

IEEE Link Task Force Autodetect

Specification for NWay Autodetect

Version 1.0
April 10, 1994

Principal sponsors of this specification are:

AMD
Ascom Timeplex
AT&T - Bell Labs
David Systems Inc.
Hitachi Cable
Luxcom
Microlinear
National Semiconductor
Olympic Technology Group
Seeq
SMC

Copyright National Semiconductor, 1994

Technical inquiries and editorial comments should be directed in writing to:

Walter Hurwitz
National Semiconductor
2900 Semiconductor Drive
P.O.Box 58090
Santa Clara, CA 95052-8090

FAX: (408) 749-8532
email: walter@berlios.nsc.com

IMPORTANT INFORMATION AND DISCLAIMERS

1. THIS SPECIFICATION IS FURNISHED FOR INFORMATIONAL PURPOSES ONLY. AMD, ASCOM TIMEPLEX, AT&T - BELL LABS, DAVID SYSTEMS, HITACHI CABLE, LUXCOM, MICROLINEAR, NATIONAL SEMICONDUCTOR, OLYMPIC TECHNOLOGY GROUP, SEEQ, AND SMC MAKE NO WARRANTIES WITH REGARD THERETO, AND IN PARTICULAR DO NOT WARRANT OR REPRESENT THAT THIS SPECIFICATION OR ANY PRODUCTS MADE IN CONFORMANCE WITH IT WILL WORK IN THE INTENDED MANNER OR BE COMPATIBLE WITH OTHER PRODUCTS IN NETWORK SYSTEMS. NOR DO THEY ASSUME RESPONSIBILITY FOR ANY ERRORS THAT THE SPECIFICATION MAY CONTAIN OR HAVE ANY LIABILITIES OR OBLIGATIONS FOR DAMAGES (INCLUDING, BUT NOT LIMITED TO SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES) ARISING OUT OF OR IN CONNECTION WITH THE USE OF THIS SPECIFICATION IN ANY WAY. AMD, ASCOM TIMEPLEX, AT&T - BELL LABS, DAVID SYSTEMS, HITACHI CABLE, LUXCOM, MICROLINEAR, NATIONAL SEMICONDUCTOR, OLYMPIC TECHNOLOGY GROUP, SEEQ, AND SMC COMMERCIAL PRODUCTS MAY FOLLOW OR DEVIATE FROM THIS SPECIFICATION AT ANY TIME
2. NO REPRESENTATIONS OR WARRANTIES ARE MADE THAT ANY PRODUCT BASED IN WHOLE OR IN PART ON THE ABOVE SPECIFICATION WILL BE FREE FROM DEFECTS OR SAFE FOR USE FOR ITS INTENDED PURPOSE. ANY PERSON MAKING, USING, OR SELLING SUCH PRODUCTS DOES SO AT HIS OWN RISK.
3. THE USER OF THIS SPECIFICATION HEREBY EXPRESSLY ACKNOWLEDGES THAT THE SPECIFICATION IS PROVIDED AS IS, AND THAT AMD, ASCOM TIMEPLEX, AT&T - BELL LABS, DAVID SYSTEMS, HITACHI CABLE, LUXCOM, MICROLINEAR, NATIONAL SEMICONDUCTOR, OLYMPIC TECHNOLOGY GROUP, SEEQ, AND SMC, NEITHER INDIVIDUALLY NOR COLLECTIVELY, MAKE ANY REPRESENTATIONS, EXTEND ANY WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTY OR REPRESENTATION THAT THE SPECIFICATION OR ANY PRODUCT OR TECHNOLOGY UTILIZING ANY ASPECT OF THE SPECIFICATION WILL BE FREE FROM ANY CLAIMS OF INFRINGEMENT OF INTELLECTUAL PROPERTY, INCLUDING PATENTS, COPYRIGHT AND TRADE SECRETS OF ANY THIRD PARTY, OR ASSUMES ANY OTHER RESPONSIBILITIES WHATSOEVER WITH RESPECT TO THE SPECIFICATION OR SUCH PRODUCTS. IN NO EVENT WILL AMD, ASCOM TIMEPLEX, AT&T - BELL LABS, DAVID SYSTEMS, HITACHI CABLE, LUXCOM, MICROLINEAR, NATIONAL SEMICONDUCTOR, OLYMPIC TECHNOLOGY GROUP, SEEQ, AND SMC BE LIABLE FOR ANY LOSSES, DAMAGES (INCLUDING DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES), COSTS, JUDGEMENTS, OR EXPENSES ARISING FROM THE USE OR LICENSING OF THE SPECIFICATION HEREUNDER.

Table of Contents

Table of Contents

List of Figures

List of Tables

Introduction.....	1
Overview	1
Proliferation of CSMA/CD-Compatible LANs	1
The Interoperability Challenge.....	3
Autodetect Application Example	3
The NWay Solution.....	4
Criteria for Auto-Negotiation Standards	5
Remote Capability Detection & Auto-Configuration.....	5
No Risk of Incompatibility with 10BASE-T.....	5
Cost Effective & Easy to Implement	5
Re-Negotiation	6
Auto-Configuration Priority Scheme	6
Manual Override	6
Management.....	6
Remote Fault Sensing.....	6
Extensibility	6
Compatibility w/ Standard Based Link Signaling Protocols	7
Arbitration Resolution.....	7
Symmetry	7
Standard Across CSMA/CD	7
Non-Criteria for Autodetect Solutions.....	8
Flow Control	8
UTP Cabling Identification.....	8
No AUI Operation.....	8
No Fiber support.....	8
Optimal End-to-Discovery	9
Why Link Based Solutions?.....	10
Introduction to NWay	11

National Semiconductor

April 15, 1994

NWay Introduction	11
General NWay Function and Features	12
General NWay Architecture and Operation	12
Detailed NWay Architecture	14
NWay Protocol	14
NWay Arbitration state machine	15
NWay Transmit state machine	18
NWay Receive state machine	20
Technology Specific state machine requirements	24
NWay Peripheral Logic Block / Management Interface	24
Data Word Bit Field Encoding	25
Protocol Discriminator	25
Acknowledge Bit	26
Capability Field	26
Reserved Bits	27
Next Page Bit	28
Priority Resolution	29
Fast Link Pulse (FLP) Description	30
FLP Logical 0 / 1 Encoding	30
FLP Timing	30
10BASE-T Link Pulse Comparison	32
FLP Electrical Characteristics	32
Management Support	33
Manual Override	33
Renegotiation	33
Remote Fault Sensing	34
MIB Statistics	34
Protocol Connection Timeouts	35
Parallel Advertisement	35
Operational Examples	36
NWay Capable <---> 10BASE-T in Link_Fail state	37
NWay Capable <---> 10BASE-T only in Link_Pass state	39
NWay Capable connection with multiple common capabilities	40
NWay Capable connection between 100BASE-T4 nodes	41
Summary	42
Appendix Overview	44
NWay Questions and Answers	45
10BASE-T Review	51
Link Pulse Template	52
DC Balance	53
FCC Compliance	55
Glossary	57

List of Figures

FIGURE 1.	Proliferation of CSMA/CD-based LANs will spawn multi-function nodes.	2
FIGURE 2.	Autodetect is crucial in determining capabilities and configuring links	3
FIGURE 3.	Top-level Block Diagram	15
FIGURE 4.	NWay Arbitration state machine	17
FIGURE 5.	NWay Transmit state machine	19
FIGURE 6.	FLP Detect Timers (FLP_test_min/max_timers)	20
FIGURE 7.	FLP Burst Timer (Link_test_min/max_timers)	21
FIGURE 8.	FLP Data Detect Timers (Data_detect_min/max_timers)	21
FIGURE 9.	NWay Receive state machine	23
FIGURE 10.	Data Word Encoding	25
FIGURE 11.	Data Encoding of Zeros and Ones	30
FIGURE 12.	FLP Pulse Timing	31
FIGURE 13.	FLP Burst Timing	31
FIGURE 14.	FLP Burst versus NLP	32
FIGURE 15.	Example: NWay Capable <---> 10BASE-T	37
FIGURE 16.	Case 1: Reception begins in between FLP bursts	37
FIGURE 17.	Case 2: Reception begins just before the last pulse in an FLP bursts	38
FIGURE 18.	Example: NWay Capable <---> 10BASE-T	39
FIGURE 19.	Example: Two NWay Capable Stations with Multiple Common Protocols	40
FIGURE 20.	Example: Two NWay Capable 100BASE-T4 stations	41
FIGURE 21.	Normal Link Pulse Timing	51
FIGURE 22.	Link_Loss_State Transitions	51
FIGURE 23.	Link Pulse Template	52
FIGURE 24.	FCC Emissions (Simulation Results)	56

List of Tables

Table 1:	FLP Timer Min/Max Values	22
Table 2:	Protocol Discriminator Encoding	25
Table 3:	Acknowledge Bit Encoding	26
Table 4:	Capability Field Encoding	26
Table 5:	Examples of Multi-capable node encoding.....	27
Table 6:	Reserved Bit Encoding	27
Table 7:	Next Page Bit Encoding	28
Table 8:	FLP Min/Max Timing	31

Introduction

Overview

Maintaining interoperability among CSMA/CD-compatible LANs is becoming more complex. A variety of new incompatible speed, wiring, and feature options are on the horizon. Since a single RJ-45 connector is used as the common connection point for this rapidly growing set of network types, the need for a universal mechanism to exchange network capabilities between nodes is becoming imperative. The industry must move quickly to standardize on an autodetection and configuration mechanism in the 802.3 specification or sacrifice interoperability among the next generation of networking solutions.

Designed to manage interoperability in multi-functional LAN environments, National Semiconductor's NWay is a single "Way" for two stations with "N" different common modes of communication to establish a common mode of operation. At power-up, NWay automatically establishes a link that takes advantage of the highest common denominator of their mutual capabilities. Subsequently, NWay's configuration signaling mechanism allows network managers to detect the entire range of remote capabilities.

Proliferation of CSMA/CD-Compatible LANs

In order to provide users with higher throughput, the Fast Ethernet Alliance has coupled the CSMA/CD Ethernet MAC with a 100 Mb/s physical layer. Two different 100 Mb/s technologies have been proposed: 100BASE-TX and 100BASE-T4. When both options are brought to market, interoperability issues will arise. These two Fast Ethernet options utilize different transmission schemes and cabling: 100BASE-TX is based on two pair Category 5 UTP, STP and fiber optic media, while 100BASE-T4 is based on four pairs of Category 3, 4, and 5 UTP media.

