

Chemical Product Design

E.L. Cussler and G.D. Moggridge

University of Minnesota (USA) and Cambridge (UK)

Abstract

Chemical product design includes deciding what to make, as well as how to make it. Because chemical products play an increasingly prominent role in the chemical industry, their design should receive increasing prominence. We propose a four step procedure as a template around which to organise ideas for designing chemical products.

I Introduction

Chemical product includes choosing what product we will make. It precedes chemical process design, which deals with how we plan to make the product chosen. In the past, most of those involved in the chemical enterprise emphasised process design, because the enterprise was focussed on perhaps fifty commodity products. For these commodities, price was the key, and efficient production was the route to low prices. This strategy correctly dominated the efforts of the chemical industry in the twentieth century.

Now the goals of the chemical enterprise have become much broader. These goals include not only the fifty or so commodity chemicals, but several thousand high value added chemicals. Many of these are pharmaceuticals. The broader goals of the chemical industry also include chemical mixtures, often with specific microstructures, like detergents and coatings. The goals include polymers processed to have characteristics desired in filters or films. These broader goals mean that chemical professionals now participate in a wider variety of business decisions. These include deciding what to make.

In this article, we suggest a four step template by which we can decide which chemical products we want to make. Because of the enormous variety of products which are possible, we should not expect this scheme to work perfectly in every case, but it can serve as a mental checklist for organising our thinking. The four steps are as follows:

1. Needs We must decide what needs our product will fill.

2. Ideas We must generate ideas which could satisfy these needs.
3. Selection We must efficiently select the best ideas.
4. Manufacture We must design the process for making our product.

The first three steps of this template are unique to product design, but the fourth step includes the more familiar topics of process design. We will give details of each of these four steps in the body of this article.

II. Needs

In defining chemical needs, we must remember that the product is not for ourselves, but for our “customers”. We must begin by asking our customers what they require. We use the term customers in a loose sense here: we do not necessarily mean those who will buy our product, but rather those who will benefit from its chemistry. These may be companies or government agencies, rather than individuals.

The consensus among marketing organisations is that the best way to get this sort of information is by face to face interviews with customers. Fewer than ten such interviews risks missing significant information, while more than fifty gives duplication. One group of customers merits particular attention, the “lead users”. These are the people most expert in the product, who will benefit most by its improvement. Frequently, these lead users will have good suggestions for improvements.

The needs recorded from interviews will be a hodgepotch of conflicting and incomplete statements, of varying relevance and practicality. Our task is to organize these needs as groups and to edit them into a cogent list. We expect to drop some stated needs, either because they appear impractical or beyond our company’s expertise or mission.

The needs should then be ranked, for example as essential, desirable and useful. The essential needs are those without which the product cannot succeed. We will aim to achieve as many desirable needs as possible, particularly if our competitors’ products fail to do so. We are unlikely to design explicitly for the useful needs, although we understand that we gain if our product can fulfil these, too. The way in which needs are organized will depend on the product being considered. This will usually be easy if we

aim to improve an existing product. The more innovative the proposed product, the harder it is to satisfactorily define the needs.

Next, we must convert our qualitative list of needs into specifications, including as much quantitative and chemical detail as possible. When this occurs, the role of marketing declines and that of engineering rises. We can develop specifications via three steps:

1. Write complete chemical reactions for any chemistry involved.
2. Make mass and energy balances important to the product's use.
3. Estimate the rates of any important changes which occur during the product's use.

These steps will lead to a set of ideal product specifications. We should examine these carefully, because they are often entirely impractical. If this is the case, we must revise our specifications to be more realistic. These revisions may lead us to abandon the project altogether. If the only way of meeting customer requirements is to break a law of thermodynamics, we should stop product development now. This type of rapid, critical examination is sometimes known as a "chicken test", after a Canadian method for testing airplane engines for their capacity to fly through flocks of geese: the Canadians threw frozen chickens into the engines.

The final step in the Needs stage is to specify a benchmark. This is an existing or idealized product against which we can measure our new design. If we cannot beat the benchmark, the product is not worth developing. In some cases, when an innovative new product is being contemplated, no obvious benchmark will be available. We may be able to envisage an idealised benchmark by analogy with similar products

Example: Water Purification for a Family The World Health Organization (WHO) estimates that 1.7 billion people do not have access to a clean water supply. Lack of pure water is the world's biggest cause of child death. In the absence of major civic projects, the point of demand is the family. Imagine that a large non-governmental organization is interested in producing a water purification unit suitable for use by individual families in third world countries. Interviews with governmental and non-governmental agencies working in such countries reveal the following list of needs:

Essential

Supply a family with safe drinking water
Inexpensive

Desirable

Rural focus
Simple to operate

Useful

Environmentally sustainable

Assess these needs, and suggest a benchmark.

