

## **Documentation Samples**

This document contains excerpted pages from three documents I have created during my career.

- **Users Guide — *The WaveWatcher Operator's Handbook***  
CIENA makes advanced repeaters (i.e., signal-boosting nodes) for high-density fiber optic networks. WaveWatcher is a program which supports network hardware nodes manufactured by CIENA. WaveWatcher provides a variety of real-time monitors on the network node activity. (Page 2)
- **Internal Hardware/Software Documentation — *The Network/Spectrum Interface***  
Straddles the boundary between a software API (application programming interface) document and register-level documentation for an ASIC (application specific integrated circuit). The electronic chip is Cognio's SAgE ASIC, which monitors and analyzes the radio frequency spectrum. The Network/Spectrum Interface (NSI) is a set of software messaging protocols used to control SAgE. (Page 13)
- **Technology White Paper — *CBF and Vector CBF***  
A white paper on composite beam forming (CBF), a technology developed by Cognio. CBF uses multiple antennas to enhance radio frequency signal range/signal reliability for radio communications. (Page 25)

(Master page numbers for this document are on the lower right corner of each page.)

*An excerpt from...*

## **The WaveWatcher Operator's Handbook**

*A Ciena Publication*

Written and illustrated by Steven C. Oppenheimer

Ciena, Inc., is the creator of the MultiWave 1600 system. MultiWave compresses and transmits, over a single fiber optic line, data that previously required sixteen fiber optic lines for transmission.

WaveWatcher is a graphical program which runs on UNIX workstations. The program monitors the performance of communications hardware sites, called "nodes", in a MultiWave network.

- The excerpt contains an overview of WaveWatcher; a discussion of the relation between on screen icons and real, physical network element; and instructions for tracking events in the network.
- The manual was written using FrameMaker 5 for Windows 95.

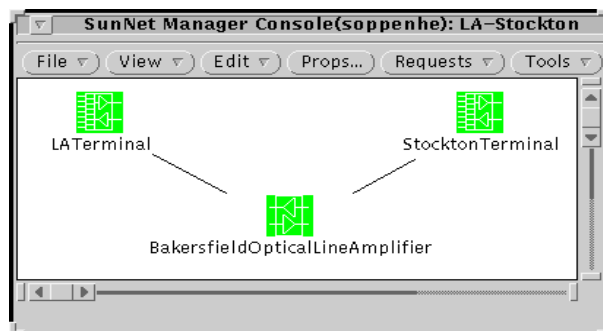
# Chapter 1 Introduction

---

## 1.2 OVERVIEW

WaveWatcher offers these features and characteristics:

- Continuous, real-time observation and management of a MultiWave network from a central location (a central office or network control center).
- The reporting of any changes in operating status (e.g., normal to major alarm) for any node in a Multi-Wave network.
- The ability to create and display a “graphical map” (a drawing) of a MultiWave network using icons that symbolize:
  - Geographic regions and rings
  - Routes (placed within geographic regions or rings)
  - Paths (placed within routes)
  - Physical components, i.e., terminals and optical line amplifiers (placed within paths).





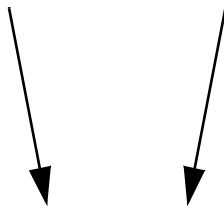
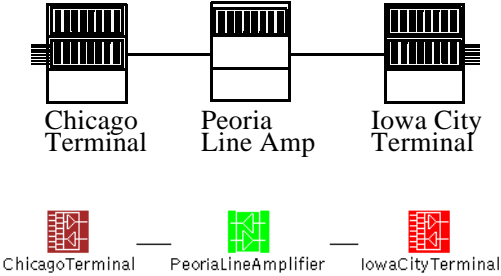
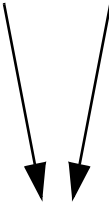
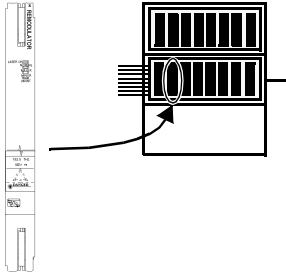
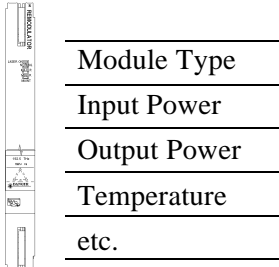
- When a more detailed analysis of a node’s performance is required, WaveWatcher provides access to MultiWave’s Craft Interface.™

# Chapter 1 Introduction

## 1.3 THE ROLE OF WAVEWATCHER™ IN MANAGING MULTIWAVE™ SYSTEMS

WaveWatcher™ works in tandem with MultiWave's™ Craft Interface to provide progressive network management capability, as shown in **Figure 1-1**:

**Table 3-1. MultiWave Network Management Hierarchy and Software**

Level Within Network (from highest to lowest)	Graphic Representation	Software Management Tool
1. MultiWave Region, Ring, Route, or Path 		WaveWatcher
2. Specific Nodes Within a Path 		WaveWatcher
3. Circuit Packs in a Node 		Craft Interface
4. Circuit Pack Attributes		Craft Interface

# Chapter 1 Introduction

---

## **1.3.1 WaveWatcher™ -- How it Fits in Your Network Management Activities**

Like a wide angle lens on a camera, WaveWatcher™ provides a high-level view of a MultiWave™ network. As such, it meets the management requirements listed in items 1 and 2 of Table 1 on page 1-3. WaveWatcher™ is accessible from a Sun workstation, and is typically used in a primary communications hub such as a central office or network control center.

WaveWatcher™ is the management tool that detects a general problem within a MultiWave™ network, and isolates that problem to a specific node. A trouble ticket might be generated by a network administrator and passed to the craftsman, who would rely on the Craft Interface to “shed more light” on the problem.

## **1.3.2 The Craft Interface vs. WaveWatcher**

Like a macrolens on a camera, the Craft Interface provides a minutely detailed view of any circuit pack within a MultiWave™ node. As such, it meets the management requirements in items 3 and 4, listed in Table 1 on page 1-3. The Craft Interface can be used at any node that reported an alarm status through WaveWatcher™.

Once the Craft Interface has traced a problem to a circuit pack (e.g., a Remodulator) within a node, an alarm list can be viewed to determine its exact cause (e.g., the laser temperature of the Remodulator exceeded its maximum threshold).

**NOTE:** The Craft Interface can be accessed from WaveWatcher; however, WaveWatcher cannot be accessed from the Craft Interface.

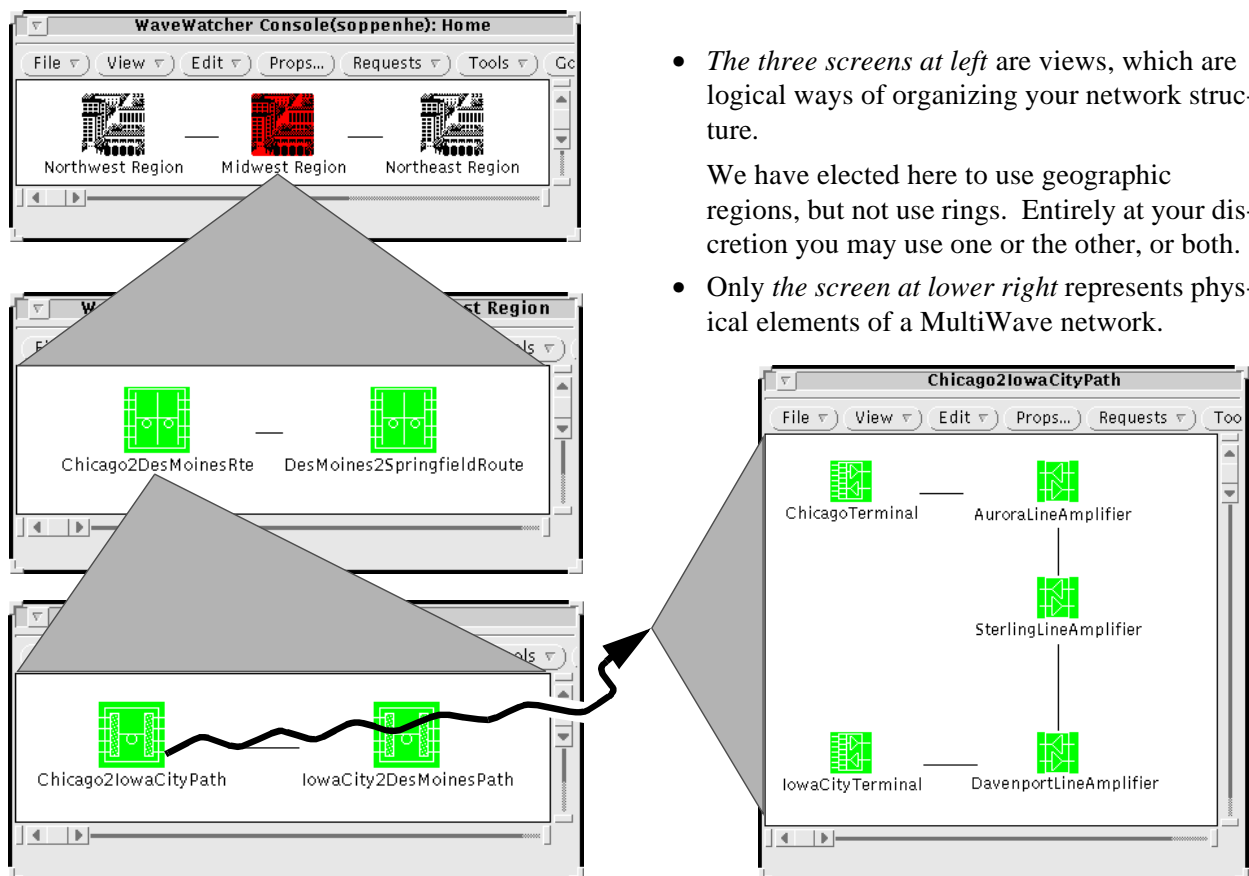
### 2. MULTIWAVE ICONS AND THE PHYSICAL NETWORK

#### 2.1 CHAPTER SUMMARY

When you use WaveWatcher, a diagram of your network will usually already be in place. (For information on creating a network diagram from scratch, please see the “WaveWatcher System Administrator’s Guide.”) In the network diagram, icons are used to represent:

- **Physical network elements**, such as optical line amplifiers (OLAs) and MultiWave terminals.
- **Views**, which are logical and geographic groupings of network elements:
  - **Paths** – have a terminal at either end, and may have one or more line amplifiers in between.
  - **Routes** – contain one or more paths.
  - **Rings** and/or **Geographic Regions** – contain one or more routes. You may place rings inside a region; or place geographic regions inside a ring.

