Bandwidth Engine® MSR576 Schematic Design Guidelines

Application Note AN-606

Version 0.1, June 2012
MoSys, Inc.

Introduction

The MSR576 device belongs to the Bandwidth Engine family of products. It is a memory dominated device with ALUs and low-latency, high-bandwidth SerDes interfaces. Using MoSys' 1T-SRAM® embedded memory technology, the Bandwidth Engine IC achieves higher speed than DRAM along with higher density and better reliability than SRAM. It offers 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. The Bandwidth Engine IC includes 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 provides schematic design guidelines for using the MoSys® Bandwidth Engine MSR576 device assuming a generic FPGA host as shown in Figure 1. It also shows the MSR576 graphical schematic symbol. MSR576 device usage in ASIC/NPU environments is similar and can be easily derived from the FPGA example given in this document. For the purpose of this document, the FPGA is also known as host controller.

This document contains the following sections:

This document should be used in conjunction with AN-603 Bandwidth Engine MSR576 Board Design Guidelines [2] which addresses signal integrity and power integrity issues.

Figure 1 Suggested interface scheme for using MSR576 device with an FPGA

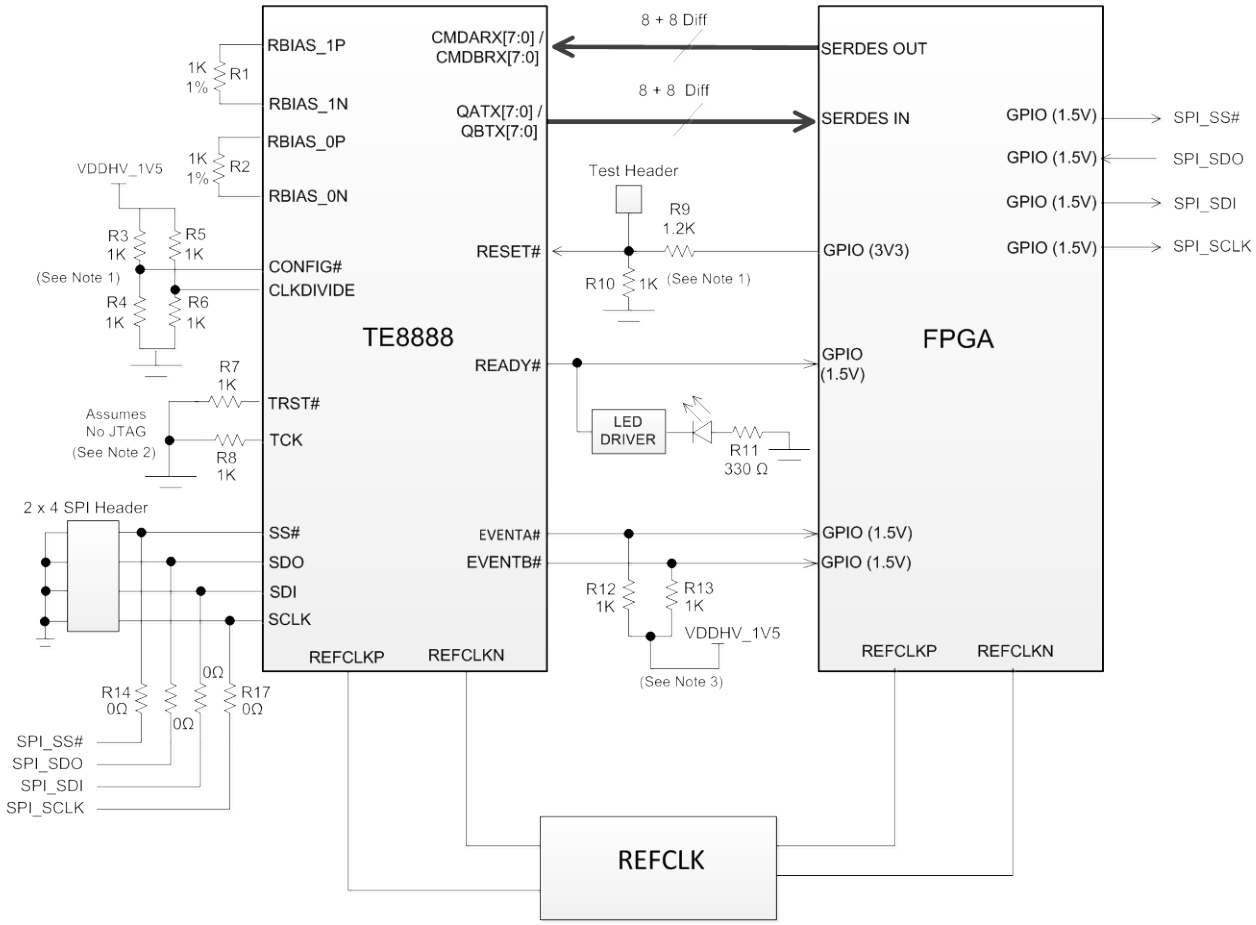

1. See RESET# Signal section for more information
2. See JTAG Port section for more information
3. See EVENT Pin section for more information

The following sections detail the PCB design requirements of different I/O groups on MSR576 device.

GCI Ports

The MSR576 device supports up to two GCI ports – A and B. Each GCI port has up to 8 RX lanes and 8 TX lanes. Thus there are up to 16 differential pairs in a GCI port. GCI Port A RX pairs are labeled as CMDARX while the TX differential pairs are labeled as QATX. Similarly, GCI Port B signals are labeled as CMDBRX and QBTX for the RX and TX pairs respectively. The GCI port differential pairs are compatible with CEI 11G SR I/O standard. The I/Os support only short reach point to point connections.

