

Meeting the White Space Challenge with Software Defined Radio

John Haine

*A Cognovo White Paper
September 2011*

1 INTRODUCTION

Mass market radio communications depends on access to radio spectrum. The award of radio spectrum to cellular operators sparked the mobile phone revolution; and the development of first Wi-Fi, then Bluetooth, were sparked by the global availability of licence-exempt spectrum at 2.4 GHz. Today virtually every smartphone sold operates on four or more frequency bands and includes Wi-Fi for local hotspot access and Bluetooth to connect to headsets and other accessories. Yet the exponential growth in data usage sparked by the coming of the smartphone is creating great pressure on the radio spectrum, especially in the magic “sweet spot” below 1 GHz that potentially combines large bandwidth with good propagation.

For precisely these reasons this region of the radio spectrum is under great pressure. By far the largest amount of the spectrum is allocated globally to UHF television broadcasting which occupies the band from ~ 470 to 860 MHz. Television broadcasting is now rapidly making the transition to digital, which gives better quality and more channels in less RF bandwidth, and as a result of this “digital dividend” some of the UHF spectrum, for example frequencies above 790 MHz in Europe and around 700 MHz in the USA, are being reallocated for cellular use.

But the introduction of digital technology brings another way to release spectrum. Even in the bands still used for broadcast, in any given area only a few of the available channels are used. If one were to colour on a map the areas covered by the TV transmitters on a given frequency, in principal the same frequency could be used for other services in the “white space” on the map, provided they did not interfere with television and themselves could tolerate interference from the TV signals. This “TV White Space” (TVWS) spectrum could theoretically release tens of megahertz in any given area provided that the equipment using it could operate anywhere in the TV bands. Suggested applications for White Space range from short-range “Wi-Fi – like” broadband internet access (but with better coverage characteristics), through wide-area specialised systems for “M2M”, to providing rural broadband access. Initiatives are under way in the United States (where the FCC has already published spectrum sharing rules); the UK (where Ofcom is consulting and has authorised trials in the Scottish Islands and around Cambridge, and announced its intention of proceeding with a “geolocation” method of controlling spectrum access¹); and in generally in Europe and Asia.

As the technology becomes established we expect TVWS communications to become another high-volume global wireless market and wireless adapters will appear in wide range of consumer products such as smartphones, tablets, and perhaps even automobiles; as well as being used for telemetry services such as meter reading and controlling a host of different types of plant.

2 STANDARDS AND SPECIFICATIONS

There are no regulatory standards for the air-interface protocol for TVWS devices, since a single protocol is unlikely to be suitable for the diversity of applications. In any case, regulators generally intend to make access to the band *licence exempt*, as is the case for example for the 2.4 GHz band.

Licence Exempt status implies that anyone can make use of the band for any purpose without a specific licence, provided that the application meets certain criteria. These normally specify the

¹ <http://stakeholders.ofcom.org.uk/consultations/geolocation/statement/>

transmit spectrum mask (both in-band and out-of-band) in terms of absolute or relative power; may limit usage (for example not allowing a fee-paying service); and usually stipulate that use for communications must cease if interference is caused to other services. For White Space the conditions are slightly more onerous: interference must not be caused to broadcast TV services or wireless microphones (which also use the band); frequencies can only be used and at a power level as authorised by a real-time geographic database; and the transceiver may also need to implement a “listen” function that verifies that the channel is free of any local transmissions above a threshold level before using the channel. These requirements are effectively a “coexistence standard” that has to be met by any equipment using the spectrum, whatever the details of the air-interface protocol. By its nature licence-exemption implies no common air interface.

A number of protocols are being proposed for White Space spectrum, some of which are listed in the table below.

