SYNOPSYS®
Accelerating Innovation

DesignWare IP for Data Centers

What Happens in an Internet Minute?

Today's residential, business and mobile networking trends are driven largely by a combination of video, social networking and advanced collaboration applications, termed "visual networking". This is where social networking annotations are layered over broadband video to create a highly interactive and immersive experience for users. These types of applications are driving significant traffic in data centers. For example, users are uploading approximately 350 million pictures onto Facebook per day — that is a lot of data and storage. In addition, due to the tremendous bandwidth increases, new applications and business models such as enterprise resource planning and mobile device management are moving to the cloud. This puts significant demands on the data center to improve efficiency and reduce space, cost and power.

Data Center Responds in Full Force

The data center industry is responding to this growth in three ways: (1) lowering power through the use of micro servers; (2) simplifying the data center network through software defined networks and rack-level converged systems as well as disaggregated racks; and (3) improving performance with cache acceleration driven by PCI Express® (PCIe®) solid state drives (SSDs).

Lowering Power

Data centers worldwide use about 30 billion watts of electricity, roughly equivalent to the output of 30 nuclear power plants. A single data center can draw as much power as a medium-sized town. The industry's answer to this challenge is a new server-on a-chip for micro servers or low-power servers (see figure 1). These chips use low power processors integrating common functions such as Ethernet, SATA, PCI Express, memory controllers and fabric switches, and memory caches.

Micro servers have the ability to significantly reduce power through integration of server motherboard functions particularly for common tasks such as web servers, search engines, on-line transactions, and video-on-demand. Innovation in low-power servers will be one of the drivers of growth in the overall IP market over the next few years.

Micro server system-on-chip (SoC) designs involve challenges such as determining trade-offs between low power vs. high performance operation; need for cache coherency; integrating high bandwidth interconnects that are scalable to many cores; implementing heterogeneous processing hardware accelerators; addressing Reliability-Availability-Serviceability (RAS); and incorporating technology that is portable to future design processes. Network access to a high performance fabric such as Ethernet along with I/O expansion using PCIe root ports and devices supporting single root I/O virtualization (SR-IOV) is critical for data center server virtualization.

Synopsys provides a range of IP that helps address the design requirements of micro server system-on-chips including:

- Lowest latency multi-port memory controllers and PHYs optimized to share main memory with compute offload engines plus network and storage I/O resources
- Low latency Ethernet digital controllers and PHYs enabling integrated network switches
- Complete PCI Express endpoint optimized for AXI ordering rules
- Multiprotocol Enterprise 12G PHY IP with support for PCI Express, SATA, and 10 and 40 Gigabit Ethernet connectivity, which are available in advanced design process nodes
- Low latency embedded memories right-sized for L1 and L2 cache with multi-bit error correction for RAS capabilities

Simplifying the Data Center Network

Traditional data centers use a tree-based network topology consisting of switched Ethernet with VLAN tagging. This topology only defines one path to the network, which has traditionally handled north-south data traffic. The transition to a flat leaf-spine data center network using 10G and 40G Ethernet links enables virtualized servers to distribute workflows among many virtual machines. A software defined network (SDN) where the control is decoupled from the data path allows a common software stack such as OpenFlow to provide a consistent industry-wide control software environment. Instead of having a proprietary software stack, designers have an OpenFlow managed network, allowing users to provision the networks very easily (virtually), without requiring physical access to the network's hardware devices. According to IDC, the SDN ecosystem and network infrastructure will grow at an exponential pace reaching marketing size of $3.7 billion by 2016.

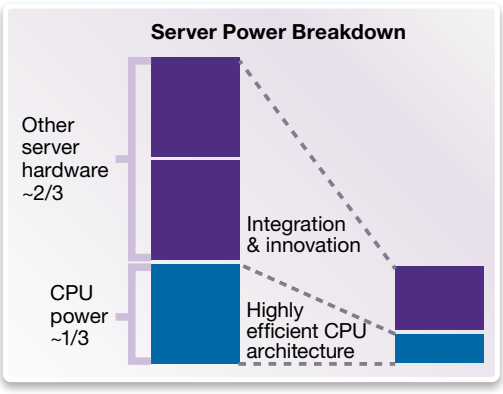

Figure 1: Reducing power with micro servers

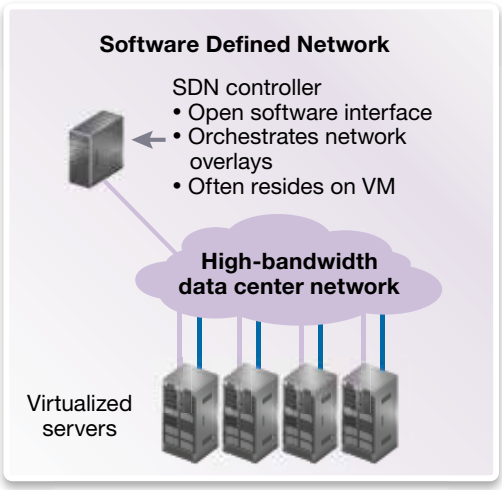

Figure 2: Simplifying the data center network with software defined network

"Moreover, software defined networking (e.g., OpenFlow) could demand additional search capability. As search-engine technology improves, merchant solutions become more attractive than proprietary ASICs, another factor in market growth."

