ATDI

■ White Papers

On Spectrum Usage Rights (SURs)

Enter

Author:
John Berry BA BSc MBA
CEng FIET FCMI

January 2008

Solutions in Radiocommunications

On Spectrum Usage Rights (SURs)

Purpose &Target Readers Overview:

This White Paper explores the concepts of Spectrum Usage Rights (SURs) as a means of managing interference experienced by a victim spectrum user as the result of actions by one or more culprit spectrum users. The concept applies equally to those operating both transmitter-receivers and receivers only. The paper initially defines terminology, identifying areas both included and omitted in current debate. It then considers practical examples where SURs are used as a method of control, commenting on their use. Finally it suggests a method of proceeding when considering engagement with the spectrum market that might use SURs as a condition of any spectrum access. This paper has been written with a would-be spectrum market player in mind although may also be useful to spectrum regulators. It does assume a familiarity with current SUR debate though someone with overall wireless networks experience will also understand the arguments.

Section 1 covers the theory behind the development of SURs as a means of defining rights and duties of spectrum use. This section debates the various elements raising questions where there is ambiguity. Section 2 discusses the development of SURs and how they are applied. Section 3 shows practical SURs and how they would be applied to practical networks by illustrating the results of some example modelling. Finally Section 4 discusses how would-be spectrum market players should use and exploit SURs to achieve their business aims.

Introduction

The following is a set of definitions that need to be explained, clarified and argued in order that the background to SURs is understood. Each concept is defined, in addition to explanations of subsidiary arguments and gaps in current thinking.

1.1 Interference. The ITU definition of interference can be paraphrased as *the effect upon the service of a radiocommunications system resulting from unwanted signals from one or more other systems.* This definition is critical to our understanding of SURs and eventually on the role of the regulator as market maker. This definition implies that for an unwanted signal to be considered an interferer, it must have some effect on the victim. The ITU also define 'harmful interference' as that which *denies the victim due service from his radiocommunications system.* So before the victim can claim that he is being interfered with by some culprit, he must first define his service because interference is relative to that.

1.2 Service & service quality. The service enjoyed by a victim[1] is a commercial choice that he makes. He could run a high availability, low error rate service such as video for which he demands a premium. Equally he may elect to run a low availability, high error rate service for which he charges little. The quoting here of error rate and availability serve to describe some of the greyscale of metrics upon which this choice may be made and is neither complete nor exhaustive. In order to generate an appropriate minimum signal to equipment noise ratio[2] in the victim receiver, the higher availability, low error rate services demand, for the same information and system characteristics, higher signal into the victim receiver to overcome fading. This choice is made finally by considering market economics but it places demands on the engineering.

1.3 Denial of service. Sub paragraph 1.2 considered signal to equipment noise as the threshold of service. Noise also comes from unwanted signals and hence the real ratio that is important is the signal to noise plus interference ratio[3]. For there to be no denial of service, a ratio must be found that avoids catastrophic effect on the victim; but this depends on his choice of service! Similarly the victim may choose to deploy robust, expensive equipment that provides a robust defence against interference giving lower signal to noise plus interference or he could deploy cheaper equipment that needs much more protection. This debate shows that limits to unwanted signals set to avoid harmful interference depend on economic choice on the part of the victim. This is the first anomaly in some thinking on interference and in the claim that SURs can be set up to be service and technology neutral. It is of course why there have been no arguments between operators under the command and control regime; all the users have so far been using the same service and where there was cross-service sharing, regulators adjudicated.

1.4 Service protection. Signal propagation follows stochastic processes and hence unwanted signals arriving at a point can only be described statistically. Variation in unwanted signal occurs over both locations and time and hence this ratio discussed above must be specified by defining the unwanted signal not exceeded for a given percentage time and locations. This concept of *'ratio maintained for…'* is defined as protection and hence the full specification for unwanted signals is that the wanted service at the victim is protected for a percentage of time and over a percentage locations. Signals propagating from culprit to victim within the horizon exhibit a small variance in time and hence yield a small standard deviation about a median value. Indeed the distribution is so narrow when compared to other effects that it is often ignored. Within the horizon, they do however exhibit significant location variations and hence yield a larger standard deviation and this becomes the major effect allowed for in providing protection.

1.5 Protection over time. In beyond the horizon cases, radio signals are not stopped by terrain or the Earth's bulge. All that happens is that median signal levels reduce as a result of the high excess losses the terrain and/or Earth's bulge inflict. Radio waves propagating in the normal atmosphere are bent slightly towards the surface of the Earth. The result is that the radio horizon is slightly extended and radio engineers talk of applying an 'effective Earth radius' to simulate the effect while still being able to represent the direct ray between transmitter and receiver as a straight line in the vertical plane. This is the case for up to 50% of the time. For small percentages of time, this beam bending is exaggerated such that the Earth is effectively flattened and the *effective Earth radius* is massively increased; so called tropospheric propagation. The excess losses are hugely reduced and the signals received are hugely increased. The result is that for small percentages of time (between 0.01% and 10%), culprits that caused no catastrophic effect at the victim considering propagation in the normal atmosphere now breach the threshold signal to noise plus interference ratio and do interfere. Some choose to ignore interference for small percentages of time, concentrating on spectrum sharing between licensees

