
Pierre Auger Observatory Communications Task

Auger Southern Observatory Communications System Technical Design Report

Version 2.0 March 2004



Preface

The Pierre Auger Southern Observatory is located close to the town of Malargue in Mendoza Province, Argentina. The Observatory has five Communications Towers which provide a variety of Voice and Data Radio Services to the Experimental Equipment.

This document was originally one chapter of a full Technical Design Report on the complete Auger Observatory. The document describes the technical details of all aspects of the communications systems used in the Observatory.

PDJ Clark – July 2006

Original Authors

V Tunnicliffe, P Walker, L de Bruijn, PDJ Clark

Institute of Integrated Information Systems
School of Electronic Engineering
University of Leeds
Leeds LS2 9JT
United Kingdom

3. COMMUNICATIONS IN THE AUGER OBSERVATORY

Due to the large coverage area and widely dispersed nature of the 1600 Cherenkov detectors that will make up the surface detector array, a communications system based on radio technology was deemed to be the only economically viable solution for the Pierre Auger Observatory.

This chapter describes the communications system and the custom digital radio hardware that was necessary to meet the design goals of the project.

The reliability of the radio network is critical to the data acquisition operations of the Observatory, and this chapter also describes the extensive network planning effort that has been undertaken to ensure that radio propagation conditions at the site are well understood.

3.1 INTRODUCTION

The Auger Observatory data communications system consists of 2 integrated radio networks organised as a 2-layer hierarchy: the individual detectors are serviced by the *surface detector wireless LAN (WLAN)*, which is a sectorised network supported by 4 data-concentration nodes. These WLAN nodes are serviced by a high capacity microwave *backbone network*. The backbone also supports communications from the Fluorescence Detector sites.

3.1.1 Microwave Backbone Network

The backbone network uses a standard 34 Mbps telecommunications architecture based on commercially available microwave point-to-point equipment operating in the 7 GHz band. The equipment consists of dish-mounted microwave transceivers installed on communications towers, together with secondary units located in shelters at the base of each tower.

As shown in Figure 3.1.1, the backbone network consists of two ‘arms’, both of which terminate at the Observatory Campus in Malargüe, at which point the data is routed to the central data acquisition system (CDAS).

The backbone network has sufficient capacity for the transfer of both the surface detector WLAN data and all data to and from the fluorescence detector sites.




Figure 3.1.1 : The backbone layer communications system

3.1.2 Surface Detector WLAN

The surface detector WLAN has been specially designed for the Auger project using custom radio hardware running proprietary network access protocols. This network operates in the 902-928 MHz industrial, scientific and medical (ISM) radio band and provides data communications to and from each of the 1600 surface detectors over a 3000 km² area. This is achieved in a manner similar to a cellular telephone system, whereby the area containing the detectors is divided into a number of sectors, and communications within each sector are coordinated by a *base station*. Sectorisation is required in order to meet legislation pertaining to maximum transmitter powers and frequency re-use within the ISM band. It also greatly distributes the data processing load of the array, and reduces the possibility that a failure at a data collection node will cause a total Observatory outage.

Figure 3.1.2 illustrates a simplified view of the sectorisation of the Observatory surface array. The base station antennas serving each sector are mounted on the communications towers, with the towers supporting up to eight sectors. The antennas are divided into two sets: an upper set mounted at the top of the tower and a lower set a few metres below. The upper set service sectors deeper into the array, whilst the lower set service sectors towards the array boundary. This is represented by the different shades in the figure. A single sector typically contains 57 surface detectors.

Sector planning is discussed in detail in §3.4.5.

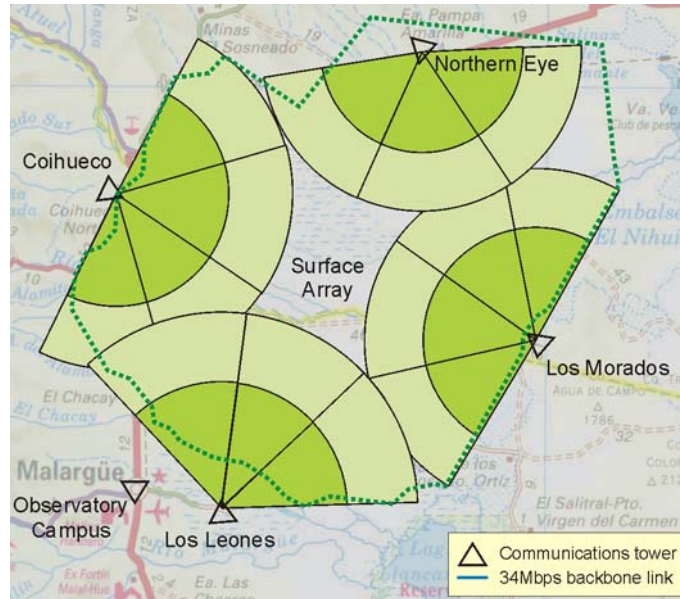


Figure 3.1.2 : Surface detector wireless LAN coverage (simplified)

3.1.3 Data Path from the Surface Detectors to Campus

WLAN air-interface and related functions at each surface detector are performed by a *subscriber unit* (Figure 3.1.3). This unit communicates with the main surface detector electronics module via a serial link, and incorporates a proprietary digital radio transceiver running appropriate firmware on its control processor. The subscriber unit is connected to a 12 dBi Yagi antenna via a short low-loss feeder, with the antenna mounted on a short communications mast integrated into the detector's solar panel support. This can be seen in Figure 3.1.4, which shows a typical detector tank installation prior to deployment into the field. The antenna is mounted at a height of 3m above the ground.



Figure 3.1.3 : WLAN Subscriber unit



Figure 3.1.4 : Surface detector with integrated WLAN mast and Yagi antenna

Data is transmitted over a path of up to 30 km to a local data concentration tower, where the signal is received via a high-gain cellular-style panel antenna. The antenna is connected via very low loss feeders to a base-station unit (Figure 3.1.5) located in a shelter at the base of the tower. A single base-station can serve up to 68 detector tanks. A base station incorporates the same digital radio transceiver platform employed in the subscriber units, with additional processing capability.

Figure 3.1.6 shows some 3 meter panel antennas being installed at the top of the Los Leones communications tower in April 2000.



Figure 3.1.5 : WLAN Base-station unit



Figure 3.1.6: Installation of WLAN panel antennas on the 41 metre tower at Los Leones

At each tower, data from several base-station units is concentrated onto an E1-ring and then processed by a custom interface termed the *micro-LSX* (Figure 3.1.7) before being passed on to the backbone microwave network for transmission to the Observatory campus in Malargüe. There the data passes from the E1 microwave network into the central data acquisition system via TCP/IP running on a conventional Ethernet network.

Figure 3.1.8 shows the second of two microwave outdoor units (parabolic dish and radio) being installed on the tower at the Observatory campus.



Figure 3.1.7 : Micro-LSX unit provides E1 to TCP/IP conversion




Figure 3.1.8 : Installation of a 7 GHz microwave antenna on the 50 metre campus tower

3.1.4 Digital Radio Transceiver Development

The need to provide many highly robust data links over distances of 30km and beyond using a minimal amount of power has presented some unique equipment performance challenges that could not be met with existing communications equipment. To meet the requirements, a low-power digital radio transceiver platform has been designed, the functionality of which is determined by re-configurable firmware stored in Flash memory and loaded into a digital signal processor (DSP) when the unit powers up. This re-configurability not only allows the functionality of the transceiver to evolve and be refined over time, it also allows a common radio transceiver platform to be used for both subscriber units and base-station units, thereby reducing the hardware development time and simplifying long-term hardware support.

The transceiver uses very low power devices and a highly flexible architecture to provide reliable long range digital communications within the strict power budget of the solar-powered surface detectors. Power consumption is less than 1.1 Watts of DC power. The transceiver circuit board is shown in Figure 3.1.9.




Figure 3.1.9 : WLAN radio transceiver circuit board

3.1.5 Radio Network Planning

In order to gain an in-depth understanding of the expected performance of the radio network prior to deployment at the site, a campaign of on-site radio frequency (RF) measurements has been undertaken. This campaign has been augmented by a detailed series of RF propagation modelling exercises using standard radio propagation models, combined with digital elevation terrain data of the site, and these exercises have produced accurate simulated radio coverage maps of the surface array area. Such coverage maps have assisted greatly in finalising the design of the topology of the radio network.

Figure 3.1.10 gives an example propagation simulation coverage map for operation at 915MHz, in which multiple transmitters are simultaneously enabled. The various colour bands denote which tower offers the strongest signal at a given location.

The radio coverage maps offer a powerful interactive tool to assist in the planning of the surface array radio network, specifically:

- The placement of data concentration towers may be readily optimised.
- Areas of poor coverage may be identified in advance of surface detector deployment. For example, Figure 3.1.11 shows an area of severe radio signal shadowing due to local terrain features close to the site of the Los Morados tower at the eastern edge of the array. Advance knowledge of these problem areas permits alternative coverage plans to be implemented.
- ‘Best-service’ sectors may be identified and assigned, i.e. which tower should a given surface detector use?
- The height of the 1600 surface detector radio masts can be minimized, thereby leading to significant project cost savings.

The accuracy of the radio coverage maps was verified by a series of on-site roving spot measurements, and also via a long-term static radio propagation measurement experiment. Figure 3.1.12 shows a series of simulated path-loss estimates compared with actual on-site measurements, and gives an example of the good correlation between the simulated maps and the actual conditions found at the site.

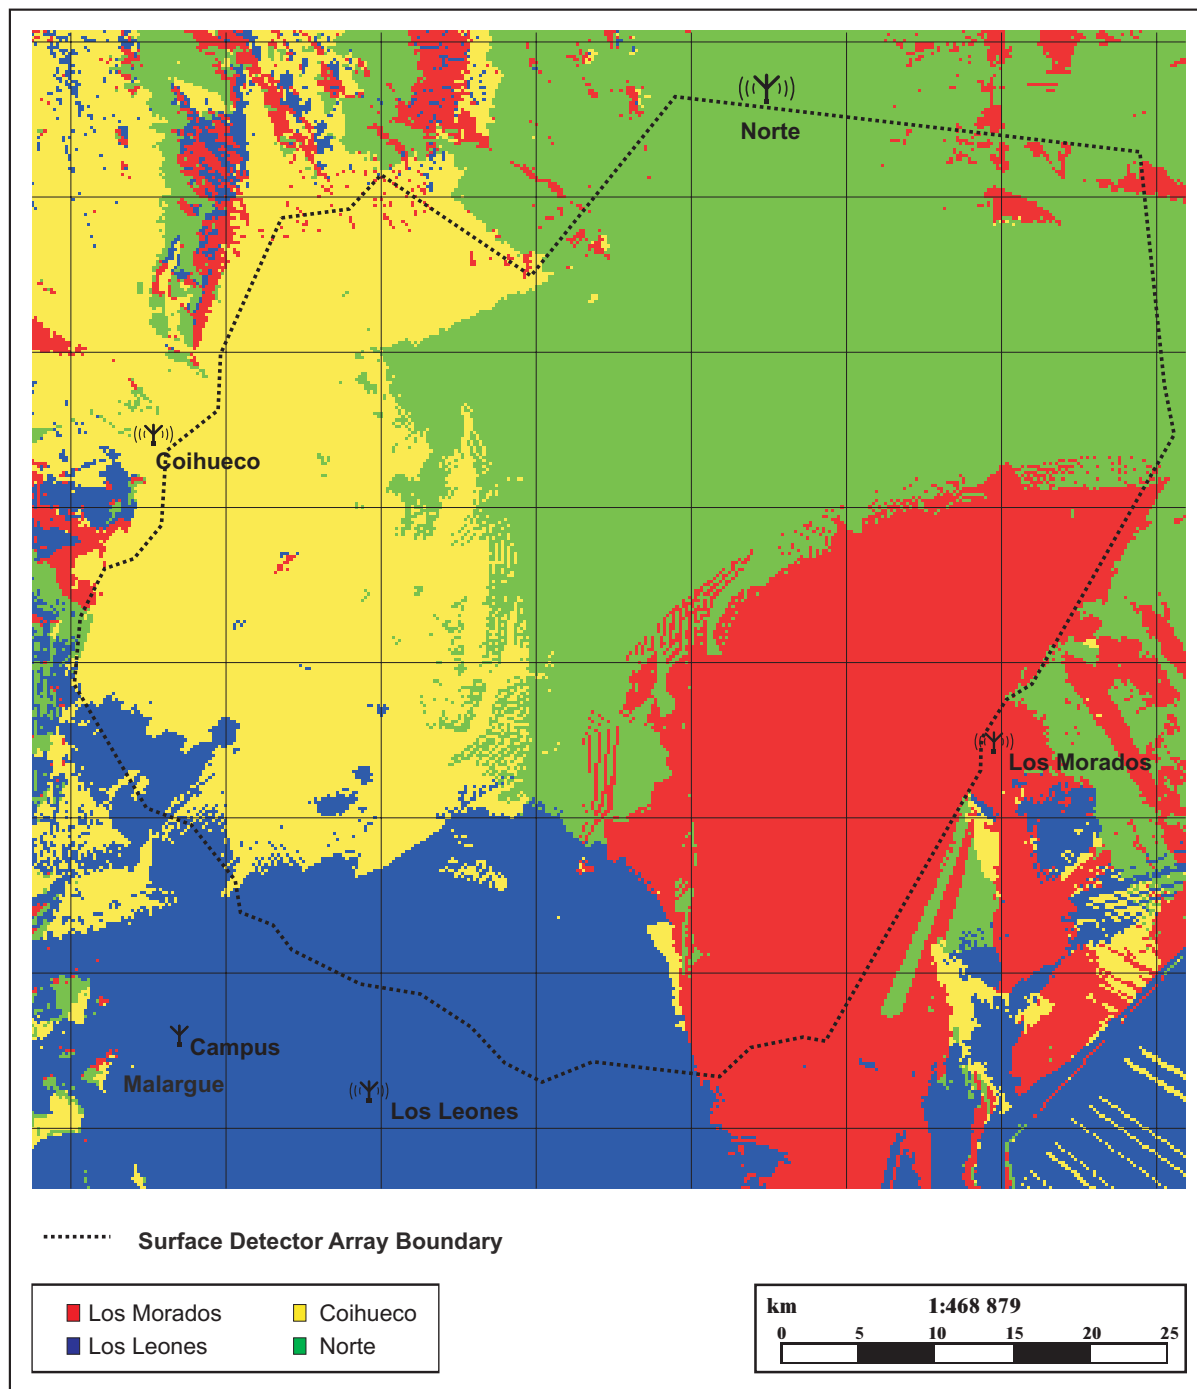


Figure 3.1.10 :Example multi-transmitter radio coverage map




Figure 3.1.11 : Example of shadowed area revealed by coverage maps




Figure 3.1.12 : Correlation of radio simulations and measurements

REFERENCES

- [1] Nitz, D.: “Triggering and Data Acquisition Systems for the Auger Observatory”, the Xth IEEE Real Time Conference, Beaune, France, 1997

3.2 COMMUNICATIONS SYSTEM REQUIREMENTS

The data-rate requirements of the surface detector array are determined primarily by the “T2” triggering rate of the individual detectors. The presence of an analysis computer at each detector greatly reduces the required bandwidth as local events within the detector must pass through several stages of discrimination before they need to be communicated to CDAS.

The design specification for the *uplink* from each detector to CDAS is a continuously available bandwidth of 1200 bps [1]. A reverse *downlink* data path is also required so that the CDAS system may request full “T3” trigger readouts from those detectors that have collected relevant data. A downlink broadcast capacity of 2400 bps to all detectors is sufficient to meet this requirement.

A bi-directional 2.048 Mbps link is available to each fluorescence detector building via direct connection to the microwave backbone network, which is sufficient for an anticipated requirement of 0.5 Mbps.

REFERENCES

- [1] Nitz, D.: “Triggering and Data Acquisition Systems for the Auger Observatory”, the Xth IEEE Real Time Conference, Beaune, France, 1997

3.3 MICROWAVE BACKBONE NETWORK

The high capacity backbone network supports communications with the surface array WLAN data concentration nodes, and also directly with the fluorescence detector sites. Each backbone link has a capacity of 16 E1, 2.048 Mbps channels, which is equivalent to 480 64 kbps voice channels. However, the backbone network within the Auger Observatory is a data-only network, and data is transferred over the E1 bearer by means of standard TCP/IP internetworking protocol. The network has sufficient capacity for the transfer of both the surface detector WLAN data stream and all data generated by the fluorescence detector eye sites. A generous margin has been built into the capacity requirement to permit support of future upgrades to the Observatory.

The backbone network is implemented using commercial ‘off-the-shelf’ microwave link equipment, operating at 7 GHz. The equipment consists of a dish-mounted microwave radio installed on a communications tower (Figures 3.3.1 and 3.3.2) and a separate shelter-based transceiver unit. A single coaxial cable provides power to the radio and links it to the shelter-mounted transceiver. From there the data is distributed across the various Ethernet ports.

This section details the rationale behind the backbone network topology, designed to provide satisfactory performance whilst minimizing costs.



Figure 3.3.1 : Dish-mounted microwave radio on the campus tower



Figure 3.3.2 : Top view of 0.6m diameter microwave dish and radio unit

3.3.1 Considerations in Microwave Systems Planning

There are three key issues that must be taken into account when examining the feasibility of a microwave link :

1. Obstruction of the path of the microwave beam by physical objects;
2. The free-space path loss as a function of distance between the transmitter and receiver;
3. Link performance during fades (e.g. due to propagation and meteorological factors).

In order to satisfy the clear path criteria, no objects (terrain, trees, buildings) should be within the first *Fresnel zone*, which is a hypothetical ellipsoid encompassing the straight-line path between transceivers, with the transceivers located at the focal points. In the vertical plane, the radius of the 1st Fresnel ellipse, at a distance of d_1 and d_2 from the transmitter and receiver respectively, is given by Equation 3.3.1, [1].

$$F_1 = 2\sqrt{\frac{\lambda d_1 d_2}{d_1 + d_2}} \quad \text{Equation 3.3.1}$$

Where λ is the wavelength in the same units as d_1 and d_2

Features outside this ellipse have little effect on the propagation in terms of fades. *Solway*ⁱ propagation simulation software was used in conjunction with a digital elevation mapⁱⁱ in order to evaluate the path loss of the various links and to ensure an unobstructed Fresnel ellipse on each path (Figures 3.3.4 and 3.3.5). All combinations of possible links within the network were evaluated to find a solution that linked all the required sites with the minimum number of paths and with the shortest lengths. Some paths were found to be unusable due to terrain obstructions, whilst others were too long. Based on these simulations the link topology shown in Figure 3.3.3 has been adopted for the microwave system [2]. It consists of 4 microwave links arranged as two ‘arms’ : one servicing both Coihueco and the Northern Eye, and one servicing Los Leones and Los Morados. Both arms terminate at the Observatory Campus in Malargüe, where the data is routed to CDAS.

ⁱ The Solway propagation analysis tool was provided by T.Tozer, Dept. of Electronic Engineering, University of York, York YO10 5DD, UK

ⁱⁱ The digital elevation map was provided by A. Filevich, Laboratorio TANDAR, Commission National Energia Atomica, Buenos Aires, Argentina

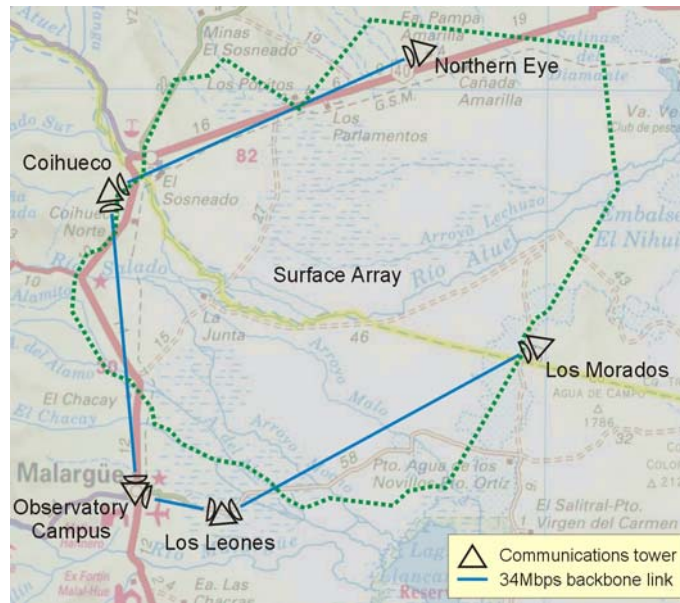


Figure 3.3.3 : Microwave backbone network link topology

The positions of the microwave towers are given in Table 3.3.1.

Location	UTM Position (WGS 84)		Altitude ⁱⁱⁱ (metres)	Tower Height (metres)
Campus	446954	6075445	1404	50
Coihueco	445324	6114107	1706	20
Los Leones	459235	6071764	1390	41
Los Morados	498900	6094560	1417	50
Northern Eye	485144	6136052	1403	50

Table 3.3.1 : Tower location information

ⁱⁱⁱ Altitude data is derived from the digital elevation map

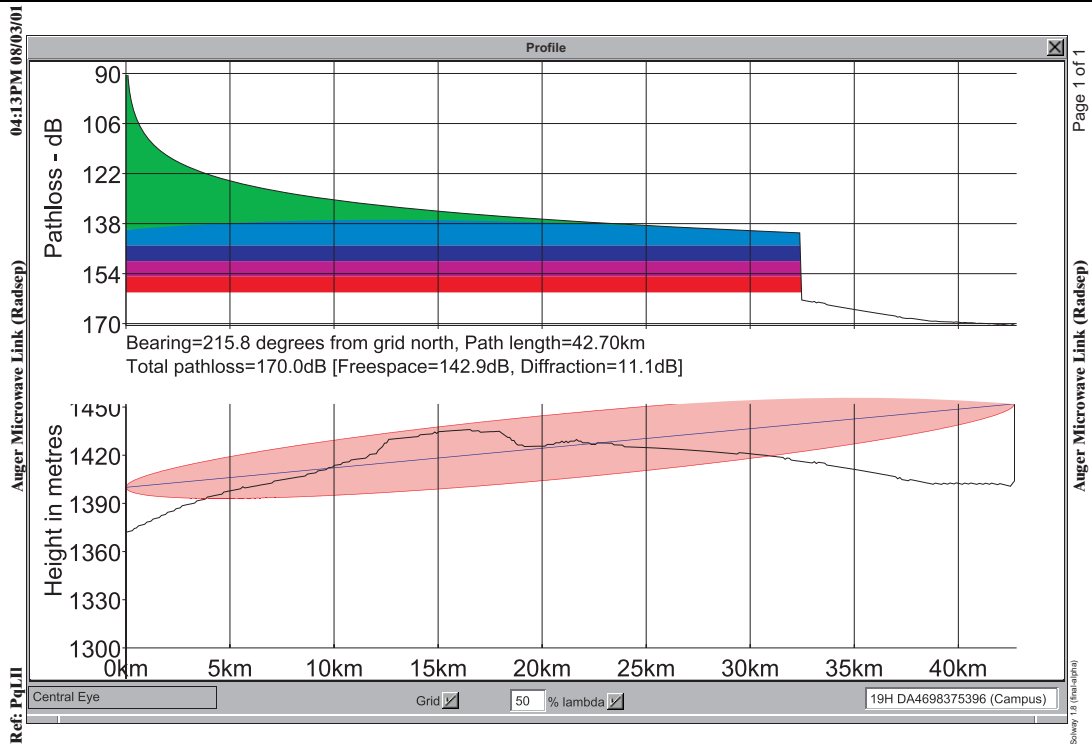


Figure 3.3.4 : Example of a terrain blocked path




Figure 3.3.5 : Example of a path with a clear Fresnel ellipse, Coihueco to Campus (39 km)

3.3.2 Backbone Network Performance

As the links carry all the data from the surface detectors and fluorescence sites back to the Campus, it is vital that their overall availability and fade tolerance is high. When discussing performance objectives for microwave links it is important to draw a distinction between *reliability* and *availability*. Short-term *unreliability* outages, e.g. due to multi-path fades, do not necessarily cause data loss. Long-term *unavailability* events, defined as greater than 10 consecutive seconds of signal outage, (e.g. due to rain fading), result in temporary loss of the link and consequent data loss. *Outages* are defined as periods in which the bit error rate (BER) of the receiver exceeds 10^{-3} . The frequency and duration of outage occurrences exhibit seasonal variations, and performance is often specified for the ‘worst month’ based on long-term regional meteorological data.

