

Regular contributor Mike Richards G4WNC takes a detailed look at the new Apache Labs ANAN-100D SDR transceiver.

The ANAN-100D



The ANAN-100D is an imposing 100W HF through to 6m bands transceiver featuring fourth-generation direct digital conversion, dual analogue to digital converters (ADC) and a modular design approach. Having evolved from the Open Source High Performance Software Defined Radio (HPSDR) project, the ANAN-100D has an impressive pedigree.

Inside Story

One of the offshoots from the original HPSDR project was the Hermes board that combined a number of modules into one board and a single, larger, FPGA. The HPSDR project was led by **Kevin Wheatley M0KHZ** and the FPGA code changes were handled by

Phil Harman VK6APH. This proved to be a great success and is used as the basis of the ANAN-10E and ANAN-100 transceivers. However, there was an opportunity to take the work further and the Apache Labs Angelia board was the result. The Angelia hardware and PCB were designed by **Abhishek Prakash** of Apache Labs while the firmware was written and is maintained by **Joe Martin K5SO**. The Angelia board has many similarities to the Hermes board but there are some major improvements that help deliver the ANAN's impressive specification. I've shown a simplified block diagram in **Fig. 1**. The front-end of the Hermes receiver is supplemented by the use of an extra LTC2208 ADC. Both ADCs

run at a sample rate of 122.88 MSPS (Million Samples Per Second) with 16-bit resolution and thus support direct digitisation of the entire 10kHz to 55MHz spectrum. The use of two ADCs allows for the creation of two simultaneous, phase-coherent receivers anywhere in the 10kHz to 55MHz spectrum. The RF feed to each ADC comes via a 31dB programmable attenuator and an LTC 6400-20 low-noise preamplifier with a fixed 20dB gain. The other major change from the Hermes board is the use of a much larger Altera Cyclone IV FPGA with 115k logic elements as opposed to 40k logic elements in the Hermes. The extra capacity is used to support the dual ADCs and, with suitable PC software, it is possible to run up to seven software receivers from each of the two ADCs! These receivers can also be placed anywhere in the 10kHz to 55MHz range.

The processing requirements for the receiver are huge and handling the two ADC data streams alone is daunting as the high sample rate results in just over 120million 16-bit samples arriving every second from each ADC. The FPGA is the only practical answer for this level of data processing because it is essentially a parallel device. This means that the FPGA can carry out a huge number of complex mathematical operations in parallel as opposed to the largely sequential operation employed