GCI Port Configurations

The MSR576 device can work with only GCI port A if 8 or less number of lanes is required. Please refer to the Selection Guide in the Bandwidth Engine MSR576 datasheet [1] for the actual lane configurations that are available. Unused differential

pairs in GCI port A and Port B can be left unconnected. All lane configurations have GCI port A and some also have GCI port B. When using both GCI ports, both ports will have the same number of RX lanes and both ports will also have the same number of TX lanes.

Lane Reordering

To provide flexibility in routing the SerDes signals of a GigaChip™ Interface link, the GCI specification has an option to logically reorder the lanes. During training, the transmitter sends each lane's logical lane identifier on the lane. The lane identifiers control multiplexers in the receiver that steer each physical lane to the desired logical lane. The MSR576 device implements a full crossbar in each GCI port, which allows any permutation of the lanes in the host-to-MSR576 link. If the system requires reordering on the lanes from the MSR576 device to the host device, the necessary multiplexers must be designed into the host's receiver. For more information about lane reordering, see Section 3.4, "Lane Reordering," in the GigaChip Interface Specification [4].

The example in Figure 2 shows the complete reversal of the lanes from the FPGA to MSR576 device.

Figure 2 CMDARX lanes reordering in MSR576 device

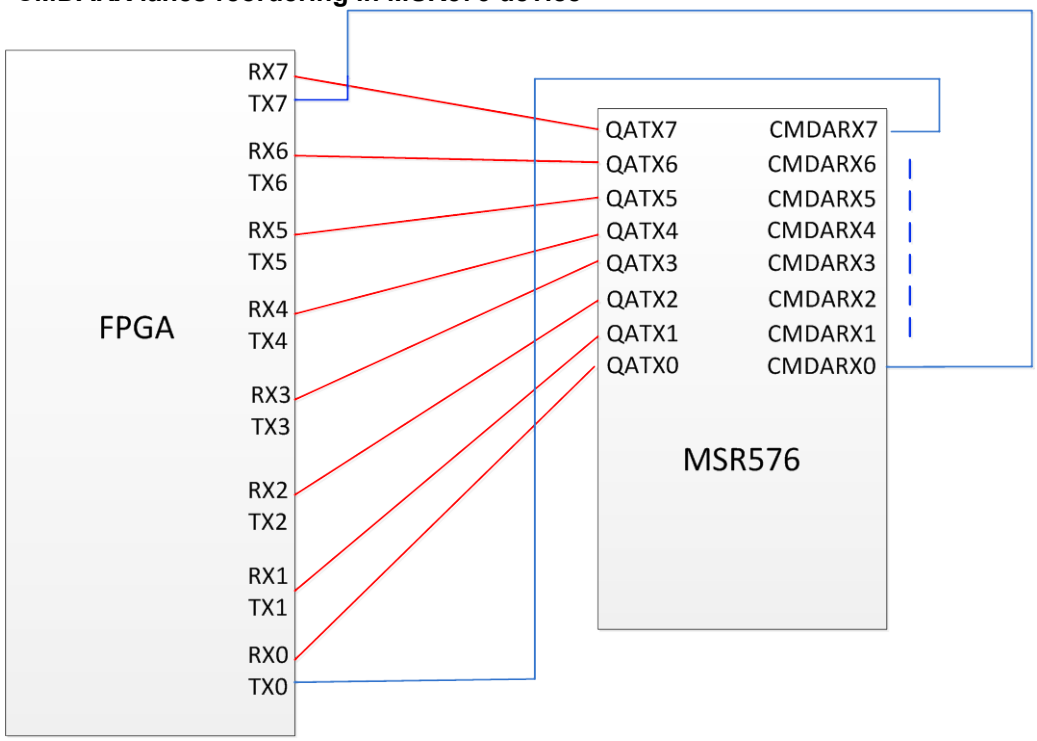

SERDES Interface

MoSys recommends connecting the SERDES interface of the MSR576 device and host controller in the partially AC coupled configuration shown in Figure 3. The

MSR576 RX lanes do not require external AC coupling capacitors because they are internally AC coupled.

As required by the GCI Specification, the PCS layer can detect polarity reversal for each lane and correct for the reversal.

Figure 3 Coupling schemes for interfacing MSR576 SERDES with the host controller FPGA

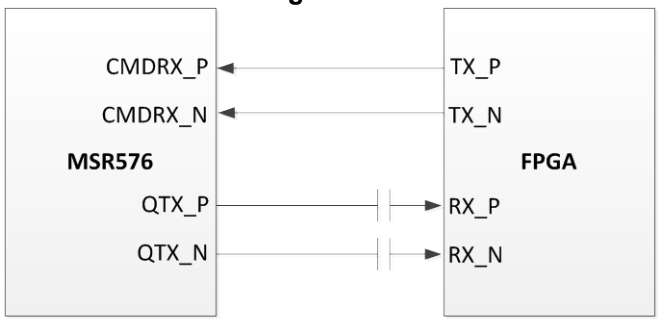

Partially AC Coupled

Spare Differential Pairs

Each GCI port has 2 spare RX and 2 spare TX lanes reserved for future use. These spare lanes can be left open. Alternatively, the spare lanes can be routed as regular differential pairs to the host controller for future versions of BE.

SERDES Bias Signals

MSR576 SERDES requires external bias resistors separately for each GCI port across RBIAS_xP/xN pins as shown in Figure 1. These RBIAS resistors require a value of 1 kΩ with 1% tolerance and they should be placed directly under the BGA footprint. Please refer to the Bandwidth Engine MSR576 Board Design Guidelines application note [2] for the RBIAS resistor layout details.