1. The term 'victim' is applied to any spectrum user whose service might be interfered. It is not constrained to those who will definitely be interfered because interference remains to be proven though modelling or ultimately though network implementation. Likewise the term 'culprit' is applied equally to a would-be or actual interferer.
2. Signal to noise has been used here as a proxy for all of the various means of expressing the ratio between the energy from the wanted signal over the energy from noise sources. In the CDMA world for example, this is cited as the Eb/N0, the energy per bit of information over the normalised noise. In every technology, some limit exists where the wanted information can no longer be discerned from the noise and hence service is no longer available.
3. Signal to noise plus interference has been used here as a proxy for all the various means of expressing the ratio between the energy from the wanted signal and the sum of the noise from natural or man-made sources plus the noise from unwanted signals from spectrum sharers. In the CDMA world for example this is cited as the Ec/I0, the energy per chip over the normalised interference (though noise has been omitted). In every technology, some limit exists where the wanted signal can no longer be discerned from the noise plus the aggregate of interference and then service is denied.

within the horizon, but this does not negate the effect. Most operators would agree that denial of service, even for very small percentage of time, is too much so the effect must always be considered.

1.6 Channels & blocks. The unit of frequency is applied to two dimensions in spectrum use: to the width of a channel in which (simplistically) the transmitted energy is contained and to the width of a block. An operator will acquire spectrum access rights to a block though an auction or other market activity. A block may accommodate one or more channels. Once acquired, the block may, in principle, be broken up into any set of channels of any width. Alternatively, channel dimensions may be set by a regulator. Two blocks can be combined to provide the traditional uplink/downlink frequency division duplex operation or a single block can make use of time domain separation of uplink and downlink. The arrangement is shown below in Figure 1.

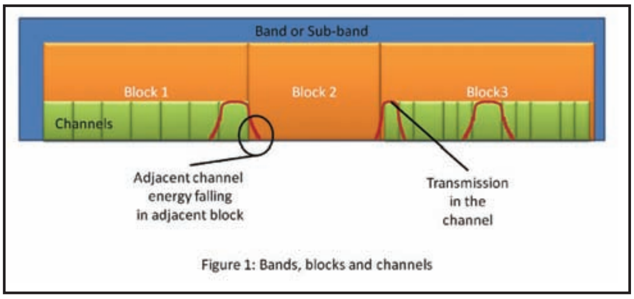

Figure 1: Bands, blocks and channels

1.7 Co-channel & adjacent channel effects. As noted above, the energy transmitted occupies a channel. In reality, the generation of signals in a channel gives rise to spurious energy in the adjacent channels and this is shown in Figure 1. While the intensity of the energy in the adjacent channels is much reduced over the main body of energy within the wanted channel, it is nonetheless significant and must be accounted for when computing interference. Similarly receivers are not perfect and are responsive to energy on channels adjacent to that carrying the wanted information. In theory, better receiver filtering can yield better rejection to this energy however the current state of the art constrains what is possible: size and cost are major determinants in achieving improvement. Unwanted energy therefore comes from other users co-channel to the victim but also from the adjacent channels (both from receiver and transmitter imperfections). All three mechanisms must be considered when considering the interference scenario existing at a point in time and space.

1.8 Differing channel bandwidths. As already established, a block can be divided up into multiple channels. Any spectrum users can choose to employ different technologies with different channel sizes. This gives rise to a complex set of overlaps where narrow channel bandwidth technologies share with wide channel bandwidth technologies as shown in Figure 2. Interference as defined above arises when the energy in the channel supporting information transmission is victim to the presence of unwanted energy – but this unwanted energy may be a plurality of narrow channel bandwidth culprits or indeed may be a small part of a wider channel. The idea of bandwidth ratio between victim and culprit channel must be considered when deciding whether or not interference is occurring. The same principles and mechanisms apply to the block; one operator may have rights to one block with others occupying the adjacent blocks. The technology-neutral situation in is far from simple as that often shown where the channels and blocks are shown with equal width.

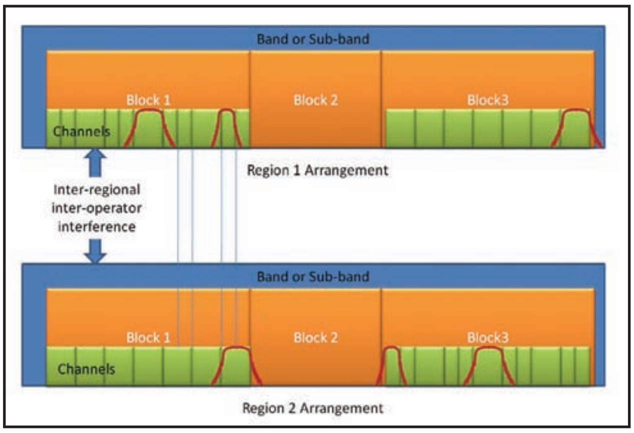

Figure 2: Mixed channelling between regions and operators

1.9 Separation. As already established, a block can be divided up into muWhether a culprit does or does not breach the signal to noise plus interference ratio depends on the electrical separation in decibels between him and his victim. Separation can come from separation in frequency or in space (geographically). The SUR is really intended to manage the geographical separation although, in its assessment of the adjacent channels, it does also consider frequency separation and the imperfections of transmitters and receivers. One can also consider frequency spacing by indicating aggregate noise powers in the wanted channel from all licensees in the spectrum around, describing this as the indicative interference level (IIL) or spectrum quality benchmark (SQB). Following on, however, from the argument about wide band/narrow band, the separation scenario depends on the relative channel bandwidths of the two operators and subsequently on the economic choices taken by him in his use of his spectrum block. Again analysis is simple when the channels are the same width supporting like services. The only good thing is that science already provides all of the modelling methods needed to handle the enhanced complexity. The main unknowns are in the station information.