Protocol	Source	Application
Weightless	Industry SIG led by Neul Ltd.	M2M
802.11af	IEEE	“Wi-Fi”-like hotspots
802.22	IEEE	Rural broadband
LTE-derived	3GPP	Rural broadband, mobile?
WiMax	IEEE	Rural broadband, mobile?
DVB-T	ETSI	Video distribution
ECMA 392	ECMA	General purpose

Neul Limited² in the UK is developing a specific standard – “Weightless” – for M2M applications using White Space frequencies which is designed from the outset to be compatible with TV transmissions and also to provide extremely good coverage so that it can reach low-power terminal devices in unfavourable locations (for example an electricity meter in a basement). Weightless is being developed by an industry group set up under the auspices of Cambridge Wireless³.

An obvious approach is to re-purpose other broadband wireless systems such as LTE, WiMax or Wi-Fi. The IEEE 802 committee has produced 802.11af, a wireless LAN enhanced for cognitive access in the TVWS band; and 802.22 designed for rural broadband access. These have the benefit that silicon technology is already available (for example from Wi-Fi or WiMax) that implement the baseband signal processing and protocol software. On the other hand they may not be ideally suited for the TVWS spectrum – for example in TVWS the adjacent channel transmit power specification is very stringent, and devices have to operate in the face of very strong TV transmissions, the primary use of the band.

² <http://www.neul.com/technology.php>

³ <http://cambridgewireless.co.uk/sigs/standards/>

There are various digital TV broadcasting standards, such as DVB-T, which could potentially be used in a “white space” mode for applications such as local video distribution.

ECMA 392 is developed in ECMA specifically for TVWS. It has elements in common with other protocols such as Wi-Fi (such as the use of OFDM) – it is not clear how much support it has in the industry.

As well as spectrum, two other factors are needed to make mass-market wireless systems successful: wide consumer application (such as cellular telephony or wireless Internet access); and low-cost silicon to implement the modem and radio. These two factors go hand-in-hand: only if many people can use a system will it generate the economies of scale needed to make the silicon cheap; and only if the silicon at least has the potential to be made at low cost will a mass market develop. The cellular phone quickly became a basic consumer good and generated a huge market which enabled low-cost silicon. Wi-Fi after a slow start got embedded in every laptop, and now every smartphone, since it enabled fast cordless Internet access using the globally available ISM spectrum at 2.4 GHz. Both these mass markets were enabled by common standards – GSM and IEEE 802.11. With White Space, by contrast, the diversity of air-interface protocols and applications could prevent rapid development unless some way can be found to unlock the supply chain for low-cost silicon devices.

3 SOFTWARE DEFINED MODEMS

The wireless modem is one of the most complex integrated circuits made. The peak computation load for an LTE terminal receiver for a 150 Mbps downlink is around 80 billion operations per second. This order of computing performance has to be accomplished in a low-cost device which consumes minimum amounts of power for a handheld product. Throughout the history of cellular the design of the air-interface protocol has leapfrogged the capabilities of processing devices, and modems have mainly been designed using a combination of conventional processing cores for control; DSP cores for some of the less intensive signal processing; and dedicated calculation logic for the most complex mathematical algorithms. This approach makes the chip inflexible: once designed for a given standard, fundamental changes in the air interface (as occurred in 3G for example with the transition from WCDMA to HSPA and then HSPA+) need new hardware to be designed; as it does also with the even more fundamental change from HSPA to LTE and then LTE-A. Multimode chips (and the cellular industry has hardly ever closed down one network when moving to the next) end up with different areas of silicon dedicated to the various standards. Apart from the inflexibility, the design process is clumsy since all the algorithms have to be designed and proved in detail before the actual chip design can start.

Today though, the capabilities of programmable processors have caught up with the evolution of wireless standards. As an example, Cognovo’s “Modem Computing Engine”⁴ (MCE) combines a number of innovations enabling virtually all aspects of different standards to be implemented in software running on a fully-programmable System on Chip (SoC).

- At its heart is the Vector Signal Processing (VSP) core. This can compute the same operation on an array of complex signal values in one cycle, a very efficient way to carry out arithmetic for modern radio modulation schemes such as OFDM. In addition the processor can carry out

⁴ <http://cognovo.com/products.html>

multiple unrelated operations in parallel in the same cycle, as determined by an optimising compiler. These two techniques (often called SIMD/VLIW) are key to performing complex algorithms such as FFTs and channel estimation extremely quickly with low power.

