Module 101

[The following is a transcript of the recorded lecture for Module 101 of the PatentX course. The recording of the lecture itself is available through https://ipxcourses.org. Stripped of the accompanying slides and other visual materials, the transcript will likely be hard to follow. It is not intended to be a free-standing document. Rather, its purpose is to assist students, who have already watched the lecture, to review the material.]

A. Obtaining a Patent

Hello, I'm Terry Fisher. This is the first in a series of lectures on Patent Law. This particular lecture is meant to provide you an overview of the patent system as a whole.

It has three parts. The first describes how an inventor goes about obtaining a patent.

The second describes how patents are exploited commercially.

The third discusses how patents are challenged and sometimes invalidated.

From this summary, you can see that I will be concentrating heavily on the procedural rules and customs that underlie the patent system. Most of the other lectures in the series will focus instead on the substantive rules of patent law. The reason why we are beginning with procedure is that it's difficult to understand either the substantive rules of patent law or the judicial decisions interpreting the substantive rules without a basic understanding of the procedural contexts in which those rules become relevant.

So let's begin.

It's easiest to get a sense of how the system works by example. Here's a relatively simple one:

In most temperate countries, including the United States, mice are ubiquitous pests. Some people regard them as cute, but most people don't want them in their kitchens, basements, restaurants, and so forth.

For centuries, inventors have been hunting for ways to catch and kill mice efficiently. A wide variety of devices have been invented and marketed. Those that have stood the test of time include the classic spring-loaded trap (traditionally made of wood and metal; nowadays more often made of plastic); glue traps (which grip the mice who walk onto their sticky surfaces until they die); and electric traps (which electrocute the mice who enter them).

Each has advantages and disadvantages as forms of pest control. None is perfect. The result is that inventors continue to try to develop better mousetraps.

One recent invention sought to improve the standard glue trap by encasing it in a 5-sided disposable bag. This innovation did nothing to enable the trap to catch more mice; rather, it made the traps easier and less unpleasant to use. The inventor was Askew Gatewood, Jr. He appears to have developed the trap in 1995.

In February 1996, with the aid of an attorney, he applied for a US patent on his invention.

To understand his application, you need to know a bit about the relevant terminology.

. In the United States and in most countries, a patent application typically contains the following things: :

1. A title.
2. An identification of the inventor or inventors.
3. Cross-references to related inventions.
4. A statement indicating that the invention grew out of governmentally sponsored research, if true.
5. A description of the background of the invention.
6. A summary of the invention.
7. A brief description of the drawings used to illustrate the invention.
8. The drawings themselves
9. A detailed description of the invention.
10. A set of claims – the most important and usually most controversial part of the application.
11. An abstract of the disclosure.
12. Finally, an oath or declaration, attesting to the applicant's belief that she is the inventor

In the US, if an inventor is not yet able to provide information of all of these sorts, she can submit a provisional application, which contains a much more limited set of materials. Provisional applications are essentially placeholders. So long as they contain enough information to describe the invention effectively, they establish the applicant's priority date (which, as we will see, is crucial in fending off challengers) – and give him or her an additional year to file a regular application.

Gatewood apparently did not need to avail himself of this option. On January 31, 1996, he filed (with the aid of an attorney) a regular patent application for his improved mousetrap.

So what happens when an application of this sort is submitted to the US Patent Office?

Here's a diagram outlining the principal paths.

The first step after the filing of the application is that it is assigned to one of the roughly15,000 examiners employed by the office.

Because of the substantial backlog of applications awaiting attention, the applicant is not likely to receive a substantive response from the examiner for over a year. On occasion, the response, when it comes, is entirely favorable.

If the examiner concludes that the application complies with all of the relevant procedural and substantive requirements, he will issue a "notice of allowance." By filing another form and paying another fee, the applicant can then obtain a patent covering all of the ground she requested.

It is more likely, however, that some or all of the claims in the original application will be rejected –

most often because, in the judgment of the examiner, they seek protection for material that either is not novel or is obvious in light of the prior art. (The meanings of those objections and other grounds on which the examiner may reject the application will be considered in detail in the fifth lecture in this series.)

If that happens, the applicant has three main options. First, she can give up. Formally, this is called abandonment; it happens automatically if the applicant does not respond within 6 months.

Second, she or her lawyer can push back – in other words, try to persuade the examiner that the original application should be accepted without any changes.

Third, she or her lawyer can amend the initial application in an effort to meet the examiner's objections. Usually this is done by narrowing some or all of the claims.

The typical response is a combination of the second and third of these options – in other words, defense of a portion of the original application combined with amendment of another portion.

The ball is now back in the examiner's court. If the examiner once again rejects the application, whether amended or not, he will likely label the rejection "final." This is a misnomer. There are various ways in which the applicant can obtain additional rounds of examination.

For example, the applicant can file what is known as a "request for continued application."

Or the applicant can file a new application, known (confusingly) as a "continuing application," which carries forward the information contained in the original application and typically adds new material.

The requirements associated with these various options are important to patent practitioners, but not particularly important for our purposes, so we will put them to one side.

Most often, at the end of this process, the examiner will allow a pruned version of the original application – and the applicant will be able to obtain a patent, though not as broad as the one she originally sought.

If the examiner persists in rejecting even the amended application, the applicant is not completely out of luck.

She can challenge the examiner's decision before an administrative tribunal within the Patent Office. That tribunal is known today as the Patent Trial and Appeal Board, commonly abbreviated PTAB.

Prior to the 2011 America Invents Act, it was known as the Board of Patent Appeals and Interferences (commonly abbreviated BPAI) – and that acronym will appear in some of the judicial opinions you may be reading in connection with these lectures. But, for simplicity, we will be using the abbreviation PTAB most of the time.

If the PTAB sides with the applicant, it will instruct the examiner to issue a notice of allowance.

