Specification

Tesla S1070
GPU Computing System

July 25, 2008 | SP-04154-001_v01

Document Change History

Version	Date	Responsible	Description of Change
0.9	July 14, 2008	GB	Initial version, preliminary
01	July 25, 2008	GB, SM	Initial release

Table of Contents

List of Figures

List of Tables

Tesla S1070 Overview

The NVIDIA® Tesla™ S1070 Computing System is a 1U rack-mount system with four Tesla T10 computing processors. This system connects to one or two host systems via one or two PCI Express cables. A Host Interface Card (HIC) is used to connect each PCI Express cable to a host. The host interface cards are compatible with both PCI Express 1× and PCI Express 2.× systems.

Key Specifications

Computing Processors

- Four Tesla T10 graphics processing units (GPUs)
- 16.0 GB of high speed memory, configured as 4.0 GB for each GPU

Mechanical Overview

- Physical Dimensions
  - System: 1.73 inches high × 17.5 inches wide × 28.5 inches deep
  - System weight without external accessories: 32 lbs
- PCI Express Cable
  - Standard: 0.5 meters in length
  - Optional: 2.0 meters in length
- Host interface Cards
  - PCI Express low profile form factor
  - Standard card requires a ×16 PCI Express slot
  - An optional card is available for ×8 PCI Express slots
- Rack Compatibility
  - Fits 4-post, 19" EIA compatible racks
  - Rack depth between posts: 28 to 32 inches
- External Connectors
  - Two cable connectors for ×16 PCI Express
  - C19 format female connector for power cord

Operating Environment

- Temperature: 5 °C to 35 °C
- Relative humidity: 10 % to 80 % non-condensing
- Maximum airflow: TBD

System Architecture

The Tesla S1070 GPU computing system is based on the T10 GPU from NVIDIA. It can be connected to a single host system via two PCI Express connections to that host, or connected to two separate host systems via one PCI Express connection to each host.

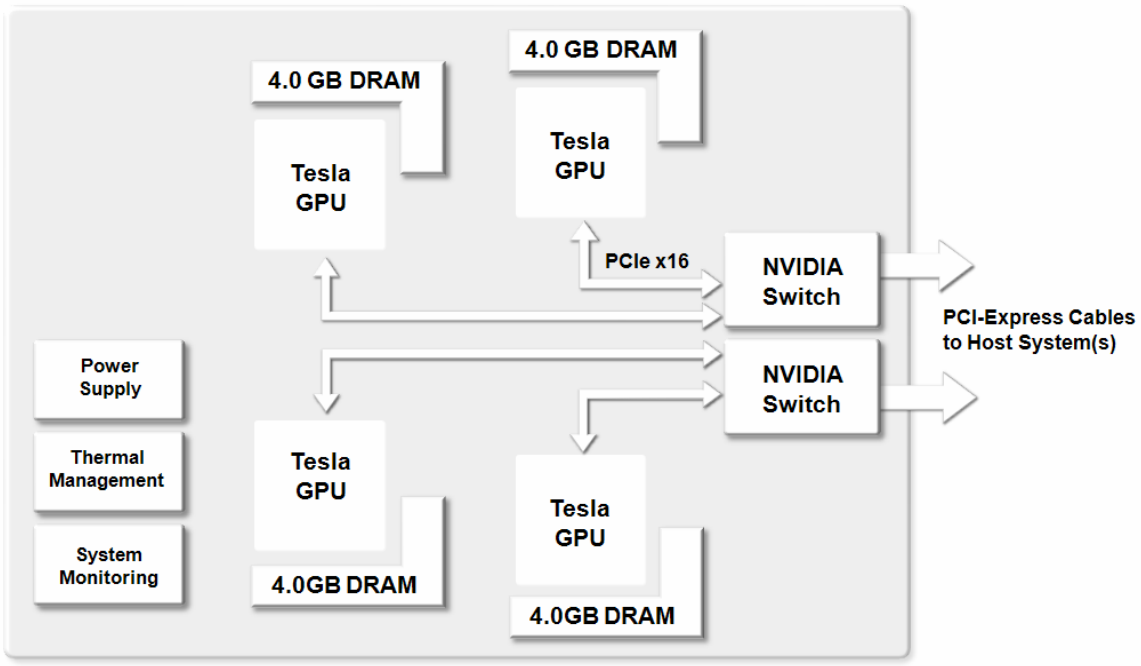

Figure 1. Tesla S1070 System Architecture

Each NVIDIA switch and corresponding PCI Express cable connects to two of the four GPUs in the Tesla S1070. If only one PCI Express cable is connected to the Tesla S1070, only two of the GPUs will be used. To connect all four GPUs in a Tesla S1070 to a single host system, the host must have two available PCI Express slots and be configured with two cables as shown in Figure 2.