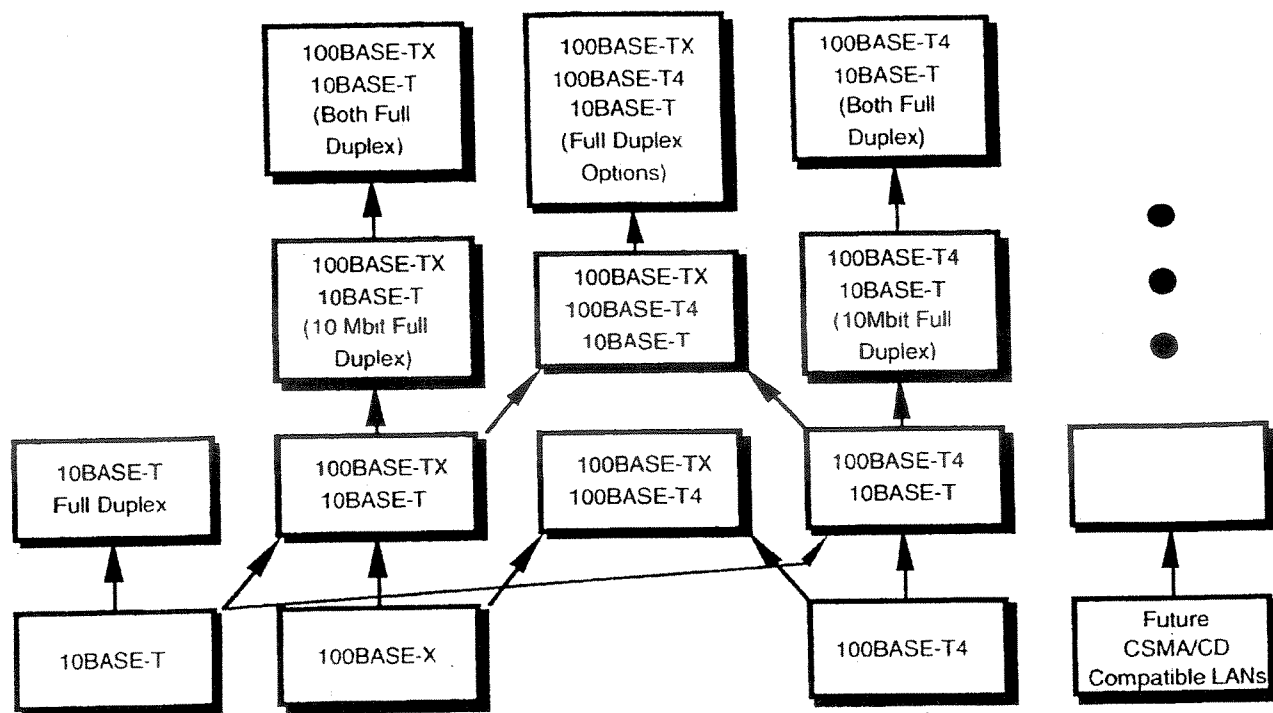


FIGURE 1. Proliferation of CSMA/CD-based LANs will spawn multi-function nodes.

Future nodes will offer multiple configuration options, making the number of permutations quite large. For one, future LAN stations will need to provide interoperability with the installed base of IEEE 802.3 LANs. Vendors supporting both of the new Fast Ethernet options have proposed multi-functional nodes that are capable of supporting both 10 Mbit/s and 100 Mb/s speeds.

These future 10/100BASE-TX and 10/100BASE-T4 multi-function cards will need an autodetect mechanism to allow them to interoperate using their common 10BASE-T mode. Vendors will likely offer single mode 100BASE-TX and 100BASE-T4 Network Interface Cards (NICs), as well as triple function 10/100BASE-TX/T4 boards to cover all bases.

Adding to the complexity of the 802.3 environment, a host of Full Duplex Ethernet products use existing cabling with a slight modification in MAC protocol to achieve a significant performance boost. Full Duplex operation is an option for both 10BASE-T and 100BASE-TX connections.

The Interoperability Challenge

Without a standard autodetection and configuration method, manual configuration will be the rule. Once installed in the enclosure of a computer or a server, it will be difficult to track the capabilities of the multi-function cards manually. With the increasing proliferation of incompatible LAN options, chaos may ensue. Foremost, LAN installers face the peril of connecting a node with unknown capabilities to an incompatible network, which can result in a network crash.

Since both hub and node products will support multiple network types, the challenge for Information Services (IS) managers and end users alike is the need to remotely detect and arbitrate the type of communication that occurs as their network environment grows and changes.

Autodetect Application Example

An ideal universal autodetection and configuration system must provide for the autoconfiguration of a link into its highest common capability. The following figure illustrates the point:

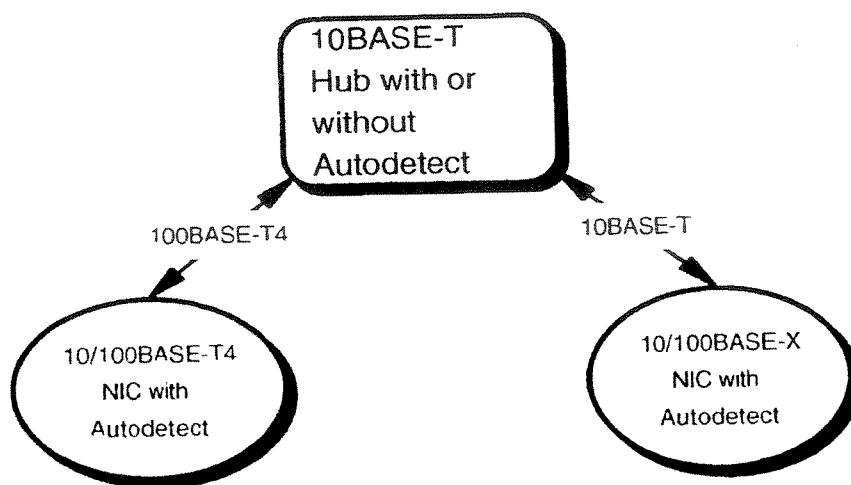


FIGURE 2. Autodetect is crucial in determining capabilities and configuring links

In the network shown in Figure 2, two different types of 10/100 Mb/s NICs are connected to a 10BASE-T hub. The link on the left could be configured in either 10BASE-T mode, or 100BASE-T4 mode. The link on the right could be configured in either 10BASE-T mode, or 100BASE-TX mode. The autodetect mechanism establishes a 10BASE-T connection

from the hub to both left and right nodes. This guarantees correct speed establishment and proper connection for the system.

The NWay Solution

NWay is designed to provide both the automatic and manual configuration capabilities illustrated in the preceding example.

Upon power up, an NWay capable hub can detect and automatically configure its ports to take maximum advantage of common modes of operation that exist, without user intervention or any prior knowledge of the capabilities of either station. In addition, NWay allows the network manager to reconfigure any pair of stations to any other common mode of operation that exists as required.

NWay provides a simple mechanism for point-to-point communication that is efficient in terms of the amount of logic gates required. Regardless of the number of protocols supported, a total of 4 state machines, 6 timers, 3 counters, and some peripheral logic is needed. An analysis of the architecture shows that the design will require approximately 400 cells to implement. NWay has a simple architecture which allows easy implementation by silicon suppliers. It should be noted that this level of complexity remains fixed as new technologies emerge and the need to coexist with the growing installed base of RJ-45 LANs multiplies.

Criteria for Auto-Negotiation Standards

The ideal auto-negotiation solution is a universal method that passes information over a 10BASE-T link, providing reliable negotiation and configuration of both ends of a point-to-point link between 10BASE-T and higher performance modes of operation for RJ-45 LANs without disruption to the network. To accomplish this goal, the following criteria must be met.

Remote Capability Detection & Auto-Configuration

The major benefit of a universal remote auto-negotiation method is the ability to leverage the maximum resources of each node. The auto-negotiation protocol must allow the nodes to discover and acknowledge their mutual capabilities to establish a connection at the highest common denominator of functionality.

Remote Capability Detection ensures that if a mode of interoperation exists, it can be discovered. Auto-Configuration ensures that the highest common denominator mode of interoperation is automatically selected and utilized.

No Risk of Incompatibility with 10BASE-T

As the lowest common denominator, the autodetection method must interoperate with 10BASE-T networks. To ensure interoperability with 10BASE-T, the auto-negotiation method must allow easy connection to an existing 10BASE-T card if both nodes support 10BASE-T. Otherwise, a node supporting auto-negotiation should do nothing that would bring down the 10BASE-T network or pollute MIB statistics.

Cost Effective & Easy to Implement

To minimize development/implementation costs, the complexity of the auto-negotiation mechanism must be low. The auto-negotiation mechanism must be silicon efficient by

maintaining a low gate count to ensure low additional cost to products. Also, the gate count should not change dramatically with the addition of future supported technologies.

Re-Negotiation

At a minimum, the autodetection mechanism must negotiate on power up or upon link reset to allow connection without disrupting the attached network. In addition, with the wide variety of possible products, a management function may have reasons to require a mode change for a given link such as an unreliable link. Thus, it is desirable that the autodetect protocol support renegotiation capability.

Auto-Configuration Priority Scheme

The priority scheme establishes a hierarchy that determines the connection type that is automatically established when multiple possible connection types exist. This allows the link with the most desirable capabilities to be chosen. The configuration signaling scheme must also provide a method by which a node may establish a link at any lower priority level connection type that might be desired by a higher level management entity.

Manual Override

To accommodate interoperation with future products that choose not to implement the auto-negotiation scheme, a manual or network management switch must be able to override the auto-negotiation.

Management

Other criteria have specified the need for management capabilities, but existence of a management agent is not required for the basic function of common mode discovery and configuration.

Remote Fault Sensing

The auto-negotiation scheme should be able to detect that the far end cable connection is not properly connected.

Extensibility

With an ever growing diversity of LAN products that support 10BASE-T as the lowest common denominator, the ideal autodetect and configuration method must be

extensible, providing expansion for future CSMA/CD-compatible devices. Given the fast pace at which technology advances, future technologies should be able to leverage this communication scheme without being forced to design new schemes, which only lead to further interoperability problems. Therefore, the auto-negotiation architecture must not change as new LAN technology options are added.

Compatibility w/ Standard Based Link Signaling Protocols

The autodetect method must also permit standard-dependent link signaling protocols. Some technologies, such as 100BASE-T4 require their own link integrity schemes for technology-dependent reasons, thus need to co-exist with the auto-negotiation method.

Arbitration Resolution

The auto-negotiation mechanism should complete mode negotiation to achieve closure within a time which ensures that popular network protocol connections are not dropped. The worst case 10BASE-T Ethernet transmit time is 300ms. Therefore, in theory, the negotiation should complete in less than 300ms to ensure this goal. In practice, the upper bound may be slightly higher.

Symmetry

The auto-negotiation mechanism should operate symmetrically so that a single implementation can be used for any auto-negotiation capable device. This means no master/slave relationship will exist between the node and hub products supporting the auto-negotiation mechanism.

Standard Across CSMA/CD

The development of a standard auto-negotiation mechanism meeting the above criteria will ensure interoperability among stations with common communication capabilities.

While implementing such a new standard should be optional, the function should be standardized by the 802.3 for all RJ-45 CSMA/CD compliant devices that implement it to prevent interoperability problems from spreading. However, vendors of multi-function products that include a 10BASE-T mode of operation are strongly encouraged to implement the standard auto-negotiation mechanism.

Non-Criteria for Autodetect Solutions

While solutions to the auto-negotiation problem should be as extensible and flexible as possible, reasonable scope and bounds to the problem need to be set. As a result, the following is a summary of issues which will not necessarily be addressed by the auto-negotiation scheme.

Flow Control

Flow control implies a more real time mechanism than is required by the set of criteria. As a result, no attempt should be made to complete signalling during a 10BASE-T IPG time. Flow control is a larger issue that needs to be addressed in the future, perhaps from a management perspective.

UTP Cabling Identification

The auto-negotiation mechanism need not attempt to discover which grade of UTP wire is connected as part of the technology configuration decision logic.

No AUI Operation

The auto-negotiation mechanism is not required to support the AUI interface for MAU applications.

No Fiber support

The auto-negotiation mechanism, need not support the basic configuration operation over fiber optic cabling.

Optimal End-to-Discovery

The auto-negotiation mechanism need only configure the highest common denominator on a link by link basis. Optimizing the end-to-end highest common denominator between stations separated by multiple links need not be supported.

Why Link Based Solutions?

To solve the auto-negotiation problem, a communication mechanism is required to transfer link state information such that the agents on each side of the link are configured into a consistent state. Since new technology products are also offering 10BASE-T, it makes sense to adapt a familiar portion of 10BASE-T to solve the auto-negotiation problem.

One communication mechanism inherent to 10BASE-T is in-band data transfer. This method had been looked into by both the Fast Ethernet Alliance and the Full Duplex Consortium before these groups decided to focus on link pulse based, out-of-band solutions. In-band solutions imply that a common communication mechanism (i.e. 10BASE-T implemented at least through the network layer of the OSI MODEL) exists which is not necessarily so. Also, there is a problem detecting hubs which support only 10BASE-T between auto-negotiation capable nodes. That is, in-band solutions can communicate beyond the other end of a point-to-point link, potentially causing interoperability problems. This makes in-band data transfer solutions problematic and undesirable to implement.

