MOTOROLA
intelligence everywhere™

Technology Overview

Bandwidth Versus Throughput

February 2004

WHITE PAPER

TABLE OF CONTENTS

LIST OF FIGURES

ACRONYMS

ACK	Acknowledgement
AP	Access Point
FTP	File Transfer Protocol
IP	Internet Protocol
LAN	Local Area Network
RF	Radio Frequency
SM	Subscriber Module
TCP	Transmission Control Protocol

Notice

The information in this publication is subject to change without notice. Motorola shall not be liable for technical or editorial errors or omissions nor for any damages resulting from the use of this material.

Each configuration tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Motorola does not warrant products other than its own strictly as stated in Motorola's product warranties.

MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. Canopy is a trademark of Motorola, Inc. All other product or service names are the property of their respective owners. © Motorola, Inc. 2004.

Introduction

The Canopy™ system enables broadband access deployment with a wireless point of presence that can be connected directly to broadband infrastructure or indirectly via a wireless backhaul link to fiber or microwave network infrastructure.

In very specific and unusual situations, protocols that were designed specifically for Local Area Networks (LANs) rather than for Internet usage can cause performance degradation when used in a wireless system. The reason behind these performance limitations can be traced directly to the protocol. While some of these protocols work very well in over provisioned high-speed wired networks, they do not translate well into the wireless arena. This is due, in large measure, to the limited bandwidth available in wireless as well as the large number of small packets the protocol requires for transmission over the network. As you will see, these protocols are extremely inefficient for use in wireless networks.

This paper describes the performance and capacity capabilities for the Canopy wireless broadband system. Please note, the examples highlighted in this document are the result of a worst case traffic analysis that was performed by Motorola for the purposes of future protocol work. They do not represent real traffic patterns that are likely to ever be seen in the real world.

Bandwidth versus Throughput

The Canopy system's performance is outstanding in typical traffic scenarios. A wireless Access Point (AP) cluster can contain from one to six Canopy modules. Each transceiver delivers an aggregate of 6.2 Mbps of bandwidth and supports 200 subscribers per sector for point-to-multipoint systems and 14 Mbps for point-to-point systems. Six multipoint transceivers have an aggregate throughput capacity of 36 Mbps.

In a Canopy system, there is a standard 2.5 millisecond repeating frame used to time transmissions. This consistent interval ensures constancy in the level of service and eliminates self interference by enabling the access points to be synchronized using the Global Positioning System (GPS). *Figure 1* depicts the Canopy air frame and *Figure 2* shows a detailed illustration of only the uplink data slot. A description of each of the functions is shown beneath the figures.