If not, the applicant can challenge the PTAB's ruling in one of two ways.

First, she can appeal the PTAB's ruling to the Court of Appeals for the Federal Circuit – commonly abbreviated as CAFC.

Second, she can bring suit against the Director of the PTO in a federal district court (specifically, in the District Court for the Eastern District of Virginia);

If that court rules against her, she can appeal to the Federal Circuit.

Most disgruntled applicants take the direct path to the CAFC, but some take the lower road – typically because in a challenge before the district court, the applicant can introduce new evidence, whereas in the direct appeal to the Federal Circuit, she cannot.

Now, a word about these two tribunals. The Federal Circuit is an unusual animal within the US judicial system. For the past 30 years, it has had exclusive appellate jurisdiction over disputes the outcome of which hinge upon the interpretation or validity of patents. It was created in 1986, partly to harmonize the disparate interpretations of provisions of the patent statute that the various courts of appeals had previously adopted – and partly to make obtaining and enforcing patents easier.

Here's the current composition of the court. As you can see, there are 12 regular judges and 7 senior judges. The latter typically have reduced case loads. Interestingly, most of the regular judges were appointed by President Obama, none by President Trump, one so far by President Biden.

Unlike most federal judges, the judges on the Federal Circuit are specialists. The large majority of the cases they hear involve patent law – so, as you can imagine, they know patent law well. In addition, several have deep backgrounds in some field of science or technology, which enables them to understand patents more easily.

Typically, the judges hear appeals in three-judge panels, but occasionally, in especially important cases, all of the judges sit altogether and issue opinions “en banc.”

Whether, on balance, the concentration of appeals in patent cases in the Federal Circuit has been good or bad for the patent system is a complex question we will take up later, when you’ve seen the court in action in various contexts.

If, through one or another of these routes, the Federal Circuit upholds rejection of a patent application,

the applicant can ask the Supreme Court to review the ruling, but it’s very rare for the Supreme Court to grant such requests.

Back to the mousetrap:

Gatewood’s application proceeded in the typical fashion.

He filed it, as I’ve indicated, in January of 1996.

It was assigned to an examiner, Jeanne Elpel.

She conducted what’s known as a prior art search -- and found several inventions that she deemed relevant – most of them embodied in previous US patents or patent applications.

On January 28, 1997, Elpel rejected the application on several grounds:

--that some of the language used in the application was excessively vague or informal;

--that the term, “traversing,” that appeared in both the written description and the claims was inaccurate;

--that 7 of the 10 claims were too imprecise;

--and that two of the 10 claims were obvious because they represented only modest extensions of a previously patented bakery bag which also contained a drawstring.

On February 3, 1997, Gatewood’s attorney filed an amended application, adjusting the language to meet Elpel’s objections and cancelling the claims that Elpel had deemed obvious.

Elpel apparently was satisfied by the changes.

On February 28, she issued a notice of allowance.

Gatewood paid a fee of $645.

And on October 7, 1997, the patent was granted.

I'll first display, briefly, the patent in its entirety – and then show you how to identify and interpret the most important parts.

Gatewood's patent, as issued, had four pages. Here they are.

If you'd like to examine the entire thing at leisure, use a computer to locate the Google Patent Search Index and type in the number 5673509.

When reading this patent – or any other patent – you should be especially watchful for the following:

The identification of the inventor or inventors

The title

The abstract

The background section

The summary of the invention

The description of the drawings

The detailed description of the invention – which may include an identification of the version of the invention that the inventor considers the best -- and may include some other examples

And last, but not least, the claims, which identify what the applicant asserts a right to.

Typically, each of the claims will have the following form: it will start with a preamble, include a transition, and then contain the body.

The transitions matter more than you might think. Three are used most commonly; here's what they mean.

- . The term "Comprising" means includes
- So the phase "comprising X & Y" means includes X and Y
- A claim in this form is known as an "open" claim
- So a claim over a gizmo "comprising" X and Y will "read on" or cover not just a gizmo that contains X and Y, but also a gizmo that contains X & Y & Z, regardless of how important Z is
- The term "Consisting essentially of" is less capacious

- So the phrase, “consisting essentially of X & Y” means includes X and Y and anything else that does not change the essential function or properties of the invention
- A claim in this form will read on X & Y & Z iff Z is unimportant to the function of the invention
- The term “Consisting” is narrower still
- So the phrase, “consisting of X and Y” means contains X & Y and nothing more
- This is known as a “closed” claim
- It will not read on a gizmo containing X & Y & Z

For obvious reasons, patent applicants prefer open to closed claims and thus use the term “comprising” in their transitions if the examiners will allow them.

These sections of the patent are known collectively as the specification. As we’ll see in the third lecture in this series, the relationship between the specification and the claims is important.

Here, then, are the aspects of Gatewood’s patent that correspond to these categories.

The inventor

The Title

The Abstract

The Background section

The Summary

The Description of Drawings and the associated drawings

The Detailed Description – the longest part of the patent

The Description of the Preferred Embodiment

And six claims.

As is typical, the first of the claims is the most general and thus the most important.

Here is its preamble.

The transition

And the body.

The body can be – and typically is – broken down into elements – about which I'll have more to say later.

Here are the elements of claim 1

Claim 2 adds an additional element and is thus narrower.

Claim 3 adds yet another.

And so forth. Claim #6, as you can see, contains 10 elements. To infringe this claim, a competitor's mousetrap would have to contain all 10 elements. For that reason, claim 6 is the narrowest and thus the least valuable for Gatewood.

That concludes our review of the mechanisms by which Gatewood could – and did – obtain a patent in the United States. Suppose that he thought or hoped that his invention would prove attractive not just to consumers in the United States, but to consumers in other countries. To protect himself against competitors in those countries, he would have to obtain a patent in each one.

