

60423660 . 1.10402

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
PROVISIONAL APPLICATION FOR LETTERS PATENT

EV188390535

A Wireless Data Packet Communications System

Inventor(s):
Marcus da Silva
William J. Crilly, Jr.
James Brennan
Robert J. Conley
Siavash Alamouti
Edward Casas

ATTORNEY'S DOCKET NO. MN1-010USP1

1 **TECHNICAL FIELD**

2 This invention relates to wireless communications and more particularly to
3 methods and apparatuses for use in wireless data packet communications systems
4 capable of supporting multiple point-to-point links, packet-by-packet steering, and
5 the like.

6
7 **BACKGROUND**

8 Conventional wireless local area network (LAN) systems typically employ
9 a micro-cellular arrangement, wherein, for example, a small base station, often
10 referred to as an Access Point (AP), is configured to communicate with wireless
11 devices attached to computing devices, such as, laptops or other portable data
12 appliances. These APs have a limited range, typically 20 to 200 feet for an IEEE
13 802.11(b) system. Thus, to cover a large area a system may require a plurality of
14 APs. This can be costly and tends to complicate the wireless system.

15 There is a need for improved methods and apparatuses that can provide
16 wireless communications.

17
18 **DESCRIPTION**

19 The following description sets forth a specific embodiment of a wireless
20 communications system that incorporates elements recited in the appended
21 exemplary claims and others. The embodiments are described herein and in the
22 attached documentation. However, the description itself is not intended to limit
23 the scope of this patent. Rather, the inventors have contemplated that the
24 invention might also be embodied in other ways, to include different elements or
25

1 combinations of elements similar to the ones described in this document, in
2 conjunction with other present or future technologies.

3 With this in mind, methods, apparatuses, and systems are provided for a
4 project code-named "Little Joe" developed by Vivato Incorporated (formally
5 known as Mabuhay Networks) having research and development offices in
6 Spokane, WA, and headquarters in San Francisco, CA.

7 Incorporated herein are the following appendices:

- 8 A. Document: Little Joe Functional Specification (113 Pages)
9 B. Document: Behavior of 802.11 Networks With Mabuhay Access
10 Points (15 Pages)
11 C. Document: Beamforming for Little Joe (25 Pages)
12 D. Document: Little Joe Link Budget (15 Pages)
13 E. HTML Document: SimpleMAC (1 Page)
14 F. Presentation Slides: Prototype Story Board (17 Pages)
15 G. Copy of various lab notebook pages 2-3,45-49, 57, 89-93, 106-107
16 (15 Pages)
17 H. Drawings: System Diagram (9 Pages)
18 I. Document: Quickfacts – DirectedPacket 1 (5 Pages)
19 J. Document: Conceptual Proposal for Little Joe Antenna (3 Pages)
20 K. Presentation Slides: DirectedPacket™ 1 Logical View (7 Pages)
21 L. Presentation Slides: Little Joe Digital Block (5 Pages)
22
23
24
25

1 **EXEMPLARY CLAIMS**

2 **1.** A method comprising:
3 using at least one electronically steerable phase array antenna in a wireless
4 data packet communications system.

5
6 **2.** The method as recited in Claim 1, wherein using said at least one
7 electronically steerable phase array antenna in said wireless data packet
8 communications system further includes:
9 using said at least one electronically steerable phase array antenna indoors.

10
11 **3.** The method as recited in Claim 1, wherein said wireless data packet
12 communications system uses packet-by-packet steering.

13
14 **4.** The method as recited in Claim 1, wherein said wireless data packet
15 communications system implements a CSMA scheme.

16
17 **5.** The method as recited in Claim 1, wherein said wireless data packet
18 communications system includes a scanning receiver.

19
20 **6.** The method as recited in Claim 1, wherein said wireless data packet
21 communications system uses array windowing.

22
23 **7.** The method as recited in Claim 1, wherein said wireless data packet
24 communications system includes a Butler Matrix.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

8. The method as recited in Claim 1, wherein said wireless data packet communications system includes at least one Little Joe device.

9. A Little Joe device.

10. A communications system having at least one Little Joe device.


11. A propagated signal from a Little Joe device.

A

Little Joe Functional Specification

[113 pages]

A-1



Document History

Contributor(s)	Description	Date
-----	-----	
-----	-----	

A-2

Table of Contents

Contents.....	3
Part 1: Radio and MAC.....	7
1. Introduction.....	7
2. System Architecture.....	8
2.1. System.....	8
2.2. Deployment Scenarios.....	10
2.2.2. Typical Building Structure.....	12
2.2.3. Traffic Load Assumptions.....	12
2.3. Interference effects.....	13
2.4. Coverage requirements.....	13
2.5. Throughput and delay requirements.....	15
2.5.1. Performance Metrics.....	15
2.6. Performance limiting issues.....	15
2.7. Technology.....	16
2.8. System Block Diagram.....	16
3. The Beamforming Network.....	19
3.1. The Butler Matrix.....	19
3.2. Ideal Beam Patterns.....	20
3.3. Receive Windowing.....	23
3.3.1. The Hamming Window.....	23
3.4. Complementary Transmit Beamforming.....	26
4. Link Budgets.....	28
4.1. Link Power Budget.....	28
4.2. Link Budget Parameters.....	30
4.2.1. Panel Transmit Power: FCC EIRP Power Limits.....	30
4.3. Characteristics of Conventional 802.11b Equipment.....	31
4.3.1. Panel Antenna Gain.....	33
4.3.2. Gain Reduction due to Scattering.....	33
4.3.3. Client Antenna Gain.....	34
4.3.4. Client and Panel Noise Figures.....	34
4.3.5. Effect of Shadow Fading.....	35
4.3.6. Effect of Rayleigh Fading.....	36
4.4. Path Loss Models.....	36
4.4.1. Propagation in Free Space (LOS).....	37
4.4.2. Propagation by Diffraction (NLOS).....	37
4.4.3. Propagation by Transmission (OBS).....	37
4.4.4. Outdoor-Indoor Path Loss.....	38
4.4.5. Indoor-Indoor Path Loss.....	39
4.5. Link Budget Results.....	39
5. Transmit Power Control.....	42
5.1. Power Control Requirements.....	42
5.2. Power Control State Transition Diagram.....	43
5.3. Power Control SNR Measurement.....	44

A-3

Little Joe Functional Specification

- 6. Media Access Control (MAC).....44
 - 6.1. ViVATO Panel MAC Operation.....44
 - 6.1.1. MAC Modes of Operation.....45
- 7. Multi-Radio Transmit Control.....46
 - 7.1. Multi-MAC Control (MMC).....47
 - 7.2. MMC Overview.....47
 - 7.3. MMC Provisioned Parameter.....48
- 8. Intra-Panel Roaming.....48
 - 8.1. Roaming Requirements.....49
 - 8.2. Beam-Switching Algorithm.....49
 - 8.2.1. Beam-Switching State Transition Diagram.....50
 - 8.3. IAPP (Seamless Roaming).....52
- 9. Channel Assignment.....52
 - 9.1. Channel Assignment Provisioned Parameters.....53
 - 9.2. Channel Assignment Internal Parameters.....53
 - 9.3. Channel Assignment Metrics.....53
 - 9.4. Channel Assignment Algorithm.....55
 - 9.4.1. Channel Assignment Preprocessing.....55
 - 9.4.2. Block-based Channel Assignment Algorithm.....56
 - 9.4.3. Emergency exit.....60
- 10. Downlink Traffic-Shaping.....61
 - 10.1. Traffic-Shaping Requirements.....61
 - 10.2. Traffic-Shaping Architecture.....61
 - 10.3. Traffic-Shaping Functional Description.....63
 - 10.3.1. Leaky Bucket Algorithm.....63
 - 10.3.2. Granularity of Operation.....64
 - 10.3.3. Dynamic Parameter Update.....64
 - 10.3.4. Operating Point Estimation Algorithm.....65
 - 10.4. Provisioned and Internal Parameters.....66
- 11. The Scanning Radio.....67
 - 11.1. Scan Mode.....67
 - 11.2. Roaming.....68
 - 11.3. State Transitions.....68
- 12. References.....69
- Part 2: Software System Architecture.....72
- 13. Introduction.....72
- 14. Software Overview.....74
 - 14.1. Software not covered by this Document.....74
 - 14.2. Software Module Overview.....75
- 15. Control Plane.....76
- 16. Drivers.....76
 - 16.1. Console Driver.....76
 - 16.2. Ethernet Drivers.....77
 - 16.2.1. Secure Management.....77
 - 16.2.2. Backhaul / Daisy-chain for the Next Panel.....78

4

A-4

16.3.	Source for Wireless Drivers	79
16.4.	Searcher and Merlin Interfaces and Functions	79
16.5.	Scheduler / Shaper	81
16.6.	RRM (Radio Resource Management).....	82
16.7.	FLASH	83
This table shows initial estimations of management module resource usage. Need clarification of data items in the table. Some of these items are daemons, running all of the time. Others are one instance per invocation. Cish is dependent on lots of other elements, such as ipchains, brctl and ifconfig.....		
17.	Management Interfaces.....	83
17.1	HTTP	85
17.2	CISH (CLI).....	86
17.3	User Manager	87
17.4	SSHD	87
17.5	Telnet	87
17.6	SNMPD	87
18.	Environmental Control	87
18.1	Temperature and Fan Control	88
19.	Wireless Control.....	88
19.1	Wireless Bridging.....	88
19.2	Centralized Bridging.....	89
19.4	Radio Configuration Manager.....	90
19.4	Inter-card Roaming Manager.....	90
	Mabuhay Enhanced Performance System (MEPS)	90
	Security.....	90
19.4	Authentication Manager.....	90
19.4	Authentication	90
19.4	802.1x / EAP Authentication Mechanism.....	91
19.	Other Control	91
19.4	DHCP Client	91
19.	DHCP Server and Network Address Translation (NAT)	92
19.	Typical Packet Walk (Bridged).....	93
19.	Typical Packet Walk (NAT and DHCP Server)	93
19.	Diagnostics.....	95
19.	Out-of-the-Box Power-up Sequence	95
19.	Utilities.....	95
19.4	FTP Client (for downloading new Images).....	95
19.	FTP Server (for Serving images to other panels).....	95
19.	Data Packet Interfaces	96
19.	Appendix A – PCI / PCI Bridge Support in Linux	97
19.4	Example PCI Based System.....	97
19.4	PCI Address Spaces	97
19.4	PCI Configuration Headers	98
19.4	Layout of the 256 byte PCI configuration header	99
19.	Vendor Identification.....	99

A-5

Little Joe Functional Specification

- 19. Device Identification.....99
- 19. Status.....99
- 19. Command.....99
- 19. Class Code.....99
- 19.4 Base Address Registers.....100
- 19.4 Interrupt Pin.....100
- 19.4 Interrupt Line.....100
- 19. PCI I/O and PCI Memory Addresses.....100
- 19. PCI-ISA Bridges.....100
- 19. PCI-PCI Bridges.....101
- 19. PCI-PCI Bridges: PCI I/O and PCI Memory Windows.....101
- 19. PCI-PCI Bridges - PCI Configuration Cycles and PCI Bus Numbering.....101
 - Type 0 PCI Configuration Cycle Figure:.....101
 - Type 1 PCI Configuration Cycle Figure:.....102
- 19. Linux PCI Initialization.....103
 - 19.4 The Linux Kernel PCI Data Structures.....103
 - 19.4 The PCI Device Driver.....104
- 19. Configuring PCI-PCI Bridges - Assigning PCI Bus Numbers.....105
 - b. PCI I/O and PCI Memory Windows.....106
 - a. PCI-PCI Bridge Numbering: Step 1.....106
 - b. PCI-PCI Bridge Numbering: Step 2.....106
 - d. PCI-PCI Bridge Numbering: Step 4.....107
 - e. PCI BIOS Functions.....108
 - f. PCI Fixup.....108
 - g. Finding Out How Much PCI I/O and PCI Memory Space a Device Needs.....108
 - a. PCI Configuration Header: Base Address Registers.....109
 - b. Allocating PCI I/O and PCI Memory to PCI-PCI Bridges and Devices.....109

A-6

Part 1: Radio and MAC

Overview

This document describes the LittleJoe 802.11b WiFi switch (WS) from a functional point-of-view. The document includes system description, the beamforming functions, radio resource management (RRM), and the media access control (MAC) functionality. Higher layer functionality, network management and other functions are described in separate documents M. Brewer, D. Lohman, et. al. "Software System Architecture Document", Vi.

The functional descriptions in this document provide details of the functions performed by any subsystem. The final architectural design is discussed in separate documents including the control and interface functions between the subsystems.

The product concept is discussed in Introduction.

The high level system architecture is presented in

System Architecture

. The beamforming network which is the core technology of this product is described in The Beamforming Network.

The radio link budgets are presented in Link Budgets

The transmit power control scheme is described in Transmit Power Control. The media access control (MAC) is described in Media Access Control (MAC). The multi-radio transmit control, multi-MAC Control (MMC) and roaming are discussed in

Multi-Radio Transmit Control

Multi-MAC Control, and Intra-Panel Roaming respectively. The channel assignment and traffic shaping algorithms are described in Channel

A-7

Little Joe Functional Specification

Assignment and Dow respectively. The scanning radio functions are described in

The Scanning Radio

1. Introduction

LittleJoe is ViVATO’s first-generation long-range packet switch built according to the 802.11 standard. It seamlessly supports 802.11b clients. The LittleJoe WiFi switch features:

- linear array of 16 antennas providing up to 29 dBi of gain
- Butler-matrix beamforming
- Complementary beamforming
- Multi-MAC controller
- Multi-channel operation
- a high sensitivity RF front end
- Agere 802.11b MAC and baseband processor chips
- Custom logic and software for integration
- Security enhancements

There are two configurations:

- DirectedPacket™ 1 (DP2310) transmits on one channel at a time
- DirectedPacket™ 3 (DP2330) transmits up to 3 channels at a time

These can be fit into two types of antennas/enclosures:

- Indoor half height: 1m wide by 0.5 m high
- Outdoor full height: 1m wide by 1m high

2. System Architecture

In this section we review the ViVATO Packet Switch concept and it’s advantages compared to existing APs in the market. We also describe ViVATO’s high-level system block diagram. The details of the subsystems are then described in different sections within the document.

A-8




17.3 System Concept

Figure 2 A typical deployment for conventional 802.11 APs is shown in *Reference deployment for 802.11 networks*. The example deployment scenario assumes a 10,000 m² coverage area covered by 9 conventional APs each with a cellular coverage radius of 20 m. The assumed coverage cells are circular each with a radius of 20 m. The reference LittleJoe deployment is shown in *Reference LittleJoe indoor deployment*.

One LittleJoe panel (with a range of about 140 m) is used (typically at the corner of a building indoors or mounted on a tall structure outdoors) to provide service to the whole coverage area. In some deployment scenarios, there are regions where LittleJoe cannot provide service due to shadowing or severe scattering. This is represented by the gray shaded circle in *Reference LittleJoe indoor deployment*.

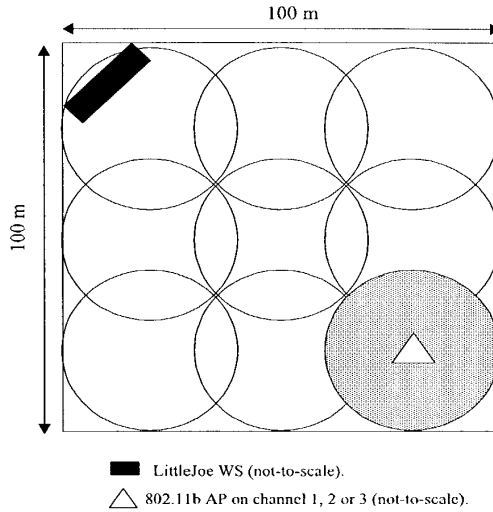
Such areas will be covered using ViVATO APs. In the above example, 9 regular APs have been replaced by one LittleJoe Packet Switch and one ViVATO AP. This reduces the cost of network deployment. In addition, it is possible to provide outdoor coverage by deploying the unit outdoors.



A-9

Little Joe Functional Specification

Figure 1 Reference deployment for 802.11



networks.

Figure 2 Reference LittleJoe indoor deployment.

17.3 Deployment Scenarios

LittleJoe is a long-range WiFi switch and therefore it can support many different applications. Nevertheless, there are five reference deployment scenarios identified:

- Indoor Office
- Indoor Warehouse
- Outdoor Campus
- Outdoor Hotel
- Outdoor ISP

17.1 Typical Mounting Conditions

There are two types of LittleJoe units: half height indoor and full height outdoor.

10

A-10

2...1. Indoor Deployment

A LittleJoe WS is placed inside the building to provide coverage throughout the whole building.

Typically a half height LittleJoe panel is installed on a corner wall inside an office building or a warehouse. The ceiling height is typically between 3 to 4m for the office building and 4 to 8m in a warehouse.

The center of the panel is typically installed at 0.25m below the ceiling. The unit should be mounted away from nearby scatterers or obstructions.

2...2. Outdoor Deployment

Typically one to four co-located full height LittleJoe panels are installed in a desirable location on campus grounds, each panel providing 100° coverage. The typical antenna height is 4 to 20m. The furthest building to LittleJoe installation site is no more than 200m away. The buildings are low-rise (6 stories or less). The scatterers are mostly local to the client and not close to the LittleJoe panel.

For the Outdoor Hotel model, the typical antenna heights would be about 4 - 8m and the building would be no more than 20m away from the panel.

The typical ranges and mounting conditions for the different deployment scenarios are summarized in *The range and antenna dimensions and heights for different de.*

	indoor office	indoor warehouse	outdoor office	outdoor Hotel	outdoor ISP
range	< 150m	200m	300m	100m	2 km
antenna size	1m x 0.5m	either	1m x 1m	1m x 1m	1m x 1m
antenna height	3 to 4m	4 to 8m	4 to 20m	4m to 8m	10 to 50m

A-11




Table 1 *The range and antenna dimensions and heights for different deployment scenarios.*

1.1.1

17.1 Typical Building Structure

The exterior walls are typically concrete or other hard construction material. There are usually many windows and metal on the exterior of the building.

The interior walls are not hard structures. They are typically made of drywall with wood or metal studs. The floors are typically made of concrete. The interior of the building has mostly open offices with soft enclosed cubicles. There may be some dry-wall enclosed offices and conference rooms with or without interior windows.

The unit is built with sufficient link budget to operate with the assumptions above. However, a site may have areas of exceptionally high path loss. These poor coverage areas may be serviced using ViVATO APs¹ connected to the LittleJoe panel through the backbone network or in-band 802.11 signalling.

1.1.2

17.1 Traffic Load Assumptions

There may be more than 200 associated users inside the building. Most users are connected to a wired network and hence do not generate any traffic unless when they require portable connectivity. Nevertheless, they may send probe request frames regularly. The network is expected to perform well with a maximum of 50 active users each with a profile of a “typical” LAN user. The traffic load assumptions for the different deployment scenarios are summarized in *Traffic load assumptions for different deployment scenarios.*

	indoor office	indoor warehouse	outdoor office	outdoor Hotel	outdoor ISP
associated users	> 200	> 200	>200	>200	400

¹A ViVATO AP (also known as the Pollen8) is an open-source AP whose software has been updated to work effectively with the LittleJoe WPS. The details will be described in separate documents.

A-12

active users	20	50	40	10	40
profile	office LAN	short transactions	office LAN	home LAN	home LAN

Table 2 Traffic load assumptions for different deployment scenarios.

17.3 Interference effects

The following are the major sources of interference to the ViVATO network:

- Microwave ovens: create a coverage hole of 5-10 m while they are transmitting. It is recommended that the panel be placed as far away as possible from microwave ovens. It is also advisable to shield the microwave ovens to reduce interference to the network.
- Cordless phones: many operate on the same band and can hence interfere severely with both the panel and client transmissions. It is advised that cordless phones in office and warehouse deployments not be used. For the ISP deployment scenario, the cordless phones will be a significant source of interference and may reduce the networks performance to unacceptable levels.
- 802.11 private LANs: the ViVATO network shares the frequency resources with other 802.11 networks.

17.3 Coverage requirements

Coverage is defined as a packet error rate of 10% or better at a specified data rate. LittleJoe should have an indoor coverage of 85% coverage at 11 Mbits/s and 95% at 5.5 Mbits/sec. ViVATO APs are used to fill in large coverage holes in rare places (no more than 10% of the coverage area) due to severe deployment conditions.

The nominal expected coverage for the different deployment scenarios are summarized in *Coverage for different deployment scenarios..*

	indoor office	indoor warehouse	outdoor office	outdoor Hotel	outdoor ISP
range	150 m	200 m	300 m	100 m	2 km
11 Mb/s coverage	85%	85%	85%	85%	75%

A-13

Little Joe Functional Specification

5.5 Mb/s coverage	95%	95%	95%	95%	80%
coverage holes	5-15%	5-15%	5-15%	5-15%	5-25%

AA



Table 3 Coverage for different deployment scenarios.

1.2

17.3 Throughput and delay requirements

The user experience should be similar to that of a deployment of multiple low-range access points shown in *Reference deployment for 802.11 networks*.

17.1 Performance Metrics

The following performance metrics measure the quality of service provided by LittleJoe compared to conventional 802.11 deployments. These are non-real-time traffic measures G. Anastasi, et. al., "MAC Protocols for Wideband Wireless Local Access: Ev:

- Packet queuing delay: the time elapsed from the time a packet is generated until it is ready to contend for access to the channel
- Packet MAC delay: the time elapsed between the time a packet is ready to contend until the beginning of its successful transmission
- Packet access delay: the sum of the queueing delay and MAC delay
- Packet transmission time: the time between the start of a transmission to its successful completion.
- Packet throughput: the number of bits transmitted in a packet divided by packet transmission time.
- Packet loss rate: the number of packets whose transmission time exceeds TCP time-out.

17.3 Performance limiting issues

The following effects are the most detrimental to LittleJoe performance:

- Nearby large scatterers
- hard structures (hard walls, metal objects, etc.)
- competing 802.11 networks
- high angle spread environment
- microwave ovens
- 2.4 GHz cordless phones

A-15

Little Joe Functional Specification

1.3

17.3 Technology

The technologies enabling the increased coverage for LittleJoe are:

- Phased Arrays
 - vertical antenna gain
 - directional receive and transmit patterns using a Butler matrix
 - taking advantage of point-to-point power limits
- Technologies to enable directional antennas with 802.11
 - complementary beamforming
 - Multi-Radio Transmit Control
 - Multi-MAC control

Additional capacity gain is obtained using:

- multi-channel operation
- seamless roaming
- traffic-shaping

1.4

17.3 System Block Diagram

This section includes the functional block diagram of LittleJoe. The single channel product (DP2310) can operate on a single channel at any time. The 3 channel product call DP2330 can operate on three different channels simultaneously.

Figure 3 The block diagram for The DP2330 product is shown in *The DP2330 high-level block diagram*.

There are 16 antenna elements each with its own RF front-end. The receive port of each circulator is followed by an LNA before the Butler matrix. The circulators ensure that:

- signals arriving at the antenna elements are directed to the receive chain and isolated from the transmit chain
- signals transmitted through the transmit chain are directed to the antenna elements and isolated from the receive chain.

A-16




Figure 3 The DP2330 high-level block diagram.

Each PA is followed by a ceramic bandpass filter. To allow for simultaneous transmit and receive on adjacent channels the LNA can handle strong input signal with no distortion and the post-PA bandpass filters have high adjacent-channel rejection.

The PA output powers can be reduced from their maximum levels in by least 24 dB (4 steps of 6 dB).

A single-channel cost-reduced model (DP2310) may use switches instead of circulators and eliminate the crosspoint switches and filters.

1.5

The Butler Matrix described in detail in

The Butler Matrix

A-17

Little Joe Functional Specification

is a network of passive hybrid power dividers and fixed phase shifters with 16 inputs and 16 outputs. The Butler matrix produces 16 orthogonal beams each pointing in a different direction. However, since the pattern is distorted at the extreme angles, only 13 ports of the Butler matrix on the radios side are used.

Figure 3 Each Butler matrix port at the radio side is connected to a splitter which splits the signal to a WLAN radio and a 13-way switch which is connected to a scanning radio (not shown in *The DP2330 high-level block diagram*).

Figure 3).

1.6 Each WLAN radio is built using the Agere 802.11b chipset². The Agere MAC controller is used and runs AP firmware. See

ViV for more details.

The receive ports of the 13 radios are connected to 13 ports of the receive Butler matrix and the transmit ports are connected to a 13x3 switch that can connect each radio to any of the 3 PAs. The PA outputs are connected to a second 3x13 switch that connects each of the 3 PAs to any of the transmit Butler matrix ports.

The beam patterns and the numbering convention at the Butler Matrix ports with and without windowing are described in Bea.

Figure 4 A scanning receiver in "promiscuous mode" and a "beam" antenna switch are used to obtain signal strength and interference information from stations on different beams and channels as shown in *The scanning radio*.

²The choice of the chipset vendor is subject to change.

A-18




Figure 4 The scanning radio.

3. The Beamforming Network

The increase in the range of the product is obtained by beamforming. The beamforming technology is based on multiple simultaneous reception of signals from different directions through a Butler matrix and transmitting a directional signal through a similar entity.

Additionally, receive windowing is applied to the Butler matrix to help mitigate the near-far problem, and complementary beamforming is applied to reduce the effect of “hidden beam”. The functional details of the beamforming network are discussed in this section.

1.7

17.3 The Butler Matrix

The Butler matrix is a network of passive hybrid power dividers and fixed phase shifters with 16 inputs and 16 outputs. The Butler matrix produces 16 orthogonal beams each pointing in a different direction. The Butler matrix is a theoretically loss-less network that provides maximal gain from the antenna aperture. *Beam and DFT numbering of the Butler matrix ports* (shows the beam and DFT numbering of Butler Matrix’s input and output ports. The beams are numbered 0 to 15 from left to

A-19

Little Joe Functional Specification

right.




Figure 5 Beam and DFT numbering of the Butler matrix ports (top view).

The mathematical description of the Butler matrix response is the same as a Discrete Fourier Transform (DFT).

If the signals at the input of the Butler matrix (radio side) are denoted by x_n where $n=0$ to 15. The output of the Butler matrix (antenna side) is

$$y_k = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-j \frac{2\pi nk}{N}} \quad k = 0 \dots N-1$$

then:

where $N=16$. The amplitude response (far-field transmit beam pattern) of the Butler matrix

$$z(\theta) = \sum_{k=0}^{N-1} y_k e^{jk\pi \cos(\theta)}$$

is:

The far field pattern for a signal at radio port l (x_l) is

$$z_l(\theta) = \sum_{k=0}^{N-1} x_l e^{-j \frac{2\pi kl}{N}} e^{jk\pi \cos(\theta)}$$

then:

As shown in Ed Casas, LittleJoe Beamforming, , for a unity input signal, this may be described by the following

A-20

formula:

$$|z_l(\theta)| = \frac{1}{N} \sqrt{\frac{1 - \cos(\phi)}{1 - \cos\left(\frac{\phi}{N}\right)}} \quad k = 0 \dots N-1 \quad \phi = N\pi \cos(\theta) - 2\pi l$$

Equivalently, the received signal (x_k) on port k on the radio side due to a signal incident from angle θ may be described by the same formula.

1.8

17.3 Ideal Beam Patterns

The *ideal beam patterns of* shows the beam patterns ($20\log(|z_k(\theta)|)$) of a 16 element Butler matrix as a function of the signal's angle of incidence θ .

There are some important observations:

- the beams get wider as we move towards the ports at the outer edges (end-fire)
- the beams at the extreme edges are very wide
- the side lobes are at almost -13 dB
- The cross-beam loss is about 4 dB

Since the beams due to signals transmitted into ports x_7 , x_8 and x_9 are very wide, they are left unused. This means that 13 radios are used. This limits the field of view from about 40° to 140° or to about 100° . Therefore each LittleJoe panel has a field of view of 100° . In other words, to provide 360° coverage, 4 panels are needed.

A-21

Little Joe Functional Specification




Figure 6 The ideal beam patterns of a 16 element Butler matrix.

A radio is attached to each Butler matrix port. Each radio has a boresight direction (the direction with peak signal level). The boresight angles for the Butler matrix ports are shown in *The boresight directions of Butler matrix ports.*

The port numbering in the second column of *The boresight directions of Butler matrix ports.* is based on the DFT index.

To avoid confusion, we have agreed on a single beam numbering convention. The numbering scheme assigns beam numbers starting from the boresight angle of zero and increasing thereafter as shown in column 1 of *The boresight directions of Butler matrix ports.* In this case, beams number 0,1, and 15 are left unused and the beam at 90° boresight is beam number 8.

For the purposes of complementary beamforming, it is required to identify antenna element 0. This is the antenna port that has zero phase shift relative to all the radio ports.

The relatively high sidelobes levels are undesirable for the receiver. 802.11 clients do not use transmit power control. This results in a wide range of received signal levels at the switch. In both indoor and outdoor deployment the ratio of propagation distances might easily exceed a factor of 10 (near-far effect). This results in path loss differences of more than 30 to 40 dB. Therefore, the sidelobes need to be reduced. This is accomplished by receive windowing described in the next section.

A-22

Butler matrix port		boresight angles
Beam Numbering	DFT Numbering	(degrees)
0	8	0.00, or 180.00
1	7	28.96
2	6	41.41
3	5	51.32
4	4	60.00
5	3	67.98
6	2	75.52
7	1	82.82
8	0	90.00
9	15	97.18
10	14	104.48
11	13	112.02
12	12	120.00
13	11	128.68
14	10	138.59
15	9	151.05

Table 4 *The boresight directions of Butler matrix ports.*

17.3 Receive Windowing

Tapering the illumination of the array (“windowing”) reduces the sidelobe levels. The resulting trade-offs are:

- an increase in the main lobe width
- a decrease in received SNR

Therefore, the decision to use windowing may depend on the environment. For instance, in a situation where an outdoor LittleJoe unit services clients in only one building, the received power dynamic range may be relatively small and hence higher sidelobe levels may not be

A-23

Little Joe Functional Specification

harmful. On the other hand, a decrease in received signal power may not be acceptable due to range requirements. Thus the selection of a window function will depend on the expected dynamic range of the received signals (the severity of the near-far problem).

One of two windows are selected at provisioning.

- Rectangular window (no windowing)
- Hamming window (default value).

The receive window type is configured by the network element management software. The default setting for the windows is the Hamming window described below.

17.1 The Hamming Window:

Figure 7 A popular window with relatively low side lobes is the Hamming window shown in *The Hamming window*.

Figure 7 The Hamming window reduces the peak sidelobe level from -13 to -41 dB but doubles the main lobe null-to-null beamwidth and increases the 3 dB beamwidth from about 6° to about 10°. The Hamming window shown in *The Hamming window*.

is described by the following formula:

$$w(n) = 0.54 - 0.46 \cos\left(\frac{2\pi n}{N-1}\right) \quad n = 0 \dots 15$$

where $N=16$. The window has a bell shape. It reduces the signal power received from the outer elements of the antenna

A-24

array.




Figure 7 The Hamming window.

The window is symmetric and is implemented using a network of attenuators which attenuate the signal from each element by the appropriate Hamming window coefficient shown in *The Hamming window coefficients*.

elements	0,15	1,14	2,13	3,12	4,11	5,10	6,9	7,8
gain	0.08	0.12	0.23	0.4	0.59	0.77	0.91	1.00
gain (dB)	-21.9	-18.4	-12.7	-8.0	-4.6	-2.3	-0.8	0.0

A-25

Little Joe Functional Specification


Table 5 *The Hamming window coefficients.*

To normalize the effect of the window coefficients to achieve the same SNR, the receiver RF circuit should increase the voltage gain to each antenna element by the factor of:

$$G_w = \frac{N}{\sum_{n=0}^{N-1} w(n)^2} = 1.64 \quad N = 16$$

This translates to an adjusted power gain of about 4.3 dB.

Figure 8 The beams at the radio side of the Butler matrix with a Hamming window are shown in *The Butler matrix receive beam shapes with a Hamming window.*



26

A-26



Figure 8 *The Butler matrix receive beam shapes with a Hamming window.*

Figure 8 As shown in *The Butler matrix receive beam shapes with a Hamming window.*

, the main beams of the Butler matrix are wider (3 dB beamwidth of about 10°) but the sidelobes have been suppressed to -41 dB. There is also approximately 1.5 dB loss in SNR due to receive windowing. The Butler matrix ports 7,8 and 9 are even worse than the non-windowed case and deemed unusable hence limiting the field of view from approximately 40 to 140 degrees.

The cross-beam loss (or the pointing loss) is about 1.5 dB.

1.9

17.3 Complementary Transmit Beamforming

As shown in *The ideal beam patterns of*, the transmit beams have very deep nulls in certain directions and the lowest sidelobe levels are around 14 dB down from the main lobe's peak.

With complementary beamforming, we intend to reduce the effect of the nulls and increase the sidelobe levels without a severe power penalty to the main beam. This is done to reduce the effect of the "hidden beam"³.

Figure 9 The complementary beam is formed by increasing the gain at the antenna element 0. This is shown in *g*.

³ The media access technique in 802.11 is Carrier Sense Multiple Access (CSMA). Forming directional transmit beams has the side effect of hiding the transmitted energy from some clients in the network; i.e, negatively impacting the carrier sense mechanism in the network. A client measures the energy transmitted from APs and other clients. If it cannot detect the presence of other transmissions, it attempts to access to the medium. Therefore, when directional beams are used, many clients detect the medium as idle when in fact it is busy. This has an effect on the performance of the network. We call this phenomenon the "hidden beam" problem.

A-27

Little Joe Functional Specification

using the DFT numbering convention. This is the only port whose output is the addition (with no phase) of all the input ports ($y_0 = x_0 + x_1 + \dots + x_{15}$).




Figure 9 Block diagram of complementary beamforming.

Mathematically, this may be described as:

$$y_i = \begin{cases} \gamma y_i & i = 0 \\ y_i & \text{otherwise} \end{cases} \quad \gamma \geq 1 \quad 0 \leq i \leq N - 1$$

To ensure the same output power as with no complementary beamforming the output voltage on all the ports should be adjusted

$$G_s = \sqrt{\frac{N}{\gamma^2 + N - 1}}$$

by the scaling factor:

It may be shown that the power penalty for the main beam

$$\Delta P = \frac{(\gamma + N - 1)^2}{N(\gamma^2 + N - 1)}$$

is:

$$\Delta P_{dB} = 10 \log \left(\frac{(\gamma + N - 1)^2}{N(\gamma^2 + N - 1)} \right)$$

or in dB:

For instance, for a 16 element array, if $\gamma=3.5$, the power loss is about 1 dB. It is desirable to have γ as a parameter that may be changed at provisioning.

The Butler matrix output due to a shows the shape of the transmit beam due to a signal at port 0 of the Butler matrix with and without complementary beamforming. The output with complementary beamforming has higher sidelobes in all directions and removes all the deep nulls except for the nulls on the main beam. The main beam's peak power is about 1 dB lower

A-28

than that without complementary beamforming.




