




TESLA M2090 DUAL-SLOT COMPUTING PROCESSOR MODULE

BD-05766-001_v03 | June 2012

Board Specification



DOCUMENT CHANGE HISTORY

BD-05766-001_v03

Version	Date	Authors	Description of Change
01	March 21, 2011	GG, SM	Initial Release
02	June 16, 2011	GG, SM	Updated memory configuration in Table 1
03	June 27, 2012	GG, SM	Updated SKU reference in Table 1

TABLE OF CONTENTS

- Overview** 1
 - Key Features 1
 - Computing Processor Description 2
 - Configuration 3
- Mechanical Specifications** 4
 - PCI Express System 4
 - Standard I/O Connector Placement 5
 - Internal Connectors and Headers 6
 - External PCI Express Power Connectors 6
- Power Specifications** 10
- Support Information** 11
 - Certificates and Agencies 11
 - Agencies 11
 - Languages 12

LIST OF FIGURES

Figure 1.	Tesla T20A GPU Block Diagram	2
Figure 2.	Tesla M2090 Computing Processor Module.....	4
Figure 3.	Tesla M2090 Bracket	5
Figure 4.	6-Pin PCI Express Power Connector	7
Figure 5.	8-Pin PCI Express Power Connector	8

LIST OF TABLES

Table 1.	Tesla M2090 Board Configuration	3
Table 2.	6-Pin PCI Express Power Connector Pinout	9
Table 3.	8-Pin PCI Express Power Connector Pinout	9
Table 4.	Configuration with External PCI Express Connectors.....	10
Table 5.	Languages Supported	12

OVERVIEW

The NVIDIA® Tesla™ M2090 graphics processing unit (GPU) Computing Module is a PCI Express, double-wide, full-height (4.376 inches by 9.75 inches by 1.52 inches) form factor computing module based on the NVIDIA Fermi GPU. This module comprises a computing subsystem with a GPU and high speed memory. The Tesla M2090 module offers 6 GB of GDDR5 on-board memory.

The Tesla M2090 can be configured by the OEM or by the end user to enable or disable ECC or error correcting codes that can fix single-bit errors and report double-bit errors. Enabling ECC will cause some of the memory to be used for the ECC bits, so the user available memory will decrease to approximately 5.25 GB for a Tesla M2090.

KEY FEATURES

GPU

- ▶ Number of processor cores: 512
- ▶ Processor core clock: 1.3 GHz
- ▶ Package size: 42.5 mm × 42.5 mm 1981-pin ball grid array (BGA)

Board

- ▶ PCI Express Gen2 ×16 system interface
- ▶ Physical dimensions: 4.376 inches × 9.75 inches, dual slot
- ▶ Board power dissipation: <= 225 W

External Connectors

- ▶ None

Internal Connectors and Headers

- ▶ One 6-pin PCI Express power connector
- ▶ One 8-pin PCI Express power connector

Memory

- ▶ Memory clock: 1.85 GHz
- ▶ Interface: 384-bit
 - 6 GB
 - 24 pieces 128M × 16 GDDR5 136-pin BGA, SDRAM

BIOS

- ▶ 2Mbit Serial ROM

COMPUTING PROCESSOR DESCRIPTION

Figure 1 is a block diagram of the Tesla T20A GPU used on the Tesla M2090 dual-slot computing processor module.

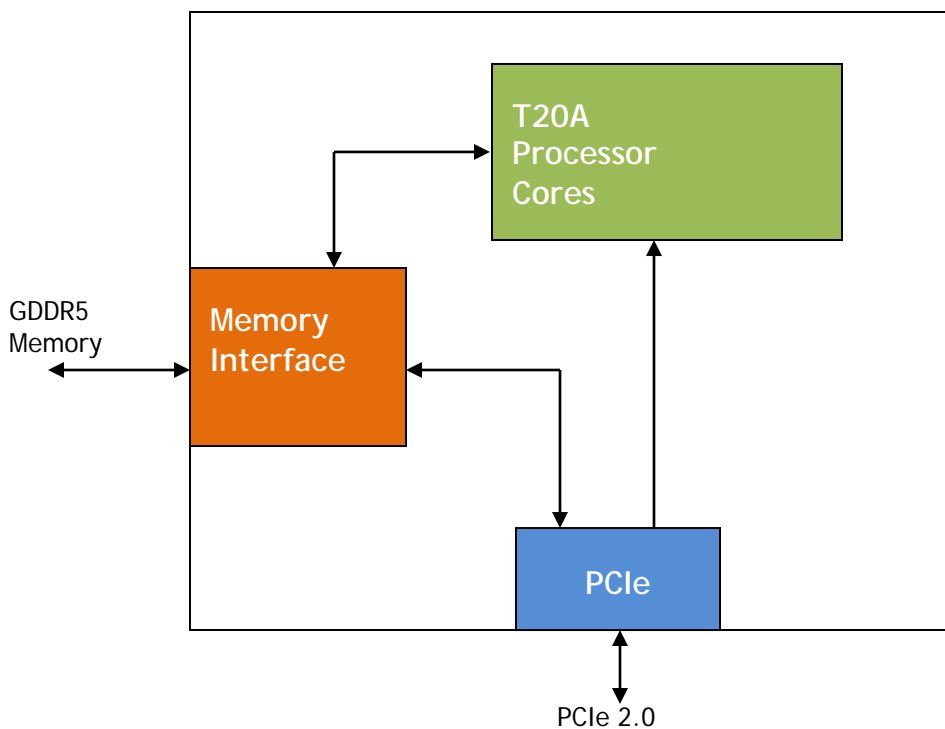


Figure 1. Tesla T20A GPU Block Diagram

CONFIGURATION

There is one configuration available (Table 1) for the Tesla M2090 dual-slot computing processor module.