- In common with many other modern computer chips, the MCE uses multiple VSP cores, rather than running one device at much higher clock frequency. With the latest silicon process nodes, this gives a much better trade-off between complexity and power consumption. It does however require a different software approach to manage tasks running on different cores. Operation of the VSPs is “orchestrated” by a sequence processor, which executes micro-coded instructions to fire off software “kernels” which execute on the VSPs. Complex logical sequences can be programmed, and the sequence processor ensures that the multiprocessor system correctly and deterministically executes the complex sequence of steps involved in a typical air-interface protocol. The sequence processor relieves a conventional ARM control microprocessor of the task of hard real-time control.
- Cognovo’s system design tool allows the design requirements to be captured using a precise system modelling language, and then to be compiled into code for the sequence processor. This makes the design process fast and efficient and greatly reduces the chance of errors.
- Finally, the methodology allows fine-grained control of the power supply applied to various parts of the SoC to make sure that only the hardware required is powered on or clocked, to minimise the overall power consumption.

Using the MCE a modem can be developed quickly and efficiently in software without the designer having to worry about the co-design of the hardware. Even if the chip design and the algorithm design start at the same time, the fact that they can proceed in parallel allows up to a year to be removed from what is usually a 3 to 4 year programme. Once a chip has been developed the time saving can be even more provided that the device capabilities are properly dimensioned. The result is that a device can reach the market and start to generate revenue much quicker, can run multiple protocols, and allows the fundamental algorithmic operation of the modem to be updated at any time, even once in the field.

The MCE is also scalable. The main target configuration, the MCE 120, is dimensioned to meet the requirements of a “category 4” LTE handset (150 Mbit/s downlink, 50 Mbit/s uplink), and has two VSP cores. This same configuration could equally be used to support for example a DVB demodulator in one core in parallel with a Wi-Fi modem in the other. Another configuration, the MCE 160, has six VSP cores and is targeted at LTE-A developments, or for use in for example femtocells or military radios. The core fabric technology can handle up to 32 VSP cores for the most extreme applications.

4 THE SOFTWARE DEFINED WHITE SPACE RADIO

Figure 1 is a block diagram showing the main blocks of a versatile Software Defined Radio for White Space applications.

