

How many patents are enough?

It is generally accepted that an organisation can have too many or too few patents. While finding and maintaining the equilibrium will always be challenging, strategic patent intelligence offers the analysis required for evidence-based decisions which are capable of withstanding robust scrutiny

By Nigel Swycher and Steve Harris

“The simplest questions are the hardest to answer” –
Northrop Frye

Worldwide, more than 1 million organisations own patents, while over 30,000 organisations spend more than \$250,000 a year obtaining and maintaining patents. These figures exclude the \$1.7 trillion invested in R&D each year, which the global patent system is designed to protect.

Despite the high stakes and levels of investment, there is no agreed methodology to help organisations measure and report on whether their patent portfolios are fit for purpose. This article explores some of the approaches taken by owners of large patent portfolios and suggests a way of harnessing strategic patent intelligence to organise available data into an evidence-based solution for companies that wish to minimise their exposure to general and specific threats of infringement.

There is no single answer to the question of how many patents are enough and it is hard to reconcile the various perspectives and approaches. It is tempting to assume that more is better, but this is true only if there are no significant costs associated with obtaining and maintaining substantial patent portfolios. Such an answer also ignores the increased scrutiny from chief technology officers, chief financial officers and the many other stakeholders who have an increased interest in patents.

Enter strategic patent intelligence. The past 10 years have seen an explosion in the quality and variety of both public and private data relating to the ownership and use (both litigation and licensing) of patents. The challenge has been making this data accessible at a speed and cost that makes it feasible to integrate into everyday decision making.

By harnessing this ability to compare a portfolio both generally against the market and specifically against other companies, this article sets out a model for balancing a portfolio against a blend of these risks. While the model assumes a defensive strategy, it could be readily adapted to meet other objectives, such as monetisation.

How many patents are enough?

This section contains a summary of the views expressed by strategic decision makers in major US and European companies in response to the question ‘how many

patents are enough?’ They include responses from a panel discussion at *IAM’s* Auto IP USA conference held in Detroit in May 2019 and a roundtable held by CIPHER the following day. Sectors represented include aerospace and defence, semiconductors, video and audio communication, industrial automation, financial services, fast-moving consumer goods, chemicals, automotive and a variety of companies collectively described as ‘technology’. Also included are views expressed at The Gathering, a US-based IP best practices group which is looking to create standardised decision-making processes in this area, as well as the results of an informal survey.

Brace yourself. It is hard to reconcile the various perspectives and approaches.

None at all

Many companies would be happy with a quiet life in which they are free to develop and commercialise their products without any third-party interference. Against this objective, it is not important to develop patents for enforcement or monetisation.

This strategy should not be equated with naivety. Those who advocate (but do not implement) this strategy recognise that it exposes an organisation to greater risk of inbound litigation, but that the assets needed to counter-assert against an aggressor are often not those of the organisation. In these situations, the deficit is commonly addressed through patent purchase. There is an active secondary market for patents, at somewhat suppressed prices, making it feasible for defendants to shop for what they need, when they need it. Variations on this theme include patent purchase programmes such as AST’s IP3 and Uber’s UP3.

So, at the extreme, if a company can buy the patents that it needs at times of distress, this obviates the need to build a native portfolio.

One good one

When push comes to shove, what is required is a patent that reads squarely onto a third-party product. This utopian view assumes that any more than one patent is superfluous. While no companies have adopted the ‘just one’ strategy, this position encapsulates numerous industry perspectives.

First, patents protect revenue. ‘One good one’ assumes that the objective is to deter competition and provide an effective means of preventing or enforcing that position.

Second, the focus on quality. While there is near universal support for quality over quantity, there is much less agreement on how to guarantee that a patent is or will be high quality.

Finally, the challenge of predicting the future. Irrespective of whether an organisation's IP strategy is offensive or defensive, patents that read onto competitor or third-party products or services are pure gold. The difficulty is knowing (whether at the time of first filing or later) who is infringing such patents.

One more than your competitor

This approach reflects the reality of détente. The automotive sector is a great example of an ecosystem which includes a relatively small number of original equipment manufacturers (OEMs) supported by a tight network of suppliers. Between them they own more than 500,000 patent families and there has been little to no litigation among this group in the past 10 years. No one has any incentive to stir up a hornet's nest by asserting one or more patents.

