

Writing an Application for a Software-Based Invention A Programmer's Guide

Copyright © 2007 by Lee A. Hollaar

September 20, 2009 version

Introduction

You've just developed a computer program that does something in a way that you think has never been used before. Since you have taken a class on intellectual property, you know that the way to protect the new technique is with a patent. But what do you do to get that patent?

If you are working for a company, you might think that all that you need to do is write up a short disclosure document, or even a memo with the code attached, and give it to your supervisor, who will then forward it through the bureaucracy until it gets to somebody who will turn it over to a patent attorney.¹ You might have to sign some forms, but most of the work will be done by others.

Wrong! If the resulting patent is going to have any value, you are going to be heavily involved in the application and prosecution (convincing a patent examiner to grant the patent) process. After all, nobody knows better than you how the new technique works, how it is an improvement over past techniques, and how it might be used in other contexts or modified while still giving the same benefits. Your work in educating the patent attorney will be invaluable. And since patent attorneys often bill by the hour, it may be desirable for you to write the first draft of the patent application.

If you are working for yourself, how much you do on drafting the patent application may make the difference between whether you file for a patent or don't because you can't afford it.

There is no way that a short note like this can give you all the information you need to produce the final draft of a patent application, file it, and successfully prosecute it in the patent office. You will still need the help of an experienced patent attorney. But it will allow you to be prepared for working with the patent attorney, reduce the costs, and can result in a better patent.

First, determine the invention

It sounds pretty obvious to say that the first step is to determine the invention (or inventions) that you want to patent. But all too often, people start out by writing a general description of the system they have just completed. That will result in a patent application that is overly long (which the examiner will hate) and yet may not be as descriptive of the actual invention as the statute requires:

¹ I'll use the term "patent attorney" to cover both patent attorneys and patent agents, rather than the normal collective term "patent practitioner." Both patent attorneys and patent agents have passed an examination to allow them to represent you before the patent office, and are regarded the same by the patent office. The difference is that a patent attorney can also represent you in other legal matters such as the assignment or licensing of the application or patent.

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same, and shall set forth the best mode contemplated by the inventor of carrying out his invention.²

Note that there are three different requirements. The “enablement” requirement that the specification teach the invention so that somebody with ordinary skills in the area of the invention or its use can recreate it or use it. The “best mode” requirement that the inventor has to tell how, at the time of the application, he thinks it should be made and used.³ And the “support” requirement is that each element of every claim must be sufficiently described in the specification so that it can be understood.

In most cases, the invention will not be the total system, but that portion of the system that makes it work better than prior systems. You do not want to think of the invention in the context of a particular application unless absolutely necessary. If you have created a database system that runs substantially faster than past database systems because you have come up with a new sorting technique, what you have invented is *not* a faster database system, but instead a better sorting technique. That way, if somebody uses your new sorting technique in a different application, he will still have to license your invention.

Quite often, a system may contain a plurality⁴ of different inventions. Some people write a specification that describes all the inventions (along with other parts of the system) and then file applications for each invention that contain that specification and claims particular to the invention of the application.⁵ This is not a good idea, since it results in extraneous material in the application for a particular invention that might cause a problem if the patent is litigated and may not result in the best description of the particular invention.

Having determined the portions of the overall system that are the inventions, the next step is to generalize each invention. This means taking it out of the context of the particular system and thinking about it in a more general way. For example, with the better sorting technique, you would stop thinking about it as part of a database system and start thinking about it as a way to sort things on a computer. This may get you thinking about different input and output forms, data structures, or other implementation details.

How general you make each invention will depend on what you know about the prior art for that invention. Your eventual goal is to come up with a claim that is as general as possible for the invention, yet is not obvious in light of the prior art. This would be a good time to get your coworkers involved in a brainstorming session for each invention.

² 35 U.S.C. § 112, first paragraph.

³ The “best mode” requirement is not part of most other countries’ patent laws, and while seeming reasonable for preventing an inventor from getting a patent while not “telling all” about the invention, in reality is often used as questionable defense by an infringer.

⁴ A handy word often used in patent claims meaning “two or more.”

⁵ The patent office will require that each application be for only a single invention. This is because they collect their fees based on applications, not inventions, and the examiner gets performance credit (called “counts”) for the number of applications handled.

But keep in mind who suggests what, since if something suggested by a coworker becomes a claim element in the application, you may need to list them as one of the inventors on the application.⁶

Second, determine the prior art

Now that the general nature of the invention has been determined, it's time to see what can be covered by a patent. What you are trying to do is describe the key aspects of the invention with enough elements or limitations so that the claim does not "read on"⁷ the prior implementations that you know about.

It is quite likely that if you are working in an area, you are already aware of much of the prior art in that area from things that you have read or other systems that you have seen. This is where you have a big advantage over the patent attorney, who must learn about the invention and its area from what you give him.

What a patent attorney often does is send a copy of your disclosure, or a brief description of your invention, to a professional searcher who will search issued patents, or patent applications that have been published.⁸ Since everything can't be searched, what is often done is to indicate a cost for the search (say, \$1,000) and the searcher starts at the most productive places and works until the number of hours paid for are done. The searcher will often talk with a patent office examiner who works in the area of the invention to get suggestions of where to search. Copies of any patents that look relevant to the invention are sent to the patent attorney, who will go over them with the inventor.