You may well be thinking: the 21st century economy is global. Shouldn't there be a way to obtain a single global patent enforceable in all countries? Perhaps, but there is not – and there is not likely to be for the forseeable future. Instead, all patents are issued by individual countries or, in a few instances, by regional organizations. With rare exceptions, patents are only effective in controlling conduct within the jurisdictions of the nations or regional organizations that issued them. So an inventor who, like Gatewood, who wants patent protection in other countries or regions must apply for and obtain patents in those other countries or regions.

How many jurisdictions would he have to reach? Well, there are approximately 197 sovereign countries in the world. I say approximately, because the autonomy of some is contested. Of that number, roughly 188 have active patent systems. Again, that number is an approximation, because the activity of a few of the systems is hard to verify. In all of those 188 countries, both residents of the country and foreigners can apply for patents. Thus, in theory, Gatewood could seek patent protection is every one.

As I mentioned a minute ago, the extraordinary cost of doing so is mitigated to some degree by the existence of a few regional organizations that in various ways facilitate the acquisition of patent protection in multiple countries. Economically, the most important of those regional organizations is the European Patent Convention.

As you can see, its membership includes -- but is broader than -- the membership of the European Union. By submitting a single application to the European Patent Office, located in Munich, an inventor can obtain a patent enforceable in all 39 of the member states. I hasten to add that the manner in which that patent is interpreted and enforced in the various member countries unfortunately is not yet uniform. Later in this lecture series, we'll discuss the

impending activation of the so-called Unified Patent Court, which will enable applicants to opt for an enforcement system that will apply uniformly throughout the zone covered by the EPC. But for now, the fact that European patents are enforced by national courts using nation-specific procedures produces substantial variation in their scope.

For the most part, the procedures used by the national and regional patent offices resemble those used by the PTO in the United States. But there are some differences. Here are a few of them.

The offices vary considerably in the speed with which they process applications. I mentioned earlier that, in the US, over a year typically elapses between the date an application is first filed and the first substantive response – the so-called first office action. But thereafter, the process tends to move fairly quickly. On average, the final decision from the PTO comes less than 2 years after the date of the original application.

The pace of the European Patent Office is similar.

The offices of some countries, such as Brazil and India, are much slower.

while the other members of the "Big 5" – China, Korea, and Japan – are faster, on average.

The offices differ even more in the frequency with which they respond positively to applications. Now, for the reasons indicated in the note at the bottom of this page, the data shown on this slide should be taken with a grain or two of salt. But nevertheless, it's apparent that, in some jurisdictions, acceptances are more common than in other jurisdictions.

The offices also vary considerably in terms of where their applications originate. As you can see, the Chinese Patent Office receives by far the largest number of applications – over one and a half million per year. 90% of those applications come from Chinese residents. The corresponding percentages are similar at the national offices of France and Italy.

By contrast, over 80% of the applications to the patent offices of Canada, Australia, Brazil, Hong Kong, Mexico, Singapore, South Africa, Israel, and Indonesia come from nonresidents.

The US and the European Patent Office are in the middle of this spectrum. In both, 55% of the applications are nonresidents.

In most countries, as you can see, the percentage of patent grants that go to nonresidents is roughly the same as the percentage of applications from nonresidents – suggesting that the offices are not discriminating in favor of their own residents.

In most low and middle-income countries, the total number of applications received by the national patent office is lower – less than 8000 in contrast to the 1.5 million received by China. In most, the percentage filed by nonresidents is extremely high. Typically, the bulk of the applicants consist of foreign corporations.

But there are exceptions. In Turkey, Bulgaria, Moldova, and Albania, for example, nonresident applicants are the minority.

Ok, that's preliminary canvas of the variety of patent jurisdictions in the world. Fortunately, there exists one mechanism, close to global in its coverage, that helps inventors secure patents in multiple jurisdictions. That's the Patent Cooperation Treaty – commonly known as the PCT.

The PCT first became effective in 1970; It has now been signed by 157 of the 188 countries that have active patent systems.

The PCT is managed by the World Intellectual Property Organization, which has collaborated in the distribution of this lecture series and the courses it supports. The executive within WIPO who oversees the PCT is Matt Bryan. He has kindly developed a short lecture explaining in detail how the system works. After you finish this portion of the lecture, you should watch Bryan's lecture.

Mr. Gatewood, like many applicants, apparently decided that the PCT route to patent protection in other countries was worth the cost.

He filed his PCT application with the US PTO in July of 1997, requesting protection in Canada, Japan, and designated member countries of the European Patent Convention

For the most part, his PCT application was just a copy of his US patent application.

At Gatewood's request, the USPTO prepared an international search report, which listed only five prior US patents as relevant prior art. WIPO published the PCT application and the associated report, then forwarded them to the three offices.

Five years after the date of Gatewood's original application, the European Patent Office notified him of its decision to grant a European patent. Its fate in Japan and Canada is unclear.

As we've seen, by using these various tools, Gatewood was able to secure a patent in the United States and in Europe. How confident should we be that he truly deserved those patents – not in a moral sense, but in the sense that he has actually complied with all of the substantive requirements for obtaining patent protection? The answer, unfortunately, is not very confident. A variety of circumstances, in combination, cause many patents to issue when they should not be. Those circumstances include:

--patent examiners are underpaid and thus sometimes underqualified;

--they are overworked – and thus unable to spend as much time examining each application as they would like. (In the US, the average patent application gets a total of roughly 19 hours of an examiner's time.)

--Finally, incentives within some patent offices encourage examiners to grant patents rather than reject them.

Most commentators lament this state of affairs. But not all. Professor Mark Lemley, most notably, has argued that it makes sense for patent offices to devote relatively little time to each application, leaving to the courts the task of scrutinizing them more carefully when they are tested in litigation. Why waste time poring over an application that will never have any commercial use?