Figure 2. Tesla S1070 Connection to a Single Host System

The Tesla S1070 can also be used with hosts that have only one available PCI Express slot. However, two host systems are required and should be connected as shown in Figure 3. Each host system will access two of the four GPUs.

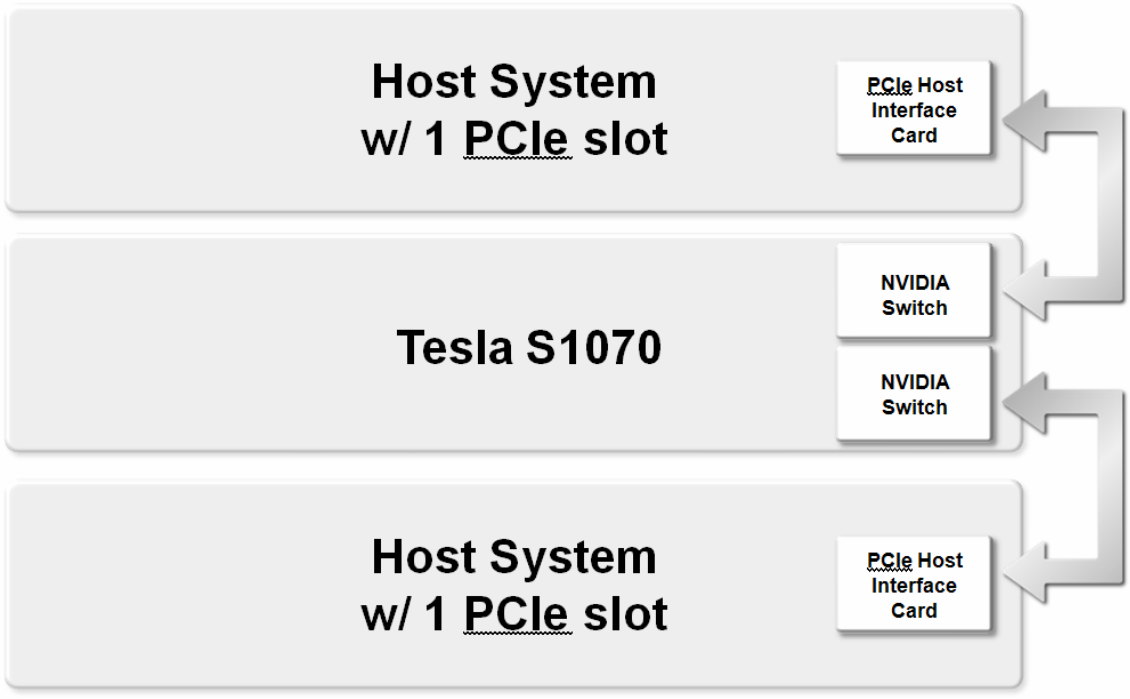

Figure 3. Tesla S1070 Connection to Two Host Systems

Configuration

There are two configurations available (Table 1) for the Tesla S1070 computing system.

Table 1. System Configuration

Specification	Description
Ordering Part Numbers	**920-20804-0001-000** (Turnkey, with standard HICs and external cables included) **920-20804-0002-000** (A La Carte, with no HICs and no cables so user can specify accessories)
GPU	T10 GPU
GPU Processor clock	1.45 GHz target (final clocks TBD)
GPU Memory clock	800 MHz (final clocks TBD)
Memory configuration	16.0 GB total configured as 4.0 GB per GPU
Memory I/O	512-bit per GPU
System I/O	Two PCIe connections. Each connection leads to two of the four GPUs.
PCI Express cables	→ A 0.5-meter cable is included in the "turnkey" kit → A 2.0-meter cable is available but must be ordered separately

System Chassis

The Tesla S1070 (Figure 4) uses a 1U form factor chassis and conforms to the EIA 310E specification for 19-inch 4-post racks with 900 mm to 1000 mm depth. The chassis dimensions are 1.73 inches high × 17.5 inches wide × 28.5 inches deep.

