TV white space for South Africa

September 13th, 2017, Published in **Articles: EE Publishers**, **Articles: EngineerIT**

by Chris Burger and Dr. Moshe Masonta, CSIR Meraka Institute

The demand for radio spectrum has continued to grow in recent years. Fortunately, technology has continued to evolve. Modern wireless devices use frequencies up to perhaps 10 GHz, with specialised microwave systems now operating as high as the E-band (around 80 GHz). Even higher frequencies, into the optical band, are being contemplated. This article provides background to the recent developments in policy formulation for TV whitespace wideband networks in South Africa.

As technology has permitted the use of higher and higher frequencies, demand has grown to absorb the spare capacity. Point-to-point, mobile and broadcast services have all continued to migrate up the spectrum, consuming newly-available spectrum as soon as it becomes feasible. Unused existing spectrum is also snatched up by new users as soon as it becomes available. In urban areas world-wide, frequency congestion remains a serious issue, and spectrum planning continues to face challenges to provide equitable allocations to satisfy the demand.

Traditional band planning

Traditional spectrum allocation awards slices of spectrum to a specific user at a specific place. In general, the allocation is full-time, and no other user is allowed to use the same spectrum in the same area. There are exceptions. Citizen's band radio is a familiar example. In South Africa, only nine channels were allocated. Users had to share those channels, often resulting in intolerable levels of interference. Modern digital communications systems such as WiFi and mobile phones circumvent some of these problems by automatically time-slicing signals to enable several users to share the same spectrum. Although the user is not aware of any other users on the same channel, individual consumer devices jostle for access to the channel on an ongoing basis. Most such time-division systems are subject to interference by other users, and each user must accept that the service is provided on a best-effort basis. Depending on the sophistication of the sharing scheme used, the system may be quite effective in serving users until the full capacity of the spectrum is reached. Beyond this point, interference results in slow response or an inability to access the service at all.

Most commercial systems require some form of quality guarantee, so spectrum is allocated to a user or a group of users exclusively. If the user does not require the spectrum full-time, the spectrum remains unused for some of the time. Clearly, the efficiency of spectrum usage suffers under this arrangement. At times, such as the middle of the night, large swathes of spectrum may remain unused, while other users may be frustrated by a lack of spectrum.

Dynamic spectrum allocation

The solution to the wastefulness of traditional spectrum allocation is to dynamically allocate spectrum according to demand. Ideally, spectrum should be allocated to a specific user at a specific location on demand, and released as soon as the user no longer actively uses the spectrum. The same spectrum can then be allocated to another user. Such schemes are labelled as dynamic spectrum allocation (DSA).

With the increasing availability of the internet and other channels of high-speed communication, it is now possible in principle to allocate spectrum almost instantaneously, and to re-allocate that spectrum equally quickly. Users can now be notified to change channels or bandwidth, and be expected to comply. With radios becoming increasingly flexible as software-defined radio (SDR) and cognitive radio systems become more pervasive, we can expect this trend to continue.

We could think of classic spectrum allocation techniques as having a time constant of years. Users are given spectrum through the issuing of licences that often have a lifetime of several years. If a user ceases to use the allocated spectrum, that spectrum remains allocated until the licence lapses. Only then can the spectrum be reallocated to another user. A time constant this long was not unreasonable in the past. Decades ago, a radio set required a bespoke crystal to be cut to allow operation on a specific channel. Its modulation scheme was cast in stone, as a specific radio set typically included modulation and demodulation building blocks for only a single modulation scheme (like AM, FM, SSB or some data mode). Now, modern software-defined radios can be instantly reconfigured for different frequencies or modulation schemes, creating a need for more agile frequency allocation.

With DSA, the feasible time constant can become shorter and shorter. With current radios, perhaps a time constant of a day or two is possible, where spectrum can be re-allocated within a day or so after its use is discontinued. As radios become

Fig. 1: A CSIR High-Performance WiFi node with a UHF Whitespace antenna provides local coverage and long-distance backbone capability.

more flexible and more versatile, the time constant will continue to shorten, until eventually it may be feasible to reallocate spectrum within a matter of seconds, and to impose limitations on the modulation schemes and bandwidth for user devices to facilitate coexistence of users active at that particular moment in that particular location. Research even shows promise for sharing on busy bands such as the GSM mobile network, allowing more users to be served simultaneously.

Internet for everyone

The South African government has resolved to provide 90% of the population with 5 Mbps of affordable broadband access by 2020. In densely-populated urban areas, this target is relatively easy to achieve, as short-range links can provide large numbers of users with adequate services. However, peri-urban and rural areas are not so easy.

The majority of WiFi services now operate at 2,4 and 5 GHz. At such high frequencies, signals decay rapidly with distance. Also, to provide useful range, antennas require high gain. Gain antennas have to be directional, so a single antenna cannot provide broad area coverage. Especially at the service-provider end, it is difficult to provide sufficiently high gain without limiting the coverage of the access point.

