Unpacking the royalty stack

Just how much device makers are on the hook for in patent licensing payouts is one of the most debated questions in the IP community. While a complete answer may remain elusive, an analytical approach can help manufacturers assess their risk

By Erik Oliver, Kent Richardson and Hannes Forssberg Malm

The issue of royalty stacking may well be the Holy Grail for device makers – unfortunately the answer does not come in a single, easily portable form. Rather, the question of how much a company will have to pay for patent licences combines technology forecasting, market adoption estimation, predicting patent law evolution and dozens of other factors. While the topic is most often discussed in relation to complex consumer electronics products, it can arise anywhere that there is a high integration of technologies per product and a large number of patents per technology.

Even if the scope is narrowed to look at what happens when a high-tech product is faced with licensing not tens, but thousands – or possibly tens of thousands – of patents (eg, royalty stacking), there are no easy answers. While new machine-learning tools can help a company to quantify the issue and provide some (partial) answers to this problem, the limited availability of data can hamper the effectiveness of such tools.

Nevertheless, asking the question still provides vital insights into managing the issues created by royalty stacking – and underlines the fact that it is impossible to completely eliminate uncertainties in this area.

Why ask the question at all?

Research estimates are that there are several hundred thousand patents (250,000-plus) relevant to a smartphone (eg, page 59 of RPX's S-1 registration with the Securities and Exchange Commission in the United States). Many of those patents are related to the cellular (wireless) technical standards, which are but a handful of the standards that a handset will implement. Other sources of patents include non-standard, but commercially necessary technologies (eg, OLED/LED screens, touch sensors and antennas). The scale of integration and, correspondingly, the potentially relevant patents is difficult to comprehend. A smartphone is just one example of a product struggling with the broader concerns discussed here.

This massive number of patents reading onto a technical standard and necessary technologies creates a significant degree of uncertainty as to how to account for the potential licensing costs for these patents– often referred to as 'royalty stacking'. Overall, the large number of patents that are potentially relevant to a given product is the heart of the issue, although there are then subsidiary problems, including:

Picture: Ruslan Ivantsov/shutterstock.com

Terminology

We use the following terms throughout the article

Term	Meaning
Manufacturer's suggested retail price (MSRP)	This is the price at which the manufacturer recommends that a retailer sell the product
Technical standard	A technical standard is an established norm or requirement about a technical system that establishes uniform engineering or technical criteria, methods, processes and practices.
Standard-essential patent (SEP)	An SEP is a patent that claims an invention that must be used to comply with a technical standard

"The question of how much a company will have to pay for patent licences combines technology forecasting, market adoption estimation, predicting patent law evolution and dozens of other factors"

- the difficulty of identifying those relevant patents (finding the 250,000-plus patents on the product out of the millions of worldwide patents);
- the significant effort (and ambiguity) in disproving, or proving, infringement of a given patent; and
- the lack of transparency about the costs of licensing the practised patent.

All of these factors also tend to make the value of patents murkier and more uncertain than in, say, pharmaceuticals, where a handful of patents drive all of the value for prescription drugs. Thus, a lack of certainty is core to the royalty-stacking problem.

Consider a product company developing a model for a new product investment decision. Such models help businesses to decide whether a product investment might produce a positive return. However, in order to develop the model, the company is forced to make assumptions about future scenarios to calculate the expected return on investment (ROI) for the new product. These calculations are designed to model a number of risks: general economic risks (eg, recessions), technological risks (eg, whether the desired product can be produced by the company or its suppliers) and business risks (eg, whether the product can be successful in the market) – to name a few.

Unfortunately, quantifying the uncertainty around the patent component turns out to be far from straightforward – even for large, sophisticated companies with the resources to conduct patent reviews and where there are commitments by licensors to grant FRAND licences.

Companies are not simply burying their heads in the sand; figuring out answers is challenging because of the sheer number of technologies that must be used (standardised as well as non-standardised but equally expected by consumers) and the difficulty and complexity of finding good data. Thus, unable to model patent costs in their ROI calculations, companies are faced with tough decisions about whether to invest in developing new products. While widely accepted, credible models exist for the non-patent components – the same cannot be said for patent risks. Uncertainty kills investments and a company may either shift its investment to products with greater certainty or end up in the red for the product. Royalty stacking is at the root of this uncertainty and is problematic at both the product level and the individual technology level. The massive pool of patents creates an issue that cuts across a variety of industries and technologies, and which arises anywhere that highly integrated products are produced. This is most visible in the consumer electronics world, as an increasing amount of technology is concentrated in our smartphones (used as an example throughout this article), watches, wearables and smart TVs. However, the implications are being felt across broader industries than just consumer electronics; companies in the automotive, home appliances and home automation and control space are facing this dilemma as well.

Unfortunately, (F)RAND commitments for standards do not help as much as one might think – given the lack of universal definition of FRAND or RAND, and the fact that there is no such thing as a 'standard' royalty rate. Further, not everyone who has patents relevant to a technical standard has made a (F)RAND commitment (eg, non-participants to the standardisation process who happen to hold relevant patents). What is more, there are a number of commercially necessary technologies in devices that are not standardised.

Two examples of royalty stacking illuminate this problem.

Cellular royalty example

Given how long mobile phones have been on the market, one might think that the cost of licensing cellular SEPs would be common knowledge. However, published data suggests royalty rates ranging from about $14 to $54 for a $400 dollar smartphone (see Alexander Galetovic, *et al*, "A New Dataset on Mobile Phone Patent License Royalties: 1H 2017 Update," Working Paper WP16011, Hoover IP2 and Ann Armstrong, *et al*, "Surveying Royalty Demands for the Components Within Modern Smartphones," Working Paper, 2014).

This variation is massive and far too broad to enable forecasting. Even multiple, signed patent licences for cellular SEPs would not solve the issue. The court in *TCL v Ericsson* still found a very broad range (between $24 and $40) for a $400-ish dollar smartphone for 4G (see *TCL v Ericsson*, 2017 WL 6611635 (CD Cal 2017) (superseded and also overruled on other grounds, Fed Cir 2019)).

Most smartphone companies operate with minuscule profits and depending at which end of the spectrum the royalties fall, cellular SEP licence fees alone can swing a product from a profit to a loss.

