

Introduction

Emerging Internet applications such as voice over Internet protocol (VoIP) and real-time streaming video perform poorly when the Internet's core network is congested. Unfortunately, adding bandwidth is not a complete, long-term solution. By prioritizing Internet traffic and the core network more efficiently, quality of service (QoS) and traffic engineering functions can address these emerging Internet applications. Network equipment manufacturers are constantly developing new solutions that solve many of the problems associated with today's Internet applications. Multiprotocol label switching (MPLS) is one such solution and is currently being standardized.

Altera® high-density, high-performance programmable logic devices (PLDs) combined with intellectual property functions, Excalibur™ embedded processor solutions, and the Quartus™ II development software offer a complete MPLS implementation solution. Altera PLDs provide a single-device solution that is flexible, offers significant time-to-market advantages, and implements functions in logic and on a processor.

This application note describes the MPLS standard and how to implement MPLS with Altera PLDs.



For more information about Altera devices, intellectual property functions, and software tools, refer to the Altera web site at <http://www.altera.com>.

General MPLS Requirements

To better address these emerging Internet applications, Internet service providers (ISPs) have the following general MPLS requirements:

- Efficient and manageable packet networks
- Products that support new services, such as virtual private networks (VPNs) and QoS
- Easier integration of internet protocol (IP) and asynchronous transfer mode (ATM)
- Additional bandwidth
- Compatibility between routers of different manufacturers

With the exception of additional bandwidth, the MPLS standard can address all of these issues. The MPLS working group, a group from the Internet Engineering Task Force (IETF), is in the process of standardizing MPLS. The IETF is made up of representatives from various networking companies.

To address ISP's requirements, network equipment manufacturers need to consider the following issues:

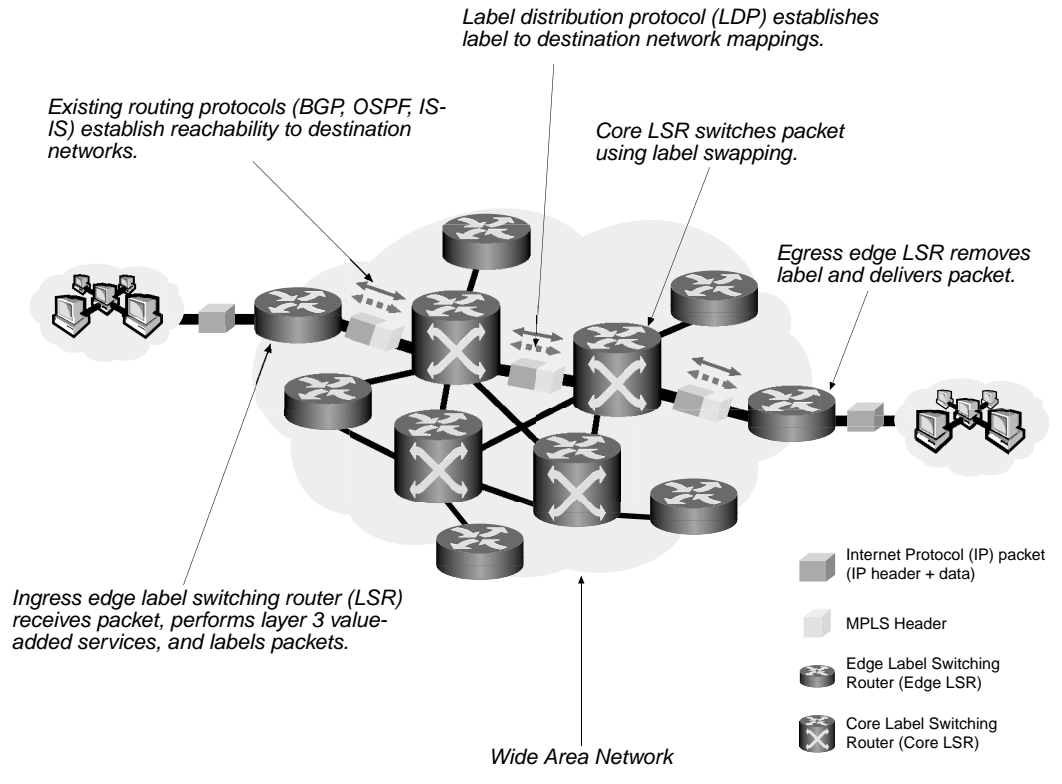
- MPLS does not standardize all ISP's requirements. Therefore, a successful MPLS solution requires a flexible platform. Implementing MPLS in an inflexible platform (e.g., ASICs) results in a short product life cycle.
- Time-to-market is important. The turnaround times imposed by ASICs will not allow the manufacturer to effectively establish a foothold in the market.
- To implement an MPLS system, designers need a combination of logic- and processor-based functions. Separating these functions in different devices results in performance degradation and an increase in design time. Also, a single-chip ASIC solution will be inflexible, increase time-to-market, and quickly become obsolete.