FIGURE 1. Example taxonomy: Autonomous systems (source: Ciper Automotive (partial taxonomy extract))



There are a range of problems with this position. First, it is an arms race. Staying ahead means continuously filing for more, a natural reluctance to let patents lapse and the need for an ever-increasing budget.

Second, it ignores the changing nature of competition. While traditional automotive benchmarking was conducted by reference to US, Japanese and European OEMs and their major suppliers, this has changed. In a sector where advances in electrification and autonomy will be determinative, there is now a cohort of technology companies and start-ups in a range of new geographies to take into the equation.

Third, it is difficult to establish a reasonable basis for comparison. Most companies approach this challenge by designing and applying taxonomies. A taxonomy in this context is most commonly a structured list of technologies to which all patents are tagged. This is essential in order to compare like with like. Two companies seldom compete with each other in all of the same areas. So, while Nestlé and Unilever compete head on in the ice cream market, they are not competitors with regard to detergent or pet food. Comparisons must also factor in differences in revenue and market share. Benchmarking is relative and must take account of scale.

As many as you can

Patents are assets. Assets have a value, so more is more. While superficially attractive, this strategy ignores cost. If a patent costs more than the benefit that it delivers, it fails on the most basic economic analysis. It also assumes the mythical unlimited patent budget.

The real-world equivalent of this is as many as the agreed quota. Patenting is commonly keyed off numerical targets, such as the number of new filings. This number is often derived from R&D headcount or spend. While a reasonable basis for ensuring a correlation between business strategy and patenting strategy, it largely ignores the patent landscape more generally. Notably, no organisations have indicated that they are measured on the overall size of their portfolio.

Building an ideal patent portfolio

The range of answers in the previous section creates a picture of the many drivers of patenting strategy. These include reputation (eg, for innovation), monetisation and contribution (eg, to pools and start-ups). This spectrum was validated in a recent survey of participants at The Gathering, which also confirmed that the most common drivers were defensive and to provide protection against threats. The risk is that a company will be forced to pay substantial royalties to other owners of relevant patents. All of this ignores the impact of NPEs where the size and shape of an organisation's own portfolio is largely irrelevant.

In this article, the ideal portfolio is one that helps to minimise the threat of a third party using its portfolio to extract royalties from the sale of a product by a company (hereinafter, A). Threats are specific when coming from an identified third party and general when created by the vast number of other owners of patents relating to technologies incorporated into company products. Most companies actively maintain a threat list. These are the identified patent owners that for competitive or other reasons are perceived to be motivated to pursue a company for patent licensing revenue.

Our starting point is designing a portfolio that minimises the exposure of A to a specific threat. The calculation to estimate exposure is relatively straightforward:

$$\text{balancing payment} = s \frac{\hat{p}_a r_b - \hat{p}_b r_a}{\hat{p}}$$

In the above calculation, p is the number of patents that cover the product area, r is the revenue and s is the royalty rate. The subscript a indicates the portion attributable to A, b to threat list company B and so on. This logic was previously described in “The role of AI in evidence based strategic IP decisions” (IAM, November/December 2018).

Minimising the risk position against a series of competitors requires the balancing payment to be zero or in A’s favour – hence:

$$s \frac{\hat{p}_a r_b - \hat{p}_b r_a}{\hat{p}} \geq 0$$

$$\hat{p}_a r_b \geq \hat{p}_b r_a$$

$$\frac{\hat{p}_a}{r_a} \geq \frac{\hat{p}_b}{r_b}$$

The calculation is run for each of the companies on the threat list. There will be negative and positive results, which is to be expected. The formula at this stage also assumes that all patents are equal, which is generally not the case (this issue is addressed later).

The strongest position would be to have an overwhelming number of patents in all areas, but the goal is to estimate the position required to achieve a balanced portfolio, such that the economically optimum number of patents is held in each area. This requires a situation where a cross-licence with the strongest competitor in the space is likely to produce a licence which can be obtained at a lower cost than the carry cost of sufficient patents in the area, when risk adjusted.