FIGURE 1. WaveWatcher Views of the Physical Network

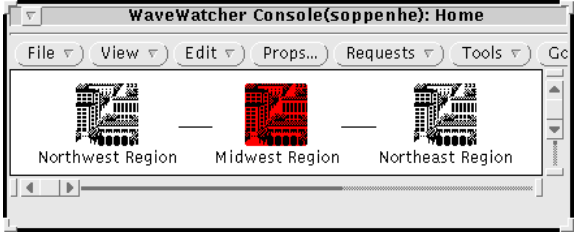
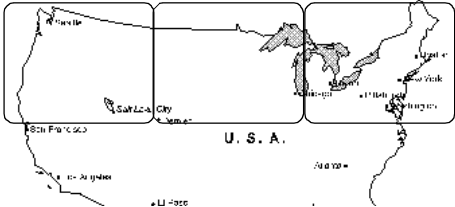
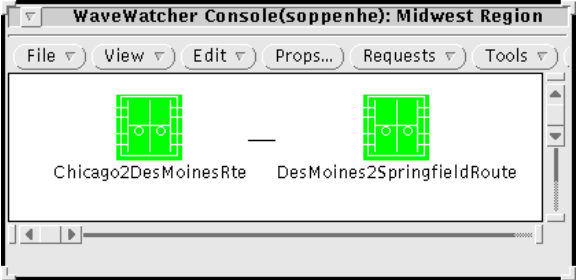
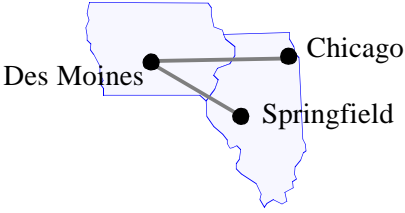
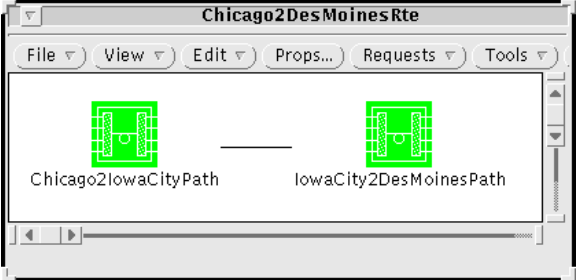
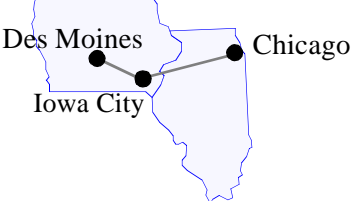
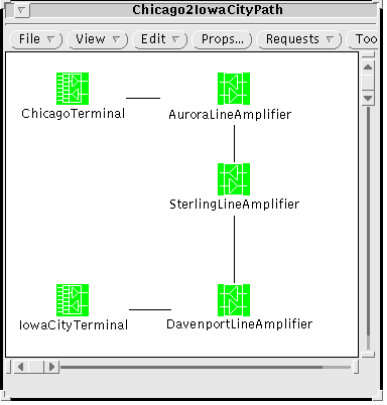
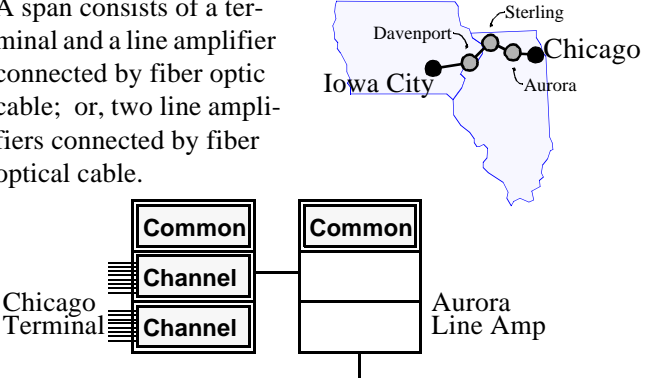


See Table 1, “WaveWatcher Icons vs. The Real Network,” on page 2-2 for more information.

# Chapter 2 WaveWatcher and the Physical Network

## 2.2 WAVEWATCHER ICONS VS. THE REAL, PHYSICAL NETWORK

TABLE 1. WaveWatcher Icons vs. The Real Network

Icons	What They Represent
<p><b>Geographic Regions and/or Rings</b></p> 	<p>High level groupings of network elements, typically covering a large geographic region.</p> 
<p><b>Routes</b></p> 	<p>A route covers more than 600 kilometers. Icons for two routes are shown at left: Chicago – Des Moines, and Des Moines – Springfield.</p> 
<p><b>Paths</b></p> 	<p>A path covers no more than 600 kilometers, and has at most two terminals. The first icon at left represents the path from Chicago to Iowa City.</p> <p>The second icon at left represents the path from Iowa City to Des Moines.</p> 
<p><b>Spans within paths.</b></p> <p>The icons here represent physical components—terminals or OLAs—within the MultiWave network.</p> 	<p>A span consists of a terminal and a line amplifier connected by fiber optic cable; or, two line amplifiers connected by fiber optical cable.</p> 

## Chapter 2 WaveWatcher and the Physical Network

### 2.3 MOVING BETWEEN DIFFERENT LEVELS IN YOUR NETWORK DIAGRAM

You will often need to navigate up and down through your network diagram.

- **Navigating Up** means viewing a higher level of the diagram, which covers a larger area and contains more network elements, paths, etc.
- **Navigating Down** means viewing a more local, detailed level of the diagram.

#### 2.3.1 Navigating Down

There are two ways to navigate down:

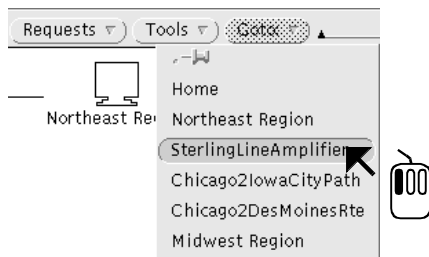
1. Double-click on an icon. (Click rapidly twice in succession.)



The WaveWatcher console shows you the “contents” of the icon. It may contain other icons, or it may be empty.


*or....*

1. Select the **Goto** menu. From this menu, select the network level of interest. (Note: This menu only displays the **Home** level, and levels you have already visited by double-clicking.)



#### 2.3.2 Navigating Up

1. Select the **Goto** menu. From this menu, select the network level of interest. (**NOTE:** This menu only displays the **Home** level and those levels you have already visited by double-clicking.)



### 3. MONITORING EVENTS AND ALARMS IN YOUR MULTIWAVE NETWORK

#### 3.1 CHAPTER SUMMARY

The WaveWatcher Console displays a picture of your MultiWave network. When unusual events occur, WaveWatcher signals you by changing the colors of the icons on screen, and by making the icons blink. The meaning of the different colors is discussed in Table 1, “Node Operating Status and Recommended Icon Colors,” on page 3-4.



#### 3.1.1 Procedures

This chapter describes procedures to:

- **Obtain the current operating status** of any node(s) within a MultiWave network. See “How Wave-Watcher Visually Alerts You To Events and Alarms” on page 3-3; “Viewing Event Histories” on page 3-6; “Viewing Current Alarms: The Fault List” on page 3-20; and “Re-setting a Blinking Network Icon” on page 3-27.
- **Review the operating event history** of the nodes in a MultiWave network. See “View Event Histories for All Nodes: The Event Viewer” on page 3-7.
- **Set filters** that limit the view of operating histories to nodes that meet specific criteria (e.g., date and time). See “Setting Filters in the Event Viewer” on page 3-13.
- **Obtain detailed information** about a node’s operating status via the Craft Interface. See “Accessing the Craft Interface” on page 3-25.

#### 3.1.2 Tools

You may also use this chapters as a guide to specific WaveWatcher tools. These tools include:

- The **Event Viewer** – Presents all the events that have occurred in your network. This includes alarms which have been asserted, alarms which have been cleared (resolved), and non-alarm events (e.g., the addition of a new node to the network). See page 3-6.
- The **Fault List** – Presents alarms which need attention *now* (i.e., which are currently in effect). Once an alarm has been cleared—the problem has been solved—the alarm drops off the **Fault List**. See page 3-20.
- **Alarm Reports for Individual Icons** – You can view detailed information on individual nodes using The Craft Interface (“Accessing the Craft Interface” on page 3-25).

## Chapter 3 Monitoring Events and Alarms in Your Network

---

### 3.2 DEFINITIONS: EVENTS AND ALARMS

#### 3.2.1 Alarms

An alarm is an abnormal condition which exists in a MultiWave element; the condition may negatively impact network traffic. An alarm condition has duration in time, and remains in effect until it is cleared. It is like your house going on fire; the alarm condition remains in effect until the fire is put out.

WaveWatcher is alerted to alarms in two different ways: (1) A network element (an agent) sends a trap signal to WaveWatcher; or... (2) During polling, WaveWatcher receives a status message which it interprets as being unsatisfactory.


(NOTE: WaveWatcher will also signal a critical event if, during polling, it receives no response at all from a node. WaveWatcher will query the node several times—two, by default—before signalling this event.)

#### 3.2.2 Events

Any *notification to you* (through WaveWatcher, for example) of a change in the status of a network node. *Some* events indicate an alarm. Other events may indicate when a node is added to the system, or when an alarm is cleared (i.e., the problem condition is resolved).

An event does not have a duration in time. It is like a quick tap on the shoulder, signaling you to turn around to see a sunrise (not an alarm); or to see your house burning down (an alarm!); or to see that the fire has been put out (the alarm has been cleared).

**Delay:** There is a small delay between *events* (notification to you, via WaveWatcher) and *alarms* (negative change in condition in the network.)



#### 3.2.3 Summary

An **alarm** is a condition, with a duration in time, which may have a negative impact on network performance. An **event** is a notification of any change in network condition. An event may or may not signal an alarm. Unlike alarms, events do not have a duration in time.

## Chapter 3 Monitoring Events and Alarms in Your Network

### 3.3 HOW WAVEWATCHER VISUALLY ALERTS YOU TO EVENTS AND ALARMS

The map in Figure 3-1 shows two MultiWave Terminal nodes, Chicago and Iowa City, linked together in a Multi-Wave network.

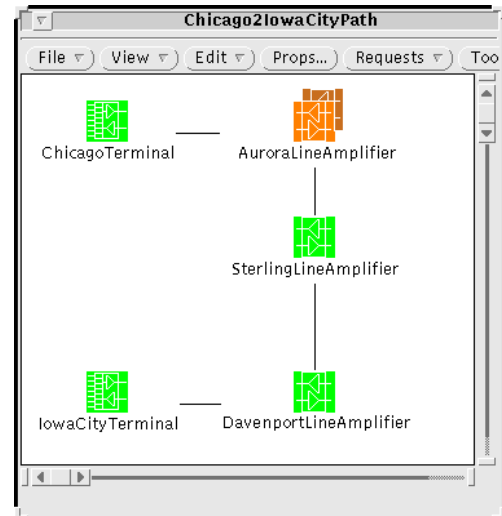
The *current operating status* of each node can be immediately determined by the *color* of its icon, as described in “Node Operating Status and Recommended Icon Colors” on page 3-4.

- **More than One Alarm in a Node:** If more than one alarm class exists within a node, the color reflects the status that is highest in severity (i.e., the worst problem).
- **Blinking and State Changes:** Any time a node *changes state*, the node will begin *blinking* from solid color to pastel. (In Figure 1 on page 3-3, the blinking is represented by the “doubled” icon.)

**Note:** This means that even a green node (a node operating normally) can blink, if it has changed from some other state to green.

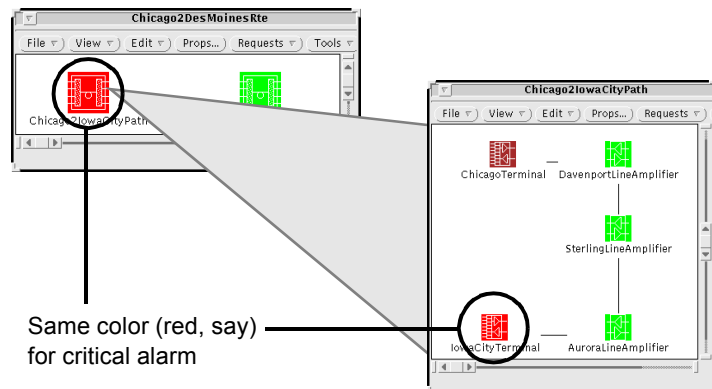
- **To Stop The Blinking:** Select {icon} → Glyph State → Normal. (See “Re-setting a Blinking Network Icon” on page 3-27 for illustrated instructions.) After the blinking has stopped, the icon will remain in the same color as before, indicating the operating status of the network element it represents.

FIGURE 1. A Node In Alarm




#### 3.3.1 Containers and Status Colors

When node icons are stored within a route container, the color of the container also reflects the status that is highest in severity.



## Chapter 3 Monitoring Events and Alarms in Your Network

To view the icons (i.e., the lower level views) within a container, double-click on the icon; or use the **Goto** option on the menu.



See Chapter 4 for more information on moving between levels in your network diagram.

**Table 3-1. Node Operating Status and Recommended Icon Colors**

If the Icon is: <sup>a</sup>	Alternate Colors: <sup>b</sup>	It Indicates:	Which Means that:
<b>Red</b>		A <i>Critical</i> operating condition.	WaveWatcher™ has lost communication with the node.
<b>Red</b>		An operating condition that is causing a <i>Major Alarm</i> .	The performance of one or more of the circuit packs in the node is negatively affecting service.
<b>Yellow</b>		An operating condition that is causing a <i>Minor Alarm</i> .	One or more circuit pack attributes in the node have exceeded normal limits, but none of the conditions are service-affecting.
<b>Orange</b>		<i>Warning</i> criteria.	WaveWatcher™ has determined that configuration information (e.g., part number) associated with one or more of the node's circuit packs is missing.
<b>Light Blue</b>		<i>Informational</i> criteria.	Events that do not impact service (e.g., attribute changes) have occurred within one or more circuit packs in the node. <b>Note:</b> Such information originates from the Auxiliary Management Module (AMM).
<b>Violet</b>		<i>Special</i> criteria.	WaveWatcher™ has found that the icon does not represent a MultiWave™ node.
<b>Green</b>		<i>Inhibited</i> operating state.	One or more circuit packs within the node have been placed in maintenance mode. The node remains in operation, but does not generate alarm information.
<b>Green</b>		Normal operation.	All circuit packs within the node are currently operating within normal standards.

