Affinity Chromatography: The Fine Print

by Pete Gagnon, Validated Biosystems

Affinity chromatography is often held up as the ideal in chromatography. Exquisite specificity. Unmatched simplicity. Load it, wash it, elute the purified product. What could be better? Affinity chromatography could be better for one thing. There is a very high price tag for both the specificity and the simplicity, and this isn't referring to the dollar cost for the media itself. Affinity chromatography involves complications that have immensely important ramifications in purification process development. The only thing more costly than the complications themselves, is overlooking them.

Product denaturation. One of the most serious and persistent concerns about affinity chromatography is the potential for product denaturation. This is especially the case with strong affinity ligands that require harsh conditions for elution. The elution method of choice is usually exposure to low pH, typically in the range of pH 2.5-3.0. Detailed studies of protein conformation under these conditions have documented permanent conformational changes as a result of such exposure. Hydrophobic residues normally protected in the interiors of the primary structural domains become exposed on the surface, increasing the tendency for the protein to become involved in nonspecific hydrophobic interactions. Elevated tendency toward aggregation is one of the results, as is easily proven by examining size exclusion chromatography profiles following affinity purification.

A usually less obvious but still common side effect of such changes is an elevated tendency toward proteolysis of the product. This tends not to be noticed unless it's severe, but it's common nevertheless. The most insidious side effect — and the most overlooked — is modification of secondary effector functions. Antibodies are a good example. The non antigen-binding parts of the molecule are richly endowed with receptors that interact with a variety of proteins, carbohydrates, and cells in the immune system. Conformational modification of these receptors can alter their functionality. All of these factors — elevated tendency toward aggregation, proteolysis, and alteration of effector functions — can alter product effectivity, safety, and pharmacokinetics — and this is before even taking into consideration that the elution conditions may have altered a protein's primary therapeutic function.

One of the ways to deal with product alteration is to develop milder elution conditions. The first objection to this suggestion is usually that "If I use a less extreme pH, then the elution will be less effective." In fact, elution pH can often be raised without any loss of effectivity, and it's certainly worth evaluating a given application to find out what the real limits are. If modifying pH alone doesn't bear fruit, there are other opportunities. Biological affinity interactions are mediated by complex combinations of hydrophobic interactions, charge interactions, hydrogen bonding, and other mechanisms. Instead of using just low pH to twist a protein out of shape so severely that it can't remain bound to an affinity ligand, it is almost always possible to target one or more of the actual binding mechanisms. This alone usually won't be sufficient to cause elution, but it almost always ameliorates the severity of the pH conditions required to elute the protein. For example, hydrophobic interactions can be weakened with up to 50% ethylene glycol. Ethylene glycol is actually stabilizing to most proteins up to this level, but at the same time it is a very effective polarity reducer. It's also

nonionic and won't interfere with downstream charge-based purification methods. Hydrogen bonds can be suspended by the inclusion of 1.0M urea. At this concentration it's conformational effects on a protein are nil, and like ethylene glycol, it is nonionic. 0.5-1.0M NaCl will help to suspend any charge interactions, without significant enhancement of hydrophobic interactions. The alternative use of a chaotropic salt like sodium perchlorate at the same concentration will likewise suspend charge interactions and may further weaken the interaction through its chaotropic properties.

Using combinations of these mechanisms often makes it possible to raise elution pH by a full pH unit; sometimes substantially more. The question invariably arises: Don't these complex formulations exert the same denaturative effects as pH alone? They don't. Think of it like this. Let's say that you have a particular oak board that you would like to remove from a wall so that you can use it for something else. It is screwed to the wall, nailed to the wall, and glued to the wall. If you just try to rip it off with a crow-bar, your chances of recovering it intact are grim. But if you remove the screws and nails, then weaken the glue by pouring hot water around the edges, you can remove it with much less force, and your chances of recovering it intact are improved proportionately. This raises the second objection to this approach: it cancels out part of the simplicity of affinity chromatography. Library research will often lighten the load by revealing the dominant mechanisms of an affinity interaction, but there will still need to be some experimentation to develop the most effective while least denaturing elution buffer.