Lower frequencies offer a far more favourable balance of coverage and range. Signals decay more gradually with distance, and antennas can provide wider coverage while still retaining enough gain for the less-lossy technology. It would therefore be advantageous to migrate some of the existing WiFi services onto lower frequencies. However, as we have seen, demand for new spectrum has historically consumed spectrum as it became available. The VHF and UHF bands, the frequency range below existing WiFi allocations, are fully committed to existing services. New demand from broadband services simply cannot be accommodated without somehow infringing on existing allocations in those bands.

Solving the congestion

Two approaches have surfaced to improve the situation.

The first approach is to use novel modulation techniques that make more efficient use of spectrum. Mobile phone networks are a prime example. Early mobile phones required an exclusive channel per user, quickly consuming available spectrum. Modern technologies such as GSM use time-division and other techniques to allow many users to coexist on a single channel. Similar breakthroughs have been made in broadcasting. Traditionally, each transmitter required a unique channel throughout its coverage area. With a limited number of channels available, congestion soon resulted. With modern digital techniques, many stations can share a single channel, or indeed even a single transmitter. Digital satellite television services provide an example that we are all familiar with. A single service can offer literally hundreds of channels using a single satellite transponder and a single receive antenna.

At last, the migration to digital terrestrial television broadcasting is in full swing in South Africa. As stations are migrated onto digital carriers, spectrum is expected to be freed up for other uses, including mobile telephony and broadband service. A subsection of the UHF band is already in the process of being vacated by television broadcasting, and will become fully available for other services, including mobile telephony. This band is known in the industry as the digital dividend.

Another approach is to re-use existing spectrum, using some form of DSA. The most common example of such DSA is TV whitespace systems, which use vacant television broadcast channels for wideband communications. Because television broadcasts require considerable bandwidth, a single television channel can accommodate wideband data, or several simultaneous narrower channels. A typical analogue television channel is 8 MHz wide, making it possible to accommodate several TV stations or perhaps 15 Mbps of data communications in a single channel.

In principle, it appears possible to simply listen on a channel to determine whether it is clear. Unfortunately, this approach has an inherent pitfall. The prospective re-user might be unable to detect signals from the broadcast transmitter, but will interfere with existing users who might be within line of sight of both the broadcast transmitter and the would-be data transmitter. Another approach is therefore required.

The solution comes in the form of a geolocation spectrum database (GLSD). This database contains information of all known transmitters and their characteristics. Using their channel allocations, permitted power, antenna gain figures, and operating schedule, the GLSD models the propagation of their signals, and effectively determines the coverage of each transmitter's signal. It is then possible to determine the spectrum occupancy at any particular point, anywhere in the coverage area. Openings, where a particular channel is unoccupied at a specific time and place, are known as whitespaces.

TV whitespace systems promise to provide considerable potential for wideband access in sparsely-distributed coverage areas. Fortunately, areas that are sufficiently sparse to require relatively long distances to be spanned between users and access points also typically have sparsely-occupied television channels, creating an excellent opportunity for spectrum re-use.

White space systems in Southern Africa

As TV white space systems appear to offer a very attractive solution to southern Africa with its sparsely-populated countryside, it is no surprise that southern African nations have been at the forefront of efforts to implement such systems. Other countries that have actively explored TV whitespace systems have included the USA, Canada, UK, Japan and Singapore.

The September 2016 issue of IEEE Spectrum carries an article titled "Bridging Africa's Broadband Divide", in which authors David L. Johnson and Chomora Mikeka outline efforts to implement trials in Cape Town and Malawi. Mikeka spearheaded a trial at Malawian schools in 2013, using vacant channels identified by simple spectrum scanning using a spectrum analyser. Some interference seemed to have resulted, but in general the trial proved a resounding success. Malawi's internet penetration was around 6% when the trial started, and the trial forms part of a strategy to markedly improve this penetration.

The trial in Cape Town broke new ground by using channels adjacent to those being used for live broadcasts, something that ran contrary to practice in most countries. Nevertheless, no reported interference to existing services resulted. The trial was mostly conducted by the CSIR, using manual analysis as the GLSD was not ready at the time. Throughputs of up to 12 Mbps could be achieved in standard 8 MHz television channels.

Since about 2010, the CSIR's Meraka Institute had been working with ICASA (the national spectrum regulator), Sentech (the national broad cast signal distributor) and others, punting the merits of white space systems. The institute developed a sensing technology demonstrator to demonstrate the feasibility of the approach. In 2011, Google approached ICASA about implementing white space systems. ICASA introduced them to the institute, knowing that the required capability was available locally.

The Cape Town trial was the culmination of this process, in which the Meraka Institute cooperated with Google and other partners to develop and demonstrate technology. The trial attracted considerable public interest, with over a hundred news articles appearing in the press during 2013. The results of the trial were sufficiently successful that the trial was cited by the USA's regulator, the FCC, as proof that adjacent-channel operations were feasible, contrary to what had previously been thought.

In 2014, the Meraka Institute conducted another trial in cooperation with Microsoft, the Department of Science and Technology, the University of Limpopo and Multisource. Other African trials have since been run in Botswana, Ghana, Kenya, Malawi, Namibia and Tanzania. Regulations at various stages of completion have since been drafted in Ghana, Kenya, Malawi and South Africa.