Designing MPLS Systems

MPLS allows efficient management of traffic flows in a network. When a packet enters an MPLS-enabled network, a label is attached to that packet. This label defines the route that the packet should take to reach the required exit point of the wide area network (WAN). Traditionally, when a packet arrives at a router, its destination address is read by the packet processing function when determining the next hop. The next hop is calculated each time the packet arrives at a new hop. In MPLS, since the route is determined at the ingress of the WAN, other important networking functions such as VPNs and traffic engineering can be implemented.

Figure 1 shows how an MPLS network processes a particular packet that does not have established mapping.

Figure 1. MPLS Network Overview



This section describes the following MPLS topics:

- Label switching—how labeled packets are forwarded through an MPLS network
- Label distribution—how labels are distributed in an MPLS network
- Traffic engineering—an important feature of the MPLS framework that allows for efficient management of network resources
- VPNs—a method of using the public Internet to connect remote networks and form a private network

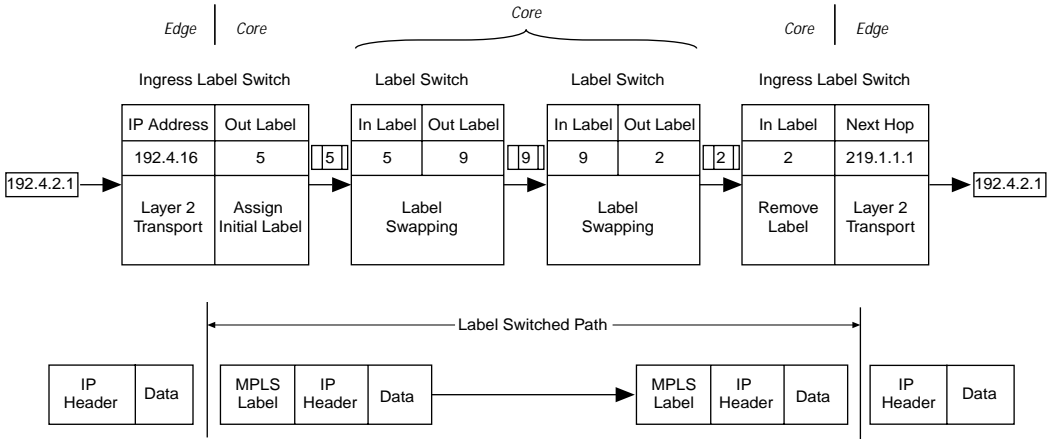
Label Switching

The basic operation of an MPLS network involves switching that is based on a label that has been attached to a packet. When a packet enters an MPLS-enabled network, the LSR analyzes the packet header and assigns a label based on the destination address and other information contained in the packet header. The label that is assigned to that packet distinguishes the route the packet should take to reach the required egress node of the MPLS-enabled network.

When a packet is assigned a label at the ingress of an MPLS-enabled network, its forward equivalence class (FEC) must be determined. An FEC is a group of packets that are forwarded in the same manner (i.e., over the same path, with the same forwarding treatment). Each LSR must keep track of how packets should be forwarded. This FEC information is contained in a library information base (LIB) that includes FEC-to-label bindings.

When a labeled packet is received at an LSR, the input port and label is read and the output port (the next hop) is determined. An outgoing label replaces the incoming label. This outgoing label is in context for the next hop's label switching operation. Figure 2 shows MPLS label swapping.