Figure 10 The Butler matrix output due to a signal at x_0 (beam 8) with complementary beamforming ($\gamma=3.5$).

4. Link Budgets


The link budgets are prepared to provide a basis for making design decisions that might affect the coverage area of LittleJoe. Additionally, they provide guidance to potential customers about the improvement in coverage area they can expect from a LittleJoe panel compared to a conventional AP. In this section, we provide a number of models for indoor and outdoor applications and provide best-case, worst-case and nominal range estimates for the different deployment scenarios described in Deployment Scenarios.

Throughout this section the term panel refers to the LittleJoe packet switch and client refers to a standard IEEE 802.11 WLAN client card.

1.10

17.3 Link Power Budget

A-29



Little Joe Functional Specification

A link budget is useful for evaluating design decisions. The link budget predicts the operating margin, which is the amount by which the received signal level exceeds the level required to achieve a sufficiently low error rate⁴ for a large-enough fraction of users.

The basic link equation

$$P_R = P_T + G_T + G_R - L$$

is:

where the variables, described in *Link budget terms* are in dBm or dB and L represents L_I or L_O depending on the panel location.

The operating margin M is the amount by which P_R exceeds the receiver sensitivity

$$M = P_R - S_R$$

S_R :

This operating margin is calculated separately for the downlink (panel to client) and uplink (client to panel). Both margins must be positive for a client to obtain service. Since the path loss is time and location-dependent (due to fading), the fraction of users within the coverage area that have positive operating margins will vary.

The link budgets use statistical models. Their purpose is to examine the sensitivity of the system performance to design changes. It is not designed to predict performance in a specific installation. Other techniques that make use of site-specific data are used for that purpose.

term	description
P_T	transmitter power
G_T	transmitter antenna gain
G_R	receiver antenna gain
L_O	mean path loss for outdoor-indoor case
L_I	mean path loss for indoor-indoor case
M	margins for shadow and Rayleigh fading
P_R	received power
S_R	receiver sensitivity

⁴In our case, a frame error rate (FER) of less than 8×10^{-2} for 1024-byte frames.

A-30

Table 6 *Link budget terms*

In

Link Budget Parameters

we describe each of the above parameters and identify known values. In Characteristics of Conventional 802.11b Equipment

we present the radio characteristics of typical 802.11b equipment. In Path Loss Models

1.11 we suggest models for the path loss and finally in

Lin, we compute link margins and estimate the range increase of LittleJoe compared to the theoretical limits of existing APs.

1.12

17.3 Link Budget Parameters

17.1 Panel Transmit Power: FCC EIRP Power Limits

The FCC limits transmitter power for in the unlicensed 2400 to 2483.5 MHz band to 30 dBm (1 W). In addition, the EIRP for point-to-multipoint devices is limited to 36 dBm. The EIRP for point-to-point devices is not limited, but must be reduced by 1 dB for every 3 dB of antenna gain above 6 dBi.

For example, a point-to-point system using an antenna with again of 29 dBi would be restricted to an EIRP of:

$$30 \text{ dBm} + 29 \text{ dB} - (29-6)/3 \text{ dB} = 51.3 \text{ dBm}.$$

A-31

Little Joe Functional Specification

A similar reduction in transmit power is not required for point-to-point systems in the 5725--5850 MHz band. In other words, for those client cards operating in this band (such as 802.11a clients), there is only a transmitter power limit of 30 dBm and no EIRP limit.




Figure 11 Diagram showing EIRP limits as per CFR 47, Part 15.247(b). Point-to-multipoint (P-MP) systems must operate below 36 dBm EIRP regardless of antenna gain and may not exceed 30 dBm transmitter power. Point-to-point (P-P) systems may increase EIRP above 36 dBm by 2 dB for each 3 dB increase in antenna gain.

17.3 Characteristics of Conventional 802.11b Equipment

Specifications for transmit power, receiver sensitivity and claimed indoor range were obtained from the data sheets for various manufacturer's 802.11b WLAN APs and client cards. The results are given in *Transmit power levels, receiver sensitivities and claimed indoor ra.*

Sensitivities for Orinoco and Prism III measured at 10^{-5} BER which is 8×10^{-2} FER for 1024-byte frame. Range estimates are at 11 Mb/s. Orinoco range estimate is for a *closed* environment. Apple range estimate is for *typical use*. Ericsson range quoted for an external antenna and an *office environment*.

Make/Model	transmit power (dBm)	sensitivity 1 Mb/s (dBm)	sensitivity 11 Mb/s (dBm)	indoor range (m)
Orinoco World PC card	15	-94	-82	25
Intersil Prism III	?	-91	-84	37
Apple Airport	15	?	?	45
Nokia C110/C111 client & A032 AP	15	?	-84	20-100
Ericsson PC Card PA11	20	-90	-84	75

A-32

Intel Pro/Wireless 2011B AP	18-20	-90	-83	30
Intel Pro/Wireless 2011B client	14-18	-87	-81	30
Cisco 350	20	-94	-85	40

A-33






Table 7 *Transmit power levels, receiver sensitivities and claimed indoor range for some conventional WLAN APs and client cards.*

17.1 Panel Antenna Gain

The actual panel antenna array gain is highly dependent on the implementation. However, an estimate can be obtained from its physical size and the antenna type. The antenna gain is related to its effective aperture

by:

$$G_R = \frac{4\pi A_{eff}}{\lambda^2}$$

Assuming A_{eff} is equal to the array's cross-sectional area (typical for a linear array with a reflector):

$$G_R = \frac{4\pi wh}{\lambda^2}$$

where w is the width of the antenna, h is the height of the antenna and λ is the wavelength.

For the indoor unit: $w=8\lambda$ and $h=4\lambda$ and hence:

$$G_R = \frac{4 \cdot \pi \cdot 8\lambda \cdot 4\lambda}{\lambda^2} = 128\pi = 26 \text{ dBi}$$

For the outdoor unit: $w=8\lambda$ and $h=8\lambda$ and hence:

$$G_R = \frac{4 \cdot \pi \cdot 8\lambda \cdot 8\lambda}{\lambda^2} = 256\pi = 29.1 \text{ dBi}$$

Increasing the frequency to 5725 MHz increases the gain of a equal-sized antenna by $(5725/2400)^2$ or about 7.5 dB (but also increases free-space loss by an equal amount).

1.12.1

17.1 Gain Reduction due to Scattering

Computations of antenna gain and sidelobe level assume a single plane wave front arriving at the antenna. In typical WLAN installations the signal will arrive via multipath scattering and there will be many angles of arrival. This will result in a reduction in gain because signals arriving from directions other than a beam's boresight angle do not sum coherently.

A-34

The exact degree of the gain reduction depends on the antenna pattern and the angle of arrival distribution. The resulting reduction in gain can be significant. For example, in L. Greenstein and V. Ereg, "Gain Reductions Due to Scatter on Wireless Paths" median gain reductions of 3 to 5 dB were observed for a 37 degree half-power beamwidth antenna at 3m heights in a suburban scattering environment. Since our antenna has a narrower beamwidth and the indoor environment exhibits more severe scattering, the gain reduction may be significantly larger.

On the other hand, M. J. Gans, R. A. Valenzuela, J. H. Winters, and M.J. Carloni, "High Data Rate In" reports that in most cases over half of the energy is contained in a single narrow angle of arrival. Similar results showing several widely-separated but discrete angles of arrival are visible in the data reported in G. German, Q. Spencer, L. Swindlehurst, and R. Valenzuela, "Wireless Indoor Ch" and J. G. Wang, A. Mohan, and T. Aubrey, "Angles-of-Arrival of Multipath Signals in."

The "separability" of the paths in M. J. Gans, R. A. Valenzuela, J. H. Winters, and M.J. Carloni, "High Data Rate In" might be accounted for by a factor of 8 difference in frequency (19GHz/2.4 GHz, which is a difference of 64 in far-field distance) and the brick interior wall construction of the building tested that increases the contribution of diffraction (single-direction) as compared to transmission and scattering (many-direction) effects.

Another effect that will reduce the gain of the antenna when there are nearby scatterers is that the wave fronts are not well approximated by plane waves and this results in an additional loss of gain and increase in sidelobe ratio.

In the link budgets we have entered some arbitrary values for the gain reduction. These values are to be refined through propagation measurements.

1.12.2

17.1 Client Antenna Gain

Amongst those vendors considered, only Ericsson provides a gain specification for their PC client card antenna (0 dBi). This number agrees with measurements reported by others for other client cards. Omnidirectional antennas used by enterprise APs have higher gains. For example, the Nokia C950 has a gain of 2.5 dBi and the Cisco AIR-ANT3213 has a gain of 5.2 dBi.

17.1 Client and Panel Noise Figures

The Intersil Prism II receiver IC specifies a noise figure of about 2 dB. Losses in the antenna switch and bandpass filter will increase this number, perhaps to 4 dB. The Orinoco "Ruby" receiver IC has a noise figure of about 5 dB.

A-35

Little Joe Functional Specification

The LittleJoe panel uses a 1 dB NF LNA and a circulator with 1 dB loss to achieve a NF of about 2 dB.

A client card chipset with an external LNA is used in LittleJoe. The difference between typical client and predicted LittleJoe panel noise figures must be included in the link budget. This difference is also used when comparing the performance of the LittleJoe panel with conventional APs.

17.1 Effect of Shadow Fading

17.3 The path loss models described Path Loss Models

estimate the median path loss for a given distance. However, different locations with the same path distance will have different path losses. These variations have been found to be normally distributed when the path loss is expressed in dB. The standard deviation of the path loss depends on the scattering environment but typical values for office environments are 3 to 6 dB.

We cannot increase the transmitter power to compensate for shadow fading since the system is already operating at maximum transmit power levels. Instead, the shadow fading reduces the fraction of the coverage area that can be serviced.

We will assume the coverage area is a circle or a “wedge” of a circle so that we can use the equations derived for circular cells W.C. jakes, ed.,

$$a = \frac{-\gamma}{\sigma\sqrt{2}}$$

$$b = \frac{10n \log e}{\sigma\sqrt{2}}$$


$$F(\gamma) = \frac{1}{2} \left(1 - \operatorname{erf}(a) + e^{\left(\frac{1-2ab}{b^2}\right)} \left(1 - \operatorname{erf}\left(\frac{1-ab}{b}\right) \right) \right)$$

:

where $\gamma = M$ is the link budget margin at the coverage boundary (dB), σ is the standard deviation of fading (dB), and n is the path loss exponent. For other coverage region shapes, a different expression must be derived or the value computed through numerical integration.

The percentage coverage computed above is an average over the whole coverage area and it may include large coverage “holes.”

A-36



Since shadow fading is caused by objects such as walls, bookcases, doors, etc. we should expect the dimensions of the “holes” to be approximately the same as the dimensions of the shadowing objects. This is unlike the Rayleigh fading where the fades have dimensions on the order of the wavelength.

1.12.3

17.1 Effect of Rayleigh Fading

A margin is usually included in a link budget to counter the effect of multipath fading. For NLOS propagation this fading is Rayleigh-distributed. The probability that a Rayleigh distributed random variable r will be R dB below the mean can be approximated

$$P(r < R) = 10^{-R/10}$$

by:

for $R < 01$. For example, the signal will be 10 dB below the mean about 10% of the time and 20 dB below the mean about 1% of the time.

For typical indoor scatterers, the duration of the Rayleigh fades (tens or hundreds of milliseconds) is slow relative to the frame duration (less than about 10 milliseconds).


The fading on antennas separated by a significant fraction of one wavelength is weakly correlated. Many clients use switching diversity to combat Rayleigh fading. The client’s receiver switches between two antennas until it finds a signal that is sufficiently strong. This squares the probability of fading to $10^{-2R/10}$.

Rayleigh fading affects data throughput rather than coverage. Not including any fade margin would result in the channel being unavailable about 40% of the time without diversity and about 16% of the time with two-antenna diversity. Allowing a 10 dB Rayleigh fading link margin would mean the signal was faded about 10% of the time without diversity and about 1% of the time with two-antenna switched diversity. Allowing a 5 dB fade margin, would mean 10% probability of fading with two-antenna switched diversity.

The effect of Rayleigh fading on data throughput is difficult to compute because of complex interactions between frame loss and contention-control mechanisms in the 802.11 MAC and congestion-control mechanisms in TCP/IP.

As far as the link budget is concerned, we assume a 5 dB fade margin for downlink transmissions. This assumes that the clients have switched diversity. We also assume a 5 dB fade margin for the uplink since the multiple antenna elements will provide some diversity gain against fading.

A-37



Little Joe Functional Specification

17.3 Path Loss Models

This section describes the models used to predict the outdoor-indoor and indoor-indoor path loss. Descriptions and experimental validation of these path loss models can be found in J. Kivinen, X. Zhao, and P. Vainikainen, "Empirical Characterization of Wi.

17.1 Propagation in Free Space (LOS)

Loss in free space is given

$$L_{FS}(d) = 20 \log \left(\frac{4\pi d}{\lambda} \right)$$

by:

in the far field ($d > 2D^2/\lambda$, $d \gg D$ and $d \gg \lambda$ where D is the largest dimension of the antenna).

1.12.4

17.1 Propagation by Diffraction (NLOS)

For NLOS (non line-of-sight) paths, propagation is mainly by diffraction and a power-law path loss formula is a good model:

$$L(d) = L_{FS}(d_0) + 10n \log \left(\frac{d}{d_0} \right)$$

where d_0 is a reference LOS distance (free space break point) and n is the path loss exponent a value that depends on the geometry of the paths. For indoor NLOS paths n is typically between 3 and 4.

1.12.5

17.1 Propagation by Transmission (OBS)

When propagation is mainly by transmission through walls for floors, a simple model is to modify the free-space path loss with a wall or floor attenuation factor for each penetrated wall and/or

floor:

$$L_{OBS}(d) = L_{FS}(d) + p_{wall}W_{wall} + p_{floor}W_{floor}$$

where p_{wall} and p_{floor} are the number of penetrated walls and floors respectively and W_{wall} and W_{floor} are attenuation constants that depend on the wall and floor construction materials.

38

A-38

For the same-floor propagation through several walls, this model can be simplified by substituting a constant attenuation per unit distance:

$$L_{OBS}(d) = L_{FS}(d) + \alpha d$$

where α is the attenuation per meter. A typical value for α in buildings with hard walls is 0.6 dB/m (resulting, for example, from W_{wall} of 4 dB per wall and a wall every 6 or 7 meters).

1.12.6

17.1 Outdoor-Indoor Path Loss

A model for outdoor-to-indoor propagation J.E. Berg, "Building Penetration Loss along Urban Street Microcells," in E. Damosso and L. Correia, eds., combines the OBS and LOS models and a correction for the angle of incidence:

$$L_O(d) = L_{FS}(S+d) + WG_e \left(1 - \frac{D^2}{S^2}\right) + \max(\Gamma_1 + \Gamma_2)$$

where S , d , and D are in meters (see *Model for COST 231 building penetration loss model*. From E. Damosso and L. Correia, eds.), W_e is a constant related to the external wall construction (about 7dB for concrete walls with unshielded windows) and WG_e is a similar constant for shallow angles of incidence (20 dB).




Figure 12 Model for COST 231 building penetration loss model. From E. Damosso and L. Correia, eds., .

A-39

Little Joe Functional Specification

The loss inside the building is modeled

using:

$$\Gamma_1 = W_i p$$

or

$$\Gamma_2 = \alpha(d-2)\left(1-\frac{D}{s}\right)^2$$

where W_i is the additional loss per interior wall (4 dB per wall), p is the number of walls passed, and α is the attenuation constant, about 0.6 dB/m.

If there is significant building penetration from several directions, the signal levels from each direction should be computed separately and summed. Note that the building penetration loss W_e may be 10 to 20 dB higher for buildings with low-emissivity (energy-efficient, "low-E") windows that have metallic coatings.

1.12.7

17.1 Indoor-Indoor Path Loss

The power-law NLOS model described in Section 4.4.2 is usually used to model indoor-to-indoor path loss.

1.13

17.3 Link Budget Results

[5] A series of Link Budget spread sheets have been prepared and are available. It is important to note that depending on assumptions about the deployment scenario, the range estimates may vary widely. Hence, in this section, we calculate the link margin gain of the LittleJoe panel compared to conventional APs (transmitting at 30 dBm) and report the resulting increase in range in *Estimated range increase for indoor and outdoor deployments co.* For results based on different deployment environments, see Siavash Alamouti, LittleJoe Link Budget Spread Sheets, May 10, 2002.

, Ed Casas, LittleJoe Beamforming, .

	indoor unit		outdoor	
	downlink	uplink	downlink	uplink
Theoretical transmit antenna gain	26.0 dBi	0 dBi	29.1 dBi	0 dBi
Required backoff	6.7 dB	0 dB	7.7 dB	0 dB

A-90

Transmit gain over 30 dBm	19.3 dB	0 dB	21.4 dB	0 dB
Actual transmit gain over 30 dBm	18 dB	0 dB	20.0 dB	0 dB
Receive antenna gain	0 dBi	24 dBi	0	27.0 dBi
window and/or pointing loss	-2 dB	- 4 dB	-2 dB	- 4 dB
NF improvement	0 dB	2 dB	0 dB	2 dB
link margin over regular APs	16.0 dB	22.0 dB	18.0 dB	25.0 dB
best-case range increase	5 times	12.6 times	7.9 times	18.0 times
nominal range increase	2.5 times	4.2 times	4.0 times	6.9 times
worst-case range increase	2.2 times	3.5 times	2.8 times	4.2 times

A-41

Table 8 *Estimated range increase for indoor and outdoor deployments compared to omnidirectional APS.*

The FCC limits transmitter power in the unlicensed 2400 to 2483.5 band to 1 Watt (or 30 dBm). However, the transmit power has to be backed off by 1 dB for every 3 dBi of antenna gain over 6 dBi. The approximate⁵ effective transmit gain for indoor and outdoor deployments are calculated in *Estimated range increase for indoor and outdoor deployments co.*

The added coverage compared to regular APs is a function of the link margin above regular APs shown in *Estimated range increase for indoor and outdoor deployments co.* The best-case range increase is based on a path-loss exponent of 2, the worst-cases is based on path loss exponent of 4 and the nominal path loss exponent was set to 3.5 indoors and 3 outdoors. Note that the results in *Estimated range increase for indoor and outdoor deployments co* indicate that the range increase is more significant on the uplink. This is due to the full antenna gain on the uplink (no transmit back-off as in downlink) and the NF improvement in the LittleJoe receiver. Some of the gain is offset by the assumption that a 4 dB windowing and pointing loss is incurred on the uplink. This is calculated based on windowing SNR loss of 1.5 dB SNR on the uplink and a pointing loss due to lack of fine steering on the downlink and. uplink.

In *Estimated range increase for indoor and outdoor deployments co*, we have used the actual transmit and receive antenna gains based on measured gains in our laboratory. The actual gain for the half height antenna was measured to be around 24 dBi (compared to theoretical value of 26 dBi), and for the full height antenna at around 27 dBi (compared to 29.1 dBi).

[5] In practice, the 802.11 clients transmit significantly lower power than the FCC limits (15dBm to 17 dBm). Therefore, there would be significant disadvantage in the uplink. The comparison in *Estimated range increase for indoor and outdoor deployments co* is based on the limits of our technology compared to omni-directional equipment. For actual range estimates with existing clients please refer to Siavash Alamouti, LittleJoe Link Budget Spread Sheets, May 10, 2002.

5. Transmit Power Control

LittleJoe does not perform transmit power control on a per-client basis. This is to be compliant with the MAC channel sense mechanism. Per-client power control may hide the transmission to one client from the other clients in the network. This would create a hidden beam problem. Therefore, the power control scheme for LittleJoe is applied to all the radios and is independent of the beam, channel, destination (client) or frame type. The power is reduced if and only if all

⁵The effective antenna gain and hence the resulting back-off depend on the implementation and may be slightly different.

A-42

the associated clients have an SNR of more than 30 dB. Otherwise, the maximum power is applied.

1.14

17.3 Power Control Requirements

The term “maximum level” used in the following refers to the maximum transmit power per antenna that is allowed by FCC.

The transmit power level:

- shall never exceed the maximum level
- shall be accurate to within 1.0 dB with probability 99.999%
- shall be set with a resolution of 0.5 dB or better⁶
- shall be set to the maximum level if the SNR of any associated client is less than 30dB
- shall be reduced by 1 dB for every 1 dB that the minimum SNR of any associated client⁷ exceeds 36 dB
- shall be able to be reduced by at least 24 dB from the maximum level at 6 dB steps (to within 2.0 dB accuracy)
- is raised to the maximum level for a duration of 10 seconds immediately⁸ following reception of a probe request frame

The relationship between (downlink) transmit power level and minimum (uplink) SNR is shown in *Transmit power as a function of me*. The figure shows that the transmit power is set according to the power requirements of the client with minimum SNR.

⁶The accuracy and resolution of power control should be included in the link budget as a 1.5 dB loss for downlink transmissions.

⁷Stations from which no frames have been received in the last 24 hours are expected to be disassociated

⁸Since the transmit power is controlled by the host, there may be significant delays in changing the power level. One or more of the initial probe responses may be transmitted at the initial (lower) power levels. If the client is unable to receive any of the probe responses (e.g. because they were transmitted at the initial power level), it will scan again and will then receive the probe responses transmitted at the higher power level. It is possible that in some cases the client will receive some responses at the initial level and some at the higher level and choose the wrong beam as a result. This situation will be handled by the roaming algorithms.

A-43

Little Joe Functional Specification




Figure 13 Transmit power as a function of measured uplink SNR.

17.3 Power Control State Transition Diagram

Figure 14 The power control state machine for each panel is shown in *The power control state machine*.




Figure 14 *The power control state machine*.

The state machine can be in one of two states:

- **Full Power:** the transmit power is set to the maximum level. This state is required to allow all clients to receive probe responses and beacons transmitted from the panel.
- **Adjusted Power:** the transmit power is adjusted using the SNR of the client with the lowest SNR (see the rules above). This is the normal operating mode and provides some power control with enough margin so as not to impact performance.

The rules governing the state transmissions are:

A-49

- *timeout*: the time allowed for associations to complete (e.g. 10 seconds) has passed since the most recent probe request
- *probe request*: a probe request frame is received on any beam

Important Note: *this design requires APs to indicate received probe request frames to the host (even during normal operation). If this is not possible, full transmit power should be applied at all times.*

1.15

17.3 Power Control SNR Measurement

The client SNR estimate is a time-averaged value. Let the SNR measurement for the last frame be denoted by SNR_{new} . The average SNR is calculated

$$\overline{SNR} = 0.9\overline{SNR} + 0.1SNR_{new}$$

as:

6. Media Access Control (MAC)

The LittleJoe MAC subsystem is composed of a host interface, thirteen MAC controller chips, a Multi-MAC controller (MMC) and thirteen baseband (modems) attached to thirteen radios. Each of the MAC controller chips are IEEE 802.11b based using standard 802.11 DCF (CSMA/CA) as the mechanism for scheduling transmissions. Each MAC controller chip independently monitors a different beam to determine the next opportunity for transmission. Once one of the MAC controller chips transmit, the other MAC controller chips on the same channel sense the transmission and wait their turn to transmit. This approach guarantees that each MAC controller/radio has an equal chance to transmit. Additionally, the MMC helps avoid collisions of downlink transmissions with ongoing uplink receptions.

1.16

17.3 ViVATO Panel MAC Operation

1.17

As described in , the LittleJoe architecture has one WLAN radio on each Butler Matrix port. Each radio operates with its own independent MAC controller.

A-45

Little Joe Functional Specification

All of the radios listen simultaneously but only one of them can transmit at a time. Each radios's MAC protocol enforces the "one transmitter at a time" rule through the use of it's CCA circuitry. Any transmission from any radio on any beam will interfere with and cause the loss of any frames currently being received on other beams on the same channel. However, since the radios on different beams can hear each other's transmissions, there should be no collisions after the initial downlink DATA or CTS frame (as a result of MACS on other beams setting their NAV timers). However, an MMC is added to avoid collision of downlink and ongoing uplink packets. The details of the MMC function is described in

Mul.

The LittleJoe panel supports unmodified 802.11b clients. No special MAC software is required at the client. Each beam appears as a different AP and a client associates with whichever beam it thinks is providing the best coverage. However, since the clients are portable and the wireless environment may change, the initially selected radio may not indefinitely be the best radio and hence the client may have to roam to another radio. To support a client transition from one BSS (radio) to another BSS (radio) within the same ESS, there needs to be a mechanism to support 802.11 MAC layer mobility. Inter-access point communication (IAPP) specified in the IEEE 802.11f draft supports that mechanism. The details of the roaming function is described in

Intr.

Since different APs in the panel may be assigned to different channels, a channel assignment algorithm is used as described in

46

A-46

Cha.

A traffic-shaping functionality described in Section 10 is implemented to limit the downlink load to ensure system stability and fairness between uplink and downlink traffic loads.

1.17.1

17.1 MAC Modes of Operation

Currently the IEEE 802.11 standard provides two modes of operations which are all part of the basic service set (BSS), which is defined as a group of clients that communicate with each other. The two modes are Independent basic service set (IBSS), and infrastructure basic service set (BSS).

The Little Joe panel will only support the infrastructure BSS mode.

Infrastructure BSS mode of operation is distinguished by the use of an access point (AP). The AP is used for all communications including communication between clients in the same service set.

If one client transfers a frame to the second client, it must take two hops. First, the originating client transfers a frame to the AP. Second, the AP transfers the frame to the client. With all communications relayed through the access point, the service area is defined by the coverage area of the AP itself.

In infrastructure BSS mode, clients must associate with an AP before communication can begin. The client always initiates the association process and the AP may choose to grant or deny service. Associations are exclusive in that a client can only be associated to one AP at a time.

LittleJoe has 13 independent 802.11 APs within one panel. We hereby refer to the APs within the panel as radios. The APs need to be coordinated to form a single entity. This is supported by the Extended Service Set (ESS) in 802.11. For a LittleJoe panel to properly cover a 100 degree area with portability support, there is a requirement for the union of multiple BSSs to be configured to be part of the same ESS. This requires that the 13 radios operate in concert to allow the outside world to use a single MAC address to talk to a client somewhere within the ESS. To support this functionality, the LittleJoe backbone must act as a single link-layer domain.

Clients within the same ESS may communicate with each other, even though these clients may be located on separate beams. For clients in the ESS to communicate with each other, the wireless medium must act as a single layer 2 connection. Radios act as bridges, so direct

A-47

Little Joe Functional Specification

communication between clients in an ESS requires that the backbone network also be a layer 2 connection.

The ESS may also function between multiple LittleJoe panels. This requires that, in addition to the radios inside the panel working as a single ESS, the Ethernet backbone may also be used as the backbone for the ESS.

7. Multi-Radio Transmit Control

The panel has 13 independent radios that are coordinated through the CSMA protocol. The protocol tries to ensure that only one transmitter per channel operates at a given time. However, there still remains a small chance of a transmit collision. Therefore, to avoid exceeding FCC transmit power limits and slight degradation in performance, a transmit control function is provided. The control function ensures that:

- only one transmitter transmits at a given time on any given channel (first-come, first-serve)
- only complete frames are transmitted (i.e.; frames that begin while another transmission is in progress are discarded)
- if not possible to determine which of several transmissions occurred first, then select any of the them arbitrarily.

17.3 Multi-MAC Control (MMC)

LittleJoe uses 13 radios. Without any coordination between these radios, transmitting a downlink packet on one radio would destroy any uplink packets being received simultaneously on another radio. This effect is referred to as “data suicide”. Allowing each radio to operate independently causes serious performance issues LJ MAC problems: Using multiple co-located APs with smart antennas, A Tu. MMC controls the transmissions of each radio to prevent data suicide. More details on MMC are discussed in Little Joe Multi-MAC Controller, (aka CCA Glue Logic), Presentation/Design re.

1.18

17.3 MMC Overview

Figure 15 The MMC operation is shown in *LittleJoe Multi-MAC Operation*.

48

A-98

There are 13 radios with independent CCA input and outputs. For clarity, we call these *busy_out* and *busy_in*. The *busy_out* indicates the state of the CCA output as detected by the baseband processor and *busy_in* indicates the state of the CCA input to the MAC processor. The *busy_out* signals are fed into the MMC. Based on the values for *busy_out* signals, the channel assignment vector, and the *busy_out* enable vector, the MMC sets the *busy_in* signals for every radio.

The MMC function is as follows:

Define a signal, *global_busy*, for each channel. *global_busy* is active if:

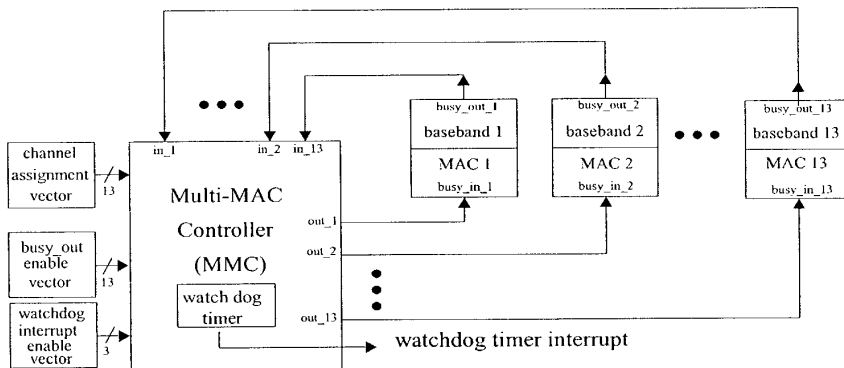
- *busy_out* for any radios operating on that channel indicates busy excluding those radios whose *busy_out* enable signal is not set.

Set the MMC output for a given radio to busy if:

- *busy_out* for that radio indicates busy, or
- *global_busy* for this radio's channel is active.

Each of three watchdog timers monitors the period that *global_busy* has been active for that channel. If *global_busy* has been active for more than *aBusyActiveThreshold*, and that interrupt is not disabled (indicated by the watchdog interrupt enable vector), an interrupt is generated to the host.

The *busy-out* enable vector is set taking into consideration the effect of interference and overlapping BSS. It ensures that external channel activity does not quiet all the radios in the panel.



A-49

Little Joe Functional Specification

Figure 15 *LittleJoe Multi-MAC Operation.*

17.3 MMC Provisioned Parameter

The only provisioned parameter for the MMC is:

- `aBusyActiveThreshold`: the threshold of *global-busy* active status beyond which an interrupt is generated to the host.

8. Intra-Panel Roaming

The LittleJoe switch contains 13 radios operating in the AP mode. A client initially associates to one radio. It selects the radio with the best signal at the association time. However, since the clients are portable and the wireless environment may change, the initially selected radio may not indefinitely be the best radio and hence the client may have to roam to another radio.

However, roaming is initiated by clients and cannot directly be controlled by the radios in the ViVATO switch. The roaming behavior is dependent on the client implementation. In most commercially available clients, roaming is triggered when the channel quality (SNR) falls below a threshold. The channel quality assessment (SNR measurement) is based on received beacon strength. To ensure that a client is associated with the best radio, ViVATO's switch forces the client to roam to the radio with the best signal quality. This is done using the beam-switching algorithm.

Also, to ensure seamless roaming between beams as directed by the client, we support the Inter-Access Point Protocol (IAPP) which is an extension to 802.11 to support interoperability, mobility, handover, and coordination between APs (or radios) in a wireless LAN.

For LittleJoe to support roaming or load balancing for a client, it is necessary to support IAPP where reassociation occurs in the MAC layer and is transparent to the upper layers.

The specification of IAPP is defined by IEEE 802.11f and this specification is used as a baseline for the implementation of handover messaging between beams. The messaging is implemented within the panel's host controller. However, at a later time, it may be viable to expose the IAPP messaging through the backhaul to other LittleJoe panels as well as other 802.11f compliant third party APs.

17.3 Roaming Requirements

50

A-50

- Beamswitching: To ensure that the clients are associated with the radio (beam) with the best signal level
- IAPP: to ensure client-initiated seamless roaming between the radios in the panel.

1.19

17.3 *Beam-Switching Algorithm*

The roaming algorithm disassociates the client once it moves out of the associated main beam. However, such movement is difficult to detect in the wireless environment and disassociation may result in packet loss and long association procedure. The effect is particularly significant for clients located between two neighboring beams. So, the roaming algorithm will disassociate the client when there is significant difference between signal qualities on different beams.

1.19.1

17.1 *Beam-Switching State Transition Diagram*

Figure 16 As shown in *The roaming state machine*.

A-51

Little Joe Functional Specification

, the roaming state machine for each client served by the panel has three states:





Figure 16 The roaming state machine.

Monitor: once a client is associated to a beam, the radio continues to collect RSSI values for each packet received from that client. It recalculates a new measure called the *aSmoothedRSSIValue*⁹ over a window size of *aRSSIWindowSize* and compare it to a threshold called the *aRSSILowerControlLimit*.

Correct Beam Test: the scanning radio is used to measure the RSSIs and calculate *aSmoothedRSSIValue* for the client on each of the adjacent ports. *aRSSIWindowSize* samples for the two adjacent ports are averaged and compared to the same parameter for the current beam to determine the best beam.

⁹The roaming algorithm has no provisioned parameters. All the parameters are internal to the algorithm and may not be altered by Network Management. In this document all internal parameters use the oblique style of the AvantGarde font: (*aInternalParameter*) and provisioned parameters use the Arial font: (*aProvisionedParameter*).

A-52i



A-5Zii

Force Roam: the client is placed temporarily onto a black list so that it cannot associate to the current beam. Then the disassociation procedure will be called.

The rules governing state transitions and the subsequent actions are:

- *Association:* this transition occurs automatically when the client associates
- *Sample RSSI:* recalculate $aSmoothedRSSIValue$ and $aRSSILowerControlLimit$.
- *Smoothed RSSI drop:* $aSmoothedRSSIValue$ drops to $aRSSILowerControlLimit$
- *Correct Beam:* scan indicates current beam is the best.
 - Action: resample with new RSSI values and recalculate a new $aLowerControlLimit$
- *Roaming Scan Timeout:* the scanning radio has been monitoring the neighbouring beams for more than $aRoamingScanTimeout$ without any decision about the correct beam.
- *Wrong Beam:* scan indicates a better beam whose RSSI exceeds the RSSI of the current beam by $aSignalDropThreshold$ dB.
 - Action: black-list the client and force a disassociation.
- *Roaming Timeout:* Timeout after $aRoamingTimeOut$.
 - Action: remove the client from the black list and remove any state information about the client.
- *Reassociation:* client initiates a reassociation to a new beam.
 - Action: allow the reassociation sequence to take place and remove state information about the client.