Solution: The first question is how much water we need to purify. One person requires five litres of drinking water per day and typically uses about 50 litres for cooking, cleaning and washing. A household might have ten persons, so a minimum reasonable requirement would be 100 litres per day. A five year lifetime will be desirable, so the total through put will be around 200,000 litres.

Because waterborne diseases are the greatest threat, we must remove protozoa, bacteria and viruses. Toxic materials are sometimes a problem, but it is probably not useful to focus on these – the variety is great, the problem is usually local and a better solution is often preventing discharge. While the WHO guidelines for microbial removal are useful to have in mind, to see these high standards as an absolute requirement might be a mistake. A less effective device will be more useful than one which no one can afford.

Cost will be critical. We are aiming at some of the poorest people in the world. Interviews in Nepal suggest that a purchase price of \$15, with an annual operating cost of \$5, is the most that the rural population can afford. This is a stringent requirement, one we may not be able to meet. Moreover, the rural focus means that the device must operate in the absence of a power supply. The product must be easy to explain and to use by those with little experience of technology. The need of “Environmental sustainability” is vague, so we ignore it for now.

Finally, we choose chlorination as a benchmark. This has the advantage of being simple, cheap and well established, but clearly has problems in terms of supply, use and discharge of chemicals. Still, our device must be more attractive than local chlorination if it is to succeed.

III Ideas

Once we have chosen specifications for our target product, we need some good product ideas. In principle, we only need one idea, the one which we will manufacture. In practice, product development requires up to one hundred ideas in order to find one truly worth pursuing. To get these product ideas, we will depend on people more than publications. The most important people are those on the team responsible for developing the specific product. We will normally assemble this team for free-ranging discussions which aim at generating possible answers. How to run such “brainstorming” discussions is carefully described in the literature, and so is not detailed here. We mention only that these discussions should initially be non-critical, and that all participants should be treated as equals. Discussions of new chemical products are sometimes curtailed for reasons as trivial as that the boss’ spouse disagrees with some of the ideas suggested.

In addition to depending on the product team, we should pay special attention to customers who already are using existing, similar products. Other human sources – consultants, private inventors, and the like – are often less useful. Literature has widely ranging value. Patents and trade information from competitors is often more useful than archival literature. Other methods for idea generation use schemes for developing chemical synthesis. In most cases, the key is the product development team.

We now have our hundred or so ideas of widely varying quality. We must somehow sort through these ideas to locate the best five or so for further developments. Evaluating all of the items thoroughly will normally take more resources than we will ever have, and will take much more time than we have. We must find fast ways to screen ideas.

We suggest proceeding with two stages. First, we should try to sort the ideas on completely qualitative grounds, reducing the number to perhaps twenty. Later, with a bare minimum of quantification, we should try to screen the surviving ideas, aiming to get just the five or so we think are best. To reduce the number of ideas to twenty, we begin with a list of all the ideas. We first can remove redundancy. Often this redundancy will occur because some ideas are special cases of others. For example, in a discussion of better barriers for landfills, one idea could be:

“The barrier should capture mercury”

A second idea might be:

“The barrier should adsorb all heavy metals except calcium”

The first idea is just a special case of the second. In addition to removing redundancy, we want to drop ideas which are obvious folly. In doing so, we should be cautious, because some silly ideas may contain dreams. Sometimes, we can benefit from keeping a separate list of these flawed dreams, just to serve as a stimulus to later development. Normally, the efforts to remove redundancy and folly will still leave us with around seventy ideas.

To reduce the number of ideas further, we should try to organise them into categories. How this can be done depends on the ideas generated: there seems to be no general strategy. Once such an outline is made, it may expose gaps, which may imply repeating the brainstorming. More often, we will find that large groups of ideas will be inconsistent with our organization's objectives or its strengths. These groups of ideas can be dropped, a major simplification. These steps commonly cut the number of ideas to meet our target of twenty. One note of caution, however: many of our best ideas will cluster under a single heading. Because we don't want to overspecialize too soon, we should consider choosing at least one idea beyond this cluster for the next stage of product development.

We must now reduce our twenty surviving ideas down to five or fewer. We will not have the resources to make detailed calculations for the twenty survivors, so we need approximate but quantitative tools which let us continue the screening, but on a more rational basis.