If all of the patents covering one area are considered, there will be an average number of patents per million dollars of revenue (the dotted red line in Figure 2) across all companies. There will also be a number of specific competitors, which will be either above or below that average for their figures.

In Figure 2, companies below the dotted red line are more exposed than average to disadvantageous licensing outcomes, whereas those above it are less exposed. The further a company is below the line, the greater the exposure in this specific product area. This does not mean that being above the line provides an adequate defence against a specific opponent, as a counterparty that is especially well stocked in patents compared to their revenue will still be able to negotiate a licence in their favour in most cases.

In the ideal situation, A would have a number of patents compared to its revenue that exceeds the ratio of every player. However, in reality, few companies have a sufficient budget to achieve this. By considering the threat level of each competitor and A’s tolerance for risk, it is possible to balance these considerations and come to a rational position to account for A’s place in the industry.

Balancing portfolio against specific threats

Using the approach from the previous section, it is possible to calculate the optimum (target) size of A’s portfolio in each technology area. In the simplest case, consider the companies on the threat list one by one and perform the below calculation:

$$\frac{t_b}{r_a} = \frac{\hat{p}_b}{r_b}$$

In the above calculation, t_b is the target or optimum number of patents to have in A’s portfolio to counter the specific threat, B. This results in a position that would just meet the inequality discussed earlier.

Not all patents are equal

The most common objection to an approach based purely on patent quantities is quality or patent strength. The survey of Gathering members identified patent quality as the primary focus of internal patent strategy. It is therefore appropriate to account for this with the addition of weightings.

If there is more information about the relative strength of A’s (w_a) and the threat’s (w_b) portfolios, a weighting can be added to the equation:

$$\frac{w_a t_b}{r_a} = \frac{w_b \hat{p}_b}{r_b}$$

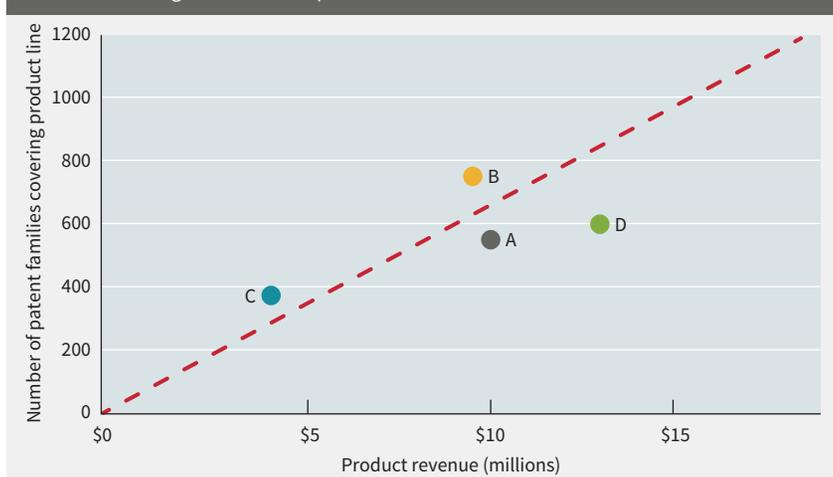
The portfolio strength weight includes numerous factors, including the quality of the patents and relevance to the product area. An average portfolio will have a quality weight of 1.0, while one that is 20% stronger than average will have a weight of 1.2.

Rearranging this to focus on target size results in the following calculation:

$$t_b = \frac{w_b \hat{p}_b r_a}{w_a r_b}$$

These factors are ultimately subjective, but by using a combination of past history, experience of

FIGURE 2. Looking for a balanced portfolio



negotiating cross-licences and conventional patent scoring methodologies, it is possible to produce a representative weight for a known portfolio. In the absence of information, taking a value of 1.0 for the weights is a reasonable assumption. It should be stressed that this is not advocating pick lists or individual claims charting, as in many situations it can be assumed that both parties have excellent (and weak) patents in broadly equivalent proportions.

There is also a level of uncertainty when it comes to weighting portfolio strength, as few of the factors that go into it can be precisely quantified, but the methodology is robust against approximation in this factor. Where there is doubt, a conservative approach would be for a company to down-weight its portfolio to give a pessimistic position.

