Chorus - ENT3003

Wire-Speed Packet Forwarding ASIC for Access Aggregation Routers

Product Brief, version 1.4

Description

Entridia's Chorus™ is a fully integrated wire-speed Internet Protocol (IP) packet forwarding ASIC targeted at IP-based aggregation systems in Service Provider networks. Chorus is part of Entridia's OPERA™ *(Optical Edge Routing Architecture)* family of products targeted at high performance MAN/LAN routers. Chorus incorporates four 10/100 Ethernet ports and an OC-3c Packet over Sonet (PoS) port. It also offers connectivity to legacy WAN networks (T1/E1, T3/E3, ISDN and xDSL) via the Control Plane Processor interface and offers glueless expansion capability via Entridia's OptiStream™ expansion bus. With deterministic ingress-to-egress latency of less than 8µs, Chorus is ideally suited for converged Service Provider systems.

The Chorus device is configured and controlled by Entridia's ENTROS™ software layer which interfaces via any real time operating system (RTOS) to all standard routing protocols including RIP, OSPF, and BGP.

The Chorus device can be enhanced to support 128K routes and flows at full line rate by using Entridia's Crescendo Forwarding Table Controller (ENT2003) to interface to external CAMs and SRAMs. Alternatively, third party lookup engines or algorithmic processors may also be used with the Chorus device.

Features

- 1.2 Gbps of network bandwidth
- 8 µs ingress to egress latency
- One OCPort™ PoS-PHY interface for 155 Mbps OC-3c Packet over Sonet links with integrated PPP framing, 8 priority levels, and support for extended scheduling capabilities
- Four 10/100 Ethernet ports with 2 levels of priority
- Support for tagged IEEE 802.1Q VLAN mode packets for Ethernet ports
- Control Plane Processor interface allows easy integration with an external microprocessor
- 6.4 Gbps OptiStream bus allows glueless cascadability to realize an OC-12 capable system
- 256 entry on-chip Flow Table with multi-field packet analysis, wire-speed firewall functionality
- 512 entry on-chip Route Table with full CIDR functionality
- 128K routes and flows *(Unified Mode)*, 64K routes and 64K flows *(Discrete Mode)* at full line-rate using Entridia's external Crescendo device
- On-Chip support for DiffServ classification
- Arbitrary tag-based packet classification
- 64-bit, 100 MHz SRAM interface for packet data and 16-bit, 100 MHz SRAM interface for status
- 520 HPBGA, 0.25 um 2.5V CMOS process

Block Diagram

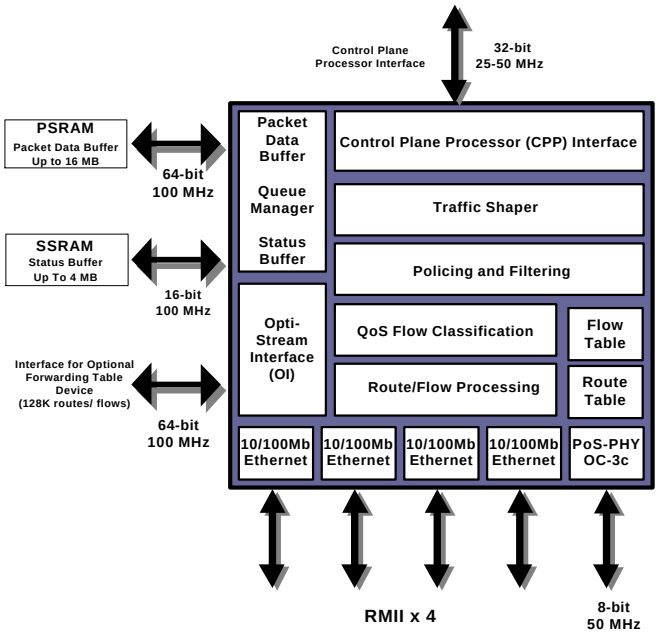

Entridia Corporation
131 Theory, Suite 150
Irvine, CA 92612
USA

Phone: (949) 823-3600
Fax: (949) 752-1462
E-mail: info@entridia.com
http://www.entridia.com

Features and Architecture

Data Link Layer (OSI Layer 2) Frame Processing

The Chorus device's data link layer frame processing includes all the traditional Media Access Controller (MAC) functionality for 10/100 Mbps Ethernet media and an OC-3c PoS physical layer. The functionality of the Ethernet Media Access Interfaces includes:

Layer 2 Multicast: The transmit side derives the appropriate multicast MAC address for a Class D IP-Multicast address.