One common method for this screening is to choose five or fewer key attributes shared by most of the surviving ideas. These attributes will include factors like scientific maturity, ease of engineering, risk of failure, and cost. We should choose factors which are different for different products. For example, even if safety is the most important product attribute, we gain nothing by choosing safety if all our potential products are equally safe.

Once these key attributes are chosen, we need to assign weighting factors to each. We will normalise these weighting factors, assuming values which sum to one. This

implies that all the products which we are still considering are capable of satisfying all attributes, at least to a limited extent. If there is one attribute which is truly essential, we should drop all ideas which can't satisfy this constraint, and continue our evaluation of the survivors.

On the basis of these attributes and their weighting factors, we now score all of our ideas relative to a convenient benchmark. We may find it convenient to assign the benchmark scores of five, and then to choose scores from each new product between one (poor) and ten (excellent). We then calculate an average weighted score for each product. The potential products with the highest scores are those which we choose for further development.

Example: A Pollution Preventing Ink A printing company prints personal checks with lithographic ink containing the carcinogenic solvent methylene chloride (CH_2Cl_2). Workers at this company also clean the presses by wetting a shop rag with the same solvent, and scrubbing down the press. This works well. The trouble is that much of the methylene chloride evaporates and so risks workers' health and censure from the environmental authorities. Also, the soiled rags have recently been reclassified as a hazardous waste.

The company clearly needs to use a different ink, one which has less negative environmental impact. Some ideas for this ink are shown in Table I. Sort these ideas to identify those most worth pursuing.

Solution The ideas easily break into four groups, as shown in Table II. The numbers in parentheses refer to the designation of ideas in Table I. The first group in Table II involves changes to the printing presses. Because the company does not want to make the enormous capital investment involved in changing the presses, this group is deferred.

The second group of ideas involves either containing the solvent or using a different solvent. These ideas are the easiest to implement, and hence the most tempting for further development. The third group of ideas implies the invention of a new ink, a more major effort than the substitution of a new solvent.

Table II. Ideas for a New Lithographic Ink with Reduced Solvent

1. Don't use a solvent.
2. Switch solvents.
3. Clean the press with robots.
4. Change the press.
5. Use an electrostatic ink.
6. Use a laser printer instead of the current design.
7. Change ink chemistry.
8. Recycle all of the solvent.
9. Clean the press with a high-pressure spray.
10. Extract the solvent from the rags used to clean the press.
11. Do the whole process in a clean room.
12. Isolate all equipment.
13. Clean the press less often.
14. Clean the press in a fume hood.
15. Print more checks at a time.
16. Mix the current solvent methylene chloride with other solvents.
17. Have workers wear a self-contained breathing apparatus.
18. Use a solvent mixture.
19. Use a solvent which dissolves the ink.
20. The solvent in the ink should differ from the cleaning solvent.
21. Use a non-volatile solvent.
22. Use partial cleaning of specific components of the press.
23. Steam clean the press.
24. Clean the press with air.
25. Put the press in a car wash.
26. Clean the press by brushing.
27. Clean the press by burning.
28. Make the lithography more like a jet printer.
29. Don't use checks.
30. Use a disposable press.
31. Use oil to trap the solvent.
32. Make checks by photocopying

Table II Sorted Ideas for a Pollution Preventing Ink The numbers in parentheses refer to the ideas suggested in Table I.

- I. Improve Current Printing
 - A. Change Press (4)
 1. Isolate Press (3, 11, 12, 14, 17)
 2. Use Laser Printer (6)
 3. Use Photocopying (32)
 - B. Change Cleaning
 1. Less Often (13, 15)
 2. Other Solvents (9, 23, 24, 25, 27)
- II. Use a New Solvent
 - A. Change CH₂Cl₂ Operation
 1. Recycle (10)
 2. Burn (27)
 3. Freeze (new)
 - B. Replacement of CH₂Cl₂ (2, 20)
 1. Non-volatile solvent (21)
 2. Oil as Solvent (31)
 3. Solvent Mixtures (16, 18)
- III. Solvent-Free Ink Chemistry (1, 7)
 - A. Electrostatic Ink (5)
 - B. "Solvent which Dissolves Ink" (19)
- IV. Don't Use Checks (29)

The final idea, "Don't Use Checks", may initially seem foolish; but remember the increasing number of electronic money transfers. The company may decide that electronic data processing which replaces handwritten checks is like the automobile which replaced the horse-drawn buggy. If so, then printing checks may be like making buggy whips. Thus this fourth idea should be carefully considered in further development.

