
Guidelines for Implementing Low-Cost Ethernet on the Motherboard

By Dr. Robert (Bob) Bridge, Director of LAN Products

Crystal Semiconductor Corp.
4210 S. Industrial Ave.
Austin, TX 78744
USA
Phone: 512-912-3248
Fax: 512-912-3871
email: rfb@crystal.cirrus.com

1. ABSTRACT

This paper provides guidelines implementing a low-cost Ethernet solution on a motherboard. In 1997, 10 Mbps Ethernet will be the preferred Ethernet technology for value-priced commercial desktop systems. An Ethernet component vendor should be chosen who can provide in-depth software support, and who has a product meeting all the latest industry standards including Wakeup-via-LAN, full duplex, auto-negotiation and PC97. Plans must be made for uniquely programming an EEPROM per motherboard. Design attention must be paid in the Ethernet subsystem to mechanical design, minimizing board space, and minimizing emissions. Meeting emission standards can be challenging since Ethernet uses unshielded cables, which can act as a broadcast antenna for the system.

2. INTRODUCTION - LAN TRENDS

In the last year, there has been a significant increase in the number of Ethernet on the Motherboard PC systems. What has caused this increase? The two most significant factors are the low-cost of the newest Ethernet controller ICs, and the fact that

Ethernet has become the defacto standard for networking commercial PCs.

In North America, between 85% and 90% of all new commercial PCs are networked using Ethernet. Competing technologies, such as Token Ring, are losing market share. Ethernet continues to gain share because it is low cost, and because of new Ethernet technologies that increase network bandwidth. 10 Mbps Ethernet switching technology, Fast Ethernet and Gigabit Ethernet are dramatically increasing the bandwidth of Ethernet LANs, and lessening marketplace interest in alternative high-speed technologies such as 100BASE-VG and ATM. 10 Mbps switched bandwidth allows network bandwidth to be increased without changing and hardware or software in the PC.

The industry has seen a dramatic increase in the sales of 10/100 LAN adapter cards. In 1997, the number of 10/100 Mbps Ethernet network interface controllers will rise to 40% of the total Ethernet market. Most of the 10/100 Mbps controllers will be used at 10 Mbps. 100 Mbps Ethernet will actually be used primarily in the LAN backbone network and servers, and will have only limited

deployment to the desktop client machines. A typical corporate network will use 10 Mbps to connect to desktop clients to 10 Mbps hubs. The 10 Mbps hubs will connect to 10 Mbps ports on a Ethernet switch. The switch will have connections at 10 or 100 Mbps to other switches and to servers. When using such a segmented and hierarchical Ethernet network of hubs and switches, using 10 Mbps bandwidth to the desktops provides ample bandwidth for typical client applications such as Microsoft Office.

In light of these trends, a typical response of PC systems vendors is to deploy different Ethernet technology in different classes of systems. High-end clients and servers ship with 10/100 Ethernet capability. Value-based, low-end commercial systems ship with 10 Mbps technology. This Ethernet differentiation is driven by the cost premium required for 10/100 Mbps technology. In 1997, 10/100 Mbps Ethernet ICs will have a component cost approximately 3-to-4 times higher than 10 Mbps Ethernet ICs.

3. CHECKLIST FOR DESIGNING ETHERNET ON THE MOTHERBOARD

The following sections provide a designer's checklist of the issues to be considered in selecting Ethernet components and implementing the Ethernet Motherboard design.

3.1 Low-risk

To minimize motherboard design risk and schedule risk, it is important to select an Ethernet vendor who provides high-quality technical support, who is shipping well-proven device drivers and ICs and who provides high-quality documentation. The vendor should be able to support any required customization of the drivers.