In spite of taking great pains to develop a nondenaturing elution buffer, you may still find that your product has an elevated tendency to aggregate, or that it exhibits elevated vulnerability to proteolysis, or that it exhibits modified primary or secondary functions. The ugly truth is — to varying degrees — that this is largely inevitable. The best studied affinity interactions consistently reveal that when binding occurs, it is accompanied by a phenomenon called induced fit. Induced fit refers to a situation where after coming into contact with one another, either the affinity ligand, its receptor, or both, undergo conformational changes that lock them into place. One of the best practical indicators that induced fit is occurring is when you can bind a product to an affinity ligand under mild conditions, but very harsh conditions are required to remove it. One of the best characterized examples of this is the binding of protein A to IgG, which binds in the hydrophobic cleft between the Cg2 and Cg3 domains. X-ray crystallographic data shows that the Cg3 domain is unaffected by the contact, but the adjacent Cg2 domain is displaced longitudinally toward the protein A and Cg3. Besides altering the local conformation, this destabilizes the receptor-rich distal third of the Cg2 domain, which in turn causes a partial rotation and destabilization of the carbohydrate domains between the Cg2 domains. The effect is apparently permanent. No matter how careful you are to develop gentle elution conditions, comparison of the purified product against a non-affinity purified control virtually always confirms measurable changes in key behavioral features of the product. This is not to say that developing mild elution conditions is not worthwhile. Doing so can make an immense difference. But it may not be able able to avert denaturation problems entirely.

Leaching. The second major concern with affinity chromatography is leaching of bioactive ligand and contaminants that may be associated with it. This issue gets underplayed but it is very serious. The inevitability of leaching is the reason the FDA insists that any biological affinity ligand used in the manufacture of a biological product meet the same application requirements as the end product itself. This extends even to how the affinity ligand is purified. For example, the protein A going into many chromatography products is purified by affinity chromatogra-

phy on immobilized human polyclonal IgG. The IgG column is potentially contaminated with virus, which can potentially leach into the purified protein A, and from there into your final product. Even to the extent that virus contamination is adequately controlled, there is still the issue that protein A is derived from pathogenic bacteria, and may contain dangerous contaminants. Some chromatography media manufacturers have acknowledged these issues by obtaining their protein A exclusively from recombinant lines of non-pathogenic bacteria, and purifying it by non-biological methods. If you are making your own monoclonal antibody columns for purification of a biological, pathogenicity of the cell line in which the monoclonal is grown is not an concern, but you still have to deal with the virus issue, and with the same media components that are scrutinized for your end product — BSA, host cell proteins, DNA, endotoxin, etc.

Unfortunately the issues discussed above represent the simplest aspect of leaching. The biological ligand itself is likely to have more significant effects on product safety and efficacy than any of the contaminants that may be co-immobilized with it. In the case of protein A for example, more than 150 publications describe its interference with virtually every immunological mechanism known. This shouldn't be surprising. In fact it would be more surprising if it were not the case; it has known major effects on the functions of IgG — a molecule that can be reasonably regarded as the hub of natural immunity. It should be equally obvious that parallel concerns apply to any affinity ligand. The whole basis of affinity chromatography rests on its ability to discriminate the unique features of a given molecule, from molecules lacking those particular features. The problem arises with the fact that any feature that makes a particular molecule unique, is probably going to be linked directly to some aspect of its function. Using antigens and antibodies as an example, the linked functions may be primary, like antigen binding, or they may be secondary, like complement fixation. This leaves you with a high probability that leached affinity ligand will interfere with either the direct function of your product, or with other functions that may modulate efficacy, alter biological clearance, or affect other aspects of pharmacokinetic behavior.

Expectedly, chromatography media manufacturers go to considerable lengths to minimize leaching. Some even claim to have developed immobilization chemistries that eliminate it. This is naive however. Breakdown of the chromatography matrix or immobilization linkage has been shown to be a very minor source of bioaffinity ligand leaching. It arises mostly from enzymatic degradation. With immobilized antibodies, proteolytic clipping of Fab or F(ab)'2 fragments is common, the enzymes coming from the column feedstream, or copurified with the original antibody preparation. With immobilized protein A, heavy proteolytic degradation is widespread, and virtually uncontrollable. This is assured by the ligand structure. Protein A consists of 5 compact, cylindrical, protease-resistant IgG-binding domains, joined together by notoriously protease-labile sequences of random coil. Other ligands exhibit their own unique breakdown patterns, but they all break down, and they all leach.

This makes removal of leachate essential. The obstacle here is that the leached ligand is bound to your product. The same affinity that allowed the column to capture your product initially, ensures that leachate will be reassociated with your product as soon as buffer conditions permit. If the molecular characteristics of the leachate are substantially different from those of your product, it may be possible to discriminate product-leachate complexes from uncomplexed product. For example, if your product binds poorly to an anion exchanger but the leached ligand binds very strongly, complexes may exhibit intermediate behavior that allows them to be selectively removed. You'll have to sacrifice some of your product in the bargain (whatever portion is

complexed with leachate) but this approach sometimes works. More often, the behavior of the complexes is sufficiently different to support a reduction of leachate, but not quantitative removal. For example protein A is more hydrophobic than IgG, but only slightly. The elution behavior of product-leachate complexes from a HIC column is so similar to uncomplexed product that a 50% reduction of leached protein A may require sacrificing of 50% of the antibody.