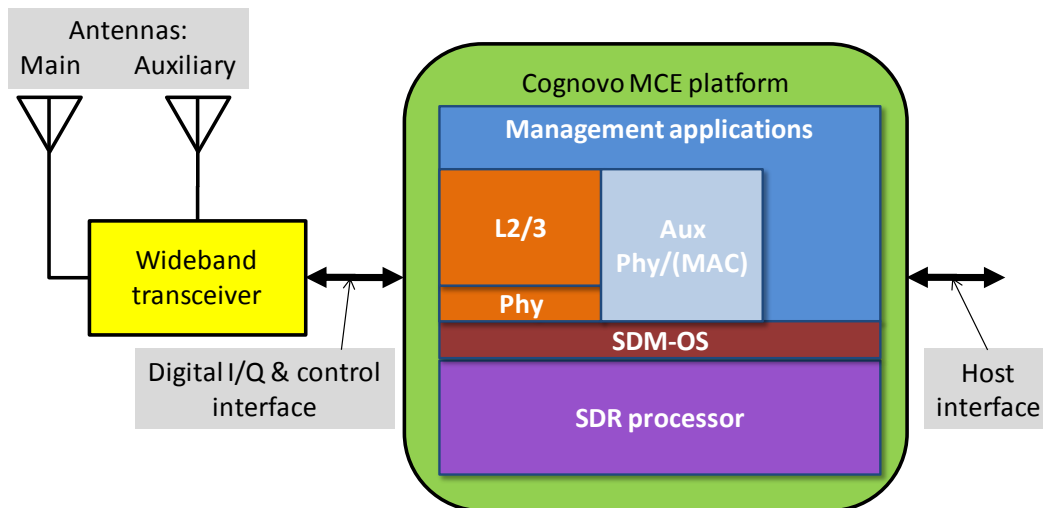


Figure 1: The White Space SDR

- The system is built around the Cognovo MCE platform. As well as the Vector processors this includes one or more ARM cores which can support various higher-layer and auxiliary protocols in addition to the selected White Space physical layer. It communicates with the radio transceiver through a standard “DigRF” interface defined for cellular applications in the MIPI Alliance.
- The MCE runs the White Space physical layer software and associated L2/L3 protocols, using the services provided by the SDMOS.
- In addition, other auxiliary air interface protocols run on the platform. An obvious example would be GPS or Galileo for geo-positioning. Other examples would be a measurement receiver for accurate identification of channel status for spectrum sensing; a DTV demodulator to help with signal identification (or indeed to provide a TV media stream); and cellular or Wi-Fi to provide temporary Internet access to the frequency management database. Clearly as more air interfaces are coded for the SDR platform the choice becomes wider. (Indeed, from one point of view the White Space protocols could be regarded as “auxiliary” to for example a cellular modem.)
- A library of management applications is provided to handle aspects such as database access and channel authorisation; and download and authentication of modem software modules. These are common services used by any White Space protocol.

5 A HARDWARE ARCHITECTURE

Figure 2 presents an outline hardware architecture for a White Space device, using a modem chip instantiating the Cognovo MCE 120 and a typical modern RF integrated circuit, in this case the Lime Microsystems⁵ LMS6002D Configurable broadband transceiver IC.