It used to be that most searches were conducted at the patent office just outside of Washington DC. Patent depository libraries, such as the University of Utah's Marriott Library, had microfilm copies of patents, but they were not easily searchable. This has changed since the patent office put its collection of patents and applications on the Internet. Now, everybody has access to the same patent information and it is not necessary to go to a patent depository library, although if you do a trained librarian will be able to help you with how to search.

There are two ways that the collection can be searched. The first is by words or phrases that occur in the patent as well as information in fixed fields such as inventor name. This works well if there are already standard terms that describe the invention but are not something that is going to occur too often in patents that aren't relevant. The problem is that since patents are by their nature about new technology, often there isn't a

⁶ The requirement is that any inventor of any aspect of any claim be listed, but that people who are not inventors must not be. (It doesn't do your boss any favor to list him as an inventor when he really wasn't, since that may result in an invalid patent.) But just because somebody worked on the implementation of the invention doesn't make him an inventor. 35 U.S.C. § 116 allows later correction of inventors if the original list was done "without any deceptive intention," so if there is any question about whether somebody is an inventor, discuss it with the patent attorney.

⁷ Another patent term, meaning that each and every element of the claim is present in something. For example, if the claim has elements A,B,C, then it reads on A,B,C,D but not on A,B. If a claim reads on the prior art, it is not patentable.

⁸ Until 1999, all pending United States patent applications were kept confidential. Now, applications are published eighteen months after they are filed, unless the applicant has requested that an application not be published and it has not been also filed in a foreign country. See 35 U.S.C. § 122.

accepted term for something at the time of the earliest patents and the applicant had to make up a term, which may or may not have been accepted over time.⁹

The other way to search is to take advantage of the classification scheme used by the patent office.¹⁰ Every patent is manually placed in one or more classes and subclasses when it is issued. The class/subclass for many things are listed in the Index to the United States Patent Classification (USPC) System¹¹ and descriptions of what is in each subclass is given in the Manual of Patent Classification.¹² When too many patents are in a subclass, the patent office splits it into a number of subclasses or even establishes a new class. For example, when the patent office first started issuing patents on software-based inventions, they were put in a couple of subclasses in the class that contained digital computers. Now, “database and file management or data structures” has its own class (class 707).

Instead of rummaging through the Manual of Classification, the best way to determine subclasses to search is to find a patent close to your invention. If there is a patented product like yours, it may have its patent numbers marked on it or otherwise available.¹³ Or you may find a close patent by doing a word search, or searching for all the patents assigned to a particular company making a similar product. That patent will indicate a number of subclasses, and you can check the description of each in the Manual of Classification to see if it is a good one for further searching.

Keep track of every article and patent you find that is closely related to your invention, since you may want to submit them for the examiner’s consideration unless you have structured your claims so they clearly avoid that art.

Then determining the claims

With a good understanding of the invention in its broadest form and the related prior art, it’s time to figure out how to claim it. We’ll first look at the claim elements in general, and then see how we can turn the elements we select into a claim as either a method or a machine. At the end of this note is a very simple patent application, and you can see how to do these claims and the language added to the specification to support them.

The first question you need to ask yourself is whether, in light of the prior art that you know, you can write a claim such that a competitor cannot avoid it without departing

⁹ “An applicant is his own lexicographer” is the phrase often used.

¹⁰ Searching by class/subclass is the traditional way of doing patent searches. At one time, the patent office had huge rooms filled floor-to-ceiling with compartments holding patents. A copy of each patent was placed in the compartments for every class/subclass to which it was related. You’d go to it, load the contents of the compartments for the subclass of interest on a cart, take them to a table, and flip through them, looking at the title, drawing, and abstract on the first page. If it looked interesting, you’d read more of the patent. It would be impossible for a library to duplicate this collection, which is why most professional searching was done at the patent office.

There are still advantages in doing searching at the patent office’s Public Search Facility in Alexandria, Virginia. The workstations there provide the same user interface as used by the examiners, which takes advantage of the higher-speed connection to the database than the Internet web search.

¹¹ <http://www.uspto.gov/web/patents/classification/uspcindex/indextouspc.htm>.

¹² <http://www.uspto.gov/web/patents/classification/>.

¹³ Patent owners do this because if they don’t mark the product, they are not entitled to any damages until they have actually notified an infringer. See 35 U.S.C. § 287(a).

from what you have invented. For a method claim, that means that the competitor either does not perform one of the steps in the claim, or performs a step differently.

It is important to remember that you now developing a patent claim, not a description of how to perform the method. You'll do that in the specification. In particular, you will not want to include some steps, even though they are important for the functioning of the method, because they are not needed to distinguish the invention over the prior art.¹⁴

Again, this a good time for brainstorming with your coworkers. And you might also get your patent attorney involved at this point.

As an example, consider the quicksort algorithm,¹⁵ a recursive sort that uses a “divide-and-conquer” technique of splitting the list to be sorted into two subsets – one containing items less-than or equal-to a “pivot element” and the other containing items greater-than or equal to the pivot element – and then sorting those two subsets and combining them. The method form of the claim might be:

A method running on a digital computer for sorting a plurality of items stored in a collection in the computer's memory, the method comprising:
selecting a pivot element; and then
partitioning the collection into a lower subset comprising items less than the pivot element and an upper subset comprising items greater than the pivot element; and then
sorting the items in the lower subset; and
sorting the items in the upper subset.

