and the ternary azeotrope) is also the amount of reflux appropriate for the number of stages to produce the desired distillate and bottoms purities.

Since the overhead ternary azeotrope is heterogeneous, it may be opportunistically decanted, producing the exact composition in the organic layer previously assumed for the azeotropic entrainer as well as an aqueous layer as determined by the liquid–liquid tie line through the ternary azeotrope (Fig. 26). By mass balance it can be shown that the total amount of azeotropic entrainer required to be mixed into the column has not yet been generated by the decant, so the recycle problem is not quite completely solved.

Now the decant aqueous layer must be processed. In any distillation region a number of overhead and underflow compositions are achievable from a given feed through simple single-feed distillative separation depending on the staging, reflux, distillate-to-feed ratio, and constraints of the region boundaries (Fig. 27). The determination of exactly all reachable compositions by distillative separation of a given feed is somewhat complex (Wahnschafft *et al.*, 1992). However, in practice, just two alternatives are generally of most interest: production of the most volatile composition in the region (the low-boiling node) as distillate (with the bottoms determined by mass balance constrained by the region boundaries), or production of the least volatile composition in the region (the high-boiling node) as bottoms (with the distillate determined by mass balance constrained by the region boundaries). Unless the feed, high-boiling node, and low-boiling node are collinear (as they were in the previous column and as they also are for binary systems), these two different choices lead to two different separation flowsheet alternatives.

By far, the most common prejudice (sometimes completely overlooking the other alternative) is to propose the most volatile composition (low-boiling node) as distillate (Fig. 28). In this region containing the decant aqueous layer, the lightest composition is the ternary azeotrope. With enough stages, all of the azeotropic composition in the decant aqueous layer may be recovered, and the underflow will contain only a binary ethanol–water mixture. If the distillate of

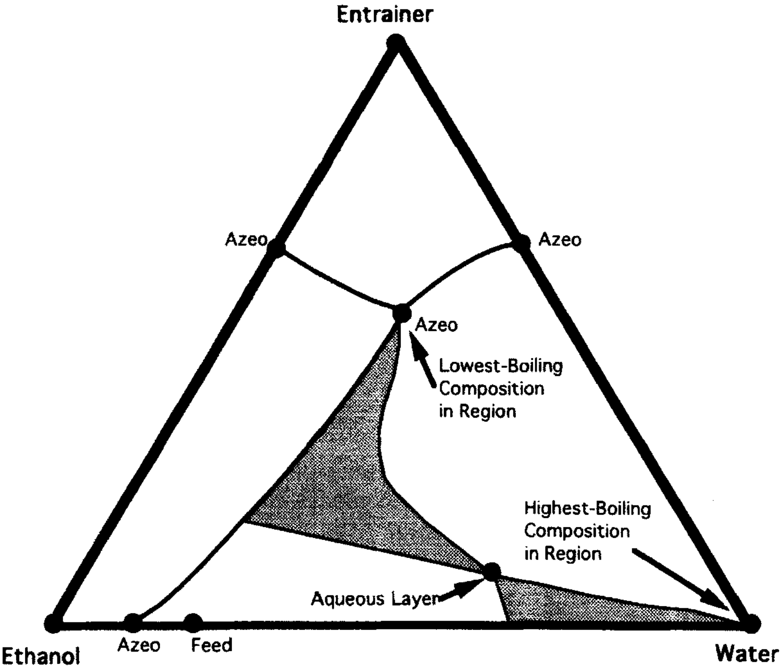

FIG. 27. Ethanol dehydration: Reachable distillative compositions from decant aqueous layer.

this second distillative separations task, whose composition is nearly identical to the distillate from the first column, is recycled to the decant task, mass balance now indicates that the full amount of azeotropic entrainer assumed to be available for the boundary-crossing mixing with the original feed has now been regenerated.

As a final step, the binary ethanol–water underflow from the second column may be opportunistically separated to produce a pure water underflow and a composition close to the original feed to which it also may be recycled. The amount of extractive entrainer must now be readjusted because of the increased amount of feed to the first column, but iteration shows that the flowsheet structure generated remains unchanged and has converged. The tasks specified accomplish the composition goals, producing pure water and pure ethanol. The entrainer remains totally within the system. The subsequent integration into equipment involves three columns and one decanter (Fig. 29).

Returning for a moment to the decant aqueous layer, the other separation alternative is to produce the least volatile composition in the region (water, the high-boiling node) as bottoms. Given enough staging, an amount of water can be removed such that the overhead composition is constrained by the distillation region boundary. This may appear not to be useful, since the resulting overhead

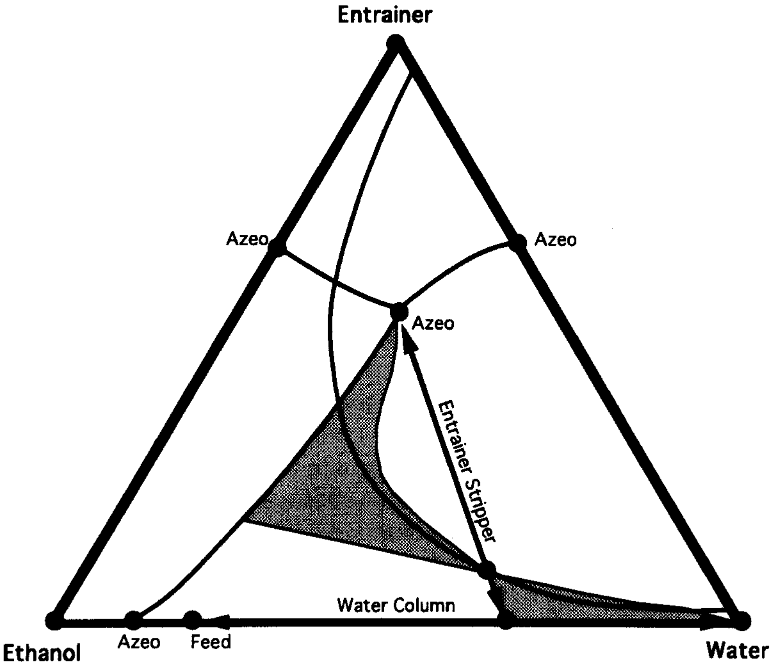

FIG. 28. Ethanol dehydration: Mass balance lines for entrainer stripper and water column.