The Linley Group

SDN SoC design challenges include maximizing performance, while providing OpenFlow acceleration using fast embedded memory for flow lookups and providing RAS features and technology portability to future design process. Implementing traffic management in hardware in order to support legacy network management as well as OpenFlow for SDN architectures are driving a new class of communications processors and network switch ASSPs.

Synopsys' DesignWare® Ethernet Digital Controllers and multiprotocol Enterprise 12G PHYs support key data center features such as Data Center Bridging (DCB) and Jumbo Packet to handle the VxLAN/NVGRE packet parsing, extra data packets required by OpenFlow and VxLAN implementations. PCI Express 4.0/3.1 digital controllers supporting up to 16 lanes deliver critical network functionality such as Process Address Space ID (PASID) for NIC interfaces and robust internal error reporting and extended advanced error reporting (AER). Synopsys' low latency embedded memories are ideal for OpenFlow tables in multi-port switches for flow management and high performance DesignWare DDR memory controller and PHYs support up to DDR4 2667 Mbps for highest performance.

Adding Value with Converged Systems and Disaggregated Racks

Converged systems combine compute, storage, network and management into a single solution. These pre-integrated, rack-level systems reduce the overall complexity, integration and cost of the system. Furthermore, having full control of the rack, gives designers the ability to re-architect at a rack level vs. system design using individual server or switch system components. Converged systems enable faster deployment, easier interoperability and consistent management. It also reduces training and support overhead as well as enables efficient system optimization. Similar to converged systems, disaggregated racks approach data center topology from a system rack level by designing systems at a rack level. New architectures and topologies can optimize the compute/CPU, network, memory and storage design compared to traditional rack architecture. "Disaggregated racks," which separates CPU, storage, power and networking resources allows individual components to be replaced as needed providing cost and upgrade flexibility for data center operators.

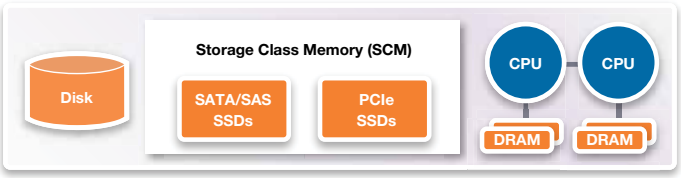

Figure 3: Reducing cost and improving performance for storage

Improving Performance through PCI Express SSDs

Server-based Solid-State Drives (SSDs) can utilize a PCI Express interface to directly connect to the server CPU and function as a cache accelerator allowing frequently accessed data, or "hot" data," to be cached extremely fast. High performance PCIe SSDs with extremely efficient input/output operation and low read latency improves server efficiency and avoids having to access the data through an external storage device. PCIe SSD server acceleration is ideal for high transaction applications such as web acceleration or database queries like SQL, Oracle or SharePoint. Flash-based PCI Express SSDs not only reduce power and cost but also minimizes area because it requires less rack space for the hard disk drives (HDDs). Server-based SSDs can also be SATA or SAS drive that can be plugged into the SATA/SAS interface on the system (see figure 3).

Accelerating Data Center SoC Designs with IP

Driven by the explosion in virtualization, next-generation data center systems require multi-terabit data processing and efficient, scalable compute servers. Furthermore, the continuing rise of internet traffic is resulting in the development of high-throughput data and telecommunications network systems using standardized APIs. Designers building integrated circuits (ICs) for these applications need a combination of high-performance and power-efficient IP solutions to help deliver total system throughput and meet the quality-of-service requirements.

IC architects and designers need optimized IP functions that are specifically designed for system requirements for data center networking, compute and storage applications. Typical IC system requirements are cache coherent interconnects, embedded protocol accelerators and standards based connectivity networking protocols. Synopsys' low latency, mission critical IP functions have been proven to work together in many designs using ARM® AMBA® 4 AXI™ and AMBA 3 AXI-based user interconnects. In addition, optimized DMA functions allow designs to be efficiently implemented.

"In Ethernet, the hot area is 10 Gigabit Ethernet, which is becoming the new standard for server connectivity. Unlike earlier ramps of Ethernet technology, 10GbE will see several types of PHYs in widespread use, not just twisted-pair cabling."