To determine a realistic level of link performance, it is useful to consider the targets set for commercial microwave systems. Short systems with 2-5 hops may be designed for something in the order of 99.999 % per-hop path reliability equating to an outage of 320 seconds/year. Such objectives are typical of those used in telephone, utility, and public safety networks. For other services much lower path reliabilities may be acceptable, approaching 99.99 % or about 1 hour/year, [1]. However, the reliability figures are often dominated by outages due to multi-path fading which are of short duration (typically 1-2 seconds per event). These short outages can often be handled transparently (to the user) by the error correction protocols. Outages due to rain fading last longer and hence can cause actual data loss, but are of low total duration. They are estimated to be approximately 60 seconds/year in the 7 GHz frequency band for 40 km links using similar equipment to that used in the backbone links. Consequently, even if the outage statistics suggest many hundreds (or thousands) of seconds of link unavailability, the actual period of data loss on a data-only network will be much lower.

3.3.3 Backbone Network Availability and Reliability Objectives

The International Telecommunications Union (ITU) provides recommendations for availability and reliability objectives. The ITU *availability* (interruptions > 10 seconds) objective for a high grade circuit is given by [3] :

$$\text{Availability} = (100 - (0.3 \times L / 2500)) \% \quad \text{Equation 3.3.2}$$

Where L is the length of the link in km.

For the Southern Auger Observatory, using the proposed microwave links, a total link length of 140 km results in an objective of 99.98 % availability. For comparison, the surface detector communications has a target of 99.8 % availability.

The ITU G.821 (High Grade) performance objective for *reliability* is specified in terms of the total number of outage seconds during the worst month on a one-way circuit [4].

$$\text{Outage} = (0.054 * L / 2500) \% \quad \text{Equation 3.3.3}$$

Where L is the length of the link in km.

The performance of the various proposed links was evaluated using software tools from three different manufacturers of microwave equipment [5][6][7]. The tools were found to agree closely with each other based on the specifications of the recommended equipment.

The short links meet the ITU recommendations for high-performance link reliability (Table 3.3.2). The longer links, however, fail to meet the recommendation by a factor of 10-15. As discussed above, however, these outages do not necessarily precipitate in actual data loss, and it is expected that they will have little or no overall impact on the availability of the microwave backbone. The compound availability (based on rain outage) and outage statistics are shown in Table 3.3.3.

Link From To		Distance (km)	Annual Unavailability (rain outage) (%)	Annual Outage (sec/yr)	Worst Month Outage (%)	ITU Reliability Target (%)
Coihueco	Campus	38.75	0.000210	1354	0.00760	0.00084
Los Leones	Campus	12.72	0.000001	7	0.00005	0.00028
Los Morados	Los Leones	46.89	0.000080	2851	0.01500	0.00102
Northern Eye	Coihueco	tbd	-	-	-	-

Table 3.3.2 : Predicted performance and implementation information

Link to Campus	Compound outage to Campus (sec/yr)	Compound unavailability (rain outage) to Campus (sec/yr)	Compound Availability of the link (%)
Coihueco	1354	65	99.9998
Los Leones	7	1	100.000
Los Morados	2858	26	99.9999
Norte	4483	44	99.9999

Table 3.3.3 : Compound outage and availability statistics for each site linked to the Campus

If it is assumed that all short term multi-path fades will be accommodated by the network protocols, the overall residual link availability is 99.999 %, equivalent to 220 seconds per year of lost data, which is on par with specifications set for the commercial telecommunications systems and exceeds the ITU recommendations for links of this length.

In reality, we may find under operational conditions that the network does suffer from slightly increased loss of data due to short term fades. However, even an order of magnitude departure from expected performance would still only constitute a network unavailability of less than 1 hour/year which is nevertheless a sensible availability target that would have very little impact on the data-taking capabilities of the experiment.

REFERENCES

- [1] Hall, M. P. M., Barclay, L. W., Hewitt, M. T.: “Propagation of Radiowaves”, IEE, London, 1996, pp68-70.
- [2] Clark, P.D.J., de Bruijn, L.: “An Optimised Microwave Backbone Link Topology for the Southern Pierre Auger Observatory”, GAP 2001-043, 2001.
- [3] ITU-R Rec. F.557-4, “Availability Objective for Radio-Relay Systems”, Geneva, 1997.
- [4] ITU-R Rec. F.634-4, “Error Performance Objectives for Radio-Relay Links Forming Part of a High-Grade Circuit”, Geneva, 1997.
- [5] PathCalc 6.1, Digital Microwave Corporation, San Jose, CA, USA.
- [6] StarLink 2.1, Harris Corporation, Melbourne, FL, USA.
- [7] M-L Perf 7A, Ericsson Microwave Systems AB, Moelndal, Sweden.

3.4 SURFACE DETECTOR WIRELESS LAN

3.4.1 Overview

The surface detector communications network is a hybrid design consisting of a proprietary wireless LAN operating in the 915 MHz ISM band together with commercial microwave equipment.

Figure 3.4.1 shows the overall network topology of the surface detector communications system. The system provides a bi-directional mechanism to transfer data between each surface detector station and the CDAS at the Observatory campus.

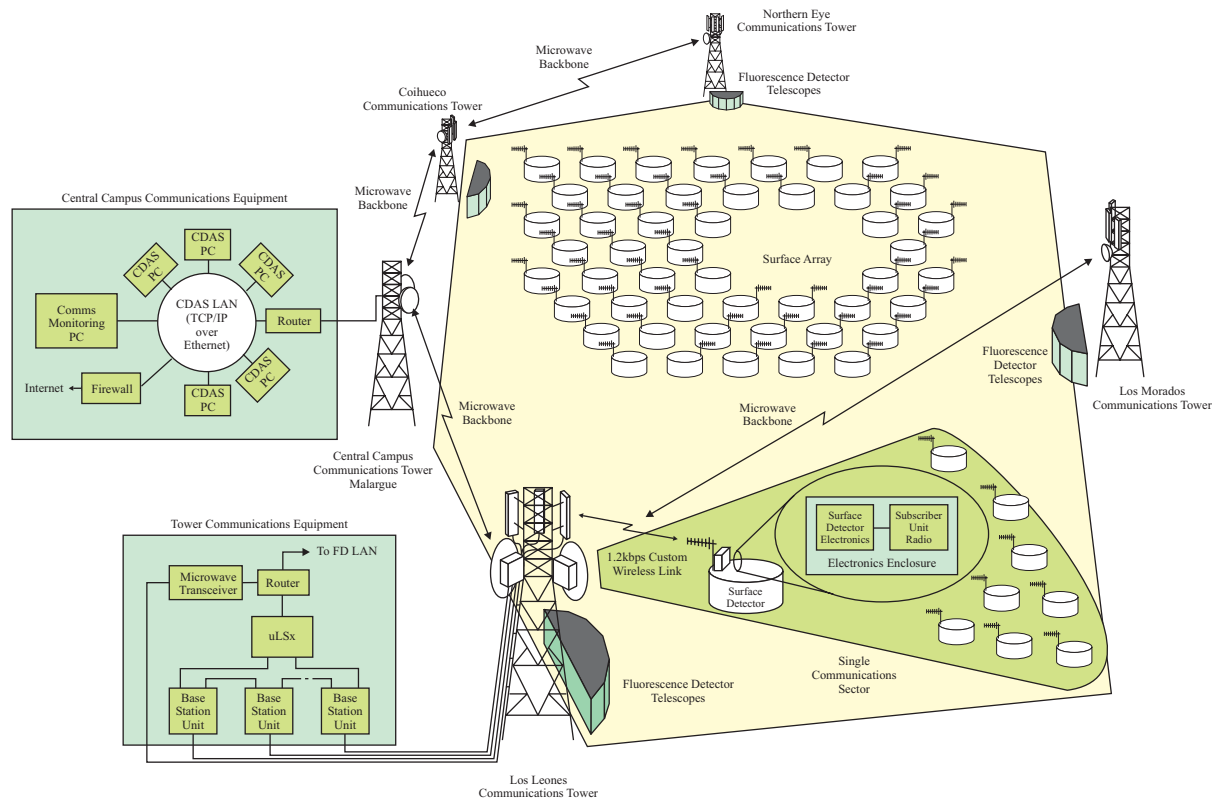


Figure 3.4.1 : Overall communications system network topology

Cosmic ray physics data recorded at each surface detector is processed by the local station electronics module and transferred to a WLAN subscriber unit via a serial link. The subscriber unit broadcasts this data to a corresponding base-station unit located at a distant communications tower up to 30 km away from the detector.

A custom on-air protocol provides a 1200 bps contention-free communications link between each subscriber unit and the base-station unit, thereby guaranteeing a continuous link between the detector and the CDAS.

Each base-station coordinates the wireless links within a sector of the observatory. One sector will serve typically 57 subscriber units, although this may be increased to 68 where necessary. A single communications tower will support a *cluster* of up to 8 base-stations, and in this manner multiple clusters of base-stations provide communications to all 1600 surface detector stations across the array.

3.4.2 Wireless LAN Link Analysis

In order to develop a surface detector communications network topology, it is necessary to perform the following:

- Determine the operational limits of the communications system;
- Establish availability criteria for the network;
- Generate a physical network layout to meet both the above objectives.

3.4.2.1 Link Budget

The decisive factor in determining the feasibility of a radio communications link is the expected fraction of power arriving at the receiver compared to the power broadcast at the transmitter. This ratio is the *link budget*, usually expressed in dB. The *path loss*, which is governed by the physical path separating the radios, should not exceed the link budget and hence a *maximum allowable path loss* figure is derived, beyond which communications system performance will be impaired.

The maximum allowable path loss, PL_{\max} , can be calculated using Equation 3.4.1, which is general to any radio communications system (Figure 3.4.2).




Figure 3.4.2 : A generalised radio link

$$PL_{\max} = G_{tx} + G_{rx} - F_{tx} - F_{rx} + P_{tx} - RSS_{\min} \quad \text{Equation 3.4.1}$$

Where PL_{\max} , F_{tx} and F_{rx} are specified in dB, G_{tx} and G_{rx} in dBi, P_{tx} and RSS_{\min} in dBm.

- G_{tx} and G_{rx} are the gains of the transmit and receive antennas respectively
- F_{tx} and F_{rx} are the losses in the feeder cables connecting the transmitting and receiving radios to the antennas
- P_{tx} is the transmit power
- RSS_{\min} is the minimum received signal strength required at the receiver to operate at a specified error rate

The above parameters are quantified in the following sections. As the radio network is bi-directional (subscriber unit \leftrightarrow base-station), it is necessary to analyse both the uplink path (subscriber unit to base-station) and the downlink path (base-station to subscriber unit).

3.4.2.1.1 Antenna Gain

The main antenna parameters are the 3 dB beam width and maximum gain. The 3 dB beam width is the angle over which the main lobe of the radiation pattern is within 3 dB of maximum gain. The maximum gain is usually expressed in dB relative to an isotropic radiator (dBi) and is limited by cost, physical dimensions and the required beam width.

Antennas are usually horizontally or vertically polarized; the latter was selected for the surface detector antennas. This is normal for terrestrial communications in the VHF-UHF range as ground reflections are less likely to cause destructive interference at the receiving antenna. The topology of the surface detector communications network – a single base-station servicing a large number of subscriber units – favours the use of high performance antennas at the base-station end of the link (with associated higher costs). This allows the use of relatively low cost antennas at the subscriber unit end, where the volumes involved are much greater.

From the subscriber unit perspective, a high gain antenna is advantageous for power efficiency reasons. Such antennas typically exhibit narrow beam widths, and this is in fact desirable as each subscriber unit only needs to communicate with a single predefined base-station. However, in order to simplify installation, the beam width should not be so narrow as to make antenna alignment critical. Furthermore, cost is a major issue for the surface detector system, and therefore a relatively inexpensive Yagi antenna is employed at each site. This antenna has a gain of 12 dBi, a beam width of 30° and a physical length of 70 cm.

The antenna requirements for the base station are somewhat different. First and foremost, the beam width should be sufficiently wide to service an entire sector. The horizontal gain, however, should be as high as possible to yield a greater maximum allowable path loss. Both criteria are satisfied through the use of a co-linear panel antenna for each base-station. The unit specified has a gain of 17 dBi and a beam width of 90°. Co-linear panel antennas are used commercially for GSM cellular mobile communications, and the selection of such a commonly used antenna helps to reduce system costs.

3.4.2.1.2 Feeder Loss

The loss of the coaxial feeder cable between a radio and its antenna is linearly related to the frequency of operation and the feeder's length, which in turn is dependent upon the required antenna elevation and the radio's physical position. In the case of the surface detectors, the mast height is 3.5m above the ground (optimised during system commissioning – see § 3.4.5), and the radio is located with the local station electronics module on top of the detector. This has permitted relatively low cost LRM400 (Times Microwave) coaxial cable to be used for production, with a total feeder loss of approximately 1 dB at 915 MHz.

Experimental Astrophysics (Ppd) 1 - 24

The Pierre Auger Project TDR

The base station unit antennas are tower-mounted in order to provide good coverage, but the improvement in coverage lessens as antenna heights are increased beyond about 50 metres [2]. Tower costs also rapidly escalate above this height. The towers are therefore restricted to 50 metres, and this allows high grade Andrews Helix feeder cable to be used, limiting the feeder loss over 50m to approximately 1 dB at 915 MHz.

3.4.2.1.3 Transmit Power

The transmit power is limited by FCC licensing regulations to a maximum effective isotropic radiated power (EIRP) of +36 dBm [1], where the EIRP is given by:

$$P_{EIRP} = P_{tx} + G_{tx} - F_{tx} \quad \text{Equation 3.4.2}$$

Hence, taking into account antenna gains and feeder losses, the transmit power, P_{tx} , is calibrated to +25 dBm for subscriber units (antenna gain = 12dBi), and +20 dBm for base-station units (antenna gain = 17dBi).

3.4.2.1.4 Required Receive Signal Strength

The required receive power is determined by several factors: the RF performance of the radio, the required bit rate, the maximum acceptable raw bit error rate and the modulation scheme employed. For Version 4.1 of the Auger WLAN radio transceiver, the receive power necessary to satisfy the bit error rate requirements is approximately -100 dBm at the antenna connector (see §3.5.1).

3.4.2.2 Maximum Allowable Path Loss

From the above data we can calculate the maximum path losses for the uplink and downlink using Equation 3.4.1.

- Maximum uplink path loss (subscriber unit to base-station)

$$PL_{max, uplink} = 12 \text{ dBi} + 17 \text{ dBi} - 1 \text{ dB} - 1 \text{ dB} + 25 \text{ dBm} - -100 \text{ dBm}$$

$$PL_{max, uplink} = 152 \text{ dB}$$

- Maximum downlink path loss (base-station to subscriber unit)

$$PL_{max, downlink} = 17 \text{ dBi} + 12 \text{ dBi} - 1 \text{ dB} - 1 \text{ dB} + 20 \text{ dBm} - -100 \text{ dBm}$$

$$PL_{max, downlink} = 147 \text{ dB}$$

The system must function satisfactorily based on the ‘weakest link’, hence, the maximum path loss in the downlink is considered. We now have a theoretical path loss figure that can be used in the following discussion of network planning.

3.4.3 Network Propagation Analysis

The design of the surface detector network has involved the use of *Solway*ⁱ RF propagation simulation software in conjunction with a series of on-site path loss measurements and a long-term propagation monitoring experiment. This combination of theoretical predictions and practical measurements prior to deployment permits a high degree of confidence in the expected overall network performance.

ⁱ The Solway propagation analysis tool was provided by T.Tozer, Dept. of Electronic Engineering, University of York, York YO10 5DD, UK

3.4.3.1 Path Loss Spot Measurements

Several series of spot measurements were performed at the site during March 1999 and April 2000, and these were compared with the path loss simulation results calculated by *Solway* to ensure that there were no significant discrepancies. Figure 3.4.3 shows the series of points at which spot measurements were taken. The points on the map indicate the placement of the transmitter:

Px : transmitter placed at Puntilla

Cx : transmitter placed at Coiheuco

Lx : transmitter at Los Leones

Mx : transmitter at Los Morados

Ex : transmitter at Coihueco, measurements taken in the El Chacay woods area

To perform the site measurements, a subscriber unit radio was configured to transmit a continuous signal on a mid-band frequency (915 MHz), and this radio was temporarily installed at several proposed tower sites using a short mast and a surface detector Yagi antenna. The antenna was oriented towards the areas where spot measurements were to be taken. For reception, an identical antenna was connected to a handheld spectrum analyser. Received signal strength measurements could then be taken directly and used to calculate the associated path loss.

Figure 3.4.4 shows a scatter plot of the measured and associated simulated path loss for the various paths. The results correlate very closely, the standard deviation being approximately 5.5 dB. This is well within the 8 dB quoted for *Solway*, further reinforcing confidence in the simulation tools used.

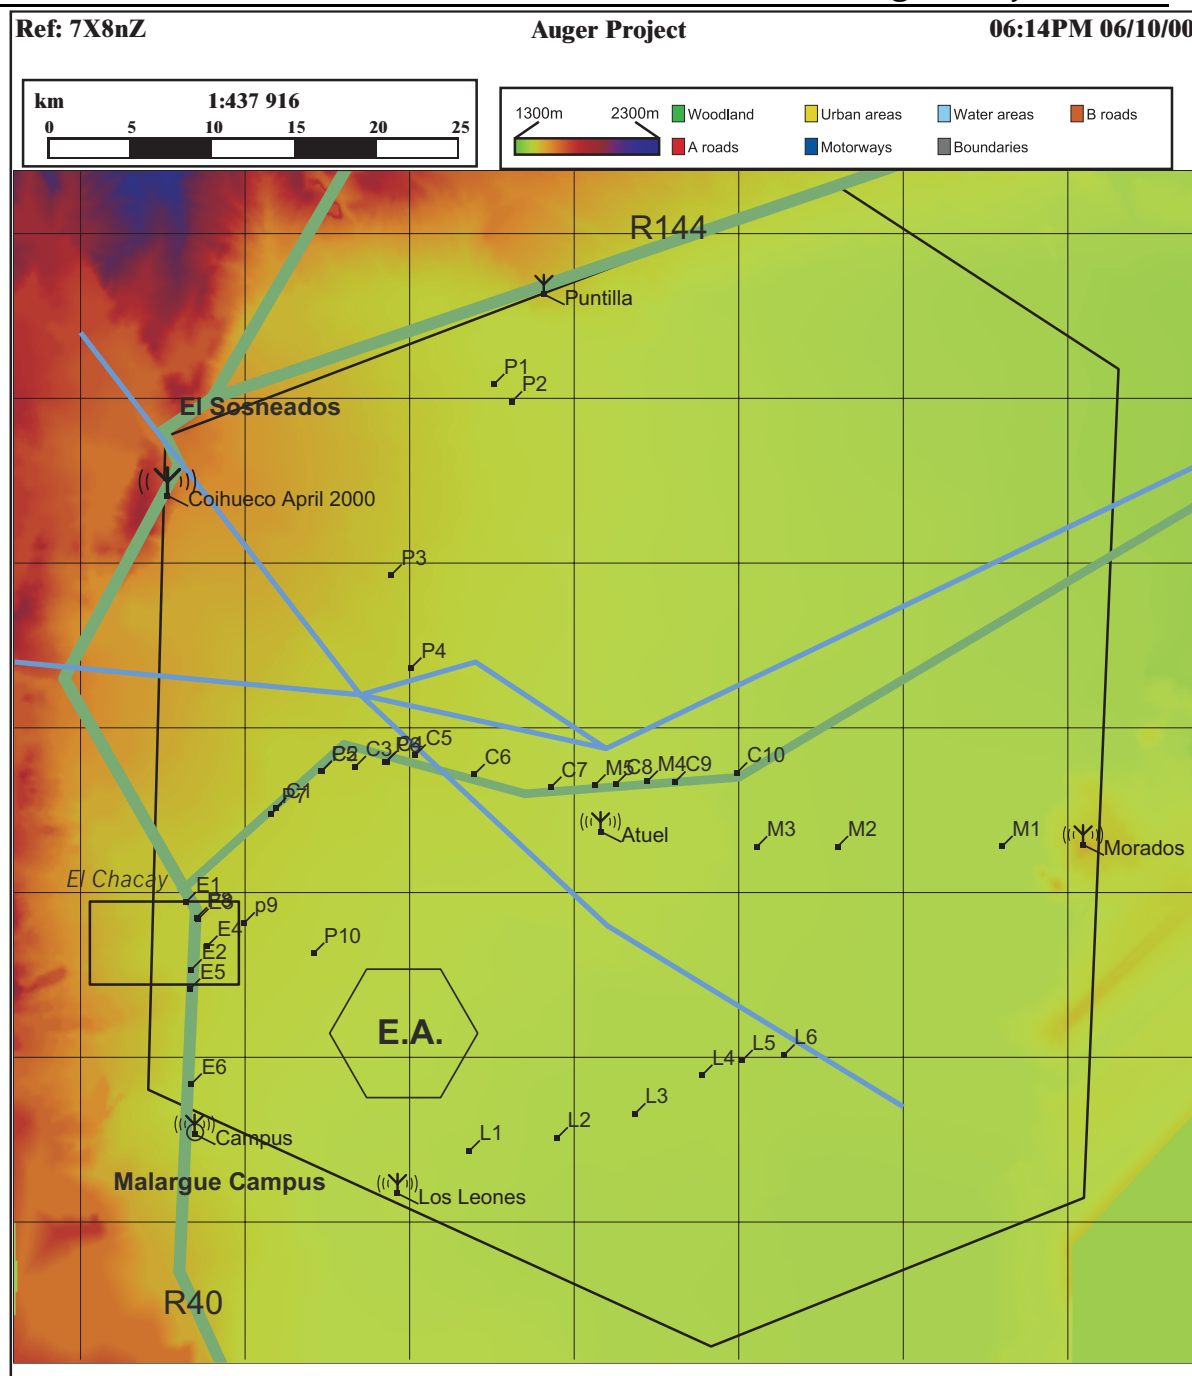


Figure 3.4.3 : Positions of propagation spot-measurements




Figure 3.4.4 : Correlation between simulated and measured path losses

3.4.3.2 Propagation Experiment

The simulation results and spot measurements show ‘average’ and ‘instantaneous’ path losses respectively, but do not provide information on the extent of path loss variation over time. A propagation experiment was designed in order to assess the distribution of signal strength, and hence path loss, as a function of time.

A number of mechanisms exist that will alter the path loss. In general these can be grouped into propagation ‘enhancements’ and ‘fades’. Of particular interest was practical data on the depth, frequency and duration of ‘fades’ as it is these that limit the overall communications system availability. The propagation data gathered allows the prediction of the percentage availability of the communications system based on the average path loss experienced.

The propagation monitoring experiment was commissioned in December 1999. It consisted of two v3.2 radio transceivers configured as base-stations, which were installed in a temporary shelter at Los Leones (Figure 3.4.5), and six radios configured as *out-stations*; these were distributed in a line extending North East from Los Leones at 5 km intervals, out to a maximum of 30 km.



Figure 3.4.5 : Los Leones base-station controller shelter (Dec 1999 - May 2001)

The out-stations and base-stations were powered from 24V battery systems, charged using solar panels (Figures 3.4.6 and 3.4.7). Accurate timing information was provided through individual GPS receiver boards, and the out-station antennas were all mounted at a height of 5 meters. The base-station antennas were mounted at 3.8 meters and 5 meters, allowing information to be gathered about the effectiveness of ‘diversity reception’; fade depths are not necessarily correlated between antennas mounted at different heights. The antennas were all of types similar to the Yagi design eventually selected for final production.




Figure 3.4.6 : Out-station configuration



Figure 3.4.7 : Out-station
(without battery box)

The system employed a fixed frequency (915 MHz) and each of the six out-stations had a predefined transmit slot in order to prevent mutual interference. During this slot a continuous wave (CW) signal was transmitted for 165 ms every second, as shown in Figure 3.4.8. Hence, a signal strength measurement for each out-station could be recorded by the base-stations every second. The data was stored on a PC and periodically archived and sent back to Leeds for analysis.




Figure 3.4.8 : Timeslot architecture for the Propagation Experiment
(Dec 1999 - May 2001)

In May 2001 a major upgrade was performed to the propagation experiment hardware and software. The six out-station radios of the propagation experiment were replaced with new v4.0 radios. The base-station radios were replaced with a single v4.0 base-station, installed in the main Los Leones communications shelter and connected to one of the tower-mounted panel antennas at a height of 39 meters. The software was upgraded to provide full digital data communications, rather than a simple carrier signal strength measurement.

Overall the upgraded system was analogous to that used in the Observatory Engineering Array, and provided valuable data regarding long-term network performance. In addition, the experiment provided a test-bed for the communications network software. Received signal strength was recorded for the out-stations and for the base-stations, once every second. When a bit error occurred, the data packet was logged to allow the bit error rate to be calculated.