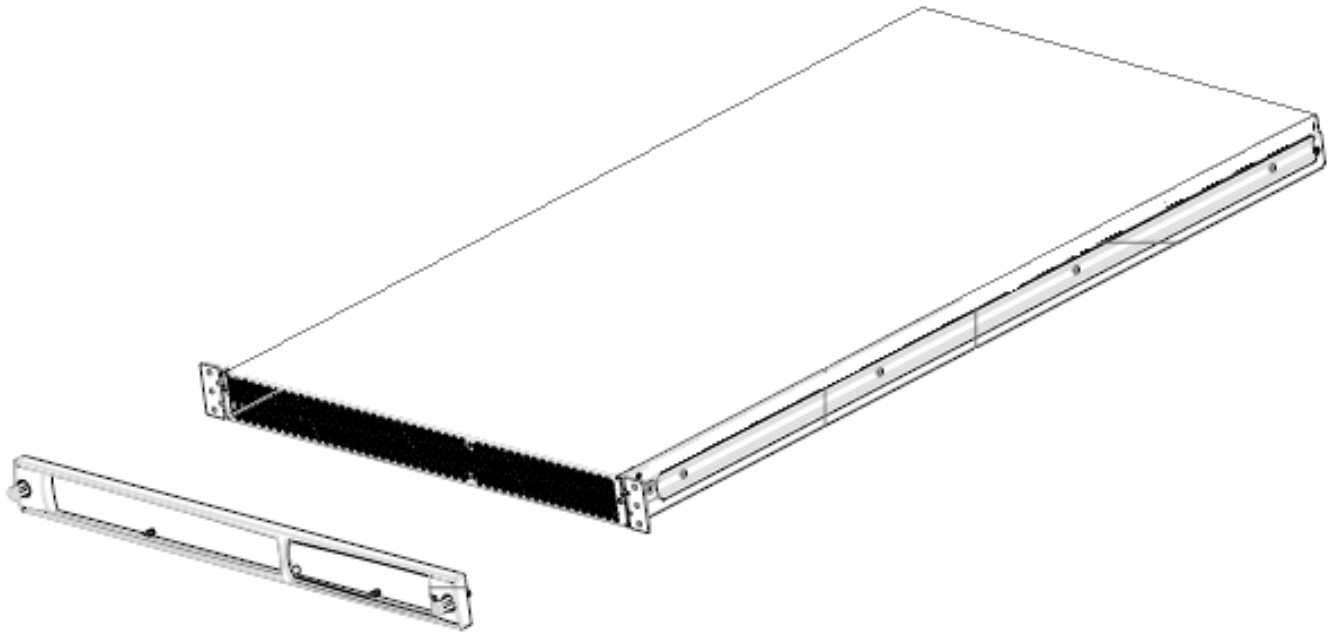

Figure 4. System Chassis Drawing

Host Interface Card (HIC)

The HIC conforms to the PCI Express low profile form factor. This card is compatible with both PCI Express 1× and PCI Express 2.0 systems. A ×8 version is also available for systems that do not have ×16 PCI Express slots. The HICs ship with a full-height bracket installed and includes a low-profile bracket.

Figure 5 shows the ×16 version of the card with the full-height bracket.

