Part I Strategy and Development

Bioseparation and Bioprocessing. Edited by G. Subramanian

ISBN: 978-3-527-31585-7

1
Process Development – When to Start, Where to Stop

Glenwyn D. Kemp

1.1
Introduction – What is Process Development?

Process development provides the vital link between R & D and manufacturing. It takes a drug discovered in the laboratory and shown to be clinically promising, and allows it to be produced in sufficient quantities to be used on the patient population. However, this is not the total extent of process development. The purification method developed also has to meet the critical quality parameters of purity and activity within a very proscriptive regulatory framework. Moreover, it has to do this reliably and reproducibly for batch after batch, and at such a cost that the drug is commercially viable.

Hence, the process development scientist has to balance the often conflicting demands of low cost and high quality within a highly constrained timeframe. This aim of this chapter is to discuss ways of balancing these conflicting demands, to look at areas where emphasis should be placed and to discuss areas where investment of time may be less productive than hoped.

A fundamental challenge for process development is that it inevitably has to occur within a constrained timeframe. This is an important paradigm and is in contrast to classical scientific training in which the desired outcome is to determine a definitive answer to a problem, however much time and effort is required. In the case of process development, the timeframe is dictated by the needs of clinical trial phases and economic management. In this context, the desired outcome is not necessarily a definitive answer, but rather a workable solution. This approach can be both counterintuitive and uncomfortable to scientists new to the field of process development. Fortunately, experience eventually accustoms scientist to that environment where best compromises are often the only achievable outcome.

Bioseparation and Bioprocessing. Edited by G. Subramanian

ISBN: 978-3-527-31585-7

1.2 The Challenges of Process Development

The single fundamental question with which to challenge any process destined for the manufacturing floor is not whether it is the best process possible, but rather to ask "Is it good enough?", i.e. is it fit for purpose. As with all good science a reliable answer can only come from fully understanding the original question. Therefore, it is worth spending a little time to consider in more detail what makes a process "fit for purpose". In order to qualify for this accolade the process should be pass the following criteria:

1.2.1 Is the Purity High Enough?

The key term here is "enough". There is no economic benefit in producing a higher purity if this is not required. However, the typical purity would be expected to be routinely somewhat higher than the minimum specification, to allow for process robustness as we shall see below. In general, a higher purity is obtained by adding more purification or product-handling steps. However, the greater the number of these steps, the lower the recovery of the product will be. It is essential to consider the overall recovery of product for all the process steps rather than to look at individual process steps: 90% recovery may seem acceptable, if this is applied to three chromatography purification steps the overall recovery becomes a reasonable 73%. Consider, however, that each centrifugation, filtration, chromatography or even hold step may potentially cause a loss and a more realistic number of operations for the process may be eight steps, giving an overall recovery of only 43%. If the average recovery drops to 85%, the overall recovery after eight steps will only be 27%. A few extra percentage lost in the odd step or two may seem trivial, but for a billion dollar blockbuster drug every 1% loss of drug will represent a potential $10 million in lost revenue! The addition of extra unwarranted purification steps will have serious economic implications. There will also be implications for the process robustness by having effectively included unnecessary steps.

The question of what is a "high enough" purity is itself fraught with difficulties. With an experimental drug it instinctively feels safer to aim for the highest purity possibly, with absolute purity being the ideal result. As discussed above, this is usually the most expensive option. Drugs produced in research laboratories are often the result of an intensive labor of love, hand crafted by highly skilled scientists with little care for cost, time or difficulty. The preparation of the drug is in itself the scientific challenge. Such a handmade bespoke drug is ideal for initial evaluation, but not necessarily an appropriate model for a mass production. There is a danger that, in setting purity criteria using the initial data from the development laboratory, this will be used as the lower baseline for purity. This may result in an impossibly high standard to be set for the final drug. It should also be borne in mind that it is relatively easy to convince the regulatory authorities to accept

a revision of purity specification for a drug if the new specification is tighter, e.g. higher purity, but it is almost impossible to argue with the regulatory authorities for an easing of specifications, especially without recourse to effectively repeating a phase III study.

The outcome of these considerations is that the drug should have a minimum purity against which it is validated, and that this purity should be based on pragmatic consideration of the potency and efficacy of the drug, and the nature and level of likely contaminants. It may also be necessary to identify and separate impurities, if only to demonstrate that they are clinically neutral. An example of this would be isomeric forms of a monoclonal antibody. In some cases these may be clinically inactive or neutral, in which case the isoforms can be copurified with the drug itself. If, however, the isoforms are shown to have an adverse clinical activity, perhaps by competitive inhibition of the active drug, then these isoforms will need to be removed. This can be a difficult separation to perform and will inevitably add to the cost, either as a result of requiring more time, additional steps or lost recovery of the active drug.

It may even be worth considering returning to the original drug screen to find an alternative antibody without the isoforms if the problem is picked up early enough in the development cycle.

1.2.2
Is the Process Robust?

As we noted above, the purification method used in the research laboratory may typically be the work of a single highly skilled and capable individual. It may require precise peak cutting, high-efficiency columns, temperamental buffer systems and be performed in a cold room. All of these factors may well be detrimental to the robustness of the process. Put simply, the robustness of a process is how far from the ideal set of conditions the process can operate and still yield product of the required purity.

It should be the aim of the process development scientist to make a process as "bomb proof" as possible. Certainly routine variations in the process parameters should not present a challenge to the process.

Process robustness can be thought of as path along a cliff top. If the path is set too far back from the cliff, the benefit of the breathtaking views will be lost; if, on the other hand, the path is set to close to the cliff, one careless footstep may lead to a catastrophic loss of enjoyment! The ideal is to place the track close enough to the cliff edge to reap the benefits of the view, but far enough away from the edge to be safe. Hence, the ideal process will be operating near to the limits of the process for maximum efficiency, e.g. sample load mass, flow rate or holding time, without undue risk of loss due perhaps to low recovery, purity or activity. The control of process robustness is by a combination of step selection, i.e. the actual methods and media used for each specific step and the parameters defined in the operation of those steps – the process specifications.

1.2.3 Are the Process Specifications Valid?

Unfortunately, it is often the case that the process development scientists are closer to the research laboratory then the production floor in their outlook. The result is usually overly tight specification which cannot be reliable repeated at full scale. For example, if the resolution required for an ion-exchange chromatography purification is very precise, the step will be inherently more sensitive to variations in, for example, column bed height, buffer flow rate, buffer pH, salt gradient mixing. While 1–2 L of buffer in a flask can easily be checked and adjusted for conductivity, pH and even temperature, this is a more challenging task for 10 000 L of buffer. If a buffer falls outside the specified combination of pH and conductivity it will have to be reworked or discarded – both expensive options.

Thus, it is imperative to know, for example, how precisely a buffer can be prepared; what are the acceptable ranges of pH and conductivity. This can then be used as a basis for setting a range value to the specification. Indeed if a buffer cannot be used within the typical range generated on the production floor it would be better to investigate alternative buffers at the process development stage than to try to reengineer the manufacturing equipment or procedures.