3.2 Compatibility with latest IEEE 802.3 Ethernet requirements

The IEEE 802.3 committee is responsible for orig-

inating Ethernet specifications. These standards are then forwarded to ISO for international ratification. The IEEE 802.3 committee remains very active as witnessed by the progression of Ethernet standards including the original thick-coax Ethernet, Cheapernet, 10BASE-T, switching, 100 Mbps Ethernet and Gigabit Ethernet. It is important for the PC motherboard to provide those Ethernet features most desired by the marketplace. As stated above, most customers are satisfied with 10 Mbps Ethernet for value-priced PCs. Europe has a strong installed base of 10BASE-2 (Cheapernet), and may require a BNC connector on the PC case. Cheapernet is no longer required by large corporate customers in North America, who only require an RJ45 connector for 10BASE-T. This same connector can also support 100BASE-TX. TX has become the defacto physical layer for 100 Mbps Ethernet.

Corporate customers who want an easy migration path to 10/100 hubs and full-duplex switches will want Ethernet auto-negotiation to be implemented in their 10 Mbps Ethernet PCs. Auto-negotiation simplifies network administration by corporate network managers since new hubs and switches can automatically learn the capabilities of the PC at the other end of the cable. For example, the hub can determine automatically if the PC is capable of full or half duplex operation, and 10, or 100 Mbps operation. If auto-negotiation does not exist, each port on a hub or switch will have to be manually configured for the correct duplex mode and speed, and manually reconfigured whenever a PC's Ethernet capabilities change.

Full duplex is a technology enabled by switching that doubles Ethernet throughput to a desktop system, even if that system remains at a 10 Mbps speed. The original Ethernet implementations were half-duplex, in which case all the stations on a network segment share one 10 Mbps bus. In half duplex, only one station on the cable can transmit or receive at the same time. With the advent of Ether-

net switching, every station can simultaneously transmit and receive at the full network bandwidth. The Crystal LAN™ CS8920 Plug and Play Ethernet controller is an example of a controller that supports full duplex and auto-negotiation.

3.3 Compatibility with Wakeup-via-LAN Standards

A major thrust of North American corporate customers is managing the Total Cost of PC Ownership. The total cost of ownership includes factors such as:

- a) the cost of acquiring PCs
- b) the cost of PC support and help-desk costs
- c) the cost of electricity
- d) the cost of end-user training

Systems vendors have developed a variety of innovative technologies to decrease the cost of ownership. Examples of such tools are HP's OpenView system and IBM's NetFinity.

To minimize power consumption, end-users usually power off their PCs at night. If a PC is powered off during the night, systems such as NetFinity can not use these idle hours to perform PC maintenance. Instead maintenance must be done during the day when users are using their machines for other purposes.

A PC administration tool might want to perform the following types of tasks at night:

- a) updating an application on all clients to the latest application revision level,
- b) performing a virus scan on all clients,
- c) finding and removing unauthorized computer games from all clients,
- d) checking all clients for a change in hardware configuration (identifying any theft of memory modules or looking for non-authorized adapter cards).

Furthermore, being able to access a powered-down PC can also enable a mobile user to download files via a remote LAN-access application from his desktop machine to his portable machine.

Access to a powered-down PC requires a Wakeup-via-LAN capability. The MagicPacket™ technology is the current industry-standard for waking up a powered down machine, and is used, for example, by HP OpenView and IBM NetFinity. Magic Packet defines a specific data pattern that is embedded in a standard Ethernet frame. This data pattern is detected by the LAN controller IC that then sends a signal to the PC's Advanced Power Management (APM) circuitry. The APM circuitry can then wake up the CPU and system. In this approach, the LAN controller IC is powered by the PC's Auxiliary power supply, so that when the PC is powered off, the LAN IC is still powered on. This technology is difficult to implement on an adapter card because it requires two non-ISA signals to be routed to the LAN controller IC. Those signals are Auxiliary Power supply input, and Wakeup indication output. With an Ethernet on motherboard implementation, these signals are easily provided to the LAN IC. For a complete Wakeup-via-LAN application schematic, contact Crystal Semiconductor.

Only a limited number of Ethernet controller ICs implement this new Magic Packet technology. One such IC is the CS8920.