*An excerpt from...*

## **The Cognio NSI Specification**

*A Cognio Publication*

Written and illustrated by Steven C. Oppenheimer

Cognio, Inc. ([www.Cognio.com](http://www.Cognio.com)) has developed a number of technologies to enhance the performance of wireless communications between computers. One of these technologies is an integrated circuit, SAgE, which monitors and analyzes the RF spectrum.


The SAgE chip is embedded in Sensors, which are placed at strategic locations in the vicinity of a wireless local area network. The Sensors, in turn, communicate over cables with a computer (the Server) which reports on the RF environment.

The SAgE chip (inside the Sensor) and the Server communicate using a system of messages known as the Network Spectrum Interface, or NSI, for short. The *NSI Specification* describes the messages which travel back and forth between SAgE (within the Sensor) and the Server. This excerpt includes selected pages from that document.

The *NSI Specification* is similar, if not quite identical, to the kind of documentation used to explain software APIs.



*Software Documentation*



## **The Network / Spectrum Interface**

Date: June 30, 2003  
Document Number: SW 1.0  
Release: 1.0

**Cognio, Inc.**  
**101 Orchard Ridge Drive, Suite 350**  
**Gaithersburg, MD 20878**  
Telephone: 240-944-1500 Facsimile: 240-631-1943

## 2.3 Sending and Receiving Messages: The Network / Spectrum Interface (NSI)

The point at which the ISM data stream leaves the MCU, and control commands enter the MCU, along with the entire set of protocols used for two-way communication (described in this document), is called the Network / Spectrum Interface (NSI). It is also referred to as the Level 1 interface, or L1 interface. (L0 is the interface between the SAgE technology and the MCU; most users of Cognio's software will not access the data stream at the L0 level.)

This interface can take many forms; but in concrete terms, it may be a cable carrying TCP/IP traffic from an 802.11 Access Point to a PC which is running software designed to accept the traffic for further analysis and processing. This software will be referred to generically, in this document, as Manager software (Manager software).

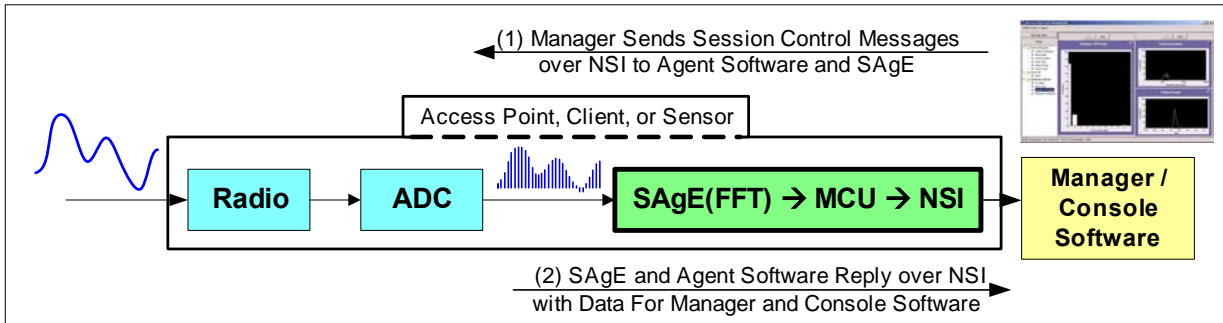


Figure 2-2: The Overall ISM Architecture Including the NSI Interface

Note that this is only a typical example. The TCP/IP traffic (or traffic using some other network protocol) could also be carried by the PCI bus inside a laptop PC, provided the PC has built-in 802.11 technology, or an 802.11 card in the PCMCIA slot. (And of course, provided the 802.11 hardware contains Cognio's ISM technology.)

If the source of the data stream is a TCP/IP connection, the Manager software would need to implement a socket, and access the correct port, to read the NSI data stream. A sample of typical code for this purpose is shown in Figure 2-2 below. (The sample is in Java, and shows client-side code.) Once the port connection to the data stream is established, the use of the data stream is determined by the Manager software itself.

```
! Open Socket and Port (Remember to first assign the correct value
! for the 802.11 device PortNumber)
Socket MyClient;
try {
    MyClient = new Socket("Machine name", PortNumber);
}
catch (IOException e) {
    System.out.println(e);
}
! Create input stream to get data from NSI
DataInputStream input;
try {
    input = new DataInputStream(MyClient.getInputStream());
}
catch (IOException e) {
    System.out.println(e);
}

! Create DataOutputStream to send control commands and
! configuration data to NSI
DataOutputStream output;
try {
    output = new DataOutputStream(MyClient.getOutputStream());
}
catch (IOException e) {
    System.out.println(e);
}
```

**Figure 2-3: Sample Java Client-Side Code to Create Socket and Port, and Create Data Streams**

The class `DataInputStream` has methods such as `read`. The class `DataOutputStream` allows you to write Java primitive data types; one of its methods is `writeBytes`. These methods can be used to read data from, and write data to, the NSI.

If the transport of the ISM data stream occurs over other low-level media, other methods are used to access the data stream. For example, if the ISM data is carried over the PC's PCI bus, a PCI device driver will typically provide access to the data.

This document indicates the messages which the Manager software can send to the NSI to initiate the flow of ISM data; and this document also describes the structure of the messages received from the NSI containing the descriptive ISM data.

## 2.4 How NSI Information Is Useful

The ultimate purpose of Cognio's ISM technology is to help wireless networks achieve the best possible RF discovery, security, performance, range, reliability and manageability. To do that, the WLAN must have built-in smarts to be aware of, understand, and adapt in real time to the changing conditions of the RF environment. ISM provides the foundation for that intelligence. There are eight broad types of information provided via the NSI:

### 2.4.1 "Raw" SAgE Data

The first five types of data can be viewed, loosely speaking, as raw data from SAgE, reflecting direct properties of the RF spectrum. Bear in mind, however, that some of this data is, in fact, "massaged" or cleaned up by Agent drivers, before being sent over the NSI.

- **RF Power vs. Frequency Spectrum** – This describes the power in the spectrum as a function of frequency (RF power per each of 256 frequency bins), over a given bandwidth.
- **Spectrum Analysis Statistics** – These hardware-generated statistics indicate the average power and maximum power in the RF spectrum over a period of time (generally on the order of 1/10 of a second), within each of 256 frequency bins. The statistics also report the duty cycle, i.e., on portions of the bandwidth that remain above a (user-defined) threshold.
- **Pulse Events** – ISM provides data on individual pulses detected by SAgE.
- **Raw "Snapshot" Data** – This is raw data of portions of the RF spectrum, in digital form. This data can help identify the location of interfering devices, and can also be used to extract detailed information such as the model of a device.
- **Pulse Histograms** – The histograms provide statistical data on the pulses, including the distribution of pulses over time; the distribution of gaps between pulses over time; and the percentage of pulses distributed among different frequencies, energy levels, and bandwidths.

### 2.4.2 Processed Data

The following data is generated by Agent software (i.e., driver or operating system level software) that processes SAgE data.

- **Events** – Based on an analysis of the above data (by the Measurement and Classification Engines), ISM can report on specific types of events, such as a Bluetooth device being turned on or off, or a cordless telephone in operation, or a denial of service attack.

While the details of this processing are beyond the scope of this document, a brief illustration in *Figure 2-4* on page 8 gives some idea of how SAgE data is used to produce more specific information – in this case, how SAgE data is used to identify specific RF sources.

### 3 Message Formats

There are two types of messages available to programmers working with the NSI:

- **Session Control Messages** – These are sent by the Manager software to the NSI to initiate tests. They are discussed briefly below. They are discussed again in detail in the final part of this manual, beginning in the section *Session Control* on page 44.<sup>2</sup>
- **Informational Messages** – These are sent by the NSI back to the Manager software, and contain the test data of interest. These messages are covered in Sections 4 through 9.

#### 3.1 Session Control Messages (Overview)

Session control messages are used to initiate, control, and terminate data collection. We mention here, in brief, some of the key features to understand about session control.

##### Note

The most important thing to know, up front, is that before you can run several different kinds of tests, you must first configure at least one of SAgE's hardware-based *pulse detectors*. This is done by sending a session control message to start running the Pulse Event test, and including the **IE\_PD\_CFG** information element in the startup message. See the following sections for more information:

- *General Radio and Session Control Issues* on page 44
- *Pulse Tests, Other Services, and Pulse Detector Configuration* on page 57
- *Starting a Pulse Event Test* on page 57
- *Pulse Detectors and Pulse Detector Configuration* on page 80

Some other general points about session control:

- All tests (Spectrum Analyzer Power vs. Frequency; Spectrum Analyzer Statistics; Pulse Events; Pulse Histograms; Spectrum Management Events; and Snapshot Messages) may be started and stopped by the Manager software via Session Control messages.
- One, several, or all eight tests can run at the same time.
- If more than one user wants to run a test at the same time, Session Control messages will contain priority information determining which user has higher control of a resource. However, even if only one user can control a resource, more than one user can watch the data stream from that resource.
- Most of the tests have various configuration parameters. These parameters are sent via Session Control messages, and they determine specific details of the test.

---

<sup>2</sup> We defer the detailed discussion of Session Control Messages to the back of the manual. It may seem backwards to discuss test results first, before showing how to set tests in motion. However, before you can know what kinds of sessions to set up, it's helpful to know the corresponding informational messages (data messages) you can receive from the NSI. So, that information is discussed earlier in the manual.

For example, in monitoring the spectrum, Session Control messages tell the NSI how wide the bandwidth should be (narrowband or wideband), and what is the center of the frequency spectrum being monitored. For spectrum analyzer statistics—which provide statistical analyses based on collected time-series of RF spectrum samples—Session Control messages tell the NSI how many cycles of Fast Fourier Transforms to include in each bundle of statistics.

So, the Session Control messages contain vital parameters which determine the specific character of the tests. In many cases—as can be seen later in this document—the headers of the Informational Messages return at least part of this same configuration information. This is strictly for the convenience of the programmer.

- In many cases, detailed test configuration parameters can be omitted from the Session Control messages. In those cases, the NSI uses default settings.

## 3.2 Informational Messages

The six informational messages which the Manager software may receive from the NSI correspond to the eight types of information described in the previous section:

- **Spectrum Analyzer Power vs. Frequency Messages** – These messages describe the total power in the spectrum as a function of frequency, over a given bandwidth.
- **Spectrum Analyzer Statistics Messages** – These messages provide a statistical analysis of the data in the RF Power vs. Frequency measurements.
- **Pulse Event Messages** – These messages provide data on individual RF pulses.
- **Pulse Histogram Messages** – These messages describe the distribution of RF pulses per unit of time, in terms of the percentage of pulses distributed among different frequencies, energy levels, and bandwidths.
- **Spectrum Management Event Messages** – These messages describe specific events, such as a microwave oven being turned on or off, or a cordless phone injecting noise into the network.
- **Snapshot Data Messages** – These messages contain portions of unprocessed data of the RF spectrum (except that the data has been processed from analog to digital). The data can help identify the location of interfering devices, and can also be used to extract ID information which can determine the brand name of certain devices.

### 3.3 Common Header Fields

All messages sent from the NSI to Manager software begin with certain common header fields. A summary of that header information is repeated in each section of this manual, in the context of the specific messages; here, however, we present a more detailed view of these common headers.