Logging was performed using a PC connected to the microwave link backbone, allowing easier access to the monitoring data.

The Propagation Experiment was decommissioned in December 2003.

3.4.4 Network Availability Specifications

An analysis of the data from the propagation experiment in terms of the cumulative distribution function (CDF) for the path losses from the out-stations and the two base-stations is shown in Figure 3.4.9. The height of the ‘shoulders’ of these curves, especially that for the base-station 1 (BS1) with its higher antenna, indicates that the total fade duration is low.




Figure 3.4.9 : Cumulative distribution of path losses from the outstations to the base-stations

It can be seen from Figure 3.4.9 that small fades (less than 2-3 dB) occur relatively frequently and therefore radios operating at the maximum average path loss limit of 147 dB are likely to experience communications disruptions. In order to derive an acceptable system fade margin it is necessary to decide on a target *availability statistic* for the surface detector array communications. In practice no communications system can be planned to have 100 % availability. Broadcasting and mobile services are normally planned for an availability of 90 % to 99 % [3]. For Auger, a higher level of availability is required, limited by cost implications and confidence in the accuracy of the fade statistics.




Figure 3.4.10 : Fade depth distributions at varying distances

The propagation experiment data shows that there are few occurrences of fades with a path loss greater than 10 dB below the mean, corresponding to a CDF of 0.998 (see Figure 3.4.10). Hence for confidence reasons, it was decided to set this as a target availability for the surface detector communications. This corresponds to less than 20 hours/year of data non-availability due to fades, and assumes that all fades, irrespective of duration, cause data loss. In practice this assumption is pessimistic as fades are not necessarily correlated between adjacent radios and so the non-availability is expected to consist of numerous short-duration, isolated occurrences, many of which will be overcome by the error correction protocol.

Therefore, in order to provide 99.8 % communication system availability, a fade margin of 10 dB is added to the maximum allowable path loss of 147 dB, yielding a target maximum path loss of 137 dB.

3.4.5 Network Planning

3.4.5.1 Determination of antenna heights

In order to meet the communications availability requirements for the array, the path loss to each subscriber unit must be within the limits calculated above, including the fade margin. The main external factors that affect path loss are:

- Surface detector mast heights
- Number, position, altitude and heights of the communications towers for the base-stations

For optimal radio coverage of the site it is necessary to strategically locate the communications towers to gain the greatest height advantage. The most cost-effective way of doing this is to make use of local terrain features, and the selected sites for the fluorescence detector buildings provide good positions for locating the communications towers. Moreover, co-siting the towers with the fluorescence detector buildings eliminates the need for separate power sources, simplifies maintenance, and permits the installation of the microwave equipment for the backbone network to the same towers.

A digital elevation mapⁱⁱ of the site was obtained and *Solway* was used to generate path loss coverage plots for various base-station positions and antenna heights, showing the distribution of path loss as a function of position. As established in the previous section, the average path loss was to be better than 137 dB to obtain 99.8% availability.

The overall objective of the planning exercise is to meet the availability requirements at the lowest possible cost. Significant savings can be made on the surface detector end of the communications system because of the large number of units involved. Lowering the surface detector mast height provides a major cost saving as it reduces the amount of mast material and coaxial feeder cable required. Experience from the Engineering Array has shown that the mast height should be at least 3.5 m above ground for safe access to the detector electronics pack. Thus a surface detector antenna height of 3.5 m was used in the simulations.

The high elevation of the site of the fluorescence detector building at Coihueco, 170 m above the terrain of the main array, has allowed a lower tower to be employed at this site, whilst still providing excellent coverage for both the WLAN and the microwave backbone. Simulations indicated that a 20 m tower would be sufficient. Los Leones, in contrast, is only 20 m above the local terrain, and consequently a higher tower (41 m) needed to be specified in order to provide adequate coverage. In the case of Los Morados, the nature of the terrain precluded siting the tower on the peak, as there was no suitable site for the fluorescence detector building close by. Consequently the tower position was moved to a plateau almost 100 m below the peak, and this loss of altitude meant that the tower needed to be increased from an anticipated 20 m to 40 m.

Early simulation work had also considered a tower positioned in the centre of the array, but the excellent coverage provided by Coihueco and Los Morados, and that anticipated

ⁱⁱ The digital elevation map was provided by A. Filevich, Laboratorio TANDAR, Commission National Energia Atómico, Buenos Aires, Argentina

by the the proposed Northern Eye site ('Norte'), means that a central tower will, in all probability, not be required.

A plot of the path loss coverage from the towers located at Los Leones, Los Morados, Coihueco and the proposed site at Norte is shown in Figure 3.4.11. The colours represent path loss in dB, according to the scale at the bottom of the figure, and so, for a maximum permissible loss of 137dB, all green, light blue and dark blue areas are considered 'good'.

With towers at Los Leones, Coihueco, Los Morados and Norte, and a surface detector mast height of 3.5 meters, it is predicted that 99.96 % of the array will have an average path loss of better than 137 dB, giving an overall surface detector communications availability for the entire array of 99.8 %.

3.4.5.2 Network Planning Tool Development

While deployment of detector tanks was confined to the area covered by the Engineering Array, the issue of base-station loading, tank antenna alignment, and marginal radio links did not exist to any significant degree; the tanks were distributed equally over two base-stations located at Los Leones, and all antennas were oriented towards the Los Leones tower, which was visible with the naked eye from the Engineering Array. However, as the deployment of tanks has ramped up, and the array has extended North and North East, a more coordinated approach to base-station allocation has been required. To this end a Windows-based Network Planning Tool has been developed to guarantee that all base-stations of the WLAN will be equally loaded, and that no tank will be serviced by a marginal link.

The Network Planning Tool is used during tank deployment and commissioning. The application takes UTM tank coordinates as its input and returns the base-station to which the tank must be allocated, together with the magnetic bearing to the appropriate tower for the tank antenna.

A screen capture of the Network Planning Tool is shown in Figure 3.4.12.

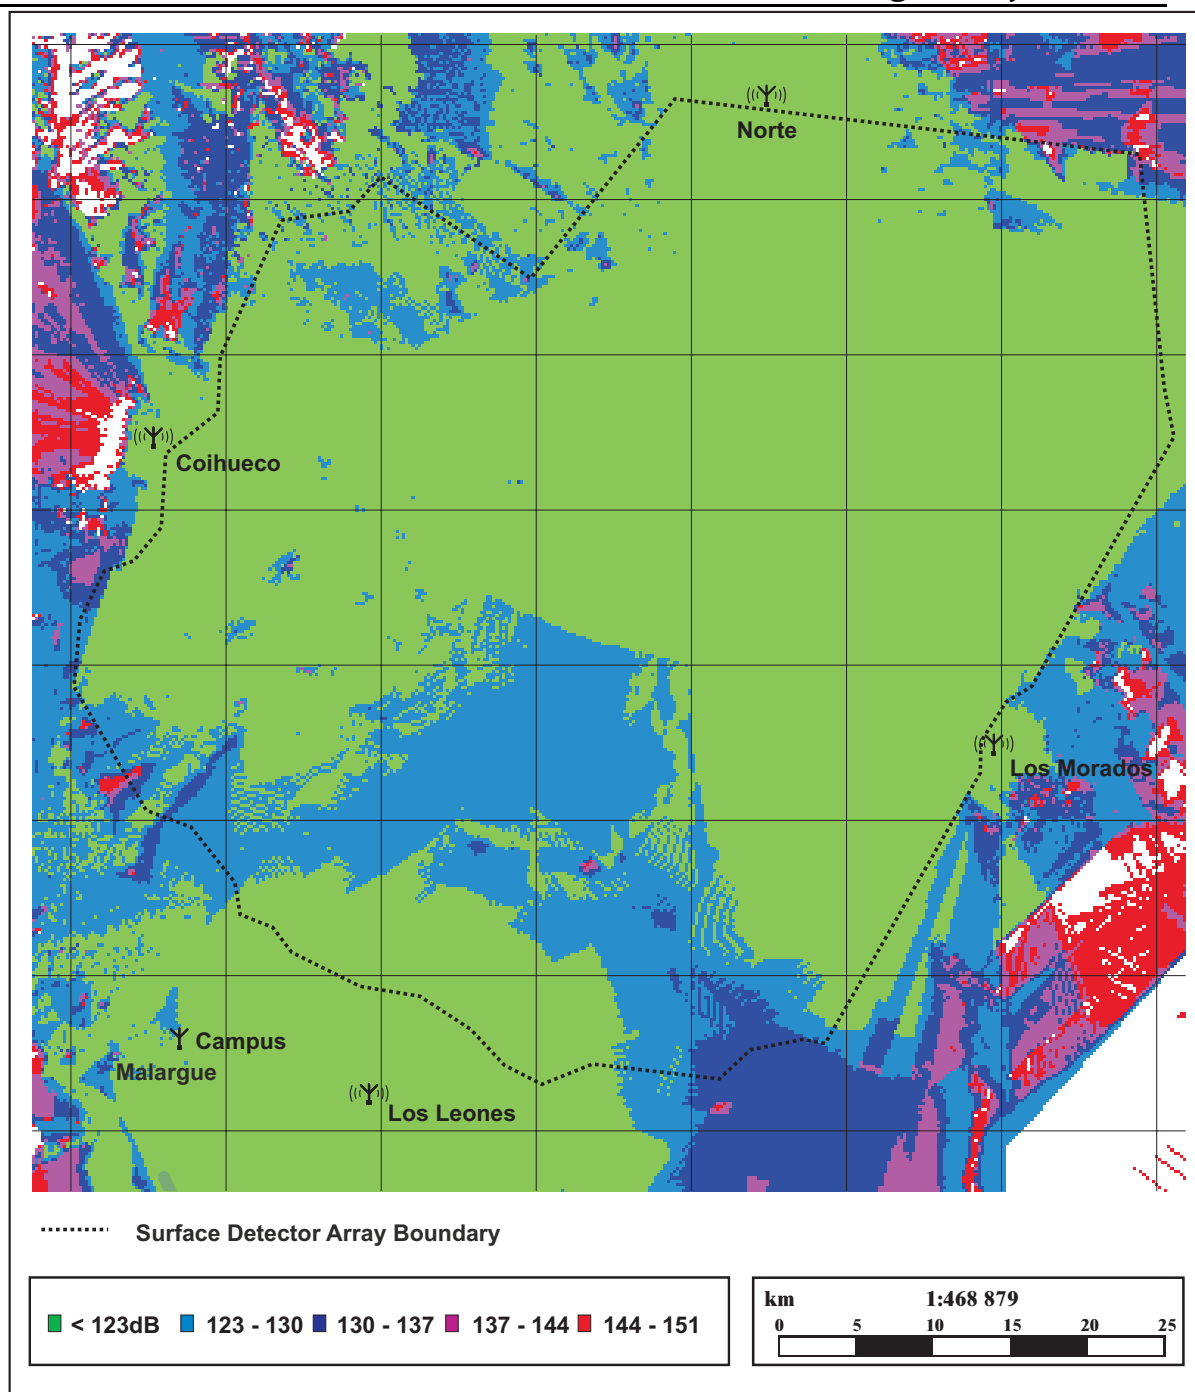


Figure 3.4.11 : Simulated path loss coverage of the site, with towers at Los Leones, Coihueco, Los Morados and proposed Norte position

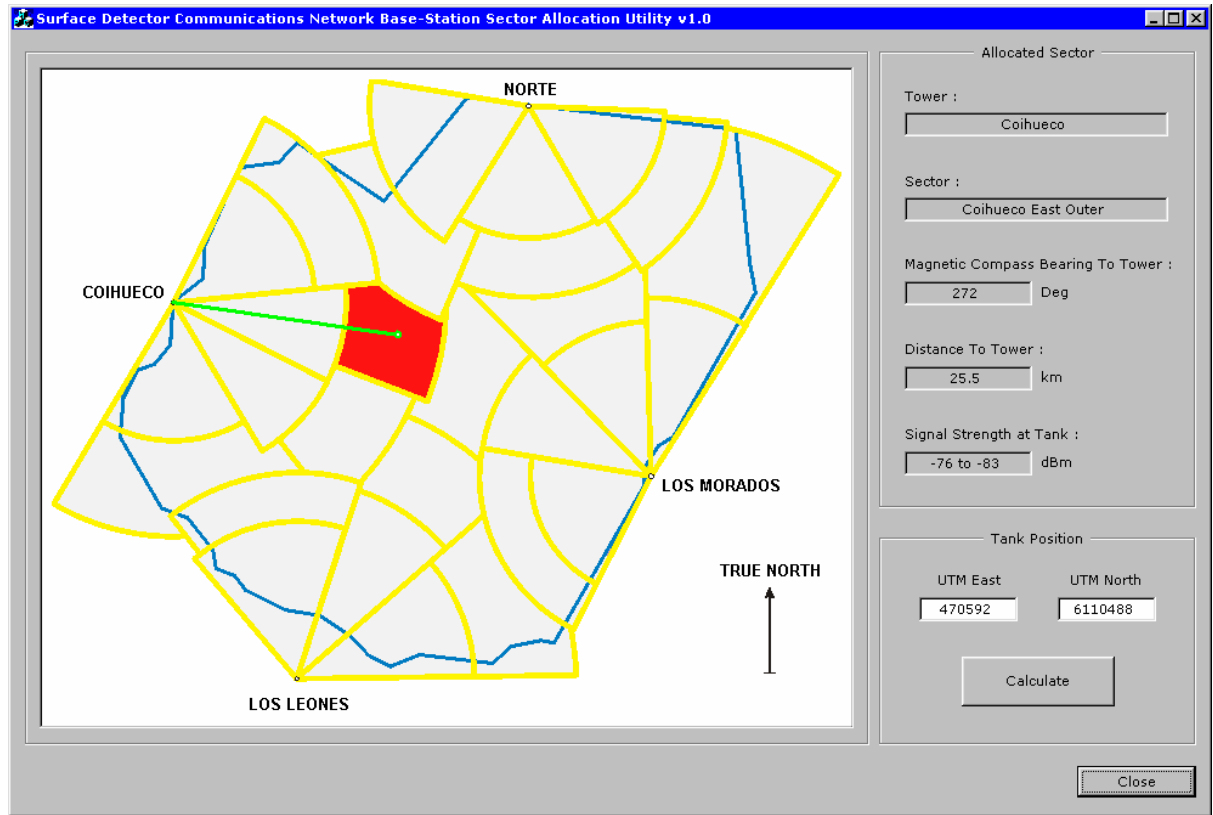


Figure 3.4.12: Screen capture of the Network Planning Tool

The sectorisation of the array, as shown in Figure 3.4.12, is based on the propagation maps and the requirement for equal areas (for equal loading). Each sector covers an area containing typically 57 tanks on a 1.5 km grid, but, as stated in previous sections, a base-station can support up to 68 tanks if required.

The application also returns the expected downlink RF signal strength based on the propagation maps, which are embedded. This can be checked on the Communications Monitoring PC at the Observatory Campus.

A recent modification to the Planning Tool allows the UTM coordinates to be read from a list of tanks contained in a file. The tool then automatically generates an output file specifying the parent base-station and the antenna bearing for each tank in the list.

Figure 3.4.13 summaries the WLAN sectors and the microwave backbone coverage of the site.

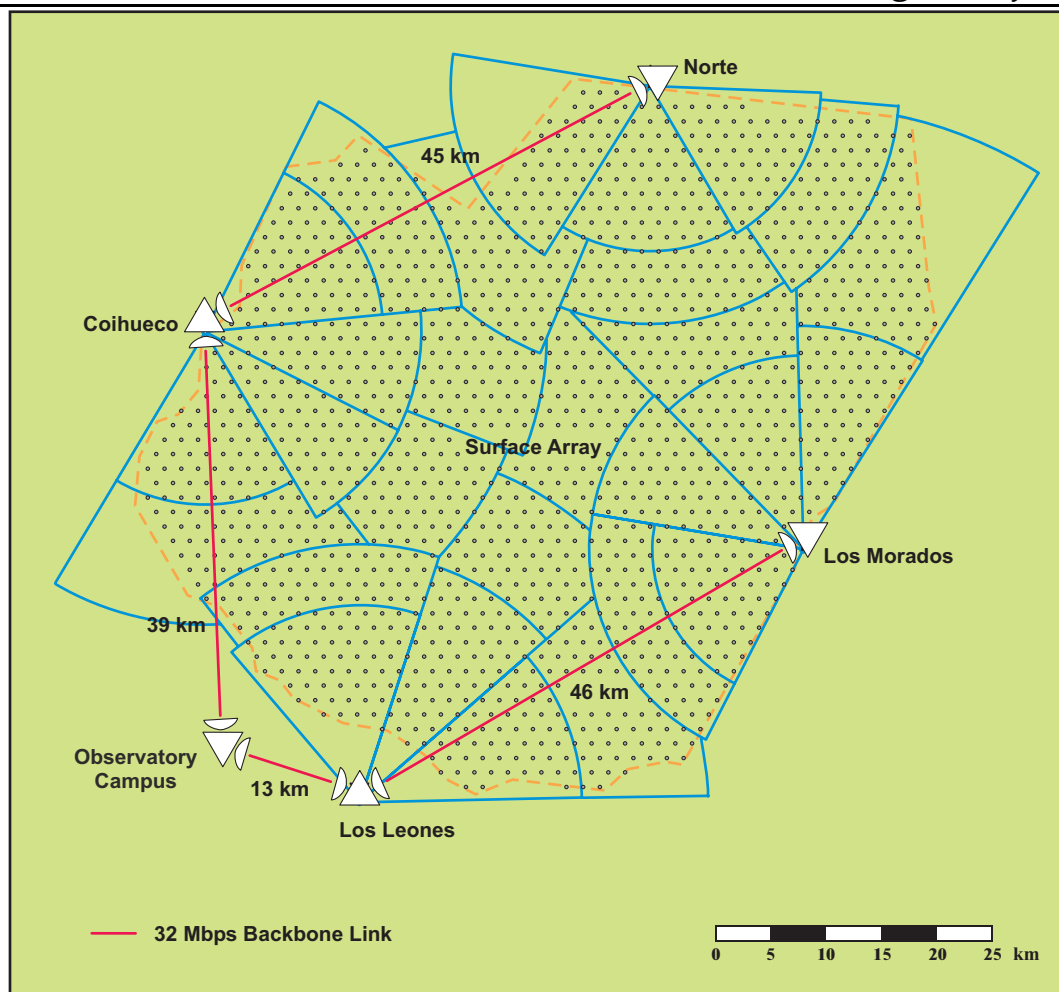


Figure 3.4.13: Combined WLAN sectorisation and microwave backbone

3.4.6 Wireless LAN Communications Protocols

The following section describes in detail the protocols used throughout the surface detector communication system, as *uplink data* flows from the detector to the CDAS and *downlink commands* flow back to the detectors from the CDAS.

3.4.6.1 Local Station ↔ Subscriber Unit Communications

This protocol transfers data between the local station and the wireless LAN radio at each detector. It is the first step in the long journey from a detector in the field to the observatory campus and is based on a custom packetised protocol operated over an industry standard RS232 link. It provides bi-directional communications between the local station and subscriber unit, and packets generated by each unit are of a standard format as illustrated in Figure 3.4.14.




Figure 3.4.14 : Local station ↔ subscriber unit communications packet structure

The *preamble* section of the packet ensures that the receiving unit correctly identifies the start of the packet. The *length* character defines the length of the packet (defined as the number of bytes between the *type* character and the *ETX* character inclusive). The *type* character is used to inform the recipient how it should interpret the information contained within the *data* field of the packet, which may contain between 0 and 249 bytes with no restricted characters. A 32-bit cyclic redundancy check (*CRC*) code is generated by the transmitting unit and used by the recipient to check packet integrity. The CRC error detecting algorithm is guaranteed to detect all 1, 2 and 3 bit errors and is based on the ethernet 0xEDB88320 standard [4]. An ETX end-of-transmission character is used to mark the end of the packet.

CRC codes are used throughout the communications system as a computationally inexpensive and robust method of detecting errors within a block of data. A *generator polynomial* is used to generate a code that is appended to the data block before transmission. The receiving unit generates a CRC from the data in an identical manner and compares it to the transmitted code. Inconsistencies in the codes identify corrupted transfers. CRC codes do not provide error-correcting capabilities and packets containing errors are discarded and re-requested, to provide an error-free link.

The local station \leftrightarrow subscriber unit communications protocol is described in defined and detailed in a separate document [5].

3.4.6.2 Subscriber Unit \leftrightarrow Base-Station Unit (On-Air) Communications

The following sections describe how the subscriber unit \leftrightarrow base-station link is negotiated, maintained and operated for each surface detector station. Two techniques are employed to provide simultaneous and continuous access for all 1600 surface detector stations:

- (i) Within a sector, the radio channel is divided into a number of discrete time periods or *timeslots* whereby each subscriber unit is given access to the base-station in turn.
- (ii) Between sectors, a frequency hopping scheme is employed that ensures that each sector is operating on a unique radio frequency at any given moment, thereby preventing interference between adjacent sectors.

3.4.6.2.1 Timeslot Architecture Overview

Each surface detector station is equipped with a global positioning system (GPS) receiver, which is used to provide the station with highly accurate timing information for use in synchronising the acquired physics data. The free availability of this timing information is also exploited within the communications system to synchronise radio traffic activity within each sector and across the network as a whole.

The GPS receiver provides a synchronous pulse every second, accurate to approximately 10 ns, and this signal is used to identify the start of a *frame* within the communications system. Contention-free communication between subscriber unit and base-station is achieved using a time division multiple access (TDMA) structure. A frame, which is one second in duration, contains 85 timeslots divided into 9 categories, as shown in Figure 3.4.15. The frame principally contains 68 uplink slots and 6 downlink slots. Within a sector, each subscriber unit is allocated 1 uplink slot, and one slot will support a 1200 bit data payload. Thus, the uplink physics data payload requirement of 1200 bps is satisfied. The downlink slots are used by a base-station for broadcast messages to all subscriber units within its sector. A further 11 slots are reserved for network management and maintenance, and packet error control.

Experimental Astrophysics (Ppd) 1 - 40

The Pierre Auger Project TDR

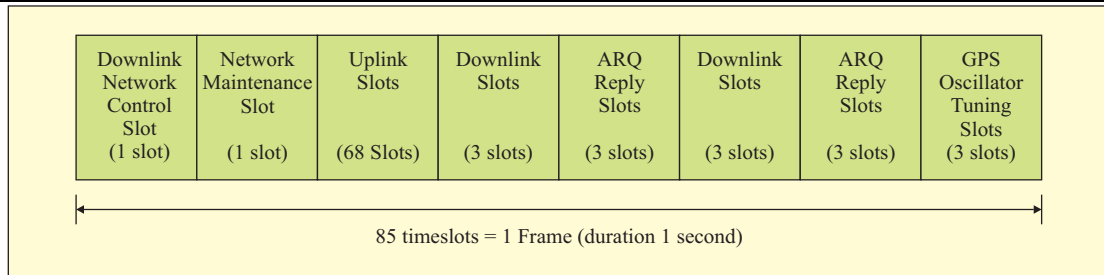


Figure 3.4.15 : TDMA frame timeslot structure

A frequency hopping spread spectrum system is employed, whereby a different transmission frequency is used for each timeslot. At the start of each timeslot, the next transmission frequency is selected from a pseudo-randomly ordered set of 51 channels. This process, which is a regulatory requirement to operate in the ISM band, effectively increases the transmission bandwidth of the signal, offering more resistance to narrow band interference.

All base-stations and subscriber units are programmed with an identical circular frequency hopping table, and each unit uses this table to select its transmission and reception frequencies. To prevent multiple base-stations and subscribers from selecting the same transmission frequency at the same time, and consequently jamming each other, they must be synchronised and co-ordinated. Synchronisation is achieved by using a simple pointer offset technique that offsets the start frequency for each sector with respect to all other sectors; this is illustrated in Figure 3.4.16.




Figure 3.4.16 : Base-station hop set synchronisation.

At each communications tower within the Observatory, a μ LSx unit presents GPS timing information to the local base-stations on its E1 ring, and all base-stations exploit this timing information to generate a synchronised and global (network-wide) pointer into the common frequency hopping table. Each base-station then adds its own unique offset to this global pointer position, based on its electronic serial number (ESN). Thus, up to 51 base-stations may be accommodated across the network before spectrum resources have to be shared.

Once the base-station has synchronised its frequency table pointer with the rest of the network, the position of the pointer at the start of the first frame is stored. This pointer is incremented at the start of subsequent frames and used to define the frequency used during the first timeslot. This ensures that all hopping frequencies are used equally, a further regulatory requirement for operation in the 915MHz ISM band.