8...1. Calculation of $aLowerControlLimit$

The parameter $aLowerControlLimit$ is calculated using both the mean and the standard deviation of RSSI. The $aLowerControlLimit$ is calculated as follows:

$$aLowerControlLimit = \overline{RSSI} - 2\sigma$$

$$\overline{RSSI} = \frac{1}{N} \sum_{i=0}^{N-1} RSSI_i \quad N = aRSSIWindowSize \text{ in frames}$$

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=0}^{N-1} (RSSI_i - \overline{RSSI})^2}$$

where $RSSI_i$ is the RSSI value reported for frame i . The N -1th frame is the most recent frame.

8...2. Smoothed RSSI Calculation

A-53

Little Joe Functional Specification

To detect client movement, the recommendation is to sample RSSI values continuously as a packet arrives and to calculate a *smoothedRSSIValue* (S). This can be calculated with the following formula.

$$S_j = 0.1RSSI_j + 0.9S_{j-1}$$

This value is then compared to the *alowerControlLimit* and if it is larger than the limit, the client enters the Correct BeamTest state.

1.20

17.3 IAPP (Seamless Roaming)

The objective of IAPP is to enable seamless client-initiated roaming between beams within the panel, between panels, and between a panel and third party APs.

Seamless roaming enables clients that are associated with the following features:

- Clients radios are able to transition across beams within a time resolution of about 100 ms without user intervention.
- No intervention required by the operating system. Therefore independent of operating system used by the client.
- Roam across multiple panels
- Roam from the ViVATO panel to 3rd party APs that conform to IEEE 802.11f.

For details of IAPP, please refer to Jim, Brennan, Seamless Roaming for LittleJoe, Internal ViVATO Publication, Se.

9. Channel Assignment

The DP2330 channel assignment includes two parts: measurement of metrics and channel assignment.

The metric measurement function resides in the host. It measures the channel activity and provides information for channel assignment and other purposes. Channel assignment function is provided in the management software package. It provides the best channel assignment based on the given measurement information. The management software functions related to channel assignment is decided by software development. This section only focuses on the algorithm.

A-54

1.21

17.3 Channel Assignment Provisioned Parameters

- *aChannelAssignmentCycle*: the time duration between changes in the channel assignment (default value 24 hours).
- *aHeavyInterference*: the interference activity threshold. If interference activity is above this value, the channel is considered as **Bad Channel** (value TBD).
- *aBadChannelThreshold*: the number of measurement periods (*aMeasurementDuration*) that a channel has interference activity above *aHeavyInterference* threshold (default value 4).
- *aJamInterference*: the interference activity threshold. Interference activity is above this value in the last measurement, **Emergency Exit** is triggered (value TBD, $aJamInterference > aHeavyInterference$).

1.22

17.3 Channel Assignment Internal Parameters

- *aMeasurementCycle*: the time duration in which a complete measurement is done (default value 24 hours).
- *aMeasurementDuration*: the time duration between two measurement points (default value 30 minutes).
- *aPeakLoadLimit*: the maximum load allowed on one channel (value TBD).
- *aChannelSixBiasFactor*: the bias factor to penalize transmission on channel 6 to reduce the intermodulation problem.

1.23

17.3 Channel Assignment Metrics

The scanning radio and traffic radios have to measure some metrics of channel activity. The measurement shall repeat in a cycle of *aMeasurementCycle*. During *aMeasurementCycle*, the metrics are measured every *aMeasurementDuration*. The desired metrics include: number of associated clients, throughput and packet error rate (PER) of each traffic radio; interference and channel utilization of each beam/(frequency) channel. The metrics are defined below.

- $N_i(t)$: Number of associated clients of the *i*th traffic radio. It is collected from the MIB, and is averaged over *aMeasurementDuration* period.
- $S_i(t)$: Throughput of the *i*th traffic radio. It is in packets/second or bytes/second, whichever is available. It is also available from the MIB, and is averaged over *aMeasurementDuration* period.

A-55

Little Joe Functional Specification

- $P_i(t)$: PER of the i th traffic radio. It is collected from the MIB, and is averaged over $aMeasurementDuration$ period.
- $D_i(t)$: delay of the i th traffic radio. It is collected from the MIB, and is averaged over $aMeasurementDuration$.
- $\rho_{ij}(t)$: channel utilization of the i th beam on the j th channel. It is the portion of time CCA is set for a given channel. It is measured by both traffic and scanning radios and is averaged over $aMeasurementDuration$. We refer to this as Channel Utilization Factor (CUF).
- $Ns_j(t)$: number of downlink packets transmitted on the j th channel. It is averaged over $aMeasurementDuration$ period.
- $Nr_{ij}(t)$: number of correctly received uplink packets transmitted by the clients associated with the i th beam on the j th channel. It is available in the MIB and is averaged over $aMeasurementDuration$ period.
- $Nn_{ij}(t)$: number of uplink packets transmitted by the clients associated with other beams, which are correctly received by the i th beam on the j th channel. It is measured by the scanning radio and is averaged over $aMeasurementDuration$ period. We call this the Self Interference Metric (SIM).
- $No_{ij}(t)$: number of uplink packets transmitted by the clients from overlapping subnets, which are correctly received by the i th beam on the j th channel. It is measured by the scanning radio and is averaged over $aMeasurementDuration$ period. We call this the Overlapping Subnet Interference (OSI).
- $Ne_{ij}(t)$: number of uplink packets with PLCP or data CRC errors in the i th beam on the j th channel. It is measured by the scanning radio and is averaged over $aMeasurementDuration$ period. We call this the Unidentified Interference Metric (UIM).
- $I_{ij}(t)$: interference of the i th beam on the j th channel. It is the portion of time CCA is set due to interference. It is measured by the scanning radio.

All the metrics are maintained in a table within $aMeasurementCycle$. When the cycle restarts, the table can either be cleared or updated with some aging factor. The table update mechanism is TBD.

It is difficult to measure $I_{ij}(t)$ when there are traffic radios on the same channel. In such cases, $I_{ij}(t)$ can be derived from other measurements. To estimate $I_{ij}(t)$ we first estimate the total number of packets from the overlapping subnets. Assuming that all downlink packet transmissions from the panel lead to CCA high, and assuming that all uplink packets have the same error probability. Then the total number of packets (with and without CRC errors) from overlapping subnets may be estimated

$$NI_{ij}(t) = No_{ij}(t) + Ne_{ij}(t) \times \frac{No_{ij}(t)}{Nr_{ij}(t) + Nn_{ij}(t) + No_{ij}(t)}$$

by:

A-56

And $I_{ij}(t)$ may be estimated

by:
$$I_{ij}(t) = \frac{NI_{ij}(t)}{Ns_i(t) + Nr_{ij}(t) + Nn_{ij}(t) + No_{ij}(t) + Ne_{ij}(t)} \rho_{ij}$$

The metrics $N_i(t)$, $S_i(t)$ and $I_{ij}(t)$ are necessary for the channel assignment algorithm. The remaining metrics are used to estimate $I_{ij}(t)$ and would therefore not be needed if there was a direct way to measure interference.

1.24

17.3 Channel Assignment Algorithm

17.1 Channel Assignment Preprocessing

9...1. Eliminate Bad Channels

A channel that has interference activity more than aHeavyInterference for aBadChannelThreshold is not used. The interference activity is averaged over intervals of aMeasurementDuration. There are typically 48 measurement intervals in one aMeasurementCycle. If the number of periods where interference activity is more than aHeavyInterference exceeds aBadChannelThreshold, then that channel is eliminated.

9...2. Estimate Total Users in the Beam

The total number of active users in the beam may be estimated by dividing the number of associated users in that beam by the percentage of time available to those users. The total number of users on beam i and channel j may therefore be described

by:
$$N_{ij}(t) = \frac{N_i(t)}{1 - \tilde{I}_{ij}(t)}$$

where:

$$\tilde{I}_{ij}(t) = \min\{I_{ij}(t), aHeavyInterference\}$$

which is the interference activity limited to the maximum allowable interference on a given beam. This ensures that the estimate does not provide large peaks due to unusual period of high interference.

1.24.1

K-57

Little Joe Functional Specification

17.1 Block-based Channel Assignment Algorithm

The block-based channel assignment algorithm assigns neighbouring beams to the same frequency channel (see Figure 1). Such assignment can help minimize the hidden beam problem.




Figure 17 Figure 1. Block-based channel assignment

The algorithm breaks the 13 beams into a maximum of 3 blocks, with each block assigned to one frequency channel (channels 1, 6, or 11), so that the peak load on each channel is minimized. In order to find the optimal solution, we have to determine:

- the boundaries between the assignment blocks (i.e. the number of beams in each block),
- the frequency channel of each block.

There are a total of possible 66 combinations that divide 13 beams into 3 blocks. For each of these possible combinations, the three blocks have to be assigned to 3 different channels. The number of channel permutations is 6. So, we have to find out the best channel-beam combination from 66x6=396 possible total combinations. Denote $L_j(t)$ as the total load on the j th frequency channel at time t . Let

$$L_j^* = \max\{L_j(t)\} \quad t \in [0, T]$$

be the peak load on the j th channel in the last measurement period, where T is the measurement cycle ($aMeasurementCycle$). Our objective is to perform an exhaustive search to find the combination that minimizes the peak load on all channels. Mathematically, this may be described as:

$$\min\{\max\{L_j^*\}\} \quad j \in [1, f_6, f_{11}]$$

It is possible that the overall network load is very light. In such cases, it is not necessary to use all three frequency channels. We define a parameter $aPeakLoadLimit$. If the total load is below this limit, only two frequency channels (preferably 1 and 11)¹⁰ are used. If the peak load on any of the two channels still exceeds the $aPeakLoadLimit$, we then use all three frequency channels.

9...1. Ignoring the Intermodulation Problem

¹⁰Due to the susceptibility of the uplink channel 1 and 11 to downlink (6 and 11) and (1 and 6), respectively, it is desirable to avoid using channels (6 and 11) or (1 and 6) on the downlink. Therefore, when assigning two channels, the first choice is to have channels (1 and 11).

A-58

The algorithm is as follows:

- Step 1: divide 13 beams into 2 blocks. There are a total of 12 possible combinations. For each combination, the channel selection can be: $f_1f_6, f_1f_{11}, f_6f_1, f_6f_{11}, f_{11}f_1, f_{11}f_6$. There are a total of 72 block-channel combinations. Assume the k th combination has the following configuration:
 - Block 1: beams 0 to b_k (0 to $N-2$) are assigned to channel C_1
 - Block 2: beams b_k+1 to $N-1$ (1 to $N-1$) are assigned to channel C_2

$$L_{C_1}^k(t) = \sum_{i=0}^{b_k} N_{iC_1}(t)$$

$$L_{C_2}^k(t) = \sum_{i=b_k+1}^{N-1} N_{iC_2}(t)$$

Then the load of channel C_1, C_2 are:

Now, let the peak load on the first block for combination k be denoted

$$PL_1(k) = \max\{L_{C_1}^k(t)\} \quad t \in [0, T]$$

by:

$$PL_2(k) = \max\{L_{C_2}^k(t)\} \quad t \in [0, T]$$

then the peak load for the busiest block (channel)

$$PL_{max}(k) = \max\{PL_1(k), PL_2(k)\}$$

is:

- Step 2: select the combination index R with the least peak load. In other words choose R

such that:

$$PL_{max}(R) = \min\{PL_{max}(k)\} \quad \forall (0 \leq k \leq 71)$$

Simply described, this chooses the combination of channels and beams that minimize the peak load on any channel.

- Step 3: check that the peak load on the channel is less than aPeakLoadLimit. If not, go to three channel assignment.

Three Channel Assignment:

- Step 1: Since the peak load on two channels exceeds the threshold, we have to assign the load to three channels. In other words we have to divide the 13 beams into 3 blocks. There are a total of 66 possible combinations. For each combination, the channel selection can be: $f_1f_6f_{11}, f_1f_{11}f_6, f_6f_1f_{11}, f_6f_{11}f_1, f_{11}f_1f_6, f_{11}f_6f_1$. There are a total of 396 block-channel combinations. Assume the k th combination has the following configuration:
 - Block 1: assign beams 0 to b_k (0 to $N-3$) to channel C_1

A-59

Little Joe Functional Specification

- . Block 2: assign beams b_k+1 to p_k (1 to $N-2$) are to channel C_2
- . Block 3: assign beams p_k+1 to $N-1$ (2 to $N-1$) to channel C_3

Then the load of channel C_1 , C_2 , and C_3 are:

$$L_{C_1}^k(t) = \sum_{i=0}^{b_k} N_{iC_1}(t)$$

$$L_{C_2}^k(t) = \sum_{i=b_k+1}^{p_k} N_{iC_2}(t)$$

$$L_{C_3}^k(t) = \sum_{i=p_k+1}^{N-1} N_{iC_3}(t)$$

Now, let the peak load on the first block for combination k be denoted

$$PL_1(k) = \max\{L_{C_1}^k(t)\} \quad t \in [0, T]$$

$$PL_2(k) = \max\{L_{C_2}^k(t)\} \quad t \in [0, T]$$

by:

$$PL_3(k) = \max\{L_{C_3}^k(t)\} \quad t \in [0, T]$$

then the peak load for the busiest block (channel) is:

$$PL_{max}(k) = \max\{PL_1(k), PL_2(k), PL_3(k)\}$$

- Step 2: select the combination index R with the least peak load. In other words choose R such

$$PL_{max}(R) = \min\{PL_{max}(k)\} \quad \forall (0 \leq k \leq 395)$$

that:

9...2. Considering the Intermodulation Problem

If considering the intermodulation problem, we would like to avoid the combinations f_1f_6 and f_6f_{11} channels. Therefore the best approach is to avoid channel f_6 all together. In other words:

The algorithm is as follows:

- Step 1: divide 13 beams into 2 blocks. There are a total of 12 possible combinations. For each combination, the channel selection can be: f_1f_{11}, f_1f_1 . There are a total of 24 block-channel combinations. Assume the k th combination has the following configuration:
 - . Block 1: beams 0 to b_k (0 to $N-2$) are assigned to channel C_1

A-60

Block 2: beams b_k+1 to $N-1$ (1 to $N-1$) are assigned to channel C_2

Then the load of channel f_i, f_{i1} is the sum of the loads of the beams assigned to those channels. In other words:

$$L_{f_i}^k(t) = \sum_{f_i} N_{if_i}(t) \quad \forall (i \in f_i)$$

$$L_{f_{i1}}^k(t) = \sum_{f_{i1}} N_{if_{i1}}(t) \quad \forall (i \in f_{i1})$$

Now, let the peak load on the first block for combination k be denoted by:

$$PL_1(k) = \max\{L_{f_i}^k(t)\} \quad t \in [0, T]$$

$$PL_2(k) = \max\{L_{f_{i1}}^k(t)\} \quad t \in [0, T]$$

then the peak load for the busiest block (channel) is:

$$PL_{max}(k) = \max\{PL_1(k), PL_2(k)\}$$

- Step 2: select the combination index R with the least peak load. In other words choose R such

that:
$$PL_{max}(R) = \min\{PL_{max}(k)\} \quad \forall (0 \leq k \leq 23)$$

Simply described, this chooses the combination of channels and beams that minimizes the peak load of the busiest channel.

- Step 3: check that the peak load on the channel is less than aPeakLoadLimit¹¹. If not, go to three channel assignment.

Three Channel Assignment:

- Step 1: Since the peak load on two channels exceeds the threshold, we have to assign the load to three channels. In other words we have to divide the 13 beams into 3 blocks. There are a total of 66 possible combinations. For each combination, the channel selection can be: $f_i f_{i6} f_{i1}$, $f_i f_{i1} f_{i6}$, $f_{i6} f_{i1} f_i$, $f_{i6} f_i f_{i1}$, $f_i f_{i1} f_{i6}$, $f_{i1} f_{i6} f_i$. There are a total of 396 block-channel combinations. Assume the k th combination has the following configuration:
 - Block 1: assign beams 0 to b_k (0 to $N-3$) to channel C_1
 - Block 2: assign beams b_k+1 to p_k (1 to $N-2$) to channel C_2

¹¹Considering the intermodulation problem, this parameter should be set higher to avoid the three channel combination as much as possible.

A-61

Little Joe Functional Specification

Block 3: assign beams p_{k+1} to $N-1$ (2 to $N-1$) to channel C_3

Then the load of channel $f_1, f_6,$ and f_{11} are:

$$L_{f_1}^k(t) = \sum_{f_1} N_{if_1}(t) \quad \forall (i \in f_1)$$

$$L_{f_6}^k(t) = \gamma \sum_{f_6} N_{if_6}(t) \quad \forall (i \in f_6)$$

$$L_{f_{11}}^k(t) = \sum_{f_{11}} N_{if_{11}}(t) \quad \forall (i \in f_{11})$$

The parameter γ is the bias factor for channel 6 (*aChannelSixBiasFactor*) used to give a larger weight to channel 6 in order to penalize it's selection.

Now, let the peak load on the first block for combination k be denoted by:

$$PL_1(k) = \max\{L_{f_1}^k(t)\} \quad t \in [0, T]$$

$$PL_2(k) = \max\{L_{f_6}^k(t)\} \quad t \in [0, T]$$

$$PL_3(k) = \max\{L_{f_{11}}^k(t)\} \quad t \in [0, T]$$

then the peak load for the busiest block (channel) is:

$$PL_{max}(k) = \max\{PL_1(k), PL_2(k), PL_3(k)\}$$

- Step 2: select the combination index R with the least peak load. In other words choose R such

that:
$$PL_{max}(R) = \min\{PL_{max}(k)\} \quad \forall (0 \leq k \leq 395)$$

1.24.2

17.1 Emergency exit

If the i th traffic radio has interference larger than *aJamInterference* in the last measurement period, it has to move to the remaining channels.

10. Downlink Traffic-Shaping

A-62

The traffic-shaping functionality is implemented to limit the downlink load to a stable operating point and this also ensure fairness between uplink and downlink traffic load.

This section specifies the traffic shaping functionality implemented in LittleJoe. The scope is limited to specifying the functional architecture, associated algorithms and mechanisms required to implement the feature.

1.25

17.3 Traffic-Shaping Requirements

In this section we specify requirements for the traffic shaping feature. Prior to that, we define the following terms:

- **Load:** the total traffic in bits/sec offered to the MAC layer (either at AP or at client) by the higher layers (e.g. TCP/IP).
- **Operating Point:** downlink load for which uplink failure rate is less than `aMaxFailureRate` (value TBD).

The requirements for the traffic shaping feature are as follows:

- Traffic shaping needs to determine the operating point for the panel and limit the downlink load to it.
- The estimation of operating point needs to be done dynamically based on changing value of uplink failure rate.

1.26

17.3 Traffic-Shaping Architecture

The traffic shaping feature is implemented at the network layer. As shown in *Functional Placement of th*, it should be placed between the backhaul and the MAC layer. All downlink traffic including traffic generated from the backhaul and internal traffic destined from one client to another client served by the panel are passed through the traffic shaper.

A-63

Little Joe Functional Specification




Figure 18 *Functional Placement of the Traffic Shaper.*

Figure 19 As depicted in *Architecture of Traffic Shaping Module*

, the traffic shaping module consists of three components.

- Traffic queues

1.26.1 These queues define the granularity of traffic shaping function. A single queue is maintained for all downlink traffic. The reasons for this are discussed in

- Gra.
- Shaper (One per queue)
 - If aggregate downlink load (offered to all radios) exceeds a threshold, system becomes unstable. The shaping function needs to estimate the load-threshold and then ensure that offered load to MAC remains below it. The leaky bucket algorithm is used for traffic shaping.
- Parameter Configuration Module (One per Shaper)
 - This module is responsible for dynamically adapting shaping-parameters to meet system performance objective (i.e. keep uplink failure rate below the specified

A-64

threshold). Section 4.3 specifies operation of this entity.




Figure 19 Architecture of Traffic Shaping Module


17.3 Traffic-Shaping Functional Description

This section specifies the functionality of traffic shaping feature. An overview of the leaky bucket algorithm and parameter update mechanism is provided. The most critical component of the specification is identifying correct values of parameters controlling leaky bucket operation. In this section, we specify the mechanisms to evaluate correct parameter values. The actual operational values are TBD.

1.26.2

17.1 Leaky Bucket Algorithm

The leaky bucket operation shown in *Functional Placement of th* is specified by parameters (R , B). The leaky bucket algorithm ensures that outgoing traffic never exceeds the specified rate R .



A-65

Figure 20 *Leaky Bucket Operation*.

Following is a description of leaky bucket operation .

- Traffic to be shaped (from the backhaul and internal client-to-client traffic) is called the incoming traffic and the shaped traffic (going to MACs) is called the outgoing traffic.
- R (*aMaxDownlinkLoad*) is specified over a time-period ΔT i.e. the output from leaky bucket during ΔT should not exceed R bits.
- If during ΔT (*aShaperWindowPeriod*), the incoming traffic is less than R bits then all traffic is allowed to pass through the bucket.
- If incoming traffic size is greater than R bits, then up to R bits are transmitted and remaining packets –called violating traffic – are queued (In some implementations, violating traffic can be dropped).
- In the next ΔT duration, the shaping criterion is applied to queued packets and newly arrived packets, i.e. up to R bits are transmitted.
- If size of queued packets exceeds the bucket depth, B (*aMaxQueueSize*) bits, then further incoming packets are dropped.

1.26.3

17.1 Granularity of Operation

Traffic shaping can be defined for either single queue or there can be multiple queues (e.g. per client queue or per beam queue). Since the combined load to all AP's in the panel is the critical factor in uplink performance, it is recommended to maintain a single queue as input to the traffic shaper.

Besides if multiple queues are maintained and load thresholds are defined for each queue, it may result in unused capacity in some scenarios.

1.26.4

17.1 Dynamic Parameter Update

The operating point for LJ system changes depending on traffic characteristics at any given time, e.g. as uplink load increases the operating point shifts to lower downlink loads. Thus it is required to continuously measure certain metrics and determine the optimal value of the operating point.

The dynamic parameter update entity is responsible for implementing this functionality. It defines only one traffic metric.

66

A-66

10...1. Traffic Metric

Following metric is measured:

- Uplink failed frames (*aNumUplinkFailedFrames*)
 - This metric is indication of number of uplink frames lost at AP(s).
 - It includes retransmitted (and corrupted) frames in uplink.
 - It is not a direct indication of failure rate seen by clients (because a client could succeed in transmitting a frame after N retries, where $N < \text{Max-Retry-Limit}$).
 - It is an indication of the extent of hidden beam problem.

The number of failed frames should be less than *aMaxFailureRate* during the measurement window (*aShapingMeasurementPeriod*).

The measurement window (*aShapingMeasurementPeriod*) for this metric is the duration for which the traffic metric is measured to determine whether operating point needs to change.

The parameter update module entity also implements the operating point estimation algorithm, which is specified in the next section.

1.26.5

17.1 Operating Point Estimation Algorithm

The algorithm is as follows:

- Select an initial value for the operating point (a table, generated from simulations provides the operating points for different system configurations, e.g. with and without CCA/CBF). This value is the *aMaxDownlinkLoad* (R) of the leaky bucket.
- Monitor the uplink failure rate (*aNumUplinkFailedFrames*) during measurement window. If the failure rate is higher than the specified threshold (*aMaxFailureRate*)
 - Reduce operating point by Δd (*aShapingDecrementStep*) (i.e. reduce the value of R by $\Delta d\%$).
 - Measure *aNumUplinkFailedFrames* during *aShapingMeasurementPeriod*.

While *aNumUplinkFailedFrames* is higher than *aMaxFailureRate*:

- Reduce operating point exponentially (i.e. reduce operating point by $2\Delta d\%$ in 2nd pass, $4\Delta d\%$ in 3rd pass and so on, where a pass is defined as *aShapingMeasurementPeriod*).
- Repeat this process for *aMaxDecrementCount* (N) times i.e. if $N = 4$, then maximum reduction in operating point will be $8\Delta d\%$ and operating point will

A.67

Little Joe Functional Specification

remain fixed at that value while *aNumUplinkFailedFrames* is higher than *aMaxFailureRate*.

If *aNumUplinkFailedFrames* drops below *aMaxFailureRate*

- Increase operating point by $\Delta i\%$ (*aShapingIncrementStep*)
- Measure *aNumUplinkFailedFrames* during *aShapingMeasurementPeriod*.
- While *aNumUplinkFailedFrames* remains below *aMaxFailureRate*
 - Increase operating point exponentially i.e. by $2\Delta i\%$ in 2nd pass, $4\Delta i\%$ in 3rd pass and so on.

1.27

17.3 Provisioned and Internal Parameters

This section lists the parameters which control the operation of the traffic shaping feature.

The only parameter provisioned by the network management element is:

- *aMaxFailureRate*: Uplink Failure Rate Threshold ($F\%$) – Provisioned (Typical value is between 2-5%)

The internal parameters are:

- *aMaxQueueSize*: Leaky bucket depth B (bits)
- *aShaperWindowPeriod*: Leaky bucket time-period ΔT – (Typical Value is 1 second)
- *aShapingIncrementStep*: Operating point exponential increment step size ($\Delta i\%$) - Provisioned
- *aShapingDecrementStep*: Operating point exponential decrement step size ($\Delta d\%$) – Provisioned
- *aMaxDecrementCount*: Maximum number of passes for exponential reduction of operating point (N) - Provisioned
- *aShapingMeasurementPeriod*: Measurement Window duration – Provisioned
- *aMaxDownlinkLoad*: Leaky bucket rate R (bits/sec) – Dynamically updated depending on estimated operating point
- *aNumUplinkFailedFrames*: Number of uplink failed frames in the current measurement period (*aShapingMeasurementPeriod*)¹²

A-68

11. The Scanning Radio

Figure 4 A scanning receiver in "promiscuous mode" and a "beam" antenna switch is used to obtain signal strength and interference information from stations on different beams as shown in *The scanning radio*.

The scanning radio has two states scan mode and roaming mode. The state machine for the scanning radio is shown in *The state machine for the scanning radio*.




Figure 21 *The state machine for the scanning radio*.

17.3 Scan Mode

While in the scan mode, the radio periodically scan the 13 beams on the 3 channels and collects activity information and fills the appropriate tables (39 columns). The table has the following contents described in detail in

Channel Assignment Metrics

Channel Utilization Factor (CUF)

- Self Interference Metric (SIM)
- Overlapping Subnet Interference (OSI)
- Unidentified Interference Metric (UIM)

The table contents will be a running average of the above metrics. The radio should scan each combination for at least one second in a one-minute period. However, the scanning process may be interrupted by the roaming function.

A-69

Little Joe Functional Specification

1.28

17.3 Roaming

In this mode, the scanning radio is placed alternately on the neighbouring beams of the radio under test. The radio will collect the RSSI value for each frame received from the client on the two neighbouring beams. Once, the radio has received at least one frame from each neighbouring beam, it can go back to scan mode.

If there are no frames received on both neighbouring beams for a period longer than *aRoamingScanTimeout*, the radio goes back to scan mode.

Additionally, we would like to ensure that, on average, no more than half of the scanning receiver's time is spent in the roaming state. The timeout in the roaming mode (*aRoamingScanTimeout*), is set to $\min(60, T/2)$ seconds where T is the average time between requests for roaming-test measurements. The Algorithm to determine the timeout (T) is described by the pseudo code below. The algorithm uses the function `time()` (which returns the current time in seconds) and two static variables: T and `last_test_time`.

The pseudo-code for the initialization of `last_test_time` is:

```
last_test_time = 0;
```

The pseudo-code for the processing required at the start of each roaming-test measurement is:

```
// update T and last_test_time
if (last_test_time == 0) then
// initialize T
T = 60;
else
// update T
T = 0.9 * T + 0.1 * (time() - last_test_time);
endif
last_test_time = time();
// set the timeout
set_test_timeout(min(60seconds, T/2));
```

1.29

17.3 State Transitions

The rules governing the state transmissions are:

70

A-70

correct beam test: one of the radios is being tested for the correct beam test as explained in

- Intr.

end beam test: a decision is made on the correct beam or the timeout period `aRoamingScanTimeout` discussed in

- Intr expires.

12. References

- [1] M. Brewer, D. Lohman, et. al. "Software System Architecture Document", ViVATo Internal Document, Version 0.3, 5/9/02.
- [2] Ed Casas, "Beamforming for LittleJoe", *ViVATO Technical Report*, Feb. 1, 2002
- [4] Ed Casas, LittleJoe Link Budget, *ViVATO Technical Report*, Feb. 25, 2002
- [5] Siavash Alamouti, LittleJoe Link Budget Spread Sheets, May 10, 2002.
- [6] Ed Casas, LittleJoe Beamforming, *ViVATO Technical Report*, Feb. 1, 2002
- [7] Jim Brennan, "LittleJoe Mac Model", *ViVATO Technical Report*, March 1, 2002
- [8] G. Anastasi, et. al., "MAC Protocols for Wideband Wireless Local Access: Evolution Towards Wireless ATM", *IEEE Personal Communications Magazine*, Oct. 1998, pp. 53-64
- [9] R. Guesalla, "Characterizing the Variability of Arrival Processes with Indexes of Dispersion", *IEEE JSAC*, vol. 9, no. 2, Feb. 1991, pp. 203-11.
- [10] W. E. Leland, et. al. "on the Self-Similar Nature of Ethernet Traffic," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, Feb. 1994, pp. 1-15
- [11] S. Deng, "Empirical Model of WWW Document Arrivals at Access Link", in the *Proceedings of ICC'96*
- [12] L. Greenstein and V. Ereg, "Gain Reductions Due to Scatter on Wireless Paths with Directional Antennas," *IEEE Communication Letters*, vol. 3, no. 6, June 1999, pp. 169-171
- [13] M. J. Gans, R. A. Valenzuela, J. H. Winters, and M.J. Carloni, "High Data Rate Indoor Wireless Communications Using Antenna Arrays," in *Proceeding of 6th International Symposium on Personal, Indoor and Mobile Radio Communications*, vol. 3, pp. 1040-1046, 1995.
- [14] G. German, Q. Spencer, L. Swindlehurst, and R. Valenzuela, "Wireless Indoor Channel Modeling: Statistical Agreement of Ray Tracing Simulations and Channel Sounding Measurements," in *Proceeding of 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 4, pp.2501-2504, 2001.

A-71

Little Joe Functional Specification

[15] J. G. Wang, A. Mohan, and T. Aubrey, "Angles-of-Arrival of Multipath Signals in Indoor Environments," in *Proceeding of Vehicular technology Conference*, vol. 1, pp.155-159, 1996.

[16] W.C. Jakes, ed., *Microwave Mobile Communications*. Wiley, 1974.

[17] J. Kivinen, X. Zhao, and P. Vainikainen, "Empirical Characterization of Wideband Indoor radio Channel at 5.3 GHz," *IEEE Transactions on Antennas and Propagation*, vol. 49, no. 8, Aug. 2001, pp. 1192-1203

[18] J. Medbo and J. E. Berg, "Simple and Accurate Path Loss Modeling at 5 GHz in Indoor Environments with Corridors," in *Proceeding of the 52nd Vehicular technology Conference*, vol. 1, pp.30-36, 2000.

[19] H. Hashemi, "The Indoor Propagation Channel," *Proceedings of the IEEE*, vol. 81, no. 7, July 1993, pp. 943-968

[20] J.E. Berg, "Building Penetration Loss along Urban Street Microcells," in *Proceeding of PIMRC '96*, vol. 3, pp.795-797, 1996.

[21] E. Damosso and L. Corcia, eds., *COST 213 Final Report - Digital Mobile Radio - Towards future Generation Systems*. European Commission, Directorate General XIII, 1999. Report Number EUR 18957 (ISBN 92-828-5416-7)

[22] J. D. Kraus, *Antennas*. McGraw-Hill, 1950.

[23] J. E. Hudson, *Adaptive Array Principles*. Peter Peregrinus and IEE, 1981.

[24] A. V. Oppenheim and R. W. Schaffer, *Digital Signal Processing*. Prentice-Hall, 1975.



A-72

Part 2: Software System Architecture

13. Introduction

The purpose of this document is to describe at a high level the software features of Little Joe and how those features map onto system tasks. The types of interaction between tasks are also described.

Little Joe is based on the Linux operating system. The kernel has been customized such that only devices physically present on our system are supported. The system will contain the following components:

- PowerPC processor (integrated with DMA, Memory Controller and PCI Bus)
- SODIMM memory (64-128Mbytes)
- Onboard FLASH (32-64Mbytes)
- Three Ethernet controllers (10/100)
 - Secure management
 - Back haul
 - Daisy chain to the next panel
- The “King Arthur” custom PCI bridge
 - Connects to 11 802.11b chipsets
- Two RS232 ports
 - Management Console
 - Debug
- Temperature sensor
- Interface into the “Merlin” beam steer component
- PCI Interface to an 802.11b Search MAC

Software other than Linux that is used is:

PPCBoot

<http://ppcboot.sourceforge.net>

PCMCIA

<http://pcmcia-cs.sourceforge.net>

Wireless Tools

http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Tools.html

CISH

<http://www.tarball.net/cish/>

Apache

<http://www.apache.org>

SNMP

<http://www.net-snmp.org>

A-73

Tools used for the target build are:

- GNU C version 2.95.3
- GNU gLibC version 2.2.4
- CVS

Linux kernel version for target is TBD. Linux kernel for host builds is 2.4.13.

A-74

14. Software Overview

The software leverages a reasonable amount of IP from the open source community ranging from the boot code to the protocol stacks themselves. The Mabuhay software IP (for the 1st release) resides in the following components:

- The integration of the Linux components onto the custom PCB.
- The King Arthur driver code.
- Merlin beam steering algorithms and its driver.
- Access Point control software.

The rest of the code is primarily open source and the job is largely an integration effort to get all of the components to operate in our target system in a manner which is compatible with the product requirements.