#### 3.3.1 Common L1 Message Header: StandardMsgHeader

The Common L1 Message Header is used by *both* control messages and informational messages . It comes at the very beginning of all messages. The **StandardMsgHeader** header provides certain general identifying information for the message:

Sub-Field Name	Offset	Size	Format	Description and Notes
msgLen	0 / 0	2	uint16	'msgLen' is the length of the message in bytes, includes <i>all</i> fields in the message – <i>including</i> the complete 'StandardMsgHeader' itself, and the 'msgLen' field itself.
msgType	2 / 2	2	uint16	'msgType' and 'sessionType' are dependent on the subsequent fields, that is, on the overall message type. All L1 informational messages have a 'msgType' of SM_MSG_L1_INFO (decimal 46). 'sessType' values are indicated for the individual message types later in this manual.
sessType	4 / 4	2	uint16	
configToken	6 / 6	2	uint16	The value of the 'configurationToken' is determined by the user (via the Manager software) when a test is set up. The purpose is to help the Manager software distinguish incoming data based on different test configurations. Example: The user starts running a particular test, say, Pulse Histograms, at a particular center frequency. The Configuration Token may be set (by a control message) to a value of '1'. All information messages returned by the NSI then have a configToken value of '1'. At a later time, a message from the Manager software to the NSI says, "run this Pulse Histogram test over a different center frequency." As part of that message, the Manager software changes the configToken to '2'. In this way, return messages (that is, test data) for the first bandwidth are distinguished by a configToken of '1', while test data for the second bandwidth are distinguished by a configToken of '2'.
timestampSecs	8 / 8	4	uint32	Timestamp in seconds, and Fractional portion of timestamp in microseconds. The use of the 'timestampSecs' and 'timestampUsecs' is message dependent. These fields are typically filled in/updated by the transport mechanism to provide timing for the message. The actual information within the message may contain other timestamps which were collected at the originating hardware target (SAgE). NOTE: The time in seconds and μseconds is counted from when testing began, and not from some universal "zero time."
timestampUsecs	12 / 12	4	uint32	
Source	16 / 16	2	uint16	'Source' and 'Destination' fields are intended to facilitate multiplexing of session routing across common transport connections. The exact use of the fields is yet to be determined.
Destination	18 / 18	2	uint16	
Reserved	20 / 20	8		Reserved for future enhancements
<b>StandardMsgHeader length : 28 bytes</b>				

### 3.3.2 Common L1 Info Message Header: InformationHeader

All L1 *Informational* messages are started with two headers: The Standard Header (**StandardMsgHeader**), which starts all L1 messages, and which was discussed immediately above; followed by the Info Header (**InformationHeader**), which is discussed here. The InformationHeader header provides specific identifying parameters for information messages:

Sub-Field Name	Offset	Size	Format	Description and Notes
transactionSeq	0 / 28	4	uint32	Sequence for this message. This starts at 1, and is incremented for each succeeding message. The increment reflects the number of data samples (transactionCnt) in the previous messages. For some types of messages the number of data points, and hence the transactionCnt, is fixed at '1'; for these message types successive messages always have their transactionSeq incremented by '1'.
transactionCnt	4 / 32	2	uint16	'transactionCnt' generally indicates the number of entries in a message, where entries are discrete units of data. Its use is message dependent. For example, for Power vs. Frequency spectrum messages, this value indicates the number of sequential "snapshots" of the RF spectrum in the message. (Each snapshot is encapsulated in a specific sequence of bytes. If the transactionCnt has a value of 10, then the message contains 10 successive snapshots of the RF spectrum; there are ten matching byte patterns which follow, each of which reports on one snapshot of the RF spectrum.)
Reserved	6 / 34	6		Reserved for future enhancements
<b>InformationHeader Length : 12 bytes</b>				

### 3.3.3 Combined Header for Informational Messages

Since all L1 Informational messages always use the two above headers together, all such messages have the following combined header format:

Field Name	Offset	Size	Format	Description and Notes
<b>StandardMsgHeader</b>	0 / 0	28		Standard Header
<b>InformationHeader</b>	28 / 28	12		Info Header
<b>SM1 Standard and Info Headers Len : 40 bytes</b>				

Additional fields, with specific sub-fields, follow; the exact structure is dependent on the nature of the particular message. These are discussed in the remainder of this manual.

## 10.2 Session Control Dialogs

Initiating a session, sending and receiving data, and terminating a session involves a dialog between user and client. Here we illustrate some typical dialogs between user and client, using message sequence charts. The charts indicate a number of special *session control messages*, such as `SESS_START_REQ`, `SESS_STOP_REQ`, and `SESS_STOPPED_RSP`, along with the informational messages documented earlier in this manual. These session control messages are actually sent as specific numeric values (such as '40' for `SESS_START_REQ`) in the `msgType` data field of the Standard Message header (see *Common L1 Message Header: StandardMsgHeader* on page 13).

The charts illustrate some representative sessions, though other types of session control exchanges are possible. (For example, tests may be reconfigured while in progress.) The session control dialogs indicate the flow of messages that you need to program between your application software (Manager software) and the NSI. Following this subsection, we describe in detail the session control messages.

### 10.2.1 A Successful L1 Session

This message sequence chart shows an exchange where a test is successfully initiated by the user; data is sent from the service to the user; and the test is terminated by the user.

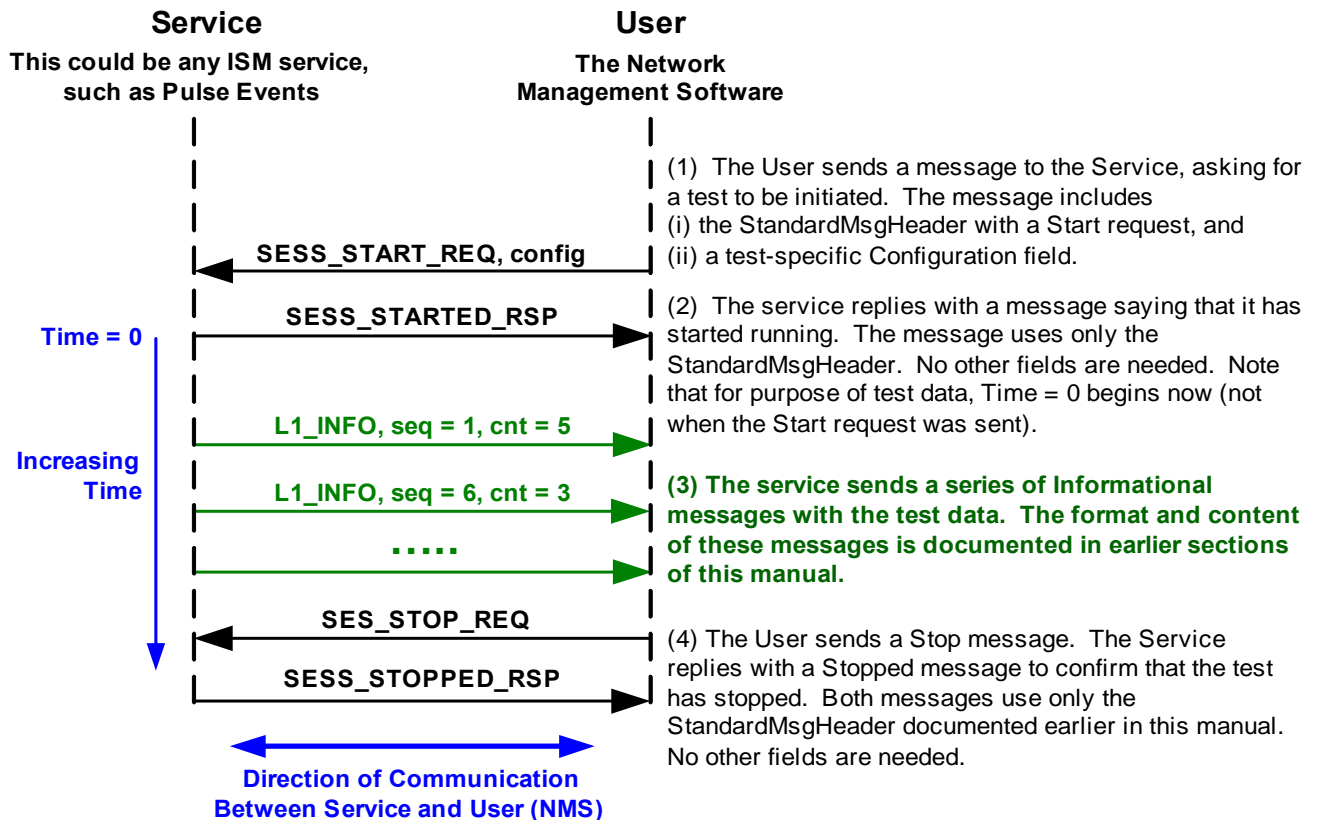


Figure 10-1: A Successful L1 Test Session Dialog

### 10.2.2 Session Starts Successfully, But Messages Are Skipped

A session may start and terminate successfully, but some data messages are lost along the way. The way to determine this is by having the Manager software keep track of the *expected* sequence number of each informational message, and compare that with the *actual received* sequence number. The sequence number is always contained in the **transactionSeq** sub-field of the **InformationHeader** field that begins each information message. **InformationHeader** also includes the **transactionCnt** sub-field, which indicates the number of data entries in the message. To calculate the expected sequence number for the next message, simply add:

$$\text{Expected InformationHeader.transactionSeq in the Next Message} = \text{InformationHeader.transactionSeq}[\text{this mssg}] + \text{InformationHeader.transactionCnt}[\text{this mssg}]$$

An example where a message has been lost is shown in the following diagram:

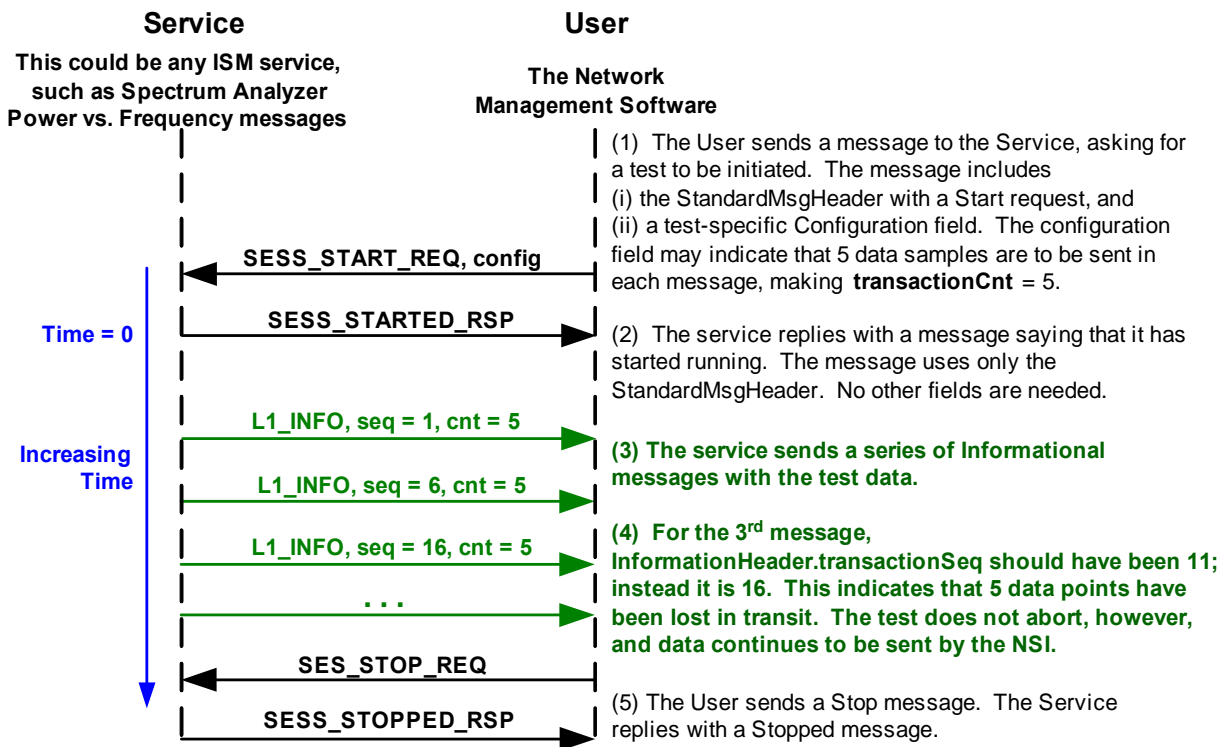


Figure 10-2: Skipped Informational Message

**Note** If there is a difference between the expected **transactionSeq** for a message, and the actual value of **transactionSeq** for the message, the difference always indicates the number of data points lost. For a Pulse Events message, if the difference is 7, it means you have not received data on 7 pulses that actually occurred.

A typical reason for lost data is a transport buffer overrun. For example, the NSI may be producing data faster than a TCP/IP transport can carry the data. The NSI will try to buffer data; but if it runs out of buffer memory, it simply dumps some data points.