(Note the use of “and then” to explicitly state the ordering of the steps, and just “and” to indicate that the steps can be performed in any order. This is helpful to people reading the claim, but not required, since unless there is some indication to the contrary, the steps of a claim can be performed in any order for there to be infringement.¹⁶)

One can see a number of things missing from this if it were instead to be an algorithm description: how do we select the pivot element, what do we do with items equal to the pivot element, and how do we sort the items in each subset. It doesn't say that we have to use the quicksort for sorting each subset (or even the same technique for each subset), or how we terminate the recursion. (Most descriptions of quicksort terminate by determining that there is only one or perhaps no element to be sorted, and returning that.)

But if we included such details, particularly in our broadest claim, it would be easy to avoid infringement. For example, if we included the common termination rule, all somebody would have to do is use another sort technique when the subset to be sorted is small. We might even want to do that, using a simple sort for lists of three or fewer

¹⁴ Sometimes, for no clear reason, patent attorneys include inexplicable claim elements such as “a power supply” or even one power supply for one part of the invention and another for another part. While this may seem harmless, it will at least result in additional effort when trying to prove infringement, since each element of the claim must be proven to be in the alleged-infringing product. And if two power supplies are mentioned in the claim, there is an argument that a device with only one power supply doesn't infringe. If something is not needed to distinguish the claim over the prior art, *leave it out!*

¹⁵ See <http://en.wikipedia.org/wiki/Quicksort>, or your favorite algorithms textbook.

¹⁶ See *Altiris v. Symantec*, 318 F.3d 1363, 65 USPQ2d 1865 (Fed. Cir. 2003).

elements to avoid the overhead of the recursion. But right now, we want to avoid easy ways to get around the claim while still claiming the essence of the invention.

Having our broadest claim, we then think of different elaborations that can be used in dependent claims:

Selection of pivot element:

- The pivot element is selected at random.
- An item of the collection is selected as the pivot element.
- An item of the collection closest to the arithmetic mean of the items of the collection is selected as the pivot element.
- An item of the collection closest to the arithmetic mean of the smallest and largest items in the collection is selected as the pivot element.

Which subset gets items equal to the pivot element:

- Items equal to the pivot element are included in the lower subset.
- Items equal to the pivot element are included in the upper subset.

How the subsets are stored:

- The partitioning forms two new collections separate from the collection.
- The partitioning rearranges the items in the collection.

How the subsets are sorted:

- The lower subset is sorted using the claimed method.
- The lower subset is sorted using a bubble sort.
- The lower subset is sorted using a heapsort.
- The lower subset is sorted using a mergesort.

What the menu does is give you the chance to determine the language that you will be using the claims you eventually write. And one of the most important things in writing a patent application is always call the same thing by the same name, no matter how boring it makes reading the application.

Of course, in the final application you are not going to include all combinations of these options. When you consider that in addition to what is listed above, each category also includes not having the element in the claim, there are 225 possible claims. If we were to specify a different sort for each subset, there would be 1125 possible claims, and if we consider that each claim also has claims for a system and an article of manufacture, we end up with 3375 claims just for this simple technique!¹⁷

There is a way to indicate a number of options for a claim element, called a Markush claim after the case that found it acceptable. For example, you could indicate how the lower subset is sorted with a single dependent claim such as “The method of claim X, where the lower subset is sorted using a method selected from a bubble sort, a heapsort, or a mergesort.” But this claim form must be used very carefully, since if prior art is

¹⁷ The patent office charges for claims in excess of 20 in an application. Even at the small-entity rate of \$26, this would result in an application fee on the order of \$100,000. And the examiner, who doesn't receive any performance credit for the excess claims but has to address them, will hate you.

found that teaches any one of the alternatives, the entire Markush element is satisfied and the claim will be rejected in examination or invalidated in litigation.

It is far better to limit the number of claims by carefully deciding what elements are really necessary than by using tricks such as Markush claims to merge alternatives together.

And now, write the application

Armed with your claim menu and your knowledge of the invention, you're now ready to write the application, or at least the first draft of the specification to give to the patent attorney. All you have to do in the specification is fully describe the broadest claim and each and every thing on your claim menu. It may take some time, but it's a pretty mechanical process.

At the end of this note is an sample application for a simple invention, that can be used as a template.

I like to start by drawing a flow-chart showing the steps of the broadest claim.¹⁸ I'll use this to guide me as I start describing the invention. Then I write the background sections, describing what existed before this invention, the problem this invention is trying to solve, and a short summary of the invention.

This is not the place to write too much. Anything you say is in the prior art will be considered prior art by the examiner, whether it is or not. And if you mischaracterize the prior art, your patent could be invalid. So just say enough to put the invention in perspective.

Many patent attorneys like to include a section indicating the goals of the invention by stating a problem and then indicating how the invention solves that problem. *Don't*. It will be too easy if the patent is litigated for a judge or jury to confuse the goals of the invention with the invention itself, which is defined by the claims. They may think of the invention as anything that meets a goal in the same way, see prior art like that, and feel the patent is invalid. (How many times have you heard a patented invention characterized by its title, abstract, or even a shorthand description? The claims in infamous Amazon "one-click" patent, 5,960,411, include a lot more than just "ordering with a single click.")

Now, it's time to start describing everything in your claim menu. Start with the first step in the drawing corresponding to the broadest claim. When you get to a point where there are possible options, describe them all. When you are done describing the options, you go back to describing drawing of the broadest claim. ("Returning to Fig. X, step Y ..." is a good way to get back to it.)