Reference Clock

MSR576 device uses the common reference clock scheme. The clock for MSR576 device and the host controller must be derived from the same source. Figure 4 shows a recommended reference clock circuit for the MSR576 device. The SI534 clock chip is available from Silicon Labs.

Figure 4 Recommended Reference clock circuit for MSR576 device

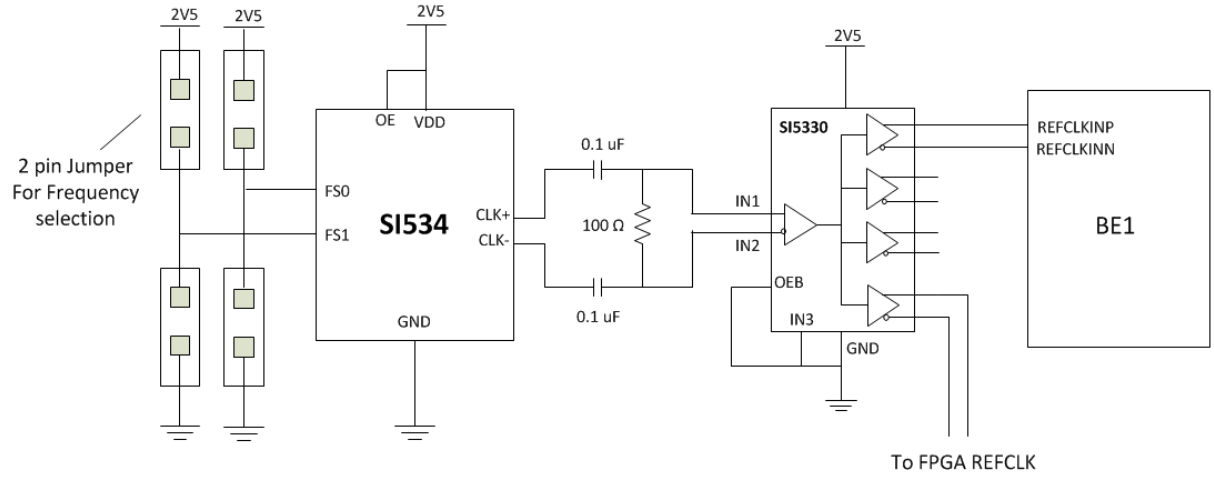

The MSR576 reference clock inputs can be directly driven by any LVDS fanout buffer as long as it meets the jitter and rise time requirements. Please refer to the Bandwidth Engine MSR576 datasheet [1] for the actual jitter and rise time specifications. Further, the maximum output common mode voltage of LVDS fanout buffer must be less than 1.3V. AC coupling capacitors should be used at the MSR576 reference inputs if the output common mode voltage of the LVDS fanout buffer is higher than 1.3V.

Table 1 lists the possible MSR576 data rates with four example reference clock frequencies.

Table 1 MSR576 REFCLK frequency and multiplier examples

REFCLK = 100 MHz		REFCLK = 125 MHz		REFCLK = 156.25 MHz		REFCLK = 312.5 MHz	
CLKDIVIDE = Low		CLKDIVIDE = Low		CLKDIVIDE = Low		CLKDIVIDE = High	
Multiplier	Data Rate (Gbps)	Multiplier	Data Rate (Gbps)	Multiplier	Data Rate (Gbps)	Multiplier	Data Rate (Gbps)
42	8.40	34	8.50	27	8.44	27	8.44
43	8.60	35	8.75	28	8.75	28	8.75
44	8.80	36	9.00	29	9.06	29	9.06
45	9.00	37	9.25	30	9.38	30	9.38
46	9.20	38	9.50	31	9.69	31	9.69
47	9.40	39	9.75	32	10.00	32	10.00
48	9.60	40	10.00	33	10.31	33	10.31
49	9.80	41	10.25				
50	10.00						
51	10.20						

SPI and JTAG Ports

SPI Port

The MSR576 SPI port is required for chip configuration and status monitoring. It can be directly connected to the FPGA or control/management plane processor with 1.5V LVCMOS compatible IOs. For debug purposes, we recommend to connect SPI signals to a 2 x 4 or 1x6 header as shown in Figure 1. The recommended USB to SPI host adapter is the Total Phase Cheetah. Zero ohm resistor jumpers in series with SPI signals (Figure 1) are for isolating SPI port from the FPGA during debug. A level translator circuit (SN74AVC4T774 or equivalent) must be used if the MSR576 SPI signals are connected to 3.3V LVCMOS IOs.

JTAG Port

The JTAG port is not normally required for MSR576 device bring up or debug. Hence, the main application for the JTAG port is boundary scan. Table 2 lists the JTAG functions needed for boundary scan. If the JTAG port is not used TRST# and TCK pins must be pulled down using 1K resistors as shown in Figure 1.

Table 2 JTAG functions

Pin	JTAG Function	Comments
TCK	Clock	Must be held Low if boundary scan not used.
TMS	Mode Select	May be left unconnected if JTAG not used
TRST#	Reset	Must be held Low except during JTAG operations
TDI	Data In	May be left unconnected if JTAG not used
TDO	Data Out	May be left unconnected if JTAG not used

Initial Board Bring-up and Debug Signals

We recommend using the following signals on the board for MSR576 device bring up and debug purposes,

VDD_KELVIN/VSS_KELVIN

VDD_KELVIN / VSS_KELVIN pins can be used to monitor on-die VDD power supply DC levels and AC noise. These pins can be connected to two test points (separated by 100 mils) on the PCB which can be probed differentially either using a real time scope (AC measurement) or a voltmeter (DC measurement). Note that the VDD_KELVIN / VSS_KELVIN pins are not shown in Figure 1.