3.4.6.2.2 Network Acquisition

To ease network commissioning and maintenance, all subscriber units are deployed with a common software build. This offers the advantage that any subscriber unit may be placed at any surface detector without regard to the wireless LAN sector that it resides in. However, this also means that it is not possible for it to have prior knowledge of the timeslot allocated to it for uplink transmission, nor the base offset to the local base-station frequency hopping table. Acquisition of the subscriber unit transmission slot is therefore negotiated, under control of the base-station, using a dedicated network control packet. Figure 3.4.17 shows the data structure of the network control packet.

Experimental Astrophysics (Ppd) 1 - 42

The Pierre Auger Project TDR




Figure 3.4.17 : Data payload of broadcast network control packets

The network acquisition process is further complicated by the fact that, due to overlap of antenna coverage patterns, a subscriber unit will successfully be able to receive network control information from multiple base-stations. The subscriber is prevented from connecting to a base-station that is not its ‘parent’ through the use of a sector member list. This is broadcast by each base-station and contains a list of the IDs of subscriber units that are permitted to connect to that base-station. The IDs are derived from electronic serial numbers allocated during testing and calibration. Each ID in the list is associated with a timeslot number, as shown in Figure 3.4.17.

A subscriber unit enters network acquisition mode when it first powers up, when commanded to reset by the local station, when the unit is removed from a member list, or when the commanded to do so over the air from the Network Monitoring PC at the campus. In acquisition mode, normal frequency hopping operation is disabled and the radio remains on a fixed frequency listening to base-station network control packet transmissions as they go past. During this initial phase, the subscriber unit has no knowledge of where the base-stations are in their hopping cycle. This approach guarantees that the static subscriber unit will eventually hear a network control packet from a base-station.

As soon as the subscriber unit decodes a network control packet, it begins to track the base-station which sent it by re-tuning its receiver channel every second, according to the hop sequence, which is now synchronised to the base-station. The base-station’s member list is spread over three frames, and so the subscriber unit must track the base-station for 3 seconds to determine whether the base-station is its parent. If the subscriber unit detects its station ID in the network control packet, it uses the associated slot number information to determine which slot it should use to inform the base-station that it wishes to join the network, and continues hopping. If the subscriber does not hear its ID in 3 frames, it stops hopping and waits on the last channel it used until it hears a network control packet from a different base-station. This ensures that the subscriber unit will receive a network control packet from the next base-station in the hop set and that eventually all base-stations will be checked for sector membership (signal strength permitting).

If the base-station successfully receives a ‘request to join’ transmission from a subscriber unit, it acknowledges the subscriber’s arrival by entering the subscriber’s ID in the *acknowledgement of successful connection* part of the network control packet. Once connected, the local station is informed that the network has been acquired and communications then begin.

The network acquisition time is dependent on the number of base-stations that a subscriber unit can detect. The maximum number of frame, F , that a unit takes to join is given by Equation 3.4.3.

$$F = AB + C + D \quad \text{Equation 3.4.3}$$

where,

A = The number of base-station units that the subscriber unit may detect

B = The number of network control packets that the subscriber needs to decode to receive the complete member list (=3)

C = The number of frames for the base-station to acknowledge successful subscriber unit connection (=1)

D = The number of hopping channels (=51)

As an example, if a subscriber unit is able to ‘hear’ the downlink transmissions from 6 base-stations, it will take a maximum of 70 frames for the subscriber unit to join the network under error-free conditions. If a subscriber unit or base-station unit fails to recover a network control packet, network acquisition time is increased.

As part of the network control software, each subscriber unit monitors the number of network control packets that fail to be delivered. If 7 consecutive packets are lost, the subscriber unit disconnects from the network, informing the local station of the disruption and recommences network acquisition in the normal manner. Communications may not take place until the network has once again been acquired.

The base-station member lists are controlled from the Communications Monitoring PC at the Observatory Campus, and so if a subscriber unit fails in the field, it is not necessary to program it’s replacement to reproduce the failed unit’s ID; the ID is simply removed from the base-station member list, and the ID of the replacement unit added.

3.4.6.2.3 Digital Data Modulation and Demodulation

Data Modulation

Once network acquisition is complete each subscriber unit provides a 1200 bps link to the base-station. The flexible architecture of the radio transceiver’s hardware permits any modulation scheme to be used, limited only by the dynamic range and conversion speed of the converters used to transform the baseband signals between the analogue and digital

Experimental Astrophysics (Ppd) 1 - 44

The Pierre Auger Project TDR

domains. The modulation process converts the raw data bits into a finite set of symbols; the number of symbols being dependent on the modulation scheme chosen.

The modulation scheme chosen for the surface detector wireless LAN is *differential quadrature phase shift keying (DQPSK)*. In this scheme pairs of data bits are encoded into 1 of 4 possible phase transitions. Figure 3.4.18 shows the *constellation diagram* for DQPSK.




Figure 3.4.18 : Constellation diagram for DQPSK

During transmission, two consecutive bits of the raw binary data stream are grouped together and used to determine the next symbol phase transition. Assuming the current symbol is (A), binary 00 will cause no phase transition in the next symbol, binary 01 will cause a $+90^\circ$ symbol transition to (B), a 11 will cause a $+180^\circ$ symbol transition to (C) and 10 will cause a -90° symbol transition to (D). The receiver attempts to recover the raw data stream by measuring the change in phase between successive symbols to determine which bit-pair was transmitted.

Receiver Symbol Synchronisation and Data Demodulation

In order for the receiver to correctly demodulate the data, *symbol synchronisation* must be performed. Symbol synchronisation is the process of selecting the energy peak within the symbol period, maximising the receiver's ability to detect symbols correctly. This is achieved using a *modulation-derived synchroniser* (MDS) and is illustrated in Figure 3.4.19.

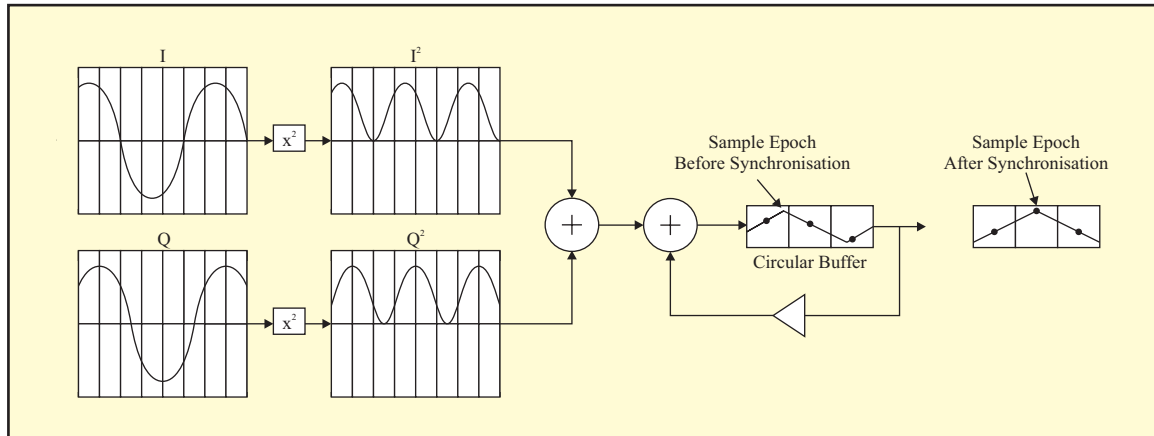


Figure 3.4.19 : Modulation Derived Synchroniser

Three analogue samples per symbol are taken during the symbol reversal period. Symbol reversals are represented by a series of A,C,A.. or B,D,B.. symbol transitions as shown in Figure 3.4.18 and provide the receiver with best-case symbol timing information. A power average over 30 symbols is made by squaring and summing the in-phase (I) and quadrature (Q) baseband sample values. The ratio of the resultant amplitudes for each of the three samples per symbol is used to estimate the timing error. This estimate is used to modify the sample timing of the receiver so that it coincides with the symbol's epoch. When the entire packet has been received, sample timing is restored so that the receiver's ability to identify timeslot boundaries is maintained. Once the receiver gain has been set and symbol synchronisation acquired, the receiver is capable of recovering the raw bit stream.

Demodulation is the process of converting the synchronised symbol transitions back into the recovered raw data stream. This process is complicated by the fact that the local oscillators used to up-convert the baseband signals in the transmitting unit and down-convert in the receiving unit are independent and therefore will not be locked in phase or frequency. This effectively causes the constellation diagram shown in Figure 3.4.18 to rotate at a rate equal to the difference in frequency between the transmit and receive oscillators and in a direction dependent on the sign of the frequency error. Assuming that the incremental error in frequency and phase is small when compared to the symbol period, the corresponding phase transition will be small and the raw data stream may be recovered by estimating the phase change between symbols. The frequency error effectively causes the phase difference between symbols to be greater or less than the transmitted symbol phase depending on the phase transition and the direction of rotation of the constellation diagram. The demodulation process uses a pair of highly processing efficient *look-up tables* (LUTs). These are illustrated in Figure 3.4.20.

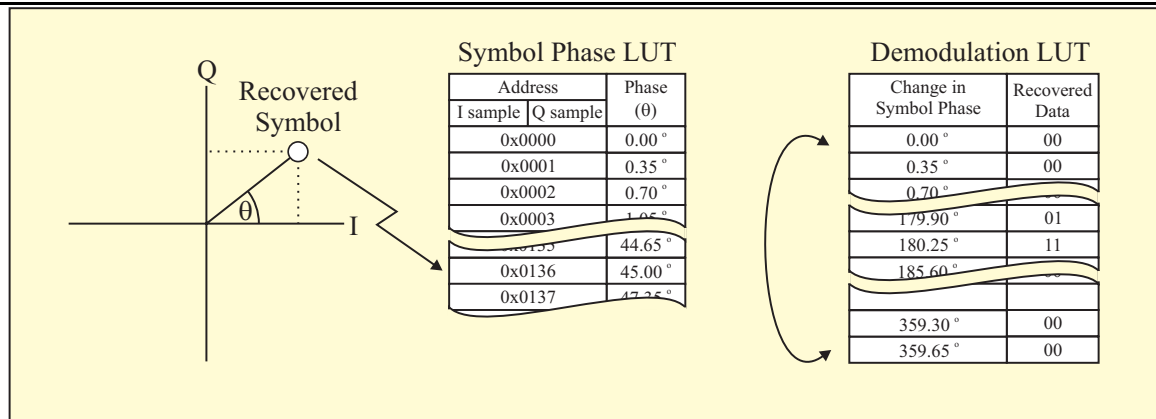


Figure 3.4.20 : Symbol phase recovery and demodulation

The *symbol phase table* is used to determine the instantaneous symbol phase at the synchronised sample epoch. The sample values of I and Q are combined and used to address the table. The value read from the table is the phase of the symbol and the change in symbol phase is calculated by subtracting the phase of the previous symbol from the phase of the current symbol. This result is then used to address the circular *demodulation table*. The table is circular so that a phase change of -90° , for instance, is mapped to exactly the same point in the table as a $+270^\circ$ transition. The value read from the table is the demodulated data and forms the next two bits of the recovered bit stream.

3.4.6.2.4 On-Air Packet Structure

All on-air packets, uplink and downlink, are of the format shown in Figure 3.4.21.

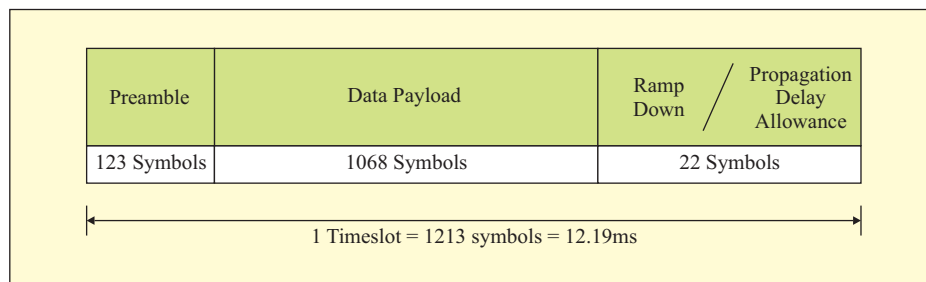


Figure 3.4.21 : On-air packet overall structure

The packet consists of three sections: The *preamble* section is used by the receiver to set its gain appropriately and to synchronise to the incoming data stream, the *data payload* contains either subscriber unit data (uplink) or base-station data (downlink), and the *ramp-down / propagation delay allowance* section enables the transmitting unit to ramp-down its PA gracefully, thus minimising spectral re-growth, and provides for the propagation delay between the transmitting and receiving units. The various sections of the packet are described below.

On-Air Packet Preamble

The preamble of the uplink packet consists of three main parts: a frequency tuning period, symbol reversals, and a Barker sequence transmission. This structure is illustrated in Figure 3.4.22.

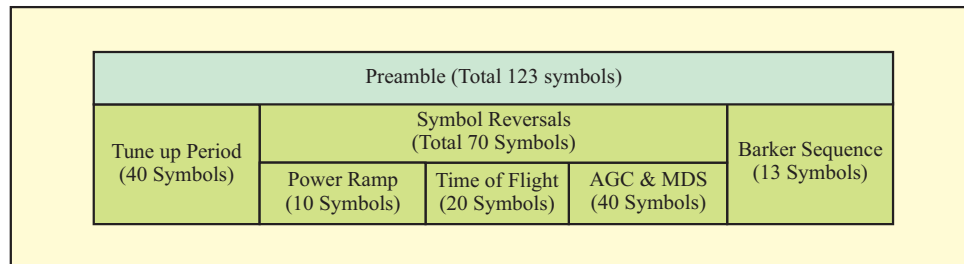


Figure 3.4.22 : On-air packet preamble structure

The tune-up period is included to give the RF local oscillators at each end of the link sufficient time to settle before transmission/reception commences. The delay is due to the loop filters in the synthesiser.

Once the local oscillators have settled, the transmitter begins to ramp up the gain of its power amplifier in a controlled manner over a 10-symbol period without modulation. This is to reduce switch-on spectral re-growth or ‘splatter’. This requirement is discussed in detail in § 3.5.1.2.1.

The received signal strength will vary, depending on the distance between the subscriber unit and the base-station unit. It will also vary in time due to propagation phenomena, and the combined variation may be as much as 70 dB. The dynamic range of the analogue to digital converter in the receiver, however, is only 40 dB. Furthermore, the full range of the converter must be utilised to minimise demodulator errors. The gain of the receiver must therefore be automatically controlled. This is achieved by taking a RMS measurement, averaged over a 10-symbol period, of the baseband voltages at the output of the receiver while symbol reversals are being transmitted and the transmission power is constant. The result of this calculation is used to access a calibrated look-up table and extract a value for receiver gain. The time of flight allowance included in the preamble is to ensure that the symbol reversals are at maximum power before signal strength measurement commences in the receiver. Time-of-flight allowance is discussed further below.

The end of the preamble is marked with a 13-bit Barker sequence [6]. The Barker sequence has very good aperiodic correlation properties, and these are exploited by the receiver to synchronise to the start of the data payload. If the sequence is not detected within the detection window, packet reception is aborted for the current timeslot and the unit waits for the next slot.

Experimental Astrophysics (Ppd) 1 - 48

The Pierre Auger Project TDR

On-Air Packet Data Payload

The data payload of the on-air uplink packet is segmented into 9 categories: packet type, packet number, station ID, communications data length, communications data, local station data length, local station data, CRC, and forward error correction (FEC – not implemented). This is shown in Figure 3.4.23 (a).

Packet Type (4 symbols)	Packet Number (4 symbols)	Station ID (8 Symbols)	Comms Data Length (4 symbols)	Comms Data (64 symbols)	SDS Data Length (4 symbols)	SDS Data (600 symbols)	CRC (16 symbols)	FEC Overhead (8 symbols)	2/3 Rate FEC (356 symbols)
----------------------------	---------------------------------	---------------------------	-------------------------------------	----------------------------	--------------------------------------	------------------------------	---------------------	--------------------------------	----------------------------------

Figure 3.4.23 (a) : On-air uplink data payload structure

The data payload of the on-air downlink packet is segmented into 10 categories: packet type, packet number, station ID, communications data length, communications data, ARQ, CDAS data length, CDAS data, CRC, and forward error correction (FEC – not implemented). This is shown in Figure 3.4.23 (a).

Packet Type (4 symbols)	Packet Number (4 symbols)	Station ID (8 Symbols)	Comms Data Length (4 symbols)	Comms Data (64 symbols)	ARQ Requests (72 symbols)	CDAS Data Length (4 symbols)	CDAS Data (528 - symbols)	CRC (16 - symbols)	FEC Overhead (8 symbols)	2/3 Rate FEC (356 symbols)
----------------------------	---------------------------------	---------------------------	--	----------------------------	---------------------------------	---------------------------------------	------------------------------------	--------------------------	--------------------------------	----------------------------------

Figure 3.4.23 (b) : On-air downlink data payload structure

In both cases the packet type field is used to identify the type of packet in the receiver, so that the appropriate interpreter can be scheduled. The packet number is incremented every frame and is used for network monitoring. The station ID field reflects the electronic serial number of the subscriber unit/base-station unit and is also used for network monitoring. The data length fields are used to specify the number of bytes that the various data fields contain. A packet is transmitted irrespective of the quantity of data available for transmission (the data fields are zero padded) and so these fields allow the subscriber units/base-station units to filter out redundant data.

On-Air Packet CRC

A 32-bit CRC field is used to check packet integrity in an identical manner to the one used in the local station ↔ subscriber unit communications protocol (§ 3.4.6.1 above). Packets that fail this check are automatically re-requested. This procedure is described in § 3.4.6.2.5.

On-Air Packet Ramp Down

In an identical manner to the way in which the power amplifier's gain needs to be ramped up in the preamble, it also needs to be smoothly ramped down at the end of the packet. This is performed on tail reversals.

On-Air Packet Time-of-Flight Allowance

Due to the topology of the Observatory, the maximum feasible path length is approximately 52 km. This corresponds to a free-space propagation delay of approximately 173 μ s. The receiver must therefore remain on for at least 173 μ s longer than the transmitter to cater for the worst case path loss. However, the receiver cannot run on beyond the frame boundary, which is synchronised to the 1 pps signal from GPS. Hence the time-of-flight allowance is incorporated into the transmitted packet such that the transmission of the data finishes 'early'. This guarantees that the receiver located at the maximum range will capture the full packet. At a symbol rate of 100 ksymbols/sec, a 173 μ s time-of-flight allowance equates to 18-symbols. An additional 4 symbols are added for PA ramp-down.

3.4.6.2.5 On-Air Communications Error Control

In order to provide *near error-free* communications, a strategy is required for correcting packets that have been identified by the CRC routine as corrupt, and this is based on packet re-transmission. The scheme is complicated slightly by the fact that the communications system is asymmetrically loaded, in that the majority of packet transactions are uplink packets from the detectors.

Two error control systems have been adopted:

1. A packet error on the uplink from a subscriber unit causes the base-station unit to request that the appropriate subscriber re-transmits the corrupted packet in a dedicated timeslot. This type of error control is termed *automatic request for re-transmission* or more commonly *automatic repeat request* (ARQ).
2. A redundant re-transmission policy has been adopted for downlink communications from the base-station to subscriber units, such that all downlink packets are transmitted 3 times per frame to maximise the chance of successful packet recovery at the subscriber. An ARQ system would be unsuitable for error control on downlink communications, since broadcast packets would have to be re-transmitted if just one subscriber unit within a base-station sector failed to decode the packet correctly. As each base-station can support up to 68 subscribers, packet failure on the downlink is 68 times more likely than on the uplink, assuming equal received signal strengths. Received signal strengths on the uplink and downlink, however, are not equal; the subscriber unit antenna system has a gain of approximately 12 dBi whereas the base-station antenna system has a gain of 17 dBi, but both units transmit at an effective isotropic radiated power (EIRP) of +36 dBm, and hence the signal received at the antenna connection of the subscriber unit will be 5 dB lower than the signal received at the base-station (channel reciprocity assumed). This reduced signal strength will further increase the packet failure rate on the downlink.

Experimental Astrophysics (Ppd) 1 - 50

The Pierre Auger Project TDR

The subscriber unit ignores the repeated transmissions if it has previously recovered the packet. In the unlikely event that the subscriber unit fails to decode all three identical packets, it is deemed lost. Packet losses are further minimised by the fact that CDAS uses any spare capacity it does not use in the current downlink packet to repeat the data destined for the local stations from previous frame periods.

When a corrupt uplink packet is detected, the base-station unit uses dedicated bandwidth on the downlink to request its re-transmission. The ARQ section of the TDMA frame is illustrated in Figure 3.4.24. An ARQ request is repeated three times to ensure increase the likelihood of the subscriber receiving it. The subscriber then re-transmits the requested packet in one of 3 dedicated reply timeslots. The entire process is repeated a second time, providing 6 ARQ reply timeslots per frame.




Figure 3.4.24 : ARQ slots structure

The base-station unit compiles a list of packets that have failed and generates the ARQ data for transmission in the ARQ downlink packet, the structure of which is given in Figure 3.4.25. When the subscriber unit receives this packet it searches for its station ID. If an ID match is found, the subscriber determines which packet should be retransmitted using the *request number* field, and then retransmits it on one of the six dedicated reply slots. The position of the subscriber's station ID within the downlink packet ID list is used to determine which reply slot it should use.

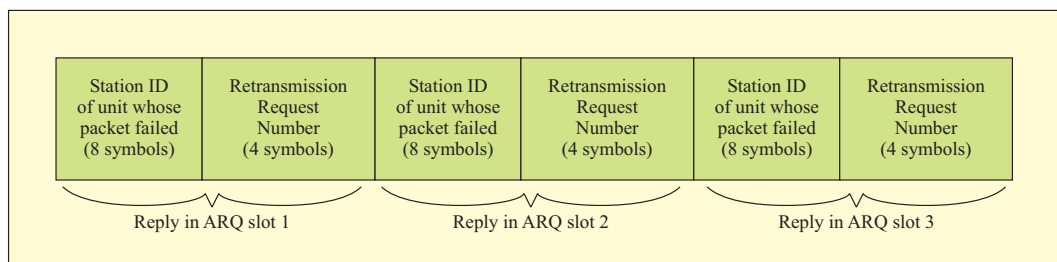


Figure 3.4.25 : ARQ downlink packet data structure

Due to memory restrictions within the subscriber unit, packets can only be held for re-transmission purposes for a finite period of time. In the unlikely event that a packet fails

to be delivered after 8 frames have elapsed since its generation, the base-station stops requesting retransmissions and the packet is deemed lost.

3.4.6.2.6 Quality of Service Implications of the ARQ System

The quality of service target for the wireless LAN is an uncorrected raw bit error rate per link of 1×10^{-6} due to thermal noise.

If we assume a subscriber unit packet length of 1400 bits and a base-station sector containing the maximum number of 68 subscriber units, we may derive a *sector uplink packet error rate* as follows :

$$\begin{aligned}\text{Individual packet error rate} &= \text{packet length} \times \text{bit error rate} \\ &= 1400 \times 10^{-6} \\ &= 0.14 \% \text{ packet error rate, PER} \\ \text{Sector packet error rate} &= \text{individual PER} \times \text{number of subscribers per sector} \\ &= 0.14 \% \times 68 \\ &= 0.095 \text{ packet errors/sector/second.}\end{aligned}$$

As 6 ARQ slots/second are available to handle error packets, the packet error rate due to thermal noise induced errors may be easily accommodated. However, RF interference and fading effects also contribute significantly to on-air errors so that an allowance of 6 ARQ slots/second offers a sensible balance between communications system efficiency and robustness, guaranteeing that the complete and irrevocable loss of data within the network will be a very infrequent event.

3.4.6.3 Base-station Unit to μ LSx Communications

3.4.6.3.1 The E1 Bearer

The surface array data recovered by the base-stations at each tower site is transferred to a μ LSx unit using an E1 interface. E1 is an industry standard 2.048Mbps telecommunications interface originally developed for digital voice communications, but which is now widely adopted for high-speed digital data communications. The E1 bearer uses a time division multiple access (TDMA) frame structure, as illustrated in Figure 3.4.26.




Figure 3.4.26 : E1 frame structure

The system uses an 8 kHz frame rate with each having a duration of 125 μs. Frames are divided into 32 equal timeslots of 3.9 μs and each timeslot, or channel, carries 8-bits of serial data and provides a 64 kbps data transfer rate. Channels 0 and 15 are reserved and carry synchronisation and signalling information. Standard E1 uses 75 Ω or 120 Ω balanced lines that may be up to 2 km in length.