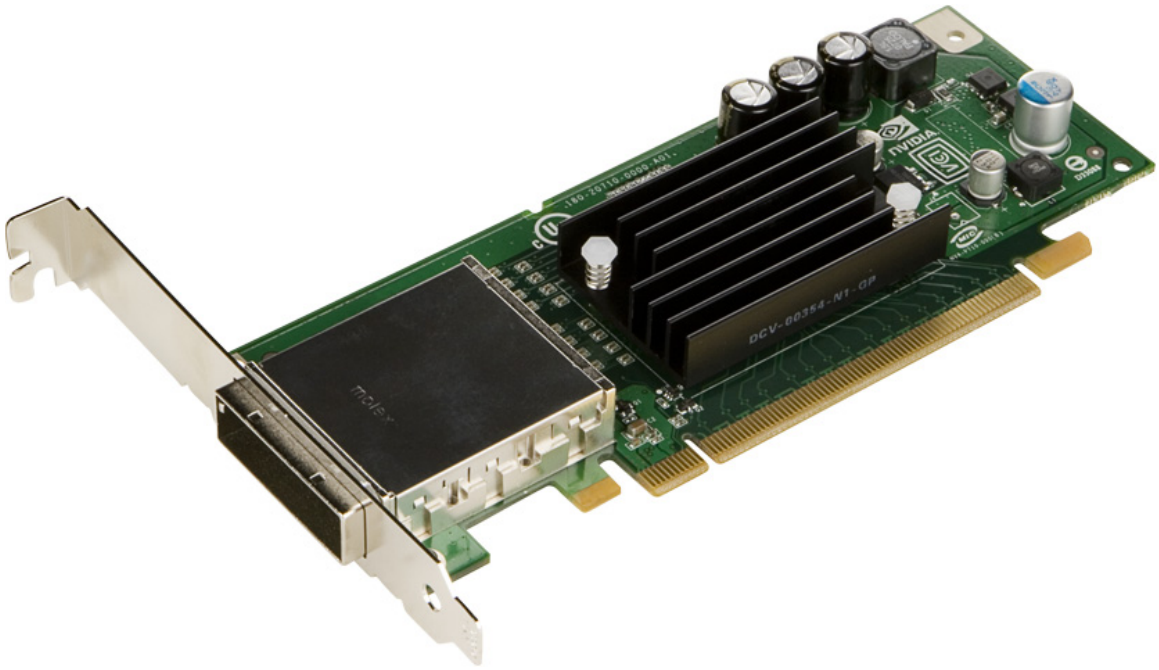

Figure 5. Host Interface Card (x16 Version)

PCI Express Cable

The Tesla S1070 uses 0.5-meter PCI Express cables as the standard connection to the host system(s). Figure 6 shows the dimensions of this cable and its connectors. A 2.0-meter version of the cable is also available as a standalone accessory and uses the same connectors as the 0.5-meter cable.

Note: For Figure 6 the dimensions are in millimeters unless otherwise labeled.

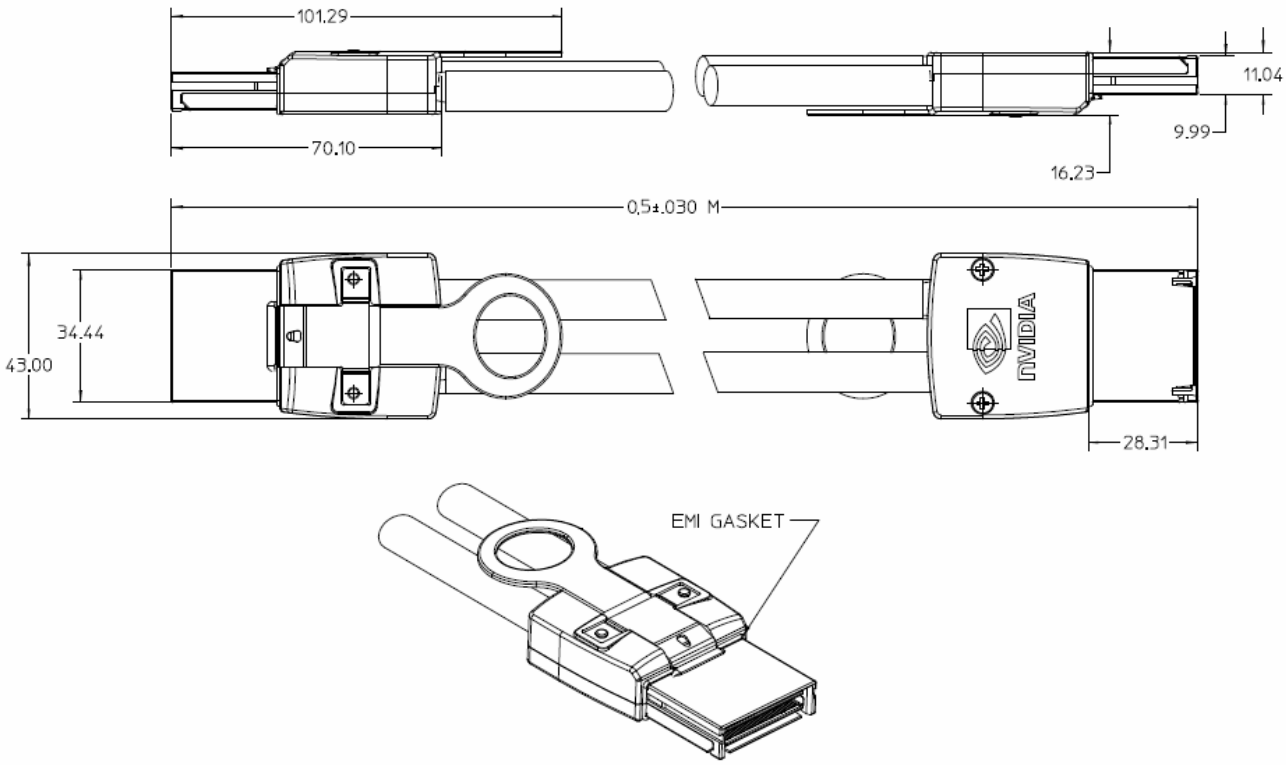

Figure 6. PCI Express Cable (0.5 Meter)

The minimum bend radius is 38.7 mm for the PCI Express cable. Figure 7 shows details of how this is measured relative to the I/O plate on the host interface card and relative to the cable/connector interface.