Lemley is certainly correct that the large majority of granted patents have no commercial life at all. Indeed, Gatewood's patents fall into this category. As best I can tell, his improved mousetrap was never commercially produced – either because it wasn't as much of an improvement as he asserted or because he was unable to get a purchase in the market. As a result, he let the US patent expire -- merely by ceasing to pay the required maintenance fee. But whether Lemley's more general contention that applications like Gatewood's don't deserve much attention is convincing is less clear.

Keep his contention in mind during the second and third segments of this lecture when I discuss the ways in which patents, once granted, can be asserted and challenged

B. Exploiting a Patent

A patent gives its owner the right to prevent anyone else from making, using or selling a product or process covered by its claims within the jurisdiction where the patent was issued.

Much more detail concerning the scope of that right will be provided in the sixth lecture in this series – but the rough contours of the right are probably clear enough.

It is commonly assumed – both by laypersons and by lawyers – that the most commercially effective way of exercising such a right is to prevent all other persons and firms from competing against the patentee, enabling the patentee to sell devices or services embodying the technology at issue for prices higher than would be possible in a competitive market.

Used in this way, a patent is a potential source of market power. Whether it's a source of actual market power depends on whether here exist in the market adequate substitutes for the technology it covers. If so, the patent gives the patentee the ability to suppress competition. It is generally assumed that, during the limited life span of the patent, the patentee can and should exploit that power to its fullest.

Here's an illustration of this strategy:

The human knee is a complex joint. For the most part, it works remarkably well. Unfortunately, it has a weakness: the anterior cruciate ligament, commonly known as the ACL. When placed under unusual stress, the ACL can tear, after which the knee will not function properly.

Several sports place participants under such stresses and thus give rise to many ACL tears.

For example, if a skier falls backward and twists – and his binding does not release his foot from the ski, his ACL is likely to tear. Expert skiers suffer this injury at least as much as beginners. For example, a study by the International Ski Federation found that, out of every 100 top racers in the world, 5 suffered an ACL injury in training or competition each year.

Football – both US-style football and the radically different form of football played in other countries – also lead each year to many ACL tears. For some reason, girls and women are more at risk than boys and men.

Viewed coldly, the large and steady stream of ACL tears represents a potential market. Most athletes who suffer this injury want to continue to participate in their sports. Surgery can restore most function to a damaged knee, but rebuilt knees are not as strong as uninjured knees. Accordingly, many athletes who have suffered this injury wear knees braces.

A wide variety of knee braces are available for this purpose.

Here are some examples currently on sale through Amazon.com.

As you can see, most are relatively inexpensive.

One is not.

Donjoy braces, made by DJ Ortho, are much more costly. Why? Partly because, in the judgment of both consumers and experts, they are much better.

That's a necessary condition for DJ Ortho to charge so much – but not sufficient. Donjoy braces have been around for quite a while – plenty of time for DJ Ortho's competitors to have reverse engineered their structure, manufactured functionally identical braces, and sold them for lower prices (while still earning a generous profit). That would have forced DJ Ortho to have reduced its prices – or forfeited its market.

The reason this has not occurred is, as you might imagine, a patent held by DJ Ortho on the key features of its brace.

Here's the abstract of the relevant patent and one of the drawings. (You now know what these things are.)

As you can see, DJ Ortho filed the patent application in 2001,

and it was granted in 2003. Because patents in the US and almost all countries last from 20 years from the filing date, the patent will not expire until 2021.

Until that time, DJ Ortho will not face any competition in the manufacture and sale of braces that have the featured claimed in the patent. Because, as we have seen, many consumers place a high value on those features, it can behave like a monopolist.

How does a monopolist behave? What follows is a highly simplified summary. In subsequent lectures, I will provide some refinements of this general account – but this will suffice for now.

For the purpose of this summary, I'll be using a basic graph. The x axis shows quantity; the Y axis shows money – for convenience, we'll use US dollars.

If it wishes to capitalize on its patent to engage in monopoly pricing, Ortho will first figure out how much it costs to manufacture each brace. It certainly won't set the price below that amount. Let's assume, for simplicity, that the cost falls here – and does not vary with quantity – i.e., that the company will not benefit from economies of scale and will not face rising costs because of the scarcity of supplies.

DJ Ortho next tries to figure out what the demand curve for its product looks like – in other words, for each potential price, how many consumers each year would buy it. A simple demand curve might look like this:

If Ortho set the price for the brace way up here at point a – this many consumers would buy it. These folks are likely to be especially wealthy, or avid athletes, or have excellent health insurance that would cover the cost of a fancy brace. If it set the price a bit lower at point b, more people would buy it. Lower still and even more would buy it.

The line connecting these and other dots determined in the same fashion is known as a demand curve.

So where should Ortho set the price?

As mentioned previously, if Ortho did not have a patent, its freedom in setting the price of its braces would be constrained by the threat of competition. If it set the price up here and sold it to this many consumers, it could anticipate that competitors would copy its brace and sell them for a lower price (to more customers). To meet this challenge, Ortho would have to lower its price, which would prompt the competitors to drop theirs. Eventually, Ortho's price for these wonderful braces would be driven down close to the marginal cost of making them. The result is that a huge number of consumers with bad knees would be able to buy the braces but Ortho would not make any money,

The patent frees Ortho from this threat. It is now able to select the price that will maximize its profits.

Picking that price is a bit more difficult that might first appear, because Ortho must consider not just its profit margin on each brace, but also the total number it would sell. On the highly simplified assumptions embodied in this graph, the profit-maximizing price is here – which will give rise to a quantity sold causing the marginal cost and marginal revenue curves to cross – and provide Ortho a generous profit.

More realistic assumptions would make Ortho's calculus more complex. But we can be sure that the profit maximizing-price would be well above the cost of making the braces.

Anecdotal evidence in support of that prediction can be obtained by comparing the prices of braces not shielded by patents, with the price of Ortho's brace.