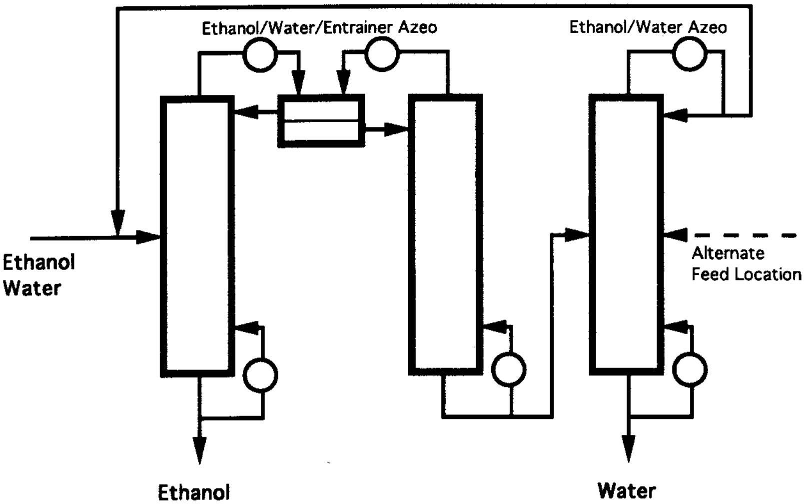

FIG. 29. Ethanol dehydration: Three-column flowsheet.

composition contains entrainer, ethanol, and water and can be separated no further. However, in this case, the composition of this overhead is similar to the mass balance line of the first mixing task, and thus could be recycled to the first column (Fig. 30). This changes the composition and amount of that overall feed from what was originally assumed when that part of the flowsheet was first addressed, but iteration confirms that the structure of the design is not affected. Further, analysis confirms that this alternative flowsheet again produces pure ethanol and pure water and completely recovers all entrainer, but requires only two columns (Fig. 31). Both flowsheets in this example have been implemented industrially, although this second two-column alternative saves about 20% in capital costs (although the energy costs are approximately equivalent). In this example, the original feed illustrated was fairly close to the ethanol–water azeotrope. It turns out that both alternative flowsheets synthesized in this example dehydrate ethanol–water mixtures of any composition, although the feed may have to be introduced more appropriately into a different column than shown.

In the preceding example, the feed composition in the first column was adjusted such that simultaneously the most volatile composition in the region was distillate while the least volatile composition was bottoms, thereby reducing the distillative separation task to a single alternative (with apparently maximum

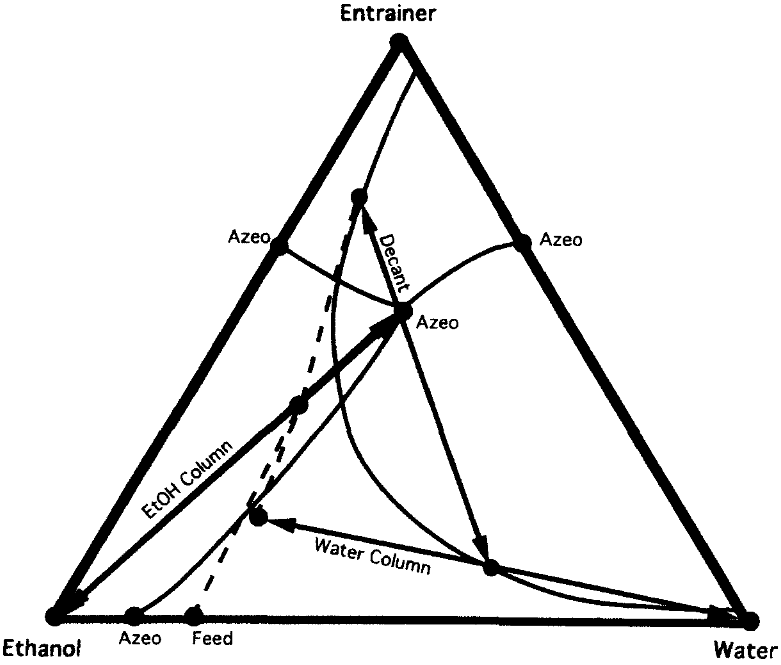

FIG. 30. Ethanol dehydration: Second alternative for decant aqueous layer.

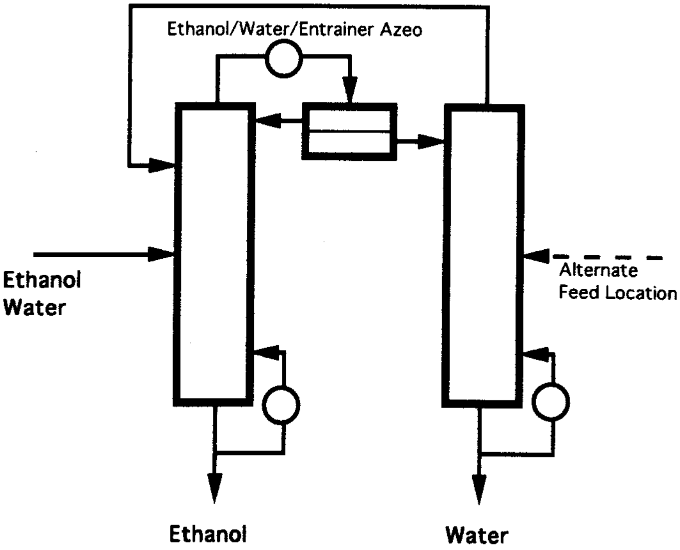

FIG. 31. Ethanol dehydration: Two-column flowsheet.