1.10 Power spectral density. While aiming to control the unwanted signal level incident upon the victim receiver, the regulator could work backwards from the limit signal to noise plus interference ratio at the victim to be protected to compute a power spectral density transmitted from the culprit. Power spectral density describes the signal power, spread across the channel bandwidth, as it 'leaves' the culprit antenna. The advantage of this is that it is simply measured at the culprit station. The disadvantage is that it only measures the power at one station and hence does not account for the density of stations – deployment topology is an important factor in interference because signals aggregate at the victim and multiple low power culprits can aggregate to give the same effect as a single very high power culprit. Power spectral density is useful but does not therefore present the whole story.

1.11 Power flux density. Taking the argument about aggregation to its limit, what we need to measure is the effect of a plurality of culprits directly on the victim. However, in making investment appraisal to decide how high to bid in a spectrum auction, for example, one needs to have an assessment of likely interference for each chosen network optional technology, topology and service type that might be deployed. Data on the degrees of freedom that might be available from a particular option against a given SUR needs to be known before a network is designed and not after. We therefore need a proxy for the victim stations against which assessment could be made. In deciding on this proxy, regulators have employed methods developed from previous spectrum management regimes. A virtual perimeter is developed around the would-be licence area. Test points are set out around this perimeter and a limit signal to noise plus interference ratio is developed assuming a test receiver at each test point. The test points are proxies for the real stations and the limit value is scaled appropriately to allow for the distance error between would-be stations and perimeter. In this case the measurement is power flux density (PFD) – signal level spread across a bandwidth measured at a point in space. This limit power flux density at the test point depends on the technology and service and hence is not technology independent. The regulator plays market maker by setting possibilities.

1.12 Adjacent channel PFD. To complete the picture, consideration needs to be made for signals from and into the adjacent channels as discussed above. The signal levels are much reduced (by receiver and transmitter filtering) and hence the distance between victim and culprit is much reduced. In this case the PFD limit is assessed close around the victim and culprit stations. Where these are not yet planned a construct can again be used.

1.13 Measurement versus modelling. All of the discussion so far implied that levels are absolute and existent. Even if varying and describable only in statistical terms, the assumption is that they can be measured. That leads the discussion to determine exactly how these levels will be determined. If the network has yet to be deployed and the block is up for auction, how does one determine the levels? There are only two methods: measurement or modelling. Measurement could be employed universally. Method could be set from previous measurement and then actual levels could be assessed by measurement after system implementation. Test stations could be installed to represent new networks though they would have to be in place for some time (for months if not years) to allow statistically valid readings to be taken. Any breaches would be resolved by negotiation using the measurements as evidence in court. In the end, victims could demand that culprits remove offending stations. The current state of art also allows the whole affair to be modelled. The state of the spectrum can be modelled with pre-auction deployments. Then the blocks for auction can be modelled with the market maker speculating on technology, topology and service post-auction. The interference potential of the blocks can be determined and hence some judgement about their value can then be made and the associated SURs set needed to effect protection. Modelling has huge advantages over measurements. Throughout science, measurements act to inform models and this has been true in spectrum management building a 50-year body of knowledge.

1.14 The model used. Modelling requires that some algorithm is used that reflects on and use of the body of knowledge to make predictions about both wanted and unwanted signal propagation. This model comprises of three parts: the environmental model, the terminals model and the propagation model as shown in Figure 3. The picture of three models is often overly simplified to focus on the propagation model alone. The first element of the model, the environment, generally comprises of a digital terrain model and a buildings and vegetation model. The fidelity and accuracy of this model varies and must be chosen to reflect the results and accuracy of modelling sought. The terminals model defines the terminal equipments likely to be used. The choice of parameters requires the regulator as market-maker to select technology, topology and service. This might be done by determining the most likely or the most extreme or the most conservative block usage. Finally, a suitable propagation model must be selected from the body of knowledge. There are many. Choice must be made considering the results and accuracy needed and also its applicability considering the terminal model. Once these three model parts have been set, the interference state can be modelled and repeated time and time again to yield a family of results that make whatever argument is needed: setting levels in a particular SUR, proving compliance with an SUR and the like.