Latent in this simple diagram is the most basic – but also most common – of the conventional justifications of the patent system as a whole – generally known as "reward theory." According to this theory, the reason why we permit – indeed, help – inventors to engage in monopoly pricing of this sort (so the argument goes) is that it provides actual and potential inventors incentives to engage in inventive activity in the first place. The lure of monopoly profits, in other words, causes people and companies to engage to socially beneficial activity.

To be sure, this way of providing incentives for innovation has a downside: it means that consumers unable or unwilling to pay the monopoly price of the product (roughly speaking, these folks) are deprived of access to it. They are thus worse off than they would be if the product were competitively priced,

To return to our example, the folks represented by the red line will have to make do with less effective knee braces – which will cause them discomfort and may lead to more injuries. From the standpoint of social welfare, that's bad. Economists sometimes refer to it as "deadweight loss" – crudely represented by this area of the graph. But we are willing to tolerate this loss – at least for the duration of the patent – in order to induce more inventive activity than otherwise would occur.

Now, I hasten to add that there are several other ways of seeking to justify the patent system – which have different implications for how the law ought to be shaped. We will consider some of them in due course. But, for better or worse, the reward theory is by far the most common.

At the start of this section of the lecture, I suggested that the strategy for exploiting a patent exemplified by Ortho's behavior is widely assumed to be the most effective way of capitalizing on a patent. Frequently, that assumption is correct. For example, monopoly pricing by the patent owner is very common – and very lucrative – in the pharmaceutical industry.

But there are other approaches to exploiting a patent that are sometimes better. One of them is illustrated by the history of the windsurfer patent.

In brief, the story goes as follows:

Sailboats have existed for centuries. They come in various shapes and sizes, but two things that almost all have in common is that the mast or masts supporting the sails are fixed in some way to the hull, and that a moveable blade inserted into the water, typically at the back end of the boat, enables the operator to steer it.

. In 1967, Jim Drake, an aeronautical engineer, and Hoyle Schweitzer, a surfer and businessman, collaborated in creating a new kind of small sailing craft that had neither of these features.

Their creation had no rudder – instead had a fixed skeg, like a surfboard. And instead of being rigidly affixed to the hull, the mast and sail were connected to the hull with a universal joint (located here), allowing operators to steer these crafts by tilting the rig forward and backward – thereby changing the relationship between the center of effort and the center of lateral resistance.

In 1968, Drake and Schweitzer applied for a US patent on the invention, which was soon granted.

The key portions of the patent, including the primary claim, are shown are your screen. As you can see, the inventors sought and obtained a patent on a:

"Wind-propelled apparatus comprising [you now know the meaning of that term]

body means adapted to support a user [in other words, a surfboard]

and wind-propulsion means pivotally associated with said body means and adapted to receive wind for motive power for said apparatus, [in other words, a rig]

said propulsion means comprising a mast, a joint for mounting said mast on said body means, a sail and means for extending said sail laterally from said mast,

the position of said propulsion means being controllable by said user, said propulsion means being substantially free from pivotal restraint in the absence of said user, said joint having a plurality of axes of rotation whereby said sail free falls along any of a plurality of vertical planes upon release by said user…." [in other words, a universal joint]

Drake did not have the time or interest to commercialize the invention. But Schweitzer did. He and his wife Diane founded a company, Windsurfing International, assigned their share of the patent to the company and set about manufacturing and marketing these devices. Soon thereafter, Drake sold his share of the patent to Windsurfing International, leaving Schweitzer in charge of the infant company.

At this point, Schweitzer made a crucial and fateful decision. Instead of trying to monopolize the market, he licensed the patent to several other companies, enabling them to make and sell competitive products. At least three considerations may have underlain this decision.

First, Schweitzer may have realized that he lacked the capacity to meet the global demand for the product.

Second and more subtly, he may have realized that demand for the product would grow faster if a variety of sailboards were readily available everywhere.

Third, by granting licenses to competitors, he may have hoped to prevent them from challenging his patent – or at least to postpone such challenges.

Schweitzer's strategy was highly successful – at least for a time. Widespread skepticism concerning both the practicability and the commercial potential of the product gradually gave way to passion and popularity.

In the late 1970s and early 1980s, sailboards began appearing everywhere – on European lakes even more often than on American waterways.

One-design racing helped amplify interest, particularly among young people.

Schweitzer's company was highly successful – partly because it made and sold lots of these things, but also because it collected large license fees from other manufacturers.

The general point lurking in this part of the windsurfer story is that monopoly pricing is not the only way of deriving income from a patent.

Under some circumstances, licensing the technology, even to competitors, may be even more lucrative.

In addition to licensing, ways of exploiting a patent other than monopoly pricing include:

assigning the patent (for a fee of course) to a company in whose hands it would be more powerful

and using it to stimulate collaborations with competitors (for example, through standard setting organizations) or with the sellers of complementary products and services.

None of these strategies is optimal for all patentees or in all settings. The best approach varies radically by context. Some of the considerations affecting the relative merits of the different approaches will come up in subsequent lectures. For the time being, just keep in mind that monopoly pricing is not always the best approach.

C. Enforcing and Challenging Patents

The strategies discussed in the preceding segment of this lecture for exploiting a patent only work if most of the patent owner's competitors respect the patent – in other words, if they refrain from engaging in activities that would infringe the patent. With considerable frequency they don't. This segment of the lecture examines what happens then.

We begin with the question of why someone might violate a patent. It turns out that there are many reasons.

One of the most common is that the violator is simply unaware of the patent. In many fields of technology, there are myriad patents. They are all publicly available, but they are not well indexed. It would be hard for a company to keep track of all of the patents held by others that

might be relevant to the company's products or services, even if the company wanted to. And, for reasons we will explore later in this lecture series, there are sensible commercial reasons why a company would not want to monitor the patent landscape closely. The upshot is that lots of patents are infringed by people or firms who are unaware of them.