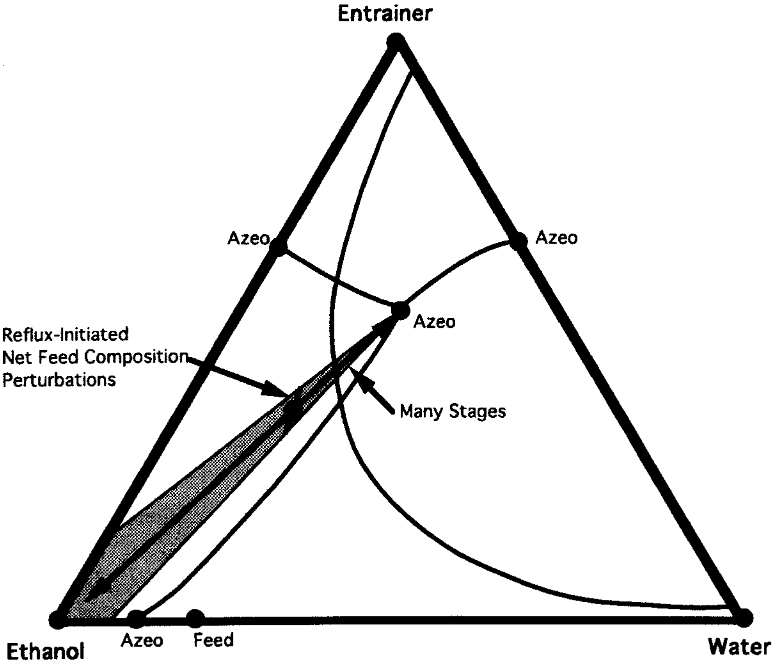

FIG. 32. Ethanol dehydration: Failure to meet ethanol product specification with excessive stages in upper part of column.

separation potential, that is, maximum volatility difference and maximum temperature difference across the column). For multicomponent systems, this is not necessarily as clever as it may appear. Since all residue curves in the region originate at the desired overhead composition and terminate at the desired underflow composition, many designs very near the desired one will be feasible, and slight perturbations in reflux conditions might result in operation associated with very different residue curves. This can be a particular problem in designs that reflux compositions different from the overhead (as in this case), as slight changes in reflux lead to overall feed composition shifts in a direction orthogonal to the column mass balance line (which does not happen with columns that reboil the underflow and reflux the overhead). In this case, slight shifts in the mass balance line can result in major shifts in the actual column composition profile, so much so that the major contaminant in the ethanol product may change from water to entrainer or vice versa. If there are excessive stages in the top of the column, the distillate composition will be pegged to the ternary azeotrope. If there are reflux perturbations that result in overall feed composition shifts, the bottoms composition may swing from the ethanol–water face of the composition triangle to the ethanol–entrainer face, and the ethanol purity specification may never be met (Fig. 32). It may appear that the problem can be resolved by relatively increasing the staging in the bottom of the column, pegging the bottoms composition to pure ethanol and letting the top composition float along the distillation region boundaries converging on the ternary azeotrope. As long as the overhead composition remains in the heterogeneous region, the flowsheet pretty much works as designed (Fig. 33). However, even in such a case, if the mass balance line moves from one side of the ternary azeotrope to another, the column composition profile will also shift dramatically and at the bottom of the column will approach the ethanol vertex from one side of the composition triangle or the other. Although the ethanol purity specification may still be met, the identity of the next plentiful contaminant will change from water to entrainer or vice versa, which may be entirely unacceptable (Fig. 34).

B. SEPARATIONS SYSTEM SYNTHESIS METHOD FOR NONIDEAL MIXTURES

In general, the steps of this separations system synthesis method for nonideal mixtures involving azeotropes include examination of the RCM representation (overlaid with vapor–liquid equlibria (VLE) pinch information, liquid–liquid equlibria (LLE) binodal curves and tie lines, and solid–liquid equlibria (SLE) phase diagrams if appropriate); determination of the critical thermodynamic features to be avoided (e.g., pinched regions), overcome (e.g., necessary distillation

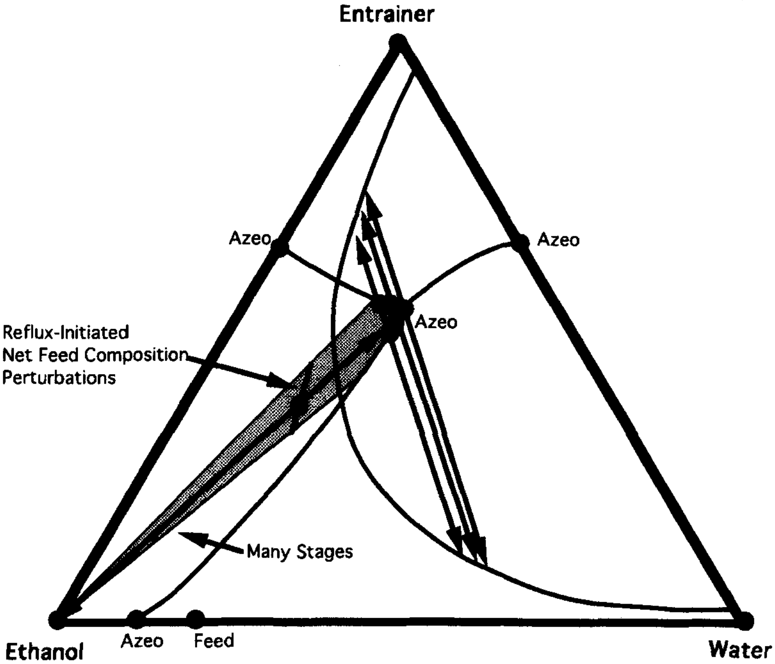

FIG. 33. Ethanol dehydration: Excessive stages in lower part of column to ensure ethanol product purity.