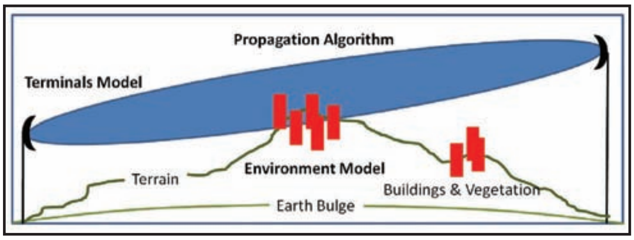

Figure 3: Three elements of the model: environment, terminals and propagation

1.15 Modelling accuracy. The model discussed above (as defined by the environment, terminals and propagation algorithm) has a definable accuracy. One can define this accuracy by referencing the body of knowledge on measurement or undertaking new measurements though these may have to be taken over some significant time period to make the results relevant. The idea of the modelling having a definable accuracy means that evidence from the modelling (in support of a claim of an SUR breach, for example) will only suggest a result to a calculable confidence. There is nothing certain, just a greyscale of confidence. There are various claims made that, provided that the same propagation prediction model is used for both SUR setting and SUR compliance, relativism will remove error. This is a dangerous claim since significant care (in setting up the whole model) would be needed for it to be usable – even the implementation of the propagation algorithm itself is often open to interpretation.
Though apparently simple, the SUR is a complex metric. Understanding SURs and how they have been set is key to being able to exploit any spectrum asset.

2. Constructing & Using SURs

2.1 History. The concepts behind SURs are not new. Under the old command and control spectrum management regime, the regulator decided all – service, technology and even topology – and certainly approved all in order to assure players of protection. As liberalisation bit the regulators were motivated to allot blocks of spectrum to operators for a particular use. The regulator then set the technology but not the topology and service. This may have been achieved by beauty contest, selecting the operator that best met the regulator's idea of the market's service requirements. The licensee had therefore some degree of freedom to deploy stations at locations and of a density that he required to serve his market. In providing this freedom, the regulator had to act some other way to protect others from interference. The argument above showed that it is possible to project this protection as a PFD limit around either the victim or the culprit either as an area or as a line. Projecting a line or polygon around a victim aimed to protect the victim system regardless of what was deployed inside and recognised the rights of the victim. Projecting a polygon around the culprit identified that provided the PFD limit was met at the polygon, the culprit operator was free to do what he pleased to meet his market and aimed to recognise the rights of all players. The polygon had either test points at the vertices or had test points at regular distances along the line. Polygons could be separated by many hundreds of kilometres or could tessellate or even overlap – the principle remained the same. The broadcast world harnessed this method of allotment planning in recent DVB and DAB spectrum administration.

Sharing spectrum (separated in space or frequency) with others using different technologies and services is likewise not new. In the past, all that happened was that the regulator selected all the services and technologies that could potentially share and subjected these to tests to determine the limit signal to noise plus interference ratios. Once complete the regulator published tables of ratios, technology versus technology, and service versus service for both co-channel and adjacent channel. These tables meant that if one was contemplating sharing, the published ratios should be used to determine whether or not the PFD at the polygon vertex or segment breached licence conditions. Clearly the difference under this regime was that the regulator decided the permissible combinations for spectrum sharing; ITU, ERO and CEPT spectrum management publications are full of these tables.

Modern spectrum management demands that regulators move to the next level of thinking: that the markets rather than the regulator decides spectrum use. The regulator cannot escape the physics though. Someone still has to determine the ratios. All that is different today is that the regulator undertakes suitable market studies to find out the most likely use of the spectrum (say) and does the calculations on this basis. That done, the licensee is given *'complete'* freedom to deploy what he wants provided that the constraints are met. He has *'spectrum usage rights'* provided that he shows compliance against the limits. These limits are calculated in exactly the same way, assuming combinations of technologies, services and topologies sharing. A single limit value is arrived at by the regulator which in his mind provides optimal spectrum use. The sorts of decisions that are then available to the licensee are:

- Low density/high power stations versus high density/low power
- High density narrow channel bandwidth versus low density wide bandwidth
- Low density/high elevation versus high density/low elevation
- Multi-transmitter simultaneous broadcast versus specific transmitter local broadcast.

Whatever option, the PFD limit remains the same.

2.2 Co-block or Overlapping Block Sharing: Construction. For the purposes of this argument let us consider the instance where a spectrum block is to be auctioned. It is free of other licensees in a specific geographic area but there are other users in an area some distance away using the same block that are to be protected. The aim is to determine a PFD that will protect the distant block while granting the auction winner certain spectrum usage rights. Those rights determine the spectrum value and need to be determined prior to bidding. The SUR PFD limit is arrived at by identifying all the likely victims and their likely deployments. The scenario is shown below.

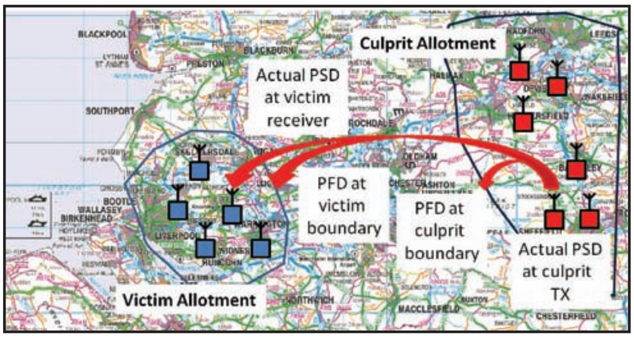

Figure 4: The various PSD and PDF relationships

The service, technology and topology of the victim is known or assumed and stations representing this system are placed in the modelling environment. The service area of the spectrum to be auctioned is set out, sourced perhaps from market studies. The auction winner, or culprit in this analysis, is to become the owner of an allotment comprising a spectrum block applied over a geographic area. Then a typical culprit system is placed within the service area. The service and technology at the victim determines the limit signal to noise plus interference that is tolerable to avoid catastrophic harm to the system. The terminal characteristics are set for both victim and culprit and a suitable propagation prediction model is selected. Then the modelling tests if the limit signal to noise plus interference is breached. The culprit station characteristics, their service and topology can then be adjusted to critically achieve the limit at the victim. Only then can the predicted PFD limit at the allotment or service area boundary can be read off. This is illustrated in Figure 4. The regulator can then change the technology, service and topology, simulating the most likely combinations. The interference condition can be tested and re-tested and the final PFD at the allotment boundary arrived at.