However, another communication mechanism inherent to 10BASE-T exists which indicates the state of the capabilities of a remote-node attached to the RJ-45: Link Pulses. Currently, there are two states the cable can indicate: Good Link or Link Fail. Good Link indicates that a remote 10BASE-T device is connected and has an active transceiver. Link Fail indicates that a remote 10BASE-T device is not available for communication. To solve the auto-negotiation problem, more states relative to the link are required. Now, the auto-negotiation capable device needs to encode further state information as to what type of transceiver is available and other configuration information. Consequently, it makes sense to use link pulses, the current method of conveying link state information, to encode the new states.

Using link pulses accommodates the two most important criteria for the autodetect scheme: interoperability with 10BASE-T and automatic configuration of a link without disruption of the network. A link pulse based autodetect scheme also allows an efficient implementation, based around the 10BASE-T physical layer, which does not require higher level OSI layer intervention.

Introduction to NWay

NWay Introduction

NWay is an interoperability solution that addresses the need to automatically configure network nodes in heterogeneous, 10/100 Mb/s multi-function 802.3 LAN environments. With the introduction of two different 100 Mb/s PHYs that utilize the same RJ-45 connector, CSMA/CD-based 802.3 LAN environments are becoming more complex. NWay provides the ability to manage this complexity by automatically setting up connections between any two nodes that share common network capabilities.

NWay utilizes a simple, robust communications mechanism that allows a pair of nodes (node-to-node, node-to-hub, hub-to-hub) to learn and store each other's capabilities. This includes the exchange of LAN capability information including the speeds, PHY types, and special features supported in a node.

NWay is designed to provide the following basic autodetection and configuration capabilities:

- Speed sensing between 10 and 100 Mb/s CSMA/CD devices
- Identification of both the 100BASE-T4 and 100BASE-TX protocols
- Identification of 10 & 100 Mb/s Half and Full duplex modes
- Automatic highest common denominator capability selection

In addition to its basic capabilities, NWay is designed as a robust architecture that provides:

- 100% compatibility with 10BASE-T
- User flexibility
- Compatibility with future 802.3 LAN standards
- A rich enough code space for future enhancements

General NWay Function and Features

NWay builds upon the existing 10BASE-T link pulse scheme to provide powerful and automatic link capability determination sensing and selection.

NWay is a simple physical layer link determination protocol. With NWay two nodes simultaneously communicate with each other to determine their respective capabilities. This ability to convey all of both node's capabilities to one another in one burst is referred to as "parallel advertisement". This allows both nodes on a link to quickly sense each others capabilities and select the most appropriate technology.

A simple yet robust handshaking state machine assures that capability determination, capability acknowledgment, and highest common denominator capability selection all occur quickly and accurately. Thus, NWay's autodetect and configuration protocol has the ability to configure point-to-point links in as little as 80 ms.

NWay's simple yet robust communication mechanism is designed to ensure the integrity of the configuration information by employing a redundant transmission scheme with handshaking to overcome any possible data corruption or noise issues.

General NWay Architecture and Operation

NWay's method of communication builds upon the link pulse mechanism employed by 10BASE-T nodes to indicate the status of the link. Traditional 10BASE-T network nodes exchange information regarding the status of the link using link pulses — short 100 ns logic high signals separated by large 16ms intervals of logic low.

NWay builds upon the 10BASE-T link pulse paradigm with an efficient communication mechanism based on quick bursts of multiple link pulses at normal link pulse (NLP) intervals.

Unlike 10BASE-T's NLPs, one of which is transmitted every 16 ms time slot, a burst of 33 Fast Link Pulses (FLP) are transmitted within the first 2ms of the same 16 ms duration. A FLP burst consists of 17 clock pulses interleaved with 16 binary data pulses.

The specific data encoded in the FLP pattern sent indicates the presence or absence of particular features in a node as well as control information for the NWay protocol. Particular bit positions indicate a node's support for 10BASE-T, 100BASE-T4, 100BASE-TX, in normal and full duplex modes. In this manner, the NWay protocol is capable of communicating all of the features that characterize a particular LAN node in a single burst transmission.

The NWay protocol controls the exchange of capability data and ensures the integrity of capability communication. During the NWay autodetect process each node begins by sending FLP bursts encoded with the node's capabilities. When a node receives the same data 3 times in succession it sets the Acknowledge bit in the data word it is sending. After a node receives 3 consecutive FLP bursts that have the Acknowledge bit set, the capabilities are fully validated.

NWay's automatic configuration protocol determines the best type of connection and establishes that link. NWay uses a standard priority table to arbitrate the type of link that is established when two or more mutual capabilities exist. The priority table establishes a link of the highest common denominator to take maximum advantage of the pair's capabilities. For example, if both nodes support 10BASE-T and Full Duplex 10BASE-T NWay will automatically select Full Duplex 10BASE-T because it provides greater data throughput.

NWay is a simple yet robust autodetect and configuration signaling technology that addresses all current and future CSMA/CD-compatible LANs. It addresses the interoperability of new 100 Mb/s devices and provides for fallback to 10BASE-T (LCD).

With NWay, even if a mode of interoperation between two nodes does not even exist, the nodes can still remotely detect each other's capabilities. Although an operational link cannot be established, the information can be stored and accessed by a network management agent.

The ability to look around the network and remotely monitor node capabilities is valuable in terms of topology mapping and fault detection. This feature will no doubt gain in importance as CSMA/CD networks take on more disparate options.

NWay's configuration capability also enables network management agents to sense node capabilities and reconfigure links as requirements change.

Detailed NWay Architecture

This section describes the NWay Protocol, Data Word Bit Field Encoding, Priority Resolution, and Fast Link Pulses in detail.

NWay Protocol

The NWay protocol is implemented with four related state machines:

- The **NWay Transmit** state machine controls the transmission of node capabilities and acknowledgment messages using Fast Link Pulses
- The **NWay Receive** state machine controls the reception of the individual clock and data pulses within a Fast Link Pulse burst.
- The **NWay Arbitration** state machine controls the overall capability exchange process using and guiding the other two state machines.
- The **10BASE-T Link Pulse** state machine controls the reception and transmission of Normal Link Pulses that are used when communicating with a 10BASE-T station.

It is also probable that individual Technology-Dependent link-integrity-test state machines will co-exist with NWay. For example:

- The **100BASE-TX Link Test-Fail** state machine controls the reset logic for the Convergence Sublayer of the 100BASE-TX standard. When in link_fail, Idle Line State is continuously sent.

The following diagram graphically shows the relationships between the four state machines and the signals that are passed between them

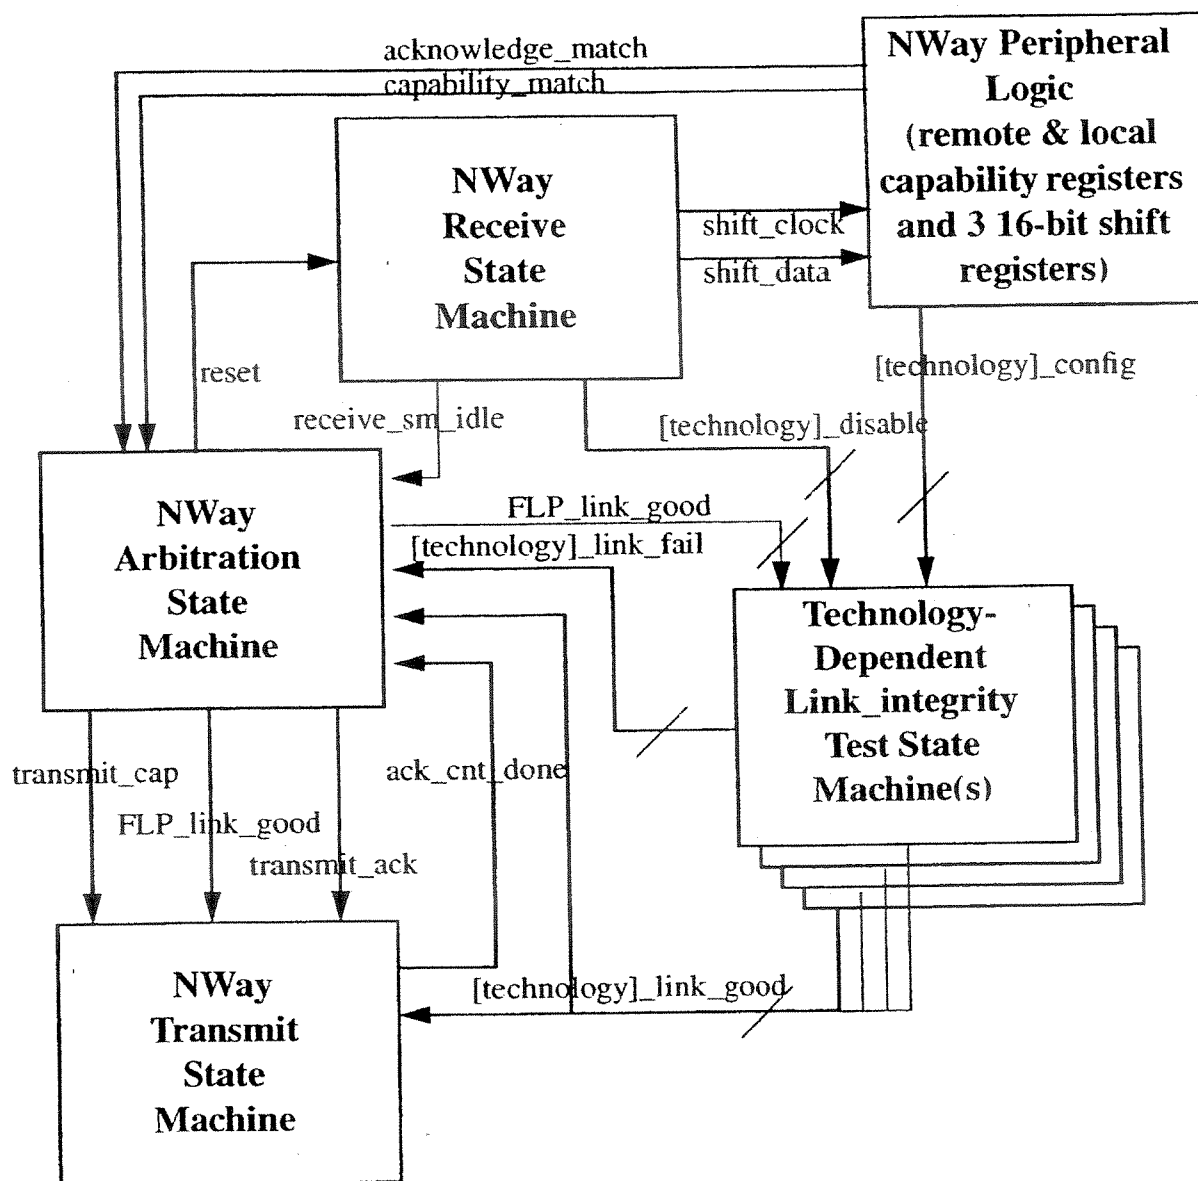


FIGURE 3. Top-level Block Diagram

NWay Arbitration state machine

The purpose of this state machine is to control the FLP transmit state machine and determine if the far-end station has identified itself either through FLPs or through a technology-dependent signalling scheme. At power on, or upon the assertion of

technology-dependent link-fail condition, this state machine moves into Capability Detect state to allow:

- the FLP transmit state machine to start sending FLP bursts containing the local station's capability to the far-end station.
- the results of FLP receive state machine to be used to determine capability match.

A capability match occurs when this state machine is in Capability Detect state, and three consecutive FLP bursts contain the same pattern ignoring the acknowledge bit. Once a capability match occurs, the FLP arbitration state machine moves into Acknowledge Detect state.