AMON

There are two AMON pins (AMON_0 and AMON_1) on the MSR576 device corresponding to the two GCI ports that monitor internal analog voltages. These AMON pins can be helpful for debugging the SERDES and can be connected to test points on the PCB. Note that the AMON signals are not shown in Figure 1.

DMON

There are two DMON differential pairs (DMON_0P/N and DMON_1P/N) on the MSR576 device corresponding to the two GCI ports. They are used to monitor internal SerDes related high-speed signals and they can be connected to SMA connectors. Note that the DMON signals are not shown in Figure 1.

Miscellaneous Signals

The miscellaneous signals in MSR576 device can be classified in two groups: Initialization and status monitors. The board design requirements for these two signal groups are described below.

Initialization Signals

The MSR576 initialization signals are RESET#, CONFIG#, and CLKDIVIDE

RESET#

There are three options for connecting the MSR576 RESET# input pin depending on the FPGA driver type. Table 3 below describes all three options and Figure 1 shows option #1:

1. If the pin is driven from a 3.3V LVCMOS driver it can be connected to the FPGA as shown in Figure 1. Both resistor R10 and R9 are required.
2. If the pin is driven from a 1.5V LVCMOS driver R10 is not needed and it is recommended that a series resistor (R9) of zero ohms be used so the MSR576 RESET# pin can be isolated from the FPGA for debug.
3. If the pin is driven by an open drain driver it must be pulled to 1.5V supply with a resistor of value 1kΩ or less. It is recommended that a series resistor (R9) of zero ohms be used so the MSR576 RESET# pin can be isolated from the FPGA for debug.

Table 3 RESET pin connection options

Option	FPGA Driver	R10	R9	Rpu (pull-up)
1	3.3V LVCMOS	1k ohm	1.2k ohm	N/A
2	1.5V LVCMOS	N/A	0 ohm	N/A
3	Open Drain	N/A	0 ohm	1k ohm maximum

For MSR576 device bring up and debug purposes, it is recommended to connect the RESET# pin to a header as shown in Figure 1. This will be helpful to reset the MSR576 device independent of the FPGA. Please consult the Bandwidth Engine MSR576 datasheet [1] and the Bandwidth Engine MSR576 Power-up and Reset AN-604 [3] for the requirements on reset duration, rise and fall times.

In addition to the above hardware RESET, it is also recommended to reset the MSR576 device when any of its supply voltage falls below – 5% of the nominal voltage. Further, the MSR576 reset input should be held in the asserted state until all the power supplies ramp above 5 % of their nominal levels. A reset monitor circuit can be used for this purpose.

CONFIG# and CLKDIVIDE

During reset, the CONFIG# pin when asserted provides a method to use the SPI port to override the default configuration of MSR576 device (such as setting up the PLL clock multiplier). A Low value causes the PLL initialization to be held off until the CONFIG_DONE bit in register rstgen_rst is set via the SPI port. A High signal causes the default values for all registers to be loaded and the PLL initialization to proceed automatically upon de-assertion of RESET#.

Asserting the CLKDIVIDE pin during reset causes external reference clock frequency to be divided by 2 (for example, a 312.5 MHz external REFCLK becomes an internal 156.25 MHz REFCLK). A Low value causes the divide circuit to be bypassed. Normally, the CLKDIVIDE pin is held Low for reference clock frequencies lower than 200 MHz and it is held High for frequencies above 200 MHz. See erratum #56 of the Bandwidth Engine MSR576 Errata List [5] for a possible CLKDIVIDE issue.

It is recommended to use both pull up (R3/R5) and pull down resistors (R4/R6) for the CONFIG# and CLKDIVIDE pins on the board as shown in Figure 1. The pull up or pull down resistors can be stuffed appropriately per the desired reset configuration. Alternatively, use of jumpers on the board as shown in Figure 5 can add more flexibility and ease for the MSR576 reset configuration.

Figure 5 MSR576 reset configuration using 2 jumpers on the board

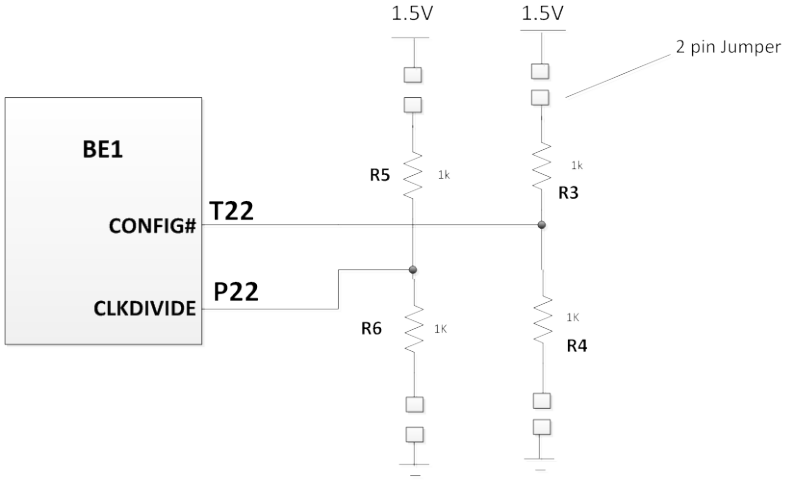

Status Monitor Signals

There are 3 status output pins on MSR576 device namely READY#, EVENTA# and EVENTB#.

READY#

The READY# output is asserted Low to indicate the device has completed configuration and the internal PLL has locked. Lane training begins automatically. When training is complete the device is ready to receive commands over the

CMDARX and CMDBRX interfaces. READY# signal is very useful for MSR576 device bring-up and debug. It is recommended to use a LED for the READY# pin for visual status of the MSR576 device. The READY# pin requires an external LED driver as it supports only 1.5V LVCMOS IO. A Reference LED driver circuit using a Micrel MIC94051 MOSFET is shown in Figure 6. The READY# output of MSR576 has an active pull-up per erratum #36 of the Bandwidth Engine MSR576 Errata List [5]. On BE-2 the READY# output will be open-drain.