The E1 system provides a drop-and-insert capability for full duplex operation. Data from a channel can be read from the E1 ring before new transmission data overwrites it, allowing channels to be allocated for both transmission and reception. If there is no data to transmit, the unit retransmits the received word. When data is available for transmission, the data received by the transmitting unit is dropped, preventing it being delivered to units downstream. The μLSx unit then routes the data back to the CDAS via the microwave backbone, which uses the TCP/IP protocol.

The E1 interface is configured as a ring as shown in Figure 3.4.27. The μLSx is configured as the master timing reference in the ring, from which all slave base-stations derive their timing signals. This ensures that synchronisation of the E1 network is maintained.

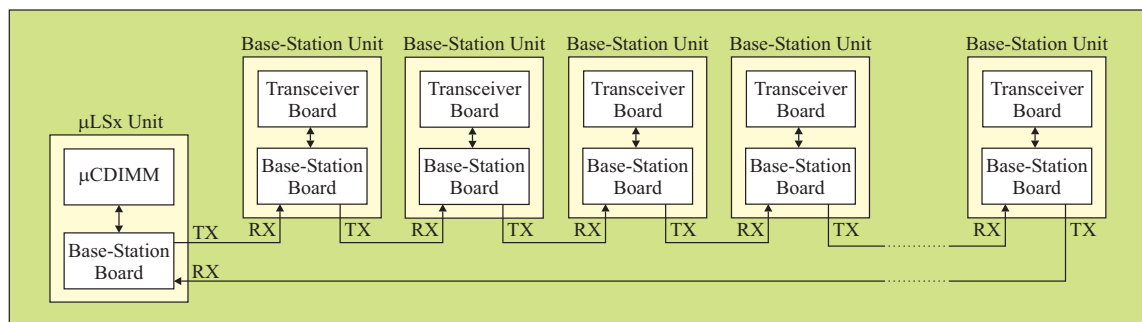


Figure 3.4.26 : Base-station / μLSx E1 ring configuration

3.4.6.3.2 E1 Network Acquisition

To ease commissioning and maintenance of the base-station equipment, all base-stations are programmed with a common software build. This means that it is not possible for the base-station to have prior knowledge of the μ LSx to which it is connected, nor the E1 channels that are available for communicating data. This connection is therefore negotiated using an E1 ring network acquisition protocol.

Using a dedicated *Network Control Stream*, the base-stations negotiate the E1 channels to be used for transmission and reception of stream data with the μ LSx to which they are connected. The base-stations and μ LSx units associate timeslots 1 and 2 with stream 0, and utilise this stream to negotiate a connection. The μ LSx periodically transmits an *E1 network control message* to the base-stations. The format of the data payload section of this packet is shown in Figure 3.4.27.

BS ID	Stream Number	Number of E1 channels to allocate	E1 channel to assign	Slot type	E1 channel to assign	Slot type		E1 channel to assign	Slot type
(1 Byte)	(1 Byte)	(1 Byte)	(1 Byte)	(1 Byte)	(1 Byte)	(1 Byte)		(1 Byte)	(1 Byte)

Figure 3.4.27 : E1 network control packet data payload

During production, each base-station is assigned a unique 8-bit serial number, and this number is used by the μ LSx to identify each unit. The μ LSx generates and transmits an E1 network control message every 25ms. The *BS ID* field is incremented each time a new network control message is generated. This ensures that all base-station within the E1 ring will be addressed. Each base-station processes the network control message, and if it detects its station ID, it uses the contents of the packet to determine how the E1 channels and streams environment should be set.

The *number of E1 channels to allocate* field in Figure 3.4.27 is used by the base-station to determine how many E1 channels to associate with a given stream, and also to determine the length of the channel information that follows. Each subsequent byte-pair within the message identifies the E1 channel to associate with the stream, identified by the *stream number* field, and the E1 channel type, identified by *slot type* field. A '0' in this field denotes receive channels, a '1' denotes transmit channels and a '2' denotes transmit and receive channels (i.e. the drop-and-insert capability of E1 is utilised where data is read from the channel prior to transmission data overwriting it). Currently, all channels are allocated as type '2' but, to increase system flexibility, provision has been made to allow individual channels to be configured as any of the types described.

Experimental Astrophysics (Ppd) 1 - 54

The Pierre Auger Project TDR

Once channel allocation is complete, the base-station unit informs the μ LSx that it has acquired the E1 network by transmitting an *E1 network acquisition acknowledgement*, which contains its station ID.

Once a base-station has acquired the E1 network, communications with the μ LSx may commence, and concentrated uplink data from the surface array is then forwarded to the μ LSx using the allocated stream. The μ LSx also routes downlink data from the CDAS and the Network Monitoring PC at the Campus to individual base-stations, using the appropriate stream.

In the event that the μ LSx is ordered to reset, a method of forcing all base-stations to disconnect from the E1 ring network is required to prevent connected units from attempting to send data down a stream that is not being serviced. A “disconnect all streams” message is issued by the μ LSx unit as it boots to ensure that no connections exist. If the base-stations detect this message, the E1 channel allocations are reset and the acquisition process re-starts. In the event that a loss of synchronisation is detected by the E1 channel processor, the base-station forces an E1 re-synchronisation, disconnects any connected stream and re-starts the network acquisition process.

The configuration of the network ensures that the distance between the base-stations and the μ LSx in each E1 ring network is only a few cm and that the entire ring network is unlikely to exceed 5 m in length. It is assumed therefore that the E1 ring network represents a *near error-free* environment and no error correction capabilities specific to the E1 ring are employed. However, the message protocol that operates over the E1 ring uses CRC to check packet integrity. In the event that a packet fails the CRC, causing an acknowledgement failure, the protocol dictates that the transmitting unit should automatically re-issue the unacknowledged packet.

Bit-error-rate (BER) tests have been conducted on an E1 ring network, 2 m in length, consisting of 2 base-stations and a Marconi Instruments 2841 digital communications analyser. 1.8×10^{11} bits of a pseudo-random-binary-sequence (PRBS) were transmitted by the analyser and, after passing through the 2 base-stations, were detected correctly without error, thereby validating the “near error-free” assumption.

3.4.6.4 μ LSx to CDAS Communications

The micro-LSx unit provides an interface between each base-station and the microwave backbone system that ultimately connects the surface array to the central data acquisition system, (CDAS) located at the Observatory Campus.

At the physical layer, each micro-LSx interface presents a standard 10Mbps Ethernet connection to a local router at each communications tower. In turn, the router connects to the microwave backbone system and hence to the Campus.

At the campus tower, TCP/IP traffic from the surface array base-station radios and the FD eye buildings is routed onto the main campus network and into the CDAS system.

As each base-station registers on its local E1 ring, the micro-LSx interface opens socket and starts a TCP/IP session for that base-station with the campus CDAS computers and forwards packets to and from the base-station and CDAS. In this manner, the CDAS system sees each base-station as a conventional TCP/IP client.

3.4.6.5 Network Maintenance and Monitoring

3.4.6.5.1 Network Status Monitoring

It is essential that continuous monitoring of the communications system is provided to ensure reliability. Consequently, a number of monitoring mechanisms have been designed into the system. The ability to pre-empt a possible failure provides service teams with valuable time to reach remote locations to investigate the problem. A dedicated communications network PC that uses the existing communications backbone infrastructure logs monitoring data and warns of possible outages. The architecture of this system is illustrated in Figure 3.4.28.

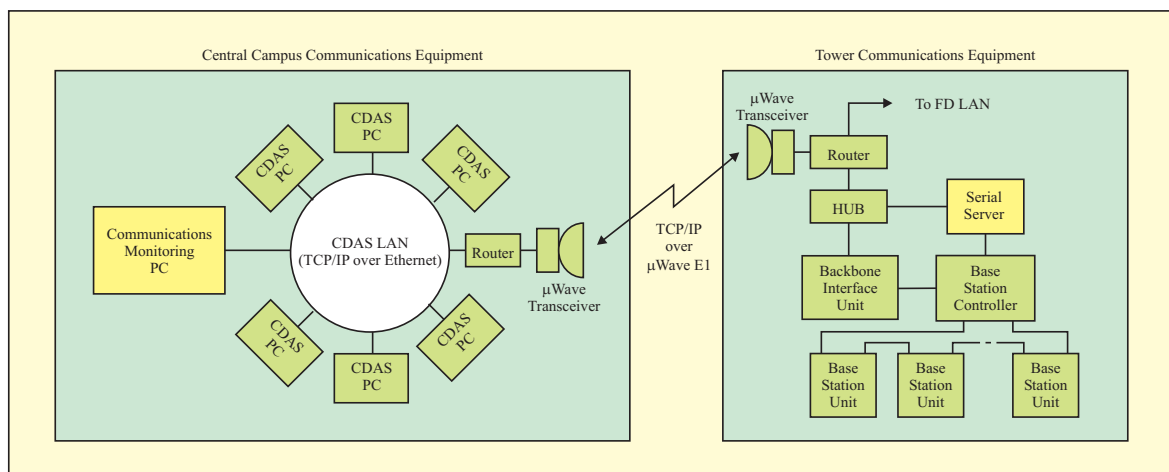


Figure 3.4.28 : PC monitoring system architecture

A PC connected into the CDAS computer network uses the TCP/IP protocol to communicate over the microwave backbone with a serial server. The serial server at each communications tower provides a RS232 port for each base-station controller unit used throughout the array. A custom protocol, identical to the subscriber unit ↔ local station protocol, is used to transfer network monitoring data from the base-station units, via the base-station controller and serial server, to the PC for logging. This architecture enables the monitoring PC to communicate with a base-station controller, a set of base-stations, an individual base-station unit, a sector of subscriber units or an individual subscriber unit for maintenance whilst remaining completely transparent to the main data traffic.

Figure 3.4.29 shows a typical screen display from the network monitoring suite. The display shows the connection status and serial numbers of 64 subscriber radios connected to one of the base-stations registered on the system (in this case number 15). Other pages

Experimental Astrophysics (Ppd) 1 - 56

The Pierre Auger Project TDR

on the suite allow the in-depth monitoring of an individual subscriber unit, the remote programming of radios in the surface array and monitoring of the packet-error rate currently being experienced by the system.

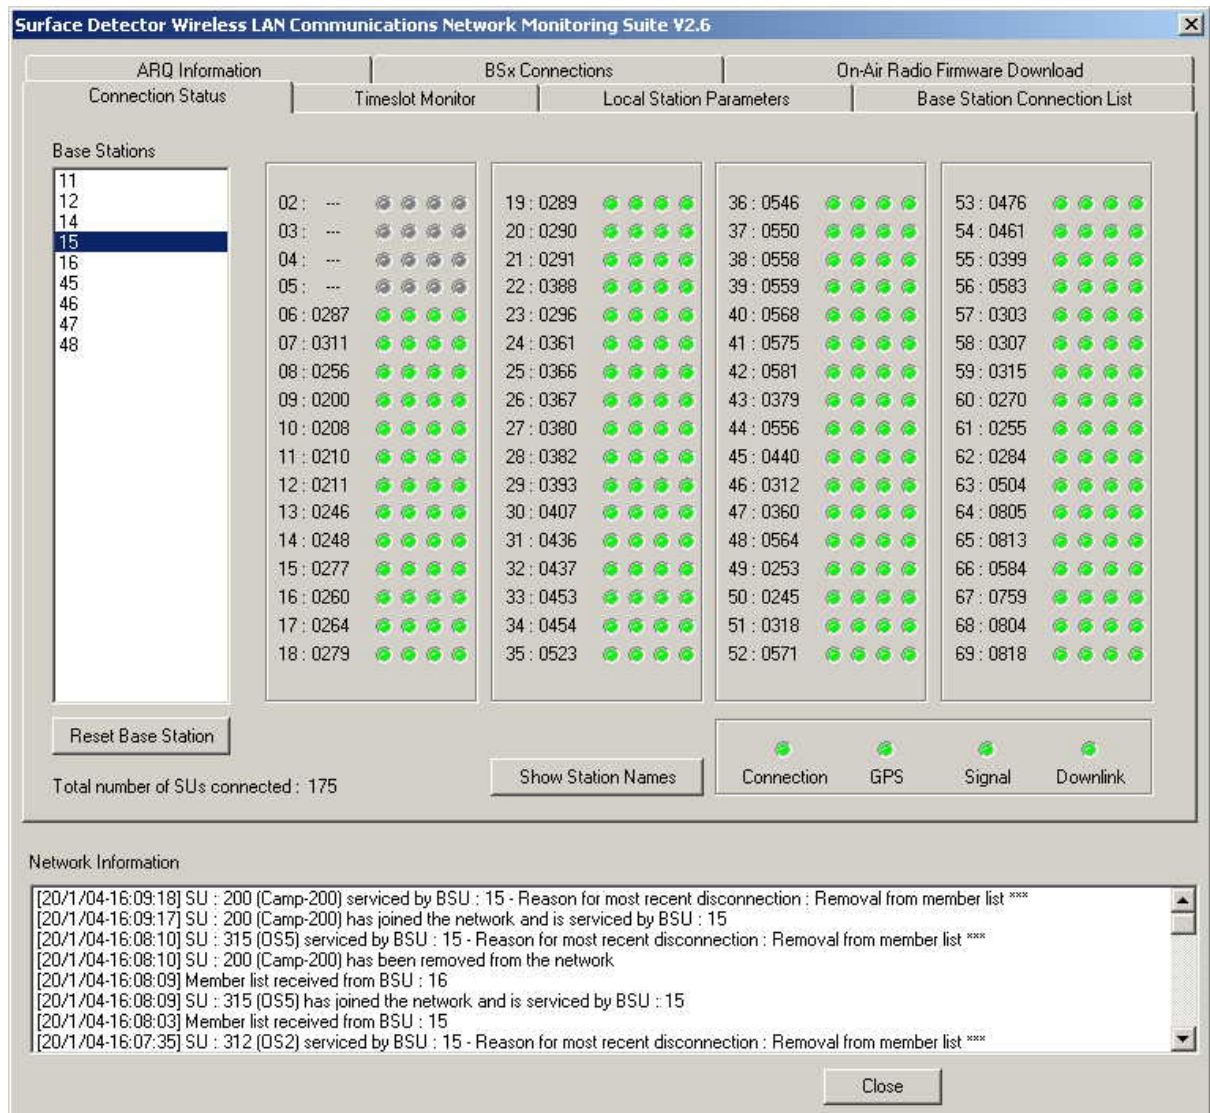


Figure 3.4.29 : Communications Network Monitoring Suite

3.4.6.5.2 Software Upgrades

During the initial stages of deployment and commissioning, the ability to upgrade the entire communications system software is of utmost importance since the logistics involved in visiting each unit becomes ever more significant as the experiment grows. A small amount of additional bandwidth has been provided to perform this task. Under control from the monitoring PC a base-station controller, base-station unit, a specific subscriber unit or sector of subscriber units may be re-programmed remotely whilst maintaining normal network operation. It may take in excess of one hour to upgrade the software and so transparent download is essential to reduce network downtime. Once download is complete the PC may command any unit to reboot using the software

upgrade. Software that is downloaded to any unit does not replace the current operating code. This prevents an error during the download procedure from incapacitating the unit, potentially breaking its lifeline with the monitoring PC.

This facility will also prove to be extremely useful should the radio system undergo any future mid-life upgrades.

3.4.6.5.3 Overall Communications Data Throughput

Table 3.4.1 summarises the maximum data throughput each link of the communications system is capable of sustaining.

Link Description	Maximum Physical Interface Data Throughput	Maximum Protocol Data Throughput	Maximum Required Data Throughput	Condition
Local Station to Subscriber Unit	38.4 kbps	36.1 bps	1.2 kbps	Unconditional
Subscriber Unit to Base-station Unit	96 kbps	1.2 kbps	1.2 kbps	A single uplink slot allocated to the subscriber. Base-station units support up to 68 subscribers.
Base-station Unit to Base-station Controller	2.048 Mbps	180.4 kbps	96 kbps	Three E1 timeslots allocated to the Base-station Unit. A base-station controller supports a maximum of 8 base-station units.
Base-station Controller to Backbone interface unit	640 Mbps	75 Mbps	768 kbps	Unconditional
Backbone interface unit to CDAS	8.2 Mbps	7 Mbps approx.	768 kbps	7 Mbps for the protocol data throughput is approximated as the TCP/IP protocol used on the backbone is contention based and therefore heavily dependent on the number of users.

Table 3.4.1 : Summary of the communication network data throughput

REFERENCES

- [1] Federal Communications Commission (FCC) Rules & Regulations, Part 15 : Radio Devices, Section 15.247 : Operation within the bands 902-928MHz, 2400-2483.5MHz, and 5725-5850MHz, October 2000
- [2] Hata, M.: “Empirical formula for the propagation loss in land mobile radio services”, IEEE Transactions, 1980, VT-29, pp.317-325
- [3] Hall, M. P. M., Barclay, L. W. and Hewitt, M. T.: “Propagation of Radiowaves”, IEE, London 1996, pp. 386
- [4] IEEE Standards for Local Area Networks : “Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specification”, International Standard 8802/3, Std 802.3-1985. The Institute of Electrical and Electronics Engineers, Inc. Distributed by Wiley-Interscience. pp27. ISBN-0-471-82749-5.
- [5] Clark, P.D.J.C, Tunnicliffe, V., Berge, O.M. : “Pierre Auger Project : RS232 communication Protocol”. v2.1, March 2001.
- [6] Fan, P., Darnell, M.: “Sequence Design for Communications Applications”, Research Studies Press LTD, Taunton, Somerset, England. 1996. pp269-272. ISBN-0-86380-201

3.5 WIRELESS LAN RADIO HARDWARE

This section describes the custom radio hardware that has been designed within the collaboration to meet the special requirements of the surface detector communications system.

3.5.1 Wireless LAN Subscriber Unit

3.5.1.1 Operational and Environmental Considerations

3.5.1.1.1 Surface Detector Data Link Requirements

The Leeds surface detector wireless LAN system is required to support two reliable data ‘pipes’ simultaneously from each of the 1600 surface detectors : a 1200 bits/s uplink from each detector to the CDAS, and a 2400 bits/s broadcast downlink from the CDAS to the detectors [1]. The wireless LAN divides into, nominally, 28 sectors, each sector containing up to 68 detectors, with each detector being connected to a digital radio. The system is similar to a cellular telephone network in that each sector is served by a single base-station unit located at a communications tower. This unit concentrates the data to and from all of the sector’s surface detector subscriber units.

The data pipes between a base-station and multiple subscriber units are supported on a time division multiple access basis, whereby each subscriber unit is guaranteed a timeslot for the transmission of data packets (see §3.4.6.2). Hence the on-air data rate from each subscriber unit is required to be considerably higher than the required average rate of each pipe. There must also be sufficient capacity for transparent error detection and correction. The wireless LAN radio network has been designed to support an on-air bit rate of 200 kbits/s; this capacity allows a packetised communications protocol to be implemented, such that data packets received with bit errors can be re-requested and re-transmitted several times within a time window governed by the surface detector trigger electronics.

3.5.1.1.2 Radio Transceiver Certification

The wireless LAN radio communications hardware for the Southern Observatory is required to satisfy the Argentinean regulations regarding radio emissions, specifically those of the Comision Nacional de Comunicaciones (CNC) [2], and has been granted certification by this body. The Leeds radio transceiver has also been designed to satisfy the closely related US Federal Communications Commision (FCC) regulations [3].

3.5.1.1.3 Subscriber Unit Operating Environment Issues

The Auger Observatory site extends over an area of more than 3000km², and substantial parts of the site can be extremely inaccessible during the wet season. During the dry season, a round trip from the observatory campus to the centre of the site can take several hours at best, and requires considerable planning. Moreover, on-site temperatures can vary from -15°C to $+35^{\circ}\text{C}$. For these reasons the subscriber unit hardware is of a zero-maintenance design using high-reliability, industrial specification components where possible. Other high reliability design techniques have been employed in the radio such as the elimination of electrolytic capacitors to the minimum number possible, and the use of special long-life parts in those places where their use could not be avoided. To further enhance reliability, all subscriber units are fully ‘burned in’ in an environmental test chamber before shipping. This procedure individually calibrates each radio and tests all internal functions under extreme operational temperature conditions.

3.5.1.1.4 Physical Description

The subscriber unit electronics are contained in a robust extruded aluminium enclosure, as shown in Figure 3.

Four-layer standard epoxy-glass FR4 printed circuit boards are used, and the majority of components are surface mount devices. The circuit board has been manufactured and populated using sub-contractors approved to ISO 9002.

The subscriber unit has two external connectors: a SMA RF connector for the antenna feeder, and a standard 9-way D-type connector for 12V DC power, serial communications with the local station, and a 1 pulse-per-second signal from the GPS engine. The unit measures 205×105×60mm and weighs 850g.



Figure 3.5.1 : Wireless LAN subscriber unit

3.5.1.1.5 Subscriber Unit Power Budget

As each surface detector is powered from a solar panel system of limited capacity, the subscriber unit has been designed to ensure that its power consumption does not exceed the 1.2 Watts dictated by the surface detector power budget. The average power consumption of the subscriber unit is 1.02 Watts (see §3.5.1.2.4).

3.5.1.2 Subscriber Unit Architecture

The functionality of the subscriber unit can be divided into the following four areas, each of which has embedded within it built-in self test (BIST) features :

1. Radio frequency (RF) transceiver;
2. Analogue baseband filtering and conversion;
3. Digital signal processing, serial communications and transceiver control;
4. Powers supplies (5 Volts and 3.3 Volts).

Experimental Astrophysics (Ppd) 1 - 62

The Pierre Auger Project TDR

These areas are highlighted in Figure 3.5.2 and described in the following sections. A functional block diagram of the subscriber unit is shown in Figure 3.5.3.