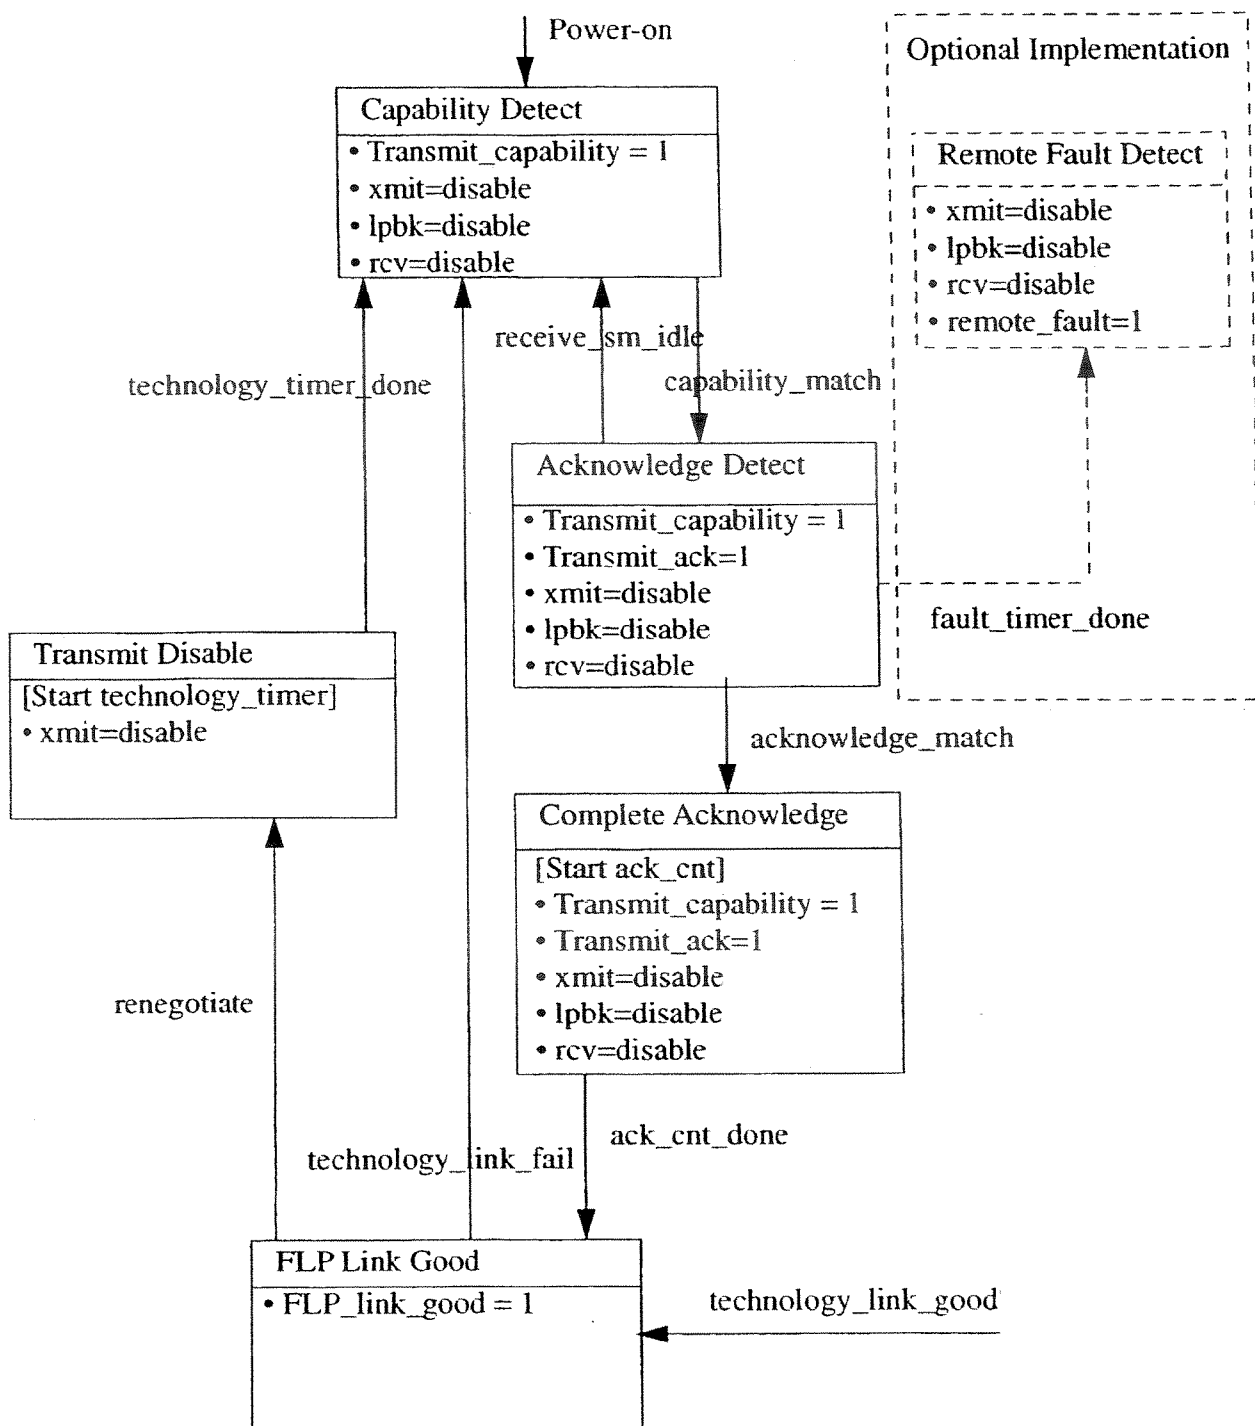
In Acknowledge Detect state, it transmits the same data pattern with the acknowledge bit set, to indicate to the far-end station that it has captured its capabilities.

In the case of unexplained loss of FLP bursts coming from the far-end station, under which the FLP receive state machine goes to IDLE state, this state machine goes back to Capability Detect state to re-start negotiation from scratch. While remaining in the Acknowledge Detect state, it waits to receive `acknowledge_match` from the far-end station. `Acknowledge_match` occurs when two out of three consecutive FLP bursts received have the acknowledge bit set. In order to guarantee that the far-end station has completed `acknowledge_match`, a minimum of 4-6 `acknowledge` patterns are transmitted. Subsequently, it transitions to `FLP_link_good` state. After this point, the technology-specific link signalling scheme takes over, if needed for the specific technology (e.g. 100BASE-TX, 100BASE-T4). If this technology-specific link signalling scheme does not successfully complete in a technology determined period of time, then `technology_link_fail` is asserted, and the NWay Arbitration state machine starts re-negotiation.

If the far-end station does not detect and transmit FLP bursts, and instead sends a technology-dependent signal, (e.g., 10BASE-T link pulses, or 100BASE-TX Idle Line State transmissions), the technology-dependent link-integrity-test state-machine will assert `technology-dependent_link_good` signal thereby halting FLP burst transmits.

In addition, upon successful establishment of technology-dependent link, if the technology-dependent link integrity test state machine goes to link fail and the arbitration state machine goes back to the Capability Detect state.

A higher-level management agent can start a re-negotiation of capabilities with the far-end station by asserting the `renegotiate` signal. The state machine moves into Transmit Disable state to wait for the far-end station to go into a link-fail state and then starts FLP burst transmissions.



Note:

- Capability_match is 3 consecutive data pattern match, ignoring acknowledgment bit.
- Acknowledge_match is 3 consecutive data patterns having acknowledgment bit set.
- ack_cnt_done is from FLP Transmit SM to indicate completion of sending a minimum of 4-16 acknowledge patterns.
- optional remote sensing capability, refer to Management support section

FIGURE 4. NWay Arbitration state machine

NWay Transmit state machine

The purpose of this state machine is to transmit a sequence of 33 fast link-pulses to the far-end station, when instructed by the arbitration state machine. The Idle state is entered at power-on and this state machine transitions to this state if the arbitration state machine is in FLP_link_good state.

The transmit_link_pulse_timer (14 +/- 8 msec) is used to separate two consecutive fast link-pulse bursts transmitted.

If the arbitration state machine is in the process of sending the capability to the far-end station, then the transmit state machine enters into the Transmit_Capability state. If acknowledgment is to be sent, then the state machine enters Transmit_Acknowledge state to initialize ack_cnt (4-16), and then it moves to Transmit_Capability state. In this state, the bit_cnt is started to keep track of the 16 bits of data to be transmitted.

The state machine then alternates between the Transmit_Clock_Bit state and Transmit_Data_Bit state to transmit clock bit and to transmit data bit (if the data bit to be transmitted is logic one) respectively. The interval_timer (62.5 +/- 15 μ s) is used to separate a clock bit and a data bit.

Once all 16 bits of data are transmitted, if the state machine is in the midst of sending acknowledgment patterns then the state machine moves to the Transmit_Count_Ack state. However, after 16 bits of data transmission, if the state machine is not in the midst of sending acknowledgment patterns, then the state machine moves to Idle directly.

The state machine returns to Idle after a 16-bit capability has been transmitted. In the Transmit_Count_Ack state, if 4-6 acknowledgment patterns have been sent, it goes back to idle, otherwise it waits for 14 +/- 8 msec and moves to Transmit_Capability state to send the next burst of 16-bit acknowledgment pattern.