17.3 Software not covered by this Document

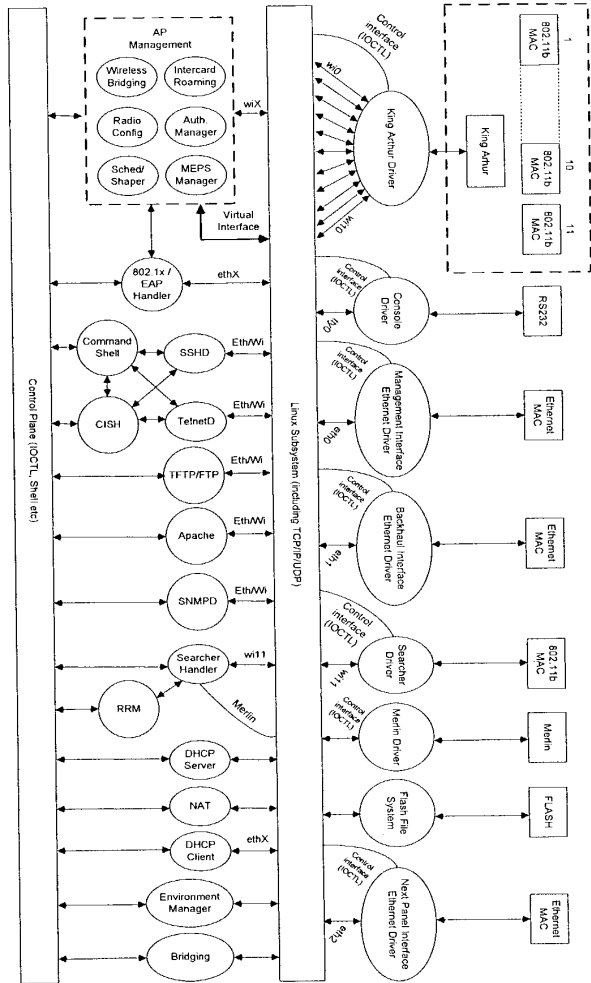
This document covers the modules and subsystems which will be present when the system is running. However, there are other important components which are not covered. These are:

- PowerPC-boot
 - Diagnostics
 - FLASH
 - RAM
 - Ethernet interfaces
 - Serial interfaces
 - Temperature calibration
 - Fan controls
 - Flash file system for Kernel
- Linux Kernel Configuration
 - Providing the necessary drivers for essential components

A-75

2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17

17.3 Software Module Overview



A-76



Little Joe Functional Specification

15. Control Plane

To simply the connectivity between modules the diagram uses a control plane which is the lingo for any one or many of the following inter-process communication techniques:

- RPC
- A shell executing a script and extracting output (textual or otherwise)
- IOCTL
- Protocol stack interface (socket or otherwise)
- Shared memory
- File


On a module by module basis a developer will decide and publish the interface to their code. The code-base is using open source software, so if the module is expected to contain valuable intellectual property that is proprietary to Mabuhay then the interface should be disjoint – for example the beam steering software should not be code compiled into the CISH command line interpreter.

16. Drivers

17.3 Console Driver

The management console interface is simply and RS232 port configured, out of the box, to be 9600 baud, 8 bits, no stop bits.

The manner in which this operates when a user connects a terminal to it should be as follows:



A-77

17.3 Ethernet Drivers


There are three Ethernet interfaces on the product. The interfaces plumb into the underlying Linux subsystem in an identical manner. However, they differ in the services offered. The following diagrams show which daemons are listening to which ports.

1.29.1

17.1 Secure Management

This port is intended as a separate “secure” out-of-band port which may be wired into an existing infrastructure to manage the device. The port is secure, not because it is encrypted (it is not), but because it is physically separate and will not participate in packet forwarding. Under no circumstances can data packets sent into the wired or wireless interfaces be forwarded to this port. Device management over this interface may be performed by CLI, HTTP(S) and SNMP.

Therefore the stack is as follows:



A driver should be developed to interface the target device to the Linux subsystem in a standard (ethX) manner. If no such driver exists for the chosen part then one will be developed and offered up to the open source community, since the driver contains no Mabuhay proprietary intellectual property.

Configuration settings of this interface should be as follows:

- The ability to set the interface to Auto-negotiate
 - The interface should advertise and be able to negotiate to 10 Mbps, 100 Mbps, Full Duplex, Half Duplex, no flow control.
- The ability to set the interface manually to:
 - 10 Mbps, Half Duplex
 - 100 Mbps, Half Duplex
 - 10 Mbps, Full Duplex
 - 100 Mbps, Full Duplex
- Out of the box:
 - Auto-negotiate enabled

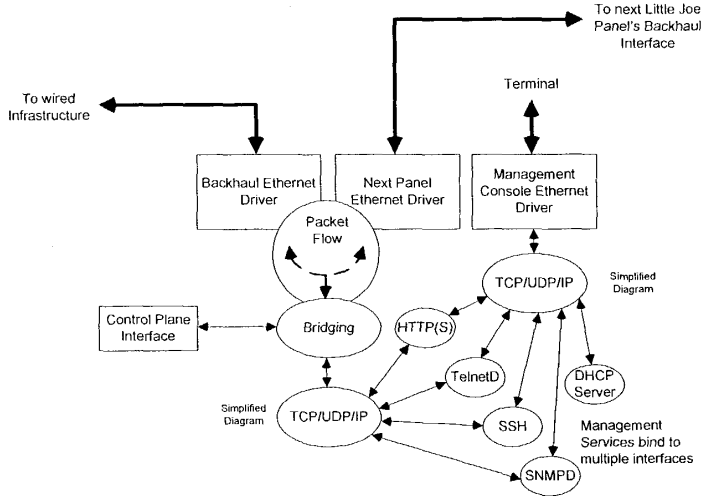
1.29.2

A-78

Little Joe Functional Specification


17.1 Backhaul / Daisy-chain for the Next Panel

This interface provides the connection from Little Joe to the wired infrastructure. Packets which may transfer into and out from this interface are as follows:



The diagram shows that management of the unit may be performed via the dedicated management port or the wired interfaces. However, via the control plane, management may be disabled to any interface for specific services. For example, each management system will have a list of "allowable" interfaces that it may serve itself to. Out of the box, CISH (which is spawned by TelnetD upon authentication clearance) may be disabled for certain interfaces, such as the wireless interface (wi13), wired backhaul, or next panel.

A typical setup of allowable services for the Backhaul interfaces may be:



The DHCP Client Service is shown. This is because a normal operating characteristic of existing access points from factory power-on is for it to gather its own IP address and perform bridging between the wired, wireless and, in the case of Little Joe, the next panel interface.

A-79

1.30**17.3 Source for Wireless Drivers**

Many of the wireless cards, such as those from Intersil and Agere are supported by source available in the open-source community. For the Intersil cards the following are available:

<http://people.ssh.com/jkm/Prism2/>

This is the so-called Host-AP mode driver sponsored by Intersil that allows a client card to become an access point, except that management functions such as association and authentication are handled by the CPU, but Beacon Generation and Probe Responses are performed in hardware.

<http://www.linux-wlan.com/linux-wlan/>

The WLAN-driver project is another driver for the same card, except this is for client-mode applications.

Both projects utilize the iwconfig utilities which use an IOCTL interface to talk to the drivers. The drivers live in kernel space and can be bound in with the kernel or loaded at runtime.

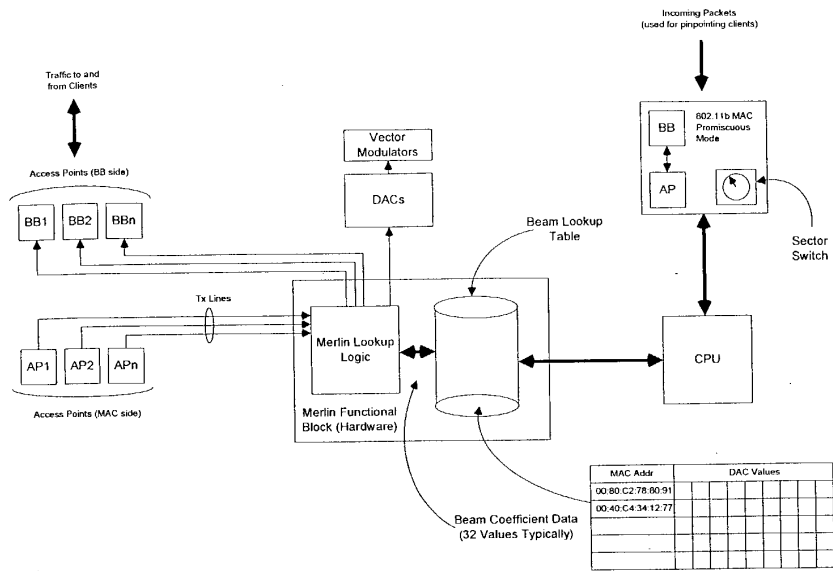
Drivers for wireless interfaces for Linux have been supplied in just object code, such as those from Cisco and Nokia, and also in source form, from the likes of Intersil and Agere.

1.31**17.3 Searcher and Merlin Interfaces and Functions**

The searcher MAC is designed to determine where clients are and build a table for the Merlin beam steerer. In order to understand the control system behind these two functions, consider the following diagram which shows how each of the wireless APs in the system interface to Merlin:

A-80

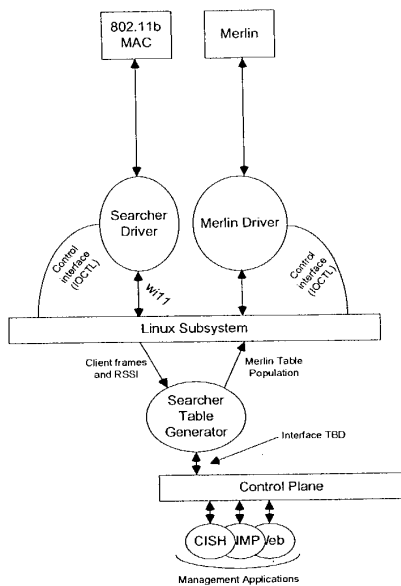
Little Joe Functional Specification



The Searcher consists of a regular 802.11b MAC and a sector switch. The MAC is set into "Promiscuous Mode" – meaning that all valid 802.11 data frames are received and can be passed to the CPU.

The sector switch allows the CPU to change the receive beam that the searcher MAC is focused on. By a process of sampling each of the sectors a map can be built of client MAC addresses and their respective receive strength signal indications (RSSI). By using simple math it is then possible to determine their location. The location consists of a series of values which can be fed into a vector modulator in order to "shine" a beam onto a specific client. These values determine both the reach (power) and the direction. Thus, when an AP wishes to transmit a frame to a client, the Merlin logic is able to see which MAC address it is to be communicated with; adjust the beam to the appropriate location, and transmit the frame. The software processes which make this happen are as follows:

A-81



17.3 Scheduler / Shaper

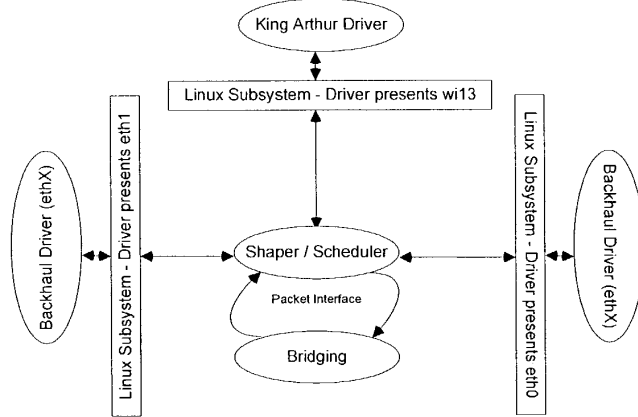
Due to the nature of 802.11 networks when combined with directional antennas, it is important to be able to identify situations which would or could adversely affect performance.

For example, it is possible that one user may wish to perform a large download, and due to this it could, when taken to an extreme, cause other clients to be starved of service.

This module identifies when such situations occur and performs the necessary packet shaping (bandwidth management) to ensure that the unit may still be able to see when another client is requesting service. There are numerous ways of achieving this and the exact algorithm is still to be determined. However, the positioning of the module in the system is clear – it must be between the back and wireless devices and be able to set up queues and bandwidth limits on a per client basis. Therefore, within the architecture it would live as follows for a typical packet flow:

A-82


Little Joe Functional Specification



17.3 RRM (Radio Resource Management)

This module takes input from the searcher which is promiscuously receiving frames in each direction, on a programmable channel.

The purpose of RRM is to identify potential interfering (or high usage) networks in order to determine if the system should be switched to a different channel. It is TBD whether it is allowable for different pointing directions are allowed to be on different channels. Therefore the control loop for decision making would be:



A-83

17.3 FLASH

This driver needs to present an interface to Linux which is compatible with a regular block device such as a hard drive. However, since Flash has a limited number of write cycles, and the number of parameters that can be changed in the system is significant and likely to get a lot larger as the product develops, changes in configuration will need to be committed through a user "save" command. Such an action will write the configuration to Flash rather than individual writes for each small change a user makes.

Note: Research is continuing into a MTD Driver and JFFS Journaling Flash File System Drivers.

Component	Filesize	Memory
Httpd	550K	3000K
Httpd.conf	50K	
libperl.so	1.2 MB	
Cish	150K	500K
Net-snmp	20K	
Total	~2MB	3500K

This table shows initial estimations of management module resource usage. Need clarification of data items in the table. Some of these items are daemons, running all of the time. Others are one instance per invocation. Cish is dependent on lots of other elements, such as ipchains, brctl and ifconfig.

Deleted: 1

2. Management Interfaces

There are three core management interfaces:

- CLI (RS232, Telnet, SSH, ASCII File)
- HTTP
- SNMP

All three management interfaces are available through the RS232 console port, Backhaul (eth0), Management (eth1), Next Panel (eth2) and the Wireless interfaces. Instances of the HTTP, SNMP, Telnet/SSH daemons will be bound to the eth0, eth1, eth2 and wix interfaces. The user will have the ability to disable management services on individual or all interfaces.

The following list is an initial draft of the features in the Little Joe AP which will be configured and monitored using the CLI, HTTP, and SNMP management services.

- 1) Basic Setup
 - a) SSID
 - b) IP
- 2) Chassis
 - a) Slots
 - b) Ethernet Ports

Formatted: Bullets and Numbering

14-84

Little Joe Functional Specification

- c) Antenna
 - i) Calibration
- d) Radios
- e) Environment
- 3) Layer 2
 - a) Ethernet
 - b) 802.11
 - c) Bridging
 - d) VLAN
 - e) Filters
 - f) Statistics
- 4) Layer 3
 - a) Basic IP Configuration
 - b) SSH
 - c) HTTP
 - d) Telnet
 - e) SNMP
 - i) MIB 1 & 2
 - ii) Mabuhay MIBs
 - iii) Dot1 and dot11 MIBs
 - iv) Traps
 - f) DHCP
 - g) NAT
 - h) NTP
 - i) FTP
 - j) Filters
 - k) Statistics
- 5) AP Manager
 - a) Associations
 - b) Local AP Roaming
 - c) Inter AP Protocol / Wireless Distribution System
 - d) Multi-AP Configuration (4 panel setup)
- 6) Security
 - a) WEP
 - i) 40-bit
 - ii) 128-bit
 - b) Authentication
 - i) Open
 - ii) Shared
 - iii) EAP
 - iv) Radius
- 7) Utilities
 - a) Users/Administrators
 - b) Preferences
 - c) Searching and Sorting

A-85

- d) Configuration File
 - i) Import/Export
 - ii) Version(s)
- e) Firmware
 - i) Uppdate
 - ii) Version(s)
- f) File
- g) Syslog
- h) TFTP
- i) FTP
- 8) Diagnosics
 - a) Events/Logs/Traps
 - b) Network
 - i) Ping
 - ii) Tracerroute
 - c) Radio
 - d) Antenna
 - i) Calibration
 - e) Merlin
 - i) Dump Beam Coefficient Data Table

17.1 HTTP

The HTTP(S) service provides a web based management interface. The complete set of CLI commands will be available through the Web screens. Access to the web interface is controlled by the management authentication process. The web interface will consist of standard HTML and CGI scripts. We are evaluating the mod_perl Apache module for an additional server-side scripting tool. The goal is to support Internet Explorer, Netscape, Mozilla and other browsers on multiple platforms. Apache is well supported and is a solid base for future features including XML, SOAP, JSP, ASP and Web Services.

The HTTPD can be configured to bind on the Backhaul (eth0), Management (eth1), Next Panel (eth2) and Wireless interfaces (wix). The HTTP service can be configured to bind on all or none of the interfaces.

The Apache HTTP server with SSL/TLS libraries provides a secure HTTPS platform.


User Interface (UI)

- Forms
- Help

The forms provide standard Web UI features including edit, list boxes, radio buttons, and tables. The UI advances the CLI giving the user for example a single page to do basic setup instead of a series of commands. Also, wizards (a series of forms) will be developed to walk the user through complex configuration tasks. Advanced UI features – Graphic view of the chassis, Network Topology, Graphical views of beam status will be considered in the future.

A-86

Little Joe Functional Specification




This diagram shows the high-level HTTP module interfaces.

1.32

17.3 CISH (CLI)

CISH is a command line (CLI) interface based using a simple Cisco-like interface, including a simple help system and command completion and command history. Command syntax is defined in tables, making it fairly easy for developers to add commands. CISH validates the command parameters, then will use the control plane interface to configure and get status information. Configuration is stored in a ASCII text file. User authentication is password only (no user ID). An ASCII Text File provides a persistent store of the configuration. The file contains the CLI commands used to configure the device. The file can be used to configure multiple devices with a common setup.



This diagram shows the high-level CLI interfaces.

A-87

17.3 User Manager

Provide a facility to add/remove users for management of the device. Users will have associated capabilities. For example, an operator may have read-only access, and a admin user have full read-write capabilities.

Deleted: user

Deleted: n

17.4 SSHD

Secure Shell from www.openssh.org – contains support for SSH1 and SSH2. Secure Shell provides encryption, authentication and tunneling capabilities. The OpenSSH package includes the ssh program which replaces telnet and rlogin, scp replaces rcp and sftp which replaces ftp.

OpenSSH does not support any patented transport algorithms. In SSH1 mode, only 3DES and Blowfish are available options. In SSH2 mode, only 3DES, Blowfish, CAST128, Arcfour and AES can be selected. The patented IDEA algorithm is not supported.

OpenSSH provides support for both SSH1 and SSH2 protocols.

Since the RSA patent has expired, there are no restrictions on the use of RSA algorithm using software.

17.5 Telnet

The Telnet daemon module will spawn a CISH shell. The telnetd service can be bound to selected interfaces. The telnetd service can be enabled/disabled.

1.33

17.6 SNMPD

The SNMP agent will be based on the NET-SNMP open source project formerly known as the UCD-SNMP package, originally based on the Carnegie Mellon University SNMP implementation (version 2.1.2.1).

The net-snmp package supports SNMPv1, SNMPv2 and SNMPv3. TDB if we need to include v3 support.

Support for MIB 1 and MIB 2 as well as Mabuhay private MIBs and Traps. Register Mabuhay private mib with iana.org.

A SNMP access function framework will provide a standard interface for SNMP set, get, getNext and Traps for other developers to integrate their modules into the SNMP agent.

18. Environmental Control

A-88

Little Joe Functional Specification

18.1 Temperature and Fan Control

The temperature module is fairly simple – it simply needs to check the temperature against predefined limits (programmable) and take appropriate action when the system is above or below temperature (or approaching it). Changes in temperature will result in the closed loop control system operating the fans (below). However, when cooling fails, the following actions can be taken:


- Issue a TRAP to an SNMP NMS
- Write error information in the local store syslog. Optionally notify a UNIX Syslog daemon.
- Issue warnings to the console, CLI and web interface.

19. Wireless Control

19.1 Wireless Bridging

The function of this module will be dependent on the actual firmware that is selected for the Access Points in the system. However, if the chosen firmware does its own wireless bridging, then this section describes its function. However, should it not, then the function will be provided through the Bridging Manager described later.

From an AP perspective, associated clients would populate a location database as follows:



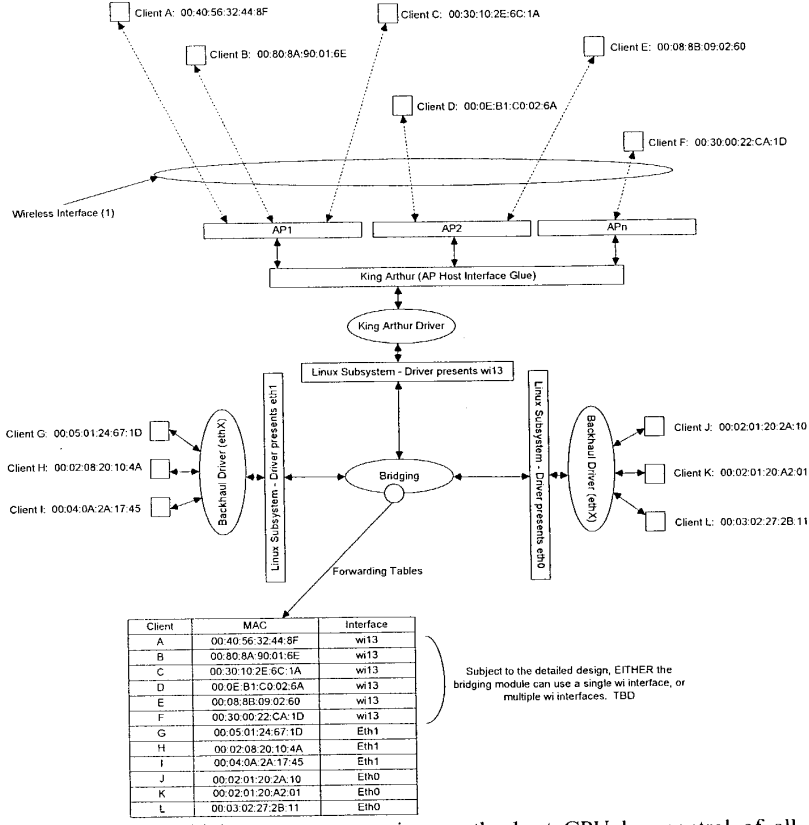
Traffic local the AP is kept local, and only traffic destined for the wired backhaul is transferred out of the AP itself.

A-89

1.34

19.2 Centralized Bridging

It is desirable to have all packets come to a central bridging manager as follows:



In this example, the bridging manager running on the host CPU has control of all of the interfaces for data forwarding. In the example the table is indexed by MAC address. However, although release 1 may not be VLAN-aware, the extensibility of the table and indexed must be considered in the detailed design.

A-90

Little Joe Functional Specification

The diagram above also shows the wireless interfaces under a single wi13 interface. The driver for wi13 would then do an additional lookup to determine the actual wireless interface to which a user may be attached.

However, an alternative which may be concluded is a better choice (dependent on final detailed design for that module), may be to bridge each of the wireless interfaces together and provide a layer 3 IP interface in the same broadcast domain. This would remove the necessity of providing two area for layer 2 lookups to be performed.

19.3 Radio Configuration Manager

SSID, SSID Broadcast, Data Rates, RTS/CTS settings, Default Radio Channel, # of Wireless clients.

1.35

19.3 Inter-card Roaming Manager

1.36

19.3 Mabuhay Enhanced Performance System (MEPS)

1.37

19.3 Security

1.38

19.3 Authentication Manager

The authentication manager will be responsible for handling incoming 802.11b authentication requests. The authentication manager receives all authentication frames and determines the authentication mechanism to use based on the bits in the authentication type field. Based on the authentication type the request is handed-off to the authentication handler which will complete the authentication process.

1.39

19.3 Authentication

Before a client device can gain access to the AP and the network, it must be authenticated. There are four AP authentication mechanisms: Shared key, Open Authentication, MAC address, and EAP.

- Open Authentication – Allows any client to authenticate with AP, but only allows data transfer if WEP keys match.

A-91

- Shared Key – 802.11b standard but not recommended because of vulnerabilities. Sends unencrypted challenge string to the client. Client responds with encrypted response, if correct the AP allows the client to authenticate.
- MAC address – The MAC address is either checked against a local AP table of allowable MACs or sent to a RADIUS server for verification. If the MAC is not in the list of allowable MACs then the device is not authenticated.

1.40

19.3 802.1x / EAP Authentication Mechanism

802.1X defines Extensible Authentication Protocol (EAP) over LANs (EAPOL). The standard encapsulates and leverages much of EAP, which was defined for dial-up authentication with Point-to-Point Protocol in RFC 2284.

Client sends authenticate request to AP, AP sends an EAPOL-encapsulated EAP request-ID to the client. The client responds with an EAPOL-encapsulated EAP response-ID message containing the user's identification. The AP then reencapsulates this same EAP response-ID message in a RADIUS access request packet and forwards this to a RADIUS server. EAP messages are relayed between the client and RADIUS by the AP, on the client side encapsulated in EAPOL, and on the server side inside a RADIUS packet.

In the final step, the RADIUS server responds with a RADIUS access accept (or deny) packet containing an encapsulated EAP success (or failure), which the AP then forwards to the client. In the case of success, the port is considered opened for data traffic and the user authenticated.

When using dynamic session keys the RADIUS access accept will include session keys, which are used by the wireless access point to build, sign and encrypt an EAPOL key message.

This is sent to the client immediately following the EAP success message. With this information, both client and wireless access point can program their encryption keys dynamically, making the encryption more difficult to crack.

WEP (Wired Equivalent Privacy)

Enable/Disable WEP.

40-bit and 128-bit key setup.

2. Other Control

19.3 DHCP Client

This module is enabled by default and operates through the wired Backhaul interface. Unless a specific IP address is specified (which from the factory can only be done through the Management Console port) the device will boot and try and get its management IP address from a DHCP Server on the Backhaul interface.

A-92

Little Joe Functional Specification

1.41

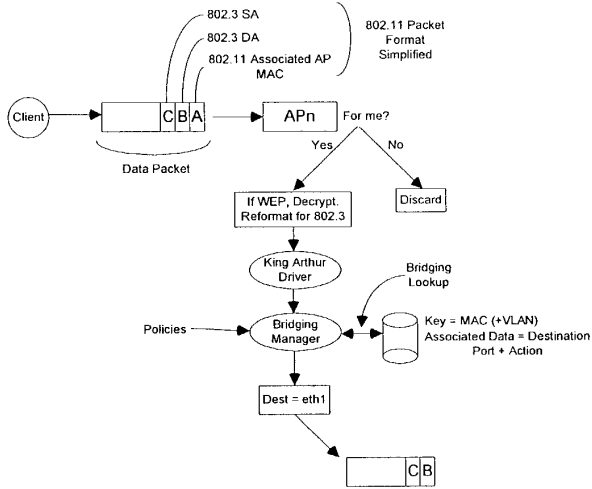
19.3 *DHCP Server and Network Address Translation (NAT)*

94

A-93

3. Typical Packet Walk (Bridged)

The following diagram shows the one way trip from an 802.11b client through to the 802.3 Backhaul Interface.




19. Typical Packet Walk (NAT and DHCP Server)

Another typical deployment for an access point, especially when there is no central DHCP Server, is for the access point to provide an IP gateway functionality whilst providing private (10 or 192 net) addresses for clients to communication over. The diagram below shows the modules that are utilized to get a packet from a client to a backhaul interface.

The NAT module maintains stateful information on each connection. Each session is translated such that a connection to a destination site is translated with a new source address (from a private address to the real IP address of the access point), and the source system is identified by a new IP source port number. When a frame is received from the destination on the backhaul interface a lookup is performed on the destination address port to see if it matches an existing connection. If a match occurs the NAT function will reform the packet with the private 10 or 192 net address and transmit the frame out of the appropriate wireless interface.

A-94

Little Joe Functional Specification

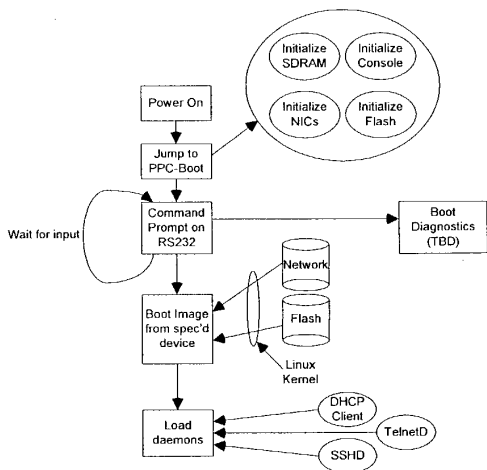


A-95

19. Diagnostics

Diagnostic tools will be provided for testing the radios.

19. Out-of-the-Box Power-up Sequence



19. Utilities

19.4 FTP Client (for downloading new Images)

This is a standard FTP client provided as part of a standard Linux distribution. An interface will be provided through both the CLI and the web server so that a customer may be perform upgrades of their embedded software.

By gaining access to the regular command shell (such as BASH) a field engineer may also use the full range of FTP facilities.

1.42


19. FTP Server (for Serving images to other panels)

If deemed necessary this would be useful for gaining access to the unit for downloading of configuration files etc.

A-96

19. Data Packet Interfaces

The following diagram depicts the Linux interface stack along with the bindings to layer 3. The wireless interfaces present themselves to the system in much the same way as any other interface. However, the data interface through which traffic may be transmitted and received is only one of those interfaces. As discussed earlier in this design, the bridging between wired and wireless and between wireless and between wired may be the same bridging module or via two separate modules. The compliance with the Linux interface architecture allows a lot of flexibility when it comes to implementation of the featureset.



A-97

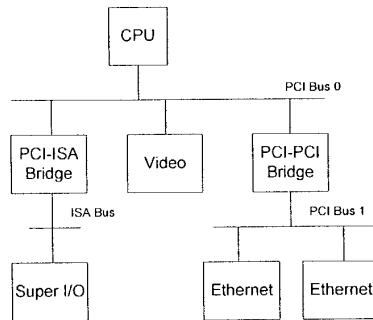
19. Appendix A – PCI / PCI Bridge Support in Linux

Peripheral Component Interconnect (PCI), as its name implies is a standard that describes how to connect the peripheral components of a system together in a structured and controlled way. The standard describes the way that the system components are electrically connected and the way that they should behave. This chapter looks at how the Linux kernel initializes the system's PCI buses and devices.

1.42.1

19.4 Example PCI Based System

This is a logical diagram of an example PCI based system. The PCI buses and PCI-PCI bridges are the glue connecting the system components together; the CPU is connected to PCI bus 0, the primary PCI bus. A special PCI device, a PCI-PCI bridge connects the primary bus to the secondary PCI bus, PCI bus 1. In the jargon of the PCI specification, PCI bus 1 is described as being downstream of the PCI-PCI bridge and PCI bus 0 is up-stream of the bridge. Connected to the secondary PCI bus are the two ethernet devices for the system. Physically the bridge, secondary PCI bus and two devices would all be contained on the same combination PCI card. The PCI-ISA bridge in the system supports older, legacy ISA devices and the diagram shows a super I/O controller chip.



1.42.2

19.4 PCI Address Spaces

The CPU and the PCI devices need to access memory that is shared between them. This memory is used by device drivers to control the PCI devices and to pass information between them. Typically the shared memory contains control and status registers for the device. These registers are used to control the device and to read its status.

The CPU's system memory could be used for this shared memory but if it were, then every time a PCI device accessed memory, the CPU would have to stall, waiting for the PCI device to

A-98

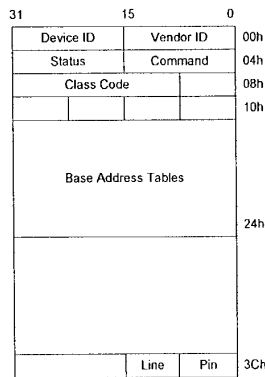
Little Joe Functional Specification

finish. Access to memory is generally limited to one system component at a time. This would slow the system down. It is also not a good idea to allow the system's peripheral devices to access main memory in an uncontrolled way. This would be very dangerous; a rogue device could make the system very unstable.

Peripheral devices have their own memory spaces. The CPU can access these spaces but access by the devices into the system's memory is very strictly controlled using DMA (Direct Memory Access) channels. ISA devices have access to two address spaces, ISA I/O (Input/Output) and ISA memory. PCI has three; PCI I/O, PCI Memory and PCI Configuration space. All of these address spaces are also accessible by the CPU with the the PCI I/O and PCI Memory address spaces being used by the device drivers and the PCI Configuration space being used by the PCI initialization code within the Linux kernel.

1.42.3

19.4 PCI Configuration Headers



Every PCI device in the system, including the PCI-PCI bridges has a configuration data structure that is somewhere in the PCI configuration address space. The PCI Configuration header allows the system to identify and control the device. Exactly where the header is in the PCI Configuration address space depends on where in the PCI topology that device is. For example, a PCI video card plugged into one PCI slot on the motherboard will have its configuration header at one location and if it is plugged into another PCI slot then its header will appear in another location in PCI Configuration memory. This does not matter, for wherever the PCI devices and bridges are the system will find and configure them using the status and configuration registers in their configuration headers.

Typically, systems are designed so that every PCI slot has it's PCI Configuration Header in an offset that is related to its slot on the board. So, for example, the first slot on the board might have its PCI Configuration at offset 0 and the second slot at offset 256 (all headers are the same length, 256 bytes) and so on. A system specific hardware mechanism is defined so that the PCI

A-99

configuration code can attempt to examine all possible PCI Configuration Headers for a given PCI bus and know which devices are present and which devices are absent simply by trying to read one of the fields in the header (usually the Vendor Identification field) and getting some sort of error. The describes one possible error message as returning 0xFFFFFFFF when attempting to read the Vendor Identification and Device Identification fields for an empty PCI slot.

19.4 Layout of the 256 byte PCI configuration header

It contains the following fields:

1.42.4

19. Vendor Identification

A unique number describing the originator of the PCI device. Interestingly, Intel's is 0x8086.

1.42.5

19. Device Identification

A unique number describing the device itself.

1.42.6

19. Status

This field gives the status of the device with the meaning of the bits of this field set by the standard.

1.42.7

19. Command

By writing to this field the system controls the device, for example allowing the device to access PCI I/O memory,

1.42.8

19. Class Code

This identifies the type of device that this is. There are standard classes for every sort of device; video, SCSI and so on. The class code for SCSI is 0x0100.

A-100

Little Joe Functional Specification

1.42.9

19.4 Base Address Registers

These registers are used to determine and allocate the type, amount and location of PCI I/O and PCI memory space that the device can use.

1.42.10

19.4 Interrupt Pin

Four of the physical pins on the PCI card carry interrupts from the card to the PCI bus. The standard labels these as A, B, C and D. The Interrupt Pin field describes which of these pins this PCI device uses. Generally it is hardwired for a particular device. That is, every time the system boots, the device uses the same interrupt pin. This information allows the interrupt handling subsystem to manage interrupts from this device.

1.42.11

19.4 Interrupt Line

The Interrupt Line field of the device's PCI Configuration header is used to pass an interrupt handle between the PCI initialization code, the device's driver and Linux's interrupt handling subsystem. The number written there is meaningless to the device driver but it allows the interrupt handler to correctly route an interrupt from the PCI device to the correct device driver's interrupt handling code within the Linux operating system.

19. PCI I/O and PCI Memory Addresses

These two address spaces are used by the devices to communicate with their device drivers running in the Linux kernel on the CPU. For example, some fast Ethernet devices map their internal registers into PCI I/O space. The Linux device driver then reads and writes those registers to control the device.

Until the PCI system has been set up and the device's access to these address spaces has been turned on using the Command field in the PCI Configuration header, nothing can access them. It should be noted that only the PCI configuration code reads and writes PCI configuration addresses; the Linux device drivers only read and write PCI I/O and PCI memory addresses.

1.43

19. PCI-ISA Bridges

These bridges support legacy ISA devices by translating PCI I/O and PCI Memory space accesses into ISA I/O and ISA Memory accesses. A lot of systems now sold contain several ISA bus slots and several PCI bus slots. Over time the need for this backwards compatibility will dwindle and PCI only systems will be sold. Where in the ISA address spaces (I/O and Memory) the ISA devices of the system have their registers was fixed in the dim mists of time by the early Intel 8080 based PCs. Even a \$5000 Alpha AXP based computer systems will have its ISA

102

A-101

floppy controller at the same place in ISA I/O space as the first IBM PC. The PCI specification copes with this by reserving the lower regions of the PCI I/O and PCI Memory address spaces for use by the ISA peripherals in the system and using a single PCI-ISA bridge to translate any PCI memory accesses to those regions into ISA accesses.