A Meraka Institute team has developed a GLSD based on the Protocol to Access Whitespace (PAWS), specifically with the intention of making their system vendor-independent. Using this open standard and a comprehensive database of existing licensed transmitters, the GLSD is able to accurately forecast the presence of commercial signals at specific locations, thereby avoiding much of the potential for interference. The CSIR GLSD was one of the first seven to be licenced for commercial operational use by Ofcom, the UK spectrum regulator, in 2016.

The chicken or the egg?

As is often the case with the introduction of new communications technologies, policy uncertainty slows down the adoption of new systems. Entrepreneurs are reluctant to fund the development of suitable chip sets until the rules have been fixed. Regulators, on the other hand, are reluctant to write regulations when technology trends are not yet clearly established. Regulations that are vague enough to allow any competing technology to be implemented will probably also be too vague to have any real regulating effect. Nevertheless, there has been significant progress. The Cape Town trials used Carlson Wireless equipment. The tight RF filtering in the equipment contributed to the success of the adjacent-channel operation. Carlson has subsequently introduced equipment fully compliant with IEEE 802.11af, based on a single chip that incorporates analogue input, logic and analogue output in a single device. Other vendors are in various stages of readiness, implementing the same standard or the alternative IEEE 802.22. Microsoft's Super WiFi, Taiwanese firm D-Link and Canadian firm 6Harmonics all offer solutions. And several nations have introduced regulations to facilitate their operation.

The South African TV whitespace regulations are the product of an interactive consultation between the Regulator and industry spanning several years. After an initial request for comments, 19 organisations responded with inputs. Some of the

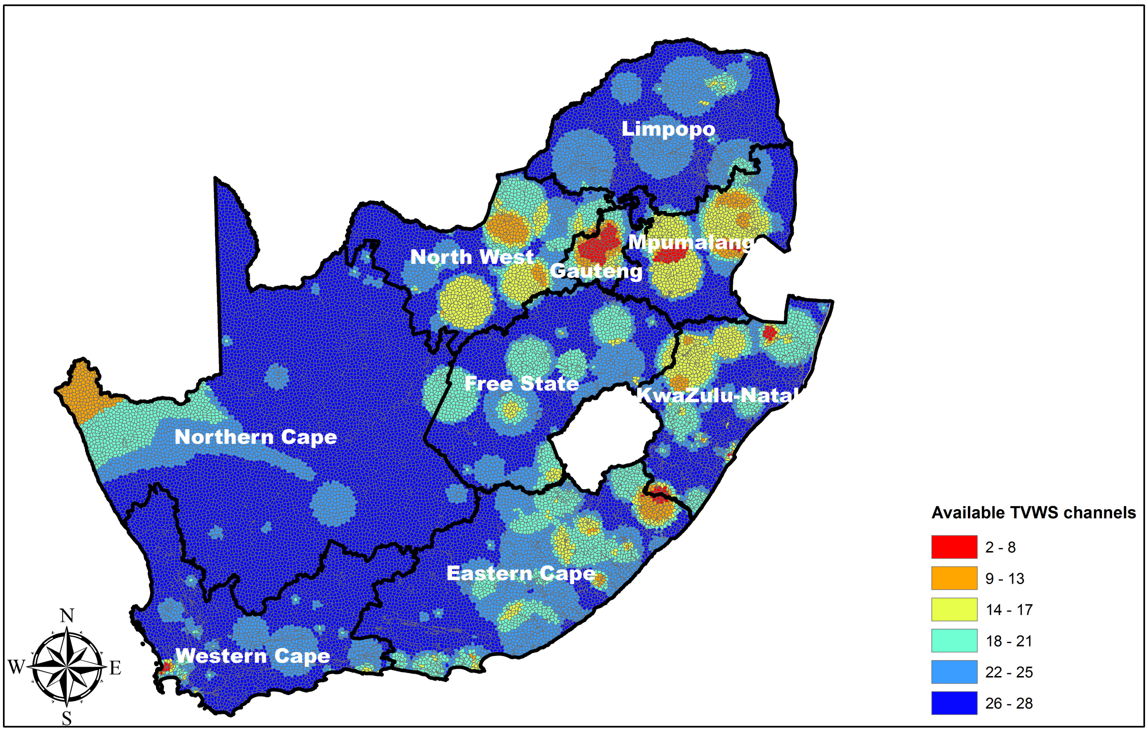

Fig. 2: Map of TV white spaces in South Africa.

contributions revolved around technical issues, with opinions also being expressed on the legislative framework and the extent to which ICASA was entitled to allocate permissions for TV white space operations. Using technical opinion from the Meraka Institute and several academic institutions, ICASA formulated draft regulations that were published during May.

Once the resulting comments have been discussed and incorporated into a final set of regulations, it is expected that South Africa will soon join the nations reaping the considerable benefits of TV white space systems. These systems should go a long way towards addressing the dearth of broadband access options in our rural areas.

Contact Chris Burger, Meraka Institute, Tel 012 841-4151, crburger@csir.co.za