Another common problem with process specifications is that of the "rangeless parameter". When a process is being defined all conditions must be given not to a specific value, but to a value within a range, e.g. pH 7.4 ± 0.2. The act of deciding what this range is will also produce support data for the robustness of the process. Not only should ranges be provided for all critical parameters specified, but it is helpful to have data to indicate a range for parameters not specified as critical, to support the assumption that they are not critical. It may also be borne in mind that ranges do not have to always be plus or minus a median value, they can also be given as "no more than" or "no less than" values. In this case it is still a good idea to have an upper or lower limit, respectively, if only for practical purposes.

In general, it is easier to reduce the range of the specification after multiple full-scale cycles than to increase the range. Therefore, it is good practice to specify ranges as large as possible from the outset. However, operation at the extremes of these ranges must still result in a final product of acceptable quality. Ideally range values would be derived from a series of experiments in which the parameters are increased or decreased successively until the process fails. This should then be repeated for combinations of parameters at the determined limits. In reality there is rarely enough time or material to allow such extensive studies. A degree of pragmatism has to be applied. The critical parameters are usually selected and become the ones to be investigated in depth. The critical parameters are usually selected from past experience or published data. The danger here is that there may be "hidden" parameters which are specific to the product or process. For example, temperature may affect the binding capacity and resolution of a hydrophobic-interaction chromatography column, but would not be expected to change either for a Protein A-affinity step. Hence, it is always a worthwhile exercise to take a step back from each process, consider all of the potential variable parameters present, and actively decide which are likely to be critical to the process

and which are not. This exercise also allows a third group of parameters to be identified – the "maybe critical" variables. In this case a few experiments at extreme values will usually indicate if the variable is critical or not.

In determining the process specifications the interaction of individual parameters must also be considered. Often the interactions of parameters are obvious and can be predicted with reasonable certainty, e.g. the effect of bed height and flow rate on resolution of a gel-filtration column. Sometimes, however, it is not obvious which parameters even interact, let alone the nature of the interaction. The areas of statistical design of experimentation (DOE) approaches are often cited for identifying interactions between different parameters. When well designed and executed these studies can indeed drastically reduce the amount of experimentation required to identify parameter interactions and to cooptimize these parameters within the process. However, DOE should be approached with an element of caution – it is essential that the assumptions made at the outset of the experimental design are understood and are sound in order to interpret the data generated correctly. Nonlinear relationships, in particular, are liable to distort the data and lead to false conclusions. Training should be undertaken before embarking on either the design or interpretation of "DOE" studies as there is great potential to make unfounded "leap of faith" conclusions.

1.2.4 Is the Process Scalable?

A fundamental question to ask at the outset of any process development exercise is "what is the anticipated final scale for this process?". This is usually a much easier question to ask than to answer. The earlier the development phase, the less accurate the answer will be; however, it is still a valid question even at the research phase. The answer will depend on a number of factors.

(i) *The potential patient population*. A drug to treat a rare congenital condition will clearly have a much lower target population of perhaps a few thousand. This can be contrasted with a drug for, say, male pattern baldness or chronic obesity, which would have potential target population of tens or even hundreds of millions. A prophylactic treatment such as a vaccine may potentially have a patient population of billions. The size of the target population will be mostly guesswork, but it should be reasonable to estimate to an order of magnitude at least.

(ii) *The required dosage – single dose or chronic treatment*.T he required dosage can usually only be estimated based on the required dosage for already established or tested drugs. Within this framework it is clear that an antibody will typically require up to several grams per patient per treatment, whereas a regulatory molecule such as a hormone or interferon will probably require several orders of magnitude lower dose. Similarly, a single dosage form

such as a vaccine will require a much lower amount of product per patient compared with a drug required for a lifelong chronic illness such rheumatoid arthritis or multiple sclerosis. Hence, the size the dose, the frequency of dosage and the duration of treatment will all be critical in determining the final requirements for the drug.

(iii) *Other potential applications*. In the past, drugs have been registered initially to treat unmet needs in small patient populations or for rare conditions in order to qualify for orphan drug status. Receiving orphan drug status gives the drug developer the benefits of protected rights to sell the drug for 7 years, tax breaks, subsidies and expedited US Food and Drug Administration (FDA) review, allowing the drug to reach the market faster and via smaller-scale, and thus lower-cost, clinical trials. Once the drug has been approved a number of much larger scale applications can then be licensed for the same drug. In extreme cases this can result in a drug that was ostensibly only going to be produced at a small or moderate scale becoming a major blockbuster. Examples include Amgen's Epogen® (erythropoietin), which was initially licensed for anemia due to kidney dialysis, and Genentech's human growth hormone. In both these only small doses are required per patient; for some other orphan drugs such as Genentech's Rituxan® this is not the case, with a current production scale at the tonne level. If the patient population expands greatly beyond the originally anticipated clinical need the assumptions made for the final process scale may be seriously compromised.

1.2.4.1 **Considerations for Scale of Operation**

Once a rough estimate of the final scale has been decided, the final scale of operation can be deduced, based on previous or published experience of expression levels and purification yields. Clearly, this will have to match the anticipated final need. In the initial period after the launch of Amgen's Enbrel® the greatest challenge was to produce enough drug to meet the high market demand. The frustrations of getting a drug licensed are nothing compared to the frustration of not being able to produce enough of a successful drug to meet demand! The final scale of production can help decide what purification steps would be feasible, e.g. centrifugation is a common method for the clarification of cell cultures in the research laboratory. In this case the centrifugation is in discreet batches. This is achievable because the volumes being handled are relatively low. Once the process is scaled up beyond a few tens of liters, batch centrifuges become less efficient. Unfortunately this is also still a range where disk-stack centrifuges are either not available or represent an overly complex and inefficient solution to the problem. At this scale, filtration is an

excellent alternative. Multilayer filters are now available starting with coarse filter pads to remove whole cells followed by subsequently finer layers designed to remove cell wall fragments and colloids by progressively finer filters. These combined membranes are even available in large-scale self-contained units for ease and handling and operator safety. At pilot scale and up to a few hundred liters these represent a cost-effective solution for initial clarification. However, for very large-scale processes handling 1000 L or more filtration pads now become unwieldy and expensive, and centrifugation once again becomes the method of choice, albeit using disk-stack centrifuges. Although filtration membranes are still required for postcentrifugation polishing, the combined centrifuge and filter train will still outperform a filtration train alone. It can be seen from this example that the optimum solution may change depending on scale and that extensive development work performed at the intermediate scale may be wasted or, worse still, lock the process into an inefficient solution. By careful forethought and consideration of the final potential scale, the process can be tailored, from the very beginning of process development, to provide the most efficient manufacturing process.

1.2.5 Is the Process Economically Viable?

It is clear that the process scientist has a challenge on their hands with each process they are given. Namely, to select the correct set of operations and the operating conditions to allow the purification of the target molecule to a sufficiently high purity, and to do so within the most proscriptive regulatory framework of perhaps any industry. So it is perhaps not too surprising that the additional challenge of cost analysis and cost reduction is not usually a high priority. It goes without saying that the yield should be as high as possible as this will make the process more economic. Fewer steps will also usually reduce costs too. However, this is often as far as a process scientist will get in cost modeling to evaluate the economic viability of a process.