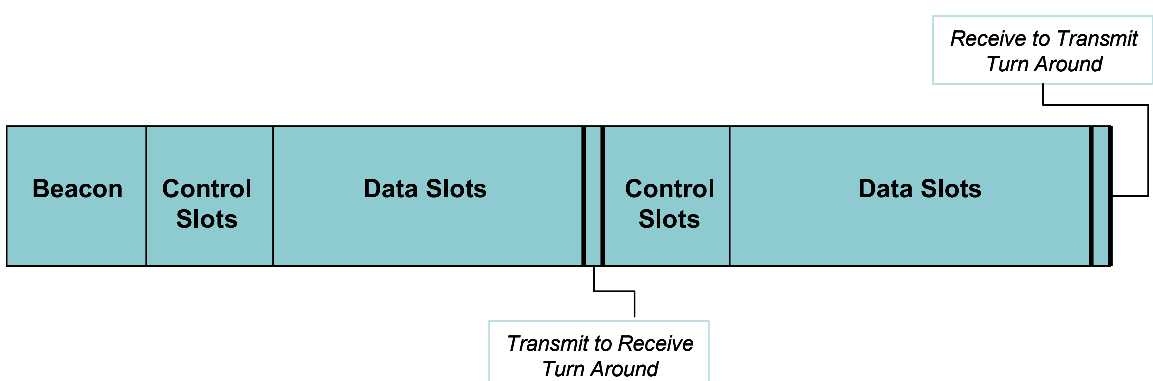

Figure 1. Canopy Air Frame

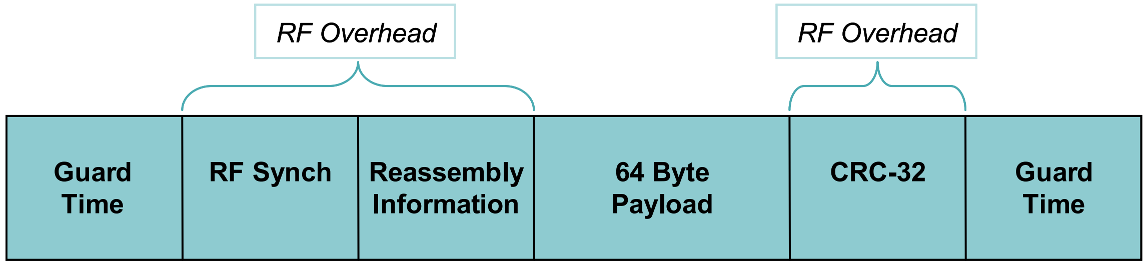

Figure 2. Uplink Data Slot

- *Beacon:* Broadcasts identity and configuration information.
- *Control Slots:* Provides for uplink/downlink scheduling and keep alive messages.
- *Data:* Carries user data traffic broken into 64 byte fragments.
- *Turn-Around:* Allows time for the radios to transition from transmit to receive mode. This time gets longer according to the distance of the farthest subscriber.
- *Guard Time:* Separates slots to ensure no overlap between SM transmissions.
- *Timing and Reassembly Information:* RF clock synchronization and reassembly header Information.
- *Error Checking:* A CRC-32 which is carried for each fragment

Special Applications

Bandwidth is a measure of the capacity of the link. Throughput is the rate at which usable data can be sent over the channel. While Canopy multipoint system's bandwidth is 10 Mbps, the throughput in terms of usable data may vary depending on:

- Types of Packets
- Packets per Second
  (Multipoint Downlink ~ 3,000 packets per second/Multipoint Uplink ~ 1,500 packets per second)
- Operating System of the Sending and Receiving Machines
- Performance of Networking Equipment between the Two Endpoints or the Network Interface of the Sending and Receiving Machines
- Performance of the Application Software

The types of packets sent over the multipoint system will have an impact on the throughput of the link based on the size of the packet being passed. This is a function of the application. In the Canopy system, all Ethernet packets are broken up into 64 byte fragments for transmission over the Radio Frequency (RF) link. Smaller packets of information may decrease the efficiency of the Canopy system's throughput for a number of reasons.

First, the increase in the number of acknowledgements that now must be sent back over the link decreases efficiency; smaller packets translate to a higher number of acknowledgments (ACK). Second, smaller packets don't always align well to the 64 byte data slots. To illustrate, think about the Ethernet as a freight train where each train car is capable of carrying 64-bytes of information. If a packet contains 80-bytes of information, it will be broken down into two train cars. The first car will carry 64-bytes while the second car carries 16-bytes. As you will note, the second car is only partially full and has 48 bytes of unused space.

In certain non-typical conditions, the Canopy system may reduce the actual throughput to less than optimum levels. This change in throughput is only experienced in very specific and unusual network configurations carrying atypical Ethernet traffic.

Figure 3 and *Figure 4* depict extreme cases for Ethernet frame sizes *(not the size of the fragments the Canopy system uses internally)* sent across a Canopy multipoint system. In both cases, the APs and Subscriber Modules (SMs) used in the test are employing their default factory configurations. Packet sizes are locked and sent at a continuous rate.

There are limits to and from one Canopy SM with the remaining Canopy bandwidth used by other SMs. The current limits are a result of certain pathological cases regarding over-the-air protocols which are being addressed to avoid confusion and enhance some specific non-Internet applications for the Canopy product. The limit to a single SM occurs when streams of ALL minimum size Ethernet packets are being sent.

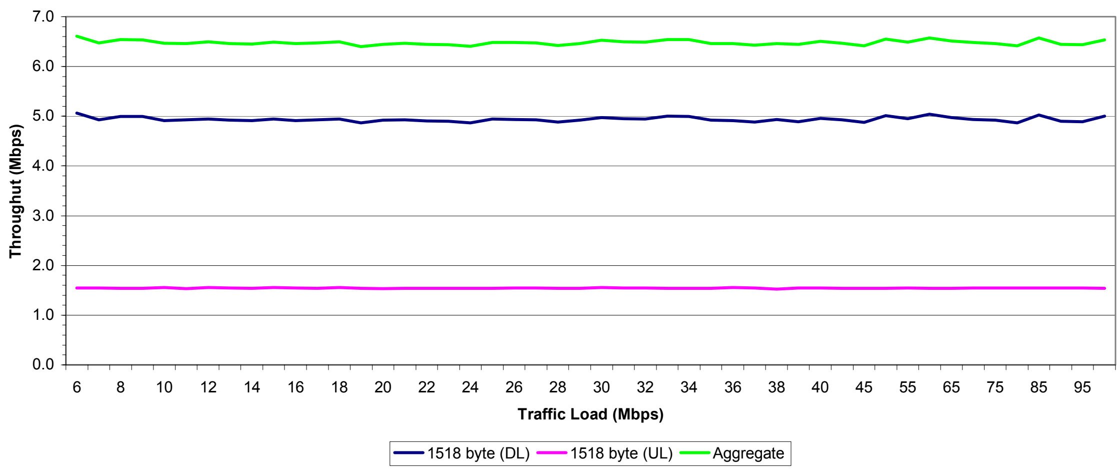

Figure 3. Throughput Profile for 1518 Byte Packets

As shown in *Figure 3*, in the 1518-byte packet case, the aggregate throughput provided through the AP is approximately 1.6 Mbps in the uplink and 4.9 Mbps in the downlink. In the 64-byte packet case, shown in *Figure 4*, the aggregate throughput through the AP is drastically lowered. This is a worst-case. The uplink can support approximately .36 Mbps and the downlink can support approximately .82 Mbps.