19. PCI-PCI Bridges

PCI-PCI bridges are special PCI devices that glue the PCI buses of the system together. Simple systems have a single PCI bus but there is an electrical limit on the number of PCI devices that a single PCI bus can support. Using PCI-PCI bridges to add more PCI buses allows the system to support many more PCI devices. This is particularly important for a high performance server. Of course, Linux fully supports the use of PCI-PCI bridges.

1.44

19. PCI-PCI Bridges: PCI I/O and PCI Memory Windows

PCI-PCI bridges only pass a subset of PCI I/O and PCI memory read and write requests downstream. For example, in the diagram at the beginning of the appendix, the PCI-PCI bridge will only pass read and write addresses from PCI bus 0 to PCI bus 1 if they are for PCI I/O or PCI memory addresses owned by either of the Ethernet devices; all other PCI I/O and memory addresses are ignored. This filtering stops addresses propagating needlessly throughout the system. To do this, the PCI-PCI bridges must be programmed with a base and limit for PCI I/O and PCI Memory space access that they have to pass from their primary bus onto their secondary bus. Once the PCI-PCI Bridges in a system have been configured then so long as the Linux device drivers only access PCI I/O and PCI Memory space via these windows, the PCI-PCI Bridges are invisible. This is an important feature that makes life easier for Linux PCI device driver writers. However, it also makes PCI-PCI bridges somewhat tricky for Linux to configure.

1.45

19. PCI-PCI Bridges - PCI Configuration Cycles and PCI Bus Numbering

1.45.1 Type 0 PCI Configuration Cycle Figure:


31	11 10	8 7	2 1 0
Device Select	Func	Register	0 0

A-102

Little Joe Functional Specification

1.45.2

1.45.3 Type 1 PCI Configuration Cycle Figure:



So that the CPU's PCI initialization code can address devices that are not on the main PCI bus, there has to be a mechanism that allows bridges to decide whether or not to pass Configuration cycles from their primary interface to their secondary interface. A cycle is just an address as it appears on the PCI bus. The PCI specification defines two formats for the PCI Configuration addresses; Type 0 and Type 1; these are shown in the figures above. Type 0 PCI Configuration cycles do not contain a bus number and these are interpreted by all devices as being for PCI configuration addresses on this PCI bus. Bits 31:11 of the Type 0 configuration cycles are treated as the device select field. One way to design a system is to have each bit select a different device. In this case bit 11 would select the PCI device in slot 0, bit 12 would select the PCI device in slot 1 and so on. Another way is to write the device's slot number directly into bits 31:11. Which mechanism is used in a system depends on the system's PCI memory controller.

Type 1 PCI Configuration cycles contain a PCI bus number and this type of configuration cycle is ignored by all PCI devices except the PCI-PCI bridges. All of the PCI-PCI Bridges seeing Type 1 configuration cycles may choose to pass them to the PCI buses downstream of themselves. Whether the PCI-PCI Bridge ignores the Type 1 configuration cycle or passes it onto the downstream PCI bus depends on how the PCI-PCI Bridge has been configured. Every PCI-PCI bridge has a primary bus interface number and a secondary bus interface number. The primary bus interface being the one nearest the CPU and the secondary bus interface being the one furthest away. Each PCI-PCI Bridge also has a subordinate bus number and this is the maximum bus number of all the PCI buses that are bridged beyond the secondary bus interface. Or to put it another way, the subordinate bus number is the highest numbered PCI bus downstream of the PCI-PCI bridge. When the PCI-PCI bridge sees a Type 1 PCI configuration cycle it does one of the following things:

1. Ignore it if the bus number specified is not in between the bridge's secondary bus number and subordinate bus number (inclusive)
2. Convert it to a Type 0 configuration command if the bus number specified matches the secondary bus number of the bridge
3. Pass it onto the secondary bus interface unchanged if the bus number specified is greater than the secondary bus number and less than or equal to the subordinate bus number.

So, if we want to address Device 1 on bus 3 of the topology defined in the section entitled "PCI-PCI Bridging Step 3" we must generate a Type 1 Configuration command from the CPU. Bridge1 passes this unchanged onto Bus 1. Bridge2 ignores it but Bridge3 converts it into a Type 0 Configuration command and sends it out on Bus 3 where Device 1 responds to it.

104

A-103

It is up to each individual operating system to allocate bus numbers during PCI configuration but whatever the numbering scheme used the following statement must be true for all of the PCI-PCI bridges in the system:

"All PCI buses located behind a PCI-PCI bridge must reside between the secondary bus number and the subordinate bus number (inclusive)."

If this rule is broken then the PCI-PCI Bridges will not pass and translate Type 1 PCI configuration cycles correctly and the system will fail to find and initialize the PCI devices in the system. To achieve this numbering scheme, Linux configures these special devices in a particular order. Section Assigning PCI Bus Number describes Linux's PCI bridge and bus numbering scheme in detail together with a worked example.

1.46

19. *Linux PCI Initialization*


The PCI initialisation code in Linux is broken into three logical parts:

1. PCI Device Driver
 - a. This pseudo-device driver searches the PCI system starting at Bus 0 and locates all PCI devices and bridges in the system. It builds a linked list of data structures describing the topology of the system. Additionally, it numbers all of the bridges that it finds.
2. PCI BIOS
 - a. This software layer provides the services described in bib-pci-bios-specification. Even if there is no BIOS, there is equivalent code in the Linux kernel providing the same functions.
3. PCI Fix-up
 - a. System specific fix-up code tidies up the system specific loose ends of PCI initialization.

19.4 The Linux Kernel PCI Data Structures

A-109

Little Joe Functional Specification



As the Linux kernel initializes the PCI system it builds data structures mirroring the real PCI topology of the system. The figure above shows the relationships of the data structures that it would build for the example PCI system described at the beginning of this appendix.

Each PCI device (including the PCI-PCI Bridges) is described by a `pci_dev` data structure. Each PCI bus is described by a `pci_bus` data structure. The result is a tree structure of PCI buses each of which has a number of child PCI devices attached to it. As a PCI bus can only be reached using a PCI-PCI Bridge (except the primary PCI bus, bus 0), each `pci_bus` contains a pointer to the PCI device (the PCI-PCI Bridge) that it is accessed through. That PCI device is a child of the the PCI Bus's parent PCI bus.

Not shown in the above is a pointer to all of the PCI devices in the system, `pci_devices`. All of the PCI devices in the system have their `pci_dev` data structures queued onto this queue. This queue is used by the Linux kernel to quickly find all of the PCI devices in the system.

19.4 The PCI Device Driver

The PCI device driver is not really a device driver at all but a function of the operating system called at system initialization time. The PCI initialization code must scan all of the PCI buses in the system looking for all PCI devices in the system (including PCI-PCI bridge devices).

It uses the PCI BIOS code to find out if every possible slot in the current PCI bus that it is scanning is occupied. If the PCI slot is occupied, it builds a `pci_dev` data structure describing the device and links into the list of known PCI devices (pointed at by `pci_devices`).

The PCI initialization code starts by scanning PCI Bus 0. It tries to read the Vendor Identification and Device Identification fields for every possible PCI device in every possible PCI slot. When it finds an occupied slot it builds a `pci_dev` data structure describing the device.


A-105

All of the `pci_dev` data structures built by the PCI initialization code (including all of the PCI-PCI Bridges) are linked into a singly linked list; `pci_devices`.

If the PCI device that was found was a PCI-PCI bridge then a `pci_bus` data structure is built and linked into the tree of `pci_bus` and `pci_dev` data structures pointed at by `pci_root`. The PCI initialization code can tell if the PCI device is a PCI-PCI Bridge because it has a class code of `0x060400`. The Linux kernel then configures the PCI bus on the other (downstream) side of the PCI-PCI Bridge that it has just found. If more PCI-PCI Bridges are found then these are also configured. This process is known as a depth-wise algorithm; the system's PCI topology is fully mapped depth-wise before searching breadthwise. Looking at reference model at the beginning of the appendix, Linux would configure PCI Bus 1 with its Ethernet and SCSI device before it configured the video device on PCI Bus 0.

As Linux searches for downstream PCI buses it must also configure the intervening PCI-PCI bridges' secondary and subordinate bus numbers. This is described in detail in Section `pci-pci-bus-numbering` below.

19. Configuring PCI-PCI Bridges - Assigning PCI Bus Numbers



For PCI-PCI bridges to pass PCI I/O, PCI Memory or PCI Configuration address space reads and writes across them, they need to know the following:

1. Primary Bus Number
 - a. The bus number immediately upstream of the PCI-PCI Bridge.
2. Secondary Bus Number
 - a. The bus number immediately downstream of the PCI-PCI Bridge.
3. Subordinate Bus Number
 - a. The highest bus number of all of the buses that can be reached downstream of the bridge.

A-106

Little Joe Functional Specification


b. PCI I/O and PCI Memory Windows

The window base and size for PCI I/O address space and PCI Memory address space for all addresses downstream of the PCI-PCI Bridge.

The problem is that at the time when you wish to configure any given PCI-PCI bridge you do not know the subordinate bus number for that bridge. You do not know if there are further PCI-PCI bridges downstream and if you did, you do not know what numbers will be assigned to them. The answer is to use a depth-wise recursive algorithm and scan each bus for any PCI-PCI bridges assigning them numbers as they are found. As each PCI-PCI bridge is found and its secondary bus numbered, assign it a temporary subordinate number of 0xFF and scan and assign numbers to all PCI-PCI bridges downstream of it. This all seems complicated but the worked example below makes this process clearer.


a. PCI-PCI Bridge Numbering: Step 1

Taking the topology in under “Assigning PCI Bus Numbers”, the first bridge the scan would find is Bridge1. The PCI bus downstream of Bridge1 would be numbered as 1 and Bridge1 assigned a secondary bus number of 1 and a temporary subordinate bus number of 0xFF. This means that all Type 1 PCI Configuration addresses specifying a PCI bus number of 1 or higher would be passed across Bridge1 and onto PCI Bus 1. They would be translated into Type 0 Configuration cycles if they have a bus number of 1 but left un-translated for all other bus numbers. This is exactly what the Linux PCI initialization code needs to do in order to go and scan PCI Bus 1.

**b. PCI-PCI Bridge Numbering: Step 2**


Linux uses a depth-wise algorithm and so the initialization code goes on to scan PCI Bus 1. Here it finds PCI-PCI Bridge2. There are no further PCI-PCI bridges beyond PCI-PCI Bridge2, so it is assigned a subordinate bus number of 2 which matches the number assigned to its secondary interface. The diagram above shows how the buses and PCI-PCI bridges are numbered at this point.

A-107



c. PCI-PCI Bridge Numbering: Step 3

The PCI initialization code returns to scanning PCI Bus 1 and finds another PCI-PCI bridge, Bridge3. It is assigned 1 as its primary bus interface number, 3 as its secondary bus interface number and 0xFF as its subordinate bus number. Type 1 PCI configuration cycles with a bus number of 1, 2 or 3 will be correctly delivered to the appropriate PCI buses.



1.47

d. PCI-PCI Bridge Numbering: Step 4

Linux starts scanning PCI Bus 3, downstream of PCI-PCI Bridge3. PCI Bus 3 has another PCI-PCI bridge (Bridge4) on it, it is assigned 3 as its primary bus number and 4 as its secondary bus number. It is the last bridge on this branch and so it is assigned a subordinate bus interface number of 4. The initialization code returns to PCI-PCI Bridge3 and assigns it a subordinate bus number of 4. Finally, the PCI initialization code can assign 4 as the subordinate bus number for PCI-PCI Bridge1.

A-108

Little Joe Functional Specification

e. PCI BIOS Functions

The PCI BIOS functions are a series of standard routines which are common across all platforms. For example, they are the same for both Intel and Alpha AXP based systems. They allow the CPU controlled access to all of the PCI address spaces. Only Linux kernel code and device drivers may use them.

1.48

f. PCI Fixup

For PowerPC based systems without a BIOS to set up PCI configuration needs to happen to:

1. Allocate PCI I/O and PCI Memory space to each device.
2. Configure the PCI I/O and PCI Memory address windows for each PCI-PCI bridge in the system
3. Generate Interrupt Line values for the devices; these control interrupt handling for the device.

The next subsections describe how that code works.

1.49

g. Finding Out How Much PCI I/O and PCI Memory Space a Device Needs

Each PCI device found is queried to find out how much PCI I/O and PCI Memory address space it requires. To do this, each Base Address Register has all 1's written to it and then read. The device will return 0's in the don't-care address bits, effectively specifying the address space required.

A-109

Little Joe Functional Specification

upstream PCI-PCI Bridge's memory ranges for any given device, it is a somewhat difficult problem to allocate space efficiently.

The algorithm that Linux uses relies on each device described by the bus/device tree built by the PCI Device Driver being allocated address space in ascending PCI I/O memory order. Again a recursive algorithm is used to walk the `pci_bus` and `pci_dev` data structures built by the PCI initialization code. Starting at the root PCI bus (pointed at by `pci_root`) the BIOS fixup code:

1. Aligns the current global PCI I/O and Memory bases on 4K and 1 Mbyte boundaries respectively.
2. For every device on the current bus (in ascending PCI I/O memory needs)
 - i. Allocates it space in PCI I/O and/or PCI Memory
 - ii. Moves on the global PCI I/O and Memory bases by the appropriate amounts
 - iii. Enables the device's use of PCI I/O and PCI Memory
3. Allocates space recursively to all of the buses downstream of the current bus. Note that this will change the global PCI I/O and Memory bases.
4. Aligns the current global PCI I/O and Memory bases on 4K and 1 Mbyte boundaries respectively and in doing so figure out the size and base of PCI I/O and PCI Memory windows required by the current PCI-PCI bridge.
5. Programs the PCI-PCI bridge that links to this bus with its PCI I/O and PCI Memory bases and limits.
6. Turns on bridging of PCI I/O and PCI Memory accesses in the PCI-PCI Bridge. This means that if any PCI I/O or PCI Memory addresses seen on the Bridge's primary PCI bus that are within its PCI I/O and PCI Memory address windows will be bridged onto its secondary PCI bus.

Taking the PCI system in at the beginning of this appendix as our example the PCI Fixup code would set up the system in the following way:


1. Align the PCI bases
 - a. PCI I/O is 0x4000 and PCI Memory is 0x100000. This allows the PCI-ISA bridges to translate all addresses below these into ISA address cycles
2. The Video Device
 - a. This is asking for 0x200000 of PCI Memory and so we allocate it that amount starting at the current PCI Memory base of 0x200000 as it has to be naturally aligned to the size requested. The PCI Memory base is moved to 0x400000 and the PCI I/O base remains at 0x4000.
3. The PCI-PCI Bridge
 - a. We now cross the PCI-PCI Bridge and allocate PCI memory there, note that we do not need to align the bases as they are already correctly aligned:
 - i. The Ethernet Device
 1. This is asking for 0xB0 bytes of both PCI I/O and PCI Memory space. It gets allocated PCI I/O at 0x4000 and PCI Memory at 0x400000. The PCI Memory base is moved to 0x4000B0 and the PCI I/O base to 0x40B0.

A-III

B

[15 pages]

Behavior of 802.11 Networks With Mabusay Access Points



Contents

1 Introduction	3
1.1 Need for Steering	3
1.2 Outline	3
1.3 References to the 802.11 Standard	4
1.4 Performance	4
2 The 802.11 CSMA/CA DCF MAC Protocol	4
2.1 Basic Operation	4
2.2 802.11 Frame Formats	4
2.3 MAC and PHY Data Units	4
2.3.1 PLCP Preamble and Header	5
2.3.2 MAC Header	6
2.3.3 Frame Types	6
2.4 802.11 MAC State Machine	6
2.4.1 RECEIVE	9
2.4.2 BACKOFF	9
2.4.3 Transmit States	9
2.5 Frame Sequences	9
3 Behavior of 802.11 Networks with Incomplete Connectivity	10
3.1 Connectivity Graphs	10
3.2 Collisions in Two-Node Connectivity Graphs	10
3.3 Collisions in Three-Node ("Hidden-Terminal") Connectivity Graphs	11
3.4 Collisions in Four-Node ("Exposed Terminal" or "Hidden-Beam") Connectivity Graphs	11

1

B-1



- 3.5 Handling of Broadcast Frames 12
- 4 Behavior of 802.11 Systems with Increased Propagation Delay 12**
 - 4.1 Increased Collision Probability 12
 - 4.2 ACK and CTS Timeouts 13
- 5 Implementing a Beamforming 802.11 AP 13**
 - 5.1 First Attempt - Polling with Fake Duration Field 13
 - 5.2 Second Approach - Polling with Client Shim 14
 - 5.3 Current Solution - DCF and Hardware-Driven Steering 14
 - 5.4 Future Possibilities - 802.11 Polling 15

2
B-2



1 Introduction

Mabuhay's *Little Joe* 802.11b WLAN access point (AP) uses a steerable antenna array. This antenna provides directionality and longer range. The use of this antenna also affects the behavior of the 802.11 MAC which was originally designed for use with omnidirectional antennas and over short distances.

This study examines the performance of the unmodified 802.11 MAC protocol in the context of a system employing multiple directional AP antennas. Two issues were looked at in detail:

- performance degradation arising out of incomplete connectivity
- range limitations due to protocol timeouts

Incomplete connectivity is due to the directionality of the AP antenna and also to the increased range of the AP compared to the range of the clients. The result is that clients are less likely to detect transmissions to other clients or transmissions from other clients.

The increased range of the system leads to greater propagation delay and may lead to increased likelihood of collisions or to protocol timeouts.

1.1 Need for Steering

Little Joe's increased transmit power levels are permitted by the FCC's point-to-point rules which require that each transmission be directed to a specific receiver.

In order for the antenna to be steered on each transmission, the beam-steering hardware must be controlled by the device that implements the 802.11 MAC protocol since this is the device that generates the frames. To minimize development time it was decided to use an off-the-shelf MAC controller for *Little Joe*.

These requirements led to the earlier solutions described in section 5 and the current approach described in 5.3.

1.2 Outline

Section 2 provides a simplified description of the 802.11 protocol. Section 3 describes the effects of reduced connectivity and section 4 the effects of increased propagation delay. Section 5 describes the past, present and possible future approaches that *Little Joe* can use to deal with these issues.

1.3 References to the 802.11 Standard

Unless otherwise indicated, references to sections (e.g. §8.1) refer to sections in the 802.11 standard, 1999 edition.

1.4 Performance

Numerical results for the MAC modifications suggested in 5.3 were obtained by computer simulation of the MAC described in section 2. The simulation software and preliminary results are described in a separate document.

2 The 802.11 CSMA/CA DCF MAC Protocol

2.1 Basic Operation

The 802.11 MAC (medium access control) is the protocol used by 802.11 stations to coordinate access to the radio channel. The MAC requires stations to wait until the channel has been free for a certain period before transmitting.

No client station may transmit during the $50\mu\text{s}$ (DIFS time) immediately following any transmission. This allows APs operating in polling (PCF) mode to take over the channel simply by leaving gaps between transmissions of less than a DIFS. After the DIFS deferral time any client or AP may transmit.

The idle channel time following DIFS is divided into short (20 us) slots and stations must begin their transmissions at the start of a randomly-chosen slot. This slot duration was chosen to be long enough that a transmission at the start of one slot will be detected by all stations before the start of the next slot. All stations synchronize their timing to the end of the immediately preceding frame.

Stations choose the slot in which to begin a transmission by using a *backoff counter*. The counter decrements at the end of each slot interval that the channel is not busy. Transmission begins when the counter reaches zero and a frame is ready to be sent.

2.2 802.11 Frame Formats

2.3 MAC and PHY Data Units

Figure 1 shows the format of 802.11 packets ("frames"). The levels of encapsulation are as follows:

- the MAC service data unit (MSDU) is the payload carried by the (802.11) MAC. This will typically be an ethernet frame.




Figure 1: PHY and MAC frame protocol data unit components.

- the MAC adds a MAC header and a 32-bit CRC to the MSDU to form a MAC protocol data unit (MPDU).
- the PLCP service data unit (PSDU) is the payload carried by the PHY (typically the 802.11b Hi-Rate DSSS).
- the PHY adds a PLCP preamble and a PLCP header to the MPDU to form a PLCP protocol data unit (PPDU).
- the frame transmitted over the air is called a PLCP (PHY) protocol data unit (PPDU).

The MAC controller may fragment an MSDU into multiple MPDUs.

2.3.1 PLCP Preamble and Header

The PLCP header and preamble are specific to a particular PHY. For 802.11b there are short and long versions.

The short preamble/header is 96µs and consists of a 72-bit preamble (56 bits for sync and a 16-bit delimiter) plus a 48 bit header supplying the PHY data rate and frame duration.

The long preamble/header is 192µs long and contains a 144-bit preamble (128 bits for sync and a 16-bit delimiter) plus the same 48-bit header.

B-5

The Intersil Prism II baseband processor is able to generate and receive either short or long preambles. Agere's *Theseus* baseband processor documentation indicates that it cannot generate short preambles.

2.3.2 MAC Header

Chapter 7 of the specification describes the various frame formats. Each frame begins with a MAC header. The header begins with a 16-bit field defining the type of frame and several control flags, a 16-bit frame duration field, and the 6-byte MAC address of the receiver. Most frames also contain additional MAC header fields.

2.3.3 Frame Types

MPDUs can be of three types: (1) data frames that carry MSDUs, (2) control frames that coordinate the transfer of other frames, and (3) management frames are used for higher-level functions (identification, routing, etc).

Since management frames are exchanged relatively infrequently, we will assume the management functions have been completed and consider only data and control frames.

For operation under the DCF there are four control frames: ACK, RTS, CTS and PS-Poll.

2.4 802.11 MAC State Machine

The operation of the MAC protocol is best described as a state machine. The 802.11 specification describes the protocol as SDL (system description language) diagrams. For our purposes the important state machines are the processes Rx_Coordination (pp. 340–343) and Tx_Coordination_sta (pp. 344–353).

Figure 2 is a state transition diagram for a simplified version of these portions of the MAC. It combines both the receive and transmit state machines. It ignores PCF (which is not widely supported). It also ignores beacons and power-save functions since these do not affect throughput/delay performance.

The following conventions are used in Figure 2:

- rectangles with upper-case labels are states (e.g. IDLE)
- diamonds show test points for state transitions. The conventions for these are:
 - lower-case labels with parentheses are predicates (functions returning logical values) controlling state transitions

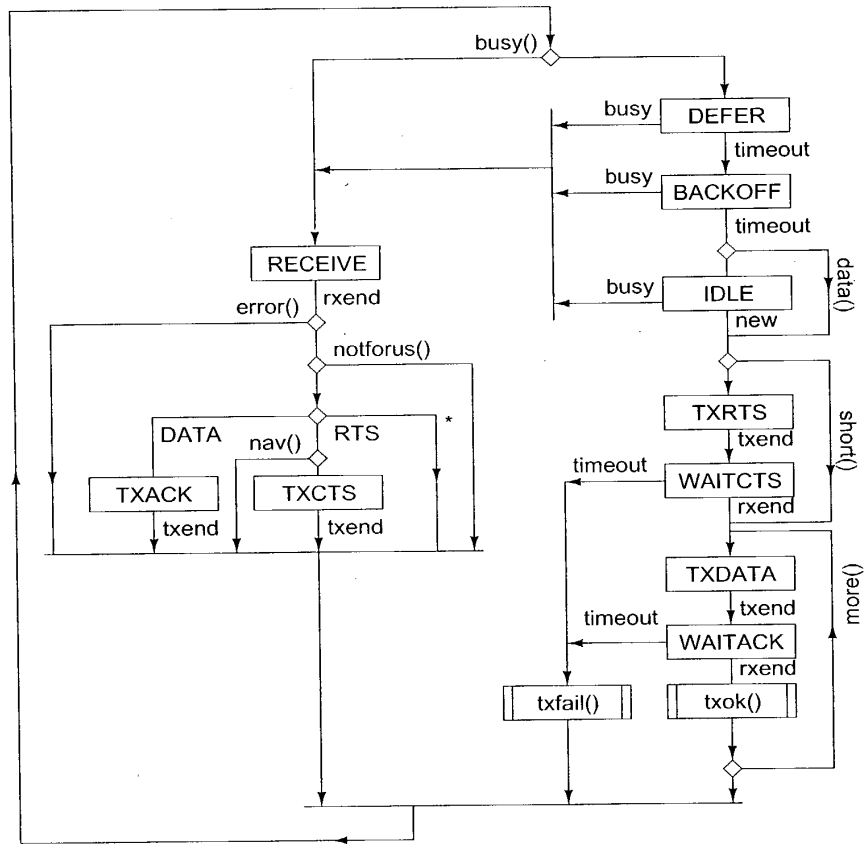


Figure 2: State transition diagram for the MAC controller.

7
B-7

- upper-case labels outside rectangles are received frame types (e.g. RTS) that label multi-way state transitions
- rectangles with lower-case function names are procedures (e.g. txok())
- lower-case labels without parentheses are PHY indications that trigger state transitions (e.g. 'timeout')

PHY indications are asynchronous events (interrupts) that terminate states. These are:

Indication	Meaning
rxend	the receiver detected the end of a frame or an error
txend	the transmitter finished sending the frame
busy	the receiver indicates the channel has gone busy
timeout	the state timeout timer has expired
new	a new frame has been queued

state timeouts depend on the state and begin when the state is entered. They are:

State	Timeout (μ s)	Notes
DEFER	20 or 364	DIFS (EIFS used if the previous frame had errors)
WAITCTS	20	response timeout (1 slot time)
WAITACK	20	response timeout (1 slot time)
BACKOFF	$20 \times$ b/o counter	if the backoff count is zero on entry to the BACK-OFF state it is set to a random value between 0 and the current value of the contention window (CW)

The predicates used are:

Predicate	Meaning
busy()	the receiver indicates the channel is busy
error()	the received frame had a CRC error
notforus()	the frame was not addressed to this station
nav()	the NAV timer has not expired
data()	there is data queued up to send
short()	the MPDU to be sent is shorter than RTSThreshold
more()	there are more fragments to be sent from the current MSDU

The state transition diagram has two processing blocks. The `txfail()` function increments a retry counter, checks that the number of retries has not been exceeded, and increases the contention window. The `txok()` function removes the bytes transmitted from the outgoing queue and resets the retry counter(s) and contention window.

The following sections give additional details about each of the states.

2.4.1 RECEIVE

This state ends if there is a PLCP CRC error, if the carrier is lost, or when the duration indicated in the PLCP header has elapsed.

If there was a CRC error the DEFER state timeout is set to EIFS, otherwise it is set to DIFS.

If the frame did not have an error and is not addressed to this station and its duration field is greater than the current NAV timer value, then the NAV timer is set to the value of the frame's duration field.

2.4.2 BACKOFF

In this state backoff counter is decremented every slot time ($20\mu\text{s}$). The backoff count is saved if this state is exited due to the channel becoming busy.

When the backoff counter decrements to zero and MSDUs are queued to transmit, the contention window and retry counts are reset.

2.4.3 Transmit States

There is a SIFS ($10\mu\text{s}$) delay before each frame is transmitted. This allows time for both stations to switch between receive and transmit.

2.5 Frame Sequences

Once a client has been authenticated and associated the only frame exchange sequences required between a DCF AP and its clients (§9.7) are:

```
{ RTS - CTS - } [ Frag - ACK - ] Last - ACK
PS-Poll - [ Frag - ACK - ] Last - ACK
```

where the - indicates a SIFS delay, RTS, CTS, ACK, and PS-Poll are the corresponding control frames, and Frag and Last are data frames. Braces ({...}) indicate frame exchanges that can happen 0 or 1 times. Brackets ([...]) indicate frame exchanges that can happen 0 or more times.

B-9

An AP uses the TIM (traffic indicator map) element of beacons to indicate the clients for which it has queued data frames. Stations use PS-Poll frames to poll for their data frames. The impact of power-save mode is not considered in this report.

3 Behavior of 802.11 Networks with Incomplete Connectivity

3.1 Connectivity Graphs

To examine the collision behavior of the 802.11 MAC protocol we can combine physical-layer aspects such as path loss, transmit power, receiver sensitivity and antenna directivity into a *connectivity graph*. Each station (AP or client) is a point in the graph and edges (lines) connect stations that can receive from each other. We assume that uplink and downlink power budgets are matched, so that each connection is reciprocal.

We also assume that collisions cause colliding frames to be lost and that the connectivity graph is fixed. In reality, receiver capture effects will allow frames to be received correctly if one signal is significantly stronger than the others. The details of the capture effect will depend on whether an individual receiver can detect these collisions and re-synchronize to a new, stronger frame. In addition, the connectivity of the network will vary over time due to fading.

Note that these connectivity graphs are purely logical constructs, they are not maps. They do *not* describe the locations of the stations, coverage regions, or any other geographical characteristics of the network.

The simplified connectivity model is used to examine the behavior of the MAC. More realistic models or actual networks must be used to estimate of the actual system performance.

A WLAN system using directional antennas reduces the connectivity of the network. As described in the following sections, this reduced connectivity increases the likelihood of collisions. It's therefore important to examine the impact of the reduced connectivity on collisions.

In the following sections we describe the causes of collisions for connectivity graph subsets of 2, 3 and 4 stations. While there may be connectivity graph subsets involving more than 4 stations that give rise to collisions, these are expected to be less likely than those involving 2, 3 and 4 stations.

3.2 Collisions in Two-Node Connectivity Graphs

The CSMA/CA protocol ensures that collisions can only happen between two stations when they choose the same backoff slot to begin their transmissions. Since




Figure 3: The three- and four-station incomplete-connectivity scenarios.

the backoff duration on the initial transmissions is chosen randomly between 0 and 31, the probability of a collision for any frame is about 1/32.

If the propagation delay between stations is longer than that allowed by the standard (about 1 microsecond) one station may not detect a frame that started in the preceding slot. This will double the two-station collision probability.

3.3 Collisions in Three-Node (“Hidden-Terminal”) Connectivity Graphs

Figure 3 shows the incomplete connectivity graph for subsets of the network that include three stations. All other three-station subsets are completely connected. Stations A and C are connected to station B only. This three-terminal incomplete connectivity graph exhibits what is known as the “hidden terminal problem.”

Since stations A and C are not connected, they do not coordinate their transmissions. A collision will happen at B if there is any overlap of the transmissions from the other two stations. The likelihood of this depends on the length of the frame relative to the inter-frame period. For long frames and short inter-frame times there will be many collisions.

This problem can be addressed through the use of the RTS/CTS control frames in the 802.11 protocol. CTS frames sent by station B will reserve the channel for the required duration and prevent collisions.

3.4 Collisions in Four-Node (“Exposed Terminal” or “Hidden-Beam”) Connectivity Graphs

Figure 3 shows the incomplete connectivity graph for a subset of the network that includes four stations. Other connectivity graphs reduce to combinations of the previous three- and two-station subsets. Stations A and D are only connected to B and C respectively. Stations B and C are also connected to each other. In addition to the hidden-terminal problem exhibited by the sub-graphs A-B-C and B-C-D, the four-terminal incomplete connectivity graph exhibits what is known as the “exposed terminal problem.”

Station D cannot hear transmissions between A and B so it is not aware of frame exchanges between them. The consequences of the exposed terminal prob-

lem depend on the use of RTS/CTS by the two leaf stations, A and D:

- If the leaf stations (A and D) use RTS/CTS, then stations B and C will set their NAV timers from the CTS frames and will not respond to RTS requests from the other leaf station. The result will be that the leaf stations will increase their contention window sizes and will retry channel access.
- If the leaf stations do not use RTS/CTS, then acknowledgment frames sent from B or C (back to A or D respectively) will cause a collision with any ongoing transmission from the other leaf station.

3.5 Handling of Broadcast Frames

An 802.11 AP uses broadcast beacon frames to distribute timing information and traffic indication map (TIM) information. The TIMs are used to support clients in power-save mode. If beacon frames were not sent then clients would not be able to synchronize their timing or use power-save mode. This would greatly reduce the battery life of portable devices and would be a significant drawback.

To comply with FCC requirements for point-to-point operation, the antenna array will be steered towards a specific client when beacon broadcast frames are transmitted. If there are clients on both sides of the center of a beam, the broadcast frames are alternately sent to the two clients that are nearest the center of the beam. If there are clients on only one side of center, then all broadcast frames are sent to the station closest to the center of the beam.

4 Behavior of 802.11 Systems with Increased Propagation Delay

4.1 Increased Collision Probability

As described above, CSMA/CA stations contending for the channel time their transmissions to begin on a slot boundary¹. Propagation delays will result in different stations having different time references for their slot boundaries.

The slot time is specific to a given PHY. It is the time required for a station to determine that a given slot is not busy and turn on its transmitter at the start of the next slot. The slot time calculation for 802.11b is shown in Table 1.

¹Frame timing boundaries are the beginning of the first symbol and the end of the last symbol.

Time	Variable	max. (μ s)
detect channel not busy	aCCATime	15
switch to transmit	aRxTxTurnaroundTime	5
propagation to other stations	aAirPropagationTime	1
perform associated processing	aMACProcessingDelay	0
Total	aSlotTime	20

Table 1: Calculation of aSlotTime for 802.11b PHY. The numbers don't add up due to an error in the 802.11 DSSS PHY specification (see §9.5.8 and §18.3.3).

4.2 ACK and CTS Timeouts

The 802.11 protocol requires positive acknowledgment in the form of an ACK control frame. After sending the last symbol of a data frame, the sender waits a maximum time of 'ACKTimeout' to detect (the end of) a correct PLCP header. A similar timeout applies to the CTS frame required in response to an RTS frame.

The 802.11 specification mentions a CTSTimeout value (§9.2.5.7) and a ACKTimeout value (in section 9.2.8), but the values are not specified in the specification.

The formal SDL description of the 802.11 protocol differs from the text of the specification. Instead of specifying the time to the end of the PLCP header, it specifies the time to the end of the ACK or CTS frame.

The SDL description of the protocol uses a timer, Tr_{sp} , to implement these timeouts. The Tr_{sp} timer is set to the duration of a SIFS time, plus the durations for an ACK frame (including the PLCP header and Preamble) plus aSlotTime. This implies that the value of ACKTimeout is one slot time (20μ s).

A 20 μ s slot time timeout would correspond to a maximum range of 3 km (round-trip time at 300 m/ μ s). This should be the maximum range of a standards-compliant 802.11 system, although it's possible that manufacturers use a longer timeout.

5 Implementing a Beamforming 802.11 AP

The design of *Little Joe* AP must provide a means for the array to be steered in the correct direction and also provide a means to achieve acceptable performance (throughput, delay) for all users in the presence of incomplete connectivity.

5.1 First Attempt - Polling with Fake Duration Field

Mabuhay originally started designing a smart-antenna product using a custom TDMA-based MAC protocol. When the company's focus shifted to 802.11 clients, mod-

ifications to the 802.11 MAC protocol were investigated. The first studies looked at a polling protocol. To keep the client from trying to access the medium between polls, the MAC header's 'duration' field was used to set the NAV to the time between polls. Unfortunately, this approach was based on the unrealistic assumption that each Mabuhay AP transmissions would only be received by one station.