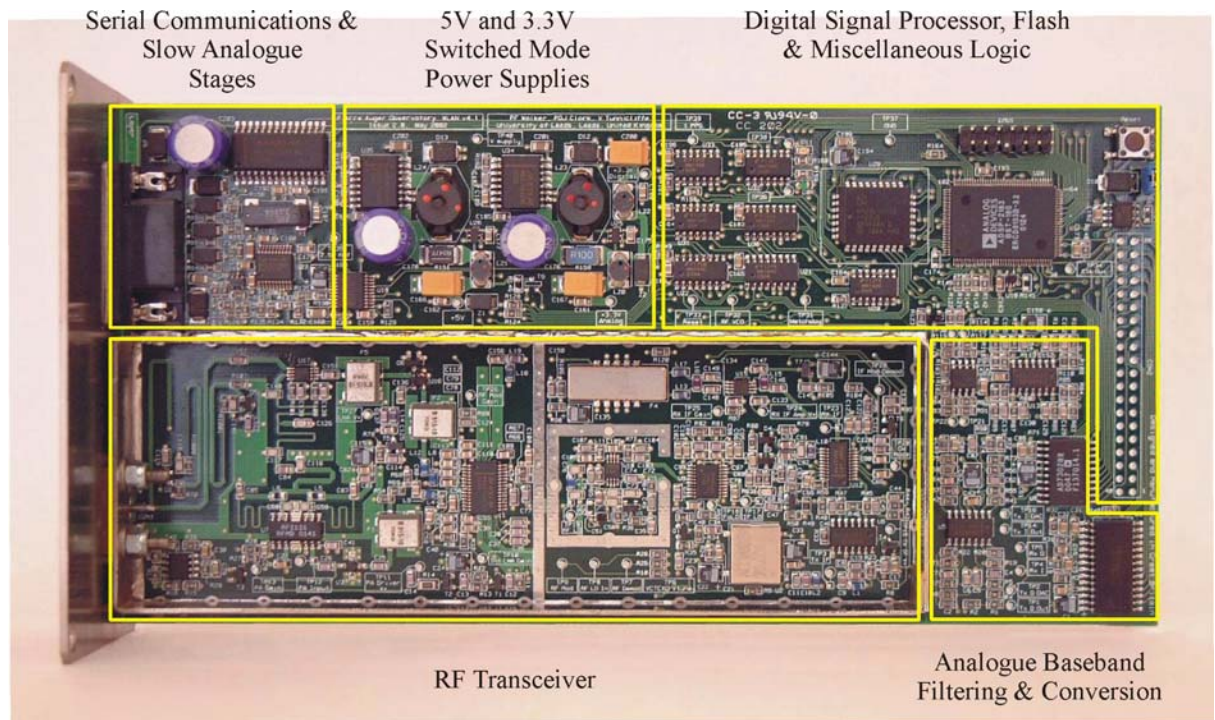


Figure 3.5.2 : Subscriber Unit PCB showing the main functional blocks

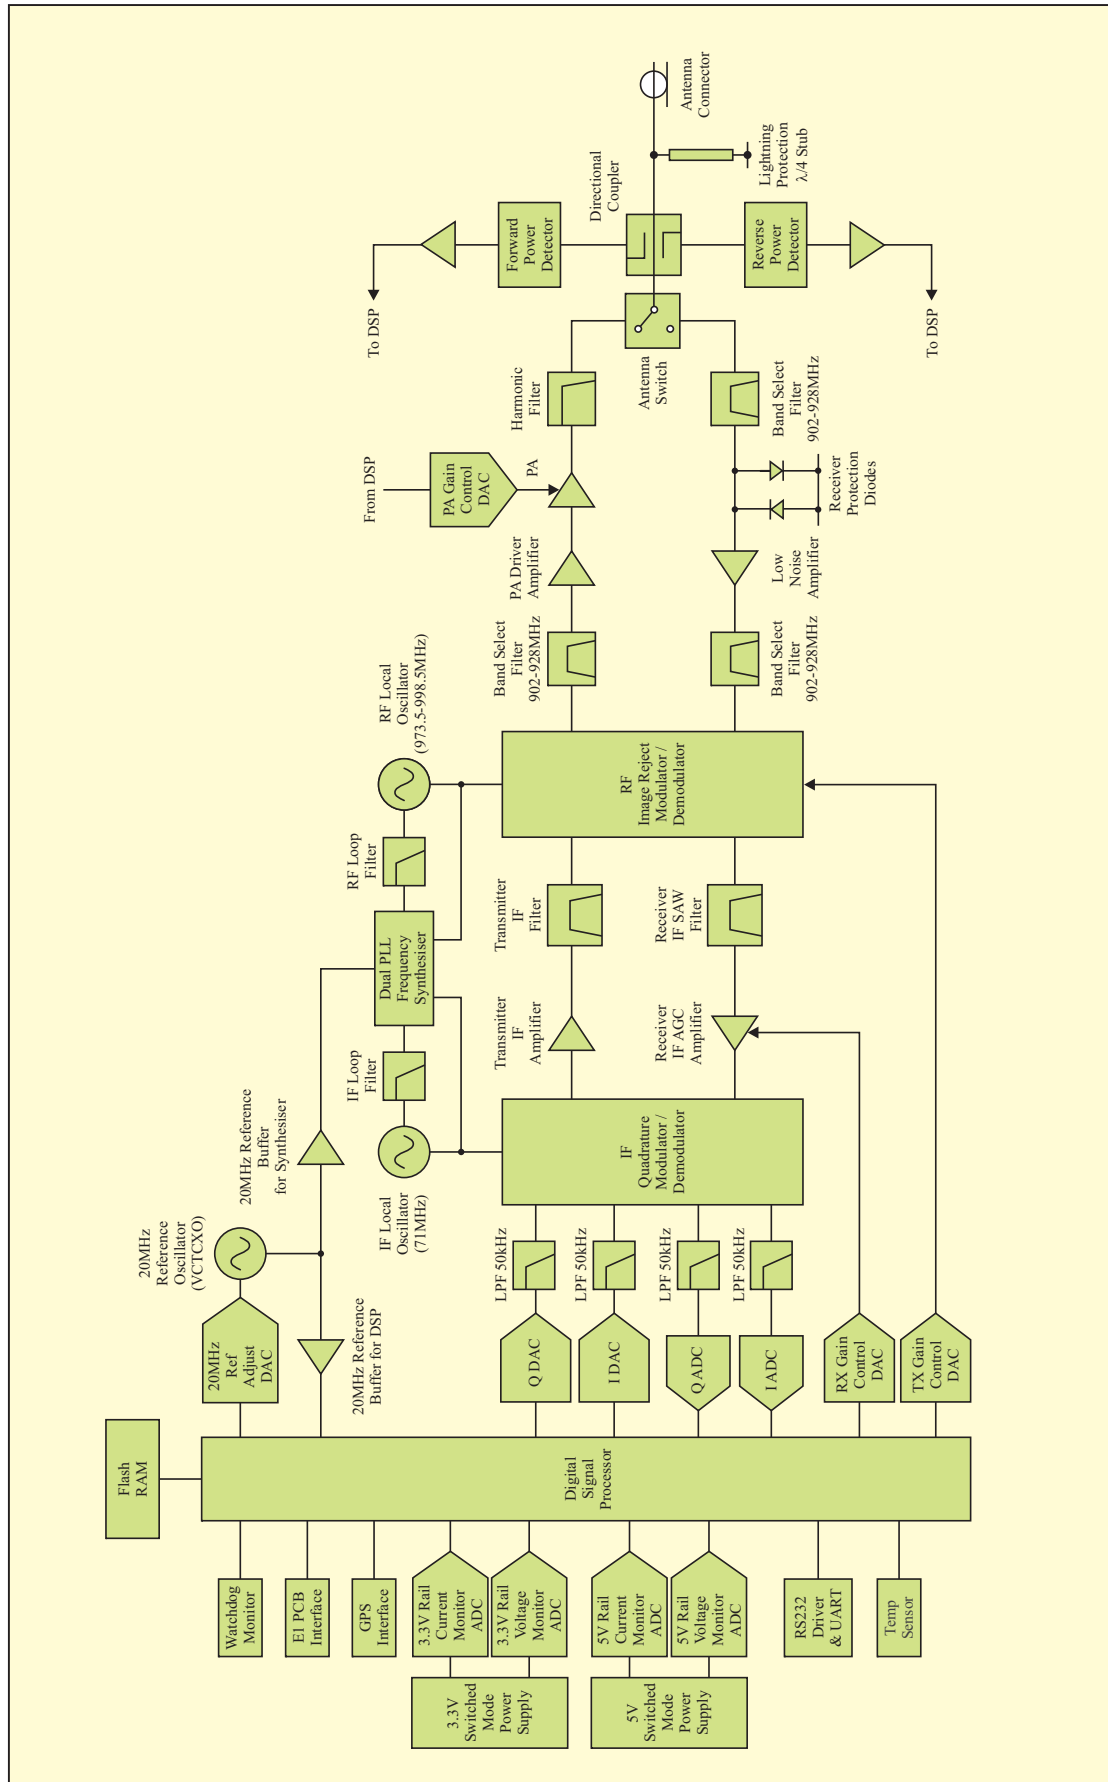


Figure 3.5.3 : Wireless LAN subscriber unit block diagram

3.5.1.2.1 RF Transceiver

The transceiver has been designed as a ‘half-duplex’ system, i.e. it does not transmit and receive simultaneously. This approach allows certain functional blocks such as the synthesiser to be shared between the transmitter (TX) and receiver (RX) sub-sections. It also eliminates the need for extensive screening between the two sub-sections. Routing of the antenna to the TX or RX section is done via a solid-state RF switch.

The transceiver operates in the 915 MHz global Industrial, Scientific and Medical (ISM) band, which extends from 902 MHz to 928 MHz. Channel spacing has been selected at 500 kHz, resulting in 51 channels being available (902.5 MHz through 927.5 MHz inclusive). In accordance with FCC regulations, the transceiver has been designed with a frequency hopping capability, whereby all 51 channels shall be used equally by each surface detector in a pseudo-random manner, and to a schedule such that the dwell-time on any one channel shall be less than 400 ms [3]. Whilst this requirement complicates the system scheduling and places specific demands on the RF hardware, frequency hopping systems have two main advantages over fixed-frequency systems. Firstly, they offer greater security against interception and jamming (for these reasons alone they have traditionally been employed in military applications), and secondly – and of greater significance for multi-user bands such as the 915 MHz ISM band – they are less susceptible to, and generate lower levels of, narrow-band interference.

The wireless LAN system employs a frequency hopping scheme whereby each surface detector changes its active frequency channel every timeslot. This operation is part of the system control software, as described in §3.4.6.2; when a subscriber unit powers-up and registers with the network it is assigned a unique index into a frequency hopping sequence, and this index is updated each timeslot. The need to change the frequency channel so rapidly places demands on the transceiver synthesiser, as it must be capable of re-tuning and settling between timeslots which, in the worst case, will be adjacent slots. The mechanism to enable such rapid hopping, however, can compromise synthesiser phase noise, which in turn will degrade receiver performance. It is therefore essential that this issue is addressed carefully to satisfy both criteria.

The architecture of the transceiver is based on a quadrature or ‘I/Q’ structure; this is the most powerful and versatile architecture for systems employing digital modulation with digital signal processing (DSP) capability. In the simplest form of quadrature modulation, a baseband data bit stream is split into odd bits and even bits and these two streams are presented to closely-matched mixers whose local oscillators differ by 90°. The outputs of these mixers are then summed. The result is a carrier whose phase relative to the local oscillators can take any of four states : +45°, +135°, +225° or +315°, hence this scheme is known as Quadrature Phase Shift Keying (QPSK). This is shown in Figure 3.5.4. In this scheme, the bit rate is equal to twice the baud rate as there are 2 bits of information contained in each QPSK phase ‘symbol’.

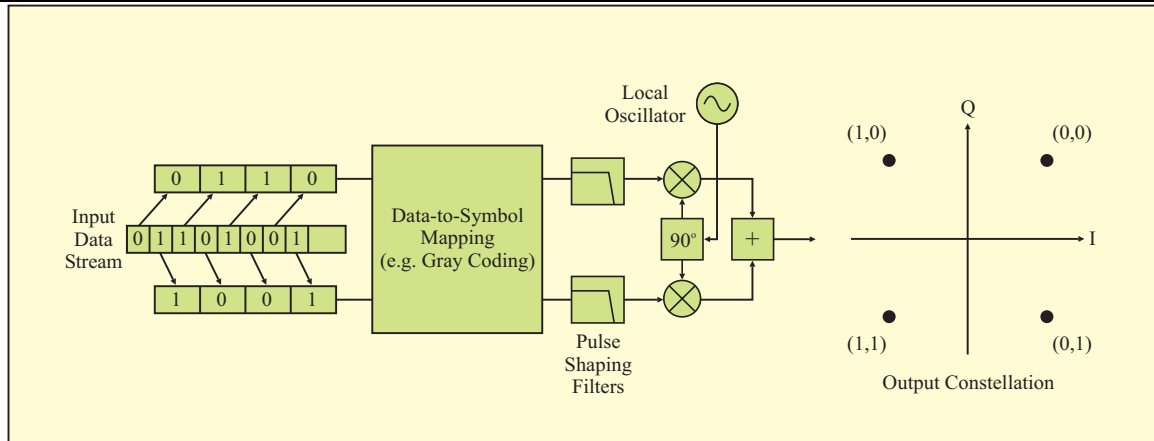


Figure 3.5.4 : QPSK modulation and output constellation

The data bit pairs extracted from the input stream are usually coded and filtered before the mixing stage, as shown in Figure 3.5.4. The coding may involve simple Gray-coding, whereby adjacent symbols of the constellation only differ by 1 bit [4]. This procedure minimises the ratio of bit errors to symbol errors as the most likely error patterns (a transition to a neighbouring position) cause only a single bit error for all possible permutations. Coding may also be differential, whereby a memory of 1 symbol is employed to allow the data to be encoded in terms of phase advances between symbols rather than absolute phases.

Filtering of the data is required to control the spectral occupancy of the modulation. Filtering, however, introduces inter-symbol interference, which has an impact on synchronisation at the receiver. This is illustrated by means of the ‘eye’ diagram of Figure 3.5.5, which represents a continuous overlay of random, filtered data bits or symbols (1 bit per I/Q symbol).




Figure 3.5.5 : Eye diagram showing inter-symbol interference (BT<0.5)

Experimental Astrophysics (Ppd) 1 - 66

The Pierre Auger Project TDR

The optimum synchronisation point for the receiver is located where the eye is widest, as this is the point of least inter-symbol interference; the tails of adjacent symbols are at such a level that their influence on the current symbol is minimal. This position corresponds to the centre of the data symbol. The extent to which the eye opens is related to the product BT , where T is the symbol period, and B is the filter bandwidth. It is also related to the shape of the filter. Assuming perfect 'brickwall' filtering, the eye begins to close when BT falls below 0.5, but this may also correspond to an improvement in spectral occupancy. In GSM, the global digital cellular standard, BT must be reduced to 0.3 to allow the spectral occupancy specification to be met. This is because the channel spacing is only 200 kHz yet each channel must support 270ksymbols/s. For the Auger wireless LAN system, which employs a channel spacing of 500 kHz and an on-air symbol rate of 100 ksymbols/s, the spectral occupancy specification is much more relaxed. Hence a BT of ≈ 0.6 can be used.

If baseband filtering permits, then quadrature modulation can be performed directly at RF. It is usually performed at a fixed intermediate frequency (IF), however, before a second stage translates this up to the required RF. Quadrature modulation enables almost any complex digital modulation scheme to be generated at the RF output. In a quadrature receiver, the RF is translated back down to baseband, once again using quadrature mixers and usually via an IF stage. Thus the original baseband data streams are recovered and recombined.

The wireless LAN transceiver is a single superhetrodyne design employing a fixed IF of 71 MHz for both modulation and demodulation functions.

RF Transceiver : Transmitter Operation

Overview of Transmitter

A block diagram of the transmitter is shown in Figure 3.5.6. The filtered baseband data streams 'A' and 'B' are applied to an I/Q modulator operating at 71 MHz, and the output of this is amplified and filtered to remove any unwanted mixing products. The IF local oscillator is supplied by a voltage controlled oscillator, which together with a synthesiser and a loop filter form a highly stable phase-locked loop. The synthesiser is programmed by the DSP running the subscriber unit software.

The filtered IF then passes to the second mixing stage to translate it to the required RF frequency. The RF local oscillator is also controlled by the synthesiser within a phase locked loop and varies from 973.5 MHz to 998.5 MHz. The mixing stage incorporates an image-reject mixer, which is essentially a single-sideband modulator, realised using a similar structure to the I/Q modulator. This provides approximately 30 dB of suppression to the unwanted term in the mixing process, in this case the frequency component at $(973.5+71)$ MHz to $(998.5+71)$ MHz. This allows the output filtering requirements to be relaxed.

The output of the RF modulator is insufficient to drive the power amplifier (PA) directly and it must pass through a driver stage of sufficient gain and linearity. The output of this stage is filtered before it is finally fed into the PA. The output of the PA undergoes

further filtering to remove any harmonics before it is routed to the output connector via the RF switch. From here, the signal passes through the RF feeder to the antenna where it is transmitted.

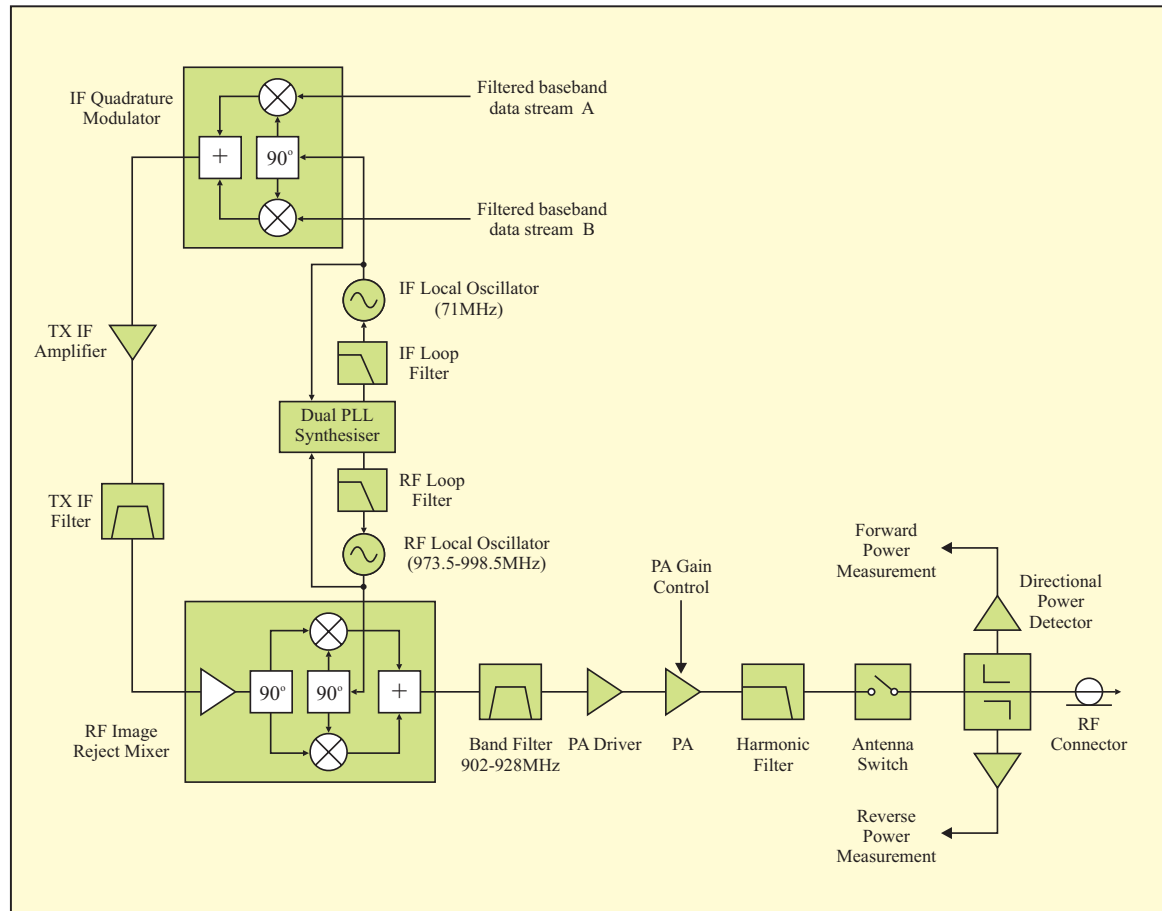


Figure 3.5.6 : Wireless LAN transmitter

Transmitter Output Power

The transmitter is capable of delivering just over 300 mW (+25 dBm) of continuous RF power into a 50 Ω load across the 902-928 MHz band. Whilst the FCC regulations permit a maximum output power from the transmitter of 1 Watt (+30 dBm), the 12 dBi gain of the Yagi antenna selected for use at each surface detector means that the maximum power must be reduced to 300 mW to meet the CNC and FCC limit of 4 Watts (+36 dBm) for maximum Effective Isotropic Radiated Power (EIRP). EIRP takes into account the maximum output power of the transmitter, the gain of the antenna and the losses in the RF feeder between the transmitter and the antenna.

Transmitter Linearity and Efficiency

Experimental Astrophysics (Ppd) 1 - 68

The Pierre Auger Project TDR

The PA output stage operates in a ‘near-linear’ Class AB mode, which means that the demands on the output filtering are less stringent. Whilst this is less efficient than non-linear modes, the transmitter still operates within the power budget. Moreover, it facilitates possible future upgrades to higher level modulation schemes to support higher data rates.

The gain of the PA must be carefully controlled to follow a well-defined ramping profile to minimise close-in spectral ‘re-growth’ when the transmitter powers up and down. Failure to do so will result in unacceptable levels of adjacent channel interference, as illustrated in Figure 3.5.7. This shows the improvement in the output spectrum of the wireless LAN transmitter operating on channel 27 when careful ramping of the PA gain is employed. At 500 kHz offset, which corresponds to the adjacent channels, the improvement over the ‘hard switched’ case is 12 dB.




Figure 3.5.7 : Effect of PA gain control on transmitter output spectrum

Transmitter RF Power Detector

Positioned between the RF switch and the RF connector is a directional power detector. Its purpose is to ascertain the integrity of the surface detector antenna system by measuring its return loss, a parameter which indicates what fraction of RF power delivered by the transmitter is reflected by the load due to an impedance mismatch. The return loss will decrease if the RF feeder degrades (due to water ingress for example) or if the antenna is damaged. This measurement can be made periodically and used to flag a degrading antenna system at the surface detector, potentially permitting maintenance to be scheduled before failure occurs.

The directional power detector operates by measuring the standing wave that arises when reflected (reverse) power becomes a significant fraction of incident (forward) power. The voltage standing wave ratio (VSWR) can be related directly to return loss at the load. The directional detector consists of a directional coupler and two diode-based power detectors. The coupler is implemented using coupled microstrip lines, whereby a small portion of the RF power propagating along the main transmission line is coupled into the detector circuits. The directivity of the coupler is approximately 20dB, which means that the error involved in measuring the return loss, and therefore the condition of the antenna system, is sufficiently low to allow a reliable good/poor decision to be made.

The power detector of each transceiver is calibrated over a -15°C to $+55^{\circ}\text{C}$ temperature range and the calibration table is stored in flash memory.

Transmitter Out-of-Band Emissions

All unwanted out-of-band and spurious emissions from the transmitter are below -50 dBc, which is the FCC regulation limit for the ISM band. Figure 3.5.8 shows the transmitter output spectrum over 2.5 GHz with the transmitter operating at full power on 902.5 MHz (channel 1). This is the worst channel from the point of view of harmonic suppression as the 2nd harmonic is closer to the upper roll-off of the transmitter output filter. It can be seen that the -50 dBc specification is easily met.

Experimental Astrophysics (Ppd) 1 - 70

The Pierre Auger Project TDR




Figure 3.5.8 : Output spectrum of transmitter operating on channel 1 (frequency-hopping disabled)

RF Screening

The entire RF transceiver is contained within a screened area of the subscriber unit PCB, and within this, the high power RF stages of the transmitter are positioned away from the low power sensitive stages such as the synthesiser. Screening also ensures that noise generated by the subscriber unit digital electronics and surface detector electronics does not degrade receiver sensitivity.

RF Transceiver : Receiver Operation

Overview of Receiver

Figure 3.5.9 shows a block diagram of the receiver. The front-end of the receiver consists mainly of a band-selection filter and a low noise amplifier (LNA). There are also two receiver input protection features: a grounded quarter-wave stub located at the RF connector, and back-to-back protection diodes on the input of the LNA. These features are described below. The impact on receiver sensitivity of all front-end components, including the antenna switch, is also discussed.

The front-end filter presents a high degree of attenuation to frequencies outside of the 902-928 MHz band, whilst allowing frequencies in this band to reach the LNA with minimum loss. A degree of selectivity is also provided by the Yagi antenna. The output of the LNA feeds into the RF down-converter, and this translates the wanted channel down to the 71 MHz IF. The RF local oscillator tunes from 973.5MHz to 998.5 MHz (902.5+71 to 927.5+71) as the down-converter uses ‘high side’ injection of the local oscillator. Thus, the wanted channel is 71 MHz below the local oscillator. Energy located at the ‘image’ frequency 71 MHz above the local oscillator must be filtered out before it reaches the mixer to prevent it falling on top of the wanted channel at IF during the mixing process; this is achieved by the front-end filter. Image suppression is further enhanced through the use of an image-reject mixer in the RF down-converter stage.

The down converted IF then passes through a narrow Surface Acoustic Wave (SAW) filter, which attenuates all but the wanted channel energy. It then undergoes amplification before entering the IF I/Q demodulation stage. This stage outputs the quadrature baseband signals, I and Q, and these are routed to the baseband filtering section before being converted for digital processing.