If only one out of each eight Ethernet frames is a more typical size, these bandwidth limits go away. Real world Internet traffic does not consist of only 64 byte packets. If this were the case, the performance of the Internet would be extremely poor due to events such as router overload. TCP acknowledgements (ACKs) are 80 bytes in length and payload packets are typically 1K bytes. Even with a seven to one mix of payload to ACK packets these limitations will no longer apply.

The TCP/IP Protocol

As the packets per second are increased within the multipoint system, there is a point at which latency must also increase as the time to process each packet of data will increase with the overhead of Transmission Control Protocol (TCP)/Internet Protocol (IP).

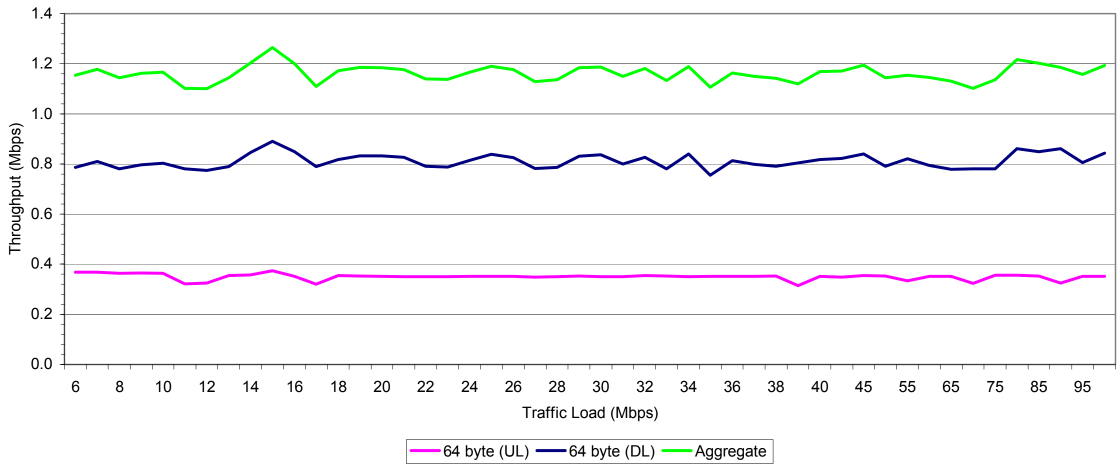

Figure 4. Throughput Profile for 64 Byte Packets

The Canopy multipoint system was architected to support 200 subscriber modules registered to a single AP. A single SM registered to an AP module will not produce the specified throughput. Rated Access Point throughput is only realized when more than one SM is registered to an AP. In a single SM scenario, there is often a considerable amount of inactive time between transmissions. This inactive period comes from the TCP protocol which requires acknowledgement of sent packets before additional transmissions can occur.

The FTP Protocol

An example of this also occurs with the File Transfer Protocol (FTP) protocol. In order to regulate communication between two units, the receiving unit sends a starting TCP receive window buffer size to the sender as part of the connection initialization process. The sender is then allowed to send up to that much traffic before receiving an ACK that the data was correctly received and processed. Most pre-Windows XP operating systems typically used a default number of less than 10 Kbps. In other words, after the sender sends 10 Kbps of data, it must wait for a TCP ACK before it can send anymore data.

When pipes were small (56K modems) it would take approximately 1.4 seconds to send 10 Kbps of data; plenty of time for the TCP ACK to get back to the server. As a result the server would continuously stream data. With the Canopy system's multipoint architecture there is an addition of 20 milliseconds of roundtrip delay (5 milliseconds in the downlink and 15 milliseconds in the uplink). Because of this added latency, a protocol such as FTP is forced to stop sending data to wait for the TCP ACK's. Canopy technology still provides the specified aggregate bandwidth, however it needs to be shared among multiple TCP sessions.

The Bottom Line

As with any wireless system, the method of packet transmission (how the wireless system transmits the information) as well as the size of Ethernet packets must be considered. Some programs transmit large numbers of very small packets containing keyboard, mouse and screen information. These small packets can seriously degrade the performance of most wireless systems; therefore, understanding a wireless system's capabilities with respect to packet rates is important in designing an efficient network. The Canopy system's performance is outstanding in typical traffic scenarios. Changes in the level of performance are only experienced in very specific and unusual network configurations.

Motorola Canopy Wireless
Broadband Products
1299 E. Algonquin Road
Schaumburg, IL 60196
www.motorola.com/canopy

MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. Canopy is a trademark of Motorola, Inc. All other product or service names are the property of their respective owners.
© Motorola, Inc. 2003.

4/020504