Not all of the Magic Packet-compatible products work the same from the system viewpoint. After a system reset, some Ethernet controller ICs will not recognize Magic Packets frames until the correct software status bit has been set in the controller register. If a machine has had its power supply interrupted at night (for example, by a lightning strike in the neighborhood), the PC may not come up in a mode where Magic Packets are recognized. Therefore an automated PC management tool could not wake up the PC. In contrast, the CS8920 can automatically enter Magic Packet recognition

mode, without requiring software intervention, by checking for lack of activity on the ISA bus.

3.4 Compatibility with Networking Software Standards

LAN device drivers reside between the Operating System (OS) and the Ethernet controller IC. The device drivers understand how to program and read the IC's control and status registers, and how to transfer user data between the IC and the PC main memory via the ISA bus. On the OS side, the drivers provide the standardized services and functions required by the OS, and hide all details of the IC hardware from the NOS.

The vendor of the Ethernet controller must provide a full set of certified device drivers. The IC vendor must also have the internal technical support infrastructure necessary to deal with any field problems that arise. The device drivers, the protocol stacks, operating systems and applications that run on top of the OS are complex. Subtle software problems can arise which need to be isolated to device drivers or operating system software or application software. If the vendor of the Ethernet IC uses another vendor's device drivers or maintains their device drivers using external contractors, then that company can not be depended upon to provide good application support. In such a situation it may be necessary to add LAN device drivers and software experts internally to your company.

Also, device driver requirements are constantly evolving. For example, Microsoft has just defined an NDIS 4 communication device standard that is being first employed with Windows NT 4.0. Therefore, look for vendors whose drivers are Microsoft and Novell certified, and have the staffing to update the drivers as required by new standards.

The minimum driver set expected by the marketplace includes:

- a) Novell-compatible DOS ODI Client for DOS, Win 3.1 and Windows for Workgroups 3.11
- b) Novell-compatible OS/2 Client
- c) Novell-compatible Netware for Novell 4.X, 3.x Servers
- d) Microsoft-compatible NDIS 2 for DOS, Win 3.1 and Windows for Workgroups 3.11
- e) Microsoft-compatible NDIS 2. for OS/2
- f) Microsoft-compatible NDIS 3 for Windows NT Server, NT Workstation, Windows '95, Windows for Workgroups
- g) Starting in July, 1997, Microsoft-compatible NDIS 4 for Windows NT Server, NT Workstation, and Windows '95
- h) DOS Packet Driver V1.09
- i) SCO UNIX compatible driver
- j) Boot PROM code for Remote Initial Program Load
- k) Setup & Installation Utility for DOS
- l) Utility for EEPROM programming
- m) Diagnostic for testing Ethernet on the Motherboard

The Crystal LAN™ drivers are an example of such a set of drivers.

3.5 Support of Boot PROM and Remote Initial Program Load

If the PC provides the option of booting the OS from a server (also described as Remote Initial Program Load, or RIPL), then the motherboard will need to contain a Boot PROM. The PROM contains a RIPL device driver. This RIPL driver contains enough of the protocol stack to allow communication with a server over the network for the purpose of downloading DOS or Windows 95 or some other OS. As part of the OS downloading, the complete Ethernet device driver is loaded onto the PC. There needs to be close cooperation between the RIPL code and BIOS, so choose an Ethernet vendor who is prepared and able to make any changes to the RIPL code required by your BIOS code.

3.6 Compatibility with Microsoft PC97 Requirements

PC 97 -compliance will become a required part of Microsoft's WHQL testing in mid-1997. Therefore, machines designed for Fall '97 Comdex need to fully meet PC97 requirements. Key networking requirements for an ISA-based Ethernet solution of PC97 include:

- a) PnP support with at least four I/O Base addresses, and at least four interrupts, and 16-bit I/O address decoding
- b) support of multiple network interface controllers in one machine
- c) ability to change all possible Ethernet resource settings from Windows 95 Device Manager
- d) NDIS 4.0 miniport driver,
- e) auto-sense of media connection, and media type

The CS8920 meets PC97 hardware requirements. Note that the CS8920 has a direct ISA bus interface, and can be accessed in either I/O mode or memory mode. Since current ISA PnP does not guarantee that there will be no memory conflicts in a system, it is recommended that only the IO mode be used in PnP motherboards.

