MemMax Scheduler

PRODUCT BRIEF

Increases Memory Usage Efficiency

- Leverages bank parallelism to hide page cycling overhead
- Groups transactions to avoid bus turnaround and rank change penalties
- Supports up to 7 address tiling schemes to optimize page and bank organization to dramatically improved system memory utilization
- Up to 85% efficiency possible

Superior Quality of Service

- Performs dynamic priority management to ensure service
- Bypasses latency-sensitive traffic mid-transaction – no Head of Line blocking
- Supports runtime configuration of Quality of Service
- Performs graceful trade-off between memory efficiency requirements and latency minimization of priority traffic

Improves Quality of Results

- Single-port, multi-threaded network interface minimizes routing congestion and improves timing convergence while maximizing throughput and protecting QoS
- Simple in-order memory controller interface reduces area and latency of controller
- Designed-in clock gating minimizes power

Shortens Time to Market

- DRAM protocol agnostic: supports DDR-1/2/3/4, LPDDR-1/2/3/4, HBM-1/2
- Integrated SRAM-compatible buffers efficiently isolate DRAM from network clocking while providing local rate decoupling
- Optimized connectivity to Sonics' on-chip networks

The MemMax® Scheduler is an intelligent DRAM scheduler designed for use with an Open Core Protocol (OCP) compliant memory controller. Ideal for high-bandwidth applications, MemMax Scheduler offers a sophisticated thread-based pipeline and advanced arbitration schemes in order to reduce interconnect over-design and redundancy. By decoupling the functionality of the System-on-Chip (SoC) from the DRAM, MemMax Scheduler encourages the adoption of DRAM technology that offers the best cost and performance value for each individual application. In addition, the MemMax Scheduler provides memory efficiencies beyond those traditionally achievable with simple scheduler or controller solutions alone. These increased efficiencies result in cost benefits for SoC integrators who can use less DRAM in their systems.