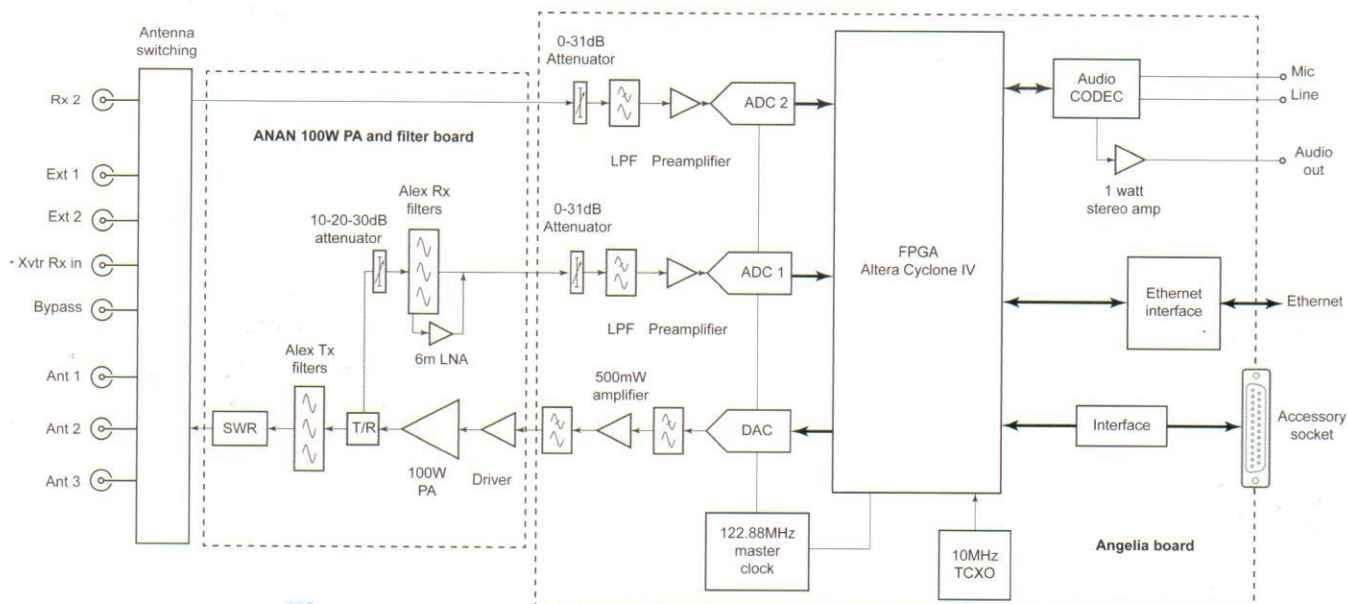


Fig. 1. Simplified block diagram of the ANAN-100D.

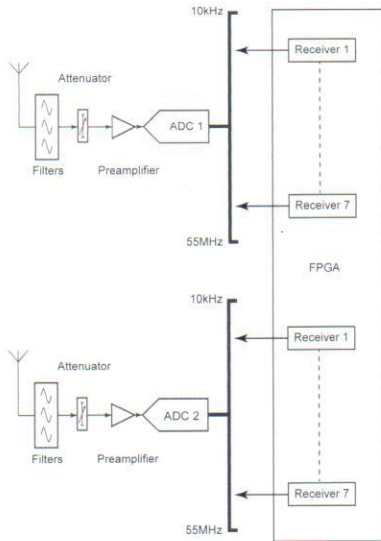


Fig. 2: With suitable PC software you can run 14 receivers with the ANAN-100D, thanks to the powerful FPGA.

by most microprocessors. One of the main roles of the FPGA in the ANAN-100D is to select down-sampled slices of the 10kHz to 55MHz bandwidth and prepare them for transmission to the PC software over the Ethernet link. The process of selecting a slice in this way is commonly known as decimation and is the digital equivalent of the way in which the mixer and local oscillator is used in a traditional superhet receiver. This is an oversimplification but it might help in visualising the process. The slice that's sent to the PC employs a lower sample rate and is therefore less demanding on both the communications link and the PC's processor. The sample rates available are currently fixed at 384k, 192k, 96k and 48k SPS (Samples Per Second). These sample rates also

broadly correspond to the tuning width available, so a 384k sample rate will allow a receiver tuning range of about 380kHz. As you can see from Fig. 2, with suitable PC software it is possible to have multiple receivers, each with their own slice of the 10kHz to 55MHz band.

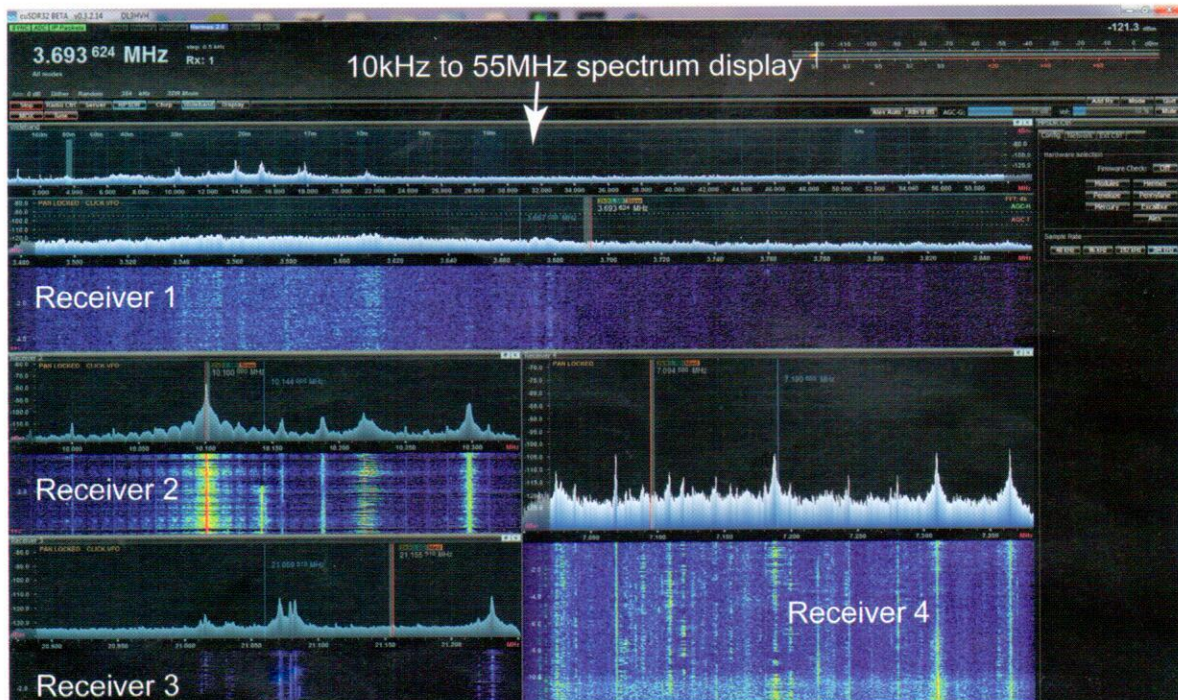
On the transmit side, the input signal from the front panel microphone socket is initially digitised using a Texas Instruments TLV320 CODEC chip before passing into the FPGA for processing. The digitised audio stream is made available to the PC software, which enables a huge range of processing options and allows for the use of an infinite number of microphone types. After processing in the PC, the transmit signal is passed from the PC software over the Ethernet as a 48kSPS IQ (In-phase and Quadrature) data stream. The FPGA provides interpolation filtering of the transmit stream to bring the sample rate up to 122.88MSPS where it is passed to an AD9744 digital to analogue converter (DAC) to produce the low level analogue transmit signal on the correct frequency. This signal then passes to the separate ANAN 100W PA and filter board that is based on the ALEX module from the HPSDR project. In addition to the 100W PA, this board also contains an SWR bridge and protection circuitry for the PA.