HEVC royalty example

Technical standards with large patent-licensing pools can be challenging as well. Consider the patent landscape surrounding high-efficiency video coding

FIGURE 1. Patent uncertainty problem – HEVC

FIGURE 2. Patent uncertainty problem – many technologies

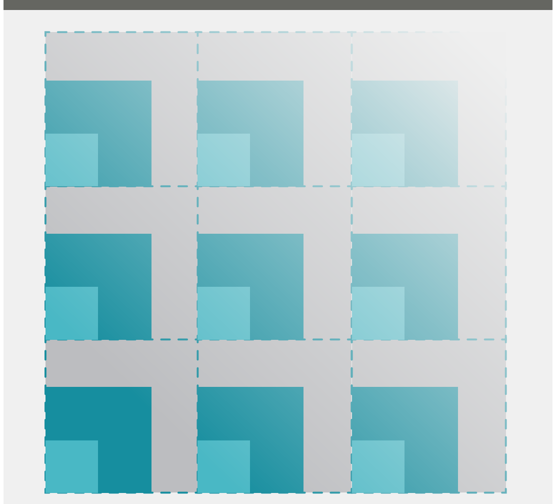

(HEVC or H.265). Looking at the membership of the pools compared to the adopted contributions to the technical standard reveals that there are several major companies (eg, Nokia, Microsoft and Interdigital) that account for 35% of the adopted technical contributions that are not affiliated with any pools (see Oliver, *et al*, "What Will TV Cost You? Putting a Price on HEVC Licences", *IAM* issue 89 (March 2018) – since this article was published, Huawei joined the HEVC-Advance pool in January 2020).

Further, there may be many owners of undeclared HEVC patents that could show up asking for licences. Figure 1 puts the pool asks ($1.60; light blue) in the context of other known licensors (dark blue) as well as the unknown licensors (grey).

This leads to the question: How much should a company model for HEVC patent licences?

Stacks of stacks

Cellular and HEVC SEPs are but two of the many technical standards that a product may need to incorporate – and the challenge does not stop there. Many non-standardised technologies in a smartphone (eg, the screen (LED, OLED), cameras, antennas and touch screens) may implicate thousands of patents. So the uncertainty – like the royalties – grows or 'stacks', as shown in Figure 2 (now with light blue representing known licence rates).

Smaller companies are even less likely to make costly capital investments without a clear line of sight to ROI; thus, the issue of royalty stacking is all the more severe for them. It is challenging to quantify these unmade or missing investments because it is rare for companies to announce rejected product development decisions. Instead, they simply invest their money elsewhere.

Building a mid-range smartphone

Consider a scenario where a consumer electronics company approaches a law firm for help modelling the patent risk as it embarks on building its first mid-range smartphone. It naively queries patent costs having heard about cellular patent fees. The law firm suggests creating a model for patent costs around the features that the company needs to include because when a high-tech product is faced with licensing tens of thousands of patents, royalty stacking becomes a serious challenge. To provide a more realistic view of what a mid-range smartphone would look like, we considered a TechInsights-provided teardown of a $450 smartphone. Using this as a reference point ensures that the analysis stays grounded in the economic realities of the handset business. Figure 3 shows some of the major components of the smartphone.

Consider that the company is planning for its phone to have similar capabilities. Given this teardown it will need:

- processor (64-bit, octa-core) and memory (LPPDR4X SDRAM);
- cellular (GSM/EDGE, CDMA and 4G/LTE) – note that 5G is not yet scoped;
- storage (NAND flash memory);
- front and rear cameras;
- touchscreen display (roughly 6.5" diagonal);
- lithium ion batteries – note that wireless charging is not yet scoped;
- sensors including accelerometer, electronic compass, light/proximity sensors and a biometric fingerprint sensor;
- additional wireless protocols:
  - WiFi;
  - Bluetooth;
  - satellite location positioning (GPS/GNSS); and
  - NFC; and
- USB-C including support for power delivery for quick charging.

> *"Cellular SEP licence fees alone can swing a product from a profit to a loss"*

The company can then cost these parts. The TechInsights data arrives at a total manufacturing cost, including assembly, of $250 – as shown in Table 1.

Given the $250 assembled cost and a target manufacturer's suggested retail price (MSRP) of $450, a high contribution margin – and high profit margin – seems to be available. However, the profit margin is not as rosy as it first appears. Figure 4 shows the economics for each unit of the smartphone. For simplicity, it has been assumed that this smartphone will be sold directly to retailers (eg, no distributor taking an additional cut) and no additional incentives to the retailer have been considered (eg, credits for activations).

FIGURE 3. Overview of the major components of a mid-range smartphone

Source: Drawing based on data in TechInsights teardown

Picture: EmBaSy/shutterstock.com

TABLE 1. Manufacturing cost breakdown (incl assembly)

Part	Cost
Integrated circuits	$120
Modules, discretes and connectors	$20
Substrates	$5
Component insertion	$7
Card test	$2
Battery sub-system	$10
Display sub-system	$33
Camera sub-system	$37
Fingerprint sensor sub-system	$3
Non-electronic parts	$11
Final assembly and test	$2
Total	**$250**

Source: Based on TechInsights teardown (numbers rounded)

FIGURE 4. Smartphone unit economics

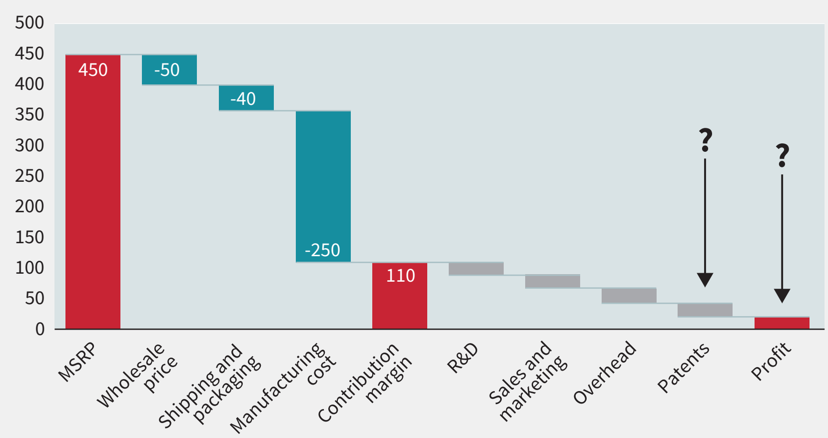

The specific amounts will vary but a typical price to retailers (excluding incentives) allows retailers between 5% and 15% of the MSRP as profit. Here it is modelled on the higher end at $50 (11%). Additionally, the cost of packaging and preparing the phone for shipment, as well as shipping can add another 5% to 15% of the MSRP. We modelled this at $40 (9%).