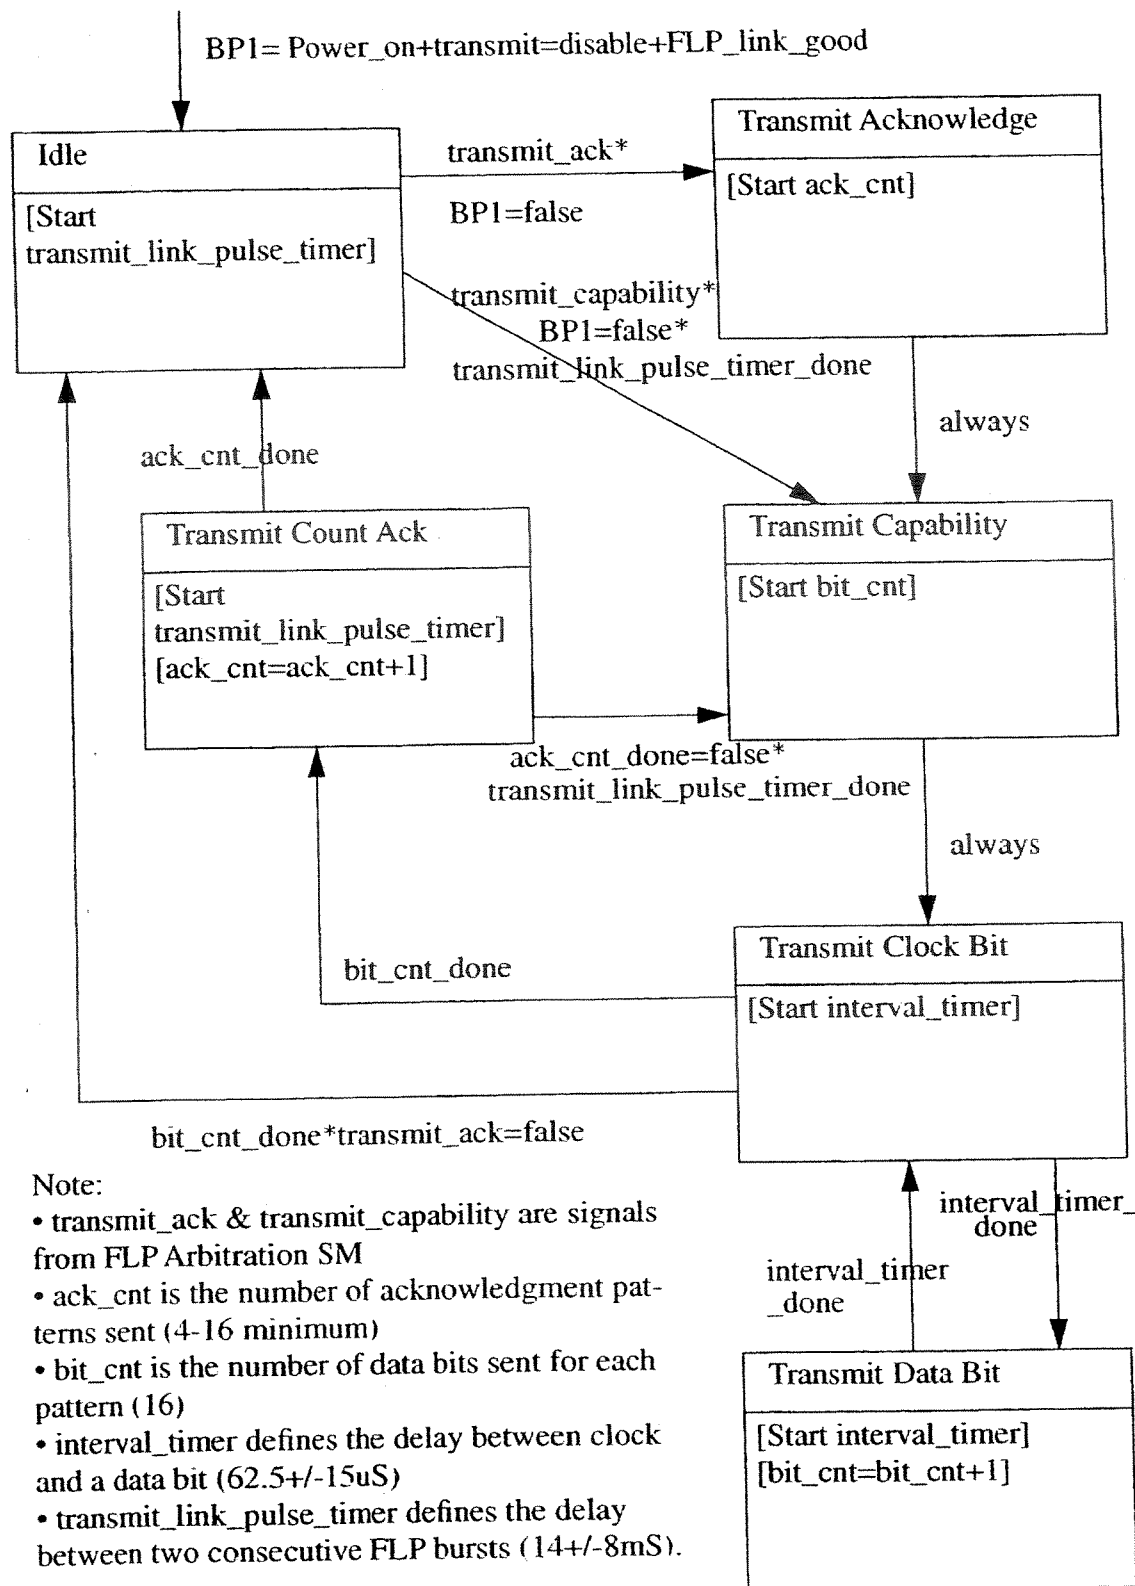


FIGURE 5. NWay Transmit state machine

NWay Receive state machine

This state machine is used to identify that the link pulses received are fast-link-pulse bursts, and if so, to store the embedded 16-bit data in a shift-register. A burst of thirty-three link pulses separated by 62.5 μ s conveys the capability of the far-end station. A link pulse is considered part of a FLP burst if it occurs within a certain time of the previous link-pulse. It is detected by the use of the two timers: FLP_test_min_timer and FLP_test_max_timer. The purpose of the FLP_test_min_timer is to mask out any noise or ringing effects on the line. The FLP_test_max_timer is used to determine whether the next link-pulse received is within a window to be of fast link-pulse category.




FIGURE 6. FLP Detect Timers (FLP_test_min/max_timers)

The Link Receive state machine goes to the idle state at power-on or when the arbitration state machine forces it to this state. When the Link Receive state machine is in idle state, a reception of a link-pulse causes a transition to the Link_Pulse_Detect state.

At this point, both the FLP_test_min_timer and FLP_test_max_timer are started. The FLP_test_min_timer has a value of 5 to 25 μ s. The FLP_test_max_timer has a value of 165 to 185 μ s. If another link-pulse is received in the window of time such that the FLP_test_min_timer has timed out, but the FLP_test_max_timer has not completed, the state machine makes a transition to Link_Pulse_Count state.

In the Link_Pulse_Count state, FLP_count is incremented. If consecutive number of fast link pulses are received, i.e., the FLP_count equals FLP_count_max (6 to 8), then the far-end station is considered NWay FLP capable, and all Technology-Specific link integrity test state machines (e.g. 10BASE-T link_integrity_test state machine) are forced into Freeze state.

The NWay Receive state machine then goes into FLP_Pass state to ignore the rest of the incoming fast link-pulse burst. This is accomplished by waiting for the idle interval after the last link pulse of the current burst, i.e., an idle period of 165 to 185 μ s has been observed. At this time, the state machine transitions into FLP_Capture state, and is ready to capture a complete burst of fast link pulses from the far-end station.

In the FLP_capture state, the link_test_max_timer (25-150 msec.) is started to insure that link pulse bursts are received within this window. When this timer expires, the state machine moves back to Idle state to indicate the absence of FLP bursts.




FIGURE 7. FLP Burst Timer (Link_test_min/max_timers)

The first link-pulse received in the FLP_capture state is interpreted as a clock bit. Two consecutive clock bits can envelop a data bit. The presence of a data bit is interpreted as a logic one for the corresponding data bit field of the shift register. The absence of the data bit is interpreted as a logic zero for the corresponding data bit field of the shift register. Two consecutive clock pulses are separated by $125 \pm 15 \mu\text{s}$. The separation of a subsequent data bit from a clock bit is $62.5 \pm 15 \mu\text{s}$. These margins are designed to be symmetric for both local and far-end stations.

In the FLP_Clock state, the data_detect_min_timer (15 to $47 \mu\text{s}$) and data_detect_max_timer (78 to $110 \mu\text{s}$) are started. If a link-pulse is received when the data_detect_min_timer has completed and data_detect_max_timer is not, it is interpreted as a data-bit of logic one. On the other hand, if the next clock_bit arrives, as indicated by data_detect_max_timer_done, then the embedded data bit is assumed to be zero. In the FLP_Clock state, the link_test_min_timer (5 to 7 msec) is also started; this timer is used to separate two consecutive fast link-pulse bursts.




FIGURE 8. FLP Data Detect Timers (Data_detect_min/max_timers)

The FLP_Data_0 state, a logic zero is shifted into the shift register. In FLP_Data_1 state, a logic one is shifted into the shift register and upon the reception of the next clock-bit, the state machine moves back to FLP_Clock state. Either upon the reception of the last clock bit, or due to a bit error corrupting the last clock bit, the state machine moves to FLP_Capture state.

The following table summarizes the timer minimum and maximum values:

Table 1: FLP Timer Min/Max Values

Parameter	Min	Typ	Max	units
FLP_test_min_timer	5		25	μs
FLP_test_max_timer	165		185	μs
Link_test_min_timer	5		7	ms
Link_test_max_timer	50		150	ms
Data_detect_min_timer	15		47	μs
Data_detect_max_timer	78		110	μs

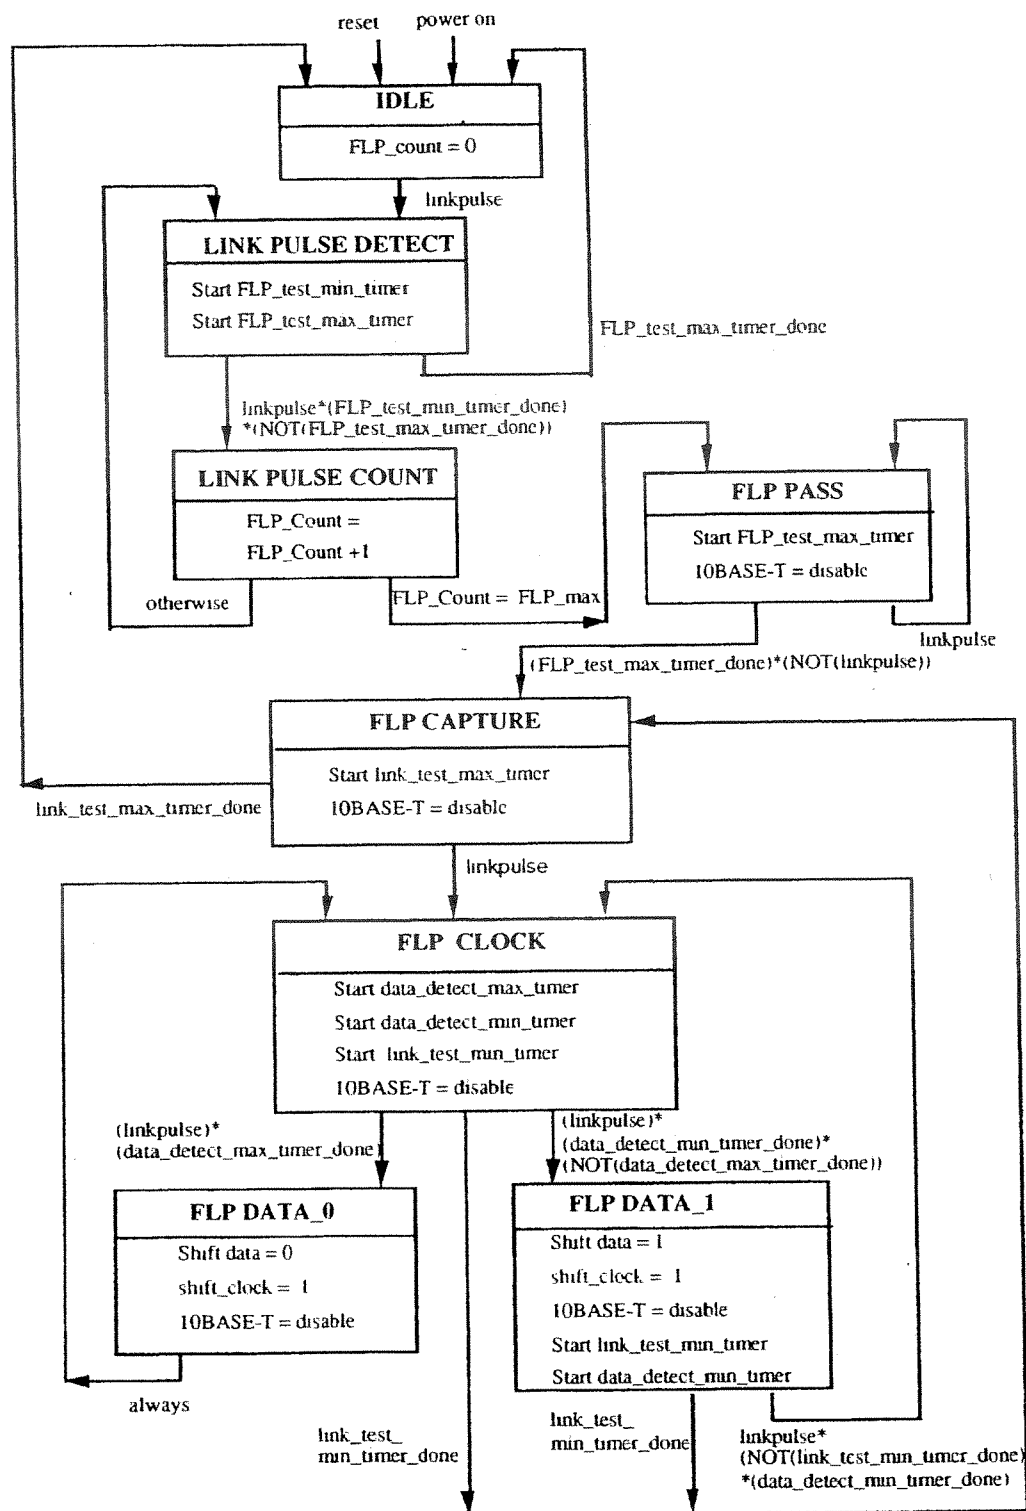


FIGURE 9. NWay Receive state machine

Technology Specific state machine requirements

The only technology-specific state machine that is required by NWay is the 10BASE-T link integrity test state machine. This state machine is required for compatibility with existing 10BASE-T nodes.