It is important that you describe everything necessary to make and use the invention. Don't leave out details that may be helpful. But you are not trying to educate somebody unfamiliar with the technology of the invention. As previously mentioned, the patent statutes say who the audience of your teaching is – "any person skilled in the art to which [the invention] pertains, or with which it is most nearly connected."¹⁹ Who that is is often an aspect of patent litigation, but for software-based patents it often ends up being a

¹⁸ 35 U.S.C. § 113 requires drawings "where necessary for the understanding of the subject matter sought to be patented."

¹⁹ 35 U.S.C. § 112, first paragraph.

person with a bachelor's degree in computer science or engineering who has taken a courses relevant to the area of the invention (computer networking, databases, etc.) and may have a year or two of practical experience. But to be on the safe side, it's better to aim your discussion at a hypothetical person a little less knowledgeable.

But you don't need to describe things that are commonly known. And this is helpful when there are a variety of ways of doing something, not because it is particularly difficult but because how it is done is dictated by the particular environment for an implementation.

For example, the way of creating a pop-up window on a screen and filling it with a graphic and a text message differs substantially between Microsoft Windows and the X window system. But anybody skilled in programming using one of those window systems will know how to do it. You can get by with just saying that the pop-up window is created as part of a particular step, and then say something like "The mechanism for creating that pop-up window differs between window systems, but would be well-known to anybody skilled in the art of implementing systems that use a particular window system." One could even add a "See ..." and list some books that indicate how to create windows in different systems.

But one has to be a little careful here. Not everything can be incorporated in a patent specification by reference to a manual or textbook.

Essential material may be incorporated by reference, but only by way of an incorporation by reference to a U.S. patent or U.S. patent application publication, which patent or patent application publication does not itself incorporate such essential material by reference. "Essential material" is material that is necessary to:

(1) Provide a written description of the claimed invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same, and set forth the best mode contemplated by the inventor of carrying out the invention as required by the first paragraph of 35 U.S.C. 112;

(2) Describe the claimed invention in terms that particularly point out and distinctly claim the invention as required by the second paragraph of 35 U.S.C. 112; or

(3) Describe the structure, material, or acts that correspond to a claimed means or step for performing a specified function as required by the sixth paragraph of 35 U.S.C. 112.²⁰

Not exactly a bright-line test, so it is better to error on the safe side by describing rather than referencing. But again, if something is commonly found in textbooks for relevant courses, it's reasonable to say that it would be known (or, at least, known when they were taking the course) to people of ordinary skill in the art.

Whenever you think a figure will help, draw one. Not only flow charts showing the steps of various routines that elaborate on the steps in other flow charts, but examples of screen displays or drawings showing the important components of the computer system

²⁰ 37 CFR. 1.57(c).

running the method or diagrams of important data structures. Anything that helps you explain how the invention works and how to use it.

The drawings don't have to be beautifully drawn, especially for those for a software-based invention. (I do recommend professionally-drawn figures for mechanical inventions, where the precise relationship between components can be best shown in a figure.) You are allowed to submit final drawings before the patent issues, but they must just be cleaned-up versions of the original figures with nothing added.

As for numbering things on a figure, I like to start out the each number with the figure number – Fig. 1 has things numbered 101, 102, 103, etc. And if in a following figure, the same thing occurs (either a step or a component part), use the same number for it. This gives a good indicate where something is being introduced in the description.

The most important rule to remember when writing the description of the invention is *Always call something by the same name*. Don't use short forms of its name and definitely don't use synonyms. It's hard enough for somebody to accurately interpret the scope of a claim without having to determine if two things with slightly-different names are really the same thing. You should even avoid the use of a pronoun like "it" instead of using the full name for something.

For anything shown in a figure, follow its name with its reference number. This is easy for a step in a flow chart, where you might just refer to it as "Step 321." For components having a number in a figure, refer to it without any article – for example, "mark of the beast 666" instead of "the mark of the beast." And if you find yourself referring a number of times to a component that does not have a number or is not shown in a figure, it's time to add to the figures.

Unlike with the description of the prior art or the goals of the invention, this is no place to hold back. Even if your application runs over the 100 pages allowed in the basic filing fee, the cost for added pages is small – currently \$135 for each additional fifty pages for a small entity. And if you don't include something that is later determined to be necessary in your specification, you can't add it without jeopardizing your filing date, maybe even losing your chance to get a patent.

A provisional application?

An option first available in 2000 is the provisional application for a patent. This allows an inventor to secure a priority date for a nominal filing fee (currently \$110 for a small entity). The application is not examined, but can be later referenced in a regular patent application which must be filed within one year of the filing of the provisional application. If a claim is supported by the provisional application, it is accorded the filing date of the provisional application.

One use of a provisional application is to give an additional year after the publication of a scientific paper before filing a United States patent application. (Remember, you have a year after publication to file a patent application in the United States. But in many other countries, the publication of a paper before filing a patent application prevents getting a patent for what was described in the paper.) If a patent application has not been filed during the year after publication, the paper is turned into a provisional application (often by just attaching a cover sheet) and filed, thereby giving another year before the filing of the regular patent application.

But that should be done only if a regular patent application can't be filed during the year after publication. It is likely that the paper doesn't describe the invention in the detail that a regular patent application should, including the required description of the best mode contemplated for making and using the invention. It may not be clear in its terminology, forgetting the rule of always calling something by the same name and instead using slightly-different terms to make the paper easier to read. This may cause problems when the eventual patent is litigated and the alleged infringer is looking for possible defenses.

Instead, a provisional application should be used even if you could file a regular application, but would like time to determine whether it is worth the cost of getting a patent or to see how the market for the invention develops. You file the application you have just written, remembering that you have to review it within a year and file it as a regular application.