There has been a tradition within the biotech industry that the products are so specialized and so exquisitely effective that they can justify the high price tag that usually accompanies a "biotech" drug. This, in turn, allows for a high cost of manufacture – after all, the market will pay the premium. This comfortable approach to cost of goods is changing. These changes are being driven both internally and externally. Internally there is by a desire to maximize profits. This becomes especially prominent as the scale of manufacture increases. Until relatively few years ago most biotech companies would have found it difficult to give an accurate breakdown for their manufacturing costs. Although the "big ticket" items would be known, such as the cost of affinity media in a production-scale column, the cost of a liter of buffer or the hourly overhead rate would be a mystery. Now many biotech companies are operating as manufacturing companies as much as research companies. This has forced them into a new mindset in which their products have become commodities. It has also resulted in senior mangers recruited from the classical pharmaceutical industry and more used to incremental

cost savings through process improvements than the "eureka" moments which previously dominated biotech.

Externally, cost reductions are being driven by a public and political pressure to reduce drug costs, and by the imminent acceptance of generic biotech drugs or "biosimilars". These pressures are leading to the "commoditization" of biotech drugs. An example of the effects of this can be seen in "old biotech" in the production of antibiotics. Many antibiotics or their precursors are produced by fermentation, and there is fierce competition within the industry to reduce costs and increase yields, driven by market forces. The barriers to entry for antibiotic production are relatively low. The molecules are well defined, and there is a large amount of generic data on production, safety and efficacy. As a result these biotech companies have to differentiate themselves, and gain a competitive edge by marginal increases in yield through strain selection/development, fermentation optimization and purification process improvements. It is not unreasonable to expect similar market forces to prevail in the future for some biotech drugs

It is within this new management paradigm that process development scientist now find themselves working. Considerations of manufacturing costs and next-generation process evolution are rapidly becoming as integral to process development as the selection of purification conditions.

Specific areas where cost reductions can be gained have been discussed already. To recap – reduce the number of purification steps, avoid overly high purity and activity specifications, and select operational steps and conditions suitable for use at the anticipated manufacturing scale. All these approaches will all have a positive impact in final cost of goods.

1.3 Strategies to Develop a Downstream Process

As discussed above, the requirement for any successful process is that it will produce enough of the target drug, at the required purity and activity in a cost-effective way that will be accepted by the regulatory authorities. In the following sections some strategies for achieving this requirement will be discussed. First, the actual methodology of process development will be considered. There are several ways to develop a process – each has it merits and its drawbacks.

1.3.1 The Bigger Test Tube Approach

This basically entails taking what has been developed already in the research laboratory and simply reproducing it at progressively larger scales. The advantage of this method is that it is very fast, the process is effectively already complete and all that needs to be sourced are larger vessels/equipment of the same nature. Also, the data previously generated will be applicable for validation of the process. Although this may seem too naive an approach, there are cases where it can have merit. If the aim is to very rapidly generate samples of drug for early clinical or preclinical trials this is probably the most cost-effective way of completing such

proof-of-principle studies. It can allow several potential drugs to be produced in parallel for final selection. If the final amount of drug required is likely to be small, e.g. cytokines, interferons or hormones, or if the process is likely to be relatively simple, as seen with some vaccines, then the larger test tube approach may suffice even as far as the final manufacturing scale. For the first years of production of Epogen Amgen continued to use the roller bottle reactors used in the original process. Scale-up was simply by adding another rack and more roller bottles. This is another advantage of this approach – scale-up is often simply a matter of replicating the production process.

The bigger test tube approach does have some significant drawbacks, however. These are more obvious the larger the disparity grows between the research and the final manufacturing scales. As noted above, some methods, such as batch centrifugation, simply do not scale-up without becoming unwieldy and inefficient. Some techniques such as preparative electrophoreses have never successfully made the transition from laboratory to pilot scale, let alone manufacturing. As equipment is scaled up, a range of challenges will occur in terms of mass transfer, homogeneity and thermal transfer. For some technologies such as fermenter design and operation, these parameters are well studied and guidelines published on how to scale-up effectively. For some technologies, especially newer technologies, such as disposable "fermenter in a bag" designs, scale-up is less clearly defined. Another caveat from this approach is in the use of esoteric or nonscalable chromatography media. An example would be an affinity media using a specific monoclonal antibody as the affinity ligand. This technique can be applied in the research laboratory with great success. However, for a production-scale process the affinity antibody would have to be produced in bulk and to almost the same degree of purity as the therapeutic target molecule itself. The cost of such a chromatography media would almost certainly outweigh the potential benefits.

Another problem can occur if the chromatography media used in the research laboratory is either not available in bulk or is simply not scaleable. Many prepacked chromatography columns sold for laboratory use contain media which is different to that provided in bulk for packing large-scale columns. The most common differences are that the analytical-scale media are less cross-linked or may have smaller particle size. The former occurs with compressible media such as agarose and results in media which are softer. These can be used successfully in small columns, where the column wall support helps protect the media from compression. On transfer to a large size column where the wall support is no longer significant, typically around 100 mm diameter or larger, the media may no longer be mechanically strong enough to support the required flow rate. This results in bed compression which, in turn, results in a higher backpressure. This then causes further media compression and a further increase in backpressure, thus creating a positive feedback loop which can quickly raise the backpressure and damage either the media or the column. The only solutions are to reduce the flow rate, which will significantly increase the process time and therefore reduce productivity, or change to a more cross-linked version of the media. This can cause problems in later scale-up stages due to the subtle changes in the background interactions between the media and feedstock components. For example, when agarose media are more highly

cross-linked they also become slightly more hydrophobic in nature, which may lead to differences in the impurity profiles of the eluted proteins.

The use of increasingly smaller particle size media is common with rigid media such as silicas where the smaller particles sizes afford greater resolution. However, bulk packing of small particle size media is more difficult in preparative-scale columns and the general solution is to use larger particle sizes. This also has advantages in both lower cost and lower operational backpressure. However, it is critical that the separation be developed on media with the same particle size to ensure that the resolution is still sufficient to achieve the desired separation. Problems can arise in obtaining suitable analytical-scale columns – usually these will only be available as custom columns from the manufacturer of the bulk media.

1.3.2 The Template Process

Most established process development laboratories use this approach as a starting point. The advantages of using a standardized template process are a reduction in process development time, familiarity with the techniques being optimized and confidence that the final process will be validatable. However, the template approach can potentially contain hidden traps. It should be stated that there is, in reality, no such thing as a true template process. That it is a process which can be simply repeated for any and all target molecules of a certain type. This was the original aim and desire of companies working with monoclonal antibodies. It was felt that if a process could be developed once, that process would then be applicable to all other monoclonal antibodies, thus generating huge time and cost savings in bringing monoclonal antibodies to the market. Unfortunately, the reality has proven more intractable. By its very nature each monoclonal antibody clone is unique and different to other monoclonal antibody clones. Even an apparently minor change in the antibody such as single amino acid variations or changed glycosylation patterns can generate an antibody that will behave differently during the purification process. These changes can be as subtle as a small shoulder on an elution peak or as dramatic as precipitation during elution. Due to this variation the use of a template process must be seen as a short cut to part way down the process development line, rather than a means of bypassing it altogether.

1.3.3 Process Development by Gradual Evolution

The shortcomings of the pure template approach usually lead to a more pragmatic approach in which a process is developed as a process of gradual evolution from a standard starting point, i.e. the template, but with the acceptance that the final process will have some specific character of its own. These changes can be driven both by the characteristics of the specific target molecule and also by changes

in the regulatory environment which may require further or more stringent processing than was previously required. An example of the latter is the increased emphasis on virus inactivation and removal since the problems experience by the blood product industry with human immunodeficiency virus (HIV) contamination. It should also be noted that the FDA and other regulatory bodies are now more open to scale-up and postapproval changes (SUPAC). This is a distinct change from earlier approaches validating biotechnology products.