boundary crossings), or exploited (e.g., decant, extraction, or crystallization opportunities); strategically addressing issues raised by the critical features first; pursuing multiple interesting compositions (mass separation agents and other compositions which are useful for mixing but which must be regenerated and recycled); opportunistic resolution of remaining concentration property differences (with the desired products or unanchored critical feature strategic tasks, guided by constraints represented by the RCM); pursuing both possible distillative alternatives when different (heaviest possible underflow in addition to lightest possible overhead); wariness of reflux (reboil) composition different from the distillate (bottoms); avoiding separating multicomponent mixtures simultaneously into lightest possible overhead and heaviest possible underflow; recycling opportunistically to similar compositions already processed; analyzing anticipated and opportunistic recycle feasibility (particularly mass balances); elimination of redundant or duplicate tasks if any; and final integration of tasks into actual equipment. These steps are briefly outlined in Fig. 35.

There are numerous other features of the separations system synthesis method based on the RCM representation that space permits only brief mention. Residue

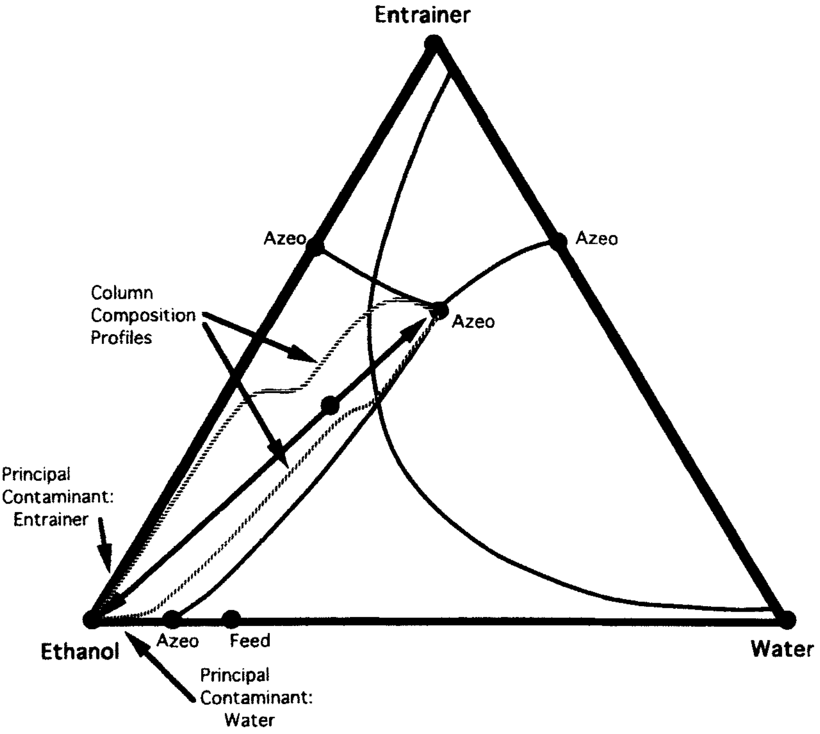

FIG. 34. Ethanol dehydration: Column composition profile shifts with small changes in organic phase reflux.

curves never quite pass through saddle compositions. As a result, the reachable composition analysis for a distillative separation usually does not include saddles; and in the systems where it comes acceptably close, the required distillate-to-feed ratio goes through a minimum at the saddle, resulting in difficult control and potential multiple solutions. Saddle products can also be reached by extractive distillation, a unique form of two-feed distillative separation in which a less volatile solvent is introduced above a more volatile feed with the objective of altering the liquid phase activity coefficients in the upper part of the column ways not otherwise possible in single-feed columns. (Azeotropic columns also typically introduce the entrainer above the feed, but they would also work if the entrainer were introduced with the feed; extractive columns will not work if the solvent is introduced only with the feed.) The behavior of extractive distillation does not follow the usual rules to determine reachable compositions. However, which saddle will come overhead in extractive distillation (generally two are possible) can be determined from the shape of the residue curves in the vicinity of the feed composition. Yet another way to reach a saddle composition is to

- **Define Goal Specifications**
- **Construct RCM Representation**
- **Identify Critical Features**
- **Identify Interesting Compositions**
- **Address Critical Features First**
- **Resolve Remaining Differences Opportunistically**
- **Pursue All Resolution Alternatives**
- **Pursue Both Distillative Separation Alternatives**
- **Avoid Separating Ternary Mixtures Simultaneously into Lightest Overhead and Heaviest Underflow**
- **Beware of Reflux Composition Different from Overhead**
- **Recycle Opportunistically**
- **Check Recycle Mass Balance**
- **Special Consideration for Pressure-Shifting Boundaries, Exploiting Boundary Curvature, Reaching Saddle Compositions, and Solvent-Enhanced Separations**
- **Eliminate Redundant Tasks**
- **Integrate Tasks into Equipment**
- **Evaluate Each Design Alternative**

FIG. 35. Separations system synthesis method.

reduce the number of components in the system. Compositions are identified as nodes and saddles, depending on the boiling points of all of the components and azeotropes in the system (Barnicki and Siirola, 1996). A pure component that is a saddle vertex of a ternary distillation region will be a node in either binary system that does not contain the other third component. Removing a component may turn saddles into nodes in the resulting reduced-order system, which might then be reached by distillation.

Although it is not possible to cross a distillation boundary in a column, the actual restriction is that the entire composition profile be in a single distillation region. Since the column feed need not necessarily (and in general does not) lie on the composition profile (it only needs to be collinear with the distillate and bottoms compositions), it is possible to operate a column in a region with a sufficiently curved boundary, even though the feed composition is in the adjacent region on the concave side of the boundary. Exploitation of distillation-boundary curvature is similar to exploiting pressure shifting of boundaries. Such schemes generally involve a great deal of recycle for mass balance and turn out not to be energetically very practical for boundaries that involve minimum-boiling azeotropes (because the large recycles are all taken overhead), but have been exploited industrially to break maximum-boiling azeotropes (where the recycles are underflows, as in the well-known example of breaking the nitric acid–water azeotrope with sulfuric acid). In other systems in which boundaries terminate in a pure component vertex, curvature can also be exploited not to break an azeotrope, but to produce an acceptably pure product (analogous to pressure-shifting an azeotrope almost to one pure component).