Two modifications must be made to the original 10BASE-T link integrity test state machine. The first modification is the requirement that the state machine powerup in the Link Test Fail Reset state.

The second modification needed to original 10BASE-T link integrity test state machine in an NWay capable device is to add another state called Freeze 10BASE-T. This state is entered upon the recognition of the far-end station's NWay FLP capability and remains in this state during NWay's negotiation phase. After the NWay negotiation is successfully completed, and the highest common denominator is determined to be 10BASE-T, the 10BASE-T Link Integrity test state machine moves into the Link_test_pass state. The state machine remains in the Freeze 10BASE-T state if the mode of communication selected is not 10BASE-T.

Additional technology-specific state machines can be added to implement technology-specific link testing.

NWay Peripheral Logic Block / Management Interface

This block contains the following:

- 1) three 16-bit shift registers to hold the received 16-bit capability information from the far-end station,
- 2) a 2-bit counter to select the current shift-register of the above-mentioned three, in which the data-bits are being shifted,
- 3) a 16-bit local capability register,
- 4) a 16-bit far-end capability register,
- 5) combinational gates to generate:
 - capability_match signal for arbitration state machine,
 - acknowledge_match signal for arbitration state machine,
 - highest-common-denominator.

Data Word Bit Field Encoding

The NWay protocol accomplishes its objective by transmitting data with Fast Link Pulses. Each FL P burst contains 16 bits of data.

The 16 bit data word contains information about the protocol and device capability. The data word also contains control information for indicating acknowledgment and next page transmission.

The 16 bit data word is encoded as follows. Bit 0 is the first bit transmitted in the serial data stream.




FIGURE 10. Data Word Encoding

Both ends of a physical network segment transmit their capability data words to each other at the same time. In this way both nodes are able to quickly determine the capabilities of the other node, determine what capabilities they have in common, and select the highest common denominator technology.

Protocol Discriminator

This field is a binary encoding of the nodes protocol group. This allows 32 protocol groups to be defined. Currently only a single protocol discriminator is specified but due to this capability NWay is expandable across all existing and future RJ-45 LAN protocols.

The following table summarizes the Protocol Discriminator encoding:

Table 2: Protocol Discriminator Encoding

D4	D3	D2	D1	D0	Description
0	0	0	0	1	CSMA/CD Compatible
Others					Reserved

Acknowledge Bit

The Acknowledge Bit is used by the NWay protocol to indicate that a station has successfully received the far_end capability data word. Setting the Acknowledge bit indicates that the capabilities of the other station have been received consistently during three consecutive FLP bursts.

The search for acknowledge bits is completed when 3 consecutive bursts are received with the acknowledge bit set.

The Acknowledge Bit is fixed in position for all protocol discriminators.

The following table summarizes the Acknowledge Bit encoding:

Table 3: Acknowledge Bit Encoding

D 5	Description
0	No acknowledgment
1	Transmitting node acknowledges correct receipt of data word

Capability Field

The Capability Field describes the networking types that are supported by the node. Each bit in the field independently indicates whether a certain networking type is supported. This allows each node to support multiple network types.

The following table summarizes the Capability Field Encoding:

Table 4: Capability Field Encoding

D 1 0	D 9	D 8	D 7	D 6	Description
0	0	0	0	1	10BASE-T
0	0	0	1	0	10BASE-T Full Duplex
0	0	1	0	0	100BASE-TX
0	1	0	0	0	100BASE-TX Full Duplex
1	0	0	0	0	100BASE-T4

To demonstrate how the bits can be used together the following table provides some usage examples:

Table 5: Examples of Multi-capable node encoding

D 1 0	D 9	D 8	D 7	D 6	Description
0	0	0	1	1	Regular and Full Duplex 10BASE-T
1	0	1	0	1	10BASE-T, 100BASE-TX, and 100BASE-T4 all supported in normal configurations.
0	0	1	0	1	10BASE-T and 100BASE-TX configurations both supported.
1	1	1	0	0	100BASE-TX, 100BASE-T4, 100BASE-TX Full Duplex supported.

The actual bit ordering of the Capability Field is not important since there is a separate priority table that is used to resolve which capability is actually selected when multiple common capabilities exist.

However, the 10BASE-T bit D6 is fixed in position across all protocol discriminator types. This will allow basic 10BASE-T capability to be universally identifiable, and if available, will provide a means for basic communication across protocols.

Reserved Bits

The Reserved Bits have been set aside for future use. When needed these bits can be defined and standardized.

Reserved Bits should be set to all zeros (0) on transmission and should be ignored and masked out on reception of the data word.

By following these two criteria, products based on the standard before a reserved bit becomes defined will be forward compatible.

The following table summarizes the Reserved Bit encoding:

Table 6: Reserved Bit Encoding

D 1 4	D 1 3	D 1 2	D 1 1	Description
0	0	0	0	Reserved for future use Mask out values on reception

Next Page Bit

The Next Page Bit provides a mechanism to transmit additional protocol and capability specific information after NWay has successfully completed the capability determination of both stations.

If additional protocol-specific information needs to be communicated to the far-end station, the Next Page bit can be set. Under this circumstance, the remaining 15 data bits may contain protocol-specific information.

The mechanism for exchange of this additional information is open for protocol-specific determination. Therefore, 100BASE-TX will decide if and how the next-page bits will be utilized for 100BASE-TX operation. Also 100BASE-T4 will decide if and how the next-page bits will be utilized for 100BASE-T4 operation. Whichever protocol is selected by the automatic protocol priority table will be responsible for Next Page interpretation.

The following table summarizes the Acknowledge Bit encoding:

Table 7: Next Page Bit Encoding

D 1 5	Description
0	Transmitting the primary capability data page
1	Transmitting the protocol specific data page

Priority Resolution

Since two nodes can have multiple capabilities in common a prioritization scheme must exist to ensure that the highest common denominator capability is chosen. The following priority is proposed (from highest to lowest priority):

1. 100BASE-TX Full Duplex
2. 100BASE-T4
3. 100BASE-TX
4. 10BASE-T Full Duplex
5. 10BASE-T

The rational for this hierarchy is straightforward. 10BASE-T is the lowest common denominator protocol and therefore has the lowest priority. Full Duplex solutions are always higher in priority than their Half Duplex counterparts. 100BASE-T4 is ahead of 100BASE-TX because 100BASE-T4 runs across a broader spectrum of copper cabling.

As new CSMA/CD-compatible LAN technologies enter the market a reserved bit will be assigned to each technology by the standards body. The new technology will be inserted into an updated priority table and made a part of the 802.3 auto-negotiation standard. The relative hierarchy of the existing technologies will not change, thus providing backward compatibility with existing NWay implementations. It is important to note that the reserved bits are forced to zeros. This guarantees that devices implemented using the current priority table will be forward compatible with future devices using an updated priority table.

Fast Link Pulse (FLP) Description

NWay's method of communication builds upon the link pulse mechanism employed by 10BASE-T nodes to indicate the status of the link. Traditional 10BASE-T network nodes exchange information regarding the status of the link using link pulses — short 100 ns logic high signals separated by large 16ms intervals of logic low.

NWay builds upon the 10BASE-T link pulse (NLP) paradigm with an efficient communication mechanism based on quick bursts of multiple link pulses at normal link pulse intervals.

FLP Logical 0 / 1 Encoding

Fast Link Pulse bursts consist of a maximum of 33 pulses which are a combination of 17 clock pulses with an optional data bit interleaved between each clock pulse. This allows 16 data bits to be transmitted in each FLP burst with a 1 being encoded as a data pulse between 2 clock pulses and a 0 being encoded as the absence of a data bit between 2 clock pulses.

The figure below illustrates the encoding of data using pulses in a FLP burst.




FIGURE 11. Data Encoding of Zeros and Ones

FLP Timing

Individual pulses in an FLP burst are the same 100ns width as normal link pulses. Both clock and data pulses are the same width.

Clock pulses are always evenly spaced apart by 125 μ s. If a data bit is present between two clock pulses it is centered between them.

The following figure illustrates the FLP Pulse Timing.




FIGURE 12. FLP Pulse Timing

The following figure illustrates the FLP Burst Timing.




FIGURE 13. FLP Burst Timing

The following table summarizes FLP Pulse and Burst timing.

Table 8: FLP Min/Max Timing

#	Parameter	Min	Typ	Max	units
T1	Clock / Data Pulse Width		100		ns
T2	Clock Pulse to Clock Pulse		125		μs
T3	Clock Pulse to Data Pulse (Data = 1)		62.5		μs
T4	Pulses in a Burst	17		33	#
T5	Burst Width		2		ms
T6	Burst to Burst	8	16	24	ms

10BASE-T Link Pulse Comparison

Since Fast Link Pulses are based on 10BASE-T normal link pulses it is easy to compare them. A single FLP burst occurs at the same interval as NLPs occur. That is the time between the start of two FLP bursts is 16ms the same as the interval between two NLPs. This is illustrated below.




FIGURE 14. FLP Burst versus NLP

The width of the individual pulses in an FLP burst is the same as the width of a NLP. Both are 100ns.

FLP Electrical Characteristics

The electrical characteristics of FLPs are very similar to the characteristics of NLPs. FLPs are transmitted single ended on the positive wire of the transmit pair. The actual pulses in an FLP burst follow the exact same Link Pulse Template as NLPs.

The DC balance of the line is not affected by FLP bursts. This is because the bursts are very short relative to the time between bursts. This gives any charge on the cable plenty of time to dissipate between FLP bursts. Due to the characteristics of FLP bursts they also fall within the same emissions level as normal 10BASE-T traffic.

Management Support

The ideal auto-negotiation scheme should allow sufficient flexibility in its mode of operation to provide not only smooth configuration to the highest common denominator, but the ability for a management agent to control the process. The following will show how the NWay solution supports management in identifying modes and controlling the configuration. However, even though these management capabilities exist, a given implementation does not require the existence of a management agent. NWay's basic auto-negotiation will take place regardless of whether a management agent utilizes these features.

Manual Override

To provide interoperability with future products which choose not to implement the auto-negotiation scheme NWay provides an override switch which freezes the auto-negotiation mechanism from taking place. The management agent then configures the desired forced mode.

Renegotiation

Auto-Negotiation must take place on power up or link reset because this accounts for the actions that a user can physically take after which the device needs to discover whether a mode change is required. However, it does not account for a power user or management agent who wishes to change the state of the link to serve their own ends. To do this NWay has provided two mechanisms by which management can signal a remote agent to renegotiate.


1. The local agent updates its local capability register and resets the `technology_link_fail` counter to zero. FLP bursts will be sent to the remote station immediately following the deferral of any data transmissions, without bringing the remote station's link down. The remote station recognizes the new capabilities and stops its own data transmission gracefully. Once the remote agent receives the

required three consecutive and consistent patterns, it sets the acknowledge bit. Configuration in the new mode occurs once the acknowledge bit has been seen set two consecutive times with the same data. On configuration the data transmission deferral of each station is released.

2. The local agent updates its local capability register and sets the technology_link_fail counter to ensure that the remote station will enter the link fail state. This will allow renegotiation to occur without relying on the remote agent to halt data transmissions.

Remote Fault Sensing

NWay can detect whether another remote NWay device has a wiring fault on the remote agent's receive pair. A local NWay agent determines that a remote fault condition exists by receiving intelligible FLP bursts which do not have the Acknowledge bit set and seeing that the Fault_Timer has expired even though it is sending FLP bursts which have the Acknowledge bit set. A management agent can then record and pass the information to an appropriate MIB.



If Remote Fault Sensing is to be supported, even though it is not required for basic operation, the Remote Fault Detect state in the NWay Arbitration State Machine must be implemented. The Fault_Timer starts as the NWay Arbitration State Machine transitions to the Acknowledge Detect state. A minimum of one second is recommended for the Fault_Timer timeout value.