Sometimes, by contrast, the violator is aware of the patent but thinks that he or she can get away with ignoring it. Small enterprises sometimes do this because they think, plausibly, that they can fly under the radar or that the patentee will conclude that it's not worth her while to pursue them. The chances of avoiding detection are especially high when the patent pertains to a procedure useful in producing products, rather than to the products themselves. If a competitor employs such a procedure in its own plant, the patentee is likely never to learn of it. The competitor, aware of this, may decide to use the procedure without permission.

Another circumstance that sometimes foments widespread patent infringement is the existence of a custom within the relevant industry of not enforcing patents. The context in which this is especially likely to be true is academic science. Research tools developed by scientists are often patented by the universities that employ those scientists. But the universities holding those patents almost never enforce them against other universities or their faculty members. Academic scientists, aware of this custom, commonly use patented research tools without obtaining licenses for them.

A very different reason for not respecting a patent is that the violator thinks that the patent is invalid. As I explained in the first segment of this lecture, the examination process that precedes the issuance of a patent is often cursory. Indeed, in many developing countries, there is no substantive examination at all. The result is that many patents even in countries like the United States, would fall if examined closely.

These are sometimes referred to as "bad patents" or "weak patents" or "junk patents." All of those labels are somewhat misleading, insofar as they suggest that the patent examiner has failed to do his job – whereas often the examiner has done the best he can with the limited time and resources at his disposal. The labels are also misleading because they suggest that there are two discrete categories of patents – good ones and bad ones. The real situation is that all patents are, to some degree, vulnerable to challenge. For some, the probability that they would survive a challenge is high; for some it's medium; for some it's low, and so forth. All patents are, in this sense, probabilistic assets.

Strictly speaking, that's also true of real property – in other words, rights to land, buildings, and so forth. A title to a tract of land is never truly invulnerable. However, in most developed countries, the recording systems work well enough that the likelihood that most titles would collapse in the face of a challenge is extremely small. With respect to patents, the probability is often much higher.

Returning to the main thread of our discussion, a common reason for intentional violation of a patent is that the violator thinks that, if push comes to shove, the patent will be held invalid.

A final, loosely related reason is that sometimes the violator believes that push will never come to shove – in other words, that the patentee will never come after him. This might be because the patentee, like the putative infringer, is aware of the vulnerability of the patent to challenge and thus does not want to test its strength in court. Alternatively, the patent may be strong, but the patentee lacks the resources to mount an infringement suit.

Suppose that, for one or another of these reasons, a competitor of a patent owner violates the patent and the violation comes to the attention of the patent owner. What can the owner do?

One option, of course, is nothing. We've already considered various reasons why the patent owner might shrug – or grit her teeth and remain silent.

If the patentee does want to respond, her first move is likely to be to send the violator a "notice" or "cease and desist letter."

Sometimes the recipient of such a letter complies – either by ceasing the activity in question or, if the patent owner suggests it, obtaining a license. If so, the patent owner usually subsides.

If the recipient does not comply, the patent owner again has the option of doing nothing – in which case the cease-and-desist letter was a failed bluff.

If the patent owner instead wants to press the issue, the usual next step is to initiate litigation. At that point, the real fight begins.

Typically, in such fights, both parties have several procedural options. It is common for each to employ more than one of these options, either seriatim or all at once.

The sets of options available in different countries vary considerably. I'll describe two models that are especially dissimilar – the US model and the German model – and then suggest how a few other jurisdictions compare to those two.

In the US, the principal option available to the patent owner is a civil suit against the putative infringer seeking damages and/or an injunction. Such a suit is initiated by filing complaint in a federal District Court.

. There are a lot of district courts. As you can see from this map, each state has at least one. The more populous states have more than one. California, Texas, and New York each have four.

So, in which district does the patentee file suit?

. The relevant statute permits such a suit to be filed in any "judicial district where the defendant resides, or where the defendant has committed acts of infringement and has a regular and established place of business."

Until 2017, the Federal Circuit construed the "residence" of a defendant corporation very broadly, which commonly enabled patentees to sue corporate defendants in districts that were either convenient or attractive to the patentees. As you might expect, districts that were especially popular among patentees included the District of Delaware, the Northern and Central Districts of California, and the Northern District of Illinois. Less predictable was the extreme popularity of the Eastern District of Texas.

Here are the data through 2015.

What explains the remarkable and growing popularity of the Eastern District of Texas? Primarily the fact that that district adopted several rules that were highly favorable to the plaintiffs in patent infringement suits.

In 2017, in the TC Heartland case, the Supreme Court repudiated the Federal Circuit's expansive interpretation of where a corporate defendant resides – reinstating an older, much narrower interpretation. Now "residence" means "place of incorporation."

. The impact of TC Heartland was dramatic and immediate. As you can see from these bar graphs, the percentage of patent infringement suit filed in the eastern district of Texas fell from around 40% to around 15% -- while filings in the District of Delaware, where most large US corporations are incorporated, grew substantially.

I don't want to leave you with the impression that patentees no longer have any options concerning where to file. If they can establish that a corporate defendant "has committed acts of infringement and has a regular and established place of business" in a district other than where it is incorporated, then the patentee can sue there. But patentees definitely have fewer choices than they did a few years ago.

Most patent infringement suits never get further than the district court – indeed, never reach final judgment. The slow pace and the high cost – and the hazard of large damage awards – prompt the large majority of litigants to settle their disputes sometime before judgment is entered.

If they don't, and if the case proceeds to final judgment, then one or the other party will win – or each will be victorious on some but not all aspects of the case.

So, returning to the diagram I introduced early in this lecture, if the inventor has obtained a patent, someone else has violated it,

the patentee has brought an infringement suit in federal district court –

the court may rule in her favor (typically, by finding the patent to be both valid and infringed)

or it may rule in favor of the defendant.