Despite that fact that each process will have its individuals twists and quirks, it is reasonable to say that the majority of therapeutic monoclonal antibodies currently being produced use an almost identical process. Starting with Protein A-affinity chromatography as a capture step, followed by cation exchange. The main variation comes in the final step. Mostly this is an anion-exchange step using either chromatography columns or charged membranes, but processes using hydrophobic-interaction or hydroxyapatite chromatography are also in place.

This observation leads on to the next method of process development.

1.3.4
The "Me Too" Process

This should be differentiated from the pure template process approach in that a "me too" process would be developed using a template from outside the current process development teams experience. Although apparently a simple approach, it is usually very difficult to obtain sufficient details to allow a complete process to be reproduced.

It is one of the greatest challenges for the process development scientist to keep up to date with the latest developments in the field. This is especially problematical for smaller or start-up companies. The key reason for this is that little or no information tends to be published on the development of current processes. From the perspective of the biotechnology companies, this may be a sound commercial decision. Certainly in the more commoditized pharmaceutical industry incremental increases in yield or reductions in overhead costs, through more efficient processes, will directly impact the economic viability of a process and therefore be regarded as a commercially sensitive. As a result, most information published on processes tend to be for "previous-generation" obsolete processes or processes for drugs which have failed in clinical trials. This information is also usually only available in presentations at conferences, rather than in scientific journals. A crossover area can be found in the commercial trade journals, where solicited articles discuss current issues. However, these are also interspersed with articles which are designed to promote new technologies or products which may not yet have widespread use or acceptance. For these reasons a "me too" process can be very difficult to achieve unless through acquiring expertise through the natural movement of staff between companies.

One advantage of reproducing a current process is that it should already be familiar to the regulatory authorities. However, this should not be seen as an opportunity to shortcut process validation. It is still the case that each process

must be validated for its own performance. As noted above, individual variations in the nature of different clones of monoclonal antibodies will still produce enough variation in response to process steps to necessitate further optimization even for a "me too" process. In this respect it would be similar to the problems of trying to implement a template process.

1.3.5 The Clean Sheet

Perhaps surprisingly, the least-encountered method for process development is the clean sheet approach, where the characteristics of the target molecule are considered and a rational theoretical purification train is postulated. This approach is the most scientifically "pure" and perhaps closest to the approach seen in a research laboratory. The key difference for a process development laboratory is that this approach allows the use new or "esoteric" large-scale techniques which may not be valid at the small scale. The downsides to this approach are clear. Novelty is not usually encouraged in process development. Being the first company to present a new purification technique to the regulatory bodies is fraught with risk. Potentially, the approval could be delayed or even not granted if there is doubt over the reliability and reproducibility of the purification process. The only justification for such a risk would be if all standard methods of purification have failed.

Adopting a conservative and risk averse approach, while making business sense, does have serious implications for research into new and novel methods or purifying proteins. It discourages process research within the biotechnology industry. Few companies are in a position to provide resources for research on novel processes which in all probability will not be implemented. On the other hand, the companies which provide the purification tools, such as chromatography media and filters, will be disinclined to develop truly novel technologies that will either not be adopted or, in the case of lower-cost methods, will harm their current sales. For entrepreneurs and inventors the situation is even less promising, as the potential adoption of a new technique is likely to require significantly longer than their funding will allow. Probably the last significant "new" technique in large-scale protein purification over the last 20 years was expanded-bed chromatography. Although widely used in the chemical industry, this method was only developed for use in protein purification in the late 1980s and was first introduced commercially around 1990s. It was almost 10 years later in 1999 when the first FDA approved process using expanded-bed chromatography was reported by GE Amersham. This was for a relatively small-scale application, requiring only around 20 L of media, for the production of a lipoprotein component of a human vaccine by SmithKline Beecham. Another supplier of expanded bed media, Upfront Chromatography, also report, on their website, the use of the technique at multitonne scale for the purification of lactoferrin and IgG from cheese whey. Notwithstanding this application, the technique remains rare in large-scale protein purification processes. It can be seen from this that the potential payback period for a successful new process is around 15 years. As a result of these commercial pressures there

is significant lack of research into new and novel methods of large-scale protein purification. This is especially the case for low-cost methods or adaptation of current technologies from other fields where the opportunities for commercial exploitation are low. Although there are some very good academic institutions working on new purification methods, the costs involved in producing suitable feedstock and operating at large scale present a real problem. Without closer cooperation and significant funding from the biotechnology industry, the academic study of large-scale protein purification is unlikely to yield any new techniques in the foreseeable future.

1.4 Process Optimization

Having discussed general approaches to process development it is appropriate, at this point, to consider some specific examples of how process steps can be optimized. These techniques and approaches are relevant to all the previously discussed approaches to process development.

1.4.1 Cell Removal/Clarification

Over recent years there has been a progressive trend towards higher titers from fermenters, and this has resulted in higher cell densities and lower cell viabilities than previously experienced. This trend is often seen in process development. In the early discovery and preclinical stages of a developing a biological drug the fermentation will not usually have been optimized. The aim at this stage is simply to produce enough of the potential drug to allow evaluation of its effects with *in vitro* and animal models. The problems tend to arise when the drug is further progressed, especially as it heads towards phase II clinical development. Typically, the upstream cell culture groups and downstream purification groups are in different laboratories, different departments, different buildings and, in some cases, different countries! The target of the upstream group is to produce the highest possible titer of product in the shortest possible time. The aim of the downstream group is to take the feedstock from the upstream group and develop a suitable purification process. The problem arises because these two groups are working in parallel simultaneously. If the upstream group are doing their job, the feedstock is likely to progressively have higher cell densities and lower cell viabilities. This will also be accompanied by changes in the impurity profiles. For example, lower cell viabilities can result in increased levels of DNA and proteases in the feedstock as the nonviable cells degrade. Physically, the feedstock may also change as a result, becoming more viscous. If a process is developed, using the original feedstock, which is not sufficiently robust it may run into serious problems when challenged with the "improved" feedstock which may manifest as increased backpressures, slower unit operation times, reduced ability to clear impurities or increased product degradation.

New feedstock components can also be a source of problems. These may have different binding characteristics to those for which the process was originally developed resulting in a failure to reach desired purity and also can interact with target protein in different ways, e.g. promoting aggregation or even copurifying.

The challenge is clearly to develop a process capable of handling a feedstock which is constantly changing. The answer lies in two areas. (i) The direction of potential changes should be anticipated and the process developed with extra robustness built in to hopefully accommodate this. (ii) There must be good communication between the upstream and downstream groups. It is this latter which provides perhaps the greatest challenge. Having a suitable management structure and means of cross-communication could be one of the best investments made in to assist process development.