IV Selection

From the good ideas, we must choose the best one to take forward for production. Because all the remaining ideas are promising ones, this decision will involve considerably more effort than we put into cutting the number of ideas from one hundred down to around five in the screening just described. The decision will involve making estimates based on chemistry and engineering. At the same time, we still wish to develop our product as quickly as possible and do not want to put resources into exploring products we will end up rejecting.

Our first step in selecting between the remaining ideas is to estimate how each will perform relative to the specifications we have set. In order to do this, we must gather more information about each idea. This will involve firming up exactly how each idea will work; we may need to do some simple experiments to achieve this; we will certainly need to carefully explore the literature. As we generate more detailed information on each idea, the idea itself will change and become clearer. Thus there may be an iteration between the Ideas and Selection stages.

Having made an assessment of each idea against each specification, we must now make an overall comparison between the ideas. In some cases, particularly where the identified needs are primarily technical, this will be easy once good estimates of performance have been made. In other cases, subjective judgements will be necessary. In such cases it is often useful to proceed by drawing up a decision matrix of attributes, as described for screening ideas. However, we previously were interested in making quick decisions to eliminate the weak ideas. Now we are considering strong ideas and want to make a careful decision. While the methodology of the decision matrix remains the same, we must now put a lot more effort into evaluating each idea.

Example: A Device Which Allows Wines to Breathe David Anderson, an alumnus of the University of Minnesota, is the owner for the best wine store in Minneapolis-St. Paul. He says that wines – especially red wines – improve if allowed to “breathe” by exposing them to air before drinking. Oxygen in the air reacts to reduce the “hard” and “soft” tannins which naturally occur. Such exposure enhances flavour, allowing both fruit taste

and aroma to be more clearly perceived. The amounts of oxygen required vary widely. Uncorking the bottle for 15 minutes before drinking is useless, but exposing wines to excessive oxygen will turn them into vinegar. One good way is to pour the wine from the bottle into a decanter, leaving any residue behind. Several pours between decanters is better. Mr. Anderson even has one friend who pours the wine into a larger glass, covers the glass with his hand, and gives the wine a good shake.

Estimate the aeration needed in a product which can let wine breathe in only a few minutes.

Solution: Aeration can make a startling difference, especially for cheap, freshly opened wine. We tried four methods of aeration:

1. Open the bottle for 15 minutes.
2. Decant into an open pitcher and let the wine sit two hours.
3. Decant the wine fast three times, entraining air.
4. Shake for ten seconds in a large glass.

Then we tested the wine.

The first method is useless, as David Anderson suggested. The other three give roughly similar taste improvements. Another method, putting wine in a blender, aerates it excessively, dramatically reducing flavour.

These results are consistent with estimates based on mass transfer, as shown in Table III. Our selected design – whatever it is – should give around one transfer unit. Any engineering solution meeting this specification will give about the same performance. However, we suspect that the success of a “wine breather” will depend strongly on the aesthetics of the aerator’s design. It is our own confidence in our lack of ability in this aesthetic area which inhibits us in developing this product.

I. Manufacture

By this stage, we have decided which chemical product we want to manufacture. We have identified a customer and that customer’s product needs. We have generated ideas to fill that need, and we have selected the best idea. We are ready to decide how we will make the product.

Table III Four Methods for Letting Red Wine Breathe

Method	<i>Mass Transfer</i>	<i>Wine</i>	<i>Time</i> (sec)	<i>NTU</i>
	<i>Coefficient</i> (cm/sec)	<i>Area/Volume</i> (cm ² /cm ³)		
Uncork bottle for fifteen minutes	10 ⁻³ (a)	3/750	900	0.004
Decant into pitcher two hours before serving	10 ⁻³ (a)	100/750	7200	1
Decant three times, entraining 0(5 cm ³) air	20 × 10 ⁻³ (b)	0.4 (b)	120	0.5
Shake 30 cm ³ wine twenty times in a 300 cm ² glass	7 × 10 ⁻³ (c)	300/30	10 (c)	0.7

(a) Estimated from the free convection caused by ethanol evaporation.

(b) Estimated from observed 5 cm³ of 0.1 cm bubbles entrained in each decantation.

(c) Estimated from penetration theory with a penetration time of one shake (0.5 sec).