There are, in addition, techniques for avoiding pinched regions, avoiding solidification in distillative separations, and for exploiting crystallization and other SLE and LLE phenomena as well as kinetic-based phenomena. Residue curves with inflections, common near heterogeneous azeotropes, can lead to multiple solutions for some distillation designs, which can result in control problems. RCMs have also been useful for the evaluation of the consistency of solution thermodynamic data and for determining the impact of potential (often hidden) saddle ternary azeotropes. They are also useful in troubleshooting column malfunctions, operational sensitivities, control problems, and even simulation convergence problems. Much RCM analysis and distillation column design technology was developed in the experimental MAYFLOWER system, and is becoming commercially available in such specialized process synthesis software as HYCON and SPLIT.

C. PRODUCTION OF DIETHOXYMETHANE

As one last example of the separations system synthesis method, a process was being developed to produce the chemical diethoxymethane (DEM), the ethanol acetal of formaldehyde. The raw materials are paraformaldehyde and ethanol, and the reaction by-produces water. Diethoxymethane, ethanol, and water form a homogeneous ternary azeotrope that is readily removed from the reaction mass (the paraformaldehyde being relatively nonvolatile). The species allocation is such that DEM is the desired product, water is waste, and ethanol is to be recycled to the reactor. A conceptual separations system to produce pure DEM, pure water, and pure ethanol from the reactor effluent is to be developed.

The problem here is that the ternary azeotrope is homogeneous. Using traditional patterns, one might consider adding yet another component to the system to form some heterogeneous ternary or possibly quaternary azeotrope that can be broken, similar to the breaking of the homogeneous ethanol–water azeotrope of the previous example. Such a system turns out to require four distillation columns. Can an alternative system not requiring the use of an external component be devised?

The DEM–ethanol–water system has three binary minimum-boiling azeotropes and one minimum-boiling ternary azeotrope, and an RCM (Fig. 36) similar to the previous ethanol–water–entrainer example. Distillation boundaries connect the ternary azeotrope with each binary azeotrope. The desired products, DEM, ethanol, and water, are each high-boiling nodes in their respective distillation regions. The system does exhibit a small region of immiscibility (Fig. 37), but both the reaction mass and the ternary azeotrope are in the single-phase region.

The reaction mass is in the same distillation region as water. Since DEM is in a different region, the critical feature distillation boundary between these two

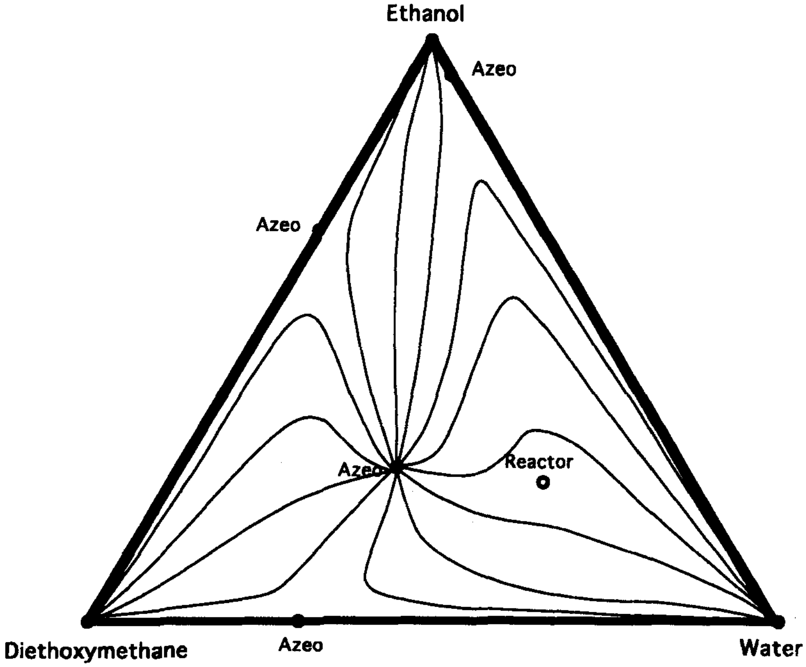

FIG. 36. Diethoxymethane: Residue curve map.

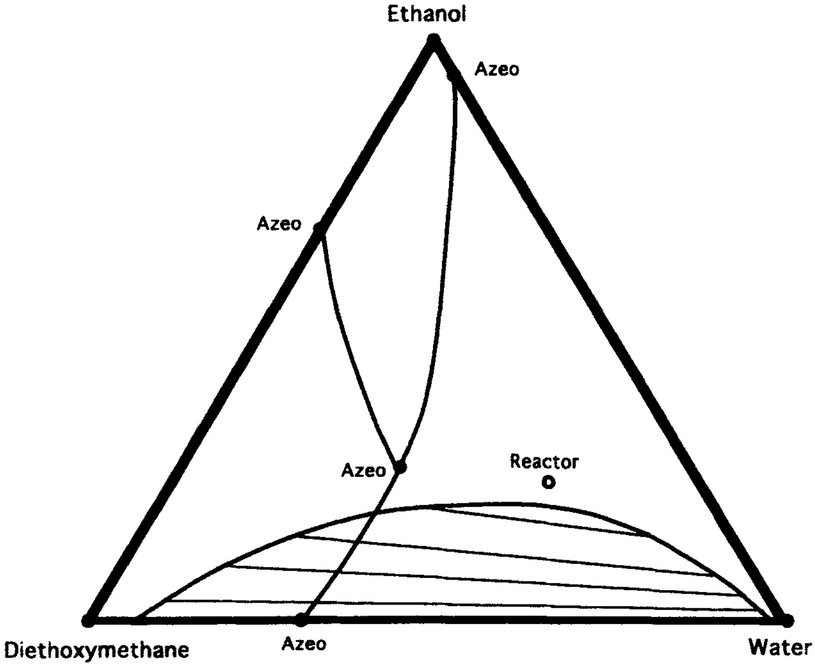

FIG. 37. Diethoxymethane: Distillation regions and heterogeneous liquid–liquid region.

regions must be crossed. Attacking this strategic problem first, it is noticed that some liquid–liquid tie lines span the distillation boundary. If the composition from the reaction can be moved into the lower part of the two-phase region, the distillation boundary may be crossed.