IETF RFC826 Ethernet Address Resolution Protocol (ARP): To associate MAC addresses with Network Layer IP addresses, the Chorus chip maintains a memory table for each Ethernet port that maps MAC addresses to IP addresses.

IEEE 802.1Q VLAN: The Virtual Local Area Networks (VLAN) tagging associates physical devices and ports to the defined VLAN, then maps that association with other LAN stations via the usage of a VLAN tag.

LLC 802.2 SNAP: The Subnetwork Attachment Point (SNAP) layer resides between the network layer and the MAC layer. The main usage for SNAP/LLC tagging is for the AppleTalk protocol.

IP to MAC address translation table: This table supports both static and dynamic entries and is maintained entirely by the Layer 2 Media Access Interface of the Chorus chip. The device manages entries in this table by examining the ARP (IETF RFC826) packets that arrive at each of the media access interfaces. Optionally, configuration software running on a companion microprocessor can augment or override this behavior and add static entries to this table.

Promiscuous mode: This register-programmable mode is used to examine all incoming frames irrespective of the destination MAC address.

Bridge mode: This register-programmable mode is used to pass the frame data to a companion microprocessor.

Glueless Reduced MII (RMII) interface to the physical layer: The RMII can support 10 Mb/s and 100 Mb/s data rates, provides independent 2-bit wide (di-bit) transmit and receive data paths, and uses TTL-compatible (3.3V CMOS) signal levels.

The functionality of the PoS-PHY interface includes:

8-bit, PoS-PHY Level 2 interface to the physical layer: Chorus has an OC-3c line-rate port that provides connection to the network. Through the port interface Chorus transmits and receives PPP encapsulated IP packets with optional 32-bit tags attached. An external 8-bit, PoS-PHY Level 2 device is required to add the appropriate framing/deframing to PPP packets transmitted from Chorus. The maximum packet size for the OC-3c port is 1,536 bytes.

The functionality of both the Ethernet Media Access interfaces and the PoS-PHY interface includes:

Management Information Base (MIB) registers and counters: These registers and counters are maintained on the chip and trigger interrupts on overflows that can be monitored and collected by management software running on a Control Plane Processor.

Network Layer Packet Processing

Field Extraction

The extractors support pre-configured offset, nibble sequences for the following protocols:

- IPv4/L4 TCP/UDP fields
- IPv6

Classify, Match, Route, or Filter based on:

- IP Source/ Destination Address
- TCP / UDP / ICMP – Protocol Fields
- TCP / UDP Source / Destination Ports
- TCP Control bits
- IP DiffServ Code Point (DSCP)

Flow Table

The Chorus device has a built-in 256 entry flow table. The flow table maintains filtering rules to determine which types of IP packets a network interface can forward or discard, and which types require further examination by software running on a Control Plane Processor. Refer to Table 1.

The on-chip flow table can be replaced to support 128K routes/flows at full line rate by using Entridia's Crescendo Forwarding Table Controller (ENT2003) device to interface to external CAMs via the OptiStream expansion bus. Alternatively, third party lookup engines

or algorithmic processors may also be used.

Each flow table entry can specify a source and destination IP subnet, source and destination port (TCP, UDP or other Layer 4 port) along with an action to take if any or all of these fields match those of the packet being filtered. The action can specify one of the following:

- **Forward**: Accept the packet, and forward it normally.
- **Priority Forward**: Accept the packet, and forward it with priority over normal packets.
- **Drop**: Discard the packet.
- **Examine**: Accept the packet and notify the Control Plane Processor of the match and the starting address of the packet data in the Packet Data Buffer memory. This enables additional filtering software running on the Control Plane Processor to further process the packet.

Route Table

The Chorus chip has a built-in 512 entry route table with full support for longest prefix match look up, essential for Classless Interdomain Routing (CIDR). As route lookup can be performed in three cycles using this approach, Chorus can guarantee wire-speed processing with deterministic latency. The Chorus chip route table also supports using the IP DiffServ Code Point (DSCP) field in making Quality of Service (QoS) policy-based destination route selection. Refer to Table 2.