(To keep things simple, I'm assuming that either the plaintiff (i.e., the patentee) or the defendant prevails on all issues.

(In this diagram and in most diagrams I'll be using in this lecture series, the patentee will be identified by the Greek letter pi, while the defendant will be identified by the Greek letter Delta.).

The loser is not yet out of luck; he, she, or it has a right to appeal to the Court of Appeals for the Federal Circuit, whose exclusive jurisdiction over patent cases we discussed earlier. The Federal Circuit will, in turn, rule in favor of the patentee or the putative infringer.

The loser in the Federal Circuit can use a petition for certiorari to ask the US Supreme Court to review the case, but the percentage of such petitions that the Supreme Court grants is very small. In the rare instances in which the Supreme Court does take the case, one or the other party will win. The decision of either the Federal Circuit or the Supreme Court will send the parties back to the District Court – at a minimum for the formality of entry of judgment, but often for further proceedings.

This is the path that most determined patentees will follow. But nowadays, there is an additional option available to patentees.

If the defendant is importing into the United States products that the patentee considers infringing (for any of a variety of reasons we will explore in lecture #6), then the patentee, instead of going to a federal district court, can ask the International Trade Commission (commonly abbreviated ITC) for an "exclusion order" – which, as its name suggests, would forbid the entry of the products in question into the US.

The big advantage of proceedings before the ITC is that they are much faster than civil suits. Among the disadvantages are that the patentee cannot recover damages in this fashion, and cannot stop domestic production or sales of the products at issue. In addition, rulings by the ITC do not bind the district courts.

Appeals from the ITC also go to the Federal Circuit and, in very rare cases, to the Supreme Court.

The defendant is not powerless in all this. As you might imagine, a defendant can assert in the course of litigation (either in a district court or before the ITC) that he did not infringe the patent at issue.

In addition, in the United States, the defendant can – and almost always does – argue that the patent at issue is invalid, for any of the reasons I have already sketched and that we will be considering in much more detail in the next four lectures.

The defendant also has various ways of taking the offense with respect to the validity of the patent.

A traditional way of doing so is to file a declaratory judgment suit in federal district court. As the name suggests, such a suit seeks a declaration that the patent is invalid or that the defendant has not infringed it. Because federal courts in the US are permitted only to decide actual cases and controversies, the defendant cannot initiate litigation in this fashion until the patentee at least makes a credible threat of filing suit. How much saber-rattling is necessary to trigger this option on the part of the defendant was recently clarified by the Supreme Court in the Medimmune case.

Why would the putative infringer file a declaratory judgment action, rather than simply wait until the patentee initiates litigation? Usually, because the infringer wants to control the choice of jurisdiction in which the fight is fought.

Alternatively, the infringer can ask the Patent Office to reconsider the validity of the patent. Until recently, such efforts to reengage the Patent Office were uncommon. But one dimension of the so-called America Invents Act (adopted in 2011) has increased them sharply.

This was the creation of the so-called *inter partes* review procedure. In brief, starting 9 months after the issuance of a patent, a third party can petition the PTAB to assess the validity of the claims in the patent – specifically, on the ground that, when tested against preexisting patents or printed publications, those claims fail the novelty or inventive step requirement. The Board will institute the proceeding (typically around 6 months after receiving the petition) if it finds that there is a reasonable likelihood that the third party will prevail with respect to at least one of the challenged claims.

The steps in the *inter partes* review procedure are set forth in this chart, prepared by Fish and Richardson – one of the leading patent law firms. As you can see, ordinarily, the board will release a final determination and an associated written opinion within a year. Parties unhappy with the PTAB's decision can appeal it to the Federal Circuit.

The most common procedural context in which inter partes review is employed is as a maneuver in infringement litigation.

The patentee files a complaint in federal district court against a party she alleges is infringing the patent.

The defendant responds by seeking inter partes review of the validity of some or all of the claims before the PTAB. (If the defendant is going to use this option, it must do so within a year of service of the complaint.)

And then asks the District Court judge to stay the infringement suit until the inter partes review is complete.

The judge has discretion in deciding whether to grant such a stay. Most will decline to do so until the PTAB has decided whether to institute the proceeding, but most will grant such stays

after the PTAB institutes a proceeding. The result is that the infringement litigation is on hold for roughly a year, while the parties fight before the PTAB.

If the defendant prevails before the PTAB, and Board cancels some or all of the claims, the District Court will be bound by its ruling. If not, the stay will be lifted and the infringement suit will proceed.

The inter partes system went into operation on September 16, of 2012. In the first few years in was used increasingly often – more often than its creators seem to have envisioned. During these early years, the PTO decided to institute a high percentage of the petitions, and the rate at which – at the end of the day – challenged patents were either wholly or partially invalidated was very high – around 85%.

Gradually, however, both the number of petitions filed and the rates of institution declined. They now seem to have stabilized. Roughly 1200 petitions are filed each year. In roughly 2/3 of the cases, the petition is instituted.

In a reasonably high percentage of the cases, the parties settle, either before or after institution.

When they don't, the rates at which some or all of the claims in the challenged patents are deemed invalid are high – not as high as in the early years, but still daunting.

OK. That, in brief, is how patents are enforced and challenged in the US. In Europe, as I mentioned, an inventor can either obtain patents from the individual national patent offices or get a single European patent from the EPO. For the time being, both the national patents and the European patents are enforced and challenged in national courts.

Of the various national court systems, that of Germany is by far the most important. More patent infringement suits are brought there than in any other country. After Germany, the most popular jurisdictions are, in order, the United Kingdom, France, and (interestingly) the Netherlands.

Many observers think that the German system is better than that of the United States – faster, more efficient, less expensive, and more predictable.