Two possible interesting compositions exist for mixing to a composition in the heterogeneous region: the DEM product and the water product. The DEM choice might be preferable for direct mixing with the reaction mass. The water choice is possible if the composition from the reactor is first somewhat concentrated in DEM (otherwise, two phases will not result). Following the second option (which has the additional benefit of retaining the nonvolatile paraformaldehyde in the reactor), again two options are possible for a first opportunistic distillative separation to concentrate DEM from the reaction mass: distill the lowest-boiling composition (the ternary azeotrope) overhead, producing a DEM-free ethanol–water mixture as an underflow, or distill the highest-boiling composition (water) as underflow, producing a composition constrained by a distillation boundary overhead. The former is chosen for the first design. The ternary azeotropic composition so produced may now be mixed with a portion of the water product. The result is two liquid phases which split into an organic-rich phase layer in the new distillation region, while the water-rich layer has a composition close to the underflow of the first column (Fig. 38).

The organic-rich layer may be opportunistically distilled to produce either the lowest-boiling ternary azeotrope overhead and a DEM–water heterogeneous mixture as underflow, or the highest-boiling DEM product as bottoms with a composition constrained by the distillation boundary close to the ternary azeotrope overhead. The latter is chosen first, producing the desired DEM product and an overhead composition close enough to be opportunistically recycled with the ternary azeo to the decant task.

The water-rich layer and the first column underflow may be combined and distilled. Again, either the lowest-boiling ternary azeotrope may be taken overhead, or the highest-boiling product water taken as underflow. If the latter is chosen, the overhead composition will lie on the distillation boundary close to the ethanol–water azeotrope (Fig. 38).

Two of the three composition goals, pure DEM and pure water, have been met; and by mass balance, sufficient water has been regenerated to serve as the mixing agent for dragging the ternary azeotrope into the two-phase region (actually, the extra water is first put in the system, then continually recycled as is the entrainer in the previous example). A third composition near the ethanol–

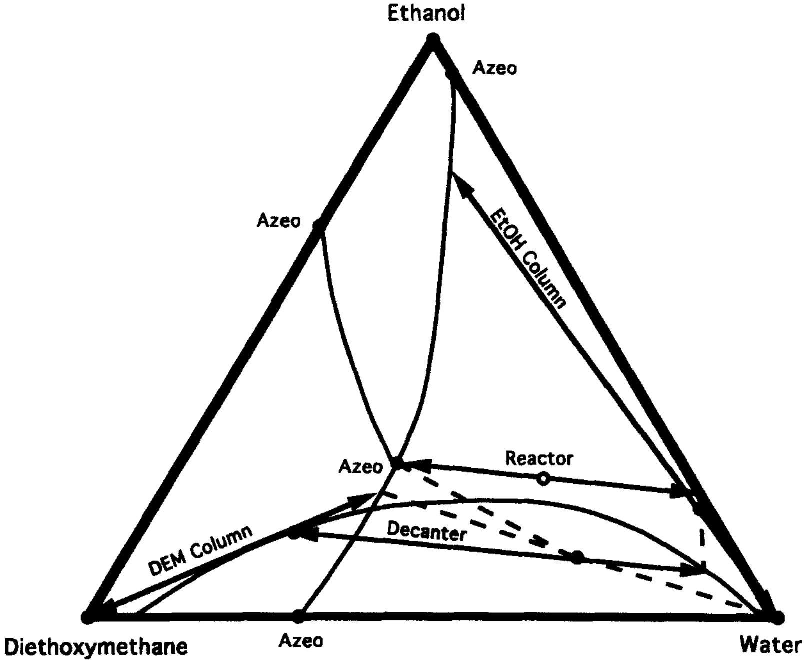

FIG. 38. Diethoxymethane: Mass balance lines plotted on distillation boundary diagram.

water azeotrope has also been generated, but this does not meet the original product specification of pure ethanol.

Getting pure ethanol will be difficult. Ethanol is in the third distillation region, and no obvious means exist to cross the relevant distillation boundaries. A new entrainer might be added, as was the case for the previous example. However, since the ethanol is only going to be recycled to the reactor, the required purity specification might be questioned. What is the effect of small amounts of contaminants DEM and water in the recycled ethanol? The acetal formation reaction is equilibrium-controlled, so there may be some deleterious effect of including products with a reactant. However, additional analysis indicated that the effect is small, and so the ethanol composition requirement was relaxed and the overhead composition recycled to the reactor directly as produced.

Following each of the other design alternatives identified above also leads to feasible solutions, but all involve more separation tasks. The tasks for the present case may be integrated among themselves and with the reaction task in a number of different ways, some of which involve fewer columns but passing more water through the reaction mass.

The conceptual design actually implemented industrially involved three columns and a decanter in addition to the reactor (Fig. 39). The savings over the straightforward scheme to form a heterogeneous azeotrope with an external entrainer was approximately 30% reduced capital and 30% reduced energy. In later refinement of the conceptual design, an extractor was substituted for the decanter, requiring even less water for mixing. The process was patented and the RCM representation was used in the patent application to help explain the process, particularly the method for breaking the homogeneous ternary azeotrope without the addition of an external component (Martin and Raynolds, 1988).