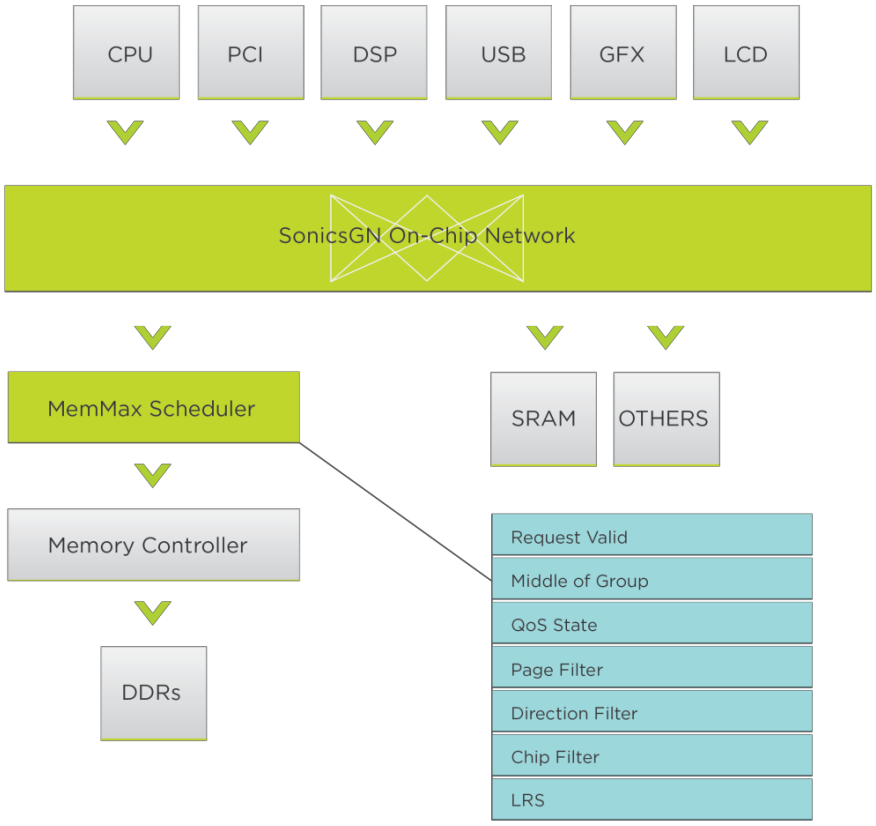

To reduce the total SoC die area and lower overall power consumption, designers can use compiled RAM to consolidate all of the flip-flop based buffers normally distributed among the various initiator cores into a single buffer within the MemMax Scheduler. When the DRAM is operated at an asynchronous frequency, this buffer also provides the crossing buffer, eliminating the need for dual buffers. MemMax Scheduler further reduces SoC costs by eliminating the excess wires required by traditional wire-intensive, multi-ported DRAM controllers and the on-chip fabrics that feed them.

While decreasing wiring area and increasing efficiencies, the MemMax Scheduler also provides a high level of Quality of Service (QoS) when faced with traffic contention. MemMax provides flexible QoS-based arbitration across the various initiator data flows mapped to each individual thread, allowing for fine-grain control over how bandwidth and latency guarantees are allocated to traffic classes. In combination with the end-to-end non-blocking nature of Sonics' on-chip networks, the QoS system ensures high throughput while reducing latency for critical traffic.

Feature Highlights

MemMax Scheduler improves overall SoC performance and reduces costs by intelligently managing access to off-chip DRAM using a variety of schemes for improving efficiency and QoS. MemMax Scheduler intelligently interleaves multiple system threads into the DRAM Controller in order to maximize DRAM utilization. In addition, high QoS levels can be achieved using MemMax Scheduler's ability to dynamically manage arbitration priorities among different threads to effectively support system data flow requirements.

MemMax Scheduler features:

Improved DRAM Efficiency

- Thread-based scheduling maximizes overall DRAM efficiency and provides levels of QoS for the traffic classes mapped to each thread
- Groups commands to maximize fast in-page operations
- Interleaves commands across banks to hide page cycling delays
- Internal bank timers delays prevent page misses from issuing to busy banks
- Optional auto-precharge at end of system burst on per-thread basis
- Groups reads and writes to minimize timing delays caused by switching data bus direction
- Groups commands based upon rank address to minimize delays caused by switching DRAM ranks
- Works well with a wide variety of controllers including: DDR-1/2/3/4, LPDDR-1/2/3/4, and HBM-1/2
- Offloads the scheduling logic from the controller
- Internal clock auto-gating can lower power consumption by several orders of magnitude

Quality of Service Modes

- Associates different service levels with each traffic class/thread (high priority/low latency, allocated bandwidth, and best effort)
- Provides flexible and programmable bandwidth rates and hysteresis per-thread
- Demotes traffic exceeding bandwidth rate to best effort to maximize throughput
- Flexibly groups commands to avoid thrashing

Transaction Processing

- Chops system bursts into short DRAM controller bursts (typically set to the DRAM burst length)
- Holds write requests until it accumulates the controller burst and guarantees acceptance of issued read bursts to avoid data storage in controller
- Optional address tiling provides up to 7 mappings of system address to DRAM rank, bank, row and column addresses to optimize throughput

Fully Configurable and Programmable

- Supports user choice of buffer sizes, modes of operation, DRAM technology, and QoS settings that best suit the application
- Integrated programming port for software tuning, scheduler re-configuration, and tiling optimization

On-Chip Network Compatibility

- By allowing fine-grained interleaving of requests from different initiators, the on-chip network eliminates the need for high-bandwidth bursting and bandwidth-matching buffers in the initiators
- Per-thread flow control is applied to all transaction phases to ensure non-blocking behavior

Open Core Protocol

- Support at system and memory interfaces enables simple swapping of alternate DRAM technology controllers (e.g., SDRAM, DDR, etc.) without requiring redesign of MemMax Scheduler or the supporting on-chip network

Cycle-accurate SystemC Models

- Fast assessment of design trade-offs
- Uses same traffic generation and transactors as RTL
- Allows concurrent application software development

SONICS