### 10.2.3 Session Rejected

Another possible outcome is that the request to start a test is rejected. This could happen for any of several reasons; for example, another Manager software program with higher priority already has control over the service. In this case, a simple exchange of messages occurs:

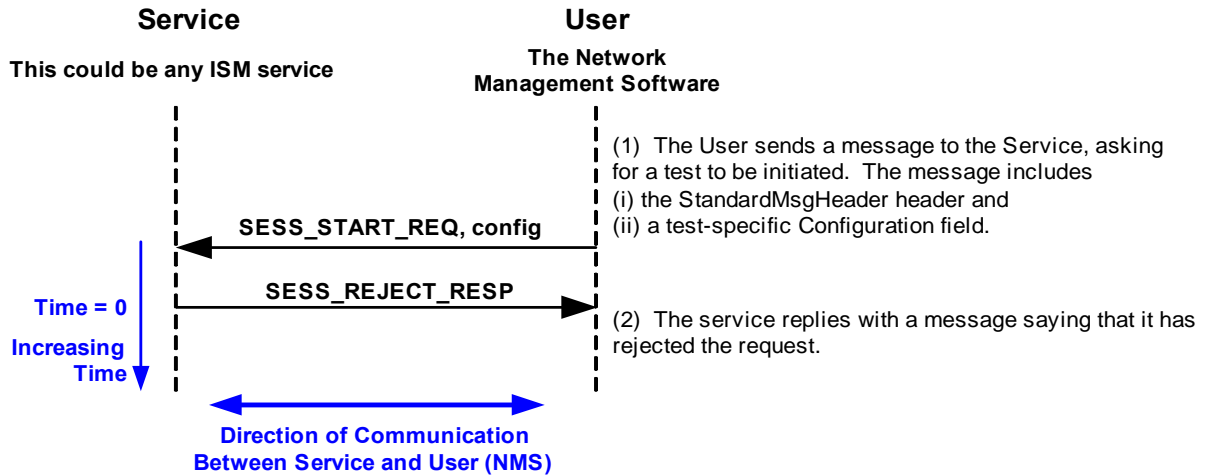


Figure 10-3: A Test Request is Rejected

### 10.2.4 Session Aborted

Sometimes ISM will terminate the session without a user request. This may happen because another user with higher priority has taken over the service.

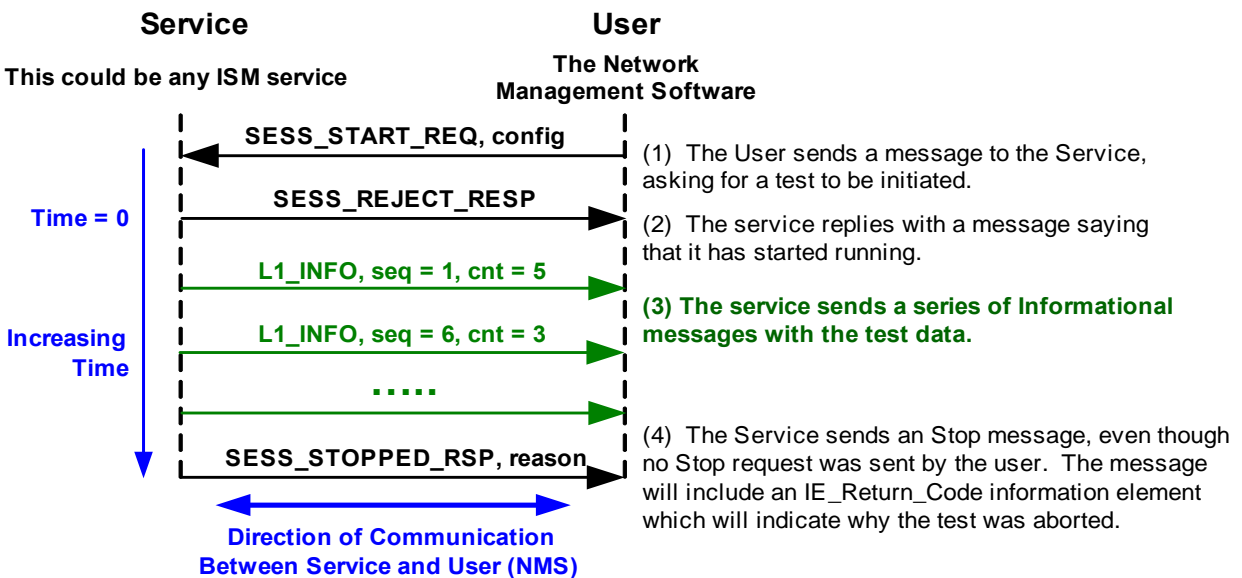


Figure 10-4: A Test is Aborted by the NSI


*An excerpt from...*

## **The Cognio CBF White Paper**

*A Cognio Publication*

Written and illustrated by Steven C. Oppenheimer

Cognio, Inc. ([www.Cognio.com](http://www.Cognio.com)) has developed a number of technologies to enhance the performance of wireless communications between computers. One of these technologies is called Composite Beamforming, or CBF. CBF, in turn, is an advanced variation on MIMO, or multiple-in-multiple-out antenna technologies, which use multiple antennas to boost RF signal gain and signal-to-noise ratio, without any need to increase the overall signal strength. This white paper introduces Cognio's CBF—including the most cutting-edge version, Vector CBF—and explains how CBF is an improvement over standard MIMO technologies.



# CBF and Vector CBF: Cognio's Enhanced MIMO Beamforming

Release 2, Version 0.33

**Contact:**

Cognio, Inc.  
101 Orchard Ridge Drive, Suite 350  
Gaithersburg, MD 20878  
marketing@cognio.com  
301-944-1500

Date: January 8, 2003

Copyright 2003 Cognio, Inc. All rights reserved.

## Contents

The Wireless Challenge: Helping Technology Catch Up To The Market .....	1
CBF Illustrated .....	3
Limitations to RF Signal Transmissions .....	6
Understanding CBF .....	11
Vector CBF: Achieving Maximum Path Utilization .....	16
CBF: The Adaptable, Standards-Compliant Solution for Maximum MIMO Benefit.....	20
Appendix A: CBF as Enhanced Beamsteering .....	24
Glossary .....	28

---

## The Wireless Challenge: Helping Technology Catch Up To The Market

New technologies have a way of taking the modern world by storm. Ten years ago, no one except scientists and engineers had heard of the Internet. Today, hundreds of millions of people use the Internet for everything from routine social communications to keeping up on the latest weather, news, and entertainment.

Similarly, ten years ago the only wireless technology that most people had contact with was their television or radio. Yet in the past five or six years cell phones have become ubiquitous, and just the past two years have seen the dramatic emergence of wireless computer networks in offices, homes, and even airports and outdoor cafés. The public market for wireless technology is expected to grow from approximately \$1.9 billion in actual sales of APs and NICs in 2001, to an estimated \$5.2 billion by 2005 (with a corresponding boom in WLAN chip sales from 10 million chips in 2001 to over 60 million by 2005).

Inevitably, along with the excitement and convenience of these new technologies comes a fair share of accompanying headaches. Sometimes these are social in nature – spam e-mails, and drivers who are more preoccupied with chatting on the cell phone than keeping an eye on the road. These problems can be addressed by education, and sometimes by laws.

Other problems, though, are purely technical in nature, reflecting the fact that newly emergent technologies often capture the market before they are fully mature from an engineering standpoint. These problems lend themselves to innovative technical solutions.

---

### In Search of Maximum Data Rate and Range

A case in point are the data rate, data reliability, and range limitations on wireless technologies. 802.11b networks have a theoretical data rate of up to 11 Megabits per second, while 802.11a/g has a theoretical rate of up to 54 Megabits per second. In practice, all three wireless networking protocols are fortunate when they can achieve half of their theoretical maximum, and the actual speed falls off rapidly with increasing distance between an Access Point (AP) and a client device (station). And even at optimum performance, current WiFi devices cannot begin to compete with a wired Ethernet network.<sup>1</sup>

WLAN data rates are also negatively impacted by other technologies that share the unlicensed band with 802.11 technologies, including Bluetooth devices, and cordless headphones and headsets. Even microwave ovens generate RF energy in the same

---

<sup>1</sup> A recent PC Magazine article compared the file transfer time for a 50 Megabyte file: Between 2 and 3 minutes for 802.11b systems, about 45 seconds for an 802.11a link, but just 8 seconds for a 100 Mbps Ethernet connection.

bandwidth, adding further static, which causes WiFi frame retransmissions... which still further slows down the real data throughput of 802.11 networks.

These data rate limitations become more pronounced precisely as wireless LANs become more popular – increased device utilization places greater stresses on APs, which must track multiple clients. Moreover, while data rate limitations merely impose longer waiting times for such pedestrian tasks as getting e-mail, viewing Web pages, or downloading files; these same rate limitations can render streaming data applications—VOIP and real-time wireless video—completely impractical.

A parallel challenge faced by wireless computer networks, in addition to data *rate* limitations, are *range* limitations. Computers can only be carried a limited distance from the fixed network APs before signal strength drops off to unacceptable levels (and data does not get through reliably). Once again, theory and practice have proven to be very different. In theory, 802.11a should be able to achieve distances (between AP and station) of up to 25 meters, while 802.11b/g should be able to achieve distances of up to 30 meters. In reality, fading due to multipath signal interference often results in substantially reduced range. The results are rate drop-offs and lost connections.

Current performance levels are simply not adequate. Whether a WLAN is deployed in an office, a factory, or a public setting, it is vital that mobile clients receive consistent levels of signals, with the maximum possible rate-at-range, and minimal interference.

*Wireless vendors who meet these technical challenges will reap the benefits of the billions to be spent on WLAN deployments.* Cognio is developing cutting-edge technologies to help our partners—vendors of 802.11 chipsets, APs and NICs—meet the technical demands of future WLAN consumers.

- *Composite Beamforming*, or CBF, is Cognio's state-of-the-art, patent-pending Multiple-In-Multiple-Out (MIMO) system solution. CBF is a fully standards-compliant solution—based on Cognio's CR2200 RF IC and custom Baseband chip enhancements—which resolves the problems of limited distance and signal fading, providing dramatic gain enhancements compared to standard 802.11 networks.
- *Vector CBF* is Cognio's unique technology which can transmit up to four 802.11 frames in the same time normally used to transmit a single frame. (The frames are actually superimposed on each other, and transmitted in parallel.) When Vector CBF technology is employed on both ends of an 802.11 link, the result can be up to a four-fold increase in data rate.

APs and stations using CBF are fully interoperable with APs and stations that are not CBF enhanced. The use of Vector CBF requires that Cognio's technology be present on both ends of the link; but it is straightforward to build APs and stations that are dual CBF/Vector CBF capable. This ensures maximum benefits when Cognio's Vector CBF technology is employed on both ends of the communications channel, while still guaranteeing full interoperability with standard 802.11 APs and stations.

This white paper describes the basic principles behind both CBF and Vector CBF, and demonstrates that these are cost-effective solutions for radically improved 802.11 rate and range performance.

CBF Illustrated

CBF, or Composite Beam Forming, is Cognio’s WLAN signal maximization technology, based on a multiple-in-multiple-out (MIMO) antenna architecture and advanced signal processing techniques. Before diving into the science and engineering behind CBF, it’s worth examining the practical benefits up front. We’ll take a look at what is revealed about CBF both by computer modeling, and by actual measurements when CBF was put to work in Cognio’s own offices.




Figure 1: CBF Simulation Results

Figure 1 above shows the results of computer modeling of rate-at-range for four different 802.11 configurations. The first configuration (in red) is with a standard 802.11 AP in communication with a standard station (network interface card, or NIC). The second configuration (blue) is for a 4-CBF AP—an AP using four antennas, which are controlled by Cognio’s CBF technology—and a standard station. The third configuration (gray) shows the results when a 4-CBF AP communicates with a 2-CBF (that is, two antenna) station. The fourth configuration (purple) shows the rate-at-range results when the 4-CBF AP is communicating with a 4-CBF station.

If we look at just a single data rate measurement of 30 Mbps (shown by the horizontal dotted line), we see that the range improves from a limit of about 35 meters for the base configuration (with no CBF) to nearly 115 meters with the optimal CBF configuration. But even if only the AP uses CBF, and the station has a standard NIC, we still see that the rate-at-range almost doubles to 70 meters.

Vector CBF yields still further improvements, as show in Figure 2 below. The figure compares the Signal-to-Noise ratio at various data rates, using different antenna schemes. (OFDM is the modulation scheme used in 802.11a/g transmissions. SD stands for “selection diversity”, a common antenna scheme discussed further below.)




Figure 2: Vector CBF Simulation Results

The key point is that while CBF already achieves significant performance gains over standard 1- and 2-antenna technologies, VCBF using two antennas achieves a data rate of up to 108 Mbps – double the maximum data rate defined under current standards. (Rates approaching 216 Mbps are possible using 4 antennas on both ends of the link.)