The PFD is defined in the following terms:
The aggregate PFD at a height of H metres above ground level should not exceed XdBW/m²/MHz at more than Z% of locations at the boundary. Watts per square metre is the standard measure of power flux – a way of assessing signal levels without having to specify the measurement antenna. When assessed across a measurement bandwidth it allows for aggregation of narrow and wide channel bandwidth unwanted signals inside the blocks. Aggregation can be carried out using one of several methods but in the simplest form this is by power summation. The boundary is defined by a polygon or test points. This SUR specification ignores time variation and hence is incomplete for sharers beyond the horizon.

2.3 Adjacent Channel Sharing: Construction. As the culprit and victim move closer (including the case where the two service areas overlap and co-block sharing is no-longer appropriate because interference would be extremely high), two other mechanisms must be controlled. The first is interference on the victim as a result of spurious emissions in the adjacent block of the culprit that appear co-block to the victim. This is caused by imperfect adjacent channel response in the culprit transmitters close to the block edge and it is right that this is constrained. The second is the converse: the interference on the victim out of block as a result of the culprit's co-channel emissions within the block. This latter mechanism depends on the victim's receiver and as a result its relevance must be questioned since the culprit is being penalised for something he cannot control. On this basis the former mechanism alone will be considered further.

Note the careful use of channel and block. In modern spectrum management blocks are auctioned. The operator then breaks the block into channels on which transmitters operate. These transmitters have adjacent channel spurious emissions occurring for many channels beyond the operating channel. The signal level received at any would-be victim will depend on how close the transmitter is to the block edge and the transmitter filtering that would form the spectrum mask pivotal in determining the energy on adjacent channels. Some erroneously refer to this channel/ block arrangement as a band and then discuss in-band and out-of-band operation. Great care is needed here and a clear definition of block and channel is important.

In the case where limits are set for adjacent block levels out of culprit transmitters, the procedure is not dissimilar to that used for geographically separated culprit and victim. The service, technology and topology of the victim is known or assumed and stations representing the victim are placed on the modelling environmental model. The service area of the spectrum to be auctioned is not in this case relevant because it is assumed to be the same as that of the victim. The only difference between victim and culprit is that they occupy adjacent blocks. As before we need to set up a signal to noise plus interference limit that is tolerable at the victim and, as previously, this can be converted to a PFD to allow an aggregate noise level to be computed as a result of adjacent block radiation from the culprit stations. This can be done simply by developing test points as proxy for the victim's receivers with assessment limited to within his licensed block. These test points can be deployed across the modelling environment at each point on the digital terrain model. There are two conditions that need to be understood: the first is that the regulator must set the adjacent channel characteristics for the victim receiver assuming technology, service and topology and the second is that it is unreasonable to expect no interference at terrain points that coincide with or are in very close proximity to the culprit transmitters. This latter point means that there will be areas around established victim receivers where planned culprit transmitters cannot go. As before the limit PFD is arrived at using the assumptions about most likely technology, service and topology to iterate a value that would just achieve the limit signal to noise plus interference ratio and hence avoid harmful interference at the victim receivers.

This time the PFD is defined slightly differently:
The aggregate out-of-band PFD at a height H metres above ground level should not exceed X dBW/m²/MHz at more than Z% of locations in a test area. The test area is defined in the SUR and is set to permit a percentage of points (where points are perhaps coincident with the matrix of the terrain model) to be reported where the PFD is breached. Test areas of tens of kilometres by tens of kilometres are discussed. Some breach must be permissible around victim receivers but otherwise the breach will depend, in this argument, on the system proposed by the culprit or auction winner. The higher the density of culprit stations, the more a breach is likely and so just as before, the culprit operator is faced with choices in technology, service and topology in order to comply with the SUR constraints. As before there is no consideration for time percentage but since both culprit and victim are within one another's horizon, this can be ignored as argued above.

2.4 Choice of Propagation Model. The propagation model comprises of three elements: the environmental model, the propagation prediction algorithm and the terminals model. Choices often focus on the propagation prediction model alone and this is an over-simplification that omits significant sources of error.

In making choices, what is of prime importance is the accuracy sought. There is a trade off between this and the model elements available. Accuracy is often described by the error achieved in modelling between predicted levels and those levels actually achieved in practice once the various networks are actually deployed. Error is not a single value but is described by a distribution of errors about a median. In describing the error we need to say something about the average error (between measured and predicted) and the width of the distribution described by its standard deviation (of error). Accuracy has a lot to do with the confidence that we seek; the higher the accuracy the greater the confidence (that the modelled results approach reality) but also the greater the confidence interval (or margin) the higher the confidence for a given accuracy[4]. This last statement shows that provided we are prepared to make use of confidence margins, absolute accuracy is not needed. The use of wider confidence margins can, if the accuracy is better than expected, lead to over protection of the victim but that is simply a facet of statistics.

