




TESLA M2050 AND TESLA M2070 DUAL-SLOT COMPUTING PROCESSOR MODULES

BD-05238-001_v02 | April 2010

Board Specification



DOCUMENT CHANGE HISTORY

BD-05238-001_v02

Version	Date	Authors	Description of Change
01	April 22, 2010	GB, SM	Initial Release (Preliminary information)
02	April 28, 2010	GB, DV	Removed "Preliminary Information" text Added Figure 3 Memory clock changed from 1.6 GHz to 1.546 GHz

TABLE OF CONTENTS

- Overview 1**
 - Key Features 2
 - Computing Processor Description..... 3
 - Configuration..... 4

- Mechanical Specifications 5**
 - PCI Express System 5
 - Standard I/O Connector Placement 6
 - Internal Connectors and headers 7
 - External PCI Express Power Connectors..... 7

- Power Specifications..... 11**

- Support Information 12**
 - Certificates and Agencies 12
 - Agencies 12
 - Languages..... 13

LIST OF FIGURES

Figure 1.	Tesla T20 GPU Block Diagram	3
Figure 2	Tesla M2050 and Tesla M2070 Computing Processor Module	5
Figure 3.	Tesla M2050 and Tesla M2070 Bracket.....	6
Figure 4.	6-Pin PCI Express Power Connector	8
Figure 5.	8-Pin PCI Express Power Connector	9

LIST OF TABLES

Table 1.	Board Configurations.....	4
Table 2.	6-Pin PCI Express Power Connector Pinout	10
Table 3.	8-Pin PCI Express Power Connector Pinout	10
Table 5.	Configuration with External PCI Express Connectors.....	11
Table 7.	Languages Supported	13

OVERVIEW

The NVIDIA® Tesla™ M20-Series 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.

This module is offered with two memory sizes:

- ▶ The Tesla M2050 module offers 3 GB of GDDR5 memory on board.
- ▶ The Tesla M2070 module offers 6 GB of GDDR5 memory on board.

Both of these products 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 2.62 GB for a Tesla M2050 and approximately 5.25 GB for a Tesla M2070.

KEY FEATURES

GPU

- ▶ Number of processor cores: 448
- ▶ Processor core clock: 1.15 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: 1.546 GHz
- ▶ Interface: 384-bit
 - Tesla M2050
 - 3 GB
 - 24 pieces 64M × 16 GDDR5 136-pin BGA, SDRAM
 - Tesla M2070
 - 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 T20 GPU used on the Tesla M20-series products.

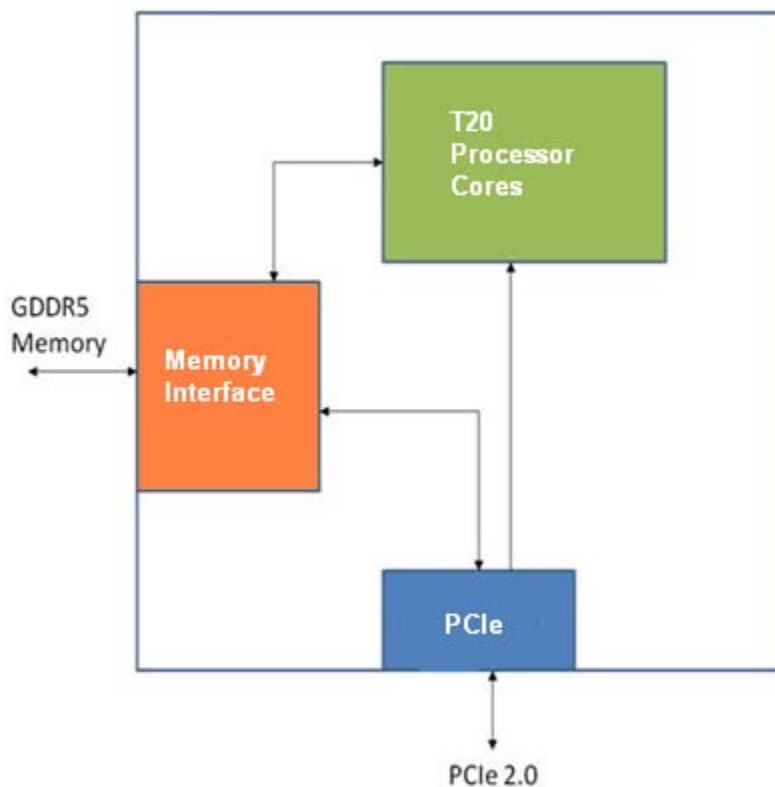


Figure 1. Tesla T20 GPU Block Diagram

CONFIGURATION

There is one configuration available (Table 1) for the Tesla M2050 and Tesla M2070.