5.2 Second Approach - Polling with Client Shim

An alternative approach was then developed based on installing a protocol "shim" on each client. The shim implements a polling protocol that runs on top of the 802.11 DCF protocol. In addition to the inconvenience of having to install custom software on each client and the increased delays due to polling, the 320 μ s (average) backoff delay that must precede every channel access under DCF reduces the throughput efficiency of the protocol.

5.3 Current Solution - DCF and Hardware-Driven Steering

Current plans (April, 2002) are for the Mabuhay clients and APs to use an unmodified 802.11 MAC protocol. The beam steering is done with dedicated hardware and thus requires no interaction with the MAC controller software.

Frames are classified as either "reply" (ACK and CTS) or "originated" frames (RTS, DATA). The amount of time between the end of the a received frame and the start of a transmitted frame is used to distinguish between reply and originated frames. Reply frames are transmitted a SIFS interval after the end of a received frame while originated frames are transmitted after *at least* a DIFS interval.

Reply frames can be steered using the MAC address of the immediately preceding received frame. Since it is the timing of "reply" frames that determines the maximum range of the system, the range of a *Little Joe* system is unchanged.

It is not possible to use the MAC address of the previous frame as the address for "originated" frames. Instead, *Little Joe* will "spoof" the interface between the MAC and PHY to obtain the MAC address before the transmission begins.

The hardware between the MAC and PHY could also allow for some simple manipulation of the signals to improve performance. The details are in the *Little Joe* architecture specification.

The performance of the 802.11 protocol in a *Little Joe* system can be improved by setting the 802.11 MAC parameters (802.11 defaults and vendor defaults).

The following MAC parameters, described in Annex D of the specification, can be configured on a per-station basis:

Parameter	Default	Recommended	
		at Client	at AP
RTSThreshold	2347	100	2347
ShortRetryLimit	7	7	7
LongRetryLimit	4	4	4
FragmentationThreshold	2346	2346	2346

The recommended values are based on the computer simulations described in another document.

5.4 Future Possibilities - 802.11 Polling


Better MAC performance might be obtained for clients that support polling (either PCF or 802.11e's HCF) because the AP could poll clients periodically and thus avoid the long backoff periods that result from multiple failed transmission attempts.

B-15

C

[25 pages]

Beamforming for *Little Joe*



1 Introduction

This report describes the evaluation of various beamforming options for Mabuhay Network's *Little Joe* 802.11 WLAN access point (AP).

Section 1 describes the relevant parts of the system and assumptions about propagation. Section 2 describes three practical beamforming methods that were considered. Section 3 describes how the different methods were evaluated and gives the results of the evaluation. Section 4 gives some recommendations.

1.1 *Little Joe*

Little Joe is the name for a high-performance WLAN Access Point product using a 16-element linear antenna array to increase range and performance.

It uses two separate RF beamformers. The first beamformer, the "searcher," is for receiving only and uses a 16-port Butler matrix whose outputs are connected to 16 standard WLAN cards installed in a PC. An application on the PC obtains the received signal levels for each received packet.

This signal level information is used to compute the complex weights for a second RF beamformer. This so-called "card 13" beamformer allows independent complex weights on each array element. This beamformer is connected to an additional WLAN card which is the one actually used for communication (it can transmit and receive).

Since this architecture cannot adjust the beamformer in real time it cannot cope with random-access transmissions from clients. This requires that the clients use

C-1




Figure 1: *Little Joe* block diagram.

a polling protocol rather than the standard random-access 802.11 MAC protocol. This might require installing a ‘shim’ in the client’s protocol stack.

This report deals with the task of converting the received signal strength reported by the “searcher” receivers into a set of weights for the “card 13” beam-former.

1.2 Antenna Array

The product operates at 2.4 GHz (12.5 cm wavelength) and uses element spacing of approximately half a wavelength, resulting in an aperture of 1 meter.

The first-null beamwidth of a *uniformly*-illuminated linear array of length L wavelengths is $115/L$ degrees [1, p.380]. The 8-wavelength *Little Joe* array will thus have a main lobe null-to-null beamwidth of about 14.4 degrees.

1.3 Propagation Mechanisms

The propagation environment will vary greatly depending on the application (e.g. indoor or outdoor mounting) and the specifics of a particular situation (e.g. building construction).

However, in most practical situations there will be no direct line of sight between the AP and remote antennas. In this case the primary propagation mechanisms are diffraction (from edges), scattering/reflection (from surfaces). Transmission through walls and windows will play a role in most outdoor-to-indoor situations.

The signal received by the array elements will be the vector sum of signals arriving by many paths. The resulting signal level will have a uniformly-distributed random phase and a Rayleigh-distributed random magnitude. In most situations some of the scatterers will be moving and the phases and magnitudes will change over time.

In some situations there may also be propagation by transmission through a low-loss medium (air, non-metallic walls and windows).

The path loss is typically modeled as inverse power law. Exponents between -3 and -4 are common.

1.4 Propagation Assumptions

Although the propagation environment will vary greatly, some assumptions are required to compare the performance of different beamforming techniques. In particular, the performance of will depend on:

- the direction of arrival of signals from the desired and undesired users(which, in turn, depends on the locations of scatterers)
- the ability of the *Little Joe* beamforming receivers to distinguish the desired signal from interference and noise
- the relative signal levels of the desired and interfering signals

In the case of outdoor-to-indoor propagation, signals from a source inside a building will undergo a significant amount of scattering before reaching the outside walls. It would be reasonable to assume that signals will appear to come from an area of approximately the dimensions of a typical room (about 5 to 10 meters). If the antenna is mounted 30 meters from the building, this represents an angle of arrival spread of about 20 to 40 degrees.

Noise and interference will limit the ability of the “searcher” receivers to accurately measure the received signal level. Accurate beamforming will be most critical when the SINR is low. Unfortunately it is also at this point when the measurements are least likely to be accurate. Therefore the performance of beamforming algorithms should be tested at the limit of receiver sensitivity as described in Section 3.2.

2 Beamforming Algorithms

2.1 Introduction

As described above, the system uses two independent beamformers: a Butler matrix to support the beamforming function and independent vector modulators on each array element to generate the receive and transmit beams.

The Butler matrix uses $N = 16$ array elements to form N beams. The pattern of each beam and the “boresight” angle of each beam is given in Appendix A. Each beam points in a different direction and its pattern has an approximately $\sin(x)/x$ shape. Only angles between approximately -60 and 60 degrees (relative to broadside to the array) can be used since the pattern is distorted at the extreme angles.

The beamforming application uses the signal strength information derived from packet receptions by the “searcher” receivers. However, not all of these receivers will correctly receive each packet. The beamforming algorithm computes the beamforming weights for a particular client (identified by its wireless MAC address) and stores it in a table which is made available to the application running the modified (polling) MAC protocol on “card 13”.

2.2 Non-Coherent Beamforming

Since the receivers will only provide signal level (amplitude) information and not phase information, it is not possible to set the beamformer weights for maximum SINR or even to combine the received signals coherently.

While it would be possible to synthesize an antenna pattern whose power-versus-angle distribution approximated that of the received direction of arrival (DoA) distribution (e.g. by using an inverse DFT), we cannot guarantee that this will result in coherent addition at the receiver. For example, the simple array shown in figure 2 has two elements as does the scattering field. We can use the beamformer weights to establish the relative signal levels (but not phases) at scatterers **a** and **b**. However, since we cannot control the relative phases at points **a** and **b** we cannot guarantee coherent signal addition at the receiver.

Since we are limited to non-coherent beamforming, a narrow beam will minimize interference. In an EIRP-limited system a narrow beam will minimize the required transmit power. The approach is to model the signal as coming from one discrete source and set the beamformer weights to maximize the gain for a signal arriving from this single direction. This requires the estimation of the DoA.




Figure 2: Simple scattering model.

2.3 Interpolation Algorithm

The relative attenuation in dB of the three central beams for angles of arrival between -8 and 8 degrees from boresight are given in Table 1 along with the levels relative to the strongest beam. For angles between -3 and 3 degrees the central beam is the strongest. Measuring the level of the second-strongest adjacent beam with an accuracy of about 2 dB would be sufficient to identify the direction of arrival to within 1 degree.

However, this table only applies for a single point source. If we had such a source, it would be sufficient to do a table lookup using any two columns to estimate the DoA. However, in most cases the source will have a broader DoA spread and thus we need a more robust DoA estimator.

The heuristic suggested here is to estimate the angle of arrival using a “center of mass” calculation over the three beams which are centered on the beam with the strongest signal. The signal powers are used as each beam’s “mass” and each beam’s boresight direction as its “position.”

More formally, let:

- $s(i)$ be the signal level on beam i (in watts)
- $c(i)$ be the boresight direction of beam i relative to the direction perpendicular to the array (in degrees, range from -90 to +90)

angle	Relative to Center Beam			Relative to Strongest Beam		
	Beam -1	Beam 0	Beam 1	Beam -1	Beam 0	Beam 1
-8	-0	-20	-25	0	-20	-25
-7	-0	-32	-38	0	-32	-38
-6	-0	-15	-21	0	-14	-21
-5	-1	-9	-16	0	-7	-15
-4	-3	-5	-14	0	-2	-11
-3	-5	-3	-13	-3	0	-10
-2	-9	-1	-14	-8	0	-13
-1	-16	-0	-18	-16	0	-18
0	-307	0	-308	-307	0	-308
1	-18	-0	-16	-18	0	-16
2	-14	-1	-9	-13	0	-8
3	-13	-3	-5	-10	0	-3
4	-14	-5	-3	-11	-2	0
5	-16	-9	-1	-15	-7	0
6	-21	-15	-0	-21	-14	0
7	-38	-32	-0	-38	-32	0
8	-25	-20	-0	-25	-20	0

Table 1: Adjacent beam levels (dB) versus angle of arrival (degrees).

- m is the index of the beam with the strongest signal (restricted to the range 1 to $N - 2$ to avoid end-effects)

then we estimate the direction of arrival, \hat{c} as:

$$\hat{c} = \frac{s(m-1)c(m-1) + s(m)c(m) + s(m+1)c(m+1)}{s(m-1) + s(m) + s(m+1)}$$

3 Performance Evaluation

To describe the operation and performance of this DoA estimator the signal level at the AP receiver was computed for some reasonable scattering environments and three beamforming methods:

1. known phase (array elements combined coherently)
2. using the weights used by the Butler matrix port with the strongest signal (essentially using the Butler matrix for both transmit and receive)
3. DoA estimation using the heuristic above

C-6

3.1 Scattering Model

Propagation is assumed to be NLOS with the signal arriving from many ($10N$) point sources that model the scattering field. Each point source has unit amplitude and a uniformly-distributed random phase. Scenarios were tested with the sources uniformly spread over arcs of 0, 10, 20, and 40 degrees, centered at boresight to beam 0 (broadside to array) at a distance of 30 meters.

In practical situations the DoA distribution of scatterers will not be the ‘pulse’ function used in this model. However, minor differences in the *shape* of the distribution should not have a large effect on the relative performance of the different beamforming algorithms.

3.2 Receiver Sensitivity Model

The receiver sensitivity model is meant to model a receiver operating at the edge of the coverage region where estimates of the signal strength on the different beams would be least accurate and thus the beamforming performance would be poorest.

To model a receiver operating at the coverage boundary, receiver sensitivity limitations and measurement errors are modeled by setting to zero the power received from any beam whose signal level is 20 dB or more below the power of the strongest beam. In addition, a Gaussian random number with a 1 dB variance is added to each measurement (after converting to dB) and then the result is quantized in 1 dB steps.

3.3 Beamforming Methods

The beamforming weights were computed in three ways:

- using the complex conjugates of the signals on the array elements (and setting the magnitude of each weight equal to 1) ¹
- using the as above, “boresight” angle of the beam with the strongest signal as the angle of arrival and setting the weights to the complex conjugates of the signals that would be received due to a signal in this direction
- estimating the DoA by means of the “center of mass” algorithm described above and setting the weights as in the previous method

¹This is equivalent to knowing the received phase on each array element and summing the coherently.

The first technique provides the highest received signal power and SNR if the weights are constrained to have unit-amplitude (this is known as “equal-gain” combining). The optimum SNR would result from using weights with amplitudes that were scaled according to the received power on each beam (“maximal ratio” combining). The optimum SINR would result from using weights computed by dividing² the desired signal vector by the noise-plus-interference covariance matrix[2]. None of these “optimum” approaches is possible with Little Joe because the phases of the desired signal (or noise plus interference) on the different antennas is not known.

3.4 Results

The received signal power for each case was averaged over 1000 trials using different pseudo-random combinations of scatterer phases. The table below gives the mean received signal level for the “optimum” method and the mean reduction in received power for the two other beam-steering methods described above. The results with windowing used a Hamming window. The results are repeatable to within about 0.2 dB.

The results for the scatterers centered at a direction of 90 degrees (broadside to the array) are:

Windowing Used	Scatterer Spread (degrees)	Received Power (Optimum) (dB)	Mean Degradation (Peak-Finding) (dB)	Mean Degradation (Center-of-Mass) (dB)
N	0	43.9	-0.2	-0.2
N	10	44.4	-1.7	-1.3
N	20	44.7	-2.2	-2.4
N	40	44.9	-3.6	-4.1
Y	0	35.1	-0.1	-0.0
Y	10	35.5	-0.5	-0.3
Y	20	35.8	-1.0	-0.7
Y	40	36.1	-1.9	-1.7

The simulations were repeated with the center of the scatterers centered between beams 2 and 3 (at 72 degrees). The results are:

² multiplying by the inverse.

C-8

Windowing Used	Scatterer Spread (degrees)	Received Power (Optimum) (dB)	Mean Degradation (Peak-Finding) (dB)	Mean Degradation Center-of-Mass (dB)
0	0	43.6	-2.3	-0.3
0	10	44.4	-1.4	-1.3
0	20	44.8	-2.3	-2.5
0	40	45.0	-3.5	-4.0
1	0	35.1	-0.6	-0.0
1	10	35.6	-0.6	-0.3
1	20	35.6	-1.0	-0.7
1	40	36.1	-1.9	-1.6

4 Conclusions

4.1 Discretely- versus Continuously-Steered Beams

The results show that in most cases there is little advantage to a continuously-steerable beamformer as compared to one that can be steered over a small (N) number of angles. This raises the question of whether the *Little Joe* architecture could be simplified by using the same Butler matrix for transmitting and receiving.

4.2 Interpolation Algorithm

The simple interpolation mechanism seems to work well in those cases where it makes any difference (narrow angle spread with the angle of arrival centered between two beams). The proposed center-of-mass algorithm should be sufficient if a steerable beamformer is to be used.

4.3 Beamforming with Narrow Angle Spreads

The degradation due to assuming a single DoA rather than using the 'optimum' weights is about 1 dB for the narrower angles of arrival. In these cases (e.g. WAN applications) DoA estimation will produce the same results as fully adaptive beamforming that optimizes SNR.

4.4 Windowing

WLANs do not use transmit power control. This results in a wide range of received signal levels at the AP. In an indoor-indoor or outdoor-outdoor deployment the ratio of propagation distances might easily exceed a factor of 10, resulting in path loss differences of 30 to 40 dB.

The increase in dynamic range of interference levels may will impact the performance of the system. One way to reduce this effect is through windowing.

Tapering the illumination of the array (“windowing”) reduces the sidelobe level but increases the main lobe width. The use of a Hamming window reduces the peak sidelobe level from -13 to -41 dB but doubles the main lobe null-to-null beamwidth to about 29 degrees[3, p. 250].

However, in a situation where an outdoor AP services clients in only one building, the interference power dynamic range would be significantly smaller. Thus the selection of a window function will probably depend on the expected dynamic range of the interference and windowing may be most useful in indoor-to-indoor applications. The tradeoffs involved in use of windowing might merit further study.

Windowing might also reduce the number of searcher receivers required, perhaps by half, since each beam would be wider. This would obviously reduce the cost of the product.

4.5 For Further Study

The assumptions about the DoA spread should be verified through measurements (or possibly by searching the literature). In applications with large DoA spreads, there may not be much advantage to doing beamforming *by DoA estimation*. This is not to say that a system doing optimum beamforming (using coherent receivers) could not perform significantly better.


Since interference was not modeled, the interference-rejection performance of the different beamforming methods was not compared. The current simulation code could be easily extended to study this. However, the WLAN interference environment is dynamically varying over time scales that are on the order of a frame duration. It would seem difficult to design an adaptive array that adapts to this type of interference, but I have not looked into this topic.

References

- [1] J. D. Kraus, *Antennas*. McGraw-Hill, 1950.
- [2] J. E. Hudson, *Adaptive Array Principles*. Peter Peregrinus and IEE, 1981.
- [3] A. V. Oppenheim and R. W. Schaffer, *Digital Signal Processing*. Prentice-Hall, 1975.

Appendix A - Butler Matrix Beamforming

Consider a uniformly-spaced linear antenna array with N elements spaced a distance d along the x -axis.



A unit-amplitude signal at a distance ($\gg d$) arriving at an angle ϕ relative to the x axis undergoes a phase shift per element of $\frac{2\pi}{\lambda} d \cos \phi$. The signal at the n 'th element is:

$$v(n) = e^{j\left(\frac{2\pi}{\lambda} d n \cos \phi\right)} \quad (1)$$

The Butler matrix performs an N -point FFT on the array elements to produce the N outputs:

$$\begin{aligned} V(k) &= \sum_{n=0}^{N-1} v(n) e^{-j\frac{2\pi}{N} kn} \\ &= \sum_{n=0}^{N-1} e^{j\left(\frac{2\pi}{\lambda} d n \cos \phi\right)} e^{-j\frac{2\pi}{N} kn} \end{aligned} \quad (2)$$

$$= \sum_{n=0}^{N-1} e^{j\left(\frac{2\pi}{\lambda}n\left(\frac{Nd}{\lambda}\cos\phi - k\right)\right)} \quad (3)$$

To simplify the notation, define the variable, $\psi = \frac{Nd}{\lambda}\cos\phi - k$ which combines the angle of arrival (ϕ) and matrix output port (k). Note that changing either variable has the same effect.

Using

$$\sum_{n=0}^{N-1} r^n = \frac{1 - r^N}{1 - r} \quad (4)$$

we obtain:

$$V(k, \phi) = \frac{1 - e^{j2\pi\psi}}{1 - e^{j\frac{2\pi}{N}\psi}} \quad (5)$$

A Matlab script (see Appendix B) plots the antenna pattern at each port (k) as a function of angle of arrival (ϕ). The result is given in Figure 3.

This is the array factor only. To obtain the overall antenna pattern the array factor must be multiplied by the element pattern (for *Little Joe*, a horizontal slot).

The physical angle of arrival, ϕ , varies non-linearly with the beam number, k . The output of the Butler matrix port k is a maximum when $\psi = 0$. This “boresight” angle for beam k is:

$$\phi_{\text{boresight}}(k) = \cos^{-1}\left(\frac{k\lambda}{Nd}\right)$$

where k must be shifted³ by multiples of N to lie between $-Nd/\lambda$ and Nd/λ .

Note that if the element spacing is less than $\lambda/2$ there will be no solution for the extreme values of k (it will not be possible to form beams at the edges of the array). If the spacing is more than $\lambda/2$ then several values of k may have the same boresight angle (the beams will overlap).

The boresight angles for a 16-element array with $\lambda/2$ spacing are given below:

³The DFT is a periodic function of k with period N




Figure 3: Array factor (dB) versus angle of arrival (degrees) for each port of a 16-element Butler matrix beamformer.

13

C-13

Beam Number (k)	Boresight Angle (ϕ)
0	90.00
1	82.82
2	75.52
3	67.98
4	60.00
5	51.32
6	41.41
7	28.96
8	180.00
9	151.05
10	138.59
11	128.68
12	120.00
13	112.02
14	104.48
15	97.18

Table 2: Boresight angle versus array index for $\lambda/2$ spacing.

14

C-19

Appendix B - Program Listings

Beamforming Simulation

```
function beamf(spread,usewin)

% DESCRIPTION:
%
% Compute the received power using different beamforming methods.
%
% INPUTS:
%
% see simulation variables below
%
% OUTPUTS:
%
% prints results
%
% AUTHOR/DATE:
%
% Ed Casas 2002/2/1

% radians to degrees

dtor=pi/180 ;

% wavelength (2.4 GHz)

L=3e8/2.4e9 ;

% array element spacing

d=0.5;

% number of array elements

N=16;

% center of scatterers DoA distribution (degrees) and scatterer
% angle spread
```

15

C-15

```

center=71.75;
% spread=40;

% distance to scatters (metres)

ds=30 ;

% number of scatterers

Ns=N*10;

% windows for searcher and beamformer arrays

if usewin
    swin=hamming(N)';
    bwin=hamming(N)';
else
    swin=ones(1,N) ;
    bwin=ones(1,N) ;
end

% RSSI measurement noise standard deviation and quantization step
% size (dB)

rmsd=1 ;
rmqu=1 ;

% number of trials to average over

Ntrial=1000;

% received power results

results=zeros(Ntrial,3) ;

for trial=1:Ntrial

    % generate random scatterer locations and phases

```

```

[x y s]=mksrc(center,spread,ds,Ns) ;

% compute signals received on each array element (1 by N)

r=rsig(x,y,s,N,L,d) ;

% window the received signal to reduce sidelobes

r=r.*swin ;

% find power levels at searcher receivers in dB

b=20*log10(abs(fft(r))) ;

% add 1dB variance measurement noise and round off to 1 dB steps

b=round((b+rmsd*randn(1,N))/rmqu)*rmqu ;

% find strongest beam and its level

[p k]=max(b) ;

% convert to Watts and zero beams <20 dB below peak (not
% received)

b=10.^(b/10) ;
b=b.*(b/b(k) >= 0.1) ;

% pc: 'peak beam' DoA

pc = bore(k,N,d) ;

% cc: 'center of mass' algorithm DoA

ki=rem([k-1:k+1]+N-1,N)+1 ;
cc=com(b(ki),bore(ki,N,d)) ;

% pw, cw: unit-amplitude weights for each DoA estimate

```

17

C-17

```

pw = exp(-j*2*pi*d*[0:N-1]*cos(pc)) ;
cw = exp(-j*2*pi*d*[0:N-1]*cos(cc)) ;

% ow: 'optimum' unit-amplitude weights

ow = conj(r)./abs(r) ;

% plot 'co-phased' values (testing)
% plot(r.*ow,'+') ; plot(r.*pw) ; plot(r.*cw) ; hold on;

% compute and save results: amplitude of beamformer
% outputs

results(trial,1) = abs(sum(r.*bwin.*ow)) ;
results(trial,2) = abs(sum(r.*bwin.*pw)) ;
results(trial,3) = abs(sum(r.*bwin.*cw)) ;

end
hold off ;

% convert to dB
results=20*log10(results) ;

% print optimum power and mean degradation due to DoA methods

fprintf(1, " %d & %2.0f & %.1f & %.1f & %.1f \\\n", ...
        usewin, spread, ...
        mean(results(:,1)), ...
        mean(results(:,2)-results(:,1)), ...
        mean(results(:,3)-results(:,1)) ) ;

```

Boresight Angle

```

function phi = bore(k,N,d)

% DESCRIPTION:
%
% Returns the boresight angle for beam k in the range [0,pi)

```

```

% radians.
%
% Peaks are at (k-1) = N*d*cos(phi).
%
% The array is assumed to lie along along the x-axis. All angles
% are measured counterclockwise from the x-axis.
%
% Note that Matlab uses 1-base arrays while almost everyone else
% assumes the DFT indices range from 0 to N-1. The Matlab FFT
% routines actually compute everything properly (as if the first
% element had index 0).
%
% INPUTS:
%
% k - beam index, 1 to N. Beam 1 has a peak response at pi/2.
% Increasing beam number decreases boresight angle.
%
% N - number of array elements
%
% d - array element spacing (wavelengths). default = 0.5
%
% OUTPUTS:
%
% phi - boresight angle for beam k (radians)
%
% AUTHOR/DATE:
%
% Ed Casas 2002/2/3

if nargin < 2
    error('Wrong number of arguments') ;
end

if nargin < 3
    d=0.5;
end

% response is modulo-N and acos() argument must lie [-1,1]

k = rem(k-1,N) ;

```

C-19


```

if size(m) ~= size(x)
    error('x and m must be of same size') ;
end

sm = sum(sum(m)) ;

if sm ~= 0
    c = sum(sum(m.*x)) / sm ;
else
    c = mean(mean(x)) ;
end

```

Scatterer Generation

```

function [x y s]=mksrc(c,range,d,Ns) ;

% DESCRIPTION:
%
% Generates a pseudo-random set of signal sources.
%
% The sources are arranged in an arc at a distance d from the
% origin and have uniformly distributed phases and unit
% amplitude.
%
% INPUTS:
%
% c - direction to center of arc (degrees, x-axis is 0).
%
% range - angular span of the arc (degrees)
%
% d - distance of arc from the origin (metres)
%
% Ns - number of discrete sources to generate
%
% OUTPUTS:
%
% x, y - coordinates of the sources, Ns by 1 column vector

```

C-21

```

% (metres)
%
% s - complex amplitude/phase of the sources, Ns by 1 column
% vector
%
% AUTHOR/DATE:
%
% Ed Casas 2002/2/1

% radians to degrees

dtor=pi/180 ;

% place scatterers over the given range of angles (Ns by 1)

% this code allows for range=0 degree and/or Ns=1

da=range/max([1 Ns-1]) ;
a=(c - (Ns-1)*da/2 + [0:Ns-1]'*da) * dtor ;
x=d*cos(a) ;
y=d*sin(a) ;

% plot(x,y,'+') ; % (for testing)

% generate random-phase signals for each source (Ns by 1)

s=exp(j*2*pi*rand(Ns,1)) ;

% s=exp(-j*2*pi*ones(Ns,1)) ; % (test fixed-phase)

% plot(real(s),imag(s),'*') ; % (for testing)

Received Signal

function r = rsig(x,y,s,N,L,dx)

% DESCRIPTION:
%
% Returns a row vector of the (complex baseband) signals received

```

```
% at the array elements from a number of sources. The array
% elements are assumed to be along the x axis with element zero
% at the origin.
%
% This version is for beamforming simulations and does not
% add path loss.
%
% INPUTS:
%
% x, y - coordinates of the sources, column vectors (metres)
%
% s - source signals, column vector (complex baseband)
%
% k - beam number (beam 0 direction is for phi=pi/2 - broadside)
%
% N - number of array elements
%
% L - wavelength (metres). default = 0.125 (2.4 GHz)
%
% dx - array element spacing (wavelengths). default = 0.5
%
% OUTPUTS:
%
% s - row vector of received signal
%
% AUTHOR/DATE:
%
% Ed Casas 2002/2/1

if nargin < 4
    error('Wrong number of arguments') ;
end

if nargin < 5
    L=0.125;
end

if nargin < 6
    dx=0.5;
end
```

C-23

```

if size(x,2) ~= 1 | any ( size(x) ~= size(y) ) | any ( size(x) ~= size(s) )
    error('x, y, and s must be column vectors of equal size');
end

% number of sources

Ns=size(x,1) ;

% compute x and y positions for each combination of source and
% array element (Ns by N)

ax = ones(Ns,1) * dx*[0:N-1]*L ;
ay = ones(Ns,1) * 0*[0:N-1]*L ;

x = x * ones(1,N) ;
y = y * ones(1,N) ;

% compute distance between each source and each array elements in
% wavelenghts (Ns by N)

d = sqrt ( ( x - ax ) .^ 2 + ( y - ay ) .^ 2 ) * (1/L) ;

% source signals at each array element (Ns by N)

r = s * ones(1,N) ;

% apply phase shift due to propagation delay to each signal
% source at each array element due to each source

r = r .* exp((-j*2*pi)*d) ;

% sum over all sources to get the resulting signal

if Ns > 1
    r = sum(r) ;
end

% example:
% octave:68> rsig([1;2], [3;4], [5;exp(-j*pi/3)],3,0.3,0.125)

```

CCCCCCCCCCCCCCCCCCCC

```
% ans =  
%  
% -3.9412 + 0.8244i -3.5510 - 1.8731i -1.9612 - 3.6918i  
% plot(real(r),imag(r),'+') ;
```

25

C-25

D

[15 pages]

Little Joe Link Budget



Contents

1	Introduction	2
1.1	Purpose	2
1.2	Requirements	2
1.3	Link Power Budget	2
2	Link Budget Parameters	3
2.1	AP Transmit Power: FCC EIRP Power Limits	3
2.2	Characteristics of Conventional 802.11b Equipment	4
2.3	AP Antenna Gain	4
2.3.1	Gain Reduction due to Scattering	6
2.4	Client Antenna Gain	6
2.5	Client and AP Noise Figure	7
2.6	Effect of Shadow Fading	7
2.7	Rayleigh Fading	8
3	Path Loss Models	9
3.1	Propagation in Free Space (LOS)	9
3.2	Propagation by Diffraction (NLOS)	9
3.3	Propagation by Transmission (OBS)	9
3.4	Outdoor-Indoor Path Loss	10
3.5	Indoor-Indoor Path Loss	10
4	Conclusions	11
4.1	Results	11
4.2	Spreadsheet Notes	11
4.3	Coverage Relative to Competitors' APs	13
4.4	Issue Requiring Further Study	14

1

D-1

1 Introduction

1.1 Purpose

The purpose of this report is to:

- provide a basis for making design decisions that might affect the coverage area of *Little Joe*
- provide guidance to potential customers about the improvement in coverage area they can expect from a *Little Joe* AP compared to a conventional AP

Throughout this report the term AP refers to the *Little Joe* access point and “client” refers to a standard IEEE 802.11 WLAN client card.

1.2 Requirements

The *Little Joe* AP will provide coverage to an indoor area 100m square (71m range if centrally mounted, 144m range if corner-mounted).

For indoor mounting, the AP antenna is 1m wide and 0.5 m high and can be wall- or ceiling-mounted.

For outdoor mounting, the AP antenna is 1m square and is mounted on the outside of a building or on a tower such that each building is within line of sight (LOS) and less than 200m away.

The building is assumed to be an office building with concrete/glass exterior walls and wood/gypsum interior walls.

Within the coverage area 90% of users are expected to obtain 11 Mb/s service. Within the coverage area holes should be no larger than 1 meter square. *Little Joe* may be augmented with conventional APs to fill in larger coverage holes.

1.3 Link Power Budget

A link budget is useful for evaluating design decisions. The link budget predicts the operating margin, which is the amount by which the received signal level exceeds the level required to achieve a sufficiently low error rate¹ for a large-enough fraction of users.

The basic link equation is:

$$P_R = P_T + G_T + G_R - L$$

where the variables, described in Table 1, are in dBm or dB and L represents L_I or L_O depending on the AP location.

¹In our case, a frame error rate (FER) of less than 8×10^{-2} for 1024-byte frames

P_T	transmitter power
G_T	transmitter antenna gain
G_R	receiver antenna gain
L_O ,	mean path loss for outdoor-indoor case
L_I ,	mean path loss for indoor-indoor case
M	margins for shadow and Rayleigh fading
P_R	received power
S_R	receiver sensitivity

Table 1: Link budget variables.

The operating margin, M , is the amount by which P_R exceeds the receiver sensitivity S_R :

$$M = P_R - S_R$$

This operating margin is calculated separately for the downlink (AP to client) and uplink (client to AP). Both margins must be positive for a client to obtain service. Since the path loss is time- and location-dependent (due to fading), the fraction of users within the coverage area that have positive operating margins will vary.

The link budgets use statistical models. Their purpose is to examine the sensitivity of the system performance to design changes. It is not designed to predict performance in a specific installation. Other techniques that make use of site-specific data are used for that purpose.

In Section 2 we describe each of the above parameters and identify known values. In Section 3 we suggest models for the path loss. In Section 4 we compute margins and fraction of coverage using these values and models.

2 Link Budget Parameters

2.1 AP Transmit Power: FCC EIRP Power Limits

The FCC limits transmitter power for in the unlicensed 2400 to 2483.5 MHz band to 30 dBm (1 W). In addition, the EIRP for point-to-multipoint devices is limited to 36 dBm. The EIRP for point-to-point devices is not limited, but must be reduced by 1 dB for every 3 dB of antenna gain above 6 dBi.

For example, a point-to-multipoint system using an antenna with a gain of 30 dBi would be restricted to an EIRP of 30 dBm + 30 dB - (30-6)/3 dB = 52 dBm.

The relationship between antenna gain and allowed EIRP is shown in Figure 1.

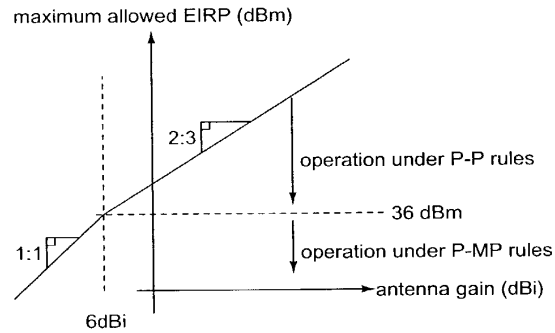


Figure 1: Diagram showing EIRP limits as per CFR 47, Part 15.247(b). Point-to-multipoint systems must operate below 36 dBm EIRP regardless of antenna gain and may not exceed 30 dBm transmitter power. Point-to-point systems may increase EIRP above 36 dBm by 2 dB for each 3 dB increase in antenna gain.

A similar reduction in transmit power is not required for point-to-point systems in the 5725–5850 MHz band.

2.2 Characteristics of Conventional 802.11b Equipment

Specifications for transmit power, receiver sensitivity and claimed indoor range were obtained from the data sheets for various manufacturer's 802.11b WLAN APs and client cards. The results are given in Table 2.

2.3 AP Antenna Gain

The AP antenna array gain has not yet been measured but an estimate can be obtained from its physical size and the antenna type. The antenna gain is related to its effective aperture by:

$$G_R = \frac{4\pi A_{\text{eff}}}{\lambda^2}$$

Assuming the A_{eff} is equal to the array's cross-sectional area (typical for a linear array with a reflector), a 1m^2 array will have a gain at 2.4 GHz ($\lambda = 0.125\text{m}$) of about 256π , or about 30 dBi.

Smaller antennas will have correspondingly smaller gains. A half-height array ($1 \times 0.5\text{m}$) would have a gain of about 27 dBi.