VII. Challenges for the Future

Many of the methods and tools being developed in the conceptual process engineering research community are still more concerned with analysis (targeting, representation, operability, etc.) than with synthesis of alternatives. More effort needs to be directed specifically to alternative generation for all conceptual design stages at all levels of detail.

Most of the applications discussed here were based on algorithmic systematic generation approaches to process synthesis. There also exist a number of heuristic rule-based or expert-systems–based approaches including PIP for total flowsheet synthesis, SPLIT for separation schemes (Wahnschafft *et al.*, 1991),

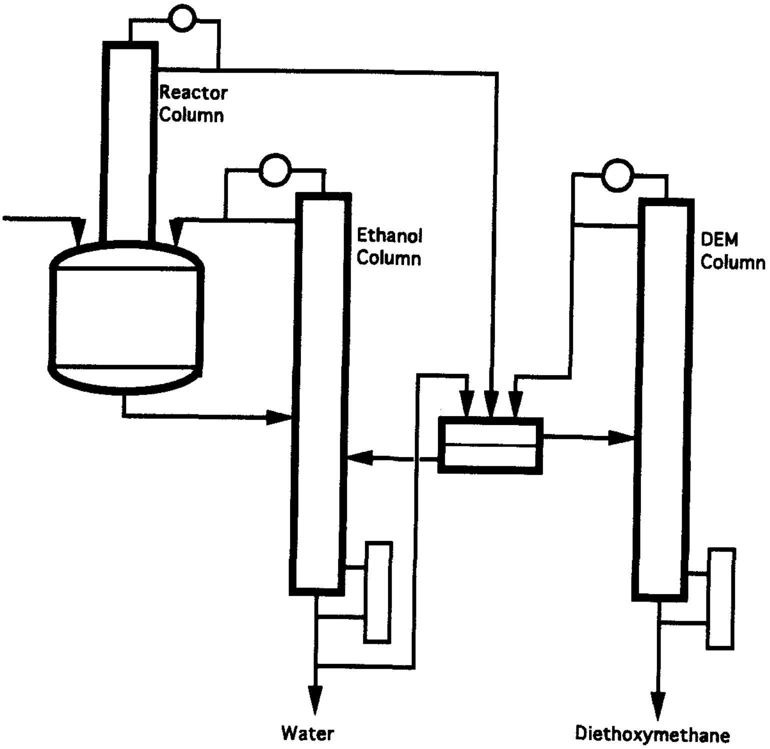

FIG. 39. Diethoxymethane: Process flowsheet.

and most of the informal methods used by conceptual design engineers. Because of combinatorial difficulties, neither the algorithmic nor the heuristic methods carry any guarantees of optimality.

The superstructure optimization approach, on the other hand, does hold out the promise of structural optimality. Given a superstructure, however, there are significant challenges remaining to improve computational efficiency to overcome combinatorial difficulties and to develop global optimization strategies for the generally nonconvex mixed-integer nonlinear problems involved. Steady progress continues to be made on these fronts (Grossmann and Daichendt, 1994). There still remains, however, the general problem of generating appropriate superstructures in the first place.

One possibility might be to employ algorithmic and/or heuristic systematic generation methods not for the synthesis of a single alternative, but for the synthesis of an entire superstructure of alternatives, for example, by following and including all alternatives encountered at all decision points during the procedure. The resulting superstructure could still depend upon the order in which various differences were attacked. Therefore, this superstructure might be subjected to evolutionary modification methods, possibly to trim away demonstrably inferior segments, but more likely to add patterns to one part of the superstructure similar to that generated for other parts. Finally, the superstructure so generated and refined could be in principle optimized to produce structurally and parametrically superior solutions.

However, no matter how completely the objective function for the optimization is formulated, there will always be additional important design criteria not included in the optimization. Examples include safety, controllability, flexibility, constructibility, maintainability, and many, many others. On these other criteria the several best solutions selected by the optimization procedure might be critiqued, possibly subjectively, before a final design is selected. The process synthesis paradigm of the future might very well then combine all of the existing approaches (systematic generation, evolutionary modification, and structural optimization), exploiting all of the existing methods (algorithmic, heuristic, and mathematical programming). The result would be a new process synthesis strategy that might be called *generate–evolve–optimize–critique* (Fig. 40).

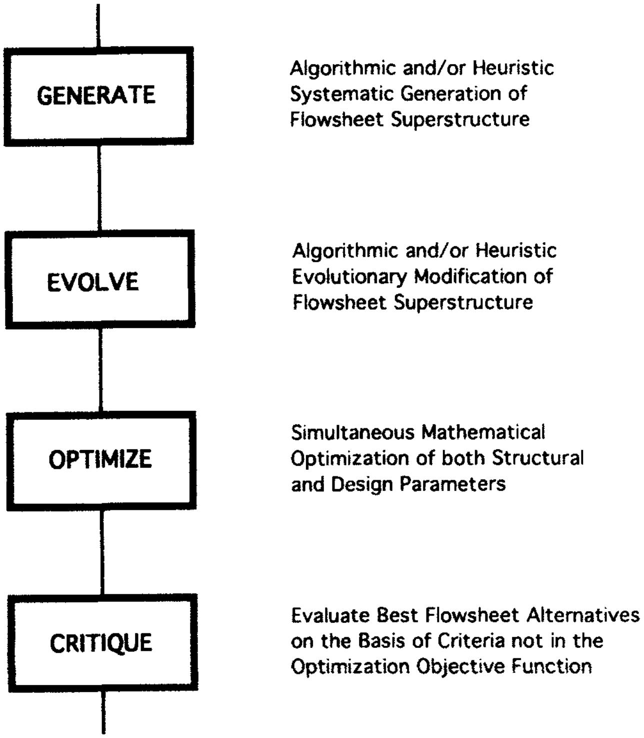

FIG. 40. Future process synthesis paradigm.