Figure 2. MPLS Label Swapping from Ingress to Egress



If layer 2 is ATM, the label is inserted into the VPI/VCI field of the ATM cell header. Similarly, if layer 2 is frame relay, the label is inserted into the data link connection identifier (DLCI) field in the frame header. If Ethernet or point-to-point protocol (PPP) is running in layer 2, a shim header is inserted between the layer 3 header and the layer 2 header. The shim header contains the MPLS label.

The MPLS standard allows for MPLS-enabled networks to be nested within each other. To accommodate this nesting, packets may have multiple labels. The label in context appears earliest in the packet (at the top of the stack) and the label at the bottom of the stack appears at the end. A label stack consists of label stack entries (i.e., the MPLS field). Figure 3 shows the MPLS label stack entry format.

Figure 3. MPLS Label Stack Entry



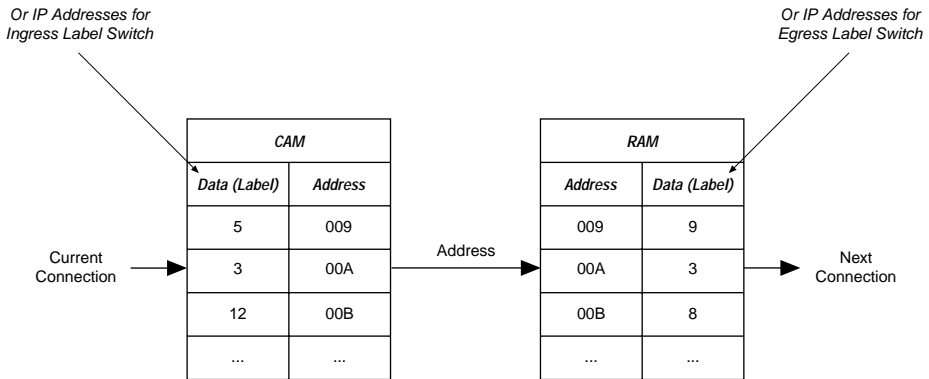
Notes:

- (1) Label: label value, 20 bits.
- (2) Exp: experimental use, 3 bits.
- (3) S: bottom of stack, 1 bit.
- (4) TTL: time to live, 8 bits.

Implementing Label Swapping in APEX CAM

Figure 4 shows how to implement a simple label swapping function using Altera’s APEX™ device family. By using Altera PLDs, designers can take advantage of the high performance and availability of predefined MegaCore® functions (library modules).

Figure 4. Using CAM & RAM for Label Swapping in APEX Devices



The number of labels that need to be stored in a LSR depends on the type of label mapping policy that is used in the MPLS network. For example, an LSR can hold approximately 1,000 labels at any given time. Using the CAM-RAM technique described above, 1,024 entries in the label swapping function would occupy 48 embedded system blocks (ESBs).



For more information regarding APEX CAM, see *Application Note 119 (Implementing High-Speed Search Applications with APEX CAM)*.

True-LVDS Solution for MPLS Data Plane Functions

To support core network traffic, MPLS data plane functions need to run at very high speeds. A fast processing function and a fast I/O interface are required. Altera devices provide many different high-speed I/O interfaces, including low-voltage differential signaling (LVDS). Altera devices are designed with the True-LVDS™ solution, which offers high data transfer rates with low power consumption. APEX 20KE devices are designed with dedicated LVDS circuitry that support transfer rates of up to 840 Mbps per channel. A total of 16 transmitter and 16 receiver channels provide up to 27 Gbps of bandwidth in a single device. With high-speed data transfer rates offered by the LVDS solution, APEX 20KE devices can operate at OC-192 speeds. The True-LVDS solution is offered across several Altera device families, including Excalibur embedded processor solution devices.



For more information regarding LVDS in Altera devices, see the following:

- [*Application Note 117 \(Using Selectable I/O Standards in Altera Devices\)*](#)
- [*Using I/O Standards in the Quartus Software White Paper*](#)
- [*Using LVDS in APEX 20KE Devices White Paper*](#)
- [*Using LVDS in the Quartus Software White Paper*](#)

Quartus II Development Tool

Users can easily implement CAM and LVDS in their designs by using the Quartus II development tool. The Quartus II development tool guides the user through the implementation steps with a comprehensive help interface. The MegaWizard® Plug-In Manager helps the user configure MegaCore functions (such as CAM, RAM, and LVDS) to user specific needs with a user friendly interface. Since the Quartus II development tool can be used with various third-party EDA tools, the Altera solution offers the best platform for developing MPLS functions.