Make/Model	Transmit Power (dBm)	Sensitivity 1 Mb/s (dBm)	Sensitivity 11 Mb/s (dBm)	Indoor Range (m)
Orinoco World PC Card	15	-94	-82	25
Intersil Prism III			-84	37
Apple Airport	15			45
Nokia C110/C111 PC card and A032 AP	15		-84	20 – 100
Ericsson PC Card PA11	20	-90	-84	75
Intel PRO/Wireless 2011B LAN AP	18 – 20	-90	-83	30
Intel PRO/Wireless 2011B LAN PC card	14 – 18	-87	-81	30
Cisco 350	20	-94	-85	40

Table 2: Transmit power levels, receiver sensitivities and claimed indoor range for some conventional WLAN APs and client cards. Sensitivities for Orinoco and Prism III measured at 1×10^{-5} BER which is 8×10^{-2} FER for 1024-byte frame. Range estimates are at 11 Mb/s. Orinoco range estimate is for a "closed" environment. Apple range estimate is for "typical use." Ericsson range quoted for an external antenna and an "office environment".

5

DS

Increasing the frequency to 5725 MHz increases the gain of a equal-sized antenna by $(5725/2400)^2$ or about 7.5 dB (but also increases free-space loss by an equal amount).

2.3.1 Gain Reduction due to Scattering

Computations of antenna gain and sidelobe level assume a single plane wave front arriving at the antenna. In typical WLAN installations the signal will arrive via multipath scattering and there will be many angles of arrival. This will result in a reduction in gain because signals arriving directions outside the main beam will be attenuated.

The exact degree of the gain reduction depends on the antenna pattern and the angle of arrival distribution.

The resulting reduction in gain can be significant. For example, in [1] median gain reductions of 3 to 5 dB were observed for a 37 degree half-power beamwidth antenna at 3m heights in a suburban scattering environment. Since our antenna has a narrower beamwidth and the indoor environment exhibits more severe scattering, the gain reduction may be significantly larger.

On the other hand, [2] reports that in most cases over half of the energy is contained in a single narrow angle of arrival. Similar results showing several widely-separated but discrete angles of arrival are visible in the data reported in [3] and [4].

The "separability" of the paths in [2] might be accounted for by a factor of 8 difference in frequency (19 GHz/2.4 GHz, which is a difference of 64 in far-field distance) and the brick interior wall construction of the building tested that increases the contribution of diffraction (single-direction) as compared to transmission and scattering (many-direction) effects.

Another effect that will reduce the gain of the antenna when there are nearby scatterers is that the wave fronts are not well approximated by plane waves and this results in an additional loss of gain and increase in sidelobe ratio.

It should be possible to model the locations of the scatterers and compute the resulting effective pattern and gain.

2.4 Client Antenna Gain

Only Ericsson provides a gain specification for their PC client card antenna (0 dBi). This number agrees with measurements reported by others for other client cards. Omnidirectional antennas used by enterprise APs have higher gains. For example, the Nokia C950 has a gain of 2.5 dBi and the Cisco AIR-ANT3213 has a gain of 5.2 dBi.

2.5 Client and AP Noise Figure

The Intersil Prism II receiver IC specifies a noise figure of about 2 dB. Losses in the antenna switch and bandpass filter will increase this number, perhaps to 4 dB. The Orinoco “Ruby” receiver IC has a noise figure of about 5 dB.

The *Little Joe* AP uses a 1 dB NF LNA and a circulator with 1 dB loss to achieve a NF of about 2 dB.

The assumption is that a client card chipset with an external LNA will be used in *Little Joe*. The difference between typical client and predicted AP noise figures is included in the link budget.

This difference is also used when comparing the performance of the *Little Joe* AP with conventional APs.

2.6 Effect of Shadow Fading

The path loss models described in Section 3 estimate the median path loss for a given distance. However, different locations with the same path distance will have different path losses. These variations have been found to be normally distributed when the path loss is expressed in dB. The standard deviation of the path loss depends on the scattering environment but typical values for office environments are 3 to 6 dB.

We cannot increase the transmitter power to compensate for shadow fading since the system is already operating at maximum transmit power levels. Instead, the shadow fading reduces the fraction of the coverage area that can be serviced.

We will assume the coverage area is a circle or a ‘wedge’ of a circle so that we can use the equations derived for circular cells [5]:

$$a = \frac{-\gamma}{\sigma\sqrt{2}}$$

$$b = \frac{10n \log e}{\sigma\sqrt{2}}$$

$$F(\gamma) = \frac{1}{2} \left(1 - \operatorname{erf}(a) + \exp\left(\frac{1-2ab}{b^2}\right) \left[1 - \operatorname{erf}\left(\frac{1-ab}{b}\right) \right] \right)$$

where $\gamma = M$ is the link budget margin at the coverage boundary (dB), σ is the standard deviation of the fading (dB), and n is the path loss exponent. For other coverage region shapes, a different expression must be derived or the value computed through numerical integration.

7

D-7

The percentage coverage computed above is an average over the whole coverage area and it may include large coverage “holes.”

Since shadow fading is caused by objects such as walls, bookcases, doors, etc. we should expect the dimensions of the “holes” to be approximately the same as the dimensions of the shadowing objects. This is unlike the Rayleigh fading where the fades have dimensions on the order of the wavelength.

Because the average “hole” dimensions would be expected to be larger than the inter-hole spacing requirement (1 meter), achieving this requirement may require significant additional power (or reduced coverage area).

Unfortunately, I have not found published models for the spatial characteristics of indoor shadow fading. Models for fade duration versus fade level are available for Rayleigh fading. A similar analysis could be done for shadow fading if it was considered necessary to quantify the impact on the link budget of limiting the coverage hole size.

2.7 Rayleigh Fading

A margin is usually included in a link budget to counter the effect of multipath fading. For NLOS propagation this fading is Rayleigh-distributed. The probability that a Rayleigh distributed random variable r will be R dB below the mean can be approximated by:

$$P(r < R) = 10^{-R/10}$$

for $R < 0.1$. For example, the signal will be 10 dB below the mean about 10% of the time and 20 dB below the mean about 1% of the time.

For typical indoor scatterers, the duration of the Rayleigh fades (tens or hundreds of milliseconds) is slow relative to the frame duration (less than about 10 milliseconds).

The fading on antennas separated by a significant fraction of one wavelength is weakly correlated. Many clients use switching diversity to combat Rayleigh fading. The client’s receiver switches between two antennas until it finds a signal that is sufficiently strong. This squares the probability of fading (to 10^{-2R}).

Rayleigh fading affects data throughput rather than coverage. Not including any fade margin would result in the channel being unavailable about 40% of the time without diversity and about 16% of the time with two-antenna diversity. Allowing a 10 dB Rayleigh fading link margin would mean the signal was faded about 10% of the time without diversity and about 1% of the time with two-antenna diversity.

The effect of Rayleigh fading on data throughput is difficult to compute because of complex interactions between frame loss and contention-control mechanisms in the 802.11 MAC and congestion-control mechanisms in TCP/IP.

3 Path Loss Models

This section describes the models used to predict the outdoor-indoor and indoor-indoor path loss. Descriptions and experimental validation of these path loss models can be found in [6, 7, 8].

3.1 Propagation in Free Space (LOS)

Loss in free space is given by:

$$L_{FS}(d) = 20 \log \left(\frac{4\pi d}{\lambda} \right)$$

in the far field ($d > 2D^2/\lambda$, $d \gg D$ and $d \gg \lambda$ where D is the largest dimension of the antenna).

3.2 Propagation by Diffraction (NLOS)

For NLOS (non line-of-sight) paths, propagation is mainly by diffraction and a power-law path loss formula is a good model:

$$L(d) = L_{FS}(d_0) + 10n \log(d/d_0)$$

where d_0 is a reference distance (typically 1m) and n is a value that depends on the geometry of the paths. For indoor NLOS paths n is typically between 3 and 4.

3.3 Propagation by Transmission (OBS)

When propagation is mainly by transmission through walls for floors, a simple model is to modify the free-space path loss with a wall or floor attenuation factor for each penetrated wall and/or floor:

$$L_{OBS}(d) = L_{FS}(d) + p_{wall}W_{wall} + p_{floor}W_{floor}$$

where p_{wall} and p_{floor} are the number of penetrated walls and floors respectively and W_{wall} and W_{floor} attenuation constants that depend on the construction materials.

9
D-9

For same-floor propagation through several walls, this model can be simplified by substituting a constant attenuation per unit distance:

$$L_{OBS}(d) = L_{FS}(d) + \alpha d$$

where α is the attenuation per meter. A typical value for α is 0.6 dB/m (resulting, for example, from W_{wall} of 4 dB per wall and a wall every 6 or 7 meters).

3.4 Outdoor-Indoor Path Loss

A model for outdoor-to-indoor propagation [9] and [10, section 4.2.9] combines the OBS and LOS models and a correction for the angle of incidence:

$$L_o(d) = L_{FS}(S+d) + W_e + WG_e \left(1 - \frac{D}{S}\right)^2 + \max(\Gamma_1, \Gamma_2)$$

where S , d , and D are in meters (see Figure 2), W_e is a constant related to the external wall construction (about 7dB for concrete walls with unshielded windows) and WG_e is a similar constant for shallow angles of incidence (20 dB).

The loss inside the building is modeled using:

$$\Gamma_1 = W_i p$$

or

$$\Gamma_2 = \alpha \cdot (d-2) \left(1 - \frac{D}{S}\right)^2$$

where W_i is the additional loss per interior wall (4 dB per wall), p is the number of walls passed, and α is the attenuation constant, about 0.6 dB/m.

If there is significant building penetration from several directions, the signal levels from each direction should be computed separately and summed.

Note that the building penetration loss (W_e) may be 10 to 20 dB higher for buildings with with low-emissivity (energy-efficient, “low-E”) windows that have metallic coatings.

3.5 Indoor-Indoor Path Loss

The power-law NLOS model given above is usually used to model indoor-to-indoor path loss.




Figure 2: Model for COST 231 building penetration loss model. From [10].

4 Conclusions

4.1 Results

The following table gives some estimates of the best-case, worst-case and best-guess values of the variables in the link budget.

The very large range of possible values makes it clear that the system will not meet requirements under all conditions. It is probably more useful to see how *Little Joe* performs relative to competing products as shown in the next section.

4.2 Spreadsheet Notes

An accompanying spreadsheet (`linkbudget.xls`) can be used to compute the operating link margin and percentage coverage on both the uplink and downlink.

The following notes explain the meanings of some of the items in the spreadsheet. On the cells with heavy borders should be modified. Other cells contain computed values.

free-space gain measured relative to the Butler matrix beam-side port. This is the free-space gain as measured at an antenna range. It will include the array factor, the element pattern, efficiency and any apparent losses due to phase/amplitude errors in the vector modulator.

The gain may be different on receive and transmit due to differences in the ability of the vector modulators to LNA and PA phase/amplitude errors.

effective aperture computed from the gain to give an approximate indication of the physical antenna size required. It is not used elsewhere in the calculations.

Variable	Best Case	Worst Case	Best-Guess	Units
frequency	2.4	5.7	2.4	GHz
AP free-space gain	30	20	24	dBi
AP antenna gain reduction	0	20	10	dB
distance	50	200	100	m
path loss exponent	3	4	3.5	-
AP transmit power	22	25	24	dBm
client transmit power	20	15	15	dBm
client antenna gain	3	-3	0	dB
receiver sensitivity	-85	-81	-84	dBm
AP LNA NF advantage	4	0	2	dB
path loss std. dev.	3	6	4	dB
downlink operating margin	49	-36	12	dB
uplink operating margin	51	-46	5	dB
downlink percentage coverage	100	2	100	%
uplink percentage coverage	100	1	98	%

Table 3: Link budget values used and results.

straddle (crossover) loss the worst-case antenna gain after beam selection and pointing. This will typically be at the edge of the outermost beam.

gain reduction due to scattering see 2.3.1. Still not quantified.

minimum antenna gain the free-space gain reduced by the above two effects.

perpendicular distance (D), outdoor distance (S) see Figure 2.

indoor distance(d) note that d here is the *indoor* distance only, not the total path distance.

Choose the Path Loss Type the spreadsheet can compute the link margin for *either* the Indoor-Indoor or Outdoor-Indoor paths. Enter 0 or 1 here to select the appropriate path loss model to use.

transmit power (downlink) the AP transmit power. It is automatically filled in as the maximum allowed by the FCC EIRP limits based on the free-space antenna gain (see Section 2.1).

transmit power (uplink) the client transmit power. (see table 2).

receive antenna gain this is the client antenna gain. Also used as the uplink transmit antenna gain.

receiver sensitivity for a given bit rate (typically 11 Mb/s). See Table 2 for other values.

LNA NF improvement the improvement in noise figure resulting from using an LNA ahead of the WLAN PC card with the receiver sensitivity stated above. It is the difference between the LNA and WLAN PC card noise figures (assuming an LNA with sufficiently high gain).

operating margin the amount by which the mean received signal level exceeds the sensitivity at the given distance (d).

path loss std. dev. standard deviation of the shadow fading. See Section 2.6.

coverage at boundary fraction of locations at the given distance (d) with positive operating margin. This value may be more relevant as a measure of coverage than the next.

coverage within cell fraction of locations within the complete coverage area ("cell") with positive operating margin.

4.3 Coverage Relative to Competitors' APs

For comparison purposes, the improvements in antenna gain, receiver noise figure and transmitter power are sufficient to estimate the increase in SNR and thus the increase in range of *Little Joe* over conventional APs. The computation is as follows:

Conventional 802.11 WLAN APs and clients operate at transmit power levels of up to 20 dBm with 6 and 0 dBi antennas respectively, resulting in EIRPs of 26 and 20 dBm. Typical receiver noise figures are about 4 dB.

Little Joe is designed to provide increased range due to: additional antenna gain (up to 30 dBi), lower noise figure (by about 2 dB), and additional transmitter power (the maximum allowed by the FCC Part 15 rules for point-to-point operation, 22 dBm for a 30 dBi gain antenna).

On the downlink, *Little Joe* realizes an EIRP advantage of 26 dB compared to a conventional AP:

	Conventional AP	<i>Little Joe</i> AP	Improvement
AP transmit power (dBm)	20	22	2
AP antenna gain (dBi)	6	30	24
EIRP (dBm)	26	52	26

On the uplink, *Little Joe* also realizes an advantage of 26 dB due to the 24 dB additional antenna gain and the 2 dB lower noise figure:

	Conventional AP	<i>Little Joe</i> AP	Improvement
AP antenna gain (dBi)	6	30	24
AP noise figure (dB)	6	4	2
net improvement (dB)	0	26	26

The attenuation, measured in dB, for indoor environments increases linearly with distance. Typically the attenuation increases 3 to 4 times as fast as the distance. Thus an EIRP increase of 26 dB results in an increase in range of about 7.5 “dB” (about 5.5 times).

The value for the AP antenna gain does not take into account gain reductions due to antenna efficiency, scattering, or pointing errors since these have not yet been quantified. A reduction of 10 dB in gain would translate into a loss of about 3 dB in distance (about 2 times).

4.4 Issue Requiring Further Study

As mentioned earlier, an important issue is the reduction in antenna array gain that may result due to nearby scatterers.

It should be pointed out that the Part 15 FCC regulations require a reduction in transmit power by 1 dB for each 3 dB of gain above 6 dBi. This gives rise to the paradoxical result that, in the case of uniform scattering field, a *Little Joe* system using a 30 dBi antenna will be 8 dB $((30-6)/3)$ worse on the downlink than a simple system using an omnidirectional 6 dBi antenna!

Another issue is the current lack of space diversity at the AP. The resulting reduction in throughput is not captured in the link budget.

References

- [1] L. Greenstein and V. Erceg, “Gain reductions due to scatter on wireless paths with directional antennas,” *IEEE Communications Letters*, vol. 3, no. 6, pp. 169–71, June 1999.

D-14

- [2] M. J. Gans, R. A. Valenzuela, J. H. Winters, and M. J. Carloni, "High data rate indoor wireless communications using antenna arrays," in *Proceedings of 6th International Symposium on Personal, Indoor and Mobile Radio Communications*, pp. 1040–6 vol.3, 1995.
- [3] G. German, Q. Spencer, L. Swindlehurst, and R. Valenzuela, "Wireless indoor channel modeling: statistical agreement of ray tracing simulations and channel sounding measurements," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings*, pp. 2501–4 vol.4, 2001.
- [4] J.-G. Wang, A. Mohan, and T. Aubrey, "Angles-of-arrival of multipath signals in indoor environments," in *Proceedings of Vehicular Technology Conference - VTC*, pp. 155–9 vol.1, 1996.
- [5] W. C. Jakes, ed., *Microwave mobile communications*. Wiley, 1974.
- [6] J. Kivinen, X. Zhao, and P. Vainikainen, "Empirical characterization of wide-band indoor radio channel at 5.3 GHz," *IEEE Transactions on Antennas and Propagation*, vol. 49, no. 8, pp. 1192–203, Aug. 2001.
- [7] J. Medbo and J.-E. Berg, "Simple and accurate path loss modeling at 5 GHz in indoor environments with corridors," in *Vehicular Technology Conference Fall 2000. IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference*, pp. 30–6 vol.1, 2000.
- [8] H. Hashemi, "The indoor radio propagation channel," *Proceedings of the IEEE*, vol. 81, no. 7, pp. 943–68, July 1993.
- [9] J.-E. Berg, "Building penetration loss along urban street microcells," in *Proceedings of PIMRC '96 - 7th International Symposium on Personal, Indoor, and Mobile Communications*, pp. 795–7 vol.3, 1996.
- [10] E. Damosso and L. Correia, eds., *COST 231 Final Report – Digital Mobile Radio – Towards Future Generation Systems*. European Commission, Directorate General XIII, 1999. Report number EUR 18957 (ISBN 92-828-5416-7).

[REDACTED]

SimpleMAC

[REDACTED]

[1 page]

[REDACTED]

[x] [Home]

E

Little Joe MAC Operation

The LittleJoe architecture has one WLAN AP card per beam. This allows each card to operate with its own independent MAC controller.

The MerlinFPGA allows each MAC to detect transmissions (via the spoofed CCA interface) and also to "hear" transmissions from other MACs (via the spoofed RX interface) even when the cards are operating on different channels. If LittleJoe does not use the MerlinFPGA (because it uses MerlinII or does no "fine" steering at all) then DP1+ performance would suffer (CTS transmissions do not propagate NAV values to the other APs).

All of the APs listen simultaneously but only one of them can transmit at a time. Each card's MAC protocol enforces the "one transmitter at a time" rule through the use of it's CCA circuitry.

Any transmission from any AP on any beam will interfere with and cause the loss of any frames currently being received on other beams. However, since the APs on different beams can hear each other's transmissions, there should be no collisions after the initial downlink DATA or CTS frame (as a result of other beams setting their NAV timers).

No special MAC software is required at the client. Each beam appears as a different AP and a client associates with whichever beam it thinks is providing best coverage.

RTS/CTS on Uplink

On the uplink, this "opportunistic" behavior gives preference to short frames unless uplink RTS/CTS is used (the CTS will be received by the other AP beams, they will obey their NAV timers and avoid interfering with ongoing uplink reception).

This is deemed as a significant issue. We have decided to set the RTS/CTS threshold to a small value (TBD) to enable the use of RTS/CTS on the uplink.

Traffic Shaping

To ensure fairness in access to the AP [Downlink Traffic Shaping] is suggested.

E-1

[REDACTED]


[REDACTED]

F

[17 pages]

Prototype Story Board

Mabuhay Networks, Inc.




F-1



60423660 . 110402

What is the prototype composed of?


- Spatial Division Multiple Access Demo
- Beam Steering Demo
- Range vs. Bit Rate Demo



F-2

Components Required for the Prototype

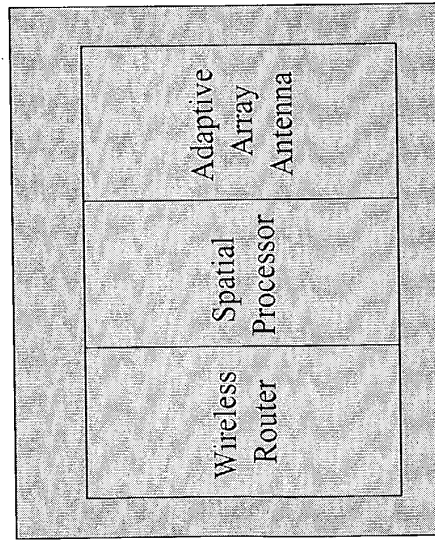

- 2 - Laptops using 802.11b NIC cards acting as subscriber stations. (CPE's)
- 2 - WebCAM USB cameras.(Connected to CPE's)
- 1 - Wireless Two Channel Router (All Demo's)
- 1 – 2 Element Adaptive Array Antenna (Beam Steering Demo)
- 1 - Spatial Processing Unit (Beam Steering Demo)
- 1 – 8 Element Adaptive Array Antenna (Distance Demo)




F-3

60423666 . 110462

Block Diagram of Prototype




Base station



F-9

Spatial Division Multiple Access using an Adaptive Array Antenna


- Purpose
 - To demonstrate that it is possible on the same channel to create multiple simultaneous beams from an adaptive array antenna that have independent data streams.
- Theory/Concepts
 - Adaptive array technology uses digital base-band processing interface to form RF wave fronts. These wave fronts will be shaped to form two beams that are separated by a given angle in which the data streams are independent.
 - Two data links will be established on the same channel and the bit rate and latency will be measured.



F-5

Spatial Division Multiple Access using an Adaptive Array Antenna (Cont.)


- Demo
 - The prototype will demonstrate that independent live video feeds from the remote stations can be established with the following highlights;
 - Two simultaneous, independent transmitting beams coming from the CPE on the same channel are used to feed both streams.
 - The RF received at the base station will be spatially separated into two streams from the same channel and decoded into separate 802.11b data streams.
 - Two CPE's will have camera demonstrating live video feeds to the base station.





F-6


Spatial Division Multiple Access Demo

60423660.110402






F-9



Beam Steering & Null Canceling Interference Demo


- Purpose
 - To demonstrate that it is possible to steer beams electronically to maximize the bit rate and range.
- Theory/Concepts
 - As a signal to noise ratio is improved on a fixed channel, it is possible to increase the bit rate for a given communications link.
 - It is the goal of Mubuhay Networks Engineering to be able to steer packets electronically on a packet by packet basis. This will allow for maximizing signal to noise ratio to a given subscriber.



1-9

Beam Steering & Null Interference Cancelation (Cont.)




- Demo
 - Measure the bit rate of one of the beams at position 2.
 - Steer the beam to position 3.
 - One of the beams will be moved electronically to position 3 which will null out the CPE at position 2.
 - Walk to position 3 with the CPE and establish connection
 - Measure the bit rate.



F-10

6243660 . 140402

Beam Steering & Nulling Demo




Fall



60423660 . 2.1.0402

Range vs. Bit Rate Demo


- Purpose
 - To demonstrate that by increasing the gain of the antenna at the base station to be similar to Mabuhay's future products, one can deduce that 1 Mb/s for a kilometer could be provided (for indoor use) using conventional 802.11b NIC cards at the subscriber unit.
- Theory/Concepts
 - Mabuhay's future products will have a 16 element adaptive array antenna with a gain of 23 dB. Due to reciprocity, the antenna gain will aid reception as well as transmission.
 - Because Mabuhay can steer a beam to a specific location, Mabuhay's product will be able to use the ISM point-to-point rules between the base station and subscriber unit. This power advantage allows for the use of a direct, narrow beam at longer range with enough power to penetrate walls at a high bit rate.





1-12

Range vs. Bit Rate Demo (Cont.)


- Demo
 - Given a 23 dB adaptive array antenna on a rotor, the demo will show that adequate bit rate (1 Mb/s) between the base station and an indoor client can be established at 1 kilometer away.
 - The adaptive array antenna will be connected to the base station's wireless router. A remote Laptop with an 802.11b NIC card will be used to establish the connection. A "Real Player" application showing a movie will be used to demonstrate both bit rate and latency.



F-13




Base Station



F-19


Example of Base Station User Interface

© 1999 Cisco Systems, Inc.




F-15

Base Station User Interface




Steer antenna



K-10

Additional Prototype Material

- Propagation Study
 - Range vs. Bit Rate
 - Understanding of different terrain in regards to bit rate
- Simulation of System with multiple beams to determine
 - Capacity
 - Throughput
 - Latency




K-17

0423660 . 110402

[15 pages]

G

Uses diversity (antenna select) to switch beams in 4° steps.



The normal diversity function chooses the antenna with the best RSSI (Received Signal Strength Indication). The same RSSI could be used as a basis to choose which one of two adjacent butler ports to use for a particular card.

This approach allows beams to be steered with approximately 5 bits of resolution

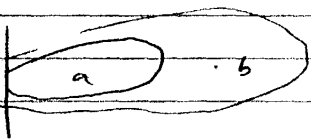
G-1

Power Control

Interference is minimized when only the minimum amount of power needed to establish a link is used. The approach on pages 1 and 2 needs a way to control the power of each packet.

Power control with CSMA in a beam formed system:

In a beam-formed system, power control can have the effect that a client unit in b can miss a transmission from a, incorrectly detecting a clear channel.




→ Can the rate determining algorithm be used in lieu of power control?

— Can the system described in the previous page be used without power control? can it be approved by the FCC and other regulatory agencies?

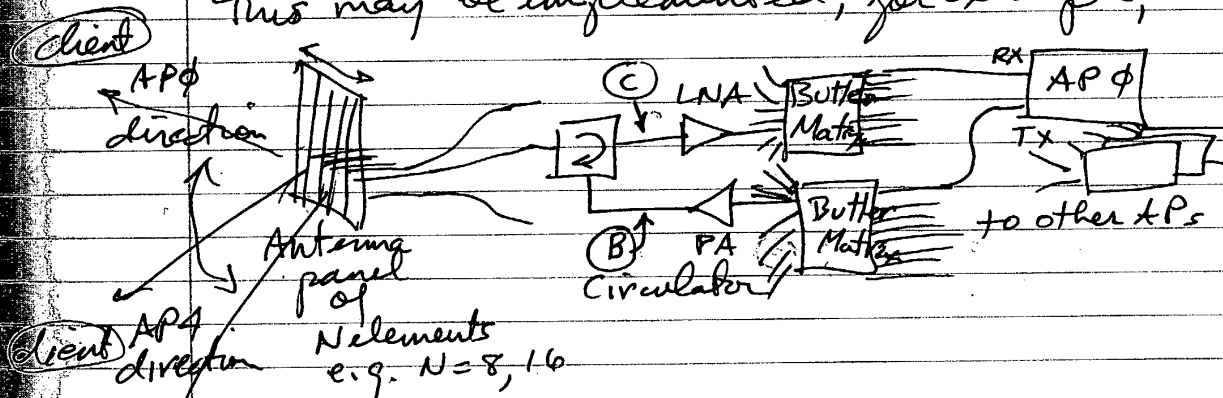
G-2

Adaptive Windowing

It is possible to build a system without knowledge of downlink (MAC-recipient address) addressee. The Butler Matrix system whereby a particular Butler port is chosen to transmit to a group of clients (one at a time) is used. For example:



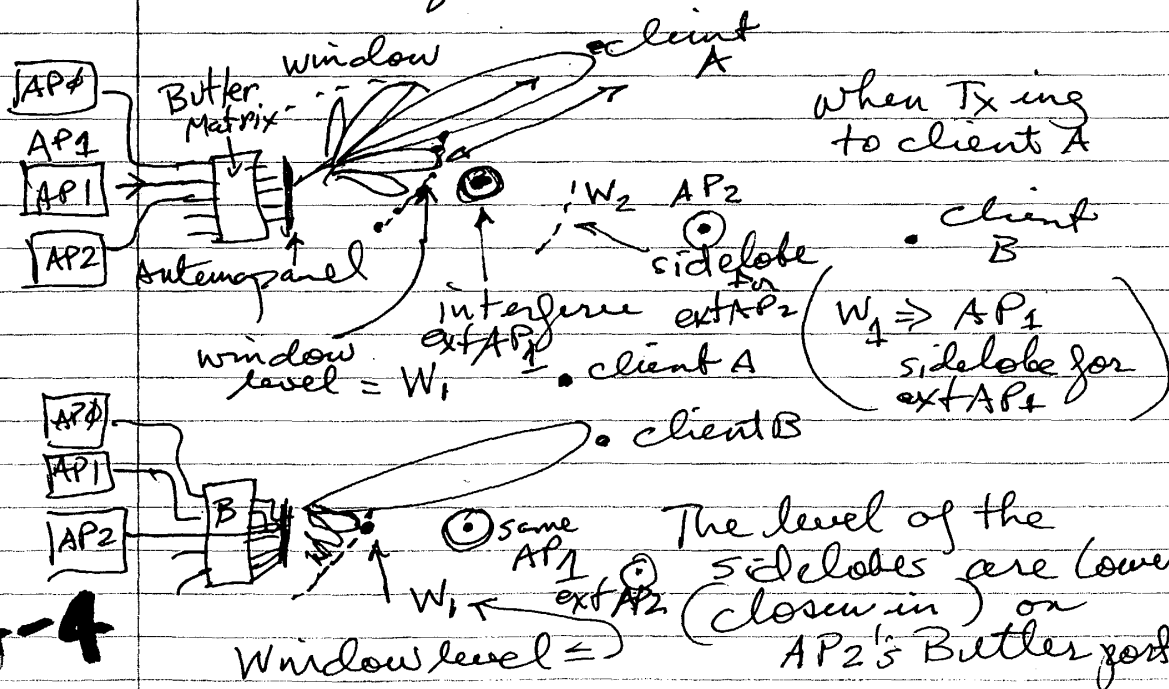
This may be implemented, for example, as:



The improved way to accomplish downlink without knowing steering information is to window adaptively on a port by port basis. For example, assume that we wish to steer a beam to a particular client. Normally we would place a vector multiplier in the path at (B) for transmit and/or (C) for receive. The presence of these vector multipliers

6-3


allow precise steering and windowing of transmissions to clients. However, knowing the values to apply to the vector multipliers is problematic because certain packets are transmitted quickly to a client before the knowledge of the MAC address of a particular client. For example ACKs are transmitted sometimes by baseband ICs before the intended MAC address of the prior uplink packet is made known to beam-steering software. The solution to this problem is to design adaptively windowed beams based on the interference environment, rather than a particular client MAC address. For example:



G-4


~~_____~~
~~_____~~ A7

Adaptive Windowing applied to the beam is learned by examining the interference environment over a period of time. The APs in the antenna panel (switch) system receive signal strength information of the various beams, over time and the minimization of interference to the ext APs (external APs). The signal strength information of interferences is used to establish window levels for the various Butler ports chosen. For example Butler port ~~1~~ ~~trying to client A~~ Butler port 1, trying to client A, requires a window to reduce sidelobes to level W_1 for ext AP₁ and W_2 for ext AP₂. Butler port 2 requires the same W_1 and W_2 values, but in a different offset from the center of the beam, due to the ~~large~~ difference in pointing angle of port 1 and port 2. Thus the windows and subsequent patterns of port 1 and port 2 will be different. For example, as shown below:




G-5

A high sidelobe is allowable for port 1, providing that the W_1 and W_2 sidelobe levels are met. For example the default sidelobe level may be -13 dB allowing uniform windows and maximum forward gain. In this case (uniform windows) SNR may be improved to about A, reducing the time of transmission, reducing chance of retransmission, and thus reducing interference. However for port 2, a low close-in sidelobe is required at W_1 region, and slightly higher at W_2 . Therefore the window will be different for port 2 than port 1. The beam may be wider. The beams windows may be calculated using an all-zero filter or other filter synthesis techniques. For example a cascaded zero filter (FIR length = A of elements) may be synthesized to match the spatial filter developed by W_1 and W_2 and the default sidelobe. The impulse response is then calculated to determine the window coefficients. The block diagram is thus:



G-6

An interference-measuring receiver is placed in the system as follows:



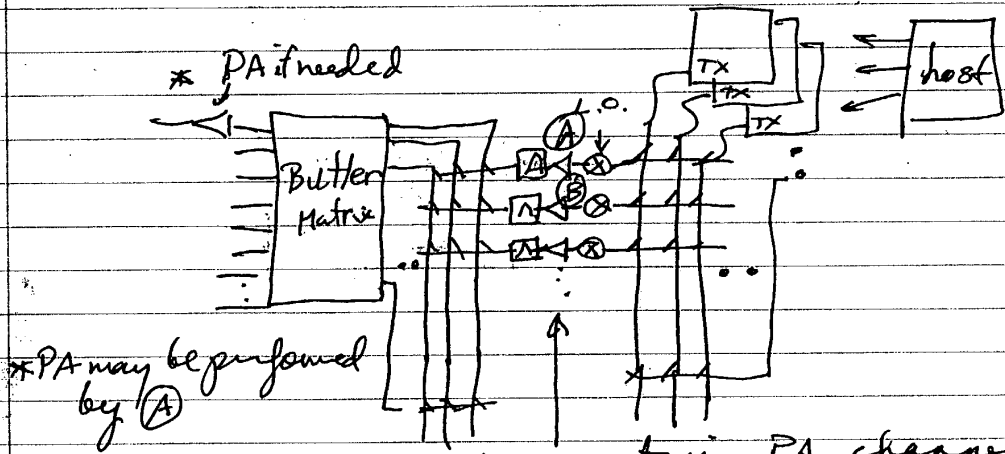
The TX complex window is chosen based on the TX port that has a signal present. The RX complex window is determined based on a best overall window and an interference radio (optional to control RX window quickly). The interference measuring radio or radios is built using a simultaneous monitor made in the 802.11 APs, or with a separate physical radio in receive-only operation. The table of windows is determined by the signals received by this radio.

G-7

Butler Systems not needing CAL

A system for Butler Matrix RX, Butler Matrix TX using APs or AP-like functions was described on page 45. In cases where adaptive windowing is not necessary, and vector multiplier cal type circuits are not necessary, then the block diagram on page 45 (lower) results.

In this case, the TX port of each "AP" is connected through a switch matrix, possibly crosspoint to the TX Butler Matrix as follows.




may contain PA, channel filter, upconversion, may be channelized, i.e. (A) carries only one channel, settable, (B) carries another, etc.

G-8

The use of common hardware in the path between the crosspoint switches allows the use of less upconversion, amplifier and filtering hardware.

Alternatively, TX connections may be made

directly to Butler Matrix ports, possibly through amplifiers and filters only. For example:



The filters may be used to reduce out of band emissions, spectral regrowth, adjacent channel power, etc.

A scanning receiver is used to determine the appropriate AP to assign a client to. Therefore the Txs may remain on somewhat static connections to the Butler Matrix. Alternatively, hand off algorithms may be used which result in no scan receiver being required.


In some APs or AP-like junctions, a "monitor" or "promiscuous" mode is available. Such a mode allows all packets received to be reported to a host. In this case, a scan receiver is also not needed.

The scan receiver connects to the Butler Matrix receive system, selectively by port, to acquire a data base of client signal levels, and interference from the various directions in the antenna's field of view. Potentially new associated clients may also be searched for.


6-9

Little Joe Protocol

Problem: Build a low cost Base station that supports 802.11 with a single beam created by a phased array antenna.



STATE machine of SEARCHER & Transceiver



There is a coupling between the transceiver and the searcher. When

transceiver is transmitting, the searcher can either be idle or calculate coefficients. When transceiver is receiving, the searcher goes into discovery mode.