VIII. Conclusions

Systematic approaches, both to the material-making innovation process and to the process synthesis steps of the conceptual engineering stages within that innovation process, have begun to have measurable industrial impact. With these techniques—some automated and some not—an increased number of higher value, lower energy, lower environmental impact, and sometimes even novel design alternatives have been synthesized and actually implemented. Energy reductions of 50% and net present cost reductions of 35% are typically achievable.

Certain features of systematic approaches to process synthesis appear to have special merit. These include architectures that are hierarchical in scope and iterative in level of detail, targeting as a useful bounding exercise, means–ends analysis as a recursive problem-solving paradigm with hierarchical strategic as well as opportunistic goals implemented by an iterative formulation–synthesis–analysis–evaluation design strategy, thinking specifically in terms of tasks to be accomplished before equipment to be employed, thinking of nonconventional properties to be exploited to accomplish identified tasks, recognizing the importance of representations to encapsulate analysis within synthesis activities, and exploring the advantages of solving related synthesis problems iteratively or even simultaneously.

Advances in problem formulation and in computational hardware and software capability offer the promise of a new generation of practical process synthesis techniques based directly on structural optimization. Soon the goal of synthesizing provably unbeatable conceptual process flowsheets may be at hand.

Acknowledgments

I would like to thank my colleagues S. D. Barnicki, R. D. Colberg, D. L. Martin, E. J. Peterson, and M. R. Shelton for their help with the preparation of this manuscript. I would also like to thank Eastman Chemical Company for allowing these experiences to be shared.

References

Agreda, V. H., and Partin, L. R. U.S. Patent 4,435,595 (to Eastman Kodak Company) (1984).

Barnicki, S. D., and Siirola, J. J. ''Enhanced Distillation,'' *in* ''Perry's Chemical Engineering Handbook,'' 7th ed. McGraw-Hill, New York, 1996.

Blakely, D. M. "Cost Savings in Binary Distillation through Two-Column Designs," M. S. Dissertation, Clemson University, Clemson, SC (1984).

Doherty, M. F., and Knapp, J. P. "Azeotropic and Extractive Distillation," *Kirk-Othmer Encycl. Chem. Technol. 4th Ed.* **8,** 358 (1993).

Douglas, J. M. "Conceptual Design of Chemical Processes." McGraw-Hill, New York, 1988.

Foucher, E. R., Doherty, M. F., and Malone, M. F. "Automatic Screening of Entrainers for Homogeneous Azeotropic Distillation," *Ind. Eng. Chem. Res.* **30,** 760 (1991).

Grossmann, I. E., and Daichendt, M. M. "New Trends in Optimization-Based Approaches to Process Synthesis," *Comput. Chem. Eng.* (1996).

Kirkwood, R. L., Locke, M. H., and Douglas, J. M. "A Prototype Expert System for Synthesizing Chemical Process Flowsheets," *Comput. Chem. Eng.* **12,** 329 (1988).

Linnhoff, B., and Ahmad, S. "Towards Total Process Synthesis," Paper 26d, AIChE Meeting, Washington, DC (1983).

Linnhoff, B., and Hindmarsh, E. "The Pinch Design Method of Heat Exchanger Networks," *Chem. Eng. Sci.* **38,** 745 (1983).

Lodal, P. N. "Production Economics," *in* "Acetic Acid and its Derivatives" (V. H. Agreda and J. R. Zoeller, eds.), p. 61. Dekker, New York, 1993.

Mahalec, V., and Motard, R. L. "Procedures for the Initial Design of Chemical Processing Systems," *Comput. Chem. Eng.* **1,** 57 (1977).

Martin, D. L., and Raynolds, P. W. U.S. Patent 4,740,273 (to Eastman Kodak Company) (1988).

Rudd, D. F., Powers, G. J., and Siirola, J. J. "Process Synthesis." Prentice-Hall, Englewood Cliffs, NJ, 1973.

Seader, J. D., and Westerberg, A. W. "A Combined Heuristic and Evolutionary Strategy for Synthesis of Simple Separation Sequences," *AIChE J.* **23,** 951 (1977).

Siirola, J. J. "The Computer-Aided Synthesis of Chemical Process Designs," Ph.D. Dissertation, University of Wisconsin, Madison (1970).

Siirola, J. J. "Status of Heat Exchange Network Synthesis," Paper 42a, AIChE Meeting, Tulsa, OK (1974).

Siirola, J. J. "Energy Integration in Separation Processes," *in* "Foundations of Computer-Aided Process Design" (R. S. H. Mah and W. D. Seider, eds.), vol. 2, p. 573. Engineering Foundation, New York, 1981.

Siirola, J. J. "An Industrial Perspective on Process Synthesis," *in* "Foundations of Computer-Aided Process Design" (M. F. Doherty and L. T. Biegler, eds.), p. 222. Cache Corporation, Austin, TX, 1995.

Siirola, J. J., and Rudd, D. F. "Computer-Aided Synthesis of Chemical Process Designs," *Ind. Eng. Chem. Fundam.* **10,** 353 (1971).

Simon, H. A. "Science of the Artificial." MIT Press, Cambridge, MA, 1969.

Wahnschafft, O. M., Jurian, T. P., and Westerberg, A. W. "SPLIT: A Separation Process Designer," *Comput. Chem. Eng.* **15,** 565 (1991).

Wahnschafft, O. M., Koehler, J. W., Blass, E., and Westerberg, A. W. "The Product Regions of Single-Feed Azeotropic Distillation Columns," *Ind. Eng. Chem. Res.* **31,** 2345 (1992).

Witzeman, J. S., and Agreda, V. H. "Alcohol Acetates," *in* "Acetic Acid and its Derivatives" (V. H. Agreda and J. R. Zoeller, eds.), p. 257. Dekker, New York, 1993.