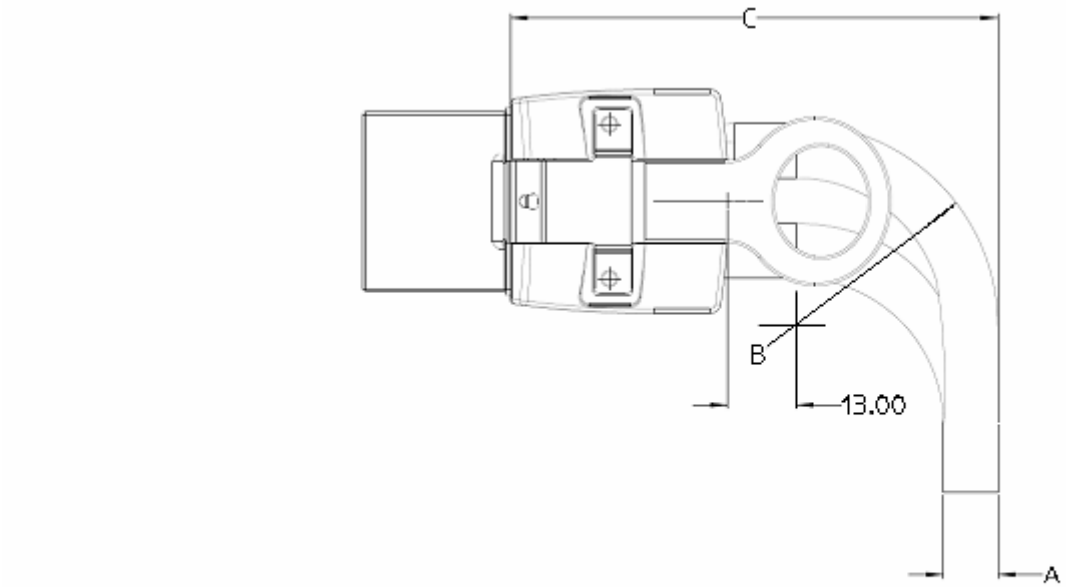

	CABLE GAGE	30
A.	CABLE DIAMETER	8.6
B.	MINIMUM OUTER RADIUS	38.7
C.	FACEPLATE TO OUTER RADIUS	94

Figure 7. PCI Express Cable Minimum Bend Radius

Thermal Testing Summary

The thermal testing results are preliminary. Thermal testing will be detailed in a thermal qualification summary report.

Table 2. Test Parameter Highlights

Parameter	Value
Fan speed	TBD
Ambient temperature	35 °C
Duration	60 minutes
Components instrumented for the test	TBD
Test applications	NVIDIA thermal stress test

Table 3. Thermal Results and Specification

Test Application	Average Power (Watts)	GPU Tjunction(°C)*	PCIe Switch Tjunction(°C)*	FET Tjunction(°C)*
NVIDIA thermal stress test	TBD	TBD	TBD	TBD

* Junction temperature as reported by NVIDIA thermal sensor. Only the highest value is reported for multiple chips. (e.g. There are four GPUs in the system, the reported value is the highest of the four.)

Table 4. Environmental Specifications and Conditions

Specifications	Conditions
Operating temperature	5 °C to 35 °C
Operating humidity	10 to 80 % RH, non-condensing
Operating altitude	5000 feet mean sea level (MSL)
Operating shock	Half sine 40g, 2 ms duration, 3 axis
Operating vibration	Sinusoidal 0.25g, 10 to 500 Hz, 3 axis. Random 1.0 Grms, 10 to 500 Hz
Acoustics	TBD dBa at 1 meter in front of system
Non-operating temperature	0 °C to 65 °C

Support Information

Languages

Language support for the Tesla 1U systems is English (U.S.) only at this time.

Certificates and Agencies

Certifications

There are no certifications planned at this time.

Agencies

- Australia Communication Authorities (C-Tick)
- Bureau of Standards, Metrology, and Inspection (BSMI)
- Conformité Européenne (CE)
- Federal Communications Commission (FCC)
- Interference-Causing Equipment Standard (ICES)
- Ministry of Information and Communication (MIC)
- Underwriters Laboratories (UL)
- Voluntary Control Council for Interference (VCCI)

NVIDIA Corporation | 2701 San Tomas Expressway | Santa Clara, CA 95050 | **www.nvidia.com**