Computer simulations are intriguing, but do they hold up in the real world? *Figure 3*, below, shows the results of measurements taken in Cognio’s development offices in Gaithersburg, Maryland, with the fixed AP near the center of the brown (central) shaded area. A laptop PC with NIC card was wheeled around the office to test the range for successful communications between AP and NIC at a data rate of 54 Mbps.

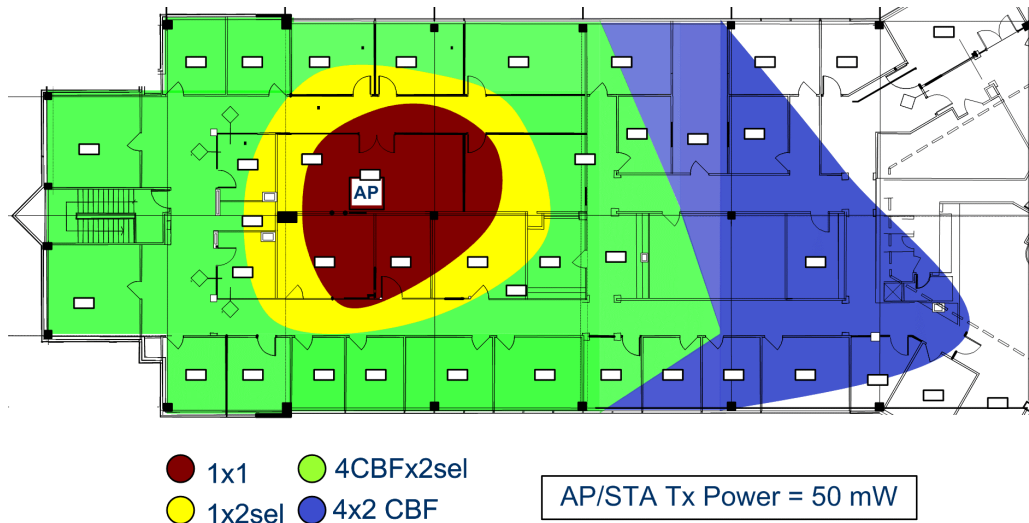


Figure 3: Coverage With and Without CBF

The inner shaded area (in brown) shows the transmission range that was achieved using a single antenna on the AP and NIC. The yellow area shows the additional coverage achieved when the NIC employed two antennas for dual selection diversity. Dual selection diversity (discussed further below) is a relatively simple scheme where two antennas are used, and the NIC selects the signal from whichever antenna is achieving the best signal. (Note that dual selection is not considered a MIMO scheme, since multiple antennas are not actually in use at the same time – it's one or the other.)

The green shows the additional coverage that was achieved when the AP used four antennas for a 4-CBF configuration. The most extensive coverage—spanning the brown, yellow, green, and blue areas—was achieved when the AP employed 4-CBF technology, and the NIC used 2-CBF technology. Two key points to note:

- The best coverage—which includes the blue—is achieved when both the AP and NIC employ CBF technology. (It should be clear that the blue area would have been much larger if Cognio's offices were longer and wider, or if the AP were placed in a more central location.)
- *Even if only the AP uses CBF technology, a remarkable improvement in range is still seen (green area).* This is important, because many customers may want to initially upgrade their networks only by replacing their APs, while not wanting to foot the bill for replacing all their NICs (especially if they have many computers).

Now that we've seen the dramatic improvements in rate-at-range that can be achieved with CBF, we'll take a quick look at the science behind it. At the conclusion, we'll examine how Cognio's CBF technology stands out not only for the excellence of the RF engineering which underlies it, but also because of its cost effective design.

## Limitations to RF Signal Transmissions

What are the basic factors that limit radio frequency (RF) signal transmission?

- **Distance Fading** – Radio frequency signals fall off by a factor of  $1/(\text{distance})^2$  in the vicinity of a transmitting antenna. Beyond a distance of a few meters, however, objects within an office (walls, metallic objects, and even people) intercept some of the RF waves, so the decrease in signal strength is even more pronounced.
- **Multipath Fading** – When an RF signal leaves an antenna, part of the signal travels directly to the target receiver. However, the RF wave is also reflected off of multiple media in the nearby environment; these reflections then find their way to the target antenna via different paths.

When two (or more) RF waves “meet” at a single point, such as an antenna, they may happen to be synchronized with each other (like people on the same end of a rope in a tug-of-war, all pulling in the same direction); in this case the waves add together, increasing their overall power, a phenomena called constructive interference. On the other hand, if the waves arrive out of synch with each other (like people on the opposite ends of the tug-of-war rope, pulling against each other), their mutual energy cancels; this is called destructive interference.

Once again, RF waves from a source (like an AP) are reflected throughout the physical environment. If the reflections, plus the original signal, happen to interfere constructively when they arrive at the target antenna (at a station), the result can actually be a boost in signal power. However, it is possible—depending entirely, and somewhat randomly, on the particular locations of AP, the station, and the sources of reflection—that the same signal arrives from multiple directions in such a way that it interferes destructively. The original signal and its reflected images partly or totally cancel each other, and the receiving antenna does not get a strong enough signal for the receiving station to detect the data.

- **Delay Spread** – An 802.11a/b/g signal is transmitted over a range of frequencies. In turn, different radio frequencies travel at different speeds through material objects. So, if the RF signal travels through physical objects (walls, filing cabinets, people) on the way from AP to station, the different frequency components of the same signal can arrive even more out of synch with each other. In this case, the problem becomes known as *delay spread*.
- **Co-channel Interference** – Each receiver in a WLAN is set to receive on a specific RF channel – naturally, the same channel that the transmitter is using. Co-channel interference occurs when a receiver picks up an RF transmission from another, nearby transmitter—not the one it intends to listen to—which is using the same frequency. (The second, undesired nearby transmitter may be part of another WLAN in the office space of an adjacent company.)
- **Interferers** – An interferer is any device which is not on the WLAN network, but which is also broadcasting in the unlicensed band. This includes Bluetooth devices, cordless phones, cordless headsets, and microwave ovens.

---

## The Unworkable Solution

The designer of WLAN hardware, then, is faced with multiple sources of signal loss. It may seem that one obvious solution is simply to boost the power of the transmitted signal. There are four reasons that this is not the answer:


- In order to prevent unlicensed band operations from becoming completely unmanageable (due to multiple transmitters blasting signals at the same frequencies), government agencies have placed strict limits on transmission power.
- Increasing the RF power does not help significantly with multipath fading. It's like putting stronger people on each of the rope in a tug-of-war – their mutual pulling still cancels out.
- Even if it were legal to increase the signal power arbitrarily, this would make each transmitter a “bad neighbor” to other members of the network and to neighboring networks. Co-channel interference would become unmanageable.
- Boosting signal strength would do nothing to address delay spread.

---

## The Current Solutions: Switched Diversity, and MRC

Cognio, of course, is not the first company to notice that distance and fading effects can reduce range and/or transmission speed for WLANs. Currently, however, there are only two solutions that exist in the marketplace, *switched diversity reception*, and *MRC*. We'll take a quick look at both.

The wavelength of unlicensed band RF signals is on the order of a few inches. Because of this, multipath fading effects can vary over distances of a few inches. For a given set of locations for AP, stations, and nearby reflecting objects, a receiver with a single antenna might have very different performance if it were moved just a few inches. At *this* spot you get destructive interference, but at this *other* spot two inches away—that is, two inches away in the right direction, if you happen to know the right direction—you get constructive interference.



Of course, no one using a laptop PC wants to constantly adjust the position of their computer to get better network performance. But because multipath effects can vary over the space of a few inches, it's not necessary to physically move a single antenna. Instead, two antennas just a few inches apart can receive very different signal strengths. One antenna might happen to be at a spot where there is virtually no signal at all, while the other antenna a few inches away receives a strong signal.

**Figure 4: Switched Antenna Diversity**

Switched diversity reception takes advantage of this aspect of multipath fading. The AP and/or the client each have at least two antennas (more can be used). The receiving circuitry monitors all the antennas, and selects the signal from the antenna giving the strongest signal; or, the circuitry may be designed to detect the signal from the antenna which yields the most favorable signal-to-noise ratio.

In practice, even if an AP and a station remain in exactly the same place, multipath fading phenomena vary in time. To compensate for this, the receiving circuitry may frequently shift between the two (or more) receiving antennas that are in use.

Although helpful up to a point—and although simple to implement from an engineering design standpoint—switched diversity has several disadvantages.

- There is no particular guarantee that *any* antenna will actually be in an optimal location, or near optimal location, in respect to multipath fading. All the antennas may be in locations where multipath fading is inhibiting signal reception.
- Switched diversity does nothing to address the overall signal decrease with distance. The farther the station gets from the AP (and vice-versa), the weaker the signal will be. Beyond a certain point, a single antenna just won't cut it.

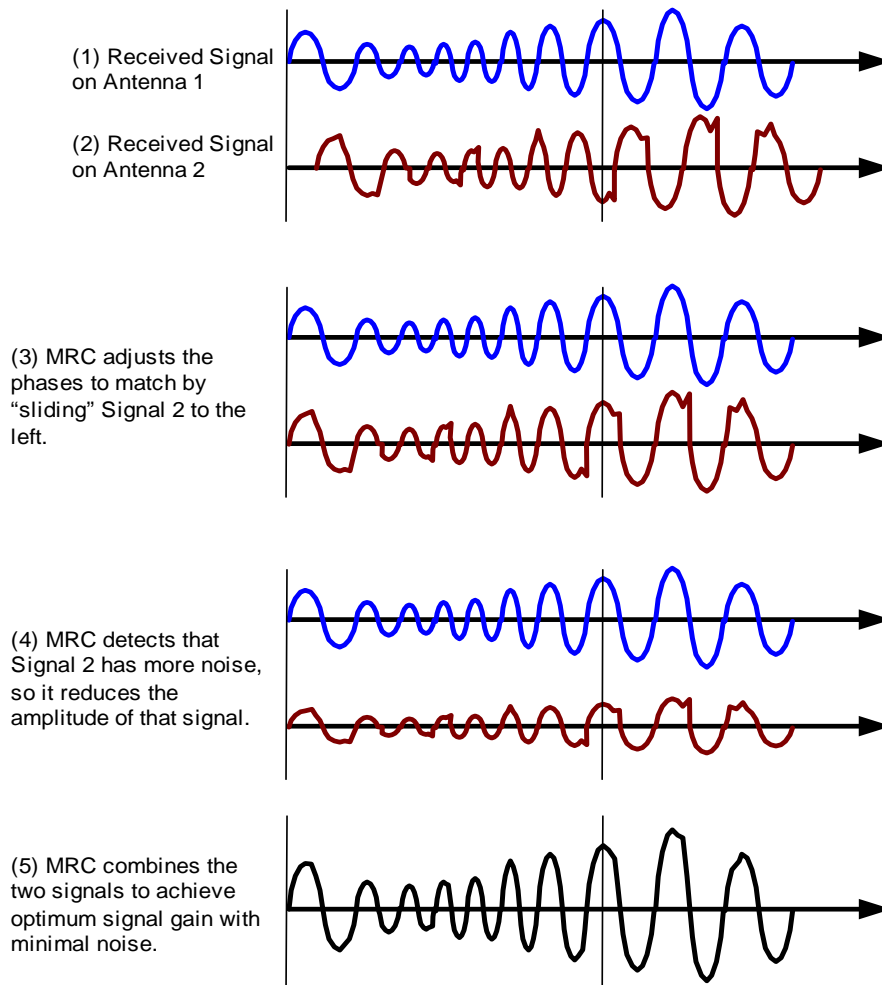
But then, another suggestion comes to mind: As long as two or more antennas are in use, why use only one at a time? Why not *combine* the signals received at all the antennas? Especially in the worst case scenario, where all the antennas are in locations where the signal is low, it make sense to combine whatever signal *is* received by each antenna. A strategy known as *maximal ratio combining*, or *MRC*, does just that.

Simply adding the signals at each antenna is actually not the best solution, for two reasons. First, the signal at one antenna may have a worse signal-to-noise ratio than the signal at the second antenna. (A worse signal-to-noise ratio means the antenna is contributing relatively more static, and less of the actual desired signal.) So, it's sometimes best to let one antenna contribute more of its received signal energy than the other antenna. We get some information from the signal that has more noise, but more information from the signal that has less noise. Mathematically, the combined signal might look something like this:  $S_{\text{Total}} = 60\% \times S_1 + 100\% \times S_2$ .