The on-chip 512 entry CIDR route table can be replaced to support 128K routes/flows at full line rate by using Entridia's Crescendo Forwarding Table Controller (ENT2003) device to interface to external CAMs via the OptiStream expansion bus. Alternatively, third party lookup engines or algorithmic processors may also be used.

Traffic Shaper

The OC-3c port has 8 levels of priority and the Ethernet ports have 2 levels of priority. Additional scheduling can be added via the Entridia OptiStream bus. The priority output queues can be selected based on the DSCP, source or destination IP addresses, and TCP, UDP or other Layer 4 ports, or a valid combination of the above parameters, making it highly flexible for system designers to implement policy-based wire-speed routing rules for the OC-3c port. Chorus has a built-in weighted round-robin scheduler with user programmable weights.

OptiStream Interface

The OptiStream feature of the Chorus chip is an enhanced glueless GTL+ (Gunning Transceiver Logic) interface which can operate at frequencies up to 100 MHz. The OptiStream expansion bus has a 64-bit wide data path that supports a sustained throughput of 6.4 Gbps. The OptiStream bus allows system designers to cascade up to four Chorus devices, forming a 16-port 10/100 Mbps, 4-port OC-3c PoS, wire-speed IP router and the attachment of an external forwarding table device such as the Entridia ENT2003 Crescendo device.

Control Plane Processor Interface for Exception Processing

The Control Plane Processor interface couples an external microprocessor (Control Plane Processor), running at 25 to 50 MHz, with internal Chorus logic and control registers. It consists of a 32-bit data bus and a 16-bit address bus.

This bus presents Chorus as a memory-like device to the Control Plane Processor. An extensive register map enables software running on the Control Plane Processor to initialize and dynamically re-configure the Chorus chip.

The CPP interface also includes an interrupt line, which is used to notify the Control Plane Processor of exception conditions and periodic events. Notification of specific conditions can be masked or enabled by programming appropriate interrupt control registers.

Configuration software running on the Control Plane Processor can install static routes (including a default gateway entry) and periodically age and rebuild other entries dynamically as a result of routing protocol messages (such as OSPF and BGP4). Also, the Control Plane Processor can maintain very large route tables in system memory and use the on-chip 512 entry table as a cache for frequently used routes.

Software Interface

Entridia provides complete software drivers and a Service Application Programming Interface (ServAPI) to enable system designers to quickly migrate their existing software to the Chorus platform.

The software drivers are targeted to the Wind River Systems' VxWorks real-time operating system (RTOS). However, the drivers use a hardware abstraction model that makes it very easy to port to other RTOSs.

To enable quick integration with existing software in routers, the Entridia driver presents the Chorus network interface as five independent network devices. The router software that deals with Network Layer communications can treat each of the interfaces as a dedicated Data Link Layer device and interoperate with them

Entridia Corporation
131 Theory, Suite 150
Irvine, CA 92612
USA

Phone: (949) 823-3600
Fax: (949) 752-1462
E-mail: info@entridia.com
http://www.entridia.com

without modification. In addition, the ServAPI is a library of C functions that simplify access to the more advanced features of the device such as filtering rules and traffic shaping parameters.

System Integration

The Chorus chip is designed to integrate into Access/ Service Provider edge routing systems without significant modification to the software currently running on those systems, while accelerating their performance and increasing functionality.

Chorus handles the following functions, which are currently handled by a microprocessor/discrete programmable logic, without any software intervention:

- **Packet buffering/Memory management**: The Chorus chip transfers all packet data into and out of the packet data buffer memory.
- **ARP table management**: The Chorus chip learns the IP to MAC address mapping, and ages and refreshes the entries at periodic intervals specified in programmable registers.
- **ARP protocol handling**: A register-programmable Proxy ARP mode enables the Chorus chip to respond automatically to ARP requests with appropriate ARP response packets to implement transparent subnet gateways.
- **Route look up**: The Chorus chip searches for the longest prefix match for a destination IP address in the route table.
- **Packet filtering rules**: The Chorus chip forwards or discards packets based on rules in the flow table.
- **Packet Data Buffer management**: The Chorus chip allocates and reclaims space for the packets in the packet data buffer memory.
- **IP Multicast**: The Chorus chip manages transmission of multicast packets to all the recipient ports without unnecessary data copying.
- **IP-based QoS**: Priority Scheduling; weighted round-robin.