Here's a diagram of the system, provided by a group of 8 scholars published recently in the European Journal of Law and Economics.

The column on the left shows the options available to a patentee seeking to enforce a patent. The column on the right shows the offsetting options available to a challenger.

The lefthand column should look familiar. The holder of a German patent or a European patent can bring an infringement suit in any one of 12 district courts in. Germany. The three most popular are those in Du¨sseldorf, Munich and Mannheim.

The party who loses in the trial court can then seek review in appellate courts. In these respects, Germany is like the United States – and virtually all national patent systems.

The upper-right hand corner should also look familiar. Within a limited period of time, any other party can initiate an opposition proceeding challenging the patent. For a European patent, the proceeding would be brought in the EPO. For a German patent, it would be brought in the German Patent Office.

Here the similarity ends. In sharp contrast to the US, in Germany, the defendant in an infringement action cannot assert that the patent at issue is invalid. Instead, the defendant must initiate a separate suit – called a nullity action – in the Federal Patent Court – specifically before a panel of five judges, three of whom are technically qualified. Very roughly speaking, you can think of these nullity actions as analogous to IPR proceedings in the US, but note that they proceed before a court, not an administrative tribunal. You can only initiate one after the opportunity for pursuing an opposition has lapsed.

The loser in a nullity proceeding can appeal to a specialized appellate court.

How often are such nullity actions successful? Surprisingly often. Here are some eye-opening numbers, collected by Tilman Muller-Stoy and his colleagues. As you can see, between 2018 and 2020, there were 221 such suits. In only 20% of them was the patent entirely upheld. In 43% it was partially revoked, and in 38% it was entirely revoked. This is not a recent development; previous studies of nullification rates came to the same conclusion.

There's a modest difference between the way in which the courts handle patents that originate in the German Patent Office and the way in which they handle European patents.

As you can see, the German patents were somewhat more likely to be upheld than the European patents. That's consistent with the general view that the German Patent Office is unusually scrupulous when reviewing patent applications – and thus that German patents are "stronger" than US or European patents. But still – only 32% of the German patents survived unscathed.

On what basis were the patents revoked? As you can see, the primary grounds were lack of novelty or inventive step – the same grounds that loom large in the US.

What happens on appeal? As you can see, in only 55% of the cases is the lower-court ruling upheld.

In 75% of the cases in which the lower-court ruling is upheld, the patent is wholly or partially revoked. In most of the cases in which the lower-court ruling is reversed, the result is total or partial restoration of the patent. In other words, it's rare for the appellate court to order revocation of a patent when the lower court did not.

So, what happens in the typical case in which the patentee initiates an infringement suit, and the defendant responds by initiating a nullity proceeding? If the judge in the infringement suit

concludes that nullification is highly likely, he or she will stay the infringement suit until the nullity proceeding is resolved. But that only happens 10 to 15% of the time. Much more often, the two suits proceed in parallel. Usually, the infringement suit is resolved first. If the patentee wins, she can enforce the judgment. The result is that the defendant's further activity may be enjoined, even though the validity of the patent remains in doubt.

The second major difference between the German and US procedures is that, in Germany, the losing party always pays the attorneys fees of the winning party. As we'll see later, that sometimes occurs in the US, but it's uncommon. The financial danger to the loser is thus higher in Germany.

In practice, the biggest difference between the two systems is that, nowadays, most infringement suits in the US are tried before juries composed of lay people, whereas in Germany judges determine the outcomes. As Bill Lee will explain, this has a huge effect on the way in which trials in the two jurisdictions are conducted.

What about duration? The following figures from the Cremer study come from the early part of this century, but are probably still relatively accurate. As you can see, in Germany, the time between the filing of an infringement suit and the issuance of judgment by the trial court was less than 2 years. A judgment of revocation typically took longer than a judgment of infringement.

The numbers in the UK and the Netherlands were similar, but the French courts took substantially longer. The courts in the US are even slower.

As I mentioned, most commentators praise the German system, but not all. Some contend that, by depriving defendants to assert the invalidity of a patent in an infringement suit, forcing them instead to initiate a nullity proceeding, the German system disadvantages defendants, especially small and medium-sized companies, and favors so-called patent trolls.

"Patent troll" is a pejorative way of referring to patentees who do not practice the invention – in other words, whose business model consists of suing parties who are actually engaged in business. The term is meant to conjure images like this – or this – beasts who exact tribute but don't contribute anything to society. A more neutral term for this type of patentee is non-practicing entity or NPE.

Whether NPEs are good or bad is a more complex question than first appears. Most research universities, for example, are NPEs. We'll return to this issue later in this lecture series.

The systems used in other countries to enforce and challenge patents typically resemble the US in some respects and Germany in other respects.

In Japan, patents may be challenged before the Japanese Patent Office in a so-called Trial for Patent Validity – which loosely resembles inter partes review in the US. However, in Japan,

unlike the US, a slight majority of such challenges are brought prior to the initiation of an infringement suit, rather than in response to one.

The system in Argentina resembles that in the US in that invalidity can be asserted as a defense in an infringement suit, but German-style nullity actions are also permitted. Argentina, unlike the US, considers patent infringement a crime as well as a civil wrong. As a result, patentees have the option of bringing suit in criminal courts, rather than civil courts – although they rarely do.

In South Africa, infringement suits are heard by the Commissioner of Patents. The title, however, is misleading. The Commissioner is a judge, appointed to appointed to the Court of the Commissioner of Patents, a specialist court of the Gauteng Division of the High Court of South Africa, Pretoria. As in Argentina, invalidity may be asserted either in a nullity action or as a defense to infringement.

In most other countries in Africa, patent litigation is currently rare, so it is difficult to generalize about the procedures used.

This concludes my review of the ways in which patents are obtained, exploited, and challenged. You will see many examples of these procedures in the coming weeks, as we explore different dimensions of the substantive law of patents.