To the extent that this shipping and packaging estimate is high, this share could alternatively be attributed to retailer incentives and/or distributors, which were not modelled separately. Reliable data on these mark-ups and discounts are challenging to obtain because retailers often receive incentives as credits for sales, activations, advertising and/or hitting certain volume targets. Further discounts may be available in the form of time-to-pay credits (eg, retailers take inventory but do not have to pay for it for five days as opposed to 20 days).

Once the parts and assembly costs ($250) are deducted, that leaves $110 of contribution margin for each phone. In a recommended product design process, modelling would then calculate the estimated profit from the phone:

- How much R&D spend remains to complete the phone?
- What are the sales and marketing costs to encourage consumers to buy the phone?
- What portion of the company's general overhead is allocated to this project?
- How much will the company need to spend on patent licences?

Of these items, patent licences are the hardest to model or – to put it another way – the most uncertain. This presents significant challenges in deciding to make an investment because it blurs a company's ability to forecast margins (see the large question marks in Figure 4). Like most handset vendors, the company here has limited flexibility in pricing its product to consumers (and retailers). Similarly, the component costs are fairly well established, and it is unlikely that meaningful discounts can be eked out.

The contribution margin has to cover R&D, sales and marketing, general overheads and patent licence spend completely and leave room for profit. Of those items, the patent landscape is the most challenging to model. Failure to do so well could result in a drastic underestimation of actual royalties and leave the manufacturer in the red. Thus, the remainder is the estimated profit for each phone.

Industry data suggests that the actual profits to handset makers are quite low. For example, Counterpoint Research showed global handset profits in the third quarter of 2019 falling 11% year on year to $12 billion and, further, that the distribution of those profits was concentrated among the top two vendors. Since those are unlikely to represent the market, focusing on the profit of vendors ranked three to six in profit share is likely to provide a more accurate snapshot. Those four vendors are earning between $3 and $18 in average profit per phone. Many vendors are earning less (eg, anyone who is not profitable is not shown on the profit graph and is, on average, losing money on their handset). (See 'www.counterpointresearch.com/global-smartphone-share'.)

How can law firms help the company to model patent licence costs? A good first step is to figure out the technical standards and major technologies required (either technically or commercially) by the client's handset. To paint a realistic picture of what standardised components go into the phone, the teardown from TechInsights was reviewed in more detail. The technical standards and other major technologies identified are grouped in Table 2.

Pause and think about this: more than 30 standards and major technologies have been identified that are likely to be implemented in a single phone. Visualise Figure 2 with 30-plus boxes. That is a lot of uncertainty. And while not every one of these standards is technically required, almost all of them are commercially essential.

This is not just a smartphone-specific problem. Royalty stacking is an issue across a broader range of products, cropping up anywhere with a high volume of technologies and patents. It is little surprise, then, that this is where many companies give up due to the seemingly insurmountable level of complexity and uncertainty.

We have found that new tools coupled with a well-defined process can improve this challenging situation. This approach reflects how tools have improved progress by providing better data, although uncertainty and royalty stacking concerns will remain. It will work best in areas that have some public data about relevant patents and prices. In those areas, adopting the process depicted in Figure 5 can assist in converting portions of the patent cost uncertainty into a more manageable risk.

The top row of the process is one that should be iterated multiple times through the broader product development process. Step 1 is designed to identify the included technologies and technical standards. While it may seem obvious, we regularly encounter products on the verge of release to manufacturing where this has not been carried out. Step 2 is where a model for each item found in Step 1 is created. Finally, the product design or pricing can be adjusted in Step 3.

The sub-parts of Step 2, Steps 2a-c, are the focus here. In Step 2a, best efforts are expended to obtain patent lists and available pricing information. The seed patents identified in Step 2a feed into the landscaping tools used in Step 2b. Lastly, in Step 2c, the results of the landscape are combined with the pricing used to model the licence fees.

We will now explore Steps 2a-b in more detail for several technical standards and then demonstrate how those results can feed into Step 2c.

TABLE 2. Technical standards and major technologies

Memory & Storage	Communication	Audio/Video Codecs	Internet/Web Browsing	Other
Volatile LPDDR SDRAM **Non-volatile** Flash Multimedia card support	**Cellular** GSM/EDGE UMTS/WCDMA CDMA2000 LTE 5G (not costed) **Connectivity** WiFi Bluetooth USB-C GPS/GNSS NFC	AMR-NB AMR-WB MPEG-4 AAC Profile MPEG-4 HE AAC Profile AAC ELD H.264/AVC H.265/HEVC MPEG-4 SP VP8 VP9	W3C standards (HTML, CSS, etc) IETF standards (IMAP, MIME, Jabber, etc)	SMS RCS USB Power Delivery MPEG-DASH Qi (wireless charging, not costed) LED screen Biometrics Antennas Touch sensors