Memory Requirements

The Chorus device is designed to work with pipelined external ZBT SRAM for the Packet Data Buffer and the Status Buffer. Chorus supports a maximum addressable SRAM capacity of 16 MB Packet Data and 4 MB for Status Buffer.

Packaging

Body Size: 40 X 40 mm
Ball pitch: 1.27 mm
Ball count: 520-pin HPBGA package *(see diagram)*

Power Requirements

Typical Power Dissipation: 7W
0.25 um 2.5V CMOS process, 3.3V I/O

123-bit Classification Key							
DestIP 32-bit	SrcIP 32-bit	Dest Port Number 16-bit	Src Port Number 16-bit	DiffServ Code Points (DSCP) 8-bit	Protocol 8-bits	Physical Interface ID 7-bit	Control Bits 4-bit

40-bit Response Data		
Priority 6-bit	Egress Tag 32-bit	Action 2-bit

Table 1. Discrete Mode Flow Table Entry Detail and Layer 3/4 Parameters

49-bit Classification Key			
DestIP 32-bit	DiffServ Code Points (DSCP) 8-bit	Physical Interface ID 7-bit	Control Bits 2-bit

72-bit Response Data		
Next Hop IP 32-bit	Network ID 8-bit	Egress Tag 32-bit

Table 2. Discrete Mode Route Table Entry Detail

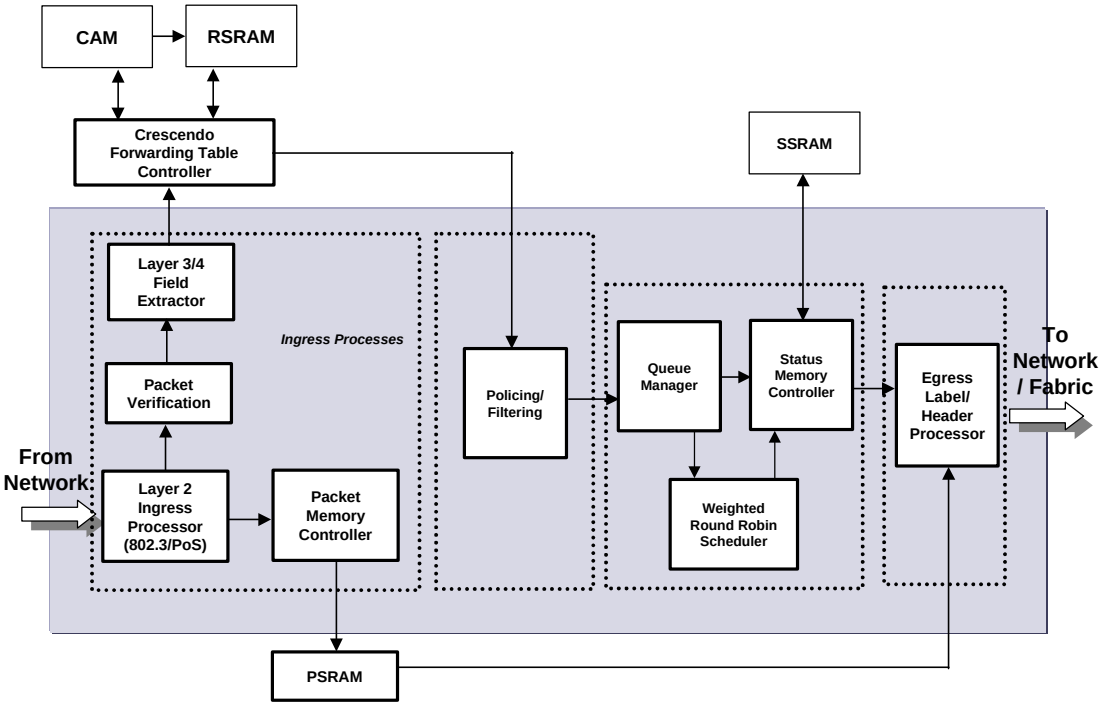

Figure 1. Chorus Packet Flow and Processing Operations

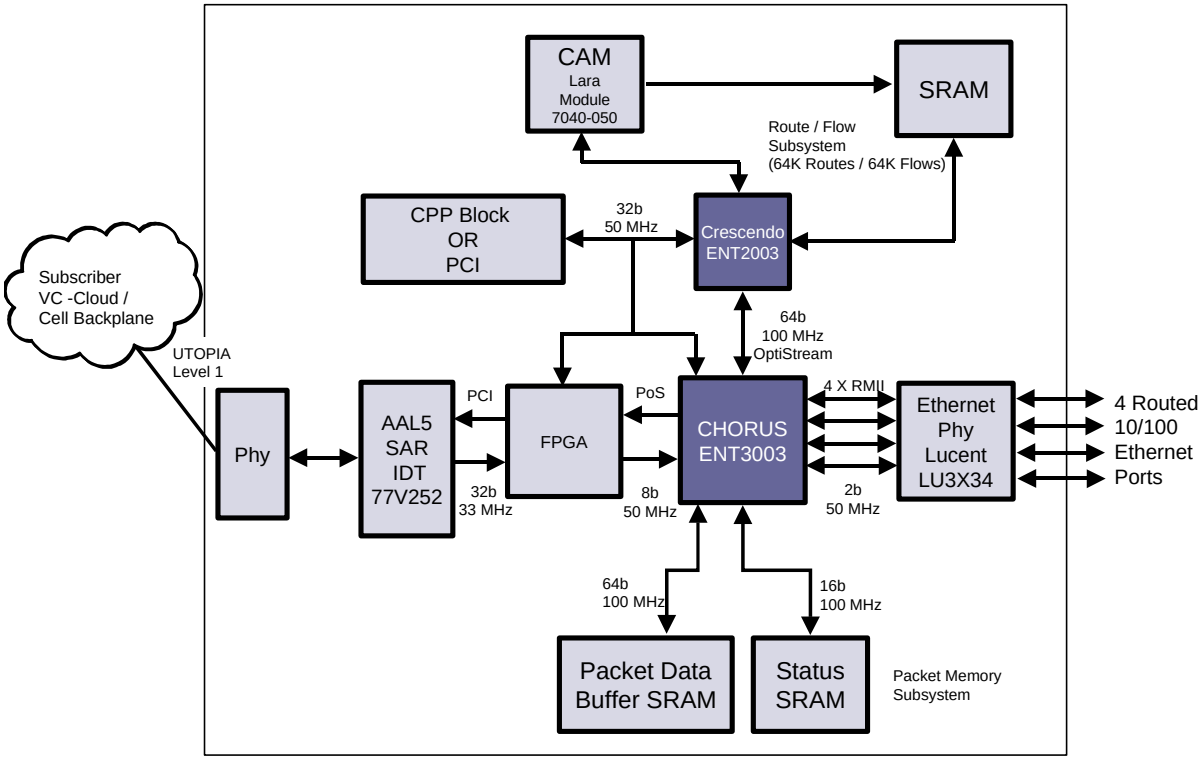

Figure 2. ATM/IP Wire Speed Multi-Service Routing Module

Entridia Corporation
131 Theory, Suite 150
Irvine, CA 92612
USA

Phone: (949) 823-3600
Fax: (949) 752-1462
E-mail: info@entridia.com
http://www.entridia.com