Taking account of all owners of relevant patents

To deal with the rest of the industry (ie, companies that are excluded from the specific threat list), the optimum position overall for all companies operating in the sector can be calculated. Calculations of this sort have historically been avoided because of the difficulty in extracting accurate data. However, advances in the automation of mapping patents to technologies have eliminated much of the manual work previously required.

Weightings are introduced to take account of A's portfolio strength as previously described. So, if the weight for A's portfolio is 1.2 (ie, 20% stronger than average patents in the product area), then by definition the mean weight of all other patents in the area must be as follows:

$$\frac{p - w_a p_a}{p - p_a}$$

Thus, the balancing equation for the area as a whole is:

$$\frac{w_a t_{general}}{r_a} = \frac{p - (w_a p_a)}{p - p_a} \times \frac{p - p_a}{r - r_a}$$

$$= \frac{p - (w_a p_a)}{r - r_a}$$

And therefore:

$$t_{general} = \frac{r_a (p - w_a p_a)}{w_a (r - r_a)}$$

In this example, the general threat represents all companies in the sector, including companies B to D. Alternatively, the general threat can be modelled as

being the remaining companies in the sector. This adds some complexity to the calculations, but only requires subtracting $\sum w_x p_x$ from p , and $\sum r_x$ from r in the final equation above, for $x \in \{a, b, c\}$ (ie, A and all of the companies on the threat list).

Balancing a portfolio for both specific and general threats

Conducting the specific and general calculations in isolation will deliver different optima for each product area: one from each of the companies on the threat list and one from the industry as a whole. These must be combined to produce a single optimum value.

The approach proposed is the addition of another set of weights, recognising that not all threats are equal (referred to as 'threat weights') to represent the likelihood and severity of each of the threats.

These weights can be employed in numerous ways – for example, if the goal is to construct the strongest possible defensive position, it is sufficient to weight the specific or general threat with the highest optimum position as 1.0 and everything else as 0.0. However, this is not a balanced approach and will typically result in a large portfolio.

A more usual approach would be to weight the threats in terms of the relative concern in that product area – for example, rating the organisation that is most likely to engage with A with a 10 and all others proportionally to that figure. There is no particular scale for the threat weights, they are simply scaled with respect to each other.

If v represents the threat weight for a specific product area, then – considering all specific and general threats – t , the target optimum number of families, is given by:

$$t = \frac{\sum t_x v_x + t_{general} v_{general}}{\sum v_x + v_{general}}, x \in \{b, c, d\}$$

In the above calculation, b, c, d are the specific threats.

Putting this all together: a worked example

As a worked example, consider A with three specific threats (B, C and D) in a particular product area. A believes its own portfolio to be 20% stronger than average, with 550 granted families and a revenue of \$10 million. The industry as a whole has 6,000 families covering this area and a total revenue of \$93 million.

In this example, the weights are set by the following rationale:

- B is seen as the biggest threat (ie, most likely to pursue), so gets the highest threat weight.
- C is seen as a relatively small threat (ie, it has a big portfolio of average quality and not much revenue).
- D is seen as a low threat (ie, it has a big portfolio of low quality and high revenues).

TABLE 1. A competitive analysis

	Quality weight	Families	Revenue	Optimum	Threat weight	Target
A	1.2	550	\$10 million	-	-	783
B	1.3	750	\$9.5 million	855	10	
C	1	375	\$4 million	781	2	
D	0.8	600	\$13 million	308	1	
General	-	6,000	\$93 million	536	1	

Note that in cases where A's revenue in some areas is zero, this model will suggest that the correct number of patent families is zero. However, this is not necessarily the case, as they have value when applied to other companies that do have revenue in that area. This is covered in later in the section on taking a holistic approach.

Projecting into the future

"It is difficult to make predictions, especially about the future"
– Danish proverb

This methodology provides one way of thinking about the correct size of A's portfolio, given the world today. While this can be used to judge the effectiveness of patenting strategy to date, without a time machine, it is not going to address any current strategic challenges faced by the business.