Receiver Noise Figure

Before receiver sensitivity can be specified, the receiver noise figure (NF) must be known. The noise figure is a measure of the degradation in signal-to-noise ratio (SNR) between the input and output of the receiver. Noise factor (F) is the numerical ratio of NF, where NF is expressed in dB. Thus,

$$NF(dB) = 10\log_{10} (F)$$

and,

$$F = \text{Input SNR} / \text{Output SNR}$$

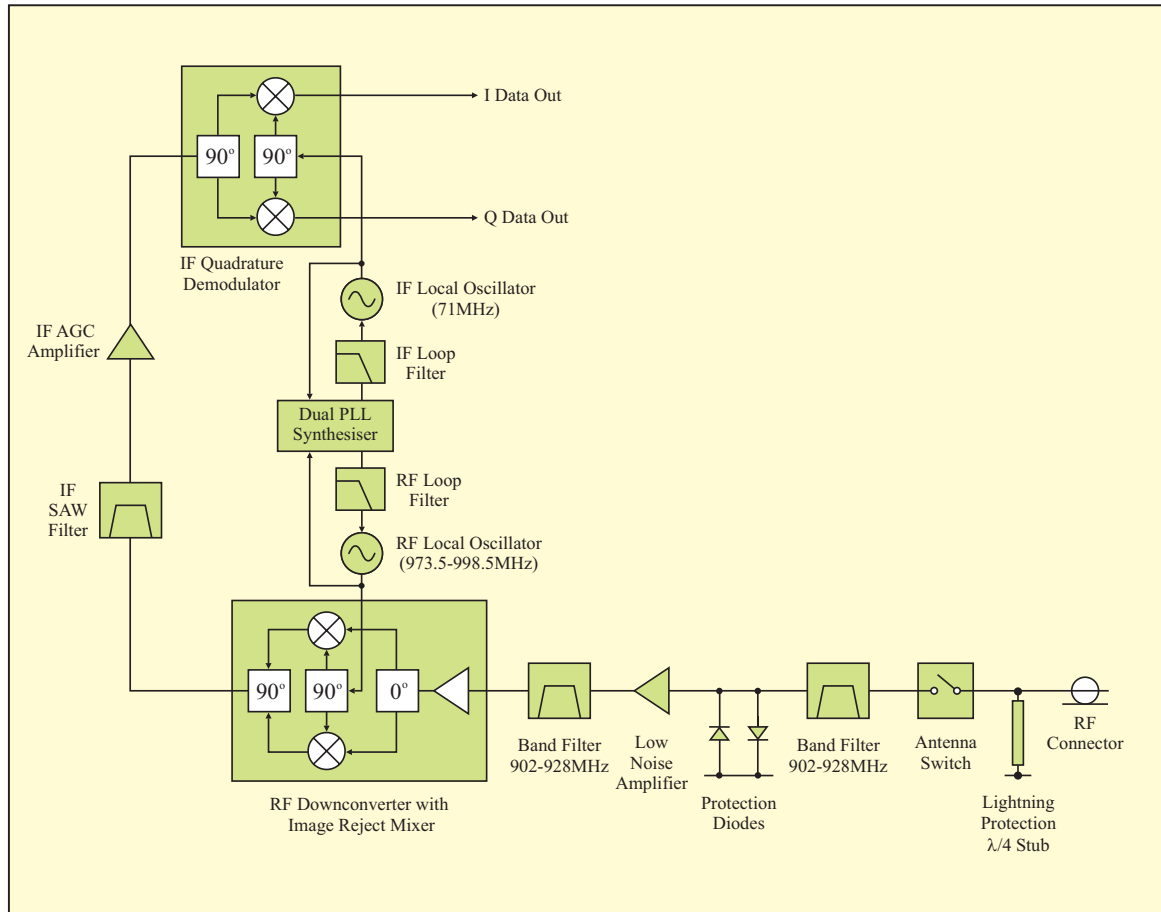


Figure 3.5.9 : Block diagram of the receiver

For a cascade of stages such as a receiver, the total noise factor is given by,

$$F_{\text{total}} = F_1 + (F_2 - 1)/G_1 + (F_3 - 1)/G_1 G_2 + (F_4 - 1)/G_1 G_2 G_3 \quad \text{Equation 3.5.1}$$

where,

F_n = the noise factor of each stage

G_n = the gain of each stage (Note : as a ratio, not in dB)

If a stage has gain < 1 , i.e. it attenuates the signal, the noise factor for that stage is the numerical value of the attenuation.

It is apparent from Equation 3.5.1 that the total noise factor of a cascade of devices is dominated by the first few stages, which, for a receiver will be the front end filter and LNA. Figure 3.5.10 shows the wireless LAN receiver as a cascade of stages to IF, each with an associated gain and noise factor (mid-band values). The gain shown for the IF

amplifier is the maximum available for this device, which is in fact a variable gain stage for receiver automatic gain control (AGC) purposes. At lower values of gain, the noise figure calculation becomes less important as the lower gain corresponds to a stronger input signal. It can be shown that noise figure degradation does not occur until the IF gain drops by ≈ 48 dB. It can also be shown that the contribution from devices beyond the IF gain stage is negligible and may be ignored for all practical purposes.

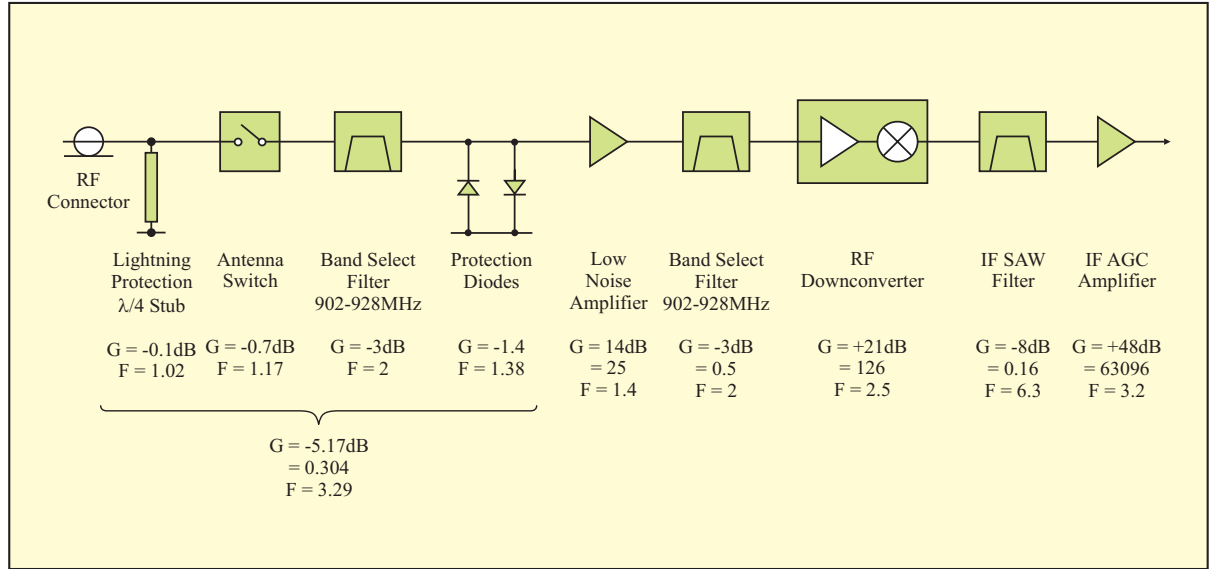


Figure 3.5.10 : Receiver gain distribution for noise figure calculation

As shown in Figure 3.5.10, the losses due to the lightning stub, antenna switch, band selection filter and protection diodes can be combined to simplify the calculation. The total noise factor can then be calculated thus :

$$\begin{aligned}
 F_{\text{total}} = & 3.29 + (1.4-1)/(0.304) + (2-1)/(0.304 \times 25) + \\
 & (2.5-1)/(0.304 \times 25 \times 0.5) + (6.3-1)/(0.304 \times 25 \times 0.5 \times 126) + \\
 & (3.2-1)/(0.304 \times 25 \times 0.5 \times 126 \times 0.16) = 5.17
 \end{aligned}$$

Hence,

$$NF = 10\log_{10} (5.17) = 7.1 \text{ dB}$$

A graph of the actual noise figure measured on a typical receiver over the 902-928 MHz band against the theoretical value is given in Figure 3.5.11. The measured value is within $\pm 0.5\text{dB}$ of the theoretical value. The discrepancy can be attributed to measurement

uncertainties and to gain variation of the devices across the band, as this is not taken into account in the theoretical calculation.

Receiver Sensitivity

The receiver has been designed to have sufficient sensitivity to satisfy the on-air bit error rate (BER) requirement of better than 1×10^{-6} [1]. The calculation for target sensitivity must take into account the type of modulation used, the equivalent noise bandwidth, and the bit rate. Modulation is QPSK, channel noise bandwidth is approximately 120 kHz (twice the bandwidth at baseband since the IF is double-sided), and the bit rate is 200 kbits/second.




Figure 3.5.11 : Receiver noise figure measurements

For a bit error rate of 1×10^{-6} using QPSK, the ratio of energy per received data bit, E_b , to noise power spectral density, N_o , is required to be 13 dB [5]. The E_b/N_o ratio is a very convenient metric for comparing different digital modulation schemes for a given bit error rate, but having specified the modulation and channel bandwidth it is more convenient to express sensitivity in terms of the signal-to-noise ratio. These parameters are related thus,

$$E_b/N_o = (S / R) / (N / B)$$

Where S = Signal power

N = Noise power

R = Bit rate (bits/second)

B = Bandwidth (Hz)

Hence,

$$S/N = (E_b \times R) / (N_o \times B) = (E_b/N_o) \times (R/B)$$

So the required SNR for a bit rate of 200 kbits/s and a bandwidth of 120 kHz is

$$SNR = 13 \text{ dB} + 10\log(200/120) = 15.2 \text{ dB}$$

Now, the noise power at 900 MHz is dominated by the internal thermal noise of the receiver. For analysis, it is convenient to treat the receiver as a noise-free device with an equivalent input noise power. At room temperature (290 K) this is given by

$$N = 10\log_{10}(kT/(1 \times 10^{-3})) + 10\log_{10}(B) + NF = -116 \text{ dBm}$$

Where k = Boltzman's constant = $1.38 \times 10^{-23} \text{ JK}^{-1}$ and NF is taken as 7 dB.

At 90°C, the noise power will rise by 1 dB.

The required input signal power to the receiver is therefore

$$S \text{ (dBm)} = \text{Equivalent input noise power (dBm)} + \text{Required SNR (dB)}$$

i.e.

$$S = -116 + 15.2 = -100.8 \text{ dBm}$$

That is, the minimum acceptable received signal strength for the wireless LAN receiver is -101 dBm, at which the bit error rate should ideally not exceed 1×10^{-6} .

Figure 3.5.12 shows a plot of the measured and theoretical bit error rate against received signal strength for differential binary phase-shift keyed (DBPSK) modulation for the wireless LAN transceiver. The measured data was obtained from a run of approximately 10 million data bits. The theoretical curve assumes perfect symbol synchronisation and flat gaussian noise across the channel bandwidth and zero implementation loss.

It should be noted that the length of the test run was insufficient for measuring bit error rates less than 1×10^{-7} reliably, but this was not considered to be an issue for concern as the target bit error rate was 1×10^{-6} , and this target had clearly been achieved at a respectable input level of -97 dBm.

It can be seen that the shape of the measured BER curve is shallower than the theoretical curve as the signal level drops and the probability of error increases. This is because the test was performed using the packet protocol employed for normal data transmission, in which whole packets are discarded if there is a single bit error in the Barker sequence used for frame synchronisation. An error count cannot be performed on these discarded packets, and consequently this distorts the overall measurements slightly.

The difference of 3-4dB between the measured and predicted bit error rates at signal levels above -100 dBm can be attributed to implementation losses, including imperfect synchronisation, and RF mixing. There is also an error of ± 2.0 dB in the signal strength measurement. The error rate performance of the receiver is well within expected norms.




Figure 3.5.12 : Measured versus theoretical bit error rates

Receiver Minimum Detectable Signal

This is a convenient way to specify receiver sensitivity without discussing bit error rate and signal-to-noise ratio requirements, and is simply the input signal power at which the ratio of (signal+noise)/noise at the output of the receiver is 2 (or, in logarithmic terms, 3dB). Clearly this will occur when the input signal power is equal to the equivalent input noise power. For the wireless LAN receiver this is approximately -116 dBm.

Receiver Phase Noise

The ultimate bit error rate performance in a communications system employing phase modulation is limited by the instantaneous phase error on the local oscillators. Every oscillator exhibits ‘phase noise’ to some degree, and this appears as noise sidebands either side of the fundamental oscillation frequency. Phase noise is specified in terms of noise power spectral density relative to the fundamental at some frequency offset. Crystal oscillators exhibit very low phase noise, a typical figure being -150 dBc/Hz at 10 kHz offset. A synthesised source, however, cannot match this, as phase noise is degraded by all elements in the synthesiser loop. A laboratory bench signal generator will typically have a phase noise of -120 dBc/Hz at 10 kHz. The phase noise specification for the phase-locked output of the RF local oscillator in a digital mobile phone, however, is much more relaxed at approximately -78 dBc/Hz at 10 kHz. This equates to a r.m.s. phase error of 5° , which is the specification for the GSM mobile telephony standard.

In addition to phase error, the phase noise sidebands on the local oscillator can result in a phenomenon called ‘reciprocal mixing’ (see below). This occurs in a receiver when a

Experimental Astrophysics (Ppd) 1 - 78

The Pierre Auger Project TDR

high-level signal in an adjacent channel mixes the noise sidebands of the local oscillator down to IF, reducing the dynamic range of the receiver.

The phase noise of the phase-locked RF local oscillator for the wireless LAN transceiver is shown in Figure 3.5.13. At 10 kHz offset the phase noise is a respectable -84 dBc/Hz.




Figure 3.5.13 : RF local oscillator phase noise

Receiver Selectivity

The receiver selectivity describes the amount by which unwanted channels are rejected, or conversely the extent to which energy on other channels impinges on the wanted channel. This specification is most easily expressed in terms of dB's of attenuation in the IF filter. Figure 3.5.14 shows the measured response of the receiver IF SAW filter. Suppression of unwanted channels is better than 35 dB for adjacent channels and typically better than 45 dB for non-adjacent channels.




Figure 3.5.14 : Receiver IF SAW filter response

Receiver In-Band Blocking

In-band blocking refers to the de-sensitisation of a receiver to a wanted signal in the presence of a large in-band signal. It is caused by gain compression and is a measure of the linear dynamic range of the receiver above the minimum detectable signal. It therefore depends on the gain distribution in the receiving system. The wireless LAN has two front-end gain modes, a high (normal) gain mode for optimum receiver noise figure, and a reduced gain mode to optimise linearity. There are therefore two blocking or gain compression figures: with the front-end gain set to high, 1 dB gain compression occurs when the in-band interferer is at -31 dBm, which is 80 dB above the minimum detectable signal. With front-end gain set to low, 1 dB compression occurs at approximately -19 dBm, which is 95 dB above the minimum detectable signal.

The blocking performance also depends on the IF amplifier gain. The blocking effects of an interferer on a close-in channel will degrade if the IF gain is high, which will be the case if the wanted signal is close to the receiver sensitivity level. This is because the IF SAW filter presents a lower attenuation to energy on close-in channels, and this energy is sufficient to drive the IF amplifier into compression. This is illustrated in case (A) of Figure 3.5.15, which shows the level of in-band interferer required on each channel to cause a 1 dB compression in the wanted signal output level, where the wanted signal is at -105 dBm on 915 MHz (mid-band). This is an absolute worst case situation as the receiver IF gain will be set close to maximum. The ripple in the SAW filter response is

Experimental Astrophysics (Ppd) 1 - 80

The Pierre Auger Project TDR

clearly visible on this plot. In contrast, the second plot (B) shows how the blocking performance improves dramatically when the receiver IF gain is reduced when the wanted signal is at a higher level of -70 dBm, which is a more typical case.




Figure 3.5.15 : Worst and typical case receiver blocking performance

Receiver Reciprocal Mixing

This occurs when a large in-band signal mixes with the phase noise sidebands of the local oscillator, causing an increase in noise at IF, and therefore a degradation in SNR on the wanted channel. Figure 3.5.16 shows the effects of reciprocal mixing on 915 MHz (midband) for a worst case scenario (receiver IF gain set to maximum). This figure shows the level of interferer required on other channels to cause a 3 dB increase in noise power at the baseband output. The level falls as the interferer approaches 915 MHz due to the increasing phase noise on the local oscillator. For adjacent and close-in channels, however, receiver blocking effects begin to dominate, making reciprocal mixing measurements on these channels invalid.




Figure 3.5.16 : 3dB noise increase due to reciprocal mixing

Receiver Image Rejection

Rejection of the unwanted receive channel image is better than 70 dB for all channels.

Receiver IF Rejection

Rejection of energy at 71 MHz is greater than 80 dB for all channels.

Receiver spurious responses

Spurious responses occur when a signal at a frequency to which the receiver is not tuned produces an output at the receiver baseband. They are caused mainly by harmonics of the local oscillators, together with non-linearities in the mixing stages. An important spurious response is the '2:2 response', which is a measure of the receiver's rejection of energy $\pm IF/2 = \pm 35.5$ MHz away from the wanted signal. This will be closer than the unwanted image frequency, but the receiver's response to this will be much lower if the 2nd harmonic of the local oscillator is sufficiently low and the mixing stage is not over-driven. This response is expressed in terms of dB above the wanted signal to cause a baseband output at the same level as the wanted. It has been measured as being >50 dB on channel 1, increasing to >80 dB on channel 51. This figure is considered to be more than sufficient for the Auger wireless LAN. There are no other spurious responses considered to be significant.

Experimental Astrophysics (Ppd) 1 - 82

The Pierre Auger Project TDR

Receiver ‘Self Quieting’ (Demodulation of Oscillator Harmonics)

There is no discernible ‘self-quieting’ effect due to demodulation of in-band harmonics of either the 20 MHz reference clock or the 71 MHz IF local oscillator.

Receiver Input Protection Features

As mentioned above, the receiver front-end has two protection features. A quarter-wave lightning protection stub, tuned to 915 MHz, ensures that the transceiver presents a solid ground potential at out-of band frequencies. This prevents high potentials developing at the input of the LNA due to high level, wideband energy such as lightning. The RF filter also presents a short circuit at DC.

Protection diodes limit the in-band RF voltages at the LNA input to ± 0.7 V.

Receiver Calibration

The receiver gain is calibrated over temperature during burn-in and the calibration table stored in flash memory. This allows each radio to report received signal strengths in dBm, which provides a powerful network monitoring facility.

3.5.1.2.2 Analogue Baseband Signal Processing

Transmitter Analogue Baseband

As indicated in the discussion of the transmitter above, the two baseband data streams for the quadrature modulator of the transmitter are filtered by pulse-shaping low-pass filters. The data streams themselves are generated by the DSP using 8-bit digital to analogue converters. The filters are 4-pole Bessel type for linear phase response, and are implemented using fast op-amps. The 3 dB cut-off frequency is approximately 60 kHz.

Receiver Analogue Baseband

The receiver has identical filters on the outputs of the quadrature demodulator for anti-aliasing prior to sampling by 10-bit analogue-to-digital converters. The sampling rate is approximately 300 ksamples/s.

The quadrature demodulator has differential outputs and so a differential stage is added before the filters. This provides 6 dB of ‘free’ power gain, thereby reducing the gain required in subsequent stages before sampling. It also allows the demodulator outputs to be dc-coupled, which alleviates problems caused by time constants otherwise introduced into the receiver gain control system. The differential stage, which is co-located with the filters outside of the RF screening can, also reduces the effects of low frequency noise impinging on the signal path between the RF demodulator and the filters by virtue of improved common mode rejection.

3.5.1.2.3 Subscriber Unit Processor and Transceiver Control

Central Processor Unit and Non-Volatile Storage

The software that controls all functions of the subscriber unit runs on an Analog Devices DSP2183 fixed-point digital signal processor. This device runs at 40 MIPS, has 16k words of program memory and 16k words of data memory. The software image is stored in a page of non-volatile flash memory and is loaded into the DSP program memory when the subscriber unit powers up. Pages in the flash memory are also reserved for calibration look-up tables and radio production history. A further page is used as a holding area by the WLAN on-air software download facility.

The software is protected against spurious lock-up via a watchdog monitor. A hardware counter on this device must be regularly reset by the DSP; if it fails to do so the watchdog will reset the DSP.

Subscriber Unit ↔ Local Station Communications Interface

The DSP has two serial ports, one of which is used for communications with the local station via a standard RS232 driver and UART, although the pin-out of the 9-way D-type is non-standard; this is to allow for additional lines, specifically: 12V DC input power, a 1 pulse-per-second line from the GPS engine, a Reset line to the local station, and a Reset line from the local station.

The subscriber unit ↔ local station interface runs at 38400 baud.

Slow Analogue Interface

The second serial port on the DSP is used to transfer digital data to and from a ‘slow’ analogue interface; this supports eight 10-bit analogue inputs and four 10-bit outputs. The input sources are dedicated to built-in self-test (BIST) features, and consist of a temperature sensor, the RF power detector circuits, and voltage and current sensors associated with the 5V and 3.3V power supplies. The temperature sensor is calibrated to $\pm 2.0^{\circ}\text{C}$, and the current sensors allow the current consumption of individual functional blocks of the transceiver to be monitored.

The slow analogue outputs are used for controlling receiver and transmitter gain stages, and also for frequency adjustment of the 20MHz reference oscillator. The 10-bit resolution of the digital to analogue converters, together with the stable 1 pps second signal from GPS, allows the RF output to be tuned to $\pm 25\text{Hz}$

Fast Analogue Interface

The analogue I and Q outputs from the receiver baseband are sampled and converted by a pair of ‘fast’ 10-bit converters, and each 10-bit word is subsequently transferred to the DSP via the parallel data bus. To preserve the phase information, sampling of the I and Q channels is performed simultaneously, and each sample held in a buffer until the DSP is

ready to read it. Sampling is driven by a timer-based interrupt on the DSP; the timer is programmed to produce sampling interrupts at approximately 300 ksamples/s.

In a similar manner, the I and Q signals required by the transmitter are generated by the DSP and each is transferred in turn over the parallel data bus to 8-bit digital-to-analogue converters. They are then held in input buffers until a DSP timer-interrupt fires, and this causes the converter outputs to be updated simultaneously, thereby maintaining the correct quadrature phase relationship in the transmitted signal.

3.5.1.2.4 Power Supplies

The subscriber unit is designed to run from a regulated 12V DC supply, though it will operate over an input range of 8-16V. There are 2 supply voltages used on the PCB: 3.3V and 5V, and these are generated using switched mode devices with an efficiency of typically 90% at full load. These devices have a switching frequency of approximately 160 kHz.

The majority of devices on the subscriber unit PCB run from 3.3V, and this rail is split into two independently-filtered power planes: a 3.3V digital plane and a 3.3V analogue/RF plane. The 5V rail, which feeds the PA and LNA, is a star-point feed to minimise common mode impedances on sensitive RF bias lines. The ground plane is 'slotted', again to reduce common impedances and ensure that digital currents do not flow under sensitive analogue stages.

The 5V and 3.3V supply rails each have a current and voltage monitoring facility for built-in self-test.

Subscriber Unit Power Consumption

DC power consumption during a full-power TX timeslot is approximately 3.36 Watts, whilst during a RX slot it is 960 mW. The average power consumption depends on the number of TX and RX timeslots, together with the processing load during the slots when the subscriber unit is off-line. A typical operational mode will have 2 TX timeslots and 4 RX timeslots, and power consumption over the remaining slots will be approximately 890 mW. This leads to an average power consumption of approximately 1.02 Watts.

3.5.1.3 Subscriber Unit Specification – Summary

A summary of the Auger wireless LAN subscriber unit is given in Table 3.5.1.

Frequency band	902-928 MHz
Modulation	DQPSK
Bit rate	200 kbits/second
Channel spacing	500 kHz
RF Frequency stability :	
GPS disciplining enabled	± 25 Hz
GPS disciplining disabled	± 2 kHz (2.0ppm)
Receiver noise figure	7 dB
Receiver sensitivity	15 dB SNR at -100 dBm input
Adjacent channel suppression	35 dB
Receiver image rejection	> 70 dB for all channels
Receiver IF rejection	> 80 dB
Receiver 2:2 response (Interferer spacing = $IF/2 = 35.5$ MHz)	
Channel 1 (902 MHz)	> 55 dB
Channel 27 (915 MHz)	> 70 dB
Channel 51 (927.5 MHz)	> 80 dB
Other spurious responses	> 80 dB
Minimum detectable signal	-116 dBm
In-band 1dB blocking dynamic range :	
Normal receiver gain mode	80 dB
Low receiver gain mode	95 dB
Phase noise	< -82 dBc/Hz at ± 10 kHz offset
Transmitter out-of-band emissions	< -50 dBc
Transmitter output power	+25 dBm into 50 Ohms
Average DC power consumption	< 1.1 Watts at 8-16 Volts DC

Table 3.5.1 : Subscriber unit specification

3.5.1.3 Subscriber Unit Testing

To ensure that reliable subscriber units are delivered to the Observatory site, each unit undergoes several phases of testing on automatic test equipment (ATE) terminals. The first is a basic functionality test performed on the circuit boards at the factory, and this is carried out on a stand-alone rig termed the Micro-ATE terminal. This set-up is duplicated at Leeds and also at the Observatory site in Malargüe. The Micro-ATE terminal consists of a PC, a known good reference radio, a 12V DC supply, and a current meter. The reference radio is used to perform a series of transmitter and receiver tests on the unit under test, such as receiver sensitivity, transmitter output power level, synthesiser tuning, etc, and it provides a relatively cheap but extremely valuable diagnosis tool; it ensures that all units delivered by the factory have satisfied the basic functionality requirements. It is also useful on site for fault diagnosis should a unit fail in the field.

Each subscriber unit has an electronic production history stored in flash memory. This contains the unit's production number (assigned by the Micro-ATE terminal the first time the unit is tested), the name of the tester, the number of times it has been tested, the date the last test was performed, and the version of the firmware stored in flash. The result of the last test (pass or fail) is also stored in the history, which can be displayed on screen at any time.