Label Distribution

A standard label distribution method is required for an LSR to distribute labels to neighboring LSRs. When an LSR assigns a label to a particular FEC, this information needs to be conveyed to its peers in the MPLS network.

A label switched path (LSP) is defined by a set of labels that are used from the ingress of the MPLS domain to the egress. Since labels map the LSP, label distribution helps by using a set of procedures to distribute labels among LSR peers.

The MPLS standard does not dictate which signaling protocol should be used for label distribution. Label distribution protocol (LDP) is the most popular protocol for label distribution; however, other signaling protocols do exist. In addition, extensions to these signaling protocols have been created to support traffic engineering for MPLS. The IETF has not decided on a single signaling protocol extension for MPLS traffic engineering. Constraint-based routing-label distribution protocol (CR-LDP) and resource reservation protocol-traffic engineering (RSVP-TE) are the two signal protocol extensions being considered for MPLS traffic engineering. MPLS-enabled product manufacturers will very likely support both protocols.

LDP is a protocol that has a set of messages defined for signaling. Four classes exist:

- DISCOVERY messages are used for finding LSRs. The discovery protocol of LDP runs on top of user datagram protocol (UDP).
- ADJACENCY messages initialize, maintain, and shut down sessions between LSRs.
- LABEL ADVERTISEMENT messages distribute label information.
- NOTIFICATION messages are used for advisory and error signaling.

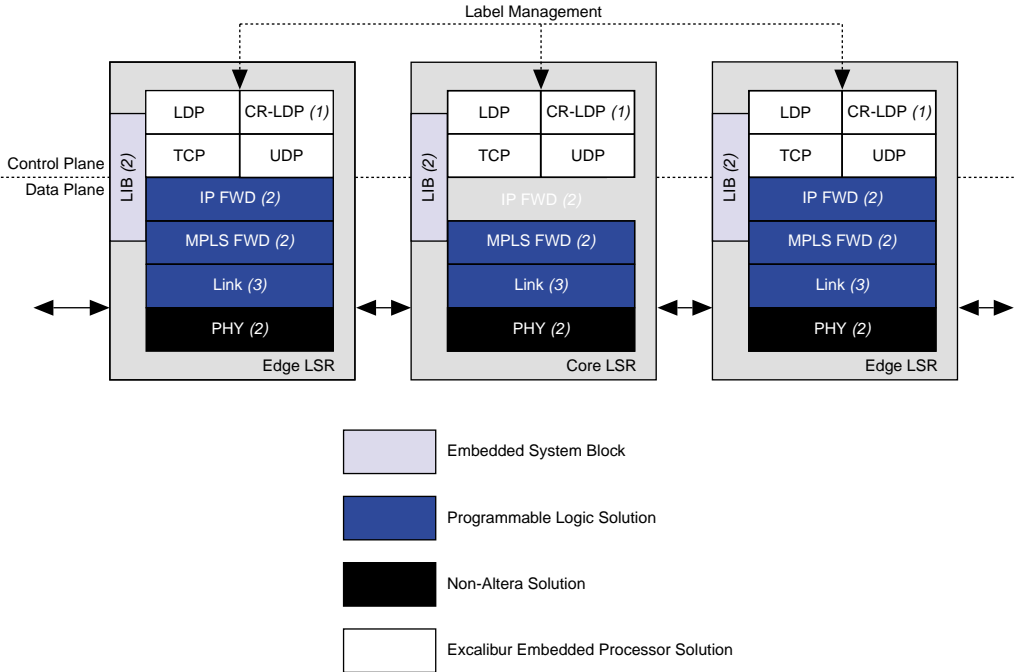
With the exception of DISCOVERY messages, LDP runs on transmission control protocol (TCP). Because of the critical nature of the information being transferred, TCP ensures reliable data transport between LSRs.