One of the benefits of the large FPGA in the ANAN 100D is the potential to include soft processors that could be used to provide stand-alone operation and other functions. In addition to the Angelia board, the ANAN 100D uses the Apache Labs designed filter and PA unit to provide up to 100W RF output on all bands. The transmit and

receive filter banks are modelled on the HPSDR Alex filter modules with large ferrite cores used to minimise IMD. In addition, all ANAN models are currently being upgraded to add adaptive digital pre-distortion to the transmitter. It may seem odd to be adding distortion to a transmitter, but this technique is widely used in commercial digital transmitters to improve the efficiency and linearity of power amplifiers. Although not available in time for the review, the new pre-distortion will deliver -50dB IMD and will soon be available via a firmware update for all ANAN transceivers.

Setting Up

As you can see from the photos, the ANAN 100D is built like a tank with its thick extruded aluminium casing and utilitarian styling. The first step is to load the SDR software and drivers onto your PC. The PC demands are not too onerous and Apache Labs suggest an Intel Pentium 4 or better running Windows 7. The software and manuals were available on the supplied CD-ROM but the latest updates were available from the Apache Labs website. The standard SDR software recommended for the ANAN 100D is the PowerSDRmRX, which is based on the well-known PowerSDR software used with FlexRadio systems. However, the ANAN-100D version has received many modifications, including the facility to run with two receivers using the dual ADCs. Communication between the PC and ANAN-100D is handled via a standard Ethernet port that's located on the front panel. For the review, I connected the ANAN-100D via a



The cuSDR software running four receivers with the ANAN-100D.

Netgear Gigabit managed switch that I have installed in the shack, but it can be directly connected to your PC or located remotely. There's no need to worry about Ethernet configuration as the hardware automatically configures the link for the fastest speed available.

Like most amateur transceivers, the ANAN-100D requires a hefty 13.8V power supply with a current rating of 25A. With Ethernet and power connected and the software loaded it was time to fire it up. The ANAN-100D uses the case as a heat sink, so it's important to make sure there's an unimpeded flow of air around the case. A quiet internal fan helps with cooling.

The first stage of the configuration process is to fire up the PowerSDRmRx software and trim the transmit levels. For the configuration adjustments, the Ant-1 output of the ANAN-100D needs to be connected to a 100W dummy load.

When starting PowerSDR software for the first time there is a delay while it runs through its FFT optimisation process. This is used to assess the processing power of the PC in order to select the appropriate FFT settings. This process took around 10 minutes on my PC. Once complete, I was able to go to the setup menu and check to see if the ANAN-100D had been successfully detected on the network. In its default state, PowerSDR is set for full network discovery and the ANAN-100D was automatically detected on my system. Although the default configuration used the router to provide the IP address automatically, you can, alternatively, set a fixed IP address.