Clarification itself cannot commonly be achieved in a single step. Thus, multitrain operations will be required. These can be subdivided into primary, secondary and sterile steps. The primary step can be via centrifugation or filtration. As noted already, centrifugation is best suited to the small scale and very large scale, with filtration being preferred at the intermediate stage. Scaling down centrifugation steps is fraught with difficulty. This is because the discreet container format of small-scale centrifugation is so different to the disk-stack centrifuge. The results of a disk-stack centrifuge can be modeled by selectively reintroducing some of the pellet from a batch centrifuge back into the process stream to mimic the higher level of solids; however, this will not truly model the redistribution of small and large solids and colloids observed. Colloids, especially, can have a severely detrimental effect on subsequent filter membranes.

It is possible to replace the centrifuge with a filter. Either in tangential flow mode or as a normal flow depth filter. For a tangential flow filter membrane, an open screen type must be selected to cope with whole cells and large cell debris. Depth filters are available in a bewildering variety of sizes and formats. Although it may be tempting to think of all depth filters as equivalent, there can be differences in performance that when scaled up can have a major impact on the process economics.

Therefore, it is best to screen several clarification filters. This is especially the case for multilayer membranes. The use of multilayer membranes gives a great improvement in ease of use, especially where fully enclosed "cartridges" are available at process scale. However, the relative porosity and capacity of each filter layer can massively impact performance. Screening is essential for these types of membranes. One approach to screening such clarification membranes is the P_{max} test. For this simple procedure, the filter being tested has a pressure gauge placed inline at the inlet and the feedstock is pumped though the membrane at a constant flow rate. As the depth filter blocks the pressure will increase until the maximum pressure rating of the filter or some other nominal limit pressure is reached. Since the flow rate is constant, either the time or the volume passed through the filter can be plotted against the pressure increase. This can be done for a number of membranes and the membrane with the highest capacity can be determined. As a further check, samples of the output from the membrane should be taken at

intervals and checked for the presence of impurities, typically done by simply monitoring turbidity. This can also be plotted against the time or volume. If a membrane is too open it may allow a greater volume of feedstock through, but this feedstock may have too high an impurity content for the subsequent filter step. Hence, this complementary check should be performed on at least the two most promising membranes.

The clarified solution produced by the above test should be retained for use, selecting the sterilizing filter in the final part of the clarification process. As with depth filters, sterilizing filters are now available in a bewildering variety of type and sizes. Sterilizing filters should be coselected with the clarification filters. The relative area of sterilization filter required is a function of the efficacy of the clarification filter layer(s). Traditionally, sterilizing filters have a higher cost than depth filters. Where this is the case, the clarification membranes should be oversized to protect the more expensive sterile filters. Multilayer depth filters and cartridge-based systems have greater ease of use, save time, ensure system integrity and have much lower hold-up volumes. These advantages make them a good choice for selection; however, they have the penalty of higher costs. In such cases it may be economically better to optimize the clarification for the multilayer membranes at the expense of a greater area of sterilizing membrane.

1.4.2 Sterile Filtration

Sterile filtration has widespread use throughout a typical process – not just for filtration of product, but also for sterilization of buffer and growth media. Typically, a sterilizing filter will have a nominal pore size of 0.22 μm. Smaller pore sizes of 0.1 μm are also available; however, the reduced pore size also reduces the process flux, reducing throughput, and will foul more quickly, necessitating a higher membrane area for a given application. Hence, these membranes tend to only be used where it is felt Mycoplasma contamination may be a potential problem. Recent advances in membrane technology have produced high flow rate sterilizing filters based on polyethersulfone chemistries. These membranes have high flux, high retention and are robust enough to be sterilized, by γ-irradiation, prior to use. Because of these advantages they are becoming the default starting point for any membrane screening study. However, older membrane products made from regenerated cellulose are still the most common sterilizing membranes in current use in full-scale processes and do have the advantage of many years of validation data behind them. Again, it should be noted that membranes from different manufacturers, made with the same base chemistry and pore sizes, might have different behaviors in terms of flux and fouling for the same feedstock. Indeed, it is the author's experience that even using the same membrane may produce different results for different target process streams. Again, it is wise to screen a few alternative membranes to see if such anomalies are present rather than simply using the same membrane used previously. If there is only marginal difference in performance, there is probably more to be gained in manufacturing

efficiency and cost reduction by consolidation of processes around a preferred membrane or supplier, although validation of a sterilizing membrane from a second supplier, if time permits, will allow for a fall back position if required in the future.

Sterile filters are usually screened and sizing data generated using the V_{max} method, based on the theoretical maximum volume which can be passed through a unit area of membrane before flow is reduced to an unacceptably low level as a result of membrane fouling. In this test the process fluid is passed through a sample of the membrane under test at a fixed pressure, usually achieved using a pressure vessel attached to an air supply via a regulator. The changing flow rate is monitored by measuring the cumulative mass of liquid passing through the membrane over a period of time. This can then be plotted as a graph of (time/volume) versus time.

In contrast to the P_{max} method described above, the V_{max} method does require slightly more specialized equipment and, moreover, the analysis of the resulting data requires a mathematical model to be used which makes assumptions about the mode of fouling present. Fortunately, most membrane manufacturers can provide help in using the V_{max} method for membrane screening and scale-up studies. The V_{max} value gives an indication for future scale-up requirements. However, this will still need to interpreted within the context of the process requirements. The V_{max} value for a given membrane can vary from as little as a $200\,L\,m^{-2}$, for a serum containing culture media which has not been prefiltered, up to well over $10\,000\,L\,m^{-2}$ for a clean buffer. In the case of the former, a larger area of membrane will be required simply to complete the filtration. In the latter case, extra membrane may well be used simply to shorten the buffer preparation time, despite the filter membrane itself being underutilized.

Often when a process is developed at the bench sterile filtration is used generously to ensure that the process fluid is not contaminated during the purification process. Ideally, in an aseptic environment, this should no longer be an issue. Despite this, it is not unusual to see many of the original "bench-top" sterilization steps incorporated into the final process. This adds unnecessary time, cost and risk to the process. Before attempting to optimize any sterile filtration step in a process, it is always worthwhile to question whether the step is still required.

1.4.3 Chromatography

Chromatography lies at the core of all biotechnology purification processes. Despite this, there remains a lack of suitable texts on process-scale chromatography. Almost all the published literature on chromatography has been written for analytical applications. It is important to appreciate that analytical chromatography has a completely different set of aims compared to preparative chromatography. Usually when the literature speaks of optimizing a separation, it means achieving the maximum number of resolved peaks. In preparative chromatography there is only one peak or real value, i.e. that of the target molecule. If all other

molecules were eluted in only two peaks, one before the product and one after, this would represent an optimized manufacturing process. The emphasis in process chromatography is merely to ensure that the target peak is sufficiently separated from the closest peaks on either side. This should be the central aim of any changes made in elution conditions. One potential opportunity from this approach is the use of ion-exchange media in isocratic mode. If the conductivity and pH are carefully selected it is possible to load an ion-exchange media with feedstock and have the target protein retarded, while nonbinding proteins pass through the column and more tightly binding proteins are retained. Although requiring more development effort, this approach removes the need for gradient formation and can allow much higher loads of target protein to be separated in each purification cycle.