Source: TechInsights teardown and ROL experience

FIGURE 5. Patent licence risk modelling process

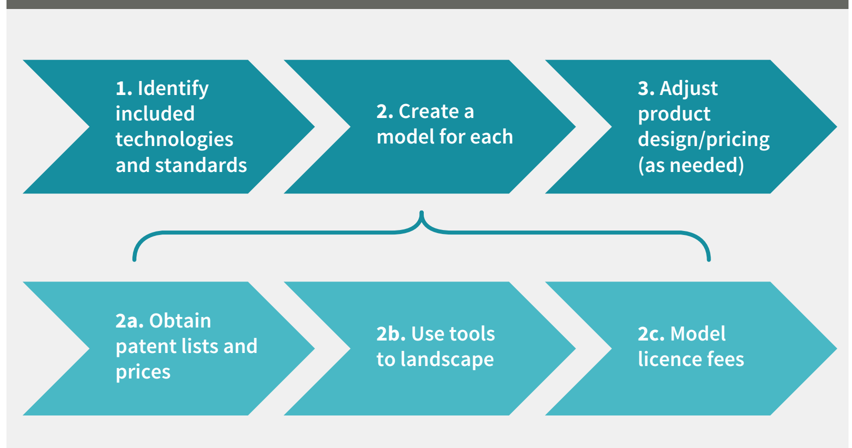

Arms around the stack – the role of landscaping tools

This section will explore to what degree automated patent landscaping tools can start to mitigate the challenge of a massive number of patents related to a single product. Three standardised technology areas within a smartphone will be used to explore the landscaping process. This will highlight both the benefits and some limitations. Overall, the tools are often helpful for identifying previously unknown patents (and owners) for a standard or technology. However, many standards bodies do not require disclosure of specific assets or limit access to the data. Both of these behaviours hamper the ability of companies to accurately assess the landscape.

To be comfortable with a patent risk assessment, a comprehensive overview of all of the patents relevant to each standard is necessary. As a practical matter, this has only been possible for a handful of companies willing to spend a fortune and with significant expertise in a given technology or standard – for a smartphone, this can be hundreds of thousands of dollars spent on patent specialists. Tools supporting machine-learning technologies are making access to this information more tractable, but there is still significant manual work to determine essentiality and infringement of the claims. But what if future tools could more reliably fill the landscape and make those determinations? This would be a huge win because a good landscape of relevant patents is key to modelling ranges for patent licence costs and, in turn, estimating profitability. That, in turn, will drive investment decisions. Thus, one can imagine future tools that do even more of the heavy lifting to reduce the issue of royalty stacking.

Commercial patent search tools primarily use keyword-style searching, which has been around for decades. However, only recently, with the advent of machine-learning tools, has it become possible to better automate the landscaping to identify candidate SEPs. With traditional methods, these tasks would have been unfeasible without significant manual labour. The newer generation of tools are able to cluster and identify similar patents from a seed list thus proving more useful. Nonetheless, even with the most advanced tools available, validity and infringement still require significant manual analysis on a case-by-case basis.

The three technologies considered here are:

- video and audio codecs;
- memory; and
- WiFi.

The two tools that we chose were the open source Automated Patent Landscaping tool (available at https://github.com/google/patents-public-data/tree/master/models/landscaping) and Cipher (https://cipher.ai/).

"The origin of the royalty-stacking problem is the huge quantity of patents relevant to integrated products (250,000-plus). Finding these patents is difficult"

Automated Patent Landscaping tool

The Automated Patent Landscaping tool is an open source project that aims to automatically construct a landscape of patents in a technology area (for more on the methodology, see Abood, *et al*, "Automated patent landscaping", *Artif Intell Law* (2018) 26:103). The source code is freely available on Github and can be installed onto any cloud computer environment to develop landscapes.

Cipher

Cipher is a commercial tool for patent analysis that uses machine learning to build classifiers. Companies can develop custom classifiers that help to group similar patents together. One feature of the classifiers is that they can be run across all patents and will help to identify similar patents. In addition, the team at Cipher has carried out custom work for some standards due to the unique challenges of identifying SEPs.

Brief comparison of the tools

Before digging into the analysis, it is worth briefly comparing the two tools used here and discussing the broader patent tool landscape. Cipher is a commercial software as a service (SaaS) tool designed for use by IP departments. Like many other SaaS applications, it is easy to get started and comes with professional support from the seller. An annual subscription is required.

The Automated Patent Landscaping tool is an open source project and requires a non-trivial amount of technical/programmer expertise to stand it up in a cloud computing environment. Further, once stood up, it is still a more programmer-oriented tool and lacks the more polished user interface of Cipher.

If your law firm or department has no significant technical resources, then Cipher is the way to go – despite the annual subscription fee. If on the other hand your company is willing to devote IT resources to standing up and maintaining it, the Automated Patent Landscape tool can deliver powerful results.

Other tools are also starting to emerge and compete in this space. Unified Patents' portal has OPAL, which is focused on three technical standards (HEVC, AVC and LTE cellular) and attempts to expand a list of patents and companies similar to what we have done here. At present, OPAL does not generalise to other technical standards or technologies (ie, you are limited to the three standards supported by Unified). In addition, Dolcera has focused on patents for cellular technical standards and has a commercial offering.

Landscaping

When it comes to the landscaping, our three chosen standards will highlight some of the challenges and thus are intentionally presented in a way to identify some of the current limits of the tools and data sources. In the audio-video codec segment, HEVC will be mapped quite extensively. However, switching to memory, the lack of data hampers the analysis. Meanwhile, for WiFi, the limited data through 2016 on declared SEPs enabled a partial analysis. Although these are examples only, they nicely illustrate the challenges that remain even when these tools are used.

Landscape 1: audio-video codecs

As highlighted in Table 2, approximately 10 audio-video codecs are common in a mid-range phone. Some of them (AMR-NB and AMR-WB) are cellular audio-encoding standards. Others are more commonly used for music (MPEG-4 AAC, MPEG-4 HE AAC and AAC ELD), while the remainder are used for video (H.264/AVC, H.265/HEVC, MPEG-4 SP, VP8 and VP9).

If a company is worried about patent licence fees and royalty stacking, why not just pick one from each category (eg, AMR-WB, MPEG-4 AAC and H.264/AVC only)? First, compared to the massive array of popular audio-video codecs and standards, the list of 10 standards is modest and some are mandatory (eg, AMR support is required for compliance with cellular standards).

Second, music and video format usage remains fragmented. Not only have content providers so far failed to agree on a single format but as new standards emerge, old files are rarely converted to newer formats. In the smartphone context, consumer demand necessitates that a device is capable of reading the most prevalent content.

Third, some companies will have much longer lists – we have seen one company planning on supporting over 100 audio/video codecs in its device. Thus, this is a fairly modest list of standards, but should provide good support for most of the websites and apps that a smartphone customer will use.