With the initial start-up routine

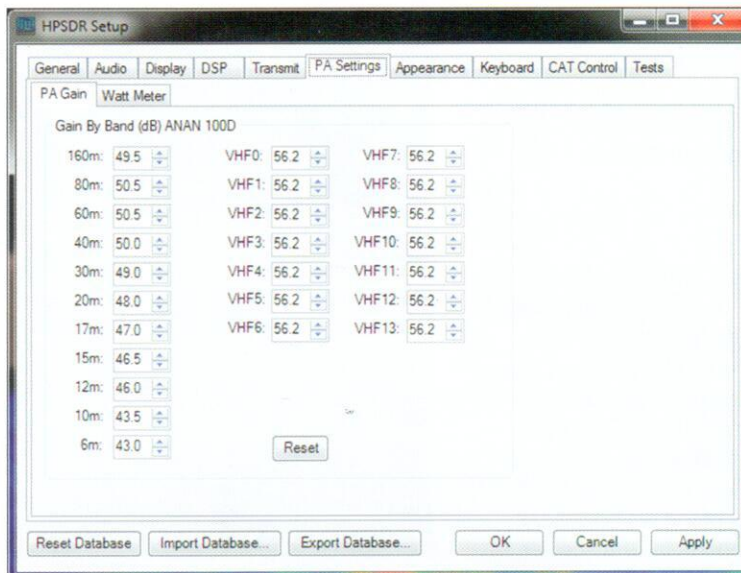


Fig. 3: Adjustment panel for trimming the transmit levels in PowerSDRmRx.

complete, I set about trimming the PA output levels using PowerSDR. The purpose here is to equalise the PA gain settings for each band so that the CW output is consistent. To do this I connected a 100W dummy load to the Antenna socket, set the Tune power to 100W and then adjusted the gain setting for each band in turn so that the forward power meter read 100W, Fig. 3.

As the ANAN-100D is supplied without a microphone, you have to adjust the transmit settings to match your station microphone. This was again carried out using a dummy load and was well covered in the user guide. I also found a very useful video on the internet describing the process. Although the video is a generic PowerSDR tutorial, the advice is appropriate for the ANAN-100D. www.w1aex.com/psdrgain/psdrgain.html


The front panel microphone socket is a stereo type that can be configured using a set of jumpers on the Angelia PCB. The most common arrangement would be to connect the microphone to the tip, PTT on the ring and ground on the sleeve. The ANAN-100D is set for a standard dynamic microphone but jumpers on the Angelia board can also be set to support electret microphones.

Morse key options are well provided for with a 3.5mm jack on the front panel that can handle straight or iambic keys.

Comprehensive Connections

To cope with the wide frequency coverage, the ANAN-100D has three BNC antenna sockets on the rear panel and these can be configured on a band-by-band basis. The software also provides the option to use a different antenna for receive, which is also selectable on a band-by-band basis. There is a separate BNC antenna socket for the second ADC input. In addition, there are two Ext, one transverter and a bypass socket that can be used to route the RF through additional filters, a pre-amplifier or a transverter. The audio output is presented via a 3.5mm jack on the front panel and this can be used to feed headphones or powered speakers. There is also a speaker jack on the rear panel, but this is a balanced output where the speaker element needs to be connected between the tip and ring of the plug.

If you have the need for precise frequency control, the ANAN can accept an external 10MHz reference via a socket on the rear panel. For the connection of external equipment such as a linear amplifier, there is a separate 'PTT out' socket along with a fully configurable 25-way accessory



summary of the available connections: 2 x digital input, 2 x analogue input, 7 x open collector outputs, left/right line out, left/right line in, left/right phones output, right balanced speaker pair and PTT input.

On the Air

It's important to appreciate that the user experience of an SDR transceiver is determined largely by the software you are running and, at the time of writing, the two main contenders for the ANAN-100D are PowerSDRmRx and cuSDR. For this review, I've mainly used PowerSDRmRx simply because it accesses more of the ANAN-100D's features and supports full transceiver operation. I also tried cuSDR because this software has a much more attractive and responsive interface and supports multiple receivers.

