RapidChip™ Technology Fast Custom Silicon through Platform-Based Design

White Paper

Introduction

The RapidChip platform features a fundamentally new way for engineers to rapidly design custom ICs. This robust silicon platform fills the void between the FPGA and standard-cell ASIC product spaces by combining the best attributes of both. Bringing together new technologies and methodologies, the RapidChip platform approach delivers packaged, tested, working silicon with complex embedded Intellectual Property (IP) in half the time required by current methods and at one-quarter their development costs.

The primary objectives of the RapidChip platform are:

- Dramatically reduce time-to-market for complex, highly integrated, high-performance custom ICs
- Dramatically reduce the engineering costs, CAD tool costs, and tooling costs associated with the development of deep submicron devices
- Deliver a very cost-effective production solution for use in medium-volume applications

Reducing time-to-market for complex IC development has been the main focus throughout the development of the RapidChip platform. This is done by addressing the areas in the construction of an IC which have the greatest impact to the design schedule and in particular to design schedule variability. Also, today's design methodology for deep submicron devices requires design engineers to be more involved with delivering a working chip than in delivering the function defined by the system architects. The RapidChip platform approach frees chip designers from these responsibilities, allowing them to concentrate on designing the differentiating intellectual property that adds value to the end system.

The RapidChip platform uses the LSI Logic 0.18 micron (G12®) and 0.11 micron (Gflx™) cell-based ASIC technologies. This results in near-ASIC performance, density, and power consumption. Also, the diffused

memories and IP cores within the RapidChip platform are the same as those used in LSI Logic cell-based ASIC product offerings.

The high cost of designing an ASIC is driving systems companies to look at FPGAs and ASSPs to satisfy their design requirements, even though neither solution is optimal in terms of production cost, performance, integration, or the ability to differentiate. The RapidChip platform dramatically reduces design costs by improving the efficiency of design engineers, delivering a methodology that incorporates low-cost design tools, and employing a manufacturing strategy that greatly reduces the mask and prototyping costs. Thus, the RapidChip platform technology dramatically lowers the barrier to designing high-performance, high-density custom ICs, enabling a new wave of innovation.

In determining the right solution to meet a set of system requirements, designers must take into account the cost of the end product. Today, ASICs offer the most cost-effective custom IC solution. The high cost of ASIC development might not make financial sense in low-volume production categories. Also, complex FPGAs might not be viable because of their very high production unit cost. The RapidChip platform offers low development and low production unit costs, allowing medium-volume IC users to take advantage of near-ASIC performance and density while meeting system product costs.

RapidChip Platform Elements

The following subsections detail the constituent hardware and software elements of the RapidChip platform.

RapidSlice

The fundamental technology used in the RapidChip platform is the metal customization of a partially manufactured semiconductor device (called a RapidSlice), in which all silicon layers have been fabricated. As shown in Figure 1, each RapidSlice™ incorporates diffused memory blocks, PLLs, IP blocks from the extensive LSI Logic CoreWare® library, and the configurable transistor fabric. The I/O ring is made up of configurable and dedicated I/Os for specific requirements.

Figure 1 Example of a RapidSlice Platform