For simplicity in this discussion, the landscaping will focus on identifying other potential H.265/HEVC patent holders using the two tools. As an aside, there is an overlap of patents between standards (eg, a patent might be an SEP for H.265/HEVC, H.264/AVC and perhaps more); depending on whether a company can get a full licence to the patent, it may be paying for a licence to the same patent multiple times – a different type of royalty stacking.

Starting at the ending first, we estimate the total number of HEVC SEP families to be between 1,515 and 1,665 patent families – this is after some exemplary adjustments to the results of the tools. But how can the tools be used to reach these estimates?

Using the Automated Patent Landscaping tool, the list of known patents held by MPEG-LA licensors, HEVC Advance licensors and patents assigned to Velos Media can be input into the tools. Those lists provided the seed and enabled the identification of similar patents, although the results have not yet been reviewed by a subject matter expert to confirm essentiality – this more detailed essentiality testing would be conducted in later phases of the analysis. For now, this level of analysis enables a first-order modelling of who else might hold SEPs (other than the holders of the initially seeded patents).

Figure 6 shows that the initial results from the Automated Patent Landscaping tool were approximately 1,100 additional families. Removing families held by members of any of the three pools reduced the set to approximately 500 families. Removing families with expired and inactive patents reduced the number to 435. The top holders of those 435 families are shown in Table 3 alongside the Cipher results. Bear in mind that this result is likely to be both under-inclusive (missing some SEPs) and over-inclusive (includes some non-SEPs). In more detailed projects for clients, these tools can be used in multiple iterations. Further, more subject matter expertise is applied to analyse the results in client projects. How these results (and the ones from Cipher) can be adjusted is discussed below.

Working with Cipher, a similar type of seeding was used to look for HEVC patents. In addition to their standard classifiers, Cipher carried out a separate experiment to explore the ability to identify SEPs in HEVC. Cipher identified the patents declared to the public pools (in this case, MPEG-LA and HEVC Advance) and took those as the positives. To further construct the classifier, Cipher performed an iterative training process where it kept adding false positives (based on positive results appearing for declaring companies that are granted in the right date range and which were not listed by the pools) to the negative set until it was returning no false positives for the member companies in the pool. The classifier was further characterised by analysing the likely ratio of real-world positives to retuned positives from running the classifier against all video patents and extrapolating from the returned positives to get a likely number of SEPs. As with the results from the Automated Patent Landscaping tool, these have not yet been reviewed by a subject matter expert to confirm essentiality.

While the number of families found by the two tools differs, the high degree of overlap of companies found in the two lists should give some confidence that the tools are on the right path to finding the relevant missing patents. Table 3 uses bold text to identify companies not found by the other tool. (Note: Huawei is still listed in Table 3 because its assets are not yet listed by HEVC Advance; Qualcomm and Ericsson are listed because Velos does not make their asset lists public and so only assets directly held by Velos are publicly identifiable.)

In a full analysis process, this would be the point where the families identified by the tools are reviewed more closely by a subject matter expert to further winnow the list. Thus, both approaches – while different – identified several patent holders that we otherwise might have been unaware of in negotiations. We will discuss how this data can be used later. For now, let us turn to memory.

Landscape 2: memory

For memory, the situation is simpler in some respects, but harder in others. Only one type of memory (LPDDR SDRAM) will be used and that is standardised by JEDEC. Unfortunately, JEDEC does not list the SEPs on its website, nor does it make them available on request. Only JEDEC members can obtain

FIGURE 6. HEVC results summarised: automated landscaping tool

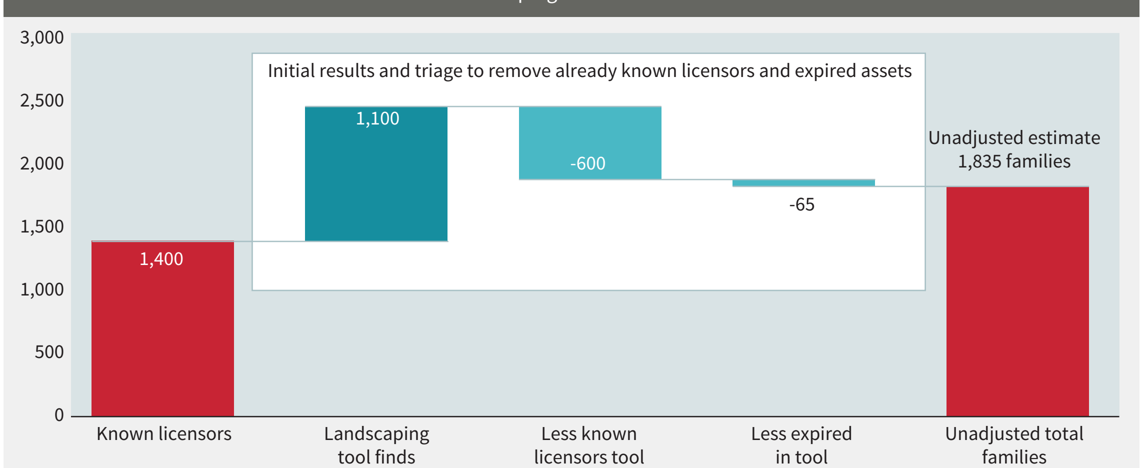

Source: ROL analysis of tool results

TABLE 3. Unaffiliated HEVC patent holders with at least two families

Automated Patent Landscaping tool

Position	Owners	Families with US grant
1	Qualcomm Inc	193
2	Sun Patent Trust	60
3	LG Electronics Inc	44
4	Google Inc	20
5	Microsoft Corporation	19
6	Huawei	14
7	Infobridge PTE Ltd	14
8	Blackberry Ltd	12
9	**Broadcom**	12
11	Interdigital Inc	7
13	**Arris**	5
10	Ericsson	4
12	Toshiba Corp	3
14	Nokia	2
15	Intel	2

Cipher

Position	Owners	Families with US grant
1	LG Electronics	62
2	Sun Patent Trust	46
3	Qualcomm	26
4	Infobridge	19
5	Huawei	17
6	InterDigital	15
7	Google	10
8	Microsoft	10
9	**Texas Instruments**	6
10	Toshiba	5
11	Nokia	5
12	Blackberry	4
13	Intel	2
14	Ericsson	2
15	Multiple entities with one family	1

Source: Cipher; Automated Patent Landscaping Tool

access to the list. This is an example of where lack of information exacerbates royalty stacking.