The Linley Group

Synopsys provides a comprehensive portfolio of high quality, silicon-proven IP that enables designers to develop SoCs for applications such as server blades, networking and storage. The DesignWare IP and system-level design solutions are optimized for high-performance, low power and low latency, and supports advanced process technologies from 28-nm to 16-nm/14-nm FinFET.

DesignWare IP	Product Features for Data Center SoCs	Data Center SoC Impact
DDR4/DDR3 Controller and PHY	Low latency, multi-port memory controller and PHY supporting DDR4/DDR3 SDRAM speeds up to 2667Mbps	Multi-port access to shared main memory enables protocol acceleration engines such as Hadoop acceleration for high-performance heterogeneous computing
PCI Express 4.0/3.1 Controller and PHY	Endpoint, root port, dual-mode and switch modes with SR-IOV, Process Address Space ID (PASID), AER, Active State Power Management (ASPM), support for L1 sub-states and hot-plug functionality	Enables IO virtualization for highly reliable server environments with support for lowest power operation. Hot plug eases maintenance and reduces cost
10G and 1G Ethernet Controller and PHY	Data Center Bridging (DCB), Energy Efficient Ethernet (EEE), TCP Segment Offload, Receive Filtering, double VLAN tagging, ARP Offload (IPV4) and VxLAN/NVGRE Packet Parsing	Ensures consistent configuration across the network. Reduces power consumption in the data center and accelerates CPU performance by offloading TCP protocol stack. Extends enterprise network by supporting overlay networks using VxLAN or NVGRE over Ethernet
40G Ethernet Controller and PHY	1G/10G/40G speeds, low latency, Energy Efficient Ethernet (EEE), with flexible filtering, Data Center Bridging (DCB), VLAN Tagging and Jumbo Frame support	Enables high bandwidth, flat leaf-spine network architecture using logical representations of data center network extending range of the physical network
USB 3.0 and USB 2.0 Controller and PHY	Low power for efficient energy consumption; small area for low silicon cost; USB 3.0 IP offers 5 Gbps data rates and support for SuperSpeed, High-Speed, Full-Speed and Low-Speed data rates	High speed USB enables local OS install/ re-install for server maintenance and debug
SATA 6G Host and Device Controller and PHY	AHCI programming model support multiple ports. Command and FIS-based port multiplier. Integrated DMA and advanced power management	SATA IP provides lower latency and higher bandwidth data transfers for local network attached storage (NAS) and enables low-power operation
Enterprise 12G PHYs	Multi-rate PHYs supporting PCIe 4.0/3.1, SATA 6G, and 10G and 40G port side and backplane interfaces including 40GBASE-KR4, 10GBASE-KR, 10GBASE-KX4, 1000BASE-KX, 40GBASE-CR4, 100GBASE-CR10, SGMII, QSGMII, XFI, and SFI (SFF-8431) protocols	Provides data center connectivity fabric for backplane and chip-to-chip applications
Embedded Memories	High Performance Core Design Kit optimized for all processor cores on SoC: CPU, GPU and DSP cores, STAR Memory system with EEC support for multi-bit error correction	Enables optimal implementation across all three dimensions: performance, power, and area. SEU mitigation enables highest reliability for improved RAS
ARM AMBA On-chip fabric	AMBA 4 AXI, AMBA 3 AXI and AMBA 2 interconnects	Optimized IP for AXI interface includes compliance to AMBA Ordering Rules, enabling high performance and cache coherency

DesignWare IP for Data Center

Summary

There are certain trends that are very clear in the data center market. Whether the goal is to reduce power, simplify the data center network or improve the performance for storage, Synopsys provides a broad portfolio of IP that is optimized for data center applications that will help you achieve your design goals faster and with significantly less risk.

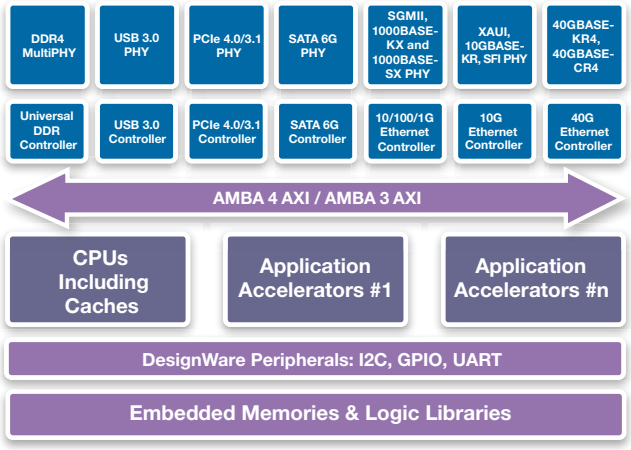

Figure 4: Typical Data Center SoC with DesignWare IP

SYNOPSYS®
Accelerating Innovation

Synopsys, Inc. • 700 East Middlefield Road • Mountain View, CA 94043 • www.synopsys.com

8/14.PS.CS4754