Typical Applications and Packaging Dimensions

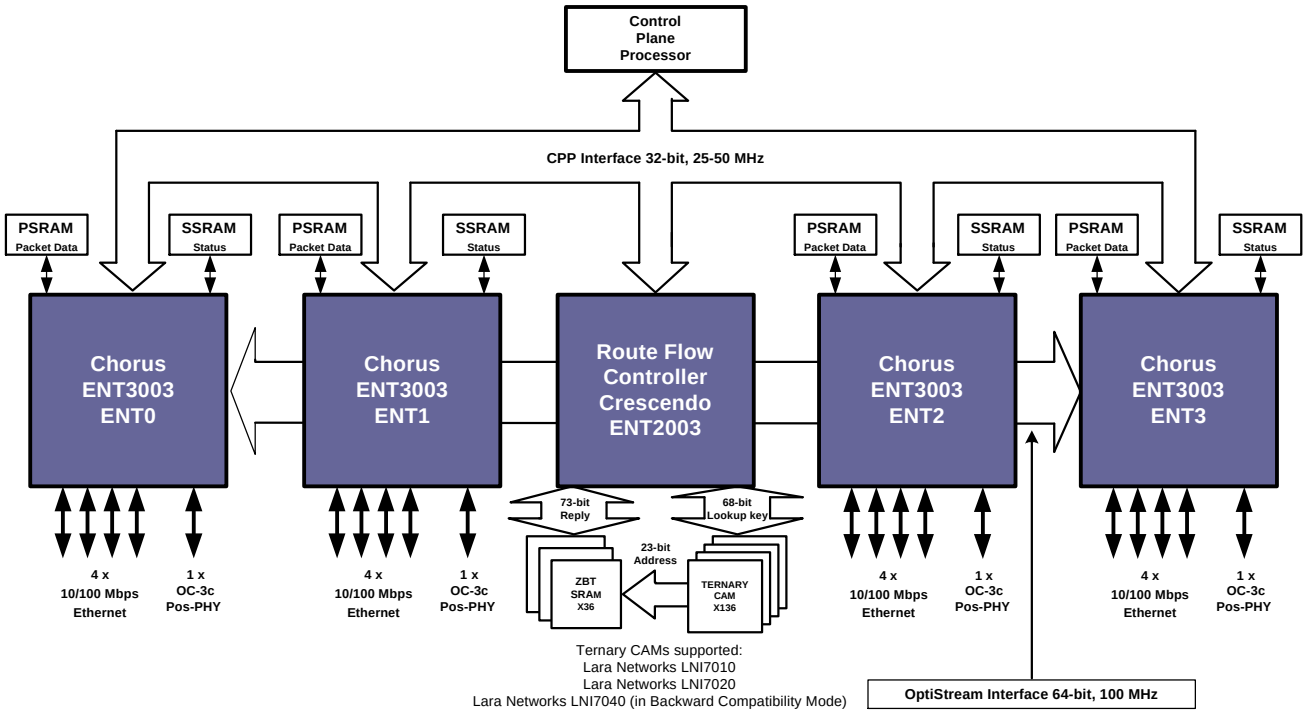

Figure 3. OC-12 and Ethernet Aggregation Router

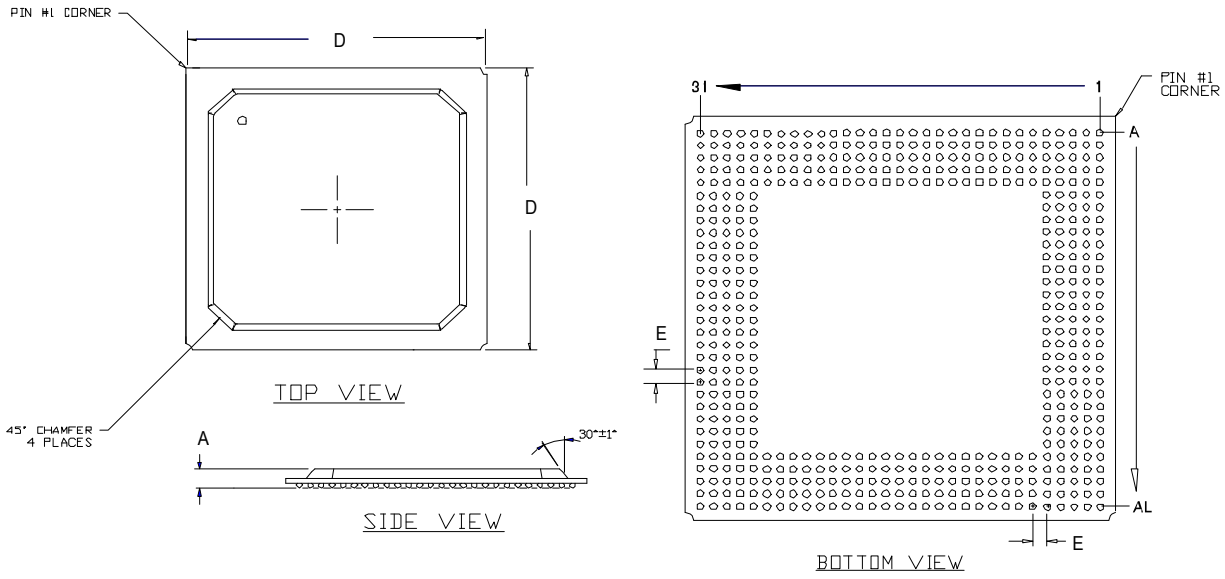

Body Size 'D' (mm)	Max Mounted Pkg Height 'A' (mm)	Ball Pitch 'E' (mm)	Pin Count / Packaging
40 x 40	1.4	1.27	520 HPBGA

Figure 4. Chorus Packaging Dimensions