What sort of accuracy might be useful in this victim-culprit relationship? This is a huge question. It has to do with the confidence and also with the accuracies possible in today's state of art. If we were planning a cellular network (where all paths between base station and mobile were within the horizon but between about 5km and 50km in length) and had carried out some optimisation of the propagation prediction model through drive testing, we might see an average error of better than 2dB and a standard deviation of error better than 7dB. Typically the terminal model is known accurately since the technology is today's and is supplied by the cellular equipment vendor. For the environmental model we might use a digital terrain model resolution of between 10 and 50 metres in the horizontal with an RMS

4. Harper, WM (1982), Statistics, The M&E Handbook Series, MacDonald and Evans, Gateshead, UK, pp 164-169.

error in the vertical of better than a few metres. We would also typically use a clutter model that describes the degree of urbanisation and vegetation on the Earth's surface with perhaps five urbanisation categories. Each category has an attributed representative height. The propagation prediction algorithm could be one of many though models based on Recommendation ITU-R P.526 with empirical excess losses for urbanisation or ATDI's Open COST 231 offers the best accuracy for mobile technology in the 100-3,000MHz range for these path lengths. This leads us to a so-called planning margin of 9dB for 90% confidence (1.25 times standard deviation). This is one benchmark.

This argument leads to 9dB margin for 90% confidence and does assume an exercise in model optimisation. If path testing is not possible (as will be the general case in spectrum management) then *'raw'* prediction algorithms that model the physics of the path must be chosen over those making use of empirical formulae. In this case the average error is likely to widen to say 5dB with the standard deviation of error of better than 10dB for the model described. This in turn leads to a margin for 90% confidence of 13dB. Were we to model interference arriving from culprits beyond the horizon, the average error might creep up to say 8dB with a standard deviation of around 15dB. The corresponding margin would be around 20dB. This illustrates that when modelling, one needs to consider the accuracy available and the confidence margin to be added to calculations made with the model. There is no point in being only 50% confident that Operator A is interfering with Operator B thereby adding no margin; higher confidence is needed before action is taken. This argument illustrates that there is a legal context to this subject. If one is only 50% confident of a breach in SUR conditions, is that sufficient to mount a court action to have Operator A remove his equipment? There are other regulatory models worthy of mentioning. Recommendation ITU-R P.1546 offers a good replacement for the mobile planning models mentioned above. This offers a similar accuracy to the un-calibrated models cited. Where small percentages of time are to be considered, there really only is Recommendation ITU-R P.452. Recommendation ITU-R P.1411 is candidate for the short path modelling needed for adjacent-block modelling over a few kilometres around the culprit stations. This latter propagation algorithm needs building structures, relative and absolute building heights and vegetation information at fidelity sufficient to discern actual building heights and street widths. This type of environmental data currently costs £250/€350 per square kilometre. The use of this type of modelling, whilst perhaps being optimum from the accuracy and confidence argument, is going to be costly for any party wishing to undertake modelling and development of evidence of SUR breach across a national roll-out. This illustrates that decisions are needed when setting out SURs. The regulator must consider how accurate the interference modelling needs to be based on his view of confidence. He must then select an appropriate model (comprising of terminal, environmental and propagation elements) to match this confidence. This is an area where care is needed to avoid battles in court of one science versus another.

3. Using SURs

This next section illustrates how SURs may be used in practice to constrain an operator. Two examples are used to explore the concept that, given a reasonably crafted SUR, an operator has some degrees of freedom before a breach occurs but conversely, SURs are fundamentally service, technology and topology dependent.

3.1 Geographic Separation. The first example shows one of two networks spaced geographically but sharing the same spectrum block. The network shown is the culprit or proposed network. The distant or victim network is off the map and spaced by perhaps a few tens of kilometres or perhaps even a few hundred kilometres depending on the percentage of time for which protection is expected.

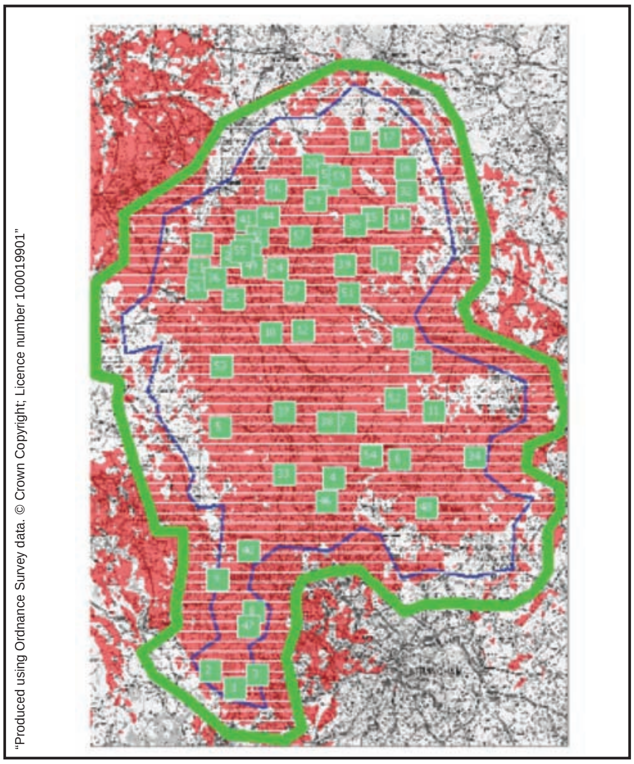

Figure 5: A network under SUR test. Blue line marks network service area. Green line marks SUR assessment boundary. Red shading shows areas where - 80dBW/m^2/MHz breached