It is a relatively simple exercise to project forward the current product and patent strategy. This projection can be used to calculate the numbers as they will be in the future. This methodology can then be used to set a strategy that will provide a strong defensive position in, for example, five years' time. Projecting forward competitor positions is more difficult, but there are clues in the patent data that can give a reasonable estimate for future positions (eg, looking at current filing and grant rates and looking at annual and quarterly reports to see evidence of the business strategy).

Need for greater sophistication

The methodology described to this point gives a rational and relatively easy-to-apply approach to estimating an ideal portfolio size in isolation. Greater sophistication takes account of more variables, at the cost of some complexity.

International filings

An important consideration is geographical coverage. All calculations are based on the assumption that the families in a portfolio have a representative and industry-normal make-up of territories. If that is not the case, this should be addressed by down-weighting the portfolio weight. Similarly, when considering specific threats, the geographical coverage of their portfolios should be considered as to whether it places limitations on the effectiveness of the portfolio (eg, China-only families against a company with no revenue in that jurisdiction).

Cost of achieving target size

The basic approach enables estimation of the make-up of the ideal portfolio, but ignores the cost of reaching that position and the cost of not doing so. This section introduces these considerations, which help to assess the impact of the gap between A's current position and the ideal.

The estimation of the cost to organically build a portfolio (or equivalently, the excess investment when the portfolio is overstocked) is easiest to calculate: the difference between the target and current size is multiplied by the typical per-family cost of obtaining and maintaining a patent and divided by the grant rate in that technology as follows:

$$cost = (t - p_a) \frac{c_a}{g_a}$$

The cost to acquire the required assets in the marketplace is more difficult to predict and requires research into the current market price for the kinds of asset required. Of course, these two approaches are not equivalent – organic growth has the advantage of being targeted towards exact technologies of interest, but is slow to achieve and uncertain. On the other hand, acquisition is much faster, and the exact make-up of the portfolio is known before money changes hands. That said, it is typically more expensive.

In order to estimate the potential benefit of growing a position from understocked to target t , for a company on the threat list, the original formula which includes the patent denominator and the royalty rate can be used:

$$balancing\ payment = s \frac{p_a r_b - p_b r_a}{p}$$

The potential benefit of growing the portfolio to t is the difference between the balancing payment at the target position and at the current position:

$$\begin{aligned} benefit &= s \frac{t r_b - p_b r_a}{p} - s \frac{p_a r_b - p_b r_a}{p} \\ &= \frac{s}{p} (t r_b - p_b r_a - p_a r_b + p_b r_a) \\ &= \frac{s}{p} r_b (t - p_a) \end{aligned}$$

Counterintuitively, this does not depend on A's revenue in the area because (for this form) the assumption is that the revenue remains relatively constant. Thus, the benefit is mostly dependent on the revenue of the specific threat company.

Note that this assumes that the revenues and patent ratios remain approximately in balance over the period required to reach t . If that is not the case, the following, more sophisticated model, can be used:

$$benefit = s' \frac{t r'_b - p'_b r'_a}{p'} - s \frac{p_a r_b - p_b r_a}{p}$$

In the above calculation, x' indicates an estimate of the value x after the portfolio has grown to t patents. An example of when to employ the more complex calculation is where it is known that one of the threat companies was pulling out of a product area or growing its market share quickly. This benefit must also be adjusted by the likelihood of a licence being entered into by each company on the threat list.

Balancing multiple products requires holistic approach

Another topic for increased sophistication of the model is considering multiple product areas at once. There will be many situations where the company portfolio is understocked in one area and overstocked in another, but where the two balance out economically such that the balancing payment in one area corrects for the balancing payment from another. In a survey conducted by The Gathering, combined under and overstocking was identified as the common case.

Action plan



Use the methodology suggested to establish whether your organisation has the ideal defensive portfolio. Start with one product or technology.

- Classify your patents into technology areas. This may be an opportunity to refresh your internal technology taxonomies.
- Identify and classify the relevant portfolios owned by companies on your threat list and the market more generally. Strategic patent intelligence platforms such as CIPHER have been optimised for these calculations.
- Gather the necessary revenue data. Most will be readily available internally as part of existing competitive or market intelligence reports.
- Plug the numbers into the formulas provided to calculate the ideal size of your portfolio, applying appropriate weightings for a subjective view of patent strength and perceived risk.
- Experiment with the extensions to the model for future proofing and geographic differences.
- Analyse the areas where you are overstocked or understocked and formulate your correction strategy.
- Assess whether this approach helps with the development of patenting strategy and communication with other parts of the organisation who ask 'how many patents are enough?'

