

Connecting the Bandwidth Engine® MSR576 to Two Hosts



Application Note AN-602

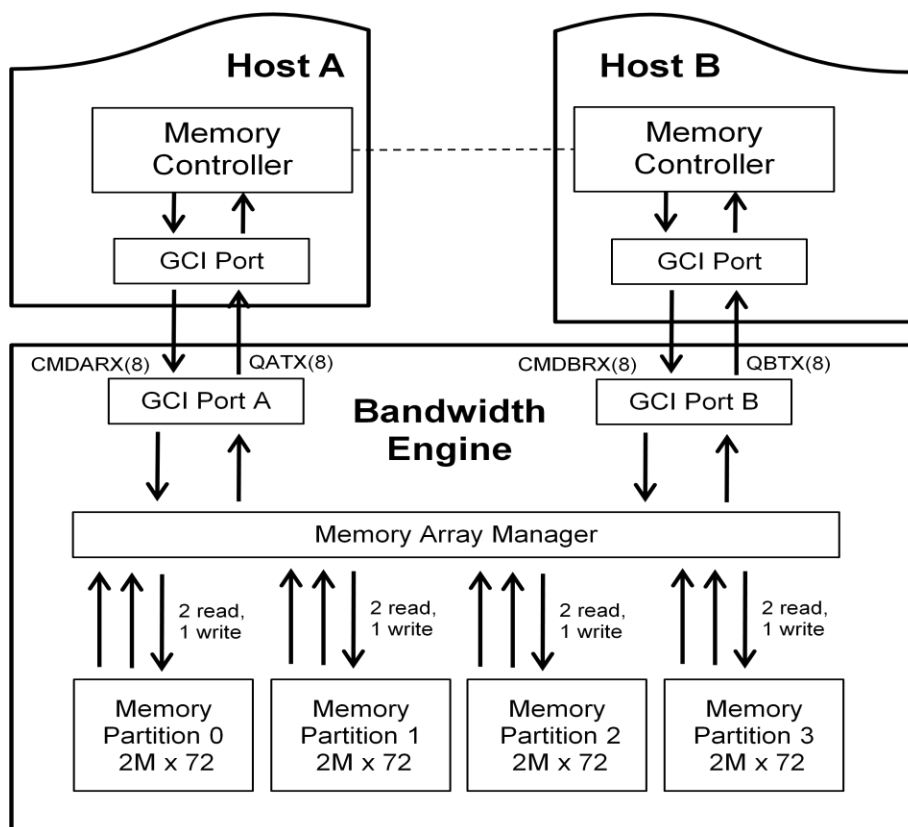
Version 0.1, March 2012
MoSys, Inc.

Introduction

The MSR576 devices belong to the Bandwidth Engine family of products. They are memory dominated devices with ALUs and low-latency, high-bandwidth SerDes interfaces. Using MoSys' 1T-SRAM® embedded memory technology, Bandwidth Engine ICs achieve higher speed than DRAM along with higher density and better reliability than SRAM. They offer higher bandwidth and lower pin count than parallel memory interfaces by means of the GigaChip™ Interface (GCI). The GigaChip Interface is an open, CEI-11G compatible serial protocol, optimized for high-bandwidth chip-to-chip communications. Bandwidth Engine ICs include ALUs that can perform operations on data within the chip, such as incrementing a counter. This feature reduces the latency and bandwidth requirements for macro operations.

This application note addresses how two hosts can be connected to a MSR576TE8888 Bandwidth Engine device as shown in Figure 1. The Bandwidth Engine block diagram shows the two GCI ports with 8 RX and 8 TX lanes each, the Memory Array Manager, and the four memory partitions within the device. This AN specifically discusses the MSR576 Bandwidth Engine device with 8 lanes per GCI port but is applicable to any Bandwidth Engine device with two GCI ports. This AN addresses the dual host Bandwidth Engine configuration while the datasheet addresses the common single host configuration.

Figure 1 Diagram of two hosts connected to a dual port Bandwidth Engine device



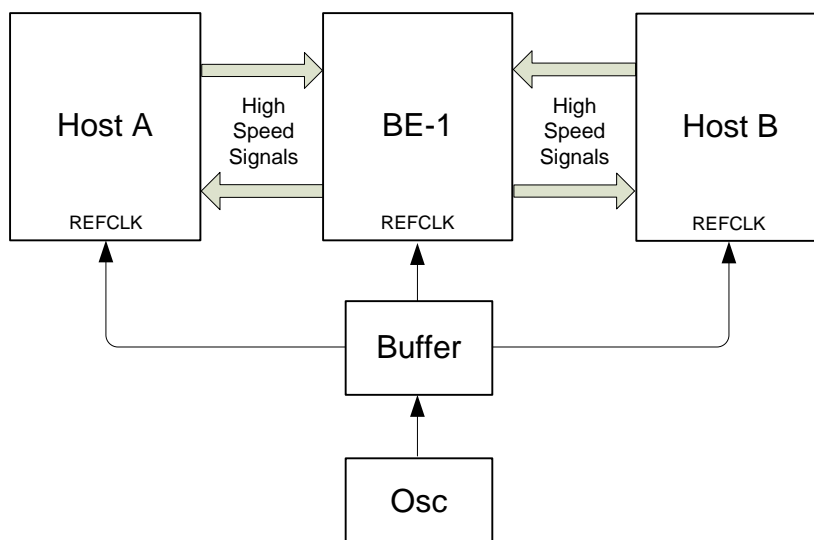
Requirements of Concern for Dual Host Configuration

The following section discusses the requirements when the Bandwidth Engine is connected to two hosts. For a detailed description of the device functionality see the MSR576 datasheet and the MSR576 User Guide.