Figure 6 Reference LED driver circuit for monitoring READY# pin status

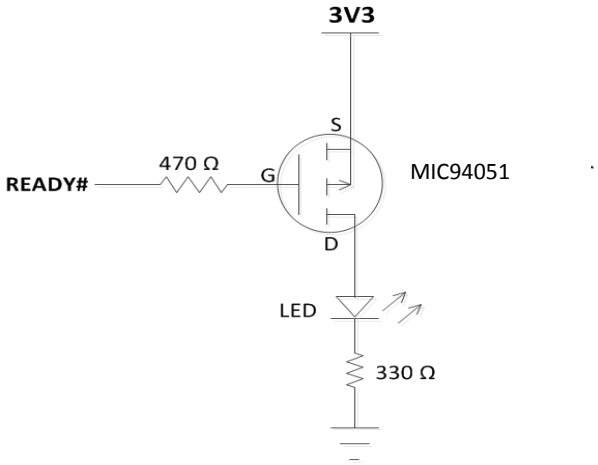

EVENTA# and EVENTB#

EVENTA# and EVENTB# pins are used to indicate an error on CMDARX and CMDBRX interfaces respectively. Asserts Low to indicate an event has been posted to a Status Register. It is recommended that the EVENT# pins be routed as short a trace as possible back to the controller. This is to minimize timing delays from EVENT# to the host controller recognizing EVENT#. See the Bandwidth Engine MSR576 Board Design Guidelines AN-603 [2] for further board routing recommendations.

The board configuration shown in Figure 7 can be used for either MSR576 device (BE-1) or BE-2 devices. MSR576 device EVENT# outputs have active pull-ups per erratum #36 of the Bandwidth Engine MSR576 Errata List [5] and do not need resistors R1 or R2. BE-2 EVENT# outputs are open drain and need a pull-up resistor to VDDHV. Table 4 shows which resistors to populate for each configuration.

Figure 7 EVENT# pin connections

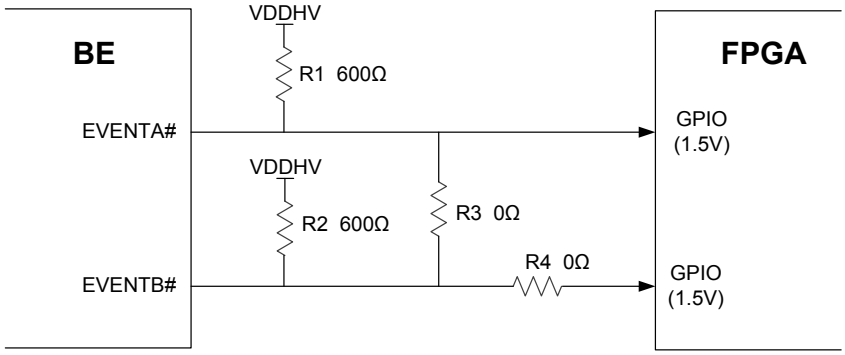

Table 4 Event pin connection options

BE Device	Number of GCI Ports	FPGA Connection	Populate These Resistors
MSR576 (BE-1)	1	One input pin	none
MSR576 (BE-1)	2	Two input pins	R4
BE-2	1	One input pin.	R1
BE-2	2	One input pin. Wire-OR configuration	R1, R3
BE-2	2	Two input pins	R1, R2, R4

Die Temperature Monitoring

It is critical to maintain the MSR576 device die temperature below 95C for proper device operation. In order to monitor die temperature, the MSR576 device provides a thermal diode interface which can be connected to a temperature monitor circuit as shown in Figure 8. The Maxim MAX 6627 device has a built in ADC through which digital equivalent of on-die temperature can be read via the SPI interface. Note that the MSR576 thermal diode interface pins (TEMPDIODEP/N) can be left open if the temperature monitor is not used.