But, unlike a regular application, it is not necessary to file claims with a provisional application. *And you shouldn't.* It is likely that when you file the regular application based on the provisional application, you will have a better idea how to structure the claims, especially to catch possible infringers. If you had filed claims with the provisional application, the new claims might be considered amendments to those claims, and be foreclosed from using the doctrine of equivalents to reach possible infringers.

Back to the claims

So, now we're at the most important part of the patent again – the claims, which will determine whether something infringes the patent or not and therefore define its value. Remember, a patent is a right to exclude others from making, using, selling or offering to sell, or importing the claimed invention. (It does not give you the right to make the invention yourself, just sue others who might infringe the patent.) So it is important that the claims cover what others might do.

But we're in good shape to write the claims, because we already have figured out the broadest claim we can make given the prior art we know, the other elements that can be added to the broadest claim from our claim menu, and have a specification that fully describes all those things using the language of the claims.

The only thing we need to do is decide what we actually want to claim, since we won't claim every option, and put those claims together. Again, we'll confine our discussion primarily to method claims for software-based inventions.

If we could know all the prior art for a particular invention – impossible because it includes all printed publications in the world, regardless of language, and what everybody in the United States knows²¹ -- our job would be simple. We would just include the broadest claim that we developed to not read on the prior art. Instead, we have to be concerned with prior art that the patent examiner may find and, more importantly, both prior art that is discovered by a defendant in patent litigation and ways that a possible infringer may try to avoid the claims of the patent.

²¹ See 35 U.S.C. § 102(a) and (b).

Helping the examiner draw the line

The conventional wisdom is that you file your broadest claim as claim 1, then have the next couple claims be dependent claims adding additional limitations to that claim. For example, if claim 1 has elements A,B,C, then claim 2 would be claim 1 plus element D (giving A,B,C,D) and claim 3 would be claim 2 plus element E (giving A,B,C,D,E). (Remember the more limitations or elements in a claim, the narrower it is, because all elements have to be found in something for it to infringe and all elements have to be found in the prior art for the claim to be invalid.)

In essence, you are asking the examiner to draw a line regarding patentability after he completes his search. If he finds prior art that shows A,B,C but not D, he will reject claim 1 but only object to claim 2, indicating that it would be patentable if it were rewritten in independent form (in other words, as A,B,C,D, no longer based on claim 1). You now know his position and you can decide what you want to do. You can take a patent whose broadest claim is A,B,C,D. You can try to convince the examiner that he really didn't find prior art showing A,B,C – which is much easier if he found those elements in different prior art references. Or, through some tricky procedures that your patent attorney will know, you can do both by paying some additional application fees.

And that would be what I would have told you to do before June 2, 2004, when the Court of Appeals for the Federal Circuit, which hears all appeals from patent cases, issued an unfortunately-misguided decision.²² To understand the problem, you have to understand the “doctrine of equivalents” and “prosecution history estoppel.”

Under the “doctrine of equivalents,” a claim may be infringed even if it does not literally read on a thing. For example, if the claim is to A, B, and C, and the alleged-infringing thing has A, B, and D, there is no literal infringement. However, if D performs substantially the same function in substantially the same way to get substantially the same result as C, then there may be infringement under the doctrine of equivalents. The extension of a claim under the doctrine of equivalents is limited by “prosecution history estoppel” – if during the prosecution of the patent application the applicant has limited a claim to avoid prior art, it cannot be extended to what was previously disclaimed.

In other words, as a patent owner you like the doctrine of equivalents and don't want to see it limited by something you do that results in prosecution history estoppel. The problem is that the Federal Circuit has held that if you amend a claim, in most cases you will not be entitled to the doctrine of equivalents on the portions you amended. (This is why you didn't include any claims if you filed a provisional application.)

And that includes amending a dependent claim to make it an independent claim by including all the limitations of its parent claim. So, if you do what was discussed above, after you amend claim 2 from “claim 1 plus D” to “A,B,C,D,” even though its scope is the same, you will no longer be able to use the doctrine of equivalents for elements A, B, or C. But the solution is simple. Just file claims 2 and 3 as independent claims, rather than dependent claims. That way, you don't have to amend them if the examiner finds them acceptable. You get to have up to three independent claims for your basic filing fee, and additional independent claims only cost \$220 each (\$110 for small entities), a small price to pay for keeping the doctrine of equivalents alive for you.

²² *Honeywell International v. Hamilton Sundstrand*, 370 F.3d 1131, 71 USPQ2d 1065 (Fed. Cir. 2004).

Selecting from the menu

The second reason for additional, and generally dependent, claims is to provide a fall-back position in the event that a broader claim is declared invalid in later litigation. Unlike during the prosecution of the patent application where you can amend a claim to address prior art found by the examiner, in litigation the only option when a claim reads on prior art is to strike the claim as invalid.

But not every claim provides a fall-back. Consider the claims in the example patent application at the end of this note. Claims 2 and 3 are the fall-back claims, so that if, for example, prior art is found that use a single instruction invalidating claims 1 and 2, claim 3 will survive. But that is questionable, since it would be hard to show that using an iterative exclusive-OR wouldn't be obvious in light of using a single-instruction exclusive-OR.