The limit PFD has been determined by working backwards from the distant victim to show that at the boundary of the culprit, if the limit is not breached, the victim will be protected. In this example the calculated value and the prediction have been arranged so that breach is shown at a good number of points along the boundary. In this example the stations occupy the full block coincident with the whole of the victim block. If we assume that the victim network uses a technology that uses the full block but the culprit chooses a narrow band technology that occupies one tenth of the victim block and then uses frequency re-use across the culprit service area, the aggregate noise at the boundary must be reduced by about 10dB. This is shown in Figure 5.

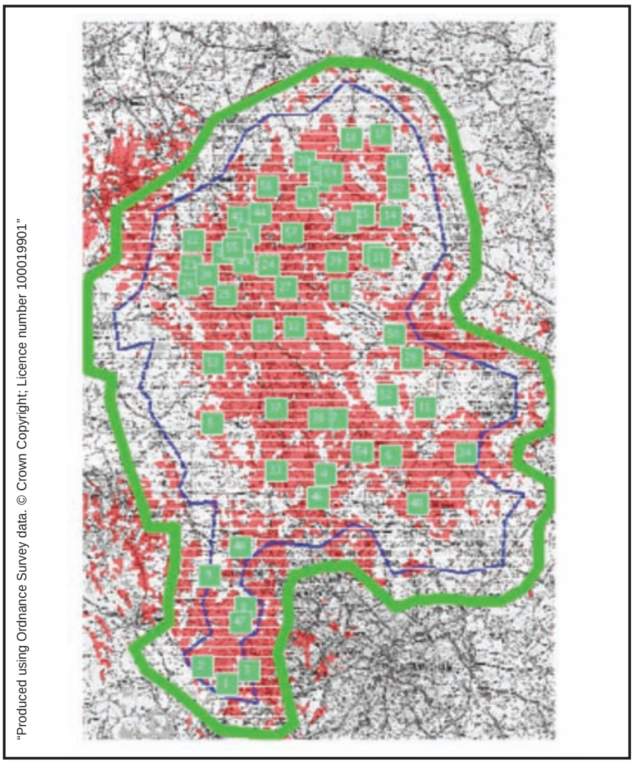

Figure 6: A narrow band network under SUR test. Blue line marks network service area. Green line marks SUR assessment boundary. Red shading shows areas where - 90dBW/m^2/MHz breached

In this instance the percentage of locations where breach occurs is much reduced. This shows that the SUR limit is very sensitive to changes in technology and hence it is important that the regulator makes the right assumptions when determining the limit value.

This *'breach at a line'* state can then be simply assessed quantitatively by summing the locations where breach occurs and reporting this as a percentage of the full boundary length. This is shown in Figure 6.

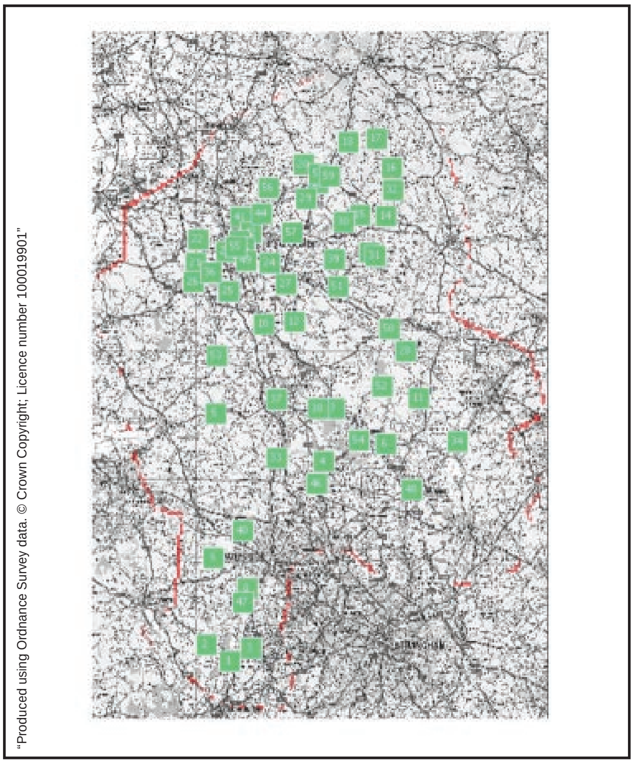

Figure 7: Filtering of PFDs at the boundary. Red intermittent line shows where - 80dBW/m^2/MHz breached. Amounts to 42% of boundary

Figure 7 shows that the limit PFD of -80dBw/m2/MHz is breached over 42% of the boundary. This could be reduced significantly by antenna tailoring to ensure that stations directed energy away from the boundary while still achieving the coverage needed. This is one of a host of methods of simultaneously achieving a balance of SUR compliance and subscriber service.