Altera PLD Solution

To implement LDP with TCP and UDP support, a software solution is required. Altera's Excalibur embedded processor solutions offer an optimal platform that can implement these functions. Altera offers two hard core solutions, a 32-bit ARM or MIPS processor that is embedded in a device with Altera programmable logic. The hard core solutions offer a direct interface to the programmable logic on the device. Altera also offers a configurable soft core RISC processor, the Nios™ embedded processor. The Nios embedded processor can be instantiated many times within a PLD, enabling multiprocessor capability on a single device.

Figure 5 shows a diagram of an MPLS protocol stack. The diagram is separated into two sections, the control plane and the data plane. The control plane can be implemented on an embedded processor while the data plane can be implemented in programmable logic.

Figure 5. MPLS Protocol Stack



Notes:

- (1) CR-LDP is constraint based LDP that is used for traffic engineering. RSVP-TE is another signaling mechanism that is used for traffic engineering.
- (2) PHY: physical layer; LIB: label information base (table of labels mapping input port/label to output port/label); IP FWD: next hop forwarding based on IP address (longest match forwarding used); MPLS FWD: label switching based on MPLS label and LIB lookup.
- (3) Link is the data link layer.

Traffic Engineering

One of the most important advantages of MPLS is that it allows traffic engineering in packet networks. Traditionally, traffic engineering was native to ATM.

Traffic engineering allows traffic to be efficiently routed, which results in better utilization of the network. This efficient routing prevents traffic congestion in certain paths. When forwarding packets, the shortest path between nodes is not always selected. For example, two packets with the same source and destination could travel different routes to reach the same destination. This efficient use of the network occurs when certain segments in the network are less utilized than others and those segments can be used to establish a different LSP.

In MPLS, different LSPs can be created and are based on a network administrator's requirements. The two most widely used methods for providing traffic engineering in MPLS networks are RSVP-TE and CR-LDP. At this time, the IETF's MPLS working group has not standardized the signaling protocol method.

Constraint-Based Routing

In normal IP routing, the data path is calculated from some measurement of efficiency (e.g., the number of hops). Constraint-based routing (CR) calculates the data path using this method, but also ensures that none of the constraints that have been set are violated along the path. CR is a method used in traffic engineering for routing that can consider information like link characteristics (delay and bandwidth), hop count, and QoS. Other constraints available include which segments of a network an LSP should or should not use.

RSVP-TE and CR-LDP are two competing protocols used for MPLS that perform CR. RSVP is an existing protocol, standardized by the IETF, that has been extended to RSVP-TE. Similarly, CR-LDP is an extension of LDP, which has been designed for MPLS. There are advantages and disadvantages to both protocols.

- CR-LDP sits on top of TCP, which ensures reliability. However, for RSVP to ensure reliability, refreshing must occur in the steady state. Refreshing will consume bandwidth and processing resources.
- TCP requires some handshaking before an LDP session can begin and results in a moderate amount of overhead. RSVP does not require connection establishment before label distribution occurs.

Because of the advantages and disadvantages of RSVP-TE and CR-LDP, designers need to keep their systems flexible enough to accommodate future changes to the protocols. Altera's Excalibur embedded processor solutions can implement RSVP-TE and CR-LDP. In addition, the Excalibur embedded processor solutions can perform remote field upgrades, which allows the system to remain current. Excalibur embedded processor solutions provide a complete solution for traffic engineering functions.

Virtual Private Networks

A virtual private network (VPN) allows the public Internet to be used as a method for connecting various networks to form a private WAN. The VPN service provider must provide data privacy and support private IP addressing use where IP address space overlaps between other network domains. This means that IP addresses used within the private network can also be used outside of that network (i.e., in the public Internet) without interference.

MPLS provides a simple solution to VPN-related issues. Because forwarding decisions are based on the MPLS label and not the destination IP address, this MPLS feature solves many VPN-related issues.

Currently, there is no standardized method of using MPLS for VPNs. As a result, network equipment manufacturers are implementing proprietary functions that only work with their own products. In the future, network equipment manufacturers must provide inter-vendor compatibility.

Similar to VPNs, other proprietary value-added services are being developed by network equipment manufacturers to offer customers a complete MPLS solution. If there is a particular proprietary feature that a customer does not like, they are either forced to continue using the system or purchase a replacement system. An MPLS implementation must be flexible because proprietary value-added services may force a customer to use an unwanted value-added service. PLDs are the only solution that allow field upgrades to hardware (programmable logic) and software over a network. In addition, field upgrades will be useful when MPLS standards are finalized or when new standards are created.