MIB Statistics

At power up all negotiation takes place during link fail, before data transmission is allowed, therefore NWay's out of band solution does not affect any MIB statistic. A management decision to reconfigure the link needs to account for ramifications to the MIB statistics.

Protocol Connection Timeouts

NWay's nominal time to negotiate ensures that protocol connections do not timeout due to a management decision to renegotiate.

Parallel Advertisement

By receiving all the capabilities of a remote station simultaneously, NWay provides a convenient mechanism to store and pass remote capabilities to a management agent.

Operational Examples

The following operational examples highlight the steps that are taken when attempting to establish a network connection between two nodes. Only three general cases are necessary to demonstrate all the different combinations of operation that will ever occur. The three cases are any NWay capable device to 10BASE-T, any NWay capable device to another NWay capable device, and any NWay capable device to a non-NWay capable device. All these examples fall into one of these three cases.

NWay Capable <---> 10BASE-T in Link_Fail state

This example highlights NWay's 100% backward compatibility with 10BASE-T. The 10BASE-T station in this example begins in the Link_Fail state and is transmitting Normal Link Pulses. The NWay capable station supports an arbitrary 100Mb/s/10BASE-T capable or 100Mb/s only capable node. Initially the NWay capable station comes up in the Link_Fail state and transmits Fast Link Pulse bursts to advertise its capabilities.




FIGURE 15. Example: NWay Capable <---> 10BASE-T

The Normal 10BASE-T node receives the Fast Link Pulses and remains in Link_Fail state because the timing of the Fast Link Pulses do not allow a normal 10BASE-T station to misinterpret FLPs as Normal Link Pulses. To understand this, there are two main cases to consider.

In the first case, the 10BASE-T only station begins receiving FLPs in between bursts. The link_test_min_timer (2-7ms) expires and the first pulse in an FLP burst gets counted as a good link pulse. However, subsequent pulses in the FLP burst arrive before the link_test_min_timer expires for a second time which resets the good link pulse counter to zero.




FIGURE 16. Case 1: Reception begins in between FLP bursts

In the second case, the 10BASE-T only station begins receiving FLPs just before the last pulse in an FLP burst. This pulse is not counted as a good link pulse because the

link_test_min_timer has not expired yet. This pulse therefore keeps the good link pulse counter reset at zero. This case then proceeds like the first case.




FIGURE 17. Case 2: Reception begins just before the last pulse in an FLP bursts

No scenario causes the good link pulse counter to count higher than 1, which is less than the 2 to 10 required good link pulse specification, so the 10BASE-T only station will **never** mistakenly enter the Link_Pass state.

The 10BASE-T Link Integrity Test state machine is a required part of NWay, regardless of 10BASE-T data transmit/receive capability. The NWay capable 100/10 node recognizes the NLPs being sent by the normal 10BASE-T node, switches over to 10BASE-T operation, and sends NLPs. The two stations then both communicate in 10BASE-T mode. If the NWay capable station is 100Mb/s only, then it must still recognize 10BASE-T NLPs being sent. In response to NLPs, FLP burst transmission is halted and no NLPs are transmitted by the 100Mb/s only station.

NWay Capable <---> 10BASE-T only in Link_Pass state

This example highlights NWay's 100% backward compatibility with 10BASE-T. The 10BASE-T station in this example begins in Link_Pass state and is transmitting Normal Link Pulses or normal 10BASE-T traffic. The NWay capable station is an arbitrary 100Mb/s/10BASE-T capable or 100Mb/s capable station. Initially the NWay capable station comes up in Link_Fail state and transmits Fast Link Pulses to advertise its capabilities.




FIGURE 18. Example: NWay Capable <---> 10BASE-T

The Normal 10BASE-T node receives the Fast Link Pulses and always remains in Link_Pass state because the link_loss_timer does not expire.

The NWay capable node's 10BASE-T specific link integrity test state machine identifies the link as being good due to receipt of 10BASE-T traffic or normal link pulses.

If no 10BASE-T data transmit/receive capability exists, then FLP bursts are halted, and no NLPs are transmitted by the NWay capable station, otherwise NLPs are transmitted and the link is established.

If some 10BASE-T Link Integrity Test state machine implementations transition from Link Pass state to Link Fail state due to FLPs, we are back to the previous NWay capable <---> 10BASE-T in Link Fail state example.

NWay Capable connection with multiple common capabilities

This example highlights NWay capable station to NWay capable station auto-negotiation when multiple common protocols are supported by both stations. In this case it is assumed that both stations support Full Duplex 100BASE-TX, 100BASE-T4, and 10 BASE-T.




FIGURE 19. Example: Two NWay Capable Stations with Multiple Common Protocols

Both nodes come up in Link_Fail state and transmit FLP bursts encoded with their capabilities. Each node learns its partners capabilities. In this case each node learns that it has Full Duplex 100BASE-TX, 100BASE-T4, and 10 BASE-T in common.

The built in priority resolution table in each node indicates that Full Duplex 100BASE-TX is the highest common denominator protocol, so both nodes switch over to this protocol and begin transmission.

NWay Capable connection between 100BASE-T4 nodes

This example highlights a connection between two 100BASE-T4 NWay capable station. In this case it is assumed that each station supports 100BASE-T4 as its highest common denominator protocol.




FIGURE 20. Example: Two NWay Capable 100BASE-T4 stations

Both stations comes up in Link_Fail state and begin transmitting FLP bursts. Each station receives the other's capabilities and engages the highest common denominator 100BASE-T4 protocol.

At this time the 100BASE-T4 protocol can proceed through its own protocol-specific link integrity test. This test can then do a specific test to insure that the two pairs of cable not used for NWay signalling are intact.

Summary

The proliferation of CSMA/CD-based LANs will spawn a new generation of multi-function hubs and nodes. The trend among hub vendors towards 'virtual networks' wherein ports can be switched between LANs will necessitate that auto-negotiation capability be embedded in each node. Unless the industry adopts an auto-negotiation solution such as NWay, users and information services managers will face difficult configuration management problems in the future.

NWay solves the problem of configuring the growing available matrix for RJ-45 CSMA/CD compatible LANs by providing a simple, robust scheme to advertise, detect, and resolve the capabilities of two stations connected with UTP. With NWay, the upgrade occurs automatically and seamlessly from 10BASE-T to desired higher performance modes without requiring the presence of a management agent.

NWay interoperates with all 10BASE-T compatible devices such that the installed network base will not be affected by the new generation of RJ-45 LAN technologies if NWay is adopted as a standard. NWay can be applied universally across node and hub products because the architecture is symmetric in nature and does not establish a master/slave hierarchy.

A low gate count and simple architecture make NWay both easy to understand and cost effective. The cost and complexity are realistically bounded by providing a rich enough code space to meet the current projected higher performance mode and future LAN technology needs. Although this proposal addresses the problem for CSMA/CD protocol LANs, NWay does not preclude other technologies from leveraging this standard communication mechanism.

NWay is also a powerful management tool which allows:

1. Manual override of auto-negotiation
2. Renegotiation on management intervention
3. Remote Fault Sensing
4. Parallel Advertisement of capabilities
5. Storage of remote capabilities

NWay's negotiation time is bounded for good links sufficiently to ensure that popular network protocol connections do not time out. Even in a noisy UTP environment, the handshaking and redundant pattern comparison ensures robust operation. Also, the DC line effects and emissions do not present a significant impact to either interoperability or compliance with FCC requirements.

Appendix Overview

The following sections present further information on different aspects of auto-negotiation and NWay:

- NWay Questions and Answers
- 10BASE-T Review
- Link Pulse Template
- DC Balance
- FCC Compliance
- Glossary

NWay Questions and Answers

The problem of maintaining interoperability is serious enough to attract the attention of a variety of network silicon manufacturers and node/hub equipment vendors. Initially developed by National Semiconductor, the NWay auto-negotiation signaling method is the result of collaboration and has been reviewed by a host of technical contributors including representatives of the following companies:

3Com
AMD
Cabletron
Cirrus Logic
DEC
Fujitsu
Intel
Level One
Microlinear
Mick Group
National Semiconductor
Sun Microsystems
Seeq
SynOptics

The following section highlights some key operational features and answers some commonly asked questions about NWay and interoperability.

Q: What is the key purpose and function of NWay?

A: As incompatible 10/100 Mb/s network products become prevalent, proper installation and configuration with minimal network disruption becomes more difficult for system administrators and installers. NWay provides the capability of automatically negotiating the network type of a point-to-point link for RJ-45 LANs. Upon initialization, NWay determines the mutual modes of interoperation between two stations and automatically configures them to take maximum advantage of their

features. NWay also allows a network management agent to reconfigure the point-to-point link as required.

Q: How does NWay's configuration signaling method operate?

A: NWay's communication mechanism is based on transmitting a burst of Fast Link Pulses (FLPs). FLP bursts provide a fast, clean communication mechanism for exchanging configuration information between the two ends of a point-to-point link. Each FLP burst encodes 16-bits of data that defines the unique capabilities of the station. A robust handshake ensures that both ends complete the configuration negotiation.

Q: What are the key features of NWay?

- A:
1. Robust automatic negotiation to the highest performance capability supported.
 2. Interoperable with 10BASE-T.
 3. Rich enough code space for forward compatibility and flexibility.
 4. Easy to implement.
 5. Cost effective.
 6. Remote Fault Sensing during configuration.
 7. Management support for remote capability storage and forced configuration.

Q: Is there any possible way that the fast link pulse scheme can interfere with 10BASE-T's Normal Link Pulse (NLP) operation?

A: NWay is designed specifically to be transparent with normal 10BASE-T link pulse operation. FLP bursts cannot falsely enable 10BASE-T link activity.

In particular, FLP bursts cannot cause a 10BASE-T node to transition from the Link Test Fail to the Link Test Pass State. The closely spaced FLPs do not allow the Link Pulse Counter to increment to two, the condition required to enter the Link Pass State. During the FLP burst transmission, each link pulse is spaced 62.5 μ sec apart. The first link pulse of an FLP burst can set the Link Pulse Counter to one. Each subsequent pulse in an FLP burst resets the Link Pulse Counter to zero because the tight FLP spacing violates the 2 ms keep out time required after a link pulse. In case the 10BASE-T node powers up in the Link Test Pass state, it will transmit traffic or

NLPs to the NWay capable station, causing the Link Integrity State Machine of the NWay capable station to move to the Link Test Pass state. This will result in either of the two following conditions:

1. NWay capable station has 10BASE-T protocol - It will cease FLP burst transmission and will establish regular 10BASE-T communication.
2. NWay capable station does not have 10BASE-T protocol - It will cease to transmit FLP bursts, ultimately causing the remote 10BASE-T node to move to the Link Test Fail state.

Q: Is there any possible way that the fast link pulse scheme can interfere with normal 10BASE-T data transmission operation?

A: No. The overriding design goal of NWay is to maintain 100% backward compatibility with the installed base of 10BASE-T nodes. NWay clearly separates data transmission from link pulse transmissions. Since Fast Link Pulse (FLP) transmission does not occur during regular data transmission there is no chance that a burst of link pulses can interrupt normal data transmission.

Q: How does NWay ensure that the point-to-point link is configured correctly?

A: NWay protocol relies on redundant transmissions to ensure that the transmitted 16-bit pattern is correctly received by the remote station. The remote station must receive an identical 16-bit pattern three times in a row before the acknowledgment pattern is transmitted. The acknowledgment pattern is the same 16-bit word with the acknowledge bit set. The reception of 2 out of 3 acknowledgment patterns completes the handshake. This robust protocol assures that the correct configuration information is exchanged.

Q: Where there is more than one mode of interoperability, how does NWay decide which mode to establish?

A: After passing configuration information via FLP bursts to determine the mutual modes of interoperation of two stations, NWay automatically defaults to the highest common denominator of the capabilities. The priority table is intended to be defined in the 802.3 auto-negotiation standard. The fastest and most capable modes

(100BASE-T4 and 100BASE-TX) are at the top of the hierarchy.