G-10

Synchronization of MAC Controller


[redacted] when used with a vector modulator

Problem: When using a system that has multiple radios on receive, [redacted] combined with a directional antennas has the problem of synchronization in which a protocol requires a transmit at the same time that a receive signal [redacted] from another station needs to be serviced. In a CSMA/CA protocol the RTS/CTS reservation scheme is typically used as a solution to solve the well known near far problem.


Proposal:

[redacted] To use the RTS/CTS scheme in which a CTS frame is sent to the ~~to~~ client stations as an ~~to~~ omni direction beam covering the complete direction of the antenna ~~with~~ with the ~~power~~ maximum power of 4 watts EIRP. [redacted]

[redacted] The CTS signal generated by the system will cause the ~~other~~ other MAC controllers to [redacted] not transmit until the end of the reserved time from the CTS frame.




G-11




Area covered by CTS frame sent at 1 or 2 Mbits/sec at 4 watts EIRP

An ~~_____~~ additional mechanism to prevent ~~_____~~ the basestation ~~_____~~ with multiple radios to from interfering with each other is to "OR" the Clear Channel Assessment (CCA) ~~_____~~ or Busy line.




In addition, ~~_____~~ Each of the ~~_____~~ CCA lines coming from the basestation to the "OR" gate has a timer associated with it. ~~_____~~ The timer is triggered the the when the line is set to busy. ~~_____~~ If the duration of the timer is set to be maximum packet duration. ~~_____~~ If the timer expires, the line will change to not busy.

G-12





The State machine for S_1 to S_N assuming each is independent.



out₁ is "ORed" with other out₁'s.

out₂ is "ORed" with a₀ and sent to MAC CCA

The state machine of S is as follows,



G-23

Proxie Beam to enable Roaming

A scheme to ^{provide a} ~~provide a~~ ^{mechanism} ~~mechanism~~ for a ~~multi-point~~ multi-point to-point systems

Problem: When a station moves outside of the optimal beam, it normally ~~trans~~ ^{beams} would reach a beacon roaming SNR ^{threshold} to begin its scanning. Unfortunately, ~~with~~ ^{with} a high gain antenna, it is possible that a station doesn't roam when it should. This means there are cases where a beam is pointing in ~~an~~ an incorrect direction. Since other stations may be dependent on a common beacon, it is not feasible to stop beconing to force Roaming ~~since~~ since every station will then begin active roaming. There needs to be a mechanism that forces a particular station to roam from the Packet switch's control. ~~to force a station to connect to the most appropriate beam.~~

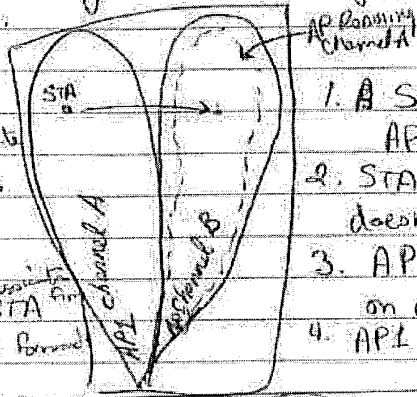
Solution:

1. A roaming AP RADIO begins to beacon at the correct port/beam with the same SSID as the station to be moved is associated with. Also, the beacon will be sent on the same frequency. This
2. The station to be moved will be disassociated on its current beam by the packet switch.
3. Only the station to be moved is allowed to associate to the roaming AP.
4. Beacons will be initiated ~~or~~ or will already be sent by the target AP port on the packet switch.


G-14

- 5. ~~Beaconing~~ Beaconing will stop on the Roaming AP radio.
- 6. The ~~Roaming~~ Request will only be accepted by the target AP port.
- 7. The Roaming AP Radio goes back to scanning mode.

- 5. AP Roaming only accept STA and lets it roam.
- 6. AP Roaming Stops beaconing
- 7. AP2 is only in ^{receiving} AP to accept STA
- 8. AP Roaming port forward sending traffic.



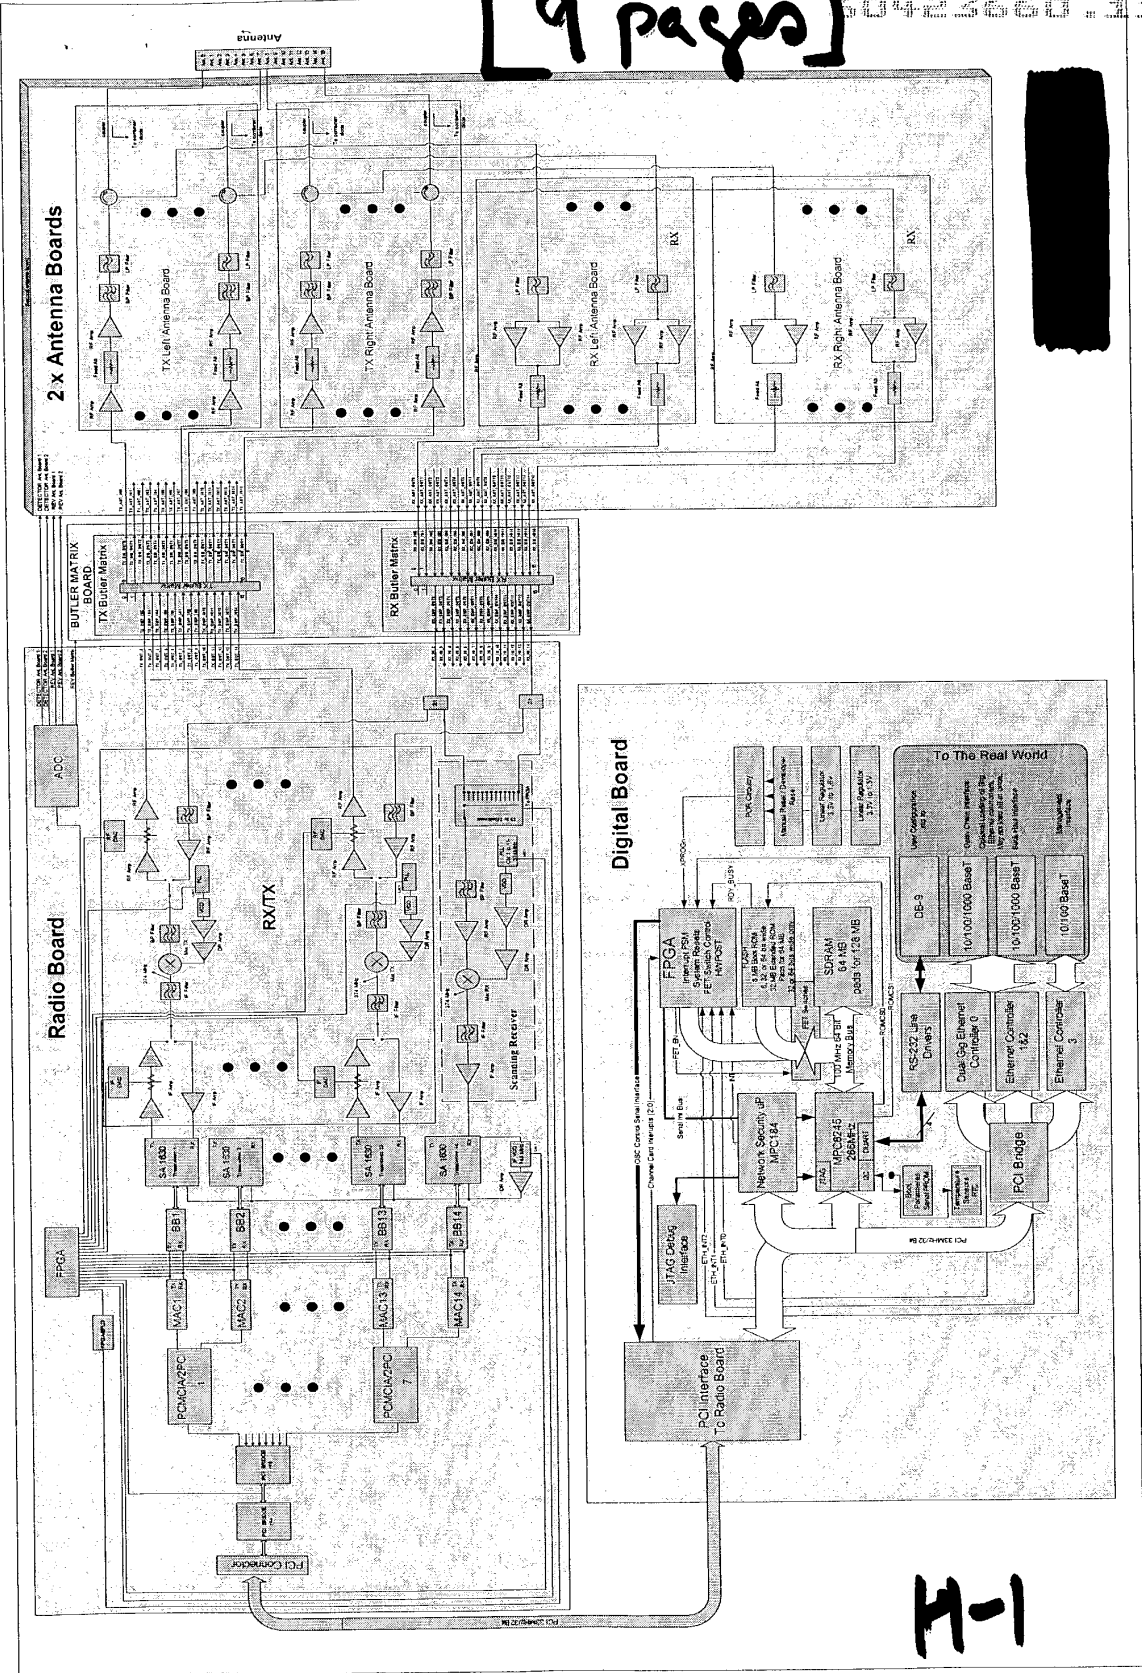
- 1. A STA is associated to AP1.
- 2. STA moves to AP2 beam, but doesn't roam.
- 3. AP Roaming creates beacon on channel A
- 4. AP1 disassociates STA




G-15

[9 pages]


H




H-1




H-2




H-3

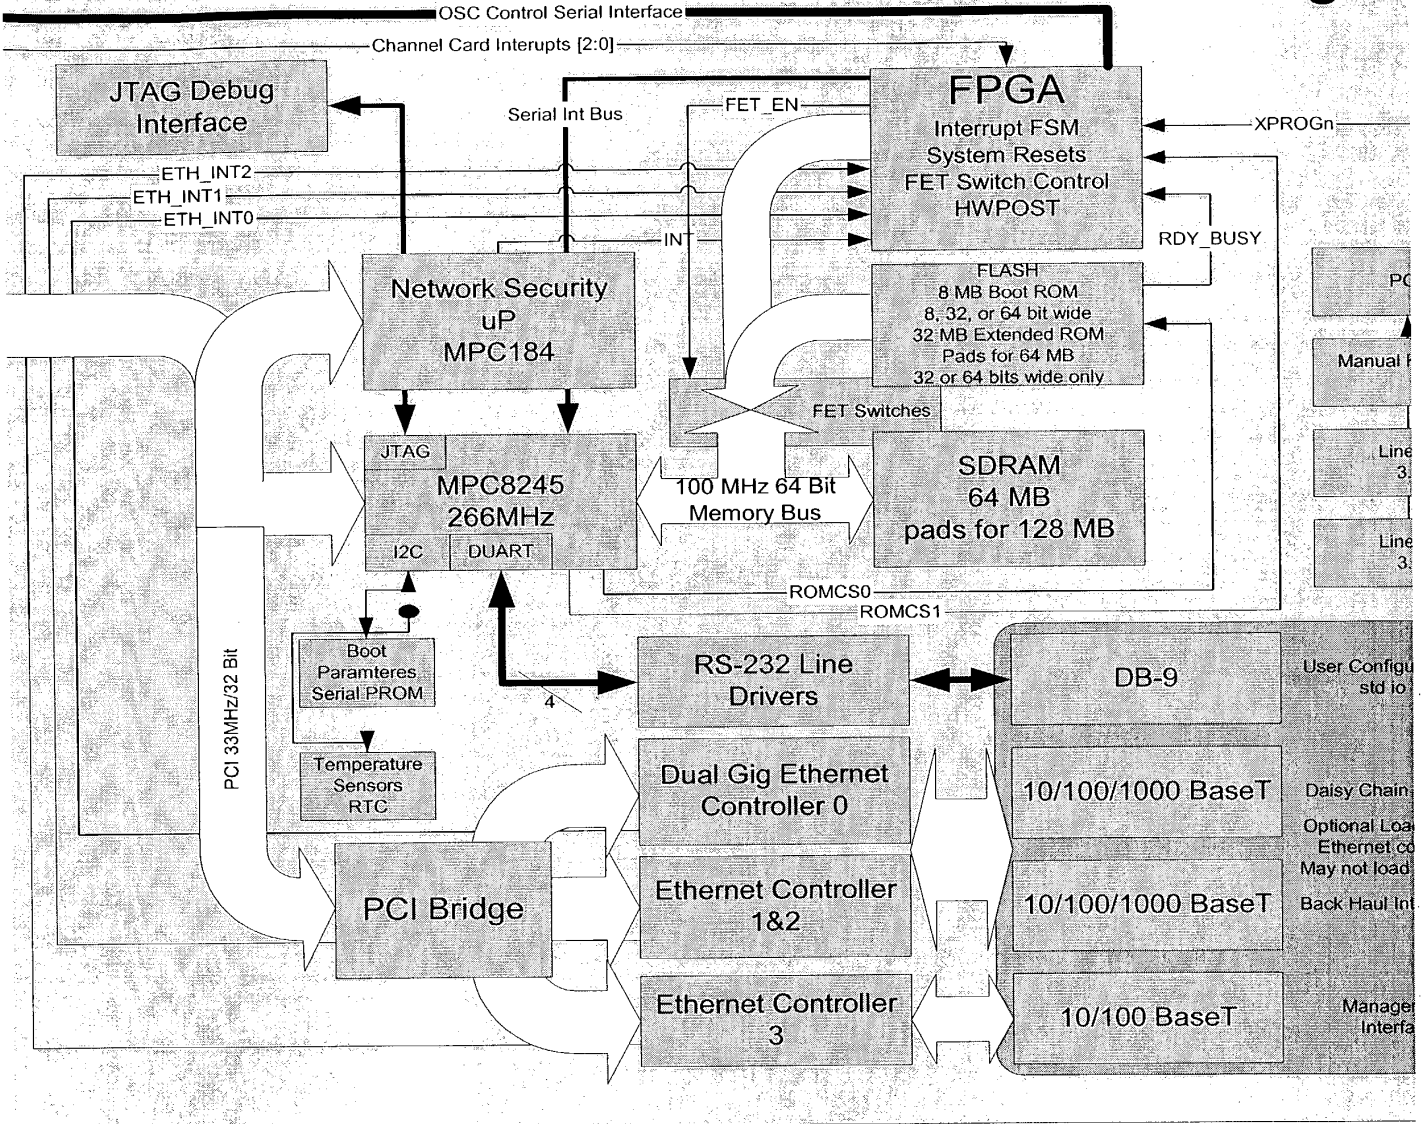




H-5





H-6



H-7

al Board



H-8

Antenna Board

LP Filter



RX

LP Filter



Antenna Board

LP Filter



RX

60423660 . 310402

[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]

H-9

[5 pages] I

Quickfacts – DirectedPacket 1

Brief

DPI is a router/bridge serving 802.11a and 802.11b clients and allows connectivity with either wired or wireless back-haul in addition to other clients. It also has a lot of the features found in the current generation of broadband routers/gateways.

All DPI features are easily configurable via a Web interface including routing, bridging, IP-Filtering, DHCP and NAT/PAT (Port Address Translation) functionality. In keeping with the "enterprise ready story", configuration of advanced features as exported by the 802.11 MIBs will be possible via the Web interface. The same configuration options is available via a menu-based (versus CLI) telnet interface. The menu-based interface should be available over the built-in RS-232 port or over the network once an IP address has been defined.

DPI must support username/password challenges for access to both Web and Telnet based configuration menus. A default "admin" user-name will be configured as a factory preset. DPI supports three levels of configuration access, read-only, configurator, and administrator with varying levels of information accessible and writable. DPI supports a variety of methods for operator authentication utilizing TACACS+ and RADIUS.

DPI supports a variety of mechanisms for wireless client authentication and access, namely 802.1X. Methods are included to optionally allow failed authentication clients to associate and be limited to a particular VLAN.

DPI supports accounting and utilizes RADIUS and TACACS+.

DPI can serve as an enterprise-level broadband router and includes support for PPPoE (Point to Point Over Ethernet), NAT and PAT, Microsoft's PPTP (Point to Point Tunneling Protocol), ACL (Access Control Lists) both MAC-based and IP as well as support for IPsec pass-through through advanced NAT transparency features. Additionally through port mapping PPTP pass-through can be achieved.

Unlike many current IP gateway products that are targeted at the residential market, the DPI is aimed at primary deployment in an Enterprise or Service Provider environment. In these deployments it is undesirable for a gateway device to be factory preset as either a DHCP server or client or to make any other assumptions about the network environment. The DPI will initialize and be a plug and play bridge using open authentication on a default ESSID.

RF Component

The product can support both 802.11a and 802.11b bands and is available in a variety of configurations.

The components that may be assembled consist of:

1. 802.11a antenna array panel.
2. 802.11b antenna array panel.
3. A hybrid 802.11a/b panel.
4. Radio summer board.
5. 802.11a radio and controller card.
6. 802.11b radio and controller card.
7. Management module.

The panel also forms the chassis, such that components 4-7 actually sit inside the panel.

Available configurations:

	802.11a/b	802.11a	802.11b
Panels	1+	1+	1+
Summer	2	1	1
Number of beams (radio)	2-7	1-4	1-3
Management Module	1 or 2 (resilient)	1 or 2 (resilient)	1 or 2 (resilient)

Each radio card is responsible for controlling one frequency on one beam. Each beam supports a single frequency of 802.11a or 802.11b.

The physical rate per beam is up to 11Mbps (802.11b) or 54Mbps (802.11a).

The system supports native 802.11 clients. A more efficient TDMA scheme is also supported by using a shim in the client. This may be downloadable via native 802.11.

The servicing scheme is based on an enhanced slotted time division algorithm that is provisioned according to the data rate agreed with the client.


Three panels covering 360 degrees can be mounted together as one system. Alternatively, same direction mounting to achieve higher capacities can be performed provided the separation between panels is sufficient for isolation.

In-band back-haul is provided by either band and includes provision for channel bonding up to three radios in 802.11b and four radios in 802.11a.


60723660 . 110402

I-1


The component layout for the three basic configurations (shown here fully equipped with the maximum number of radios) is as follows:



802.11a Four Beam Product




802.11a/b Seven Beam Product



Physical Layout

The enclosure and the internal components is anticipated to be layed out as follows:




The picture shows the antenna array components at the back of the enclosure, with the radios which drive the product on the right hand side. They plug into a Compact PCI midplane for connectivity to the redundant Packet Processor Modules on the right.

60423000 . 220402

1-2

Thus the logical view of the components is:



Wired Component

Each panel has two Ethernet fiber ports to connect the panels together. In a 360 degree configuration there would be a total of 6 ports, 4 of which would be used to daisy chain the panels and one of the remaining ports is used as a wired side-haul. If an interface to a side-haul is other than Ethernet then an external converter must be provided.

Networking Component

Authentication

A client will authenticate using 802.1X and compatibility will be guaranteed with Windows XP clients. Therefore the panel control system will support a RADIUS client and support 128 bit WEP keys and TLS.

This can be achieved using VLAN technology. A VLAN is able to keep client traffic separated according to a policy dictated by the administrator.

Other authentication methods such as PPPoE may also be utilized.

Authenticated Client Facilities

An authenticated customer will be able to:

1. Bridge to other clients according to policies created at provisioning.
2. Bridge to an IP gateway for WAN access.
3. Send and receive traffic to and from the IP gateway resident in the panel control system.
4. Obtain / renew / release IP addresses via DHCP (if enabled).
5. Apply rate limiting according to their provisioned rate.
6. Have 802.1p priorities honored.
7. Be able to access the provisioning server for changes in service.

Unauthenticated Client Facilities

What a customer can access without authentication is based on policy set by the administrator. However, the following is possible:

Unauthenticated users will normally not be able to gain access to premium services and will only be able to access specific services. Such services may be a provisioning server which, for example, may allow them to enter billing information in order to gain access. In order for this to operate, the user may also request addresses via DHCP, or alternatively be connected only to a device or service capable of providing IP masquerading for their gateway and DNS.

In a campus scenario access may be granted to the internet or intranet, but not information kept on campus servers.

Traffic from unauthenticated users will be restricted in as much that although all unauthenticated users may be members of the same VLAN, they will be unable to see each others traffic and will only be able to converse with predefined servers or gateways.

Provisioning Information

If a client is connected to a provisioning server through either authentication failure, or because a client wishes to change their service level, the following features will be available to them:

1. Ability to select the type of authentication they would like (PPPoE or 802.1X).
2. Ability to select parameters pertinent to the above.

S-T

3. The client selects their desired bit rate (this may need to be intelligently verified to check that it is available)

In a campus environment the MIS department will populate the provisioning server with information about the type of service the user is entitled to - such as the VLAN or VLANs they will be made a member of.

In a WISP environment the type of service will typically be a connection to the IP gateway (which may or may not perform NAT depending on the type of deployment), with no facility to connect to other clients in a layer 2 manner. However, through an enhanced provisioning server this facility can be added. As far as the DP-X is concerned, it simply gets provisioning data which tells it what VLANs a client is a member of, regardless of the deployment scenario.

Monitoring Facilities

In order to provide meaningful information to administrators to proactively maintain exist and deploy new coverage, a comprehensive set of monitoring facilities will be provided.

Wired side-haul and back-haul interfaces will be augmented with Remote Monitoring facilities including:

1. Basic statistics (bytes / packets sent / received etc).
2. Statistical history.
3. Alarm control (notification of abnormal traffic patterns to a central control point)
4. Conversation tracking (which users are talking the most at a macro level).

The radio interfaces connect to potentially hundreds or thousands of users. Therefore, side-haul and back-haul information at a more macro level is useful for tracking the overall health of an installation. However an "interface" can be dynamically created to monitor individual users or a group of users (using an "interface" like this allows reuse of existing MIBs).

The same monitoring characteristics on side/back-haul are then made available on a more granular level and also the following facilities are added:

1. Packet Capture
2. Host Matrix (helps a customer monitor their layer 2 domain and configure their network for optimum use and panel positioning).

VLAN Support

VLANs are supported in this product to support the services it is trying to provide, namely the client traffic separation; the ability to allow a client to see a provisioning server but not other un-provisioned clients on the same VLAN, and the ability to allow multiple customers to join the same layer 2 domain.

In addition, and especially in a campus environment, the client may wish to join multiple layer 2 domains and use 802.1Q tags in order to specify which VLAN the traffic is intended for. If a client sends traffic using tags for VLANs of which it is not a member, the traffic is dropped.

IP Gateway Functionality

In the initial release, the following features will be supported:

1. Traditional IP gateway (with Longest Prefix Match lookups)
2. Network Address Translation with DHCP non-routable addresses.
3. IP Masquerading

Since the DP-X initial deployment will be campus environments and will be at the edge of the network, so there is no immediate need to be an active participant in the network.

However, in later releases protocols such as OSPF and RIPv2 may be added.

Layer 2 Protocol Support

The Spanning Tree Algorithm and Protocol (STaP) is useful in enterprise environments for preventing logical and physical network loops (which layer 2 is otherwise unable to prevent) and providing network resiliency.

The question is when to actively participate in layer 2 protocols in a way which affects network connectivity. For example, if the DP-X is owned and operated by the same customer that owns both the clients and the backhaul network into which the DP-X is connected, then it becomes highly desirable to have the device participate in the Spanning Tree in the network. The customer is able to control the parameters governing the protocol and is therefore as comfortable with DP-X as they would be with any other layer 2 capable device. Currently the most interoperable solution in this space is 802.1s.

If the DP-X is deployed in a WISP environment with, perhaps, multiple business sites connected via common layer 2 domains, then there is a requirement that client STaP cannot affect other unrelated clients connected via the DP-X network. There is an additional requirement that the DP-X be able to carry client STaP frames transparently. Therefore the DP-X will allow the Spanning Tree software to be placed into a mode where it only runs on DP-X controlled links, such as back-hauls and side-hauls. If a customer accidentally creates a loop then, best case, their own STaP will detect and block

I-4

the loop, or, worst case, a loop is created and the DP-X is protected via rate limiting and broadcast limiting. Network monitoring software would also highlight that there is a problem with particular clients.

Technical Summary

Packet Forwarding and Topology Management:

802.1D-1998	Layer 2 Forwarding behavior.
	Single Spanning Tree.
	Handling of priority encoded frames.
802.1w	Fast Spanning Tree Enhancements.
802.1s	Multiple Spanning Tree Support.
802.1Q	VLAN tag support.
	Independent VLAN learning and forwarding model.
draft-ietf-brIDGE-learn	MIB Support for 802.1w.
RFC1493	Basic bridge management support.
	VLAN and STP extensions for wireless.
	Private MIB

Accounting

RFC2138	RADIUS Core.
RFC2139	RADIUS accounting.
RFC2975	Introduction to Accounting Management.
draft-grant-tacacs-02	TACACS+ informational draft.

Management Techniques:

SNMP v2	Remote management via wired and wireless interfaces.
SNMP v3	As above but secure.
Secure FTP	Download of upgrades from known server over WAN interfaces.
TFTP	Download of upgrades from trusted local server.
HTTPS	Remote management via a secure web browser.

Monitoring and Interface Management:

RFC3819	RMON (all 9 groups, for all fixed wired interfaces and configurable for specific clients or groups of clients). RFC2819 obsoletes 1757.
RFC1573	Interface Management (MIBII).
	Private MIB
	Antenna Management.

IT

60443652 . 110403

[3 1042] J

Conceptual Proposal for Little Joe Antenna

Basic feature set:

- Single Transmit Beam
 - The product may transmit on one of three channels in any given direction for any given packet.
- Two panels, one for Receive and one for Transmit
 - The product is capable of receiving and transmitting at the same time.
 - The product is NOT capable of receiving and transmitting on the same channel at the same time.
 - Thus, if a client on channel 1 transmits a frame into the panel at the same time the panel is transmitting to another client in a different sector, the panel will NOT receive their frame.
- The product is able to receive in all directions within a 120 degree radius on 3 channels simultaneously.


Enhanced Feature Set:

- Two more beams may be added.
 - The beams must be synchronized such that each beam, when active, talk on the same channel. This is because if they did not, the receive sensitivity would be greatly reduced. Three transmit radios active at the same time would mean that the product would be deaf to incoming traffic.

Product Layout

The product will contain 48 802.11b or 64 802.11a MACs. It needs this number for receive sensitivity. The Butler Matrix antenna array contains 16 sectors. Each sector can receive on three channels; thus 16 x 3 is 48.

The product in the example below transmits on a single frequency.



Transmit Module

The transmit side of each 802.11 MAC is wired (digitally) into a switching matrix capable of buffering. The need for this is based on the fact that the same 802.11 MACs are used for receive, and each of these MACs is capable of generating its own acknowledgement for packets it receives.


Thus, the firmware is modified so that ACKs are generated digitally and are sent to a transmit switch (D) where the ACK is buffered until the beam is in a suitable position to transmit the frame. Switch D buffers ACK responses from each channel via A, B and C.

The transmit module is also capable of transmitting data frames from the CPU of the product.


60423666 . 44402

J-1


The CPU of the product sits behind the radios in the fashion shown below:



The CPU only needs to connect (for transmit) to one radio for each channel. Data frames are created by the CPU, a channel is selected and frame transferred to the 802.11 MAC responsible for transmission of frames for that channel. A further breakdown shows how the direction of the beam is also selected:



ACK Mechanism



As it can be seen, the transmit side of the 802.11 MAC has been modified so that it may perform lookups based on the destination MAC address of the frame that needs to be transmitted. The returned associated data is the direction and power that needs to be applied to the frame. Thus the switch and packet buffer stores not only the frame, but also this associated data which is used by the beam former.


Should the lookup fail, then the associated data returns defaults. In the case of an ACK this data will be the input RSSI and default direction for that sector of the antenna, or in the case of a packet generated by the CPU, drop the frame indicator (since the MAC will be unaware of what power to apply to the frame for transmission).

Receive Side

Copyright © 2002

5.2

Receive MAC Architecture



MAC SA	RSSI	Assoc.	Assoc. RSSI	Direction
A	5	Yes	6	45
B	6	No	9	47
Default Direction				44

Client is associated with this MAC at signal strength 5

Client is associated with a different MAC at signal strength 9

Used for transmits and ACKs for report into Beam Form

Used when a lookup fails for Outgoing Packets and ACKs. It is consulted with the RSSI of the input frame to determine the client direction and the power needed to talk to them.


This diagram aims to describe the basis for "learning" where clients are in terms of power and direction. This information can then be populated into tables in each 802.11 MAC and used for automating the ACK process.

When a packet is received it may be received on multiple 802.11 MACs. The receiving 802.11 MAC will pass the frame to the CPU and then try and generate an ACK. It performs a lookup (which in this example will fail) and it then sends an ACK to a default direction on the power level appropriate to reach them (which is derived from the receive power for the frame received).



The CPU will read this frame and identify that it has received several copies. The frame can be forwarded according to normal rules, but the CPU must decide which MAC will be ultimately responsible for handling data for this client. This is determined by looking at the receive strength for each frame, and the corresponding MAC is selected. In the MAC which is selected an entry is built into the lookup table with an absolute direction

and power level for the client which is derived by the CPU triangulating the results of the multiple copies of the frame that it has received.



In the MACs that are not selected to handle the frame, an entry is built which states that the MAC is not to associate with this client (do not pass frames to the CPU for this client) along with the signal strength for the MAC with which the client is associated. This is used to generate indications from the MAC to the CPU to indicate that the client has moved and is now being received stronger on this MAC rather than the MAC it is associated with.



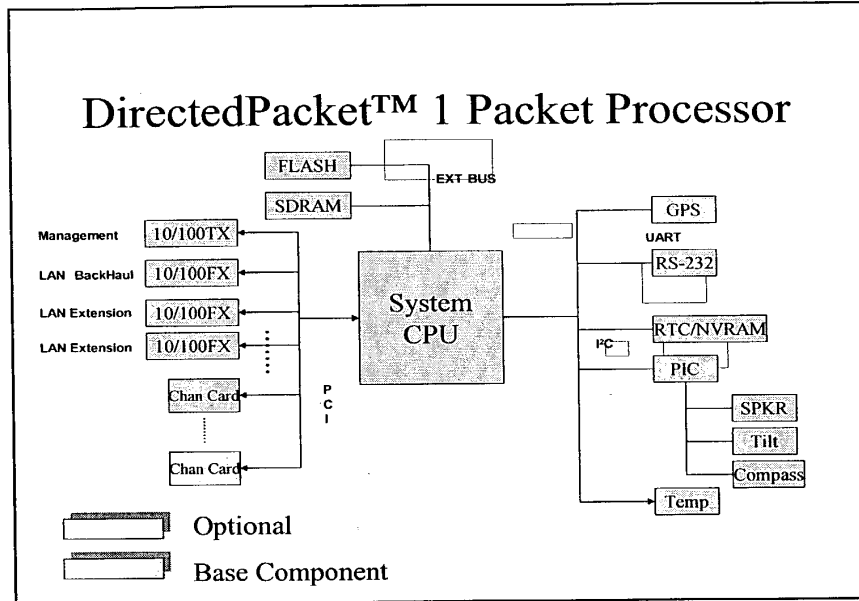
[7 pages] ~~_____~~ K



K-1



k-2



- ## Scalable
- No client number limitation
 - NAT Gateway functionality
 - Capacity by adding Channel Cards
 - 3 for 802.11b
 - 4 for 802.11a
 - Capacity by adding additional panels
 - Capacity through “Beam Bonding”
 - In-Band Backhaul: Bridging via IEEE WDS (Wireless Distribution System)
 - Three 10/100 Base FX SC connectors

K-3

Manageability

- Industry standard CLI
- Independent 10/100 management
- SNMP v1, v2, v3
- Web based management
- HP OpenView support
- Extensive MIB support (Bridging, Routing, Enterprise extensions)

Provisioning

- Zero Configuration
 - 802.1X client authentication
 - Out of box operation with popular OS
 - User driven provisioning scenarios
 - Hot Spot
 - Enterprise
- 802.1X interfaces to global database for authentication/provisioning information
 - LDAP
 - RADIUS
- DHCP server/client functionality
- Accounting via RADIUS

K-4

Bridging

- 802.1 p/Q VLAN tagging and priorities
- Full 802.1D bridging support
- IEEE WDS (Wireless Distribution System)
 - Panel can bridge to other panels or capable AP
- 802.1s,w Loop detection
- Beam Bonding: Combine channels for larger aggregate throughput

Routing

- Supports common Dynamic Routing Protocols
 - RIPv2, OSPF
- VRRP (Virtual Router Redundancy Protocol)
- Static Routes
- IPv6 capable

K-5

QoS & Enforcement

- CBQ (Class Based Queuing)
- Radius Authentication
- Rate Limiting: CIR (Committed Information Rate)
- ACL (Access Control Lists)
 - IP or MAC
- 802.1X access control


Resiliency

- Topology resilient
 - OSPF and RIPv2
- Connection resilient
 - STP (Spanning Tree Protocol)
- Router resilient
 - VRRP (Hot Standby Router Protocol)
- Multiple Beams provide redundancy
- Optional redundant Packet Processor

K-6

[5 pages]

[REDACTED] L



[REDACTED]

- 802.11b N-type connector
- 10/100 Base FX SC connector
- Adaptive Array provides AP and VIA functionality

L-1

[REDACTED]

Extend the Network

- Steerable Array provides reach
- Side-haul via 802.11b N-type connector
- Side-haul via 10/100 FX connector

Manage The Network

- Industry standard CLI
- SNMP v1 with v2 community extensions
- HTTP web interface
- Java plug-in for NMS (HP OpenView ...)

L-2

Bridge The Network

- 802.1 p/Q vlan tagging and priorities

Route The Network

- IPv4 forwarding
- IP ToS (Type of Service) Aware

L-3

Make the Network Resilient

- Topology resilient
 - OSPF and RIPv2
- Connection resilient
 - STP (Spanning Tree Protocol)
- Router resilient
 - VRRP (Hot Standby Router Protocol)

Trouble-Shoot the Network

- Tilt meter with Audio
- GPS
- Compass
- Temperature
- Thermostatically controlled fan

L-9

Shape the Network

- DiffServ
- CBQ (Class Based Queueing)

Enforce the Network

- Radius Authentication
- CIR (Committed Information Rate)
- ACL (Access Control Lists)
- 802.1X Port based access control

L-5