Table 1. Board Configurations

Specification	M2050	M2070
Generic SKU reference	900-21030-0050-000	900-21030-0070-000
Chip	Tesla T20 GPU	Tesla T20 GPU
Package size GPU	42.5 mm x 42.5 mm	42.5 mm x 42.5 mm
Processor clock	1.15 GHz	1.15 GHz
Memory clock	1.546 GHz	1.546 GHz
Memory size	<ul style="list-style-type: none"> •3 GB total •2.62 GB available with ECC enabled 	<ul style="list-style-type: none"> •6 GB total •5.25 GB available with ECC enabled
Memory I/O	384-bit GDDR5	384-bit GDDR5
Memory configuration	24 pcs 64M × 16 GDDR5 SDRAM	24 pcs 128M × 16 GDDR5 SDRAM
External connectors	None	None
Internal connectors and headers	<ul style="list-style-type: none"> •8-pin PCI Express power connector •6-pin PCI Express power connector 	<ul style="list-style-type: none"> •8-pin PCI Express power connector •6-pin PCI Express power connector
Board power	< = 225 W	< = 225 W
Thermal cooling solution	Passive heat sink	Passive heat sink

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla M2050 and Tesla M2070 computing processor boards (Figure 2) conform 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 M2050 and Tesla M2070 Computing Processor Module

STANDARD I/O CONNECTOR PLACEMENT

As shown in Figure 3, the Tesla M2050 and Tesla M2070 include a vented bracket. If you are an OEM who qualifies for bracket modifications, you have the option of receiving your modules with no bracket installed.




Figure 3. Tesla M2050 and Tesla M2070 Bracket

INTERNAL CONNECTORS AND HEADERS

The Tesla M2050 and Tesla M2070 modules support 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 M2050 and Tesla M2070 modules are performance-optimized, high-end products and use power from the PCI Express connector as well as external power connectors. The boards 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.

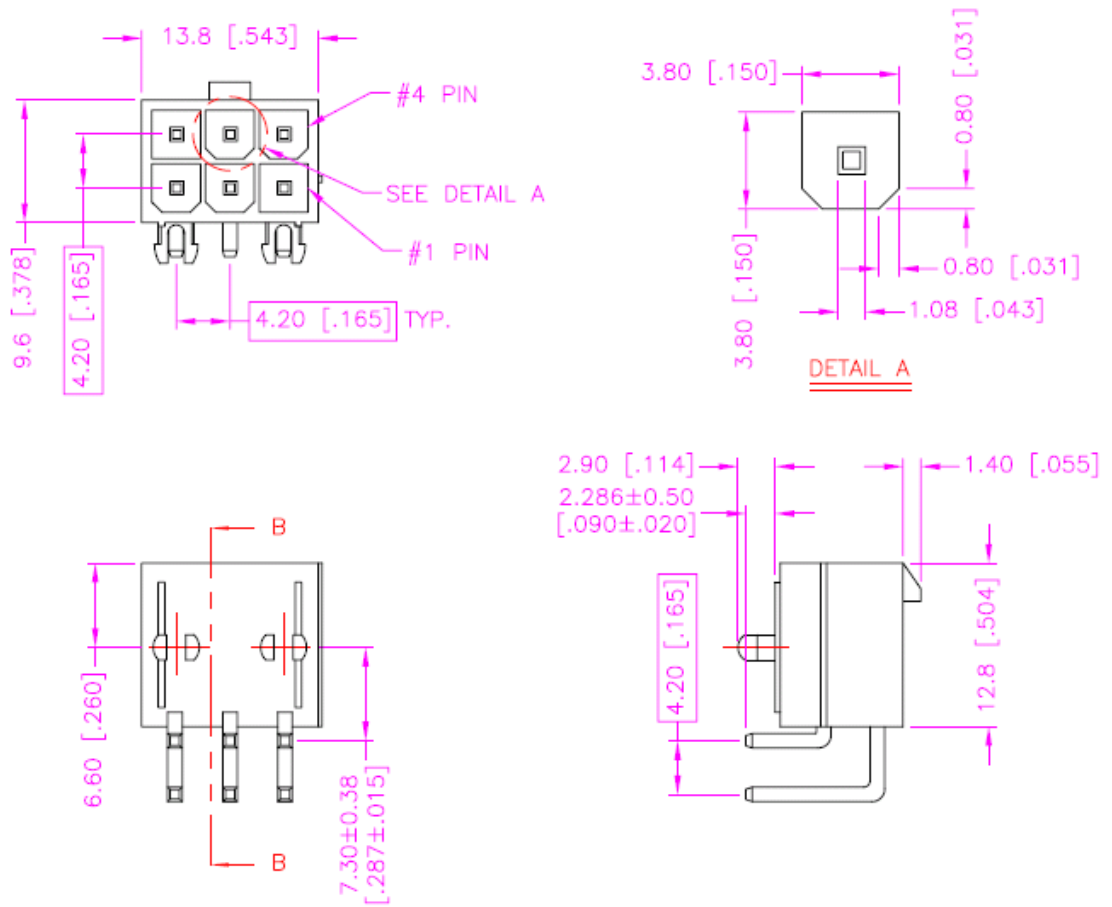


Figure 4. 6-Pin PCI Express Power Connector

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 M2050 and Tesla M2070 require 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 spec
8-pin connected	Not connected	Valid configuration - board will operate to spec
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 M2050 System Design Guide* and the *Tesla M2070 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 are ROVI-enabled 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 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

© 2010 NVIDIA Corporation. All rights reserved.