The very wide range of chemical products possible means that we will need to consider a wide range of manufacturing methods. To provide some initial guidance for manufacture, we find it useful to think of four types of product. The first type is some sort of device, especially for a medical application. Examples include the artificial kidney and the osmotic pump for drug delivery. The manufacture of these devices normally depends on mechanical engineering more than on chemistry and chemical engineering. We will not discuss mechanical devices because this topic is thoroughly described elsewhere.

The other types of chemical products can be roughly classified as shown in Table IV. The first type, commodity chemicals, made in amounts exceeding 10⁷ kg/year,

normally have molecular weights less than 100 daltons. Their manufacture often uses gaseous reagents, supported catalysts, and purification via distillation. The process engineering involved has cost as its primary focus. This type of chemical product was the focus of chemical industry growth in the last century.

The two other types of chemical products shown in Table IV are much harder to describe. The more obvious type, the ten thousand or so high value added chemicals, are exemplified by drugs. These compounds have molecular weights typically in the range of 500 – 700 daltons. Their chemical structure is normally well defined, though often one of many available isomers. The finished products are crystalline, of exceptionally high purity. Once such a high value added product is identified, it is normally manufactured as quickly as possible, because the first manufacturer will usually command two thirds of the market. The need for purity and speed, plus the small amounts to be manufactured, dictate the use of dilute solutions and generic equipment. Manufacture is more like gourmet cooking than like petroleum refining.

Table IV Different chemical products are manufactured differently

	<u>Commodity Chemicals</u>	<u>High Value Chemicals</u>	<u>Microstructured Products</u>
<u>Key Factor</u>	Cost of Product	Speed of Manufacture	Function, from microstructure
<u>Examples</u>	Ethylene, Ammonia	Penicillin, Viagra	Paint, Detergent
<u>Amount made</u>	$> 10^7$ kg/year	$< 10^4$ kg/year	$< 10^6$ kg/year
<u>Typical Molecular Weight</u>	100	600	> 100 ; often $> 10^4$
<u>Phase during Synthesis</u>	Gas	Dilute Solution	Concentrated Solution or Melt
<u>Chemical Reactor</u>	Dedicated, Continuous, Plugflow	Generic, batch, stirred tank	Ranges widely
<u>Common Separation</u>	Distillation, Absorption	Extraction, Adsorption	Ranges widely

The third type of chemical product, which is the most heterogeneous, commonly has a useful microstructure. This large group of products includes many “speciality chemicals”. Manufacturing ranges widely, and product properties are often a function of product history. For example, polymer properties are affected by molecular weight

distribution, and paint quality changes if the paint is frozen. The process engineering required is normally a compromise between the first two types.

At this point, we should have sufficient information to convince ourselves that we had selected the best idea and that it has a high probability of success. Inevitably, there will be gaps in our knowledge of the details of exactly how the product will work. Before going further, we must fill in this missing information. This missing information may be obtained by literature searches, by consulting experts and by conducting experiments. It may be tedious, time consuming and expensive, which is why we put it off until we decide to proceed. Because of the diversity of the information which is missing, we cannot effectively generalise about how to find it.

Finally we specify the process by which we will make the product. This is where chemical product design and traditional process design finally merge. Typically, there are differences of emphasis in the process design for a high value-added chemical product as compared to a commodity chemical. For commodity chemical manufacture, minimizing cost and maximizing efficiency the goal. Economies of scale, continuous processing and good heat integration are normally required. For high value-added chemical products, we are usually producing much smaller quantities. The emphasis is speed, not optimization. Batch reactions in non-dedicated equipment are the norm. Because these products usually involve delicate large molecules, separations tend to focus on extraction, adsorption and crystallisation rather than distillation.

VI Conclusion

Process design is a well established subject, with effective heuristics to aid practitioners and teachers. As diverse chemical products increase in importance, we need to place the design of these products on a similar footing. As we have outlined, there are significant differences in the design for specialty products for commodity products. For this reason, a simple mapping of conventional process design onto new challenges in product design is unlikely to be effective. In this article we have outlined a possible four step scheme for tackling chemical product design. We hope that this scheme will provide a basis for further development.

Acknowledgements

This article is based on our book Chemical Product Design (Cambridge University Press, 2001) and on our entry of the same title in the online edition of the Kirk Othmer Encyclopaedia of Chemical Technology. These in turn drew on K.T. Ulrich and S.D. Eppinger Product Design and Development (3rd edition McGraw-Hill, 2004). The work was partially supported by the US Air Force (F49620-01-1-0333), the National Science Foundation (CTS 0322882) and the Petroleum Research Fund (39083-AC9).