PowerSDRmRx is bristling with features and could easily take over the entire article. As I mentioned earlier, PowerDRSmRx is a custom version of PowerSDR that's been designed specifically to work with the ANAN range of transceivers. I've shown a screen shot of the main interface in **Fig. 4**. Here you can see the main tuning display with combined spectrum and waterfall displays. This is my favoured way of viewing the spectrum, but PowerSDRmRx provides a number of alternative display tools that are easily selected. Tuning can be achieved in a number of ways from simple click-tuning with the mouse through to memory recall from the software's memory system. Band selection is provided for all the WARC bands along with region-specific transmit protection to prevent inadvertent out-of-band transmission. As you would expect with an SDR rig, the filtering is exceptional and very versatile. Mode specific pre-sets are available and fine for most situations, but I could manually set just about any combination of bandwidth and offset either by using the sliders or by dragging the display. I found

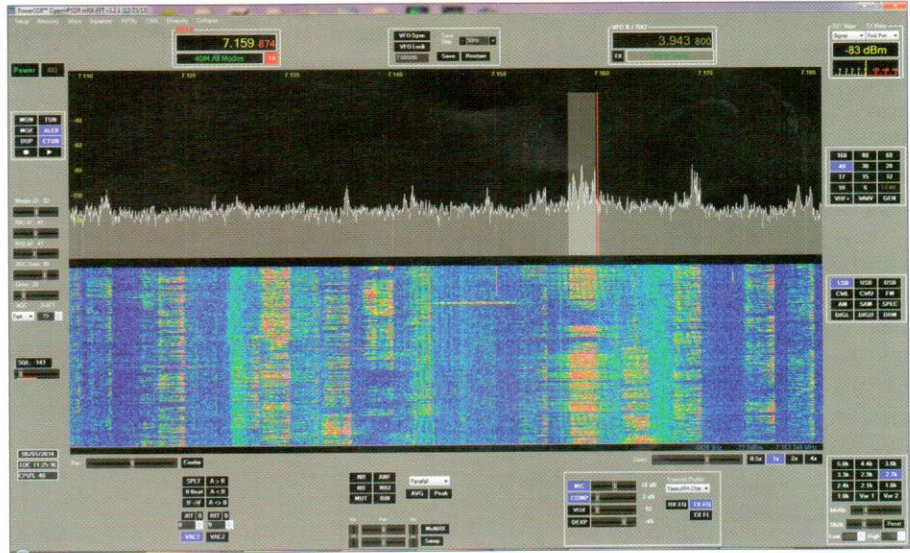


Fig. 4: PowerSDRmRx main display.

this latter option particularly powerful as I could drag the upper or lower edge of the pass-band to exclude an interfering signal. In addition to the main pass-band filtering, PowerSDRmRx includes DSP noise reduction, a noise blander and an automatic notch filter. The default settings for all these filters can be customised via the Setup menu so it's very easy to adjust these to match your personal preference.

For the transmit side, I used a Yaesu MH31 dynamic microphone for SSB as this seems to suit my voice well. The transmit adjustments in PowerSDRmRx are very comprehensive and include a 10-band graphic equaliser as well as VOX and compression controls. The software also provides the facility to save transmit profiles. This was very useful as I could have an unlimited number of profiles for different situations. After following the microphone setup guidance, I was receiving very good audio reports on air.

Tuning around the bands was a delight once I'd mastered the tuning system. The ANAN-100D has an exceptionally sensitive receiver with a noise floor of around -140dBm (0.23µV) and a very clean transmit




Fig. 5: PowerSDRmRx phase adjustment panel for diversity reception.

signal. This was a tremendous benefit during the recent 10m openings and I was able to copy some very weak stations successfully. Although I didn't have a need for it during the review, the hardware and software supports full duplex operation across any two bands.

The range of adjustments available via PowerSDRmRx and the ANAN-100D

The rear panel of the ANAN-100D, showing the comprehensive range of connections.



is almost overwhelming so I was grateful for the facility to save all my settings to disk. PowerSDRmRx employs an XML database to store all the configuration settings, so I was able to store a number of configurations to cater for different operating modes.