3.7 Low Hardware Cost

The total cost of the Ethernet solution includes the cost of following major components: the Ethernet Controller IC, an EEPROM (see section 3.12), a quartz crystal, external 10BASE-T filters, external buffer memory, boot PROM, and an RJ45 connector. Note that more integrated Ethernet ICs, such as the CS8920, contain internal 10BASE-T filters and internal buffer memory thereby eliminating up to \$2.50 in external component costs.

3.8 Boards Space

Motherboards are often difficult to design due to the large number of components required to pro-

vide all the functionality demanded by today's customers. Therefore, an important criterion for putting Ethernet on the Motherboard is minimizing the total amount of board space needed for the Ethernet components. The CS8920 requires half the board space of many Ethernet solutions, less than 2 square inches (13 square cm) of space for the Controller IC, EEPROM, quartz crystal, transformer and discrete components. If Boot PROM or Wakeup-via-LAN is used, additional board space will be needed.

3.9 High Throughput Performance

The standard benchmark for Ethernet performance is the Novel "Perform 3" test. Perform 3 measures throughput over a network between a client and a server. At 10 Mbps, most high-performance ISA-based Ethernet solutions outperform most PCI solutions. This is due to the fact that ISA-bus IO transfers can have lower latency than PCI-bus DMA transfers. An example of performance results is shown below.

3.10 Meeting EMC (ElectroMagnetic Compliancy) Requirements

In addition to the being concerned about electromagnetic emissions coming directly from the motherboard, keep in mind that the Ethernet cable provides an additional emission path to the outside of the PC case. The cable can act as an EMI (ElectroMagnetic Interference) broadcast antenna.

Meeting emissions standards can be made easier by choosing an Ethernet controller with high-performance internal filtering of the transmitted 10BASE-T signal. Some Ethernet controllers, such as most NE2000-compatible drivers, output two unfiltered square waves that must be feed into an external filter. The output of the filter passes through the transformer before reaching the RJ-45 connector. Other Ethernet controllers do some low-order internal filtering. Other controllers, such as the CS8920, provide high-order internal filtering

(for example, fifth-order Butterworth filter). The more internal filtering that is performed, the less high-frequency energy that will be present on the motherboard traces. In addition, external filters take up valuable board space and cost extra money. Two-layer adapter boards based upon the CS8920, and using no external filters, have passed FCC part 15 class (B) with 9 dB of margin.

When choosing a 10BASE-T transformer, you may choose between isolation transformer and isolation transformer with a common mode choke (CMC). That decision should be based upon the common mode noise that exists on the 10BASE-T lines in a particular system. A common mode choke reduces common mode noise emitted by the 10BASE-T lines. A CMC may be required in certain applications to meet EMC requirements and to meet 10BASE-T common mode output voltage noise specification. The physical dimensions of the isolation transformer and the isolation transformer with a CMC are the same. Both are typically available in a 16 pin DIP or 16 pin SOIC package.

3.11 Layout Guidelines

Ethernet Controller ICs are mixed signal device having digital and analog. While doing the PCB layout and signal connections, it is important to take the following precautions:

- a) Provide a low inductive path to reduce power and ground connection noise.
- b) Provide proper impedance matching especially to the Ethernet analog signals.
- c) Provide low inductive path, wider and short traces, for all analog signals.

The placement of the Ethernet controller should be such that the routes of the analog signals and the digital signals are not intermixing. No signal should route beneath the Ethernet controller on any plane.

A multi-layered printed circuit board (PCB) typically has separate ground and power (VCC) planes.

Often discrete components like resistors and capacitors are placed on the solder side of a printed circuit board. In this case, the power-supply bypass capacitors for the Ethernet controller can be placed on the solder side of the PCB. Each bypass capacitor should be placed beneath the Controller and closest to its corresponding power pin pair.