1.4.3.1 **Binding Capacity and Column Loading**

Column loading is an area which often causes problems in developing a chromatography separation. Again, in analytical chromatography, the problem is very simply solved by massively underloading the column. In large-scale operation this is not an economically desirable solution. Published data on binding capacities may give a rough indication of the load which can be reasonably expected, but these figures should not be used as the basis for a process. The usable dynamic binding capacity of a media must be determined for the specific target protein. For example, different monoclonal antibodies can exhibit large variance in dynamic binding capacity on the same Protein A media, even if the antibodies share 90% or more of a common amino acid sequence. Care should also be taken to differentiate between dynamic and static binding capacity. The static binding capacity represents the total amount of a specified protein that can be bound by a unit volume of the media. This figure is often quoted for ion-exchange media. Achieving saturation is not practical in the real-world; instead, a comparison of dynamic binding capacities under similar conditions using the same protein is the only reliable way to compare different media. Unfortunately, there is no standardized test used by all media suppliers. Thus, comparison of media solely based on data published manufacturers should be done with great caution.

For affinity chromatography media, such as Protein A capture columns, the question of how much to load can be answered reasonable simply. The column can be loaded to breakthrough point to determine the dynamic binding capacity and then loaded to around 80–90% below this value. In determining the breakthrough capacity, there a few parameters which need consideration. First, the dynamic binding capacity will be dependent upon the loading flow rate. The faster the loading flow rate, the less time the target molecule will have in the column to diffuse into the pores of the media. There will be a critical flow rate above which the dynamic binding capacity will rapidly decrease. The actual value of this flow rate will vary depending upon the mass transfer properties of the media being used. Media with very efficient mass transfer properties, such as controlled pore glass-based media or some highly cross-linked “open pore” agarose media, have very efficient mass transfer properties. These media can be used at high flow rates.

As a general rule highly cross-linked agarose media can be used at linear flow velocities of up to 500 cm h^{-1} in 20-cm beds. Controlled pore glass, on the other hand, being rigid and mechanically strong, can be operated at much higher flow velocities and in much longer bed heights, up 1000 cm h^{-1} in 40-cm high beds have been reported by the manufacturer. It is not necessary to use such extremes of bed height and flow rate, but the option of increasing either of these parameters beyond the fairly modest limits of agarose media does provide the process developer with much greater flexibility and ensures that the final-scale process can be operated well within the envelope of operation of the media, thus increasing process robustness.

Another consideration in determining the dynamic binding capacity is the selection of a breakthrough endpoint. Dynamic binding capacities are typically quoted at 10% breakthrough. If the shape of the breakthrough curve is steep, this will be only slightly higher than the actual point at which breakthrough occurred. However, for most media the breakthrough is not immediate, but gradual; the more gradual the breakthrough, the more flattering the 10% breakthrough figure will be compared to the actual capacity before target protein is lost. One method reported to increase the available capacity is to initially load a column at a high flow rate and then, as the binding capacity of the column is approached, to slow down the flow velocity by a factor of 2 or more and complete loading at the lower flow rate. This will allow greater time for diffusion of the target molecules in and out of the media in the search for the relatively few remaining binding sites. Although an effective approach to increase the binding capacity, the author is not aware of this method currently being used in a validated process.

For an ion-exchange column the amount to be loaded is a more complex issue. Overloading column will cause changes in the elution profile. Therefore, scale-up should be based on data from smaller columns that have had equivalent loading. Due to the relatively low cost of ion-exchange media, compared say to Protein A-affinity media, the traditional solution has been to oversize ion-exchange columns and only load to around 30–40% of the theoretical capacity. As processes increase in scale, however, this results in ever larger and more unwieldy columns. In the largest commercially available columns, with a diameter of 2 m, the amount of media required now represents a significant cost (630 L for a 20-cm bed). As a result of this one area of growing interest is the use of membrane-based ion-exchange devices.

This is especially true of flowthrough applications for ion-exchange media. An example can typically be seen with anion-exchange columns used in the final polishing step of many monoclonal antibody processes. The p*I* of most monoclonal antibodies allows them to pass directly through the anion-exchange media without binding, whereas several key impurities, such as DNA, endotoxin and most host cell proteins, bind to the anion-exchange media. In this case the relative amount of impurity is very low, typically less than 2%, compared to the amount of monoclonal antibody. Due to this the required binding capacity is also very low and this application is thus well suited to the use of an anion-exchange membrane system.

1.4.3.2 Throughput as a Chromatography Optimization Parameter

Determination of the dynamic binding capacity should be regarded as the first step in optimization a chromatography column. Unfortunately is it often also regarded as the final step. As noted above, the dynamic binding capacity will increase as the loading flow rate is decreased. Hence, the highest binding capacity will occur at the lowest flow rates. If time is not critical, the highest capacity can be achieved by loading over a period of hours. Consequently, it is not unusual to see laboratory-scale purification processes reported in the literature with the sample being loaded overnight, usually performed in a cold room to prevent sample degradation. However, this is not the most efficient mode of operation. In a manufacturing process the cycle time should be short enough to allow flexibility in scheduling, preferably allowing the operation to occur within one working shift with time for set-up and cleaning. The question is how to balance the compromise of a high dynamic binding capacity and a short process time. The answer lies in considering the throughput. Throughput is simply the mass of material purified per unit time. It can be seen that a high-capacity process with a very long cycle time will have a poor throughout. Similarly, very rapid process will also have a low throughput if speed is gained at the expense of too great a loss in capacity per cycle. In general, the loss in capacity is marginal compared to the reduction in cycle time until a critical flow velocity is approached. Because of this, the most efficient mode of operation is to have a relatively short residence time in the column, in the order of a few minutes for most media. In order select the optimum residence time, giving the highest throughput, the dynamic binding capacity should be determined over a range of flow rates. For each of these flow rates the cycle time should also be noted or calculated. The throughput can then be calculated as the mass of protein purified per unit time and the optimum value selected to give the most efficient process cycle. Throughput can also be further normalized to yield a values for the mass or protein produced per unit time per unit column volume. This is then a value which represents the productivity of a specific media or column. This value is useful when screening different column geometries (i.e. bed heights versus diameters).

1.4.4 Ultrafiltration

Along with chromatography, ultrafiltration is another ubiquitous process step encountered in biotechnology. It is commonly used for buffer exchange and product concentration, both between purification steps and for final formulation. Where possible, consecutive steps should be selected to minimize buffer changes of the need for concentration to ensure superfluous ultrafiltration steps are not present. Although there is some degree of separation, ultrafiltration has a low resolution, only being able to reliable separate molecules with an order of magnitude size difference, because of this ultrafiltration is not usually selected purely as a purification step.

It is not uncommon to encounter an ultrafiltration step which uses the "typical operating conditions" given in the manufacturer's literature. However, consideration of the desired outcome and some optimization will usually yield benefits in process efficiency either by reducing the cycle time, increasing product recovery or both.

1.4.4.1 Optimizing Tangential Flow Ultrafiltration

The first choice for a tangential flow step is to select the appropriate membrane. The pore size (or molecular weight cut-off size) quoted by a filter manufacturer is a nominal value, i.e. it represents an average pore size for the membrane. However, the distribution of pore sizes will affect the performance of the membrane both in terms of speed and product loss. For example, a process using a nominal 50-kDa molecular weight cut-off membrane from one manufacture may have poor product retention but be much faster than a 30-kDa membrane from the same manufacturer. However, a 30-kDa membrane, having a less homog eneous structure and thus wider distribution of pores sizes, perhaps sourced from another manufacturer, may give a better balance of retention and speed. Unfortunately, most manufacturers do not make such data widely available. It is therefore worth screening membranes from different manufactures even with the same nominal pore size, if time permits.