Because of this effect, it is worth calculating the net position against each specific threat company and, in some cases, it would be possible to reach an optimal position with fewer patents than the analysis of individual product areas would suggest.

This view is also important when considering the contribution of patents held by a company in areas where it generates no revenue but the threat company does. This will have $r_a = 0$, so the hypothetical cross-licence would result in a net payment to the company, but the product area at a time methodology would suggest that an optimal number of patents in the area is zero. By considering this in the context of other technologies, it will be clear where that corrects for some understocking in other areas, for some specific threats.

It is not possible to consider a holistic position for the general threat, as there is no representation of the balance of revenue and patents within the companies that fall into the general position.

It is possible to build a mathematical model for the net position against the threat list as a whole and solve the system of linear equations in the model, or use a successive approximation technique to discover optimal positions. However, this is excessive – the level of subjectivity in the weighting makes the precise solution uninteresting and the solution would typically be an irrational position to hold.

Challenges to this approach

Asking how many patents are enough is controversial. It implies that IP teams do not have the answer, and terminology such as under and overstocking exacerbates this impression. The idea that achieving a strategic balance can be helped with formulas meets resistance from those who think that portfolio development is more art than science and that patents should not be treated as an asset class.

The above approach seeks to blend the best of both worlds, harnessing the fact that it is now possible to automate the mapping of patents to technologies (making the essential data accessible) with the recognition that quantification on its own ignores the need to introduce subjective weightings. This does not undermine the quality of the results, but rather introduces flexibility into the model to explore many

permutations – more than one of which will likely deliver an acceptable outcome.

Some of the maths and modelling is sophisticated. While this is true, none of it is outside the bounds of a competent business analyst. When assessing the effort and reward, consider the existential threat of not answering the question – namely, the task of explaining why a portfolio does not include the patents that a company needs and justifying substantial payments to others that could have been avoided.

Building a portfolio takes time. Even if it were true that a company owned too many or too few patents, how practical would it be to rectify that imbalance? Many of the strategies in this regard are well known – for example, invention harvesting, patent pruning and patent acquisition.

Inherent in this approach is the reality that not having the right patents in the right areas will likely result in the payment of royalties to others. To minimise and manage this risk means identification of third parties and their portfolios. There are legitimate grounds for concern that this level of proactive risk management triggers wilful infringement considerations. This is manageable, but does require caution, particularly when dealing with those outside the privilege afforded to legal teams, such as board presentations. A common solution is the retention of consultants to perform the detailed calculations, such that sensitive information is kept away from the company.

Adding it all up

The question of how many patents are enough is foundational to an organisation's IP strategy. The current state of play does not lack sophistication, but there is insufficient consistency for there to be a consensus on what is best practice.

The approach described in this article has compelling logic where the risk can be articulated as the cost of understocking at both the specific (identified third parties) and general levels (by reference to the global population of relevant patents). It starts from the position that portfolios must be proportional to market share and competitors' revenues. Integrating market data in this way makes perfect sense. Those who challenge this approach believe that quantification ignores both patent strength and the fact that risk must be tempered with likelihood of occurrence. Both of these subjective weightings have been factored into the approach, with a level of transparency and flexibility that promotes more inclusive decision making. It is also possible to build into the model predictions for the future and geographic differences to calculate the cost of getting to the ideal.

Evidence-based solutions of the sort described do expose strategic IP decisions to scrutiny. This is not a bad thing if the consequence is more respect and understanding of patents across the organisation and a portfolio that is optimised for the threats that would otherwise translate into financial exposure. There are a number of organisations that currently employ this technique or ones closely resembling it in order to set their patenting strategy, and it is an interesting step forward in the development of IP strategy in a world where patents are increasingly important. **iam**

Nigel Swycher is CEO and **Steve Harris** is chief technology officer of CIPHER, a strategic patent intelligence platform