The Micro-ATE Terminal also generates a test report which, in the event of a Fail, contains fault-finding tips, such as “check components R56, C33, U1”, etc.

Two screen shots from the Micro-ATE Terminal are shown in Figures 3.5.17 (a) and (b).

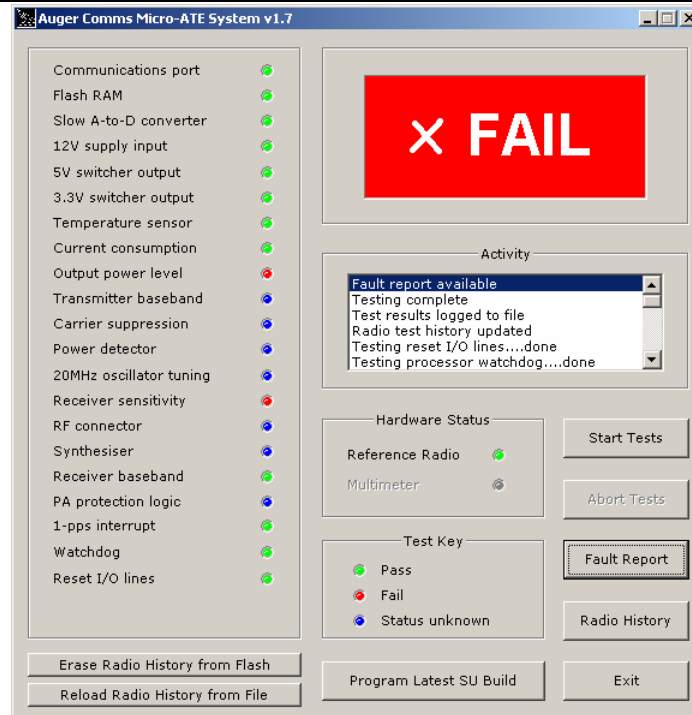


Figure 3.5.17(a) : Screen shot from the Micro-ATE Terminal:Unit Failure

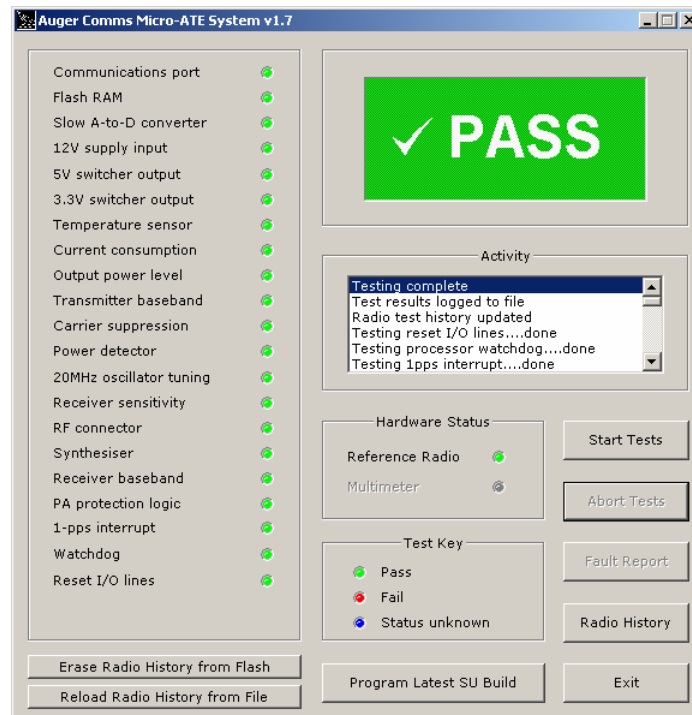


Figure 3.5.17(b) : Screen shot from the Micro-ATE Terminal: Unit Pass

After initial testing on the micro-ATE terminal, subscriber units are “burned-in” for 48 hours at 60 degC. They then undergo extensive testing in an environmental test chamber for approximately 8 hours as part of the “full-ATE” phase. This process tests and calibrates the units in batches of 8 from –15 to +55 degC. The RF tests are more extensive than those performed by the micro-ATE terminal, and include phase noise measurements, I/Q phase imbalance, carrier suppression optimisation, measurement of oscillator reference spurs, and so on, in addition to full calibration of transmitter output power and receiver gain. Units must pass full-ATE before they can be assigned an electronic serial number, after which they are given a final Network Connection test before being shipped to the Observatory site.

3.5.2 Base-Station Radio Unit

3.5.2.1 Base-Station Radio Unit Overview

The processing demands of a base-station unit far exceed those of a subscriber unit. The additional processing burden is due to the increase in duty cycle experienced by a base-station unit, which is continuously transmitting, or receiving packets. This is in contrast to the subscriber units, which are inactive for the majority of the time. The subscriber unit can reduce the peak processor loading during its active timeslots by spreading non real-time tasks into the inactive periods, in contrast to the base-station, which must constantly handle on-air traffic.

An additional *base-station support board* has been designed that is used in conjunction with a subscriber unit transceiver board to form a base-station radio. The architecture of a base-station is illustrated in Figure 3.5.18.

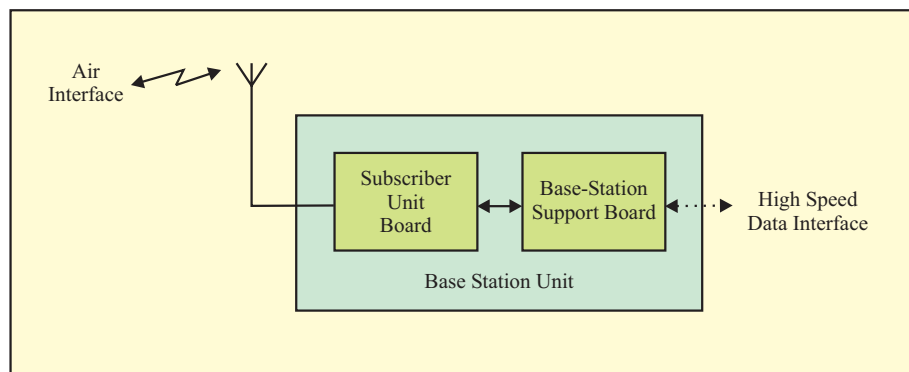


Figure 3.5.18 : Base-station radio unit architecture

3.5.2.2 Base-station Support Board Hardware Description

The base-station support board PCB is the same size as the subscriber transceiver PCB, and they both stack together to fit into the standard aluminium subscriber unit case. A custom base-station front panel allows access to the normal subscriber unit connectors in addition to the base-station connectors (Figure 3.5.19).

Figure 3.5.20 shows a photograph of the PCB with various sub-sections indicated, and Figure 3.5.21 illustrates the system architecture of the board.



Figure 3.5.19 : Base-station unit




Figure 3.5.20 : Base-station support board PCB

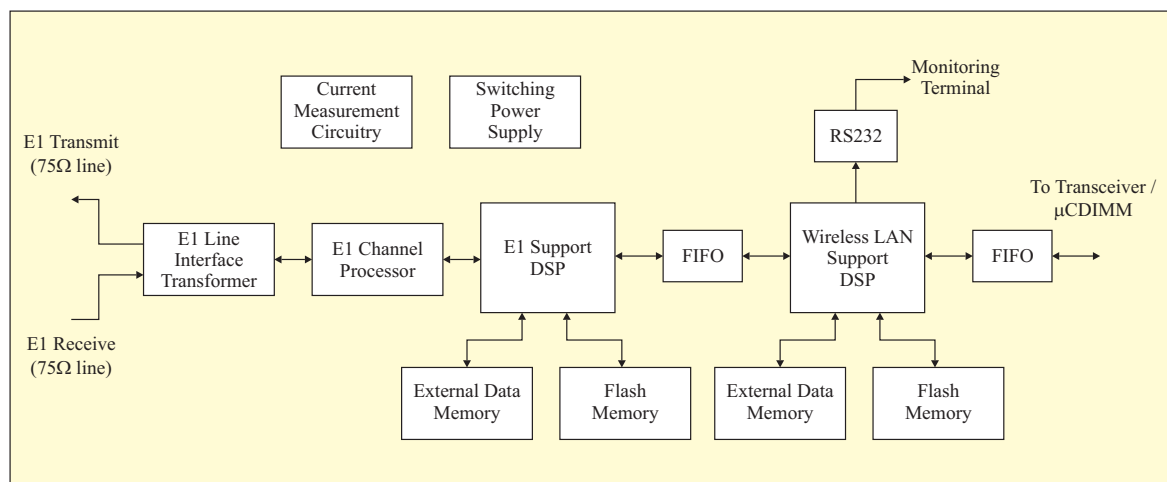


Figure 3.5.21 : Base-station support board architecture

Base-station units are required to concentrate the data traffic from multiple subscriber units and therefore require a high-speed data interface that cannot be provided by the standard RS232 port on the subscriber unit board. High-speed data communications is provided by an industry standard E1 interface which provides a 2.048 Mbps data transfer rate. The interface provides sufficient bandwidth to support up to 7 base-station units on an E1 ring connected to a single uLSx unit. An E1 channel processor on the support board manages the E1 interface, synchronising to the data stream and monitoring for alarm conditions. The device is connected to 75 ohm balanced line coaxial media via an impedance matching E1 line transformer. An overview of the E1 bearer system is presented in § 3.4.6.3, *Base-station Unit to μLSx Communications*.

The *E1 support DSP* manages the software protocol used to transfer packets of data to the E1 channel processor, whilst the *wireless LAN support DSP* runs the main base-station scheduler. Both DSPs are fitted with 1 Mbit of external static RAM data memory to compliment the 256 kbits internal data memory of the devices. This additional memory resource enables the E1 support DSP to run the E1 streams environment, and would allow a processor-demanding FEC coding scheme to be implemented in the future, if required.

Both DSPs are also fitted with 4Mbits of flash ram for non-volatile storage of bulk data and the operating code.

The wireless LAN support DSP has two external interfaces. A FIFO interface provides a high speed 640 Mbps connection to the subscriber unit board used in the base-station, and an identical FIFO interface is used to connect the wireless LAN support DSP with the E1 support DSP. A standard RS232 port is also fitted, and this enables a monitoring terminal to be connected directly to the base-station to provide communications network status information.

An on board switching power supply unit efficiently provides the unit with the necessary operating voltages. Current measurement circuitry provides the unit with the hardware required to perform accurate current measurements, enabling built-in self-test routines to be used to ensure the unit is operating correctly.

3.5.3 The μ LSX Unit

The micro-LSx unit provides an interface between a number of base-stations at a tower and the local TCP/IP network. The unit contains two PCBs; a base-station PCB and a uCDimm TCP/IP PCB. A photograph of the latter is shown in figure 3.5.23. The uCDimm board uses a commercially-available single board processor module running an embedded linux variant to implement standard TCP/IP connections with the campus CDAS system for each base-station radio.

The architecture of the complete 2-board unit is shown in figure 3.5.22. The base-station board controls the local E1 ring (refer to figure 3.4.26) and passes data to the uCDimm board via 2-way FIFO memory. The uCDimm board opens one socket connection per base-station with the CDAS system and manages the routing of data packets to and from the CDAS system and each individual base-station. A single micro-LSx interface can handle traffic from between 4-to-6 fully-loaded base-station radios, requiring around 8 interfaces to be used to handle the traffic from the 28 base-station radios that cover the full surface array.

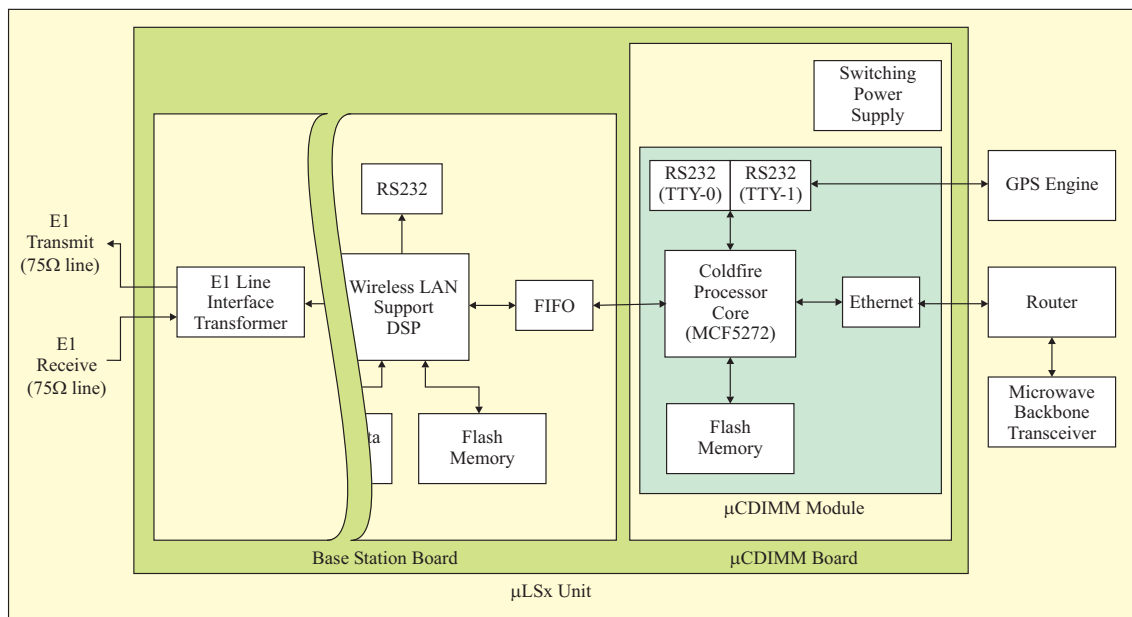


Figure 3.5.22 : Micro-LSX Architecture



Figure 3.5.23 : Micro-LSX PCB

REFERENCES

- [1] Nitz, D., “Triggering and Data Acquisition Systems for the Auger Observatory”, 10th IEEE Real Time Conference, Beaune, France, 1997.
- [2] Comision Nacional de Comunicaciones, <http://www.cnc.gov.ar/>
- [3] Federal Communications Commission (FCC) Rules & Regulations, Part 15 : Radio Devices, Section 15.247 : Operation within the bands 902-928MHz, 2400-2483.5MHz, and 5725-5850MHz, October 2000
- [4] Freeman,R.L., “Radio System Design for Telecommunications”, Published by J.Wiley & Sons, Inc. ISBN 0-471-16260-4, 2nd Edition, Chapter 14, p.734.
- [5] Proakis, J.G., “Digital Communications”, pp269-278, 3rd Edition, McGraw-Hill, 1995.

3.6 VOICE COMMUNICATIONS SYSTEM

The surface detector array covers 3000km² of sparsely populated, relatively inhospitable and remote pampas. Observatory personnel operating in the field may be several hours drive from main roads and many kilometers from colleagues elsewhere within the array. There are clear requirements for personnel to be able to communicate with their colleagues at the Campus, eye buildings and elsewhere on the pampas and hence a voice communications system has been provided for use by personnel working on-site. It enables teams in the field to exchange information between themselves and the Campus and also functions as a backup system, providing communications in the event of emergency situations that may arise when working in the field.

As of March 2003 the system consists of repeaters located at the Observatory Campus and at Coihueco, together with a number of handheld units and vehicle sets. The architecture is flexible, and will support additional repeaters to eventually provide full coverage of the Southern Observatory array area.

3.6.1 Brief Outline of Operation

The system employs ‘off-the-shelf’ voice communications hardware, operating at 473 MHz. The existing repeater antennas are mounted at the top of the 50 meter Campus tower (Figure 3.6.1) and the 20m Coihueco tower (Figure 3.6.2). This elevated position significantly increases the range that may be achieved from the low powered handheld radios. A typical radio is shown in Figure 3.6.3.

Each handheld radio and vehicle set uses a pair of frequencies to form a channel, one for transmission, f_1 and one for reception, f_2 . All the radios use the same frequencies. The repeater, however, reverses normal operation and uses the handheld reception frequency for transmission, and the transmission frequency for reception.

This means that the repeater ‘hears’ anything transmitted by any handheld in its range on f_1 and then re-transmits or ‘repeats’ this signal with much greater range and power on frequency f_2 , on which it is heard by other mobiles. Figure 3.6.4 illustrates the principle where handheld #1 is talking to handheld #2 indirectly, via the repeater.



Figure 3.6.1 : Voice communications monopole antenna at the Observatory Campus




Figure 3.6.2 : Voice communications monopole antenna at Coihueco



Figure 3.6.3 : Voice communications handheld radio




Figure 3.6.4 : Operation through a repeater

3.6.2 Multi-site Repeater Operation

The microwave backbone network enables the linking of several repeaters in order to provide seamless operation throughout the site, effectively creating a single repeater zone as illustrated in Figure 3.6.5. However, it is undesirable to operate all the repeaters on the same channel as zones of mutual interference form in regions of the site where the signal strengths from 2 or more repeaters are approximately equal. Differential delays in the audio paths to the individual repeaters would cause serious destructive interference in the signals heard by radios within these regions, leading to severely impaired communications.

To avoid this problem, each repeater operates on a unique “channel”, which is in fact a channel pair, and this must be used by a radio operating within its area. This can be operated in an automatic mode as follows: The repeaters periodically transmit a sub-audible signal on their channel. The handheld radios continuously scan the available channels for these signals in order to find the strongest local repeater, i.e. in Figure 3.6.5 the strongest repeater for the indicated handheld is Coihueco. This operation is transparent to the user, who just perceives a single repeater zone as they move throughout the site.

When the handheld transmits, the repeater at Los Morados will receive the transmission and broadcasts it as normal. Concurrently, the Los Morados repeater relays the audio signal via the microwave links to the other repeaters at Campus and Coihueco, which transmit the original transmission on their respective channels, thereby covering the entire

site. Figure 3.6.6 shows the parts of the site that will be best served by each repeater. This simulation was performed on Solway, the propagation analysis tool used for WLAN planning.

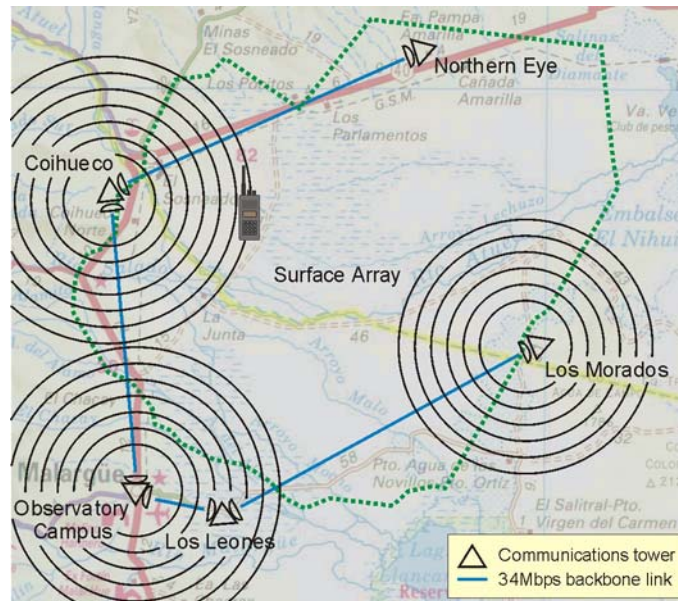


Figure 3.6.5 : Voice communication repeater links through the microwave backbone

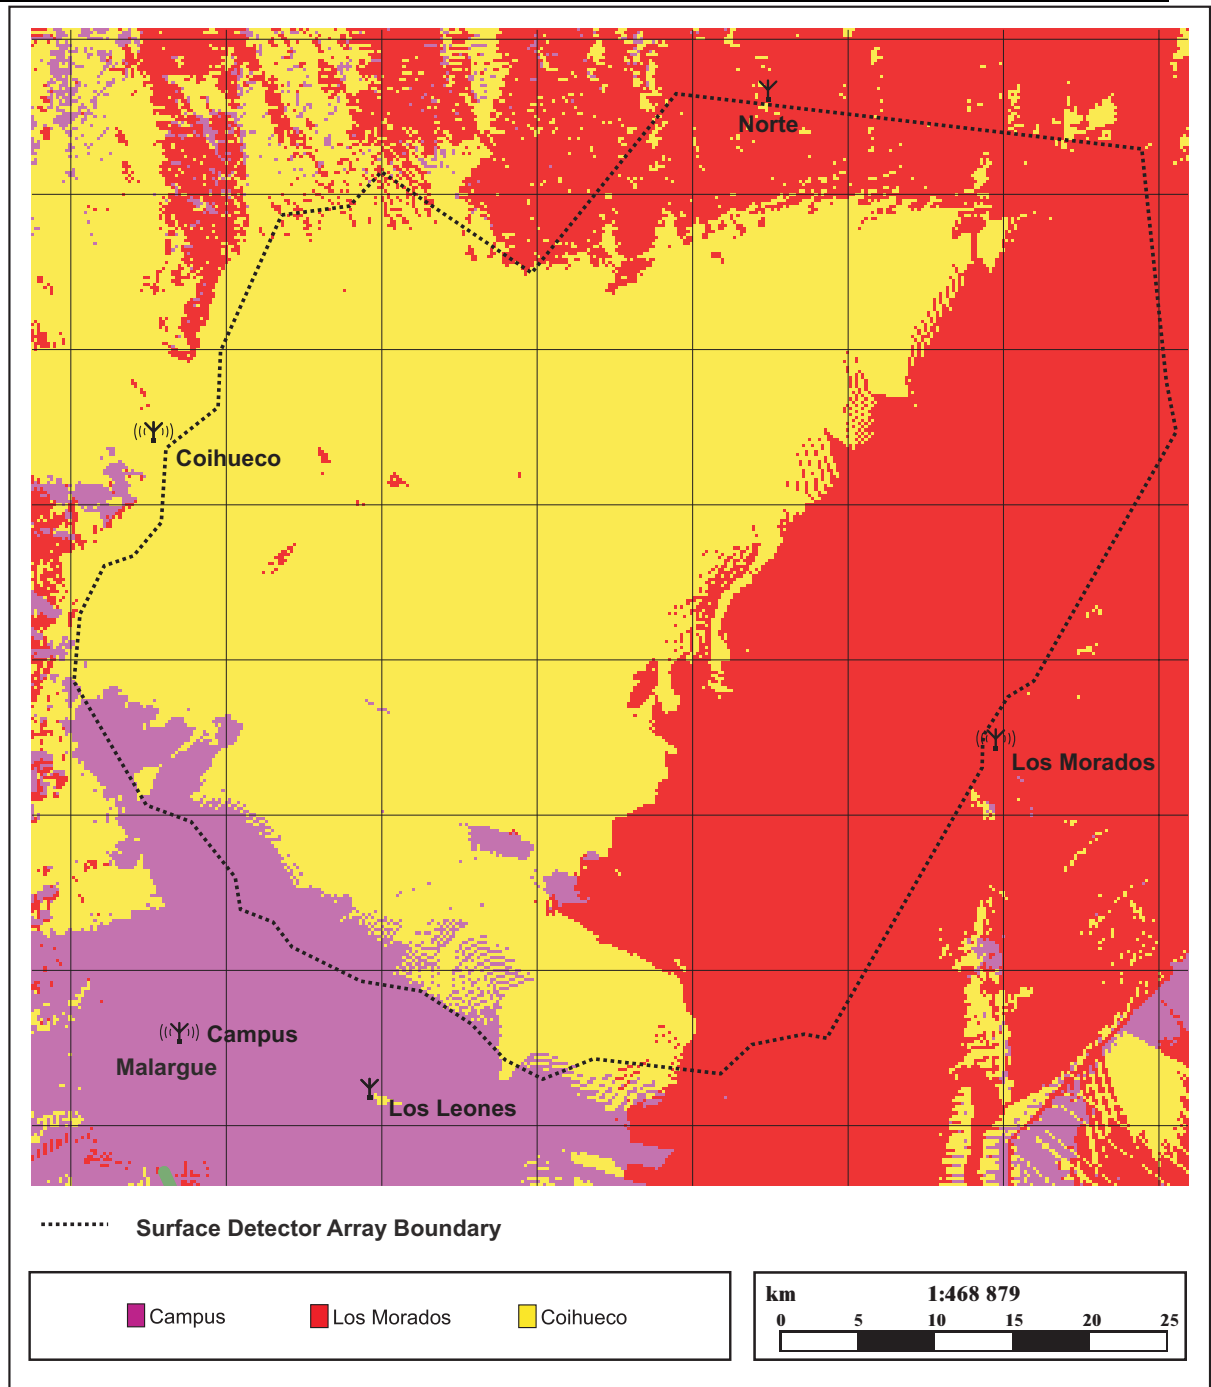


Figure 3.6.6 : Repeater coverage zones