Also: Just as two versions of a radio signal (the original version, and a reflected version) arriving at the same place can be out of synchronization with each other; the same wave, arriving at each of two different antennas, can also be out of synchronization. So, before combining the signals, the electronics in the receiver lines the waves up – high parts of the wave match high parts, and low parts match low parts. These adjustments to the signals on each antenna—adjusting the magnitude, and adjusting the synchronization—are called *weighting* the antenna signals. The formula for MRC, in abstract, looks something like this:

$$S_{\text{Total}} = W_1 \times S_1 + W_2 \times S_2 + W_3 \times S_3 + \dots$$

There is one term for each antenna; “ $S_n$ ” is the signal at an antenna; and each “ $W_n$ ” is a complex mathematical weighting term that includes both a relative strength, and also a kind of sliding or phasing factor that re-aligns any signals that were not synchronized.



**Figure 5: Phase and Amplitude Adjustments in MRC**

How does the receiving electronics “know” how to align the signals from two antennas? And how does it recognize which signal has less noise and more actual signal content? It turns out that all 802.11 frames include a standard sequence of bytes, variously called a “training sequence” or “synchronization sequence”. The receiver electronics can align these standardized bytes in order to align the two signals; and it can also determine how much signal there is, compared to how much noise, based on these standard byte sequences.

Human hearing provides a partial analogy to MRC: When someone first enters a room and begins speaking, you may not be facing them. But the fact that you recognize their voice at all—their voice is a familiar “standard”, based on past acquaintance—enables you to adapt your listening process to their presence. When your ears hear the person off to one side, the biological signal processing machinery in your brain uses the signal from both ears to determine the direction of the sound. You then turn to face the person for improved hearing.

While MRC does not physically move the antennas, the process of adjusting the phase and signal strengths of the incoming signals—with unique weight adjustments for each antenna—has a similar effect. However, because RF signals spread out throughout a

public space, and because they reflect in such complicated ways, a receiver using MRC is really optimized to receive from a particular *location* in the environment, rather than along a particular direction.

The overall result—once the weight-adjusted antenna signals are finally added together—is that the antenna array is optimally tuned to detect signals coming from a particular transmitter, at a particular location in the current environment.

MRC, which combines signals from two or more antennas, is a significant improvement over any one-antenna solutions (including dual selection diversity), but MRC still has two limits:

- MRC does not take into account the fact that 802.11 signals are sent at not just one frequency, but at multiple frequencies at the same time. It turns out that the optimum weights have to be adjusted not just for each antenna, but also for each frequency received on each antenna.
- As long as we are using two antennas for reception, why not use both for transmission? But it turns out that, even if we had calculated the optimal antenna weights for use in signal reception, these are *not* the same as the optimal weights that should be used for signal transmission from multiple antennas.

Which brings us, finally, to Cognio's Composite Beamforming, or CBF.

**Understanding CBF**

When more than one antenna is used in a wireless device, there is a general “trick” for getting the most performance out of those multiple antennas: Before combining the signals, multiple the signal on each antenna times the optimum “weight” for that signal. In turn, the “weight” is a mathematical factor which reflects not only the amplitude (or strength) of the signal; but also how the signal is aligned or synchronized with signals from other antennas.

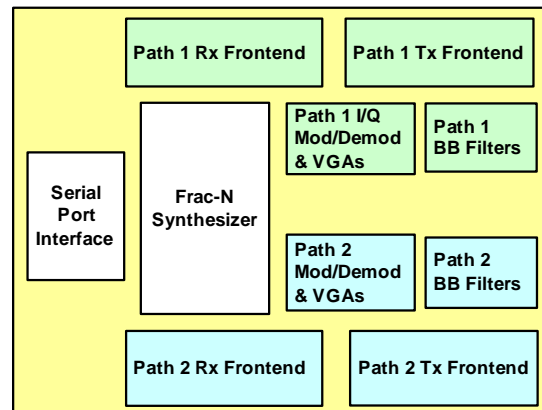
Optimum weighting ensures that the signals from the several antennas combine with both maximum signal power and minimum signal noise. When this is achieved, the result is that the receiver obtains close to the maximum theoretically possible signal strength (given the current environment, and the allowed transmission power). The goals of Cognio’s Composite Beamforming, or CFB, are three-fold:

- Ensure that the optimum weights are used not only during signal reception, but also during signal transmission.
- Determine the optimum weights as rapidly as possible, in a manner which does not interfere with, slow down, or impede standard 802.11 operations.
- Ensure that the optimum weights are applied on a per-frequency basis, as well as a per antenna basis.

One of the reasons that, until now, more advanced multiple antenna strategies have not been used in WLAN architectures is because they are technically complex to design. The math is complicated—those “S” (signal) and “W” (weight) terms shown earlier conceal a lot more detailed math buried within them—and engineering a practical and cost-effective chip design is far from straightforward. Cognio has succeeded in meeting all of these engineering challenges, as described in this section.

**The CR2200 and Transmit MRC**

The central feature of Cognio’s CBF technology is that optimum signal weights are applied not only during signal reception, but also during signal transmission. We call this aspect of CBF “transmit MRC”, and Cognio is the first and only company to offer this technology in the WLAN market. Transmit MRC, in turn, is made possible because of Cognio’s advanced CR 2200 radio chip.



**Figure 6: CR2200 Simplified Block Diagram**

Cognio's CR2200 is a highly integrated, industry leading MIMO RF transceiver targeted for mainstream 802.11 wireless applications (for both 2.4 and 5 GHz systems). With CBF, the CR2200 enables 802.11 WLAN products to achieve vastly enhanced range and rate performance for both signal reception *and transmission*. The CR2200 delivers these rate-at-range gains within a cost-optimized framework that is 100% standards-compliant with all existing and emerging 802.11 specifications.

As enabled by the CR2200, there are two crucial advantages to using MRC for signal transmission:

- Even if MRC is being used by the receiving device, when MRC is also used by the transmitting device, the receiver achieves even better signal gain.
- When MRC is used on the transmission side of the link only (while the receiver has only a single antenna, or uses only dual selection diversity), the receiver will still enjoy major gains in range and rate performance.

Two points may seem puzzling:

- First, the reader may wonder how a single-antenna *receiver* can benefit when MRC is being used only by the *transmitter*? The technical answer is that the transmission channel is *symmetric*, so that an RF pattern which works best in the receive direction also works best in the transmit direction. A more intuitive answer, in a nutshell, is this: By adjusting the weights of the signals sent from its multiple antennas, the transmitter can, in essence, focus or target maximum signal power on the receiver. If this seems hard to visualize, the reader is referred to Appendix A of this white paper, where we discuss a closely related (but slightly less advanced) concept called “beam steering.”
- A second puzzling point might be: Why is it that other companies are offering MRC on the receive side only, but Cognio's CBF features transmit MRC as well? The answer, again, lies in the radio chip. It's easy to build a radio chip which can receive signals from two antennas; alternatively, a station or an AP can be built with two radio chips (one for each antenna), and the signal from each radio can be combined in the Baseband chip.




Figure 7: A 4-CBF AP

Building a radio chip that can transmit on two antennas at a time, using two radio paths—such as the CR2200—is technically more demanding. In particular, handling the power load involved in dual-antenna radio transmissions is a technically daunting task, and to date Cognio is the only vendor in the 802.11 market that has succeeded in meeting this technical challenge.

It's notable that, in Cognio's CBF technology, two CR2200 radio chips can work in parallel to support 3-antenna or 4-antenna CBF solutions.

## Iterative Antenna Signal Processing for Rapid, Optimal Weight Convergence

The same technical issues which made CBF challenging for Cognio’s engineers to implement also make it challenging for us to explain in a white paper – especially a white paper which is not going to dive deep into the mathematics. However, it’s useful to describe the process in general terms:

We already saw above that, on the receiving end, a receiver can determine the optimum signal weights for the incoming signal, based on analysis of the standardized training sequence in every 802.11 frame. The situation is a little more complicated on the transmitter side. Since the transmitter does not know how the RF world “looks” to the receiver—put another way, it has no idea of the transmission conditions between itself and the receiver—it can initially have no idea what weights to use for the signals it sends.

Cognio’s CBF solution therefore involves an iterative process (also called an *adaptive algorithm*), which works something like this:

- The transmitter (say, an AP) sends a frame over two (or more) antennas, using arbitrary weights on each of the two antenna signals.
- The receiving CBF device (a station) calculates the optimum weights to receive the signal it gets. (This is a standard, well-known MRC calculation.)
- The receiver then does some fancy Cognio math—the details are not important—and based on the optimum antenna weights it calculated for the *received* signal, it calculates a *likely, pretty darn good* set of weights for the signal it will next *transmit*. The receiving device now becomes a transmitter, and sends a frame.
- The initial transmitter, the AP, is now the receiver. It gets the new frame back, and calculates the optimum receive weights (again, a standard MRC calculation).
- Now the AP does Cognio’s fancy math (again based on the receive weights it just used), to arrive at a *pretty darn good* set of weights for its next transmit signal.

Back and forth, back and forth, frames travel between the AP and the station. Each time, the receiving device uses the weights it has just calculated—to optimize the signal it just received—as a basis for a *new* set of calculations to determine an improved set of transmit weights. *Cognio’s engineers have demonstrated that this iterative process rapidly converges (within three or four exchanges) to the theoretically optimum values for weighting the transmitted signal on each antenna.*

There are a couple of crucial strengths to Cognio’s CBF technology:

- *Arriving at the optimal signal weights for CBF signal transmissions does not interfere with, or in any way impede, standard 802.11 operations.* The signals which are traded back and forth are standard frame transmissions—data frames, management frames, etc.—that would be transmitted anyway in the course of normal 802.11 activities. Absolutely no frame overhead is used by this process, and as a consequence there is no impact—no negative impact, at least—on performance. (There is, of course, the tremendous positive benefit of the rate and range gains from the optimized weight adjustments.)

- Neither end of the link has to know whether or not the other end of the link is using CBF. In other words, the network as a whole does *not* have to be somehow redesigned to become “CBF-aware”. The CBF solution is fully compliant with, and interoperable with, non-CBF devices and standard network protocols.
- If, in fact, only one end of the link (say, the AP) is using CBF, while the other end of the link is using a more standard technology (single antenna, or simple MRC), the CBF side of the link still hones in on the optimum achievable weights for signal transmission.

In short, CBF is a no-lose, all win proposition, whether used by a single network device or multiple network devices.

---

### Applying Optimum Weight Adjustments on a Per Frequency Basis

In addition to adjusting the weight on transmitted signals on a *per antenna* basis, a key technical achievement in CBF is that the weight adjustments are custom-tuned for *specific transmit frequencies*.

- 802.11a signals are sent using OFDM, or orthogonal frequency division multiplexing. This means that an 802.11a signal is simultaneously transmitted on 52 separate sub-carrier frequencies. As we noted above, *delay spread* is the result of these different frequencies traveling at slightly different speeds, and so arriving at the receiving antenna out of synchronization (technically, “out of phase”) with each other.

Now recall: CBF depends on adjusting the weights of the transmitted signals from each of several antennas. These weights include adjusting the relative phases of the signals from each antenna. If these weight adjustments are made by analog circuitry in the RF IC, then (by the very nature of analog electronics) the entire signal—all 52 frequency components—carried on a single antenna must receive the same weight adjustment.

Cognio’s CBF, however, performs the weight shifts using digital signal processing (DSP) techniques in the Baseband IC. (The technology to do this is part of the IP cores which Cognio provides to our partners as part of CBF.) Digital signal processing allows for much greater flexibility in adjusting signals. Specifically, Cognio’s CBF technology custom-adjusts the phase shift, *on a frequency-by-frequency basis, for each of the 52 sub-carrier frequencies* in an 802.11a signal.

Put another way: During signal reception, each separate frequency receives its own, most advantageous weight adjustment to maximize the gain of the received signal. During signal transmission, the optimized weight adjustments are calculated on a *per-frequency* basis for the transmitted carrier frequencies as well. (And of course, these calculations are carried out separately for each of the transmitting antennas as well.) This results in maximum possible transmission signal gain.

- 802.11b signals are sent using DSSS (direct sequence spread spectrum) or CCK (Complementary Code Keying). In this case, the carrier frequency is spread out over a range of frequencies. Technically, this is different from OFDM, because OFDM uses 52 distinct frequencies, while DSSS/CCK has a continuous “smear” of carrier frequencies.