Another common example, seen in lots of issued patents, is to have dependent claims that simply provide something specific for a general thing in the parent claim. For example, if the parent claim has “a computer network” as part of an element, the dependent claims may specify an Ethernet, a token ring, an Arcnet, and whatever other network types the claims drafter could think of. The idea is that if prior art is found that uses a token ring, then that dependent claim and then parent claim will be invalidated but the dependent claims specifying the Ethernet and the other types will still be there.

But by the very listing of all those network types, the patentee is saying that there is nothing particular to any of the network types when it comes to the invention. And that means that if it is known for a token ring, putting on an Ethernet is obvious.

If you haven't had to spend time in the specification explaining how to do something for a particular technology, don't bother using that technology to form a dependent claim in hopes that it will save your patent if related claims are invalidated. It won't.

So, when should we form dependent claims from our menu and include them as a litigation fall-back? First of all, when they add a new element, rather than just elaborate on an existing element. This is similar to why we include multiple claims for the examiner to consider in determining patentability. But we don't have to worry about losing the doctrine of equivalents if we amend a dependent claim to include its parent claim's elements, since we aren't going to be amending at this point. So we can use a dependent claim and save a little money.

Determine who would infringe

But the most important thing to keep in mind is who will infringe a claim, and how hard is it to avoid infringing the claim. After all, if you can't use a patent to stop competitors from using your technology or get them to license the technology, all you have gotten for your time and costs is a government certificate with a pretty ribbon.

We've been talking about claiming the software-based invention as a method. But the problem with using only method claims is that the infringer will be the end user, not your competitor who is taking your sales with their product using your technique.²³ While they may be indirectly liable if they are inducing others to infringe or contributing

²³ They would be an infringer to the extent that they use the invention themselves or during their testing, but that's likely minimal compared to the uses of their customers.

to that infringement, there are pitfalls in depending on secondary liability. For example, if their product has a substantial noninfringing use, they may not be a contributory infringer.²⁴

That's why we include claims to make competitors direct infringers, in particular article of manufacture claims so that making and selling compact discs containing software implementing the patented method is itself a direct infringement.

Another problem with method claims is that for them to be infringed, all steps of the method must be performed in the United States.²⁵ If it is possible for at least one of the steps to be performed in a foreign country, which high-speed data communications makes increasingly likely, then you need to carefully consider whether claiming that step is necessary to describe the invention and distinguish it from the prior art. If not, then do not include it in the claim.

If the step is necessary, one possibility might be to include it (along with all the other steps) in the preamble of the claim. So, instead of a claim that looks like:

A method comprising step A, step B, step C, and step D.

it would look like:

A method using the results of step A, step B, and step C,
comprising step D.

Arguably, only step D needs to be performed in the United States. This claim form goes against the normal advice not to include things in the preamble lest they be considered limitations to the claim. In this case, you want to use the preamble to be used to distinguish over the prior, just not have to be performed in the United States. Because the preamble here clearly gives meaning to the claims, rather than just provide an indication of an intended use, there is a strong argument that it should be considered a limitation on the claim when determining whether the claim is valid in light of the prior art. (There is no court case, to my knowledge, that has considered this claim form, since it's something that I thought up right after the Blackberry decision on the infringement of a method claim.)

In contrast to a method claim, a system claim may be infringed if the system is being supplied to an end-user in the United States. Since it is simple to include both article of manufacture and system claims in a patent application, perhaps only including some standard language and appropriate dependent claims, it would be foolish not to close well-known infringement loopholes.

It is important to keep in mind not only when and where claims are infringed, but by whom. For example, consider this claim that involves both a server and a client:

1. A method for negotiating a secure communications session, comprising:
 - (a) transmitting a request to a server;
 - (b) in response to the request, supplying from the server a server certificate, the server certificate including the server's public key;

²⁴ See 35 U.S.C. § 271(b) and(c). While it isn't common for physical devices to have a variety of uses (something isn't both a desert topping and floor polish, so that supplying it might not be contributory infringement of a floor polish claim), it isn't uncommon for a computer program to have a variety of features in addition to the one covered by a patent.

²⁵ See *NPT v. Research in Motion*, 418 F.3d 1282, 75 USPQ2d 1763 (Fed. Cir. 2005).

- (c) generating at the client a unique client key and communicating the unique client key to the server using the server's public key; and
- (d) thereafter communicating information using a crypto-algorithm that employs a derivative of the unique client key and the server's public key.²⁶

Who infringes the claim? Not the person running the server, because he doesn't perform the client steps. And not the person using the client, because he doesn't perform the server steps. You might be lucky to have an infringer that both runs the server and uses the client, or you might try arguing inducement on the part of one party or the other, but don't count on a finding of infringement.

In contrast, consider these two claims based on the claim above:

- 2. A method for negotiating a secure communications session, comprising:
 - (a) receiving a request from a client;
 - (b) in response to the request, supplying from the server a server certificate, the server certificate including the server's public key;
 - (c) receiving from the client a unique client key communicated using the server's public key; and
 - (d) thereafter communicating information using a crypto-algorithm that employs a derivative of the unique client key and the server's public key.
- 3. A method for negotiating a secure communications session, comprising:
 - (a) transmitting a request to a server;
 - (b) receiving from the server a server certificate, the server certificate including the server's public key;
 - (c) generating a unique client key and communicating the unique client key to the server using the server's public key; and
 - (d) thereafter communicating information using a crypto-algorithm that employs a derivative of the unique client key and the server's public key.

Claim 2 is directly infringed by the person running the server, and claim 3 is directly infringed by user of the client.