Four signals are used for 10BASE-T communication: two differential transmit signals and two differential receive signals. An isolation transformer is placed between the transmit and receive traces and a RJ-45 (modular phone jack) connector. The isolation transformer should be placed as close as possible to the RJ-45 connector. Both transmit and receive signal traces should be routed so they are parallel and of equal length. The signal traces should be on the component side and should have direct and short paths. A ground trace should be run parallel to the transmit traces.

More complete design and layout guidelines can be found in the CS8920 Technical Reference Manual.

3.12 Managing IEEE MAC Addresses and Programming EEPROMs

Each node of an IEEE 802 Local Area Network has a unique address for the media access control (MAC). This makes it possible for that particular node to have unique identity for data communication, and prevents confusion on the Ethernet cabling as to which nodes are communicating with each other. This address, known as the IEEE physical address (or MAC address), consists of 48 bits of data. This address is assigned to the CS8920 by the manufacturer of the motherboard, and must be uniquely programmed into an EEPROM connected directly to the CS8920's serial bus. To insure uniqueness of the address, 24 bits of out of the 48 bits of the physical address are assigned to the manufacturer by the IEEE standards committee. This 24 bit address is known as Organizationally Unique Identifier (OUI). The remaining 24 bits of the address are assigned by the manufacturer.

This address is used in the Ethernet frame to identify the source address and destination address of that Ethernet frame. Note that higher-level protocols, such as TCP/IP, also contain addresses, but these addresses serve a different purpose. The TCP/IP addresses allow Internet end-points to be identified. The MAC address allows the LAN end-points on one LAN to be identified.

The motherboard manufacturer must obtain a block of IEEE addresses from the IEEE, and program one address from that block into each EEPROM. Crystal provides a software utility that tracks and assigns the addresses, and which also can program the EEPROM.

For further information and an application for an OUI, please contact the IEEE at the following address:

IEEE Registration Authority,
IEEE Standards Department,
445 Hoes Lane, PO Box 1331
Piscataway, NJ 08855-1331, USA
Telephone: (908) 562-3813
FAX: (908) 562-1571

In addition to storing the IEEE MAC address, the EEPROM contains state information used to configure the CS8920 IC, the device driver and the system PnP resources. The Crystal LAN™ EEPROM program utility will program all of the contents of the EEPROM, including the following Plug-and-Play information:

- a) Vendor ID,
- b) Device Identifier,
- c) Logical Device ID,
- d) Device Serial Number, and
- e) list of available IRQ and I/O base-address resources.

A 128 word (128 X16 bit) EEPROM is large enough to store this information and typically costs less than \$0.45 US.

3.13 Mechanical Design Considerations

Most end-users and MIS managers find it helpful to have at least one LED that indicates successful LAN attachment, at the hardware level. This LED can be mounted on either the back panel, next to the RJ45 connector, or somewhere on the front bezel, for example, close to the hard drive activity LED. Note that the CS8920 has direct drives for the four LEDs. If only one LED is supported, then it is recommended that the designer use either the Link Status LED (which shows that 10BASE-T link pulses are being received) or the Local Activity LED (which shows when this PC is receiving Ethernet frames addressed to this PC or is sending Ethernet frames). The Local Activity LED output signal will be available on the CS8920A.

When mounting the RJ45 connector in your PC, make sure that the exterior end of the connector is not recessed below the surface of the PC case. Also make sure that the area of the release tab is far enough above the desk top so that a finger can reach the tab. In several of the PCs in Crystal's LAN System Verification lab, a screwdriver is required to depress the RJ45 release tab, because the release tab is too hard to reach with a finger..

4. CONCLUSIONS

Ethernet is being added as a standard feature on many commercial desktop motherboards. The Ethernet motherboard sub-system can be implemented with minimal schedule and technical risk as long as the designer understands the guidelines described in this paper.

Crystal LAN is a trademark of Crystal Semiconductor Corporation. Magic Packet is a trademark of AMD. Other product and company names herein may be the trademarks of their respective owners.