While not quite the same as the full landscape that we made for HEVC, Cipher's classifiers can be trained to identify DRAM-related patents. As a word of caution, the HEVC landscape was somewhat easier because it is well defined. DRAM, in contrast, presents a more complex case because the patents could be DRAM-related in many ways: the circuit itself, the protocol for operating the memory or manufacturing-related patents. This makes working through the landscape more challenging.

In this case, a more modest goal is to identify large patent holders of DRAM patents and help to de-risk the company's choice of DRAM vendor. In this case, this result will be highly over-inclusive of patents and patent holders. For simplicity, only the top holders by quantity are shown in Table 4. In addition, number 35, Rambus, appears due to its long association with patent licensing in the DRAM industry. For example, back in 2010, Rambus and NVIDIA signed a deal for a 1% royalty rate for some types of memory controllers and 2% for other types of memory controllers. More recently, Rambus' publicly announced rates with the memory manufacturers have had a single fixed fee and commentators have pegged the rates for memory at closer to 0.6%.

Despite being less refined than the HEVC landscape, this is an insightful list because most of the top companies are well associated with the memory industry. What if the smartphone company decided to buy memory from a no-name brand? It is possible to see the patent licensing risks from the companies listed more clearly. If the smartphone company needed to start clearing the risk in more detail, a subject matter expert could review these patents to filter down to a subset applicable to the specific LPDDR4 memory and corresponding memory controller used in the planned phone.

TABLE 4. Top holders of DRAM patents (via Cipher)

Position	Company	Number of DRAM families
1	Samsung Electronics Co Ltd	5,170
2	SK Hynix Inc	4,394
3	Intel Corporation	3,645
4	Micron Technology Inc.	2,972
5	Taiwan Semiconductor Manufacturing Co	2,313
6	International Business Machines Corp	2,304
7	GLOBALFOUNDRIES Inc	1,763
8	Renesas Electronics	1,626
9	Toshiba Corp	1,452
…	…	…
35	Rambus Inc	405

Source: Cipher

Landscape 3: WiFi

The last area to landscape in this example is WiFi. Standardisation of WiFi has taken place under the auspices of the Institute of Electrical and Electronics Engineers (IEEE). The various WiFi standards have historically been referred to by the 802.11 working group name together with letters for different generations (eg, 802.11b, now Wi-Fi 1; 1999, through 802.11ac, now Wi-Fi 5; 2014). Soon 802.11ax (now Wi-Fi 6) will be the most current generation. A significant amount of backward compatibility is commercially necessary (eg, it may be possible to drop Wi-Fi 1 support, but a device coming out tomorrow with Wi-Fi 6 support could not drop Wi-Fi 5 – or probably even 4 or 3 – support and be commercially successful). This is similar to mobile phones that need to include significant backwards compatibility to previous generations. Chips implementing WiFi can be bought at retail for as little as $5 and will often include support for other functionality (eg, Bluetooth (IEEE 802.15)).

WiFi patent licensing has been, and remains, a heavily litigated area with major lawsuits and vast disparities in awarded prices. That in turn has produced large disparities in the implied total cost of the stack for WiFi. For example, at the time of writing, news of Cal-Tech's award of $1.60 per phone against Apple and Broadcom had just been announced concerning patented technology in Broadom's WiFi chips (see *The California Institute of Technology v Broadcom Limited* (CD Cal, 2:16-cv-03714), Jury Verdict from 29 January 2020). Cal-Tech was not under a FRAND commitment; however, there are still concerns about the proportionality and fairness of the verdict relative to the licensing revenues awarded in other cases, the cost of the technical standard as a whole and how that might fit into the profitability of devices.

Before digging into the landscape, a quick recap of some recent, major WiFi patent royalty litigations may be helpful. Table 5 shows that there is a vast disparity in the per-patent rates and the implied WiFi stack rate across these cases.

This broad range underscores the importance of better modelling the scale and sources of patent risk for WiFi. One of the first challenges is that, while there are many IEEE declarations relating to WiFi patents, only approximately 300 entries listed specific, usable patent numbers (see Bekkers, R, Catalini, C, Martinelli, A & Simcoe, T (2012), "Disclosed Standard Essential Patents Database", Intellectual Property Disclosure in Standards Development, Proceedings from NBER conference on Standards, Patents & Innovation, Tucson (AZ), 20 and 21 January 2012; 2016 dataset used; hereinafter dSEP). Expanding on the point about "usable patent numbers", one practice at the IEEE is for a company to file a blanket FRAND pledge. This is good in that the

TABLE 5. Select WiFi patent verdicts

Case	Number of patents	RAND rate	Per patent rate	Implied WiFi stack	Stack as % of $5 chip
Innovatio	19 (three families)	$0.0956	$0.00503	$1.80	36%
Microsoft	24 US patents (five families)	$0.0347	$0.00145	$4.34	87%
Ericsson	Three patents (three families)	$0.1500	$0.05000	$150.00	3,000%
Cal-Tech	Three patents (one family)	Not RAND - $1.6000	$0.53333	$1,600.00	32,000%

Source: Armstrong et al (2014) page 29 supplemented with Cal-Tech verdict

TABLE 6. Top WiFi-related patent holders

Automated Patent Landscaping tool

Position	Owners	Number of families
1	Nokia Corp	31
2	Qualcomm Inc	27
3	AT&T Inc	21
4	LG Electronics Inc	18
5	Samsung	11
6	Ericsson	9
7	**Intellectual Ventures**	8
8	Apple Inc	8
9	InterDigital Inc	8
10	Huawei	8
11	Arris Enterprises LLC	7
12	Google Inc	6
13	Cisco Systems Inc	6
14	Blackberry Ltd	6
15	Broadcom	4

Cipher

Position	Owners	Number of families
1	Huawei	14,070
2	Qualcomm Inc	12,248
3	Ericsson	9,354
4	Samsung	7,589
5	LG Electronics Inc	7,298
6	ZTE Corp	7,041
7	Nokia Corp	5,558
8	Intel Corporation	3,700
9	NEC Corp	2,597
10	**China Mobile Comm Corp**	2,369
11	Sony Corp	2,217
12	Fujitsu Limited	2,165
13	NTT Docomo Inc	2,125
14	**OPPO**	2,074
15	ETRI	1,880