Another thing to keep in mind is how you will be able to determine that somebody is infringing your patent. To the extent that the claims involve a particular external data structure, or acts that can be readily observed, it makes it considerably easier to form the good faith belief necessary to file an infringement suit and avoid sanctions. (You will be able to confirm the infringement when you have a chance to examine the alleged infringing program's source code during later discovery.) But if whether the claimed method is being used can only be determined through reverse engineering, it will be far more difficult to determine if there is infringement before filing suit.

Process, machine, article of manufacture?

In 2007, when I first wrote this note, my advice on how to claim was pretty simple: claim as a method, and use dependent claims to cover the machine and article of

²⁶ From Lemley, et al., *Divided Infringement Claims*, 33 AIPLA Quarterly Journal 255, Summer 2005.

manufacture claims that reach other infringers. So, for example, there would be a triplet of claims such as:

1. A method operating on a digital computer system comprising:
[step A], ... , [step Z].
2. A digital computer system programmed to perform the method of claim 1.
3. A computer-readable medium storing a computer program implementing the method of claim 1.

That changed on October 30, 2008, when the Court of Appeals for the Federal Circuit issued its opinion in *Bilski*.²⁷ While *Bilski* was about the claiming of a business method that did not involve a computer,²⁸ when the Federal Circuit said that Bilski's method was not patentable, it stated a rule on when a process or method was patent-eligible ("statutory subject matter") that also affects computer-implemented methods.

The new rule is that a process is patent-eligible if it "transforms a particular article into a different thing" (a classic process) or "is tied to a particular machine or apparatus." The opinion put off for later what makes a machine "particular," including "whether or when recitation of a computer suffices to tie a process claim to a particular machine." The case is now before the Supreme Court, which will hear oral arguments on November 9, 2009, and issue an opinion sometime in 2010. The Court may clarify how to claim software-based techniques, or say the boundary is whether the technique is "abstract" or "technological" as some have urged, creating even more uncertainty as new court cases try to define those terms.

The way of addressing this uncertainty should be simple: claim the invention as a machine, which avoids *Bilski*'s uncertainty about when a computer-based process is patentable. Instead of:

A method operating on a digital computer system comprising:
[step A], ... , [step Z].

the claim would look like:

A digital computer system programmed to perform the method
comprising: [step A], ... , [step Z].

There is a slight difference between these claims, in that the first one is infringed when the method is actually performed, while the second is infringed when the program implementing the method is loaded into the memory of the computer, thereby "making"

²⁷ *In re Bernard L. Bilski and Rand A. Warsaw*, 545 F.3d 943, 88 USPQ2d 1385 (Fed. Cir. 2008).

²⁸ The claim at issue read:

A method for managing the consumption risk costs of a commodity sold by a commodity provider at a fixed price comprising the steps of:

- (a) initiating a series of transactions between said commodity provider and consumers of said commodity wherein said consumers purchase said commodity at a fixed rate based upon historical averages, said fixed rate corresponding to a risk position of said consumer;
- (b) identifying market participants for said commodity having a counter-risk position to said consumers; and
- (c) initiating a series of transactions between said commodity provider and said market participants at a second fixed rate such that said series of market participant transactions balances the risk position of said series of consumer transactions.

the claimed invention. But both are directed at the end user, with the company actually infringing the method brought in as an inducer or contributory infringer,²⁹ or covered by an article of manufacture claim such as:

A computer-readable medium storing instructions that program the digital computer system to perform the method of claim 1.

Unfortunately, there has been a tendency for courts and the patent office to consider a machine claim characterized by performing a method as a claim to the method itself, and apply the *Bilski* test.

The patent office has provided some guidance, but since it has to follow *Bilski*, it also doesn't provide a clear line.

For computer implemented processes, the "machine" is often disclosed as a general purpose computer. In these cases, the general purpose computer may be sufficiently "particular" when programmed to perform the process steps. Such programming creates a new machine because a general purpose computer, in effect, becomes a special purpose computer once it is programmed to perform particular functions pursuant to instructions from program software. To qualify as a particular machine under the test, the claim must clearly convey that the computer is programmed to perform the steps of the method because such programming, in effect, creates a special purpose computer limited to the use of the particularly claimed combination of elements (i.e., the programmed instructions) performing the particularly claimed combination of functions. If the claim is so abstract and sweeping that performing the process as claimed would cover substantially all practical applications of a judicial exception, such as a mathematical algorithm, the claim would not satisfy the test as the machine would not be sufficiently particular.

This suggests two possible avenues for claiming a computer-implemented method. First, if there is some special characteristic of the system used to implement the method, such as two computers communicating over a network, make that clear in the claim (and support it in the specification) so that it becomes what is "particular" about the system.

Second, make it clear in the specification how the computer is programmed to perform each step of the method. Don't, for example, just say that the step is to sort an array, but indicate how that sorting would be performed. Otherwise, even if your claim passes as statutory subject matter, it could be rejected because the specification doesn't teach how to program the system to give the special purpose computer.

Conclusion

Forgetting for a moment the *Bilski* problem, it's really not that hard to draft a patent application for a software-based invention *if* you know the invention and the prior art. And as the programmer who created the invention, you are in the best position to know that. So, if you want a good patent, you should plan on playing a major role.

But even if you have some experience with patents, you should include a patent attorney or patent agent as part of your team for drafting and prosecuting the application.

²⁹ See 35 U.S.C. § 271(b) and (c).

Besides providing a different perspective on the application, they will know how to respond to the patent office to get the best patent.