Figure 8 Reference Die temperature monitor circuit for MSR576 device

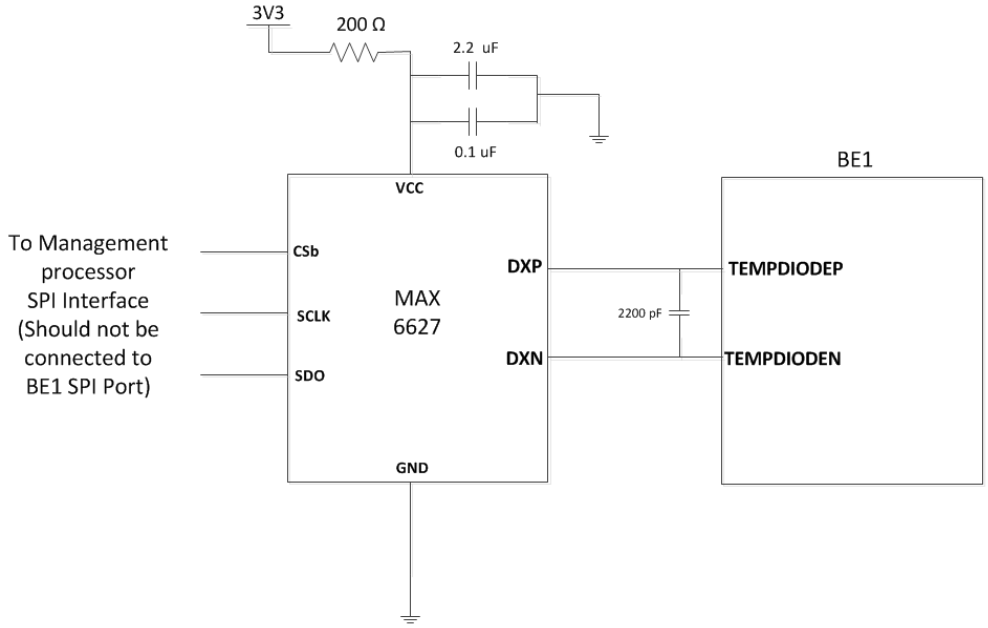

Power Supply

Figure 9 shows the recommended power supply scheme for MSR576 device. As shown, digital VSS and SERDES VSS (VSS_SDS) can be merged on the board as 1 VSS supply. We recommend using a linear Low Drop Out (LDO) voltage regulator for the SERDES 1.1V supply. Further, SERDES power supply domains should be filtered using ferrite beads unless a dedicated voltage regulator is used (which is the case for the 1.1V supply shown in Figure 9). See AN-603 for power routing guidelines.

Figure 9 Recommended Power Supply Scheme for MSR576 device

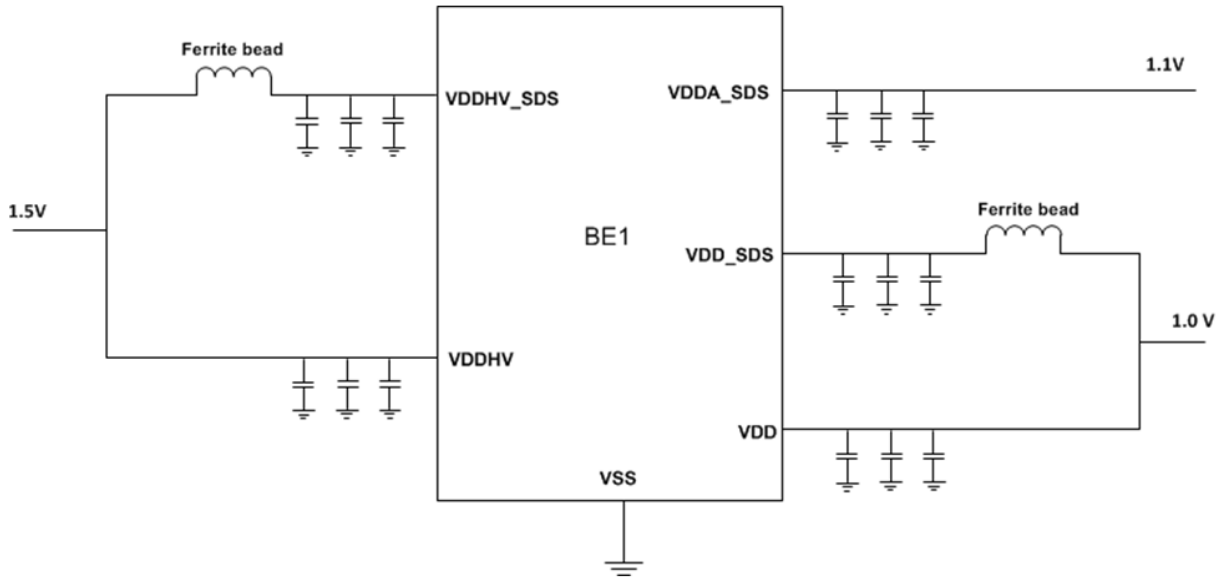

Ferrite Bead Selection

Important factors for selecting ferrite beads are current capacity, DC resistance and AC impedance. The AC impedance is usually specified at 100 MHz. The Ferrite beads that have higher AC impedance tend to have higher DC resistance as well. This DC resistance can cause excessive IR drop leading to the device failure. The following equation can be used to calculate required DC resistance of the ferrite bead (R_{fb})

$$Ferrite\ Bead\ Resistance < \frac{0.03 \times Vnom}{I} \quad (1)$$

Where Vnom is the nominal voltage of the power supply and I is the DC current consumed by the MSR576 device. In the equation 1, it is assumed that the power supply provides nominal voltage and the voltage drop across the board is less than 2%.

Although ferrite beads with high AC impedance may seem better for noise suppression, it may cause power supply network to resonate. A recommended ferrite bead for filtering the VDD_SDS and VDDHV_SDS domains is given in table below,

Table 5 Recommended ferrite bead

Characteristic	Value
Manufacturer Part number	BLM18KG300TN1D
Manufacturer	Murata (www.murata.com)
Current rating	5A
DC Resistance	10 mΩ
AC resistance	30 Ω @ 100 MHz
Body size	0603

MSR576 Current Requirements

The current consumption on the MSR576 device depends on the data rate and number of lanes. See the Bandwidth Engine MRS576 datasheet [1] for the power supply current requirements.

Decoupling Requirements

Decoupling capacitors are required on the board to meet the power supply noise specifications. There are two methodologies to estimate decoupling capacitor requirements on the board. The simplest approach is to add a high frequency decoupling capacitor (typically 0.1 uF) per power supply pin. These high frequency decoupling capacitors are required to be placed under the MSR576 BGA footprint for maximal effectiveness. Please refer to Bandwidth Engine MSR576 Board Design Guidelines [2] for further details on the placement of high frequency decoupling capacitors.

In addition to high frequency decoupling capacitors, bulk decoupling capacitors are required for low frequency decoupling. Typically, a network of capacitors in parallel of values ranging from 1 uF to 470 uF forms an effective bulk decoupling solution.