3.2 Frequency Separation. The second example describes two networks, each deployed to meet a requirement for 65% of locations coverage across an area in a European city – one using rooftop base stations and the other with base stations in street furniture such as street lights. Each network uses the same technology and service. The rooftop network needs seven sites where the street furniture network needs twelve sites for the same coverage. The compliant networks are shown in Figure 8.

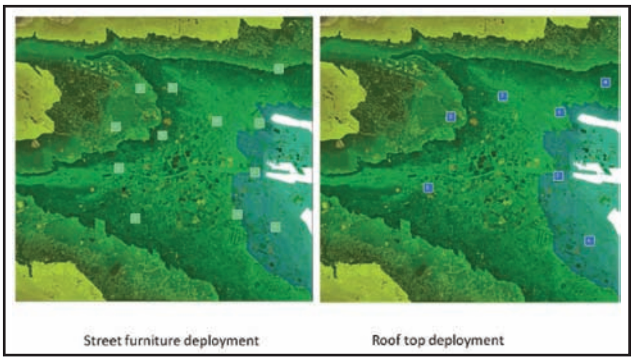

Figure 8: Two WiMAX deployments providing the same service are covered but with differeing topology

The SUR in this second example requires a PFD of -65dBW/m²/MHz in no more than 3% of the test area. The value of -65dBW/m²/MHz is arrived at by assuming a wanted minimum median operating field strength in the streets of 50dBuV/m commensurate with a WiMAX technology. It assumes a 20dB protection ratio giving a maximum nuisance field of 30dBuV/m. Assuming an adjacent channel rejection of 50dB, this then gives maximum adjacent channel field strength of 80dBuV/m. Applied across a 1MHz bandwidth and converting gives a PFD limit of -65dBW/m²/MHz. There are lots of assumptions here and in reality a more thought through figure would need to be arrived at. However, the value does give some useful working argument.

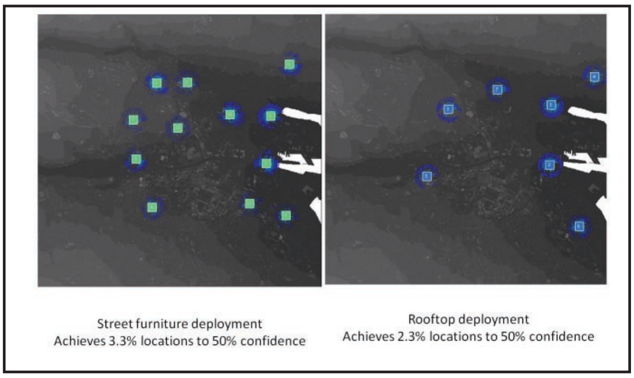

Figure 9: SUR requires a PFD of -65dBW/m²/MHz for 3% locations across the test area

As illustrated in Figure 9, a street furniture deployment breaches the SUR while a rooftop system does not – even though their coverage performance is about the same. This illustrates the degrees of freedom available and while this modelling has been confined to a change in topology, one could imagine that there could be significant change permissible across technology, service and topology and all can vary together to give high flexibility within an upper ceiling.

This modelling has been done at 50% confidence. Modelling to 90% confidence means being more pessimistic with the limit values (adding a margin) to yield a limit PFD of -75dBW/m²/MHz. The result is that the areas of breach increases dramatically meaning that fewer base stations are permissible. In this instance also, aggregation of signal levels makes a difference where previously, with non-intersecting PFD predictions, aggregation contributed little.

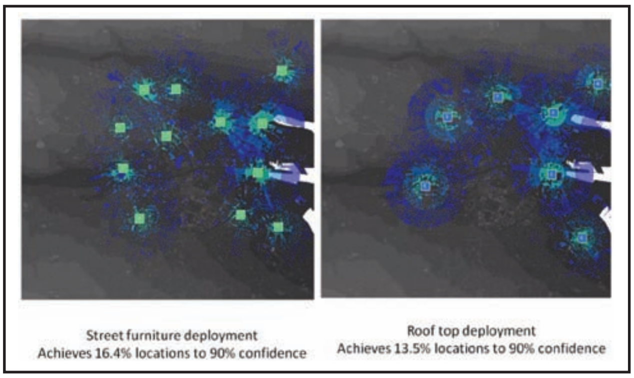

Figure 10: SUR requires a PFD of -65dBW/m²/MHz for 3% locations across the test area

Finally, looking at this SUR, one can see graphically the effect of the victim's service on the culprit's judged performance against constraints. Figure 10 shows that when the constraint is written that the PFD limit is applied only to the streets (simulating a victim system where the target user is mobile or nomadic with no indoor use) the percentage breach drops considerably to 12.9% for street furniture and 9.3% for rooftop deployment. It is clear that it is fundamentally important how the permissible breach locations are described. The figures have been arrived at by filtering the results to ignore PFD breach over buildings[5].

5. There is a slight error in this since the environmental data did not include the ability to discriminate PFD breach over water. This is simply corrected in practice but it does not materially affect the argument and hence has been left uncorrected in these results.