Reference Clock

Both hosts and the Bandwidth Engine need to have a single mesosynchronous differential reference clock (REFCLKp/REFCLKn). The ideal method is to use the exact same clock from the same source to drive all three devices. This configuration is shown in Figure 2 below. For other configurations contact MoSys Applications Engineering.

Figure 2 Common clock source



Coordinated Reset and Configuration

The RESET#, CONFIG#, and READY# pins/functions require coordination between the hosts.

GCI Ports

The data interfaces of the two GCI ports need to be set to the loosely coupled mode using Framer Configuration Register 3. See the MSR576 User Guide for bit settings.

Port Deskew

The deskew across both ports needs to be turned off using PCS Configuration Register 3. This then allows deskew only within each port. See the MSR576 User Guide for bit settings.

Partition and Bank Conflicts

When each host has exclusive access to specific memory partitions then there is no possibility of access collisions with the other host. With exclusive partition access the partition and bank conflicts are the same as a single host configuration. For a detailed description of the partition and bank conflicts see the MSR576 datasheet and the MSR576 User Guide.

When both hosts share access to a memory partition but do not share access to specific memory banks then they must each use only one of the read paths for that partition and they must share the one write path for that partition (as shown in Figure 1).

Training and Timing Wheel Alignment

The GCI receivers of the Bandwidth Engine's two GCI ports must be aligned to slot 0 of the memory partitions' timing wheel when they come out training. The Bandwidth Engine accomplishes this by using the link offset feature defined in the GigaChip interface specification. This is done without intervention from either host. See the MSR576 User Guide or GigaChip Interface Specification for more details on link offset.

If one port of the MSR576 device decides to enter retraining, then the Bandwidth Engine will force the other port to also enter retraining without any intervention from the second host.

Bandwidth Engine Functionality for Each Host (GCI Port)

The following standard Bandwidth Engine functions are available to each host:

- Output Data FIFO for smoothing any excess READ data into gaps offered by subsequent WRITES or NOPs
- Replay Buffer for each Q TX interface for resending frames in the event of a bit error on the link
- Frame ID counter and logic for recovery of CRC errors on each CMD RX interface
- EVENTi# pin for notification of errors on each CMD RX interface
- Advanced debug capabilities using Trace Buffer circuit

Basic Options for Dual Configuration

Below lists two options for connecting two hosts to one Bandwidth Engine. They are shown in order of increasing complexity.

Option 1: Private Partitions

This is the simplest and cleanest option to implement, because there is no possibility of collision of accesses between the two host devices.

Each host has access to separate memory partitions in the Bandwidth Engine device. For example Host A could be assigned exclusive access to partitions 0 and 2 and Host B could be assigned exclusive access to partitions 1 and 3. The 4 partitions in BE could be allocated either 2 for each host, or 3 for one host and 1 for the other. With this configuration the two host devices essentially have un-constrained access to each of the assigned partitions, with the standard constraints of writes and reads to banks within the assigned partitions.

Option 2: Shared Partitions with Private Banks

This option has the benefit of dividing the memory access into 64 memory banks in each of the 4 partitions (256 total banks) versus the 4 partitions of Option 1 which offers better granularity for memory allocation. However, this option does have a small risk for any shared partition where certain CRC errors on one port have the potential to cause a bank conflict for the other port. This event is expected to be very rare but there is a risk that it can happen. For example, this could happen when a bit flips in the Opcode or Address field of a frame for Port A creating a CRC error for Port A. If that flipped bit creates an apparent bank conflict (i.e. both hosts trying to read from the same bank in the same partition at the same time) then both Port A and Port B would experience an error state. Port A would then need to initiate a replay out of its Replay Buffer. Port B should treat this as a system critical error and should have the software needed to repeat the whole operation again.

Each host has access to separate memory banks in the Bandwidth Engine. For this configuration any non shared partitions (exclusive access) would be handled as described in Option 1. For memory banks in shared partitions each host must have exclusive access to specific memory banks. Each host must also use only 1 read path in a shared partition. Another concern for this configuration is associated with constraints for writes into any memory partitions. Because each partition has only one write path (as shown in Figure 1) which must be shared by both hosts the host must be designed to avoid write collisions between ports. A simple way to ensure that a write collision is avoided will be to limit writes to every other turn of the wheel. For example, the normal rule of writing to a specific partition of an 8/16 RX lane device is once every 4 frames. This would not be allowed for shared partitions and each host would only be allowed to write to a specific shared partition every 8 frames.

Version History

Date	Version	Changes
March 2012	0.1	Initial release.

3/1/2012

MoSys, Inc.
3301 Olcott Street
Santa Clara, CA 95054
USA

Phone 408-418-7500
Fax 408-418-7501
Web <http://www.mosys.com>

The information provided in this document is subject to change without notice. MoSys, Inc. makes no warranties either express or implied with regard to the accuracy or completeness of the information contained herein.

1T-SRAM, Bandwidth Engine, and MoSys are registered trademarks of MoSys, Inc. in the U.S. and/or other countries. GigaChip, the GigaChip logo, and the MoSys logo are trademarks of MoSys, Inc. All other marks mentioned herein are the property of their respective owners.

Copyright © 2012 MoSys, Inc. All rights reserved.