A usually more successful approach is to use a separation method carried out under conditions in which the leachate is dissociated from the product. For example, if you have small leachate fragments bound to a fairly large product molecule, you may be able to separate them by size exclusion chromatography under buffer conditions similar to those used to elute the original affinity column. The chief risk with this approach is that the prolonged exposure of the product to these conditions may be prohibitively denaturing. Cation exchange has proven to be a good alternative. Most proteins will bind to cation exchangers if the pH is low enough, for example a pH known to dissociate the leachate from the product. As long as the leachate has elution characteristics sufficiently different from the product, you may be able to fractionate them from one another. The reason that this scenario is usually better than size exclusion is that when proteins are immobilized on cation exchangers, they are sterically restricted and thereby resistant to the conformational changes they would undergo in the same conditions in free solution. Many proteins that denature rapidly in low pH solutions remain stable for hours when bound to a cation exchanger at the same pH; ample time for removing the leachate.

If for some reason you can't quantitatively remove leached ligand, you still need to quantify it and validate that the maximum levels occurring in your product have no adverse affect on product safety or efficacy. Even this is more complicated than it would seem on its face. If there is leached ligand in the system, it is bound to your product. The association may sterically occlude antigenic sites that are necessary to obtain accurate measurement. This is not just a theoretical possibility. It's a proven problem. To obtain accurate measurement, you need to have a similarly configured standard curve. Protein A affinity for IgG varies with species, class, and subclass of the antibody. The stronger the affinity of the protein A for the particular subclass, the shallower the curve. This means, that you have to use non-protein A purified antibody of the same subclass as your product to make your standard curve. The accuracy of the curve is also affected by conformation of the protein A: single binding domains, doublets, triplets...whole molecule. This is an especially troublesome point because, to be accurate, the conformation of the protein A in the standard curve should be the same as in the sample. Unfortunately there's no direct way to know what the conformation is in the sample. There are other complications as well. From a practical perspective, the only way to obtain unassailably accurate measurements of leached protein A, is to separate it entirely from your product, and measure the amount of recovered protein A versus a standard curve of purified protein A. Similar difficulties apply to other affinity ligands.

Cost. Some of the other negative features of affinity chromatography are better recognized. The most obvious is cost. Commercial bioaffinity ligands generally range from 7-15 times the cost of ion exchangers on the same base matrix. If you are making your own affinity media, the direct cost may be reduced, but by the time you include your development costs, it will probably end up being more expensive than commercial media. This cost ripples through your overall purification costs. Documented comparisons show a 1-step affinity purification to be 9 times more expensive per unit of product than a 2-step ion exchange process. In the case of products intended for in vivo applications, the cost will be compounded by the

need for additional steps to remove leachate and other trace contaminants. At a recent symposium, a representative of a highly competent pharmaceutical company stated that they were able to achieve better than 99% purity with their initial bioaffinity purification step...then added that three additional chromatography steps were required to meet the full scope of regulatory requirements.

Reduced operating life is another limitation with biological affinity media, and another elevated cost factor. Most cannot withstand the harsh conditions used for cleaning and sanitizing nonbiological media, or if they can, they can do so for a much lesser number of process cycles. Many are also subject to biological degradation from the feedstream. Protein A, as mentioned above, is highly labile to proteolysis, and especially under the neutral to slightly alkaline conditions at which feedstreams are normally applied. As significant as initial procurement and replacement costs are, they are trivial in comparison with the elevated validation costs that accompany the issues of product denaturation and ligand leaching. Beyond validation, especially if you are in a competitive market, reduced product performance may be the most costly factor of all.

Do all these factors mean that you should avoid the use of affinity? No. But they do mean that you shouldn't accept it's superficial simplicity at face value. Knowing the potential effects on product safety and efficacy, knowing that it will require method development — both to control denaturation and to remove leachates, and considering the elevated validation costs because of those factors, ask yourself this question: Does the inclusion of an affinity step in your process offer fractionation performance so much better than nonaffinity alternatives that it can justify the liabilities? If the answer is yes, then it's a candidate.

Several parts of this article are adapted from the book Purification Tools for Monoclonal Antibodies (ISBN 0-9653515-9-9). Most of the key technical points are supported by citations in Chapter 9.