⁵ http://www.limemicro.com/download/LMS6002D_Product_Brief.pdf

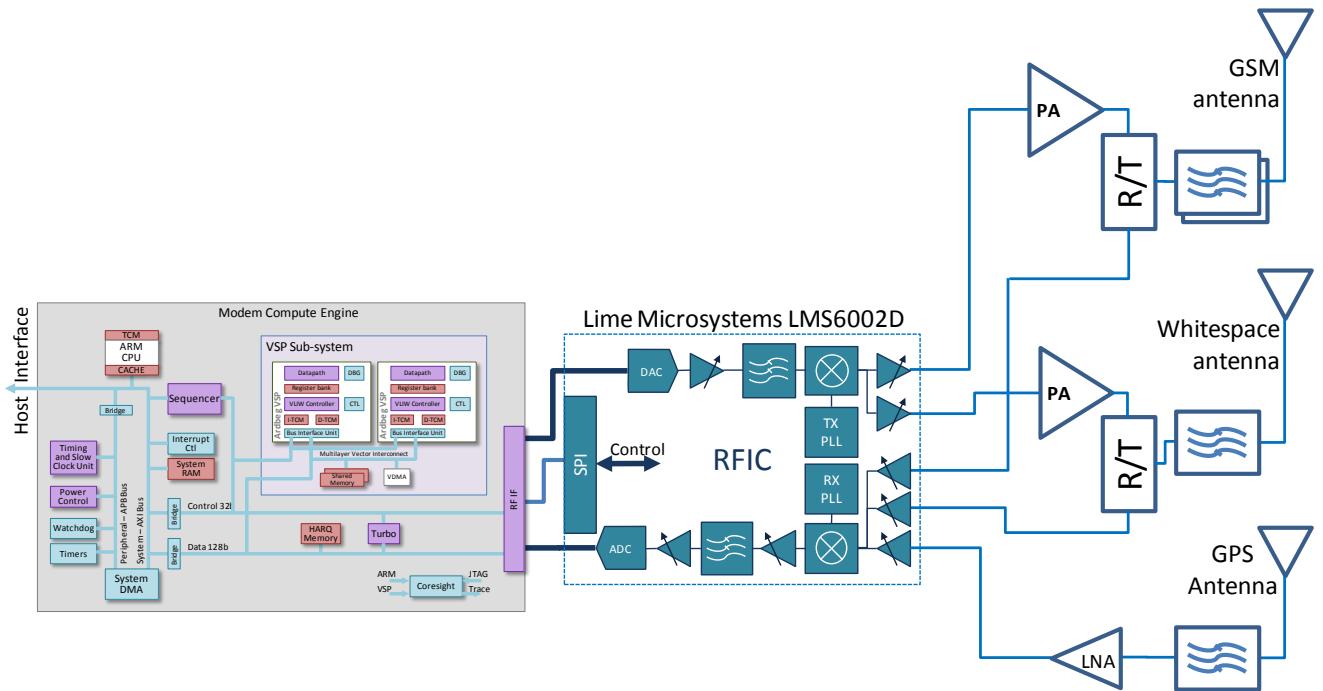


Figure 2: Outline Hardware Architecture

- The LMS6002D device is designed to cover a wide frequency range, from 375 to over 4000 MHz. This enables it to cover the UHF TV broadcasting bands plus all the 3GPP cellular bands as well as the 2.4 GHz ISM and 1.575 GHz GPS frequencies. It has a range of programmable modulation bandwidths, and separate synthesisers allowing arbitrary duplex spacing as well as time-division duplexing. All the digital-to-analogue conversion is done on-chip. It has three receiver inputs and two transmitter outputs to simplify RF switching and routing. Applications include cellular, WiMax, professional radio applications and connectivity.
- The RF transceiver will directly interface to the baseband chip. This will perform all the necessary digital processing, including the L2/L3 protocols hosted on the ARM Cortex control processor.
- In the example shown the multiple RF inputs and outputs of the Lime device are used fully to enable three modes of operation.
 - First, a receive-only channel at a spot frequency of 1.575 GHz is provided for GPS and other GNSS services. This would be used to support location of the device to enable it to obtain a suitable frequency allocation from an online geographic database. Thus the SDR modem enables the manufacturer to eliminate a stand-alone GPS device. A separate low-noise amplifier (LNA) may be needed to meet the stringent sensitivity requirement of GPS.
 - The main RF functionality is of course a TV band transceiver to access the “White Space” spectrum. Here the flexibility of the Lime solution comes into its own, allowing it to fit into a range of different channelisations to meet the needs of the White Space protocol and the rasterisation of the local TV service. White Space is inherently “time division duplex” so simultaneous transmit/receive operation is not required. As shown, a single block RF filter would be used to limit the transmit/receive band to the TV frequencies – however if necessary a tracking filter might be required to limit the band to the channel in use to reduce susceptibility to strong signals.

- Finally the example shown includes a dual-band GSM front-end operating in the 900/1800 or 800/1900 MHz bands. This is used as an adjunct to “white space”, to access the frequency assignment database to obtain a channel assignment. GSM is a good choice of air interface for this function as GSM has good coverage in most markets and (using GPRS) provides perfectly adequate bandwidth for occasional database look-ups. If necessary it would even be possible to use SMS for look-up as this service provides a somewhat better link budget which could be useful in marginal coverage locations. (An intriguing alternative would be to use a low-bandwidth but high coverage TVWS system such as “Weightless” to provide database access – such a possibility is only available through using SDR.)
- The different operational modes will be activated as required (in this example) by executing appropriate software in the SDR processor.

6 CONCLUSIONS

White Space communications, using Cognitive Radio techniques to enable re-use of the digital TV spectrum whilst protecting broadcast services, promises to be a key technique to provide extra bandwidth for broadband Internet services.

To fulfil its promise and develop a mass market, low-cost hardware will be needed. It is unlikely that a common air interface will be adopted since this spectrum is essentially “licence exempt”.

With the latest developments in processors for Software Defined Radio and versatile radio-frequency integrated circuits already available, it will be possible to make a generic silicon platform for White Space that can be adapted for any protocol and also to operate in the several modes that will be necessary to support additional functions such as geo-location and spectrum database access.

SDR promises to be a key technology to unlock the White Space market.