In terms of signal processing, however, a fairly similar technology is applied by Cognio’s CBF. At the risk of over-simplifying a bit: CBF, in essence, takes the continuous range of transmit frequencies, and treats them as a bundle of very narrow frequencies bands that are directly adjacent to each other. Here again: Calculations are made in the Baseband chip for the correct, and unique, weight adjustments to make to each of these adjacent frequency bands.

Whether 802.11a or 802.11b is in use, the essence of the matter is the same: WLAN transmissions involve not just a single frequency, but multiple frequencies. By applying digital signal processing techniques to the signal in the Baseband IC—before the signal gets anywhere near the RF IC—CBF can carefully adjust the weights for each of the separate frequencies. During signal reception these adjustments compensate for delay spread as effectively as possible; during transmission they result in the maximum possible signal arriving at the receiver *for all of the component frequencies within the signal*.

---

## Vector CBF: Achieving Maximum Path Utilization

So far we've seen a progression of ways to take advantage of multiple antennas:

- The first strategy, selection diversity, simply selects the antenna that is receiving the best signal.
- The second strategy, MRC uses the signals received at both antennas. By adding the power from both signals (with suitable weight adjustments), the receiving device gains more signal with relatively less noise. This increases rate-at-range.
- CBF uses calculations similar to (but more advanced than) those used by MRC, in order to ensure that the transmitting device also adjusts (weights) the signals that are sent by each antenna. Making suitable weight adjustments to the signal transmitted by each antenna (a process we also referred to as transmit MRC) ensures the following: Of the limited energy that is permitted for transmission, as much of that energy as possible actually arrives at the desired receiver, rather than being randomly radiated off in other directions.

The first two strategies are in general use in the WLAN marketplace, but only Cognio supports the more powerful and effective strategy of CBF. As already indicated, one of the great benefits of CBF is that it is effective even if used only on one end of the communications link. This ensures that WLAN's will demonstrate major performance improvements, even when CBF-enhanced WLAN devices are operating along side standard 802.11 devices.


However, Cognio has not stopped there. We are in the process of developing an additional technology, *Vector CBF*, which makes the ultimate, maximum possible use of the opportunities afforded by Multiple-In Multiple-Out antenna technology to enhance data rate-at-range.

---

## Multiple Antennas Allow Parallel Data Transmission

To understand the basic principles behind Vector CBF, a simple visualization will be helpful. Suppose we imagine a pair of network devices, say an AP and a station, each with just one antenna. Suppose, also, that the two antennas were physically connected by a wire (so that they did not broadcast at all, but communicated over the wire just like a regular wired network). It should be intuitively clear that there is some maximum amount of data—a certain number of frames—that can be passed over that single wire in a give period of time.

Now, visualize the same AP and station, only this time each device has four antennas. And again, suppose that instead of transmitting signals over the air, we actually connected the antennas directly by wires, four wires creating a one-to-one link between the antennas. Now, each wire can carry a frame at the same time as the other wires. In other words, *data can be carried in parallel*, so in the same amount of time we can carry four times as much data as before.




**Figure 8: Four Frames Sent in Parallel Over Four Connected Antennas**

While, of course, in a real wireless network, there are no physical connecting wires between the antennas, it turns out that the physical principle remains valid: Multiple antennas (if used on both ends of the link) allow data to be carried in parallel. Because the wireless channel is subject to interference in ways that wired connections are not, using four antennas does not always guarantee four times the data density; but it is feasible to achieve up to four times the data density that can be accomplished with only one antenna at each end of the link. (If one end of the link has just two antennas—and the other end has two or more antennas—the achievable data density will be up to two times the single antenna data density.)

---

### Compensating for Intersymbol Interference

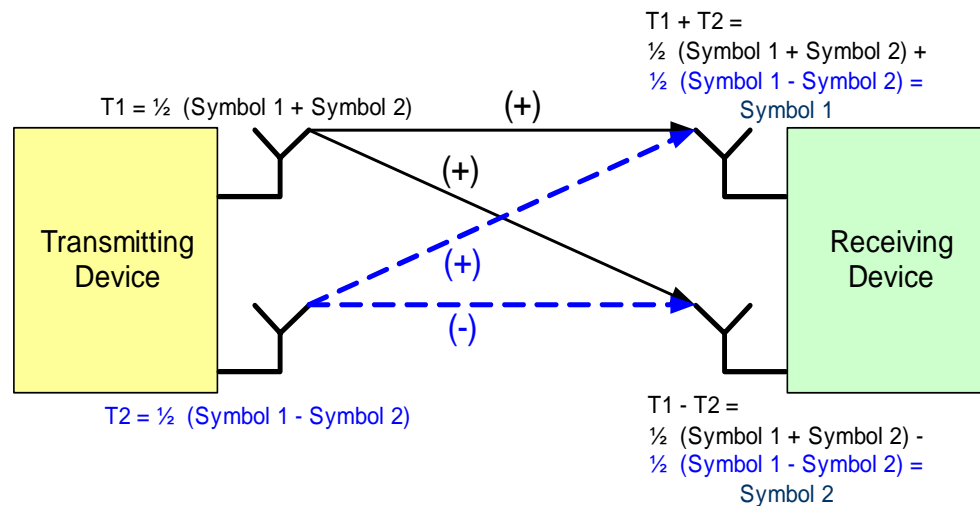
Because the multiple antennas on each end of the link are not connecting by a physical wire, each individual transmitting antenna sends its signal throughout space; the result is that each receiving antenna picks up the signals—and therefore the frames—from all the transmitting antennas. This creates a problem for the receiving device, which must somehow unscramble the four frames that have become scrambled together during transmission through the air. The problem is illustrated below (using just two antennas on each end, for clarity).



**Figure 9: Intersymbol Interference**

In the diagram we use the term “symbol” in place of “frame”, because the data in frames is actually transmitted in the form of symbols. (Symbols are composed of sine and cosine waves, which may bring back distant memories of funny looking wavy lines from high school trigonometry class.) But, whether the term “symbol” or “frame” is used, the problem is the same: Each receiving antenna receives a mixture of symbols.

The solution to this problem actually involves *deliberately* mixing the separate symbols, so that each transmission antenna is carrying a combination of symbols. This is illustrated in conceptual terms here, but we emphasize—and discuss further, below—that the illustration is radically simplified.



**Figure 10: Vector CBF**

*Figure 10* requires a little explanation. The plus (+) and minus (-) signs in parentheses, near the arrows which represent the transmission paths, symbolize the transmission properties of the RF communications channel between the transmitter and the receiver. For our example, we have conveniently *chosen* very simple numbers (in essence, +1 and -1) to represent the channel. In the real world, *the properties of the channel are determined by the physical environment*, and Cognio’s Vector CBF technology has to be smart enough to *figure out* the set of numbers which describe the channel.

At the transmitting antennas, we deliberately combine the symbols we are sending – adding them at the first antenna, and subtracting one from the other at the second.

At the receiving antennas (Warning: oversimplification follows), the transmitted symbols (T1 and T2) have *automatically* combined to yield the desired symbols at each antenna. That is, *because of the physical properties of the transmission channel itself*, T1 + T2 automatically combine at the upper receiving antenna to give Symbol 1. Similarly, T1 – T2 automatically yield Symbol 2 at the lower receiving antenna.

Conceptually, this is what Vector CBF is all about: We deliberately mix the transmitted symbols at each of the transmitting antennas so that, at the receiving antennas, the actual symbols of interest can be distinguished.

## Glossary

Term	Meaning
Access Point (AP)	An 802.11 wireless network will typically have one or more Access Points, which are hardware modules dedicated specifically to providing 802.11 connectivity for the stations on the network.
Composite Beamforming (CBF)	This is Cognio’s variation on Maximal Ratio Combining, or MRC. Using two or more transmission antennas, CBF adjusts the weight of the outgoing signal at each antenna. These weight adjustments result in the maximum possible signal arriving at the desired target. During signal reception, CBF also adjusts the weights of the signals received at each antenna to extract the maximum possible data from the signal. In English, CBF is an inexpensive way to get a lot more range out of a WLAN system, using a handful of cheap antennas.
Intelligent Spectrum Management (ISM)	Cognio’s proprietary, but fully-standards-compatible technology for using advanced analysis of the RF spectrum to achieve the maximum possible network reliability, range, and performance with a relative minimum of human attention.
Interference	<p>“Interference” has two senses, which much be distinguished by its use in context:</p> <ol style="list-style-type: none"> <li>1. In relation to the basic physics of radio waves, <i>Interference</i> is a process where two radio waves—upon reaching the same point in space at the same moment—combine their amplitudes, to give one resulting amplitude. A charged particle at that exact point in space—such as an electron inside a radio antenna—will experience the combined force of the two waves. If (at that singular point in space, at that moment in time) both waves are pushing in the same direction, they will combine constructively; the charged particle will experience their added force pushing in the same direction. This is called <i>constructive interference</i>. If each wave happens to be pushing in opposite directions, this <i>destructive interference</i> results in the charged particle feeling a weaker force; the stronger push in one direction will prevail only slightly over the weaker force in the other direction. (If, at that point in space, the force of both waves just happens to be exactly equal but in opposite directions, a charged particle would feel no net force at all.)</li> <li>2. In relation to overall network operations, <i>Interference</i> is any RF signal which interferes with the signal we actually want to hear. On an 802.11 network, RF interference—also referred to as “noise”—can come from Bluetooth devices, cordless phones and headsets, microwave ovens, radar, and other, neighboring networks. In this context—when two or more <i>different transmitters</i> of RF waves interfere with each other—it doesn’t matter much if the interference (at the physical wave level) is constructive or destructive. All that matters, and what causes the problem, is that one signal which is not wanted arrives at the receiver along with the desired signal.</li> </ol>
Multipath Interference	Multipath interference results when an initial wave, from a <i>single transmitter</i> , travels to a receiver both along a direct line-of-sight, and by one or more reflected paths. If the different versions of the “same” wave arrive out of phase, the destructive interference makes the signal largely cancel itself out.
Multiple-Input-Multiple-Output (MIMO)	The use of two or more antennas for RF transmission and reception in order to improve transmission range and data rate, and to reduce errors in data transmission.

Term	Meaning
Network Interface Cards (NICs)	<p>These are cards in a WLAN stations (such as personal computers) that provide the wireless functionality and connectivity at the hardware level. They may be installed as a card, in a PCI slot in a personal computer, for example, or as a PCMCIA card. More recent computers may have no actual "NIC"; the necessary hardware may be built onto the motherboard, and the WLAN antennas may be built into the case. For convenience, we typically still speak of the NIC as the WLAN hardware in a personal computer. NICs communicate with the network via Access Points, or APs, which serve as the backbone of the network.</p>
Omni-directional Antenna	<p>A simple, standard antenna that radiates nearly equal RF energy in every direction.</p>
Rate-At-Range	<p>The effectiveness of a WLAN signal transmission is determined both by how far the signal can go, and by the data rate. A rate-at-range specification indicates the maximum reliable transmission distance at a various data rates; <i>or</i> it may indicate the maximum reliable data rates at various distances.</p> <p>Saying that a technology—such as Cognio’s CBF—results in improved rate-at-range means that, for a given distance, a higher transmission rate can be achieved; it also means that for a given transmission speed, an AP and a NIC can be further apart while still achieving the same speed. Rate-at-range measurements are, of course, significant only in relation to some kind of baseline measurement of a non-enhanced WLAN technology.</p>
RF Waves	<p>RF waves are a form of electromagnetic waves. Electromagnetic waves are fields that are propagated through space when charged particles (like electrons in transmitting antennas) vibrate. As the fields travel, they have the effect of exerting a force on other charged particles they encounter – such as electrons in receiving antennas. The direction and magnitude of the force varies in a wavelike fashion (varying from up to down or left to right) as the wave travels. When RF waves reach an antenna the induce a voltage, or electric force, and the force creates a current (motion of electrons) inside the antenna.</p> <p>Electromagnetic waves of different wavelengths—the measure of the distance from one peak in the wave to the next— have different affects on matter. For example, electromagnetic waves with wavelengths of about <math>10^{-7}</math> meters form visible light. Electromagnetic waves with wavelengths from <math>10^{-3}</math> meters to <math>3 \times 10^4</math> meters are suitable for carrying radio signals through the air, and are called radio waves, or RF waves.</p> <p>The unlicensed band RF waves, with a frequency of about 2.4 GHz or 5 GHz, have a wavelengths on the order of a 12 cm or 6 cm (5 inches or 2.5 inches). As a result, multipath interference effects can vary over a matter of inches. (See “Multipath Interference”.)</p>