Summary of PLD Implementation

There are several key advantages to using Altera solutions for implementing MPLS.

Flexibility

MPLS is still being standardized and inter-vendor compatibility needs to be supported. Therefore, network equipment manufacturers must keep their products flexible enough to accommodate changes. When the IETF standardizes MPLS, most network equipment manufacturers will support the IETF standard.

For MPLS system designers, products are distinguished by value-added services. The system platform must be flexible when implementing new value-added functions because the functions must address an ISP's long-term requirements. Excalibur embedded processor solutions offer a combination of programmable logic and embedded processors that provide the most flexibility for implementing MPLS systems. Altera devices are capable of remote field upgrades, a key feature that is only possible with PLDs. Upgrades to ASIC solutions are nearly impossible due to their high, non-recurring engineering (NRE) costs and long turnaround times.

Performance

Wire-speed performance is required for devices in the data path. Altera's high-performance programmable logic is capable of supporting MPLS data path functions. The Excalibur embedded processor solutions provide an array of options to meet performance for control plane functions.

Time-to-Market

To effectively establish a foothold in the market, a short design cycle is required. Implementing an MPLS system with Altera PLDs shortens time-to-market because they eliminate any turnaround time that is normally associated with ASICs.

Hardware/Software implementation

To effectively implement an MPLS system, designers will need to implement a combination of logic and processor-based functions. In general, label distribution functions will need to be implemented in software and data path functions need to be implemented in hardware. Having separate chips implement these functions can cause problems, such as performance degradation, less board space, and increase in design time. On the other hand, a single-chip ASIC solution will render the implementation inflexible, increase time-to-market, and will quickly become obsolete. Altera's Excalibur embedded processor solutions offer an elegant solution to these problems as they integrate programmable logic and processors onto a single device. The configurable Nios RISC processor can be instantiated multiple times into programmable logic, providing a multiprocessor platform in a single device. The ARM and MIPS processors are hard core processors that are embedded in a device surrounded by programmable logic.



For more information regarding Altera's Excalibur embedded processor solutions, see the following data sheets:

- [*ARM-Based Embedded Processor Device Overview Data Sheet*](#)
- [*MIPS-Based Embedded Processor Device Overview Data Sheet*](#)
- [*Nios Soft Core Embedded Processor Data Sheet*](#)

Conclusion

When designing MPLS products, value-added services and competing protocol standards have forced network equipment manufacturers to stay flexible without compromising performance. Also, network equipment manufacturers must consider time-to-market and the platform on which MPLS will be implemented. The Altera portfolio of IP functions, advanced device features, and Excalibur embedded processor solutions, address all of these issues when designing an MPLS system.

References

Callon, R., Rosen, E., and Viswanathan, A. "Multiprotocol Label Switching Architecture," work in progress. Available from (<http://www.ietf.org/html.charters/mpls-charter.html>). July 2000.

Conta, Farinacci, Fedorkow, Li, Rekhter, Rosen, and Tappan. "MPLS Label Stack Encoding," work in progress. Available from (<http://www.ietf.org/html.charters/mpls-charter.html>). July 2000.

Davie, Bruce, and Rekhter, Yakov. *MPLS: Technology and Applications*. 2000.

Leon-Garcia, Alberto, and Widjaja, Indra. *Communication Networks: Fundamental Concepts and Key Architectures*. 2000.

MPLS Resource Center-Resources for Multiprotocol Label Switching. Available from (<http://www.mplssrc.com>). October 2000.

Web ProForum Tutorial. Available from (<http://www.iec.org/tutorials/mpls/index.html>). July 10, 2000.



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>
Applications Hotline:
(800) 800-EPLD
Customer Marketing:
(408) 544-7104
Literature Services:
lit_req@altera.com

Altera, APEX, Excalibur, MegaCore, MegaWizard, Nios, Quartus II, True-LVDS, and specific device designations are trademarks and/or service marks of Altera Corporation in the United States and other countries. Altera acknowledges the trademarks of other organizations for their respective products or services mentioned in this document. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

Copyright © 2001 Altera Corporation. All rights reserved.



I.S. EN ISO 9001