Although the approach described above is simple to implement, it may lead to using more capacitors than required. The second method to derive the on-board decoupling requirements is to use the target impedance method which optimizes the number of capacitors required. In this method, an impedance target is defined separately for each power supply domain. The impedance target is met by adding appropriate number of capacitors typically using EM simulators such as Power-SI or SIWAVE. Please refer to Bandwidth Engine MSR576 Board Design Guidelines AN-603 [2] for the target impedance specification for MSR576 power domains.

Heatsink Guidelines

A heatsink is recommended for the MSR576 device which comes in a lidless flip-chip BGA package. The selected heatsink and thermal interface material (TIM) need to ensure that the maximum operating case temperature as specified in the MSR576 datasheet is not exceeded. The TIM is required to be used with the selected heatsink and must be applied directly to the backside of the die before the heatsink is attached. The heatsink type and TIM need to ensure that the heatsink can be removed in case the device is returned to MoSys.

Thermal Interface Material (TIM)

The selected TIM needs to allow stress relief between the heatsink and the back side of the MSR576 die. Thermal pads, greases, pastes, and liquids are allowed. Curing epoxy is not allowed because any mechanical stress applied to the heatsink would be directly passed on to the MSR576 die bumps and package balls. Table 6 contains a list of acceptable TIMs.

Table 6 Best practice thermal interface materials

Type	Vendor	Product	Conductivity	Format
Phase change	Chomerics	T777	3.5 W/m°K	Pad
Phase change	Laird Technologies	Tpcm 780	5.5 W/m°K	Pad
Phase change	Laird Technologies	Tpcm 780SP	5.5 W/m°K	Liquid
Phase change	Laird Technologies	Tpcm 580	3.8 W/m°K	Pad
Thermal grease	Laird Technologies	Tgrease 980	3.8 W/m°K	Grease
Thermal grease	Laird Technologies	Tgrease 880	3.1 W/m°K	Grease
Thermal grease	Laird Technologies	Tgrease 2500	3.8 W/m°K	Paste

Heatsink Attachment Methods

Bolting the heatsink to the printed circuit board is highly recommended. This technique allows maximum stress relief from mechanical shock because any shock to the heatsink is transferred directly to the PC board and not to the MSR576 die/package. Alpha Novatech and Cool Innovations are two vendors that offer heatsinks that bolt to the board and are compatible with MSR576 and BE-2 devices.

Heatsinks that clip-on the device are acceptable but do not offer as much stress relief as a bolt to the board design. Malico and Cooliance are two of the vendors that offer heatsinks that clip to BGA packages and are compatible with MSR576 and BE-2 devices.

Downward Pressure on MSR576

When downward (Z axis) pressure is applied to the MSR576 both the device and the board must be supported to prevent any flexing or bowing.

Up to 14.5kg is allowed in the Z axis for short duration pressure like heatsink mounting or electrical test.

Up to 5kg is allowed in the Z axis for long term static pressure like pressure from a heatsink clip or from a heatsink bolted to the board.

Heatsink Removal

The selected heatsink and TIM need to allow for heatsink removal without damaging the MSR576 device. MoSys will not accept MSR576 devices returned with a heatsink still attached.

Schematic Symbol

MoSys provides schematic symbol (Figure 10) of the MSR576 device that can be used in customer board schematics. The schematic symbol is available in the industry standard ORCAD format as well as in other formats. Please contact MoSys Applications to get the MSR576 schematic symbol with the format required. Use of the MoSys provided schematic symbol is recommended as it has been verified and validated.