Q: What is the size and complexity of the NWay solution?

A: NWay provides an efficient mechanism for point-to-point communication that is efficient in terms of the amount of logic gates required. Regardless of the number of protocols supported, a total of 3 state machines, 6 timers, 3 counters, and some peripheral logic is needed. An analysis of the architecture shows that the design will require approximately 400 cells to implement. In addition, NWay has been architected as a simple mechanism which allows easy comprehension by silicon suppliers.

It should be noted that this level of complexity remains fixed as new technologies emerge and need to coexist with the growing installed base of RJ-45 LANs. Alternative schemes that require separate and complex state machines for each pattern will be more complex and expensive to implement, especially as new LAN options are added.

Q: How does NWay provide for the future expansion of LAN capability options?

A: NWay was designed with a rich enough code space to accommodate present and future RJ-45 LAN devices without the need to redesign the core communication mechanism. In addition to the three reserved bits for future capabilities, NWay's data structure features a Next Page Bit, which provides a mechanism to transmit additional protocol specific information. In addition, the Protocol Discriminator field will allow non-CSMA/CD RJ-45 LANs to coexist with CSMA/CD LANs, keeping 10BASE-T as the multi-function product lowest common denominator.

Q: How will NWay maintain compatibility when new options are added?

A: As new CSMA/CD-compatible LAN technologies enter the market a reserved bit will be assigned to each technology by the standards body. The new technology will be inserted into an updated priority table, as detailed in the Data Bit Field description and made a part of the 802.3 auto-negotiation standard. The relative hierarchy of the existing technologies will not change, thus providing backward compatibility. It is important to note that the reserved bits are forced to zeros. This guarantees that devices implemented using the current priority table will be forward compatible with

future devices using an updated priority table.

Q: What does the parallel advertising feature buy you?

A: NWay allows a station to advertise its complete set of CSMA/CD capabilities in one burst, an NWay feature referred to as its "parallel advertising" capability. Advertising the complete set of capabilities in one pass allows NWay to immediately discover all capabilities of the remote station and save them for passage to a management agent.

Q: What if there isn't a lowest common denominator between nodes?

A: Obviously, a link will not be established. Yet if both have the NWay autodetect mechanism, there can be a configuration management advantage. Each node is capable of remotely sensing and communicating the capabilities of the other end of their link to another entity, such as a management agent.

Q: Does NWay support renegotiation?

A: Yes. NWay's reconfiguration capability is an important feature that allows any mutual mode of interoperability to be utilized. Although a pair of stations default to their highest common denominator capability, other modes of interoperation may exist. For a variety of reasons, the network management agent may wish to reconfigure the link.

Q: What impact will embedding NWay have on intelligent hubs?

A: Hub makers are already moving towards more intelligent products with features like port switching that allow them to reconfigure ports. Hubs that support multiple capabilities (100BASE-TX, 100BASE-T4) will need an autodetect configuration mechanism to determine the capabilities that are attached to each port. Without autodetect, the system administrator will have the duty of manually tracking each station's capabilities and individual requirements.

Q: Do the fast link pulses cause any EMI/RFI noise problems for FCC qualification?

A: The emissions due to unipolar FLPs are no worse than the case of short legal packets in CSMA/CD. A comparison of FLP bursts and CSMA/CD frequency spectrum simulation results show that the energy content of the FLP burst is no

higher than that of 10BASE-T CSMA/CD. As a result, NWay will not cause products to have additional difficulty passing stringent FCC requirements.

Q: Do the fast link pulses cause any DC balance problems?

A: The voltage/charge of each FLP dissipates to almost zero (.000045 of the initial value) before the next FLP within a burst. This number is based on worst case load models for transmit and receive ends. This also assumes that the medium for transmission and reception is compatible to standard 10BASE-T. The minimum spacing between link pulses must be 12 μ s to avoid appreciable DC effects. The spacing between FLPs is a minimum of 31.25 μ s, which provides a margin that guarantees a robust solution.

10BASE-T Review

A 10BASE-T Link Pulse (NLP) is a signal applied on a network segment to determine its integrity.

An NLP has a pulse width of 100ns nominally with a 200ns maximum. These pulses are transmitted every 16ms +/- 8ms while a link is active and not transmitting regular packet data.



FIGURE 21. Normal Link Pulse Timing

Link pulses are used to control a stations entry to and exit from the Link_Loss_State. Link_Loss_State is entered when no NLP pulses are for more than link_loss_timer time which ranges from 50ms minimum to 150ms maximum.

Link_Loss_State is exited when a minimum of 2 NLP pulses are detected or up to 10 NLP pulses are detected.




FIGURE 22. Link_Loss_State Transitions

Link Pulse Template




FIGURE 23. Link Pulse Template

DC Balance

The medium for the transmission and reception of fast link pulses (FLPs) must be 10BASE-T compatible as described in section 14 of the ISO/IEC 8802-3 specification. This implies that the standard 10BASE-T hardware will be able to handle the fast link pulses.

This appendix investigates the charging effects, baseband wander and the issue of DC balance in general during the transmission of FLPs over 10BASE-T compliant hardware/cable.

As per the 10BASE-T specification the worst case load models for the transmit and receive ends of the cable are specified as shown in Fig. A-1. The worst case tolerance limits on resistance and inductance values are also specified in this figure. Any standard 10BASE-T pulses transmitted on such hardware must also satisfy the waveform template specified in 10BASE-T specification which is reproduced here as Fig. A-2.

The FLPs transmitted over the above specified 10BASE-T hardware must comply with the voltage template specified using the worst case models for transmit and receive ends under the worst case FLP spacings. The theoretical worst case pulse spacing within an FLP burst is 31.25 μ s. Time constant calculations will be presented to analyze the charging effects. The DC baseband wander calculation will be presented to show the overall effect of these pulses on the signal levels on the cable.

TIME CONSTANT CALCULATIONS:

Worst case modeling of the twisted pair interface and cable gives the following time constants:

Minimum Time constant: $\tau = L/R = 180/115 = 1.5\mu$ s.

Maximum Time constant: $\tau = L/R = 220/76.8 = 3.0\mu$ s.

It takes approximately 4 time constants to dissipate accumulated charge to $e^{-4(\tau/\tau)} = 0.018$ of the initial value where $4\tau = 4 \times 3\mu\text{s} = 12\mu\text{s}$.

This implies a minimum pulse spacing of $12\mu\text{s}$. The margin provided in this scheme is $(31.25\mu\text{s} - 12\mu\text{s})/3\mu\text{s} = 19.25/3 = 6\tau$. So, the fast link pulses provide the minimum 4τ spacing plus a 6τ margin to avoid line charging. In the total of 10τ spacing between pulses the voltage/charge dissipates to 0.000045 of its initial value. Hence line charging is not an issue for FLPs.

BASEBAND WANDER:

The DC component of the unipolar FLP's can be calculated as follows:

DC Component = (Duty Cycle) * (Max. Voltage Swing)

$$= (100\text{ns}/62.5\mu\text{s}) * (2.5\text{V}) = 4\text{mV}$$

This DC component of 4mV corresponds to 0.16% of the voltage swing of the FLPs. This amount of DC wander is insignificant because it is considerably below the receiver tolerance/sensitivity limits.

The FLP's are designed to work reliably on all 10BASE-T compatible transmission and reception media.

FCC Compliance

Simulation results of the frequency spectrum of continuous Fast Link Pulses versus short continuous legal packets in CSMA/CD (10BASE-T) are presented in Figure 24, "FCC Emissions (Simulation Results)". The emissions from Fast Link Pulses are within the same range of those of CSMA/CD.




FIGURE 24. FCC Emissions (Simulation Results)

Glossary

10BASE-T	IEEE 802.3 Physical Layer Specification for Ethernet over two pairs of Unshielded Twisted Pair (UTP) media.
100BASE-TX	Proposed 100Mb/s CSMA/CD LAN over 2 pairs of Category 5 UTP.
100BASE-T4	Proposed IEEE 802.3 Physical Layer specification for Ethernet over four pairs of Category 3, 4, and 5 UTP or STP wire.
Acknowledge Bit	A bit used by the NWay protocol to indicate that a station has successfully and consistently received the far-end capability data word. With the simple, yet robust NWay protocol, this bit is only set after the data word has been received consistently three times in succession.
Auto-Negotiation	The ability for a local station to detect the network capabilities of a remote station. In addition to autodetect, NWay allows the local station to automatically configure the link.
Capability Field	Five bits of NWay's 16-bit data field that are used to indicate the network capabilities of a node, such as support for 10BASE-T, 100BASE-TX, 100BASE-T4, as well as Full Duplex capabilities.
Category 5 UTP	A data grade of unshielded twisted pair copper wiring with superior electrical characteristics.
Category 3 UTP	A voice grade of unshielded twisted pair copper wiring that is used by current 10BASE-T installations.

CSMA/CD

Carrier Sense Multiple Access/Collision Detection (CSMA/CD) is the Media Access Control (MAC) protocol for IEEE 802.3 LANs.

FLP

Fast Link Pulses are the basis of NWay's efficient configuration signaling mechanism that is based on sending a series of narrow 100 ns pulses between stations to exchange network capability information. Upon Link Reset, a local station sends bursts of fast link pulses (FLPs) to the remote station to convey the local station's network capabilities, and vice versa. The FLP burst consists of 33 pulses (17 clock pulses and 16 data bits) that are sent in succession spaced 62.5 μ sec apart over a 2 ms period. This FLP burst is followed by a 14 ms pause, making the overall FLP cycle a total of 16 ms.

Full Duplex

A type of Ethernet networking using 10BASE-T media and a switching hub. With a simple modification to the MAC at both ends of a point-to-point link, the Full Duplex agents can transfer data in both directions simultaneously. Full Duplex mode is becoming an increasingly popular way to increase performance on point-to-point links.

IPG

Inter-Packet Gap. The minimum idle time required between packets. For 10BASE-T Ethernet the IPG is 9.6 μ s.

Link Pulse

Out of band communication mechanism used in 10BASE-T networks to indicate link status. 10BASE-T nodes exchange information using a single 100 ns wide unipolar pulse which is sent every 16 ms \pm 8 ms when the network is idle to establish or maintain connection.

MAU

Media Attachment Unit. A physical device that is used in Ethernet to transmit signals from the PLS onto the medium. This is sometimes referred to as a transceiver.

MAC

Media Access Control. A protocol, such as CSMA/CD, that dictates how a device may gain access to the physical media.

MIB

Management Information Block.

Next Page Bit

A bit in NWay's 16-bit data structure that indicates that there is additional information regarding the node's capabilities that resides in an additional one or more 16-bit data transmissions.

NLP

Normal Link Pulse. Refers to link pulses as they are used in standard 10BASE-T networks. A single 100 ns pulse is sent nominally every 16 ms to indicate link status.

Protocol Descriptor

Five bits in NWay's 16-bit data structure that are used to indicate and convey the network station's protocol, such as CSMA/CD.

Priority Table

The look-up table method used by Nway to arbitrate between possible network connection types (100BASE-TX, 100BASE-T4, 10BASE-T) when more than one common network capability exists. The priority table defines the hierarchy of connection types from the highest common denominator to the lowest common denominator that will be automatically configured.

PHY

Physical Layer Protocol. The PHY is a sublayer of the lowest layer of the OSI MODEL which is responsible for encoding data for transmission.

PLS

Physical Signaling sublayer of the Physical Layer of the OSI MODEL. This sublayer is responsible for the encoding of data for transmission.

Renegotiation

The capability of the local station to request a different but common mode of communication.

UTP

Unshielded Twisted Pair copper wire.

Special Thanks

We would like to thank the following people who were instrumental in the development of this document:

Paul Ahrens
Bill Bunch
Walter Hurwitz
Srinivas Kola
Shankar Mukherjee
Prasun Paul
Ramin Shirani