Sources: Cipher; Automated Patent Landscaping Tool; Bolding indicates that a company was nowhere to be found in the results of the other tool

FIGURE 7. Improved understanding of HEVC landscape

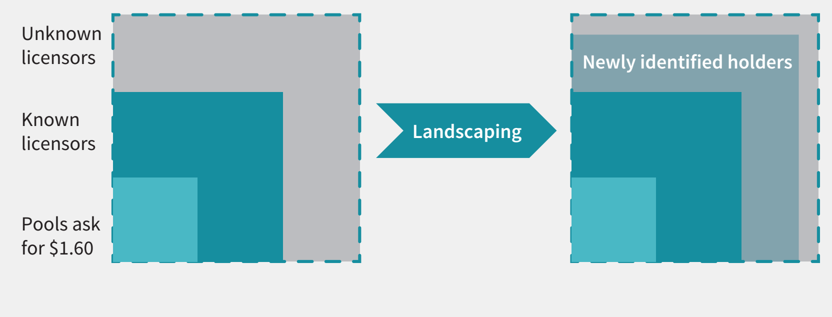

company will comply; however, it is bad for the royalty-stacking problem as it makes it challenging to identify relevant patents.

Unfortunately, most of the litigants in these cases did not list their specific asset numbers in their IEEE declarations. Also, the dSEP data set has no specific patents that were in many of the litigations listed above. To be fair, for some of the plaintiffs, the dSEP data does reflect a blanket FRAND patent declaration for them. Looking more closely at Wi-Fi 5 (current generation), the dSEP data lacks any specific patent numbers and even for Wi-Fi 4 (prior generation), only 56 patents are listed. It is possible that going directly to the IEEE patent declaration search might produce additional, more current data; however, the dSEP team performed significant work to clean up the disclosures and repeating that effort for post-2016 declarations is outside the scope of this article.

This landscape presents an odd situation that lands somewhere between the prior two cases. Some data is available, but it is quite limited. Patent owners have only identified 56 seed patents for the Wi-Fi 4 standard, presenting a high risk that the Automated Patent Landscaping tool may produce less useful results as compared to HEVC. Nonetheless, for exploration purposes, we used those listed patents as seeds and, in contrast, ran a more general WiFi-classifier from Cipher to get a sense of the companies in the space. The summary on a family level (after removing expired and lapsed patent families) for both tools is shown in Table 6 (company names in bold are those that appear in one list, but not the other).

The results here are a good example of what can happen without putting in the work to build the seed lists (Automated Patent Landscaping tool) and/or refine the classifiers (Cipher). The Cipher classifier was not designed to focus on SEPs but rather on general WiFi-adjacent patents and technologies. As a result, it massively overestimates the number of potential SEPs, with tens of thousands of results presented. Meanwhile, the Automated Patent Landscaping tool is suffering from too few input seeds to fully explore the landscape for SEPs and is reporting a woefully low number of relevant families (fewer than 200). Also, both tools failed to find some of the patents involved in the litigations in Table 5. Notably, neither found the Cal-Tech patent, which, if the $1.60 per device royalty holds up, could be disastrous for many of the smartphone companies at the lower end of the profitability range at $3 per phone.

With the bad news out of the way, consider the positives. Both tools highlighted a lot of the same companies and using both tools helped to find additional companies that one tool alone would not have found – as shown in bold. The list of companies is logically sensible (eg, the companies found are operating in the space or adjacent to it).

Depending on what a company plans to do and the stage of the product planning lifecycle, the list of top companies might be a useful stopping point. If on the other hand the goal is to build the detailed model (eg, the process of Figure 5), then it would be time to roll up your sleeves with technical specialists to further refine the seeds/classifier in order to get a more detailed and refined list. This example also highlights that this process is iterative. Steps 2a-c of Figure 5 can be repeated multiple times to get a more refined result.

Landscaping summary

Having walked through the general approach with three examples with differing levels of available information and different approaches to using tools, let us consider how the results could be used. Recall that we are focused on products faced with licensing possibly tens of thousands of patents and how tools can help to quantify the scope of the royalty-stacking problem and the issues that arise. We proposed a framework (Figure 5) and showed how Steps 2a-b could be applied with automated tools to better estimate the landscape in some cases.

The most detailed results were for HEVC. In this case, the massive array of potentially relevant patents was filtered down. Further, between 240 and 435 families that were previously not publicly listed were identified. Figure 7 visualises the impact of this landscaping work for the HEVC portion of the audio/video codec analysis.

Using the results

We ideally wanted to estimate the cost of patent licences for a product. We then scaled back to quantify the landscape by using tools to better understand it and thus estimate the royalty stack. In this section, we will explore how to use the landscaping results (Step 2c of Figure 5).

The origin of the royalty-stacking problem is the huge quantity of patents relevant to integrated products (250,000-plus). Finding these patents is difficult. However, the previous three examples show how tools can help to varying degrees by finding patents. The subsidiary issues remain: the significant effort (and ambiguity) of disproving, or proving, infringement of a given patent and the lack of price transparency to build good models for patent costs.

One specific use case for the information obtained from the tools is to negotiate with a licensor. A general consensus around proportionality has been emerging from worldwide case law around SEP licensing. Each case has its own nuances but taken together they stand for a principle of proportionality:

- *Microsoft v Motorola* and *TCL v Ericsson* in the United States;
- *Unwired Planet v Huawei* in the United Kingdom; and
- *Huawei v Samsung* in China.

The basic formula is shown in Figure 8 and can be put in simple terms: a single licensor's revenue should be generally proportional to the number of patents (or patent families) that it holds relative to everyone else holding patents (or patent families) in an area.

When looking at the formula, three questions should be asked:

- How many total SEPs are there?
- How many SEPs does the licensor have?
- What is the overall rate?

Taking these questions in order, let us look at how the landscaping results can help to build a better model.

How many total SEPs are there?
Continuing with HEVC, we can start with the pool data (MPEG-LA, HEVC and Velos), which provides approximately 1,400 families. The Cipher analysis identified about 240 additional families and the Automated Patent Landscaping identified about 435 additional families.