Rig Control and Audio Routing

With many amateurs using software logging systems and data modes programs, a computer aided transceiver (CAT) system has become an essential item for any rig. Most CAT systems still rely on a serial COM port to make the link between PC and transceiver. When using SDR-based transceivers such as the ANAN-100D, the link can be entirely in software with no physical cables. With PowerSDRmRx this linkage is achieved by using a virtual Com port pair and virtual audio cables. Virtual in this case refers to software emulation of a physical cable. For the rig control, the virtual COM port pair can be provided using free software and I used Com-0-Com. Once installed and set up, I created COM23 and COM24 as a linked pair. PowerSDRmRx was then set to use COM23 and control software such as FLDIGI was set to communicate on COM24. Audio routing was done using Virtual Audio Cable (VAC) by **Engine Muzychenko**. This provides for digital audio routing between applications and I used it to feed the data modes tones to the PowerSDRmRx transmit section and to send the audio back to FLDIGI. This all worked extremely well, allowing full control and data modes operation with no interface unit or additional cables. PowerSDRmRx also includes controls to switch the VAC on and off and to control the levels.

Diversity Reception

The provision of two phase-coherent receivers in the ANAN-100D opens up the world of diversity reception. Diversity reception provides noise reduction and signal enhancement by precise mixing of the signals from two different antennas. The trick is to be able to vary the phase and amplitude of the signal from one antenna either to enhance the wanted signal or to cancel out some of the noise. For this to work effectively, you first need two identical, phase-coherent receivers, two spaced antennas and software to control the mixing phase and amplitude. The ANAN-100D provides the receivers while the PowerSDRmRx provides the mixing facilities. That just leaves the antennas for you to provide. For the review I used my Wellbrook AI 1530S+ Imperium loop antenna as

Specification Summary

Receiver

Architecture	Direct Down Conversion
Number of ADCs	2
Number of receivers	7 + 7
ADC resolution	16-bit
Frequency coverage	10kHz to 55MHz
Dynamic range	125dB
Diversity reception	Yes
Noise floor	-140dBm (500Hz b/w @ 14.2MHz)
Preamplifier	20dB, 20dB LNA for 6m
Attenuator	Software selectable 31dB in 1dB steps
Spurious and image rejection	Greater than 100dB
Audio amplifier	1W @ 8 ohms

Transmitter

Architecture	Direct Up Conversion
Number of DACs	1
RF output	1 to 100W
Coverage	160m to 6m
Modes	All – limited only by PC software
Modulation	Digital low level
Carried and unwanted sideband suppression	Greater than -90dBc
Harmonic suppression	Greater than 50db (HF) and 60dB (VHF)

General

FPGA Logic Elements	115k
Flash RAM	128MB
SRAM	32-bit
Master clock	122.88MHz
Clock phase noise @ 10kHz	-149dBc
10MHz reference TCXO	± 2.5PPM
Frequency steps	1Hz
Antenna ports	3 + 1
Dimensions	135 x 160 x 80mm
Weight	4.5kg (approx.)

the main antenna. The phase and gain balance between the two receivers is handled using a neat Phasing control in PowerSDRmRx (**Fig. 5**). With everything set up, all I had to do was drag the centre point around the display to find the best noise reduction/enhancement. Although tricky to set up properly, I was generally able to see a signal improvement of about two S-points on many of the weak signals.

ANAN Range

Although I've concentrated on the ANAN-100D here, there are three broadly similar rigs in the ANAN line-up. The ANAN-10E and ANAN-100E both use the Hermes board with a single 40k FPGA and single LT 2208 ADC. They are both able to support multiple receivers and the ANAN-10E features 10W RF output while the ANAN-100E uses the same Apache Labs 100W PA as the 100D.

Summary

The ANAN-100D is an impressive hardware package that represents the state of the art in SDR hardware design. Not only is the performance excellent but its versatility is exceptional and the facility to upload new FPGA code

The D version with the much larger FPGA and dual ADCs has the greatest potential and it's hard to imagine how anyone would need more than this in an HF/6m rig. At the moment there is no SDR software that makes full use of the ANAN-100D's potential but that is changing rapidly and you can expect to see some exciting new software as the ANAN transceivers become established. The only downside to the ANAN-100D's versatility is the requirement to spend some time setting up the rig for your station. To do this successfully, a degree of technical knowledge is essential. That shouldn't be a problem for most amateurs because getting to know your rig is an important part of the learning experience. However, if you're looking for a simple plug and play rig, this may not be for you.

The ANAN-100D is available from Waters and Stanton Ltd. priced £2,999.95 inclusive of VAT at 20%. The ANAN-100E is priced at £2299.95 while the ANAN-10E retails for £1549.95. My thanks go to Waters and Stanton for the loan of the review model and to Abhishek Prakash of Apache Labs for his technical support.

www.wsplc.com