In tangential flow filtration systems, the early selection of the appropriate recirculation flow rate will have far reaching impacts as the process is scaled up to manufacturing levels. The permeate flux, i.e. the rate at which liquid passes through the membrane, is proportionally related to feed flow rate, i.e. the rate feedstock passes over the membrane.

The appropriate feed flow rate will maximize permeate flux while minimizing pumping requirements and maintaining a gentle environment for the product. It is critical to consider this relationship and its impact on product quality during process development, especially for later process scale-up.

Typically, a feed flow rate near the highest recommended by the filter manufacturer is selected. The main benefit of high feed flow rate is a higher permeate flux at any given transmembrane pressure. This is because the higher flow rate produces a greater sweeping action across the membrane surface, reducing the stagnant gel layer at the membrane and thus increasing mass transfer. The benefit of a higher permeate flux can be either a more rapid process operation or a reduction in the required membrane area for a given process time. Typically, the selected operating conditions will allow the process to be completed within a desired amount of time to fit in with scheduling requirements of the process. Usually the process is developed to minimize the surface area of membrane required. In the case of very long tangential flow filtration processes it may be worth increasing the surface area to reduce the overhead costs and allow more flexible process scheduling. It is worth noting that increasing the surface area will make the process step faster, but this may be at the expense of hold up volumes or reduced recoveries. Another mechanism by which the gel layer depth can be controlled is in the use of screens within the filter device to promote turbulence. With more

turbulence resulting in a smaller the gel layer and thus higher permeate flux. However, these screens are also sensitive to the nature of the feedstock, with the most efficient "turbulence promoters" creating a higher backpressure and being more prone to blockage. Therefore, care should to be taken to match the screen type with the feedstock, especially in the case of high concentration feed or feed prone to precipitation.

The disadvantages of using high feed flow rates are that this will increase the number of pump passes during the process cycle, assuming a smaller membrane area has been selected in preference to process time reduction. In practice, few protein are so shear sensitive that they will suffer denaturation from the additional pump passages; however, it is still worth considering, especially for nonglobular, multimeric or particularly large proteins. If it is proposed to use a current system for scale-up or manufacturing care should be taken that the flow rate selected at small scale will still be achievable at all future scales without significant investment in new hardware.

It should also be noted that, as concentration increases, the feed side pressure drop in the membrane device will also increase; this may become excessive in processes where the concentration factor or final product concentration is high.

1.4.5 Virus Removal

In the 1980s, HIV contamination of blood products brought the issue of virus inactivation and removal to the forefront of the process development. Since then the regulatory authorities have progressively tightened the requirements for virus removal. The current position is that phase I clinical material must have a validated virus removal step. In practice this means that there must be data to show that at least one model virus will be safely removed by the process used to manufacture the phase I drug. By the time a drug reaches phase II there must be a more complete set of data to demonstrate reliable removal of a panel of model viruses. The model viruses are selected to present the process with a set of realistic worst-case challenges. They will usually contain at least a model retrovirus, such as murine leukemia virus (MLV), and a small and robust parvovirus, such as mouse minute virus (MMV). Blood plasma products would also have a virus such as Sindbis as a specific model for hepatitis C virus. In order to validate virus removal or inactivation, a three-pronged approach is taken.

(i) *Raw material compliance*. All raw materials are evaluated for potential virus contamination. In the case of some raw material such as mammalian cell lines it is assumed that there will be an intrinsic viral load due to endogenous viruses. Other feedstock may have to undergo specific physical or chemical treatments to give some assurance of viral inactivation. This consideration has also been the driver to remove fetal calf serum and other animal-derived products from growth media recipes.

(ii) *Individual process step validation*. Each process step which is deemed to have a robust virus-reducing effect is validated to demonstrate this effect. Virus reduction can be by virus elimination, physical removal of the virus or virus inactivation. In the latter, the original virus particles remains, but are no longer capable of infection. These approaches can be measured by two different methods. Absolute removal can be measured by polymerase chain reaction (PCR)-based assays in which the amount of viral genome present is measured. Virus inactivation is measured by viral viability assays, in which the titer of active virus is determined by incubation with host cell. In both cases the amount of virus removed is quoted as a log reduction value (LRV). This is $\log_{10}$ of the ratio of the total virus load before clearance and the total virus load after clearance, e.g. a LRV of 4 would indicate that the step reduces the virus load in a test spike by a factor 10 000.

(iii) *Cumulative process clearance*. The virus clearance capabilities of individual steps can be added together to give a cumulative virus clearance capability for a process. To make this more robust the steps being added together should be orthogonal, i.e. to say they should have different modes of action (e.g. a chemical inactivation and physical filtration step). To obtain the overall clearance, the LRVs for each step are simply added together.

1.4.6 Specific Considerations for Virus Removal

1.4.6.1 Protein A-affinity Chromatography

Chromatography can be used as a method of virus removal. The mode of removal can be through passive partition (e.g. by size exclusion), by active partition (e.g. on an ion-exchange column) or by a chemical inactivation because of the chromatography conditions. Protein A-affinity chromatography provides a good example of the use of a chromatography step to remove virus. There are two mechanisms in action: physical partition through the column (monoclonal antibodies will bind to the Protein A ligand, while the virus will pass through the column unhindered) and chemical inactivation, due to the low pH of the elution buffer.

There some points which need to be considered with this step. There is the possibility of the media becoming fouled over prolonged periods of use. This may change the surface of the media, especially near the top of the column, and result in a media which will interact with and retain the virus particle. It would not be unreasonable to perform follow-up virus removal studies on used media to demonstrate no loss in the ability of the column to partition virus. Also, the chemical inactivation step is a result of holding the virus at a low pH, typically around pH

4.0. The hold step must be validated, and then specified at a specific pH or lower and at a given protein concentration or lower. Higher pH and/or higher protein concentration in the elution pool will affect the degree of chemical inactivation which occurs. The minimum hold time must also be specified. All of these parameters should be considered when designing the validation study so they will reflect a suitable safety margin for the final process

The two methods of virus assay mentioned above are particularly useful for the Protein A chromatography step. Assay by PCR will enumerate both active and inactive virus remaining in the elution pool. Comparison of this figure with the amount of virus in the spiking feedstock will show the amount of virus removed by partition alone. Analysis of viable virus in the elution pool will indicate the amount of virus reduction from both the partition and the chemical inactivation. Subtracting the latter from the former will give a log reduction value for the chemical inactivation alone. This method then provides log reduction values for the two separate orthogonal methods of virus removal.

1.4.6.2 **Virus Removal by Filtration**

A common physical method used to remove virus particles is absolute filtration. The filters used to remove virus obviously have a very small pore size, typically around 20 nm. This gives rise to the two most common problems with this technique – slow speed and tendency to block. Both of these issues are addressed by placing the virus filtration step late in the purification train. At this stage the concentration of product is high, thus reducing the volume to be handled, and also the product is pure, reducing blockage due to nondrug moieties. Although this is a generally valid approach, the volumes being handled may still be appreciable at manufacturing scale, resulting in a significant cycle time. Much work has been performed by the filter manufacturers to try to improve the flux of virus removal membrane, helping to reduce the process time issue. However, the blockage issue remains a problem. This is compounded by the virus spike method used to validate the filters themselves. In order to validate virus removal a high titer spike of virus is passed through the membrane. Such high-titer virus preparations can be highly variable in quality. At worst they can be so heavily laden with particulate cell fragments that they cause filter failure long before the filter itself would fail if presented with the "normal" feedstock. The result will be a significant overspecification of membrane area required for virus removal. Given the high cost of virus removal membranes, this presents a significant threat to the process economics. The answer lies in the use of suitably purified virus spike solutions.