Figure 10 MSR576 Schematic Symbol

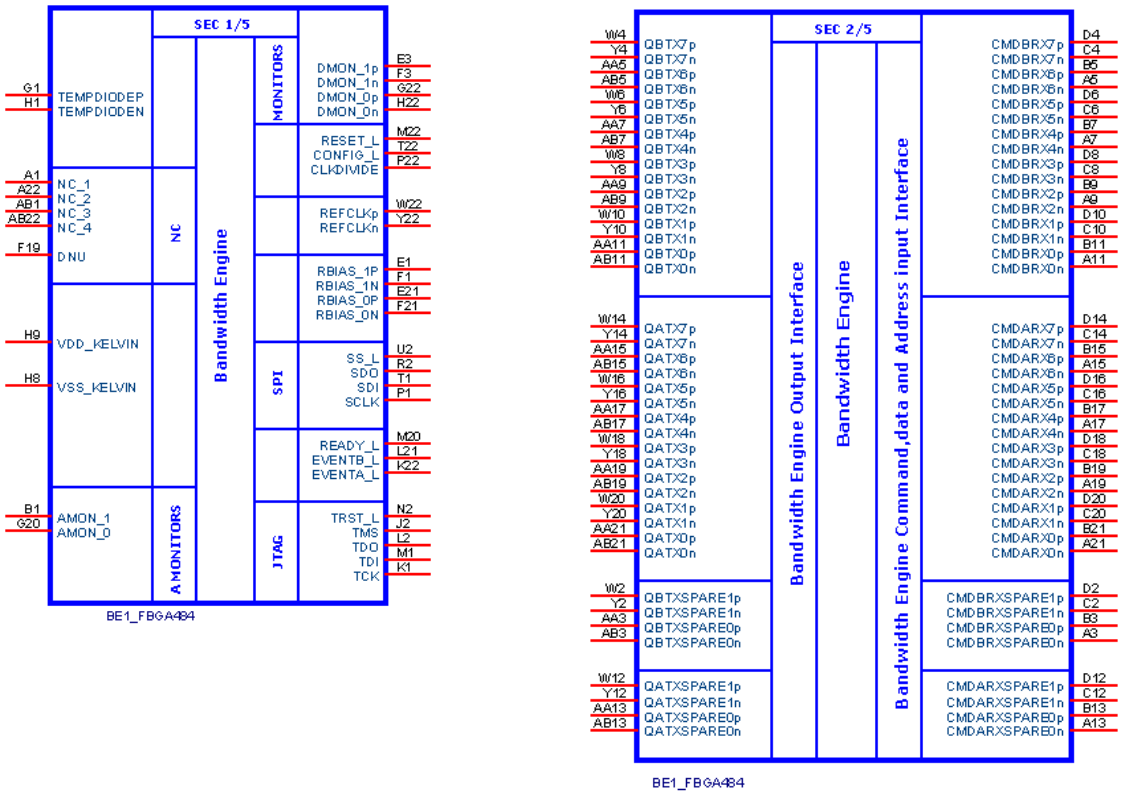

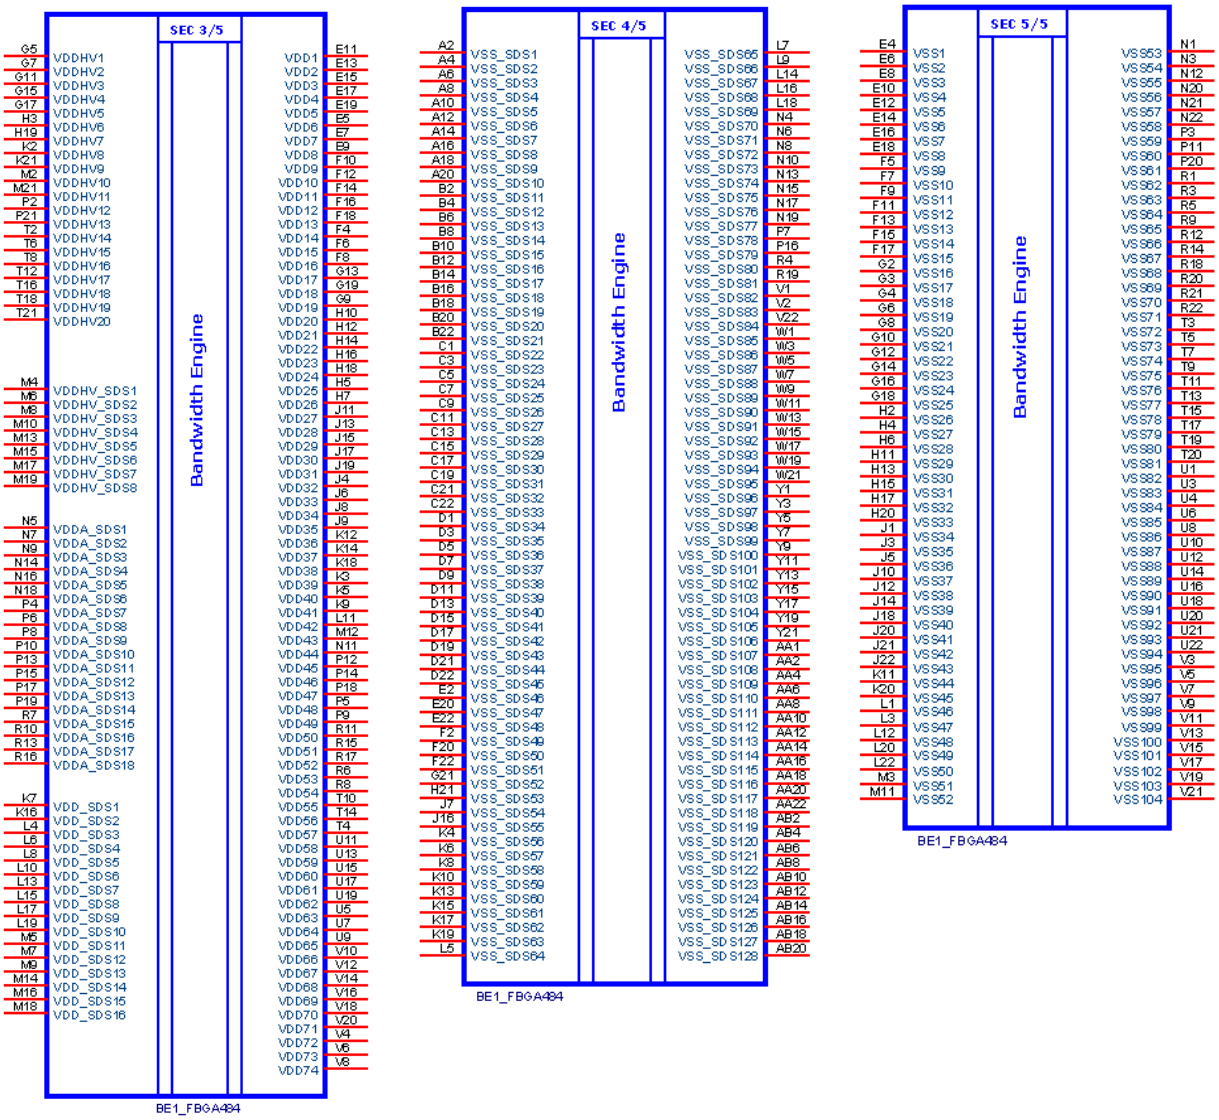

References

[1] Bandwidth Engine MSR576 Datasheet --- Version 1.0, June, 2012

[2] AN-603, Bandwidth Engine MSR576 Board Design Guidelines --- Version 0.1, March 2012

[3] AN-604, Bandwidth Engine MSR576 Power-up and Reset --- Version 0.1, July 2011

[4] GigaChip Interface Specification --- Version 1.0, January 2012

[5] Bandwidth Engine MSR576 Errata List --- Version 1.1, June 2012

Version History

Date	Version	Changes
June 2012	0.1	Initial release

6/13/12

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

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