Table 1. Tesla M2090 Board Configuration

Specification	Description
Generic SKU reference	<ul style="list-style-type: none"> 900-21030-0040-100 (ships without bracket) 900-21030-0045-100 (ships with bracket)
Chip	T20A GPU
Package size	42.5 mm × 42.5 mm
Processor clock	1.3 GHz
Memory clock	1.85 GHz
Memory I/O	384-bit GDDR5
Memory configuration	24 pcs 128M × 16 GDDR5 SDRAM
External connectors	None
Internal connectors and headers	<ul style="list-style-type: none"> 8-pin PCI Express power connector 6-pin PCI Express power connector
Board power	<= 225 W
Thermal cooling solution	Passive heat sink

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla M2090 computing processor board (Figure 2) conforms to the PCI Express full height (4.376 inches by 9.75 inches) form factor. Figure 2 is shown without the bracket.




Figure 2. Tesla M2090 Computing Processor Module

STANDARD I/O CONNECTOR PLACEMENT

As shown in Figure 3, the Tesla M2090 includes a vented bracket. If you are an OEM who qualifies for bracket modifications, you have the option of receiving your module with no bracket installed.




Figure 3. Tesla M2090 Bracket

INTERNAL CONNECTORS AND HEADERS

The Tesla M2090 module supports the following internal connectors and headers.

- ▶ 8-pin PCI Express power connector (can be used with a 6-pin power cable)
- ▶ 6-pin PCI Express power connector

External PCI Express Power Connectors

The Tesla M2090 module is a performance-optimized, high-end product and uses power from the PCI Express connector as well as external power connectors. The board can be used in two different ways.

- ▶ One 8-pin PCI Express power connector or
- ▶ Two 6-pin PCI Express power connectors

Figure 4 and Figure 5 show the specifications and Table 2 and Table 3 show the pinouts for the 6-pin and 8-pin PCI Express power connectors.




Figure 4. 6-Pin PCI Express Power Connector




Figure 5. 8-Pin PCI Express Power Connector

Table 2. 6-Pin PCI Express Power Connector Pinout

Pin Number	Description
1	+12 V
2	+12 V
3	+12 V
4	GND
5	Sense
6	GND

Table 3. 8-Pin PCI Express Power Connector Pinout

Pin Number	Description
1	+12 V
2	+12 V
3	+12 V
4	Sense1
5	GND
6	Sense0
7	GND
8	GND

POWER SPECIFICATIONS

The Tesla M2090 requires power from the PCIe connector as well as one or two auxiliary power connectors.

Table 4. Configuration with External PCI Express Connectors

8-Pin Power Connector	6-Pin Power Connector	Result
Connected (either 8-pin or 6-pin)	Connected	Valid configuration - board will operate to specification
8-pin connected	Not connected	Valid configuration - board will operate to specification
6-pin connected	Not connected	Insufficient power - board will not operate
Not connected	Connected	Insufficient power - board will not operate
Not connected	Not connected	Insufficient power - board will not operate



Note: Detailed information about power draw by rail is available to authorized system partners in the *Tesla M2090 System Design Guide*.

SUPPORT INFORMATION

CERTIFICATES AND AGENCIES

Agencies

- ▶ Australian Communications Authority and Radio Spectrum Management Group of New Zealand (C-Tick)
- ▶ Bureau of Standards, Metrology, and Inspection (BSMI)
- ▶ Conformité Européenne (CE)
- ▶ Federal Communications Commission (FCC)
- ▶ Industry Canada - Interference-Causing Equipment Standard (ICES)
- ▶ Korean Communications Commission (KCC)
- ▶ Underwriters Laboratories (cUL)
- ▶ Voluntary Control Council for Interference (VCCI)

LANGUAGES

Table 5. Languages Supported

	Windows Server 2008 and Windows Server 2008 R2	Linux
English (US)	X	X
English (UK)	X	
Arabic	X	
Chinese, Simplified	X	
Chinese, Traditional	X	
Danish	X	
Dutch	X	
Finnish	X	
French	X	
French (Canada)	X	
German	X	
Italian	X	
Japanese	X	
Korean	X	
Norwegian	x	
Portuguese (Brazil)	X	
Russian	X	
Spanish	X	
Spanish (Latin America)	X	
Swedish	X	
Thai	X	

Note: NVIDIA's CUDA software is only supported in English (U.S.)

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

ROVI Compliance Statement

NVIDIA Products that support Rovi Corporation's Revision 7.1.L1 Anti-Copy Process (ACP) encoding technology can only be sold or distributed to buyers with a valid and existing authorization from ROVI to purchase and incorporate the device into buyer's products.

This device is protected by U.S. patent numbers 6,516,132; 5,583,936; 6,836,549; 7,050,698; and 7,492,896 and other intellectual property rights. The use of ROVI Corporation's copy protection technology in the device must be authorized by ROVI Corporation and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by ROVI Corporation. Reverse engineering or disassembly is prohibited.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, CUDA and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2011, 2012 NVIDIA Corporation. All rights reserved.