Another consideration for the methodology used in validation of the membrane is the time at which the virus spike is applied. Traditionally, the spike is applied first and then the membrane is used to filter feedstock afterwards. A dirty spike sample can block the membrane and artificially reduce the volume of feedstock the membrane is able to handle. A more representative method is to run the required amount of feedstock through the membrane, to represent a used membrane, and then spike with virus. This has the double benefits of allowing a more

representative feedstock to be used to size the area of membrane required and to produce virus removal validation data on "used", i.e. worst-case, membranes.

1.4.7 Lifetime Studies

The real test of a process is not how well it works not on the first run, but how well it performs on the final runs. Repeated exposure to feedstock can have a progressively deleterious effect on chromatography media. The problem is somewhat bypassed for filtration since most filtration media are designed to be used only once. However, ultrafiltration membranes can, and indeed often are, used for repeated cycles too. It is important to validate the potential lifetime of any step of a purification process that will be exposed to multiple cycles of use. The effect of a shorter than expected lifetime can be catastrophic to the overall economic viability of a process. Conversely, if the media can be used for more times than expected this can give an added bonus to the profitability of the process. Either way, it essential to be aware of the expected operation lifetime.

An argument could well be made that the best time to generate validation data is at the end of a series of production cycle. Unfortunately, this is both too late and too expensive an exercise at full scale. During a series of processing runs data can and should be collected to assist in retrospective validation of a process. However, this is at best only support data. Instead, a lifetime study is an essential part of process development. It is widely acknowledged that lifetime studies are labor intensive and, from a scientific point of view rather boring. This is especially the case for a successful process – with a long lifetime and no problems. Lifetime studies are is also not a great way to advance a scientific career. Months of work may effectively generate only one graph showing, hopefully, that not much is happening! Despite this, lifetime studies are absolutely critical to the validation of a process, and to also give and idea of the eventual economic viability of the process.

Lifetime studies are made easier by automated fast protein liquid chromatography systems; however, there still remain some caveats to be aware of. The feedstock should ideally be from multiple sources rather than one single source, although this may be difficult. Using one feedstock throughout a study can skew the results if source represents either a better or worse than average case. Feedstock should be stored and used in aliquots that model the final process. A single lot of feedstock standing for a period of weeks or even months is unlikely to still be representative by the end of the study. Ideally, aliquots to represent single campaign batches should be stored separately and frozen if necessary. It is likely that the real feedstock will not be frozen in the full-scale process. If this is the case the validity of freeze–thawed aliquots of feedstock as representative must be validated too. It is the experience of the author that some feedstock, left on the bench for only a few hours, will show significant signs of aggregate formation which, while not being visible as cloudiness, will affect the elution profile of a chromatographic separation.

1.5
Future Trends in Process Development

There are some clear trends in process development, some of which are closer to general implementation than others.

1.5.1
Disposable Process Lines

The use of disposable bags for buffer preparation has been common for a long time in pilot plants. There has also been strong growth in the use of disposable bag fermenters. Manufacturers are now increasingly able to supply disposable assemblies for even more of the purification train, including disposable pipework and tank liners, and filtration assemblies. It may even be envisaged that disposable chromatography columns will one day be available to allow a fully disposable purification train removing concerns over cross-contamination or post-use cleaning. Although the use of fully disposable manufacturing systems may still be some way off, the ease of use and fast turnaround for different product lines makes this approach highly desirable for pilot process development laboratories.

1.5.2
Nanoscale Screening

High-throughput screening techniques have been developed and are now well established for drug screening. There is a clear move towards adapting these techniques for the subsequent process development stages too. Small-scale screening can be used to determine which media could potentially be used to purify the target molecule and also to get a good indication of the optimum conditions for purification. For example, in a single afternoon the experimenter can determine not only which media to use, but also the best buffer, pH and conductivity for loading the media, and similarly the best buffer, pH and conductivity for elution. This massively reduces the total time typically required to generate a near-optimal purification method. A further advantage of this approach, using multiwell plates or a similar format, is that only very small amounts of sample are required allowing the work to be done even earlier in the development cycle.

1.5.3
High-titer Feedstocks

As noted, there has been a progressive trend towards higher product titers as cell lines and culture conditions are better understood and optimized. In the late 1990s, a monoclonal antibody titer of around $0.1\,g\,L^{-1}$ in the clarified cell culture supernatant was not untypical. Currently, titers around $1–2\,g\,L^{-1}$ are the norm and there are reports in most cell culture conferences of titers of $5\,g\,L^{-1}$ or higher at the development stage.

As the titer approaches such high levels the relative amount of product, antibody, to contaminant, cell culture protein, changes. It is possible at these very high titers that the most common first step used in monoclonal antibody production, i.e. Protein A-affinity chromatography, will become obsolete as the starting product will already be sufficiently concentrated and pure. The role of the first step of the purification process then shifts from capture directly to purification.

1.5.4 High-concentration Formulations

Another area becoming increasingly prevalent is the use higher-titer formulations. Some biological drugs such as antibodies require significant mass to be introduced into the patient in each dose. Concentrating the drugs ever higher makes administration easier. It is already not unusual to see antibody concentrations in excess of $100\,g\,L^{-1}$. Although the antibody may well be stable at this concentration, this approach does generate some distinct process problems for the final concentration and formulation step. On concentration the titer and the viscosity increases. Very high concentration factors can cause problems with ultrafiltration membranes. A membrane which is optimal for lower viscosity may have such a low permeate flux at high viscosity that it becomes virtually unusable. Also, the dead volume in a system, i.e. the cumulative volume of material in the pipework and inside the membrane housing, may be such that, at high concentration factors, a significant part of the product cannot be recovered. Additional steps may have to be introduced to increase recovery, such as using clean air to flush out the membrane holders and pipework.

From these considerations, it can be seen that the role of the process development scientist is not only a vital bridge between R & D and manufacturing, but also between past conservatism and exciting new trends. Fundamentally, it will be the process development scientists who will have to figure out how to apply ever more cost-effective methods to produce ever-larger amounts of the next generation of biotechnology products.

Further Reading

For further introduction to process development, especially with regard to the purification of monoclonal antibodies, the reader is directed towards the following excellent sources.

Sofer, G., Hagel, L., *Handbook of Process Chromatography – A Guide to Optimization, Scale-up and Validation*. Academic Press, London, **1997**.

Subramanian, G. (Ed). *Antibodies. Volume 1: Production and Purification*. Kluwer, New York, **2004**.

Subramanian, G. (Ed.). *Bioseparation and Bioprocessing (Volumes I and II)*. Wiley-VCH, Weinheim, **1998**.

US Food and Drug Administration Center for Biologics Evaluation and Research. Points to consider in the manufacture and testing of monoclonal antibody products for human use. *J Immunother* **1997**, *20*, 214–243.