Thus, the first approximation is that there are in total between 1,640 and 1,835 HEVC families that are SEPs. This estimate accepts every asset licensed by the pools as both valid and standard essential while doing the same for the tool results.

Usually some adjustments are useful in modelling. One example set of proposed adjustments is to:

- subtract 5% for failures of the pool patents (eg, some are not going to be held valid and standard essential);
- add 10% to the results from both tools for missed patents (because there are false negatives, such as under-inclusion);
- subtract 10% from the results from Cipher (because they did more regressions to avoid false positives); and
- subtract 30% from the results from the Automated Patent Landscaping Tool (because there has been less work in this example to trim the list of false positives, for example over-inclusion).

Applying these adjustments creates a new estimate of between 1,515 and 1,665 patent families. For a fuller model, multiple adjustment ranges can be modelled. Additionally, a more detailed patent-by-patent analysis can be conducted.

How many SEPs does the licensor have?
Next, it is possible to look at particular licensors. For this example, consider Company X with 20 families. This is a straightforward look-up from the data output from the tools. If the number of total patent families is being adjusted, this number could be similarly adjusted – or held constant – to consider Company X in the most favourable light possible. Also, in actual negotiations with a company, the specific assets would be reviewed in more detail by subject matter experts for true essentiality, infringement and validity assessments.

What is the overall rate?
Finally, it is possible to model overall rates for the standard or develop a negotiating posture for one of the licensors. Both possibilities are explored here. Setting the price is always a touchy subject because to quote an HBO executive: "I've never met a writer who thought she was paid too much or a producer who thought they paid too little."

Let us start by thinking about a hypothetical price negotiation with Company X. Company X has proposed $0.10 per device for its HEVC patents. To give it the benefit of the doubt, the highest number of patent families found by the tools is ascribed to it (20) and the lower total number of SEPs is modelled (1,515). This would mean that Company X holds about 1.3% of the HEVC patent families.

How does Company X's rate compare to the pools in terms of proportionality? MPEG-LA is charging $0.20 for around 800 families or $0.02 per 100 families. By contrast, Company X would have the handset vendor paying $0.10 for 20 families or $0.50 per 100 families. Barring a significant miscalculation of the number of patent families held by Company X, this seems to be massively disproportionate – a great negotiation discussion point.

How about estimating the overall rate? Here again, models can be created. For example, the historical rates for prior video coding standards can be explored. In the case of AVC or H.264, an extremely large percentage of the patents were in the MPEG-LA pool. By estimating what percentage of all patents were in the pool and applying a multiple for the improvements of HEVC over AVC, some guideposts can be found.

For example, MPEG-LA's pool for AVC charged $0.20 and covered about 1,100 patent families. If that was 80% of the patents and HEVC is modelled as two times better, then the $0.20 for MPEG-LA can be scaled to an estimated FRAND rate for all of HEVC (eg, $0.50 total for all of HEVC). This is using AVC as the comparable. A lower number could also be modelled (eg, still 80% of the patents), but because of the number of substitutes (eg, AVC itself instead of HEVC, as well as AV1, VP8 or VP9 and Moore's law) the rate should in fact be lower (eg, 0.5 times adjustment: $0.13 total).

These numbers can be used to allocate money to individual licensors or pools. Continuing the Company X example with these numbers, with 20 families, the offer would be between $0.002 and $0.007 (ie, less

FIGURE 8. SEP formula

$$Overall\ rate \times \frac{\#\ SEPs\ of\ licensor}{Total\ \#\ SEPs}$$

Based on: *Microsoft v Motorola* (US); *TCL v Ericsson* (US); *Unwired Planet v Huawei* (UK); *Huawei v Samsung* (CN)

Action plan

- Your product design methodology should include a review of technical standards and major technologies to include or exclude.
  - Be thoughtful about which technical standards and technologies to include in your project. We often see unnecessary standards included because they are perceived by engineering to be free.
  - Consider whether older standards can achieve the purpose (eg, do you need H.265/HEVC in your product or would H.264/AVC work just as well at a lower cost?).
- Use modern machine-learning tools (some are open source) to better identify the universe of patent holders.
- Build a model of your expected patent licence costs to better ensure profitability.
- Accept that even the best models will have some variability and downside risks – it is not possible to remove all uncertainty.

than a penny) per device. Compared to the $0.10 that Company X is asking, that is a 100-times negotiating gap to close: a problem outside the scope of this article.

Modelling risk

This article shows how royalty stacking, the large number of patents potentially relevant to a given product, is a significant issue. Specifically, if a company produces products that integrate tens of thousands of patents, this overwhelming number of patents leads to difficulty in identifying the relevant patents and the need for significant effort to prove or disprove infringement, and is exacerbated by a lack of transparency about the costs associated with technologies and standards.

We used the lens of planning a new smartphone to explore these challenges and make the problem concrete. We then showed how newer patent landscaping tools can address some of these challenges and help to chip away at the royalty-stacking problem by building an automated landscape from limited data. The open source Automated Patent Landscaping tool and Cipher are two examples of tools to support these workflows. The heavy lifting of analysing the patents found by the tools remains daunting. Nonetheless, as shown, the landscape results can assist in building data-driven licensing models, despite opaque and contentious pricing issues.

In summary, this article has explored different parts of the issue of royalty stacking by providing:

- a clearer (but imperfect) picture of who holds patents across the multitude of technical standards and major technologies in a product in a manageable timeframe;
- more practical estimates of royalties based on the case law trend towards proportionality;
- steps that can be taken to ensure that the model includes a higher percentage of all relevant patent holders – this reduces the chances that our products end up being unprofitable; and
- insight into data-driven tools, which can improve negotiations with patent holders

Like any other approach, this one is not perfect, but perfection is not the objective here. It is impossible to eliminate all uncertainty; however, the more the uncertainty can be converted into a modellable risk, the better. Further, by establishing some ranges, the uncertainty is reduced and higher-quality business investment decisions are possible. iam

Erik Oliver and **Kent Richardson** are partners and **Hannes Forssberg Malm** is an associate at Richardson Oliver Law Group LLP

The authors would like to thank TechInsights for providing the teardown used as the basis of the example in this article; and Cipher for its analysis work on the three technical standards.