Method and System for Interchanging Stored Data

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims benefit of Provisional Patent Application Ser. No. 60/123,456, filed on March 9, 2009, by the present inventor.

BACKGROUND – FIELD OF INVENTION

This invention relates to data storage and manipulation techniques, and in particular to the interchanging of the data stored in two locations in a computer memory.

BACKGROUND – PRIOR ART

There are many times when it is necessary to interchange the data stored in two different locations in a computer memory. For example, it is often the case in sorting techniques that two elements in the list to be sorted need to be interchanged.

Fig. 1 illustrates a common technique in the prior art for interchanging the contents of two storage locations A and B. First, the contents of A is copied to a temporary location. Then, the contents of B are copied to A. Finally, the contents of the temporary location are copied to A.

However, if storage locations A and B are large, and there is limited memory capacity on the computer, it may not be possible to allocate sufficient memory for the temporary location.

SUMMARY OF THE INVENTION

The present invention interchanges the data stored in two different locations in a computer memory without requiring temporary storage. It does this by using an exclusive-OR Boolean operation on the data three times, resulting in the interchange of the data.

These and other features of the invention will be more readily understood upon consideration of the attached drawings and of the following detailed description of those drawings and the presently preferred embodiments of the invention.

DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a prior art method for interchanging data.

Fig. 2 illustrates the method of this invention for interchanging data.

DETAILED DESCRIPTION OF THE INVENTION

Fig. 2 illustrates the three steps fundamental steps of the method of the present invention. In the figure, the operator “xor” represents the exclusive-OR of its two operands. The exclusive-OR differs from the more-common inclusive-OR (or often, simply “OR”) operator in that when both its operands are 1, its result is a 0.

The following table defines the xor operator for a single bit of storage:

X	Y	X xor Y
0	0	0
0	1	1
1	0	1
1	1	0

For storage locations larger than a single bit, the result of the exclusive-OR of those storage locations is the exclusive-OR of each bit in one location with the corresponding bit in the other location. Depending on the size of the two storage locations (which must be the same size, since they are being interchanged) and the instructions available on the particular machine running the present invention, this can be done by a single computer instruction or by iteratively-performing the exclusive-OR on corresponding segments of the two storage locations.

Returning to Fig. 2, in step 201 the contents of storage locations A and B are exclusive-ORed and the result is stored in location A. In step 202, the contents of storage locations A (which contains the result of the exclusive-OR operation in step 201) and B are exclusive-ORed and the result is stored in location B. Finally, in step 203, the contents of storage locations A and B (which contains the result of the exclusive-OR operation in step 202) are exclusive-ORed and the result is stored in location A.

The following table illustrates the operation of the method of the present invention, with the original contents of A and B selected so that all four possible combinations of corresponding bits are shown.

	A	B
Original contents	0011	0101
A xor B -> A	0110	0101
B xor A -> B	0110	0011
A xor B -> A	0101	0011

Computer programs implementing the method of this invention, particularly for the central data server, will commonly be distributed to users on a distribution medium such as floppy disk or CD-ROM. From there, they will often be copied to a hard disk or a similar intermediate storage medium. When the programs are to be run, they will be loaded either from their distribution medium or their intermediate storage medium into the execution memory of the computer, configuring the computer to act in accordance with the method of this invention. All these operations are well-known to those skilled in the art of computer systems.

The term “computer-readable medium” encompasses distribution media, intermediate storage media, execution memory of a computer, and any other medium or device capable of storing for later reading by a computer a computer program implementing the method of this invention.

It is to be understood that the above described embodiments are merely illustrative of numerous and varied other embodiments which may constitute applications of the principles of the invention. Such other embodiments may be readily devised by those skilled in the art without departing from the spirit or scope of this invention and it is our intent they be deemed within the scope of our invention.

CLAIMS

I claim:

1. A method for interchanging the contents of storage locations A and B in a computer memory, comprising:
 - replacing the contents of A with the exclusive-OR of the contents of A and the contents of B, and then
 - replacing the contents of B with the exclusive-OR of the contents of B and the contents of A, and then
 - replacing the contents of A with the exclusive-OR of the contents of A and the contents of B.
2. The method of claim 1, where the exclusive-OR is performed by a single computer instruction.
3. The method of claim 1, where the exclusive-OR is performed iteratively.
4. A digital computer system programmed to perform a method comprising:
 - replacing the contents of A with the exclusive-OR of the contents of A and the contents of B, and then
 - replacing the contents of B with the exclusive-OR of the contents of B and the contents of A, and then
 - replacing the contents of A with the exclusive-OR of the contents of A and the contents of B.
5. The digital computer system of claim 4, where the exclusive-OR is performed by a single computer instruction.
6. The digital computer system of claim 4, where the exclusive-OR is performed iteratively
7. A computer-readable medium storing a computer program implementing a method comprising:
 - replacing the contents of A with the exclusive-OR of the contents of A and the contents of B, and then
 - replacing the contents of B with the exclusive-OR of the contents of B and the contents of A, and then
 - replacing the contents of A with the exclusive-OR of the contents of A and the contents of B.
8. The computer-readable medium of claim 4, where the exclusive-OR is performed by a single computer instruction.
9. The computer-readable medium of claim 4, where the exclusive-OR is performed iteratively.

ABSTRACT

A method and system for interchanging the data stored in two locations in the memory of a digital computer without requiring the use of a temporary storage location. The technique involves using an exclusive-OR operation three times, resulting in the interchanging of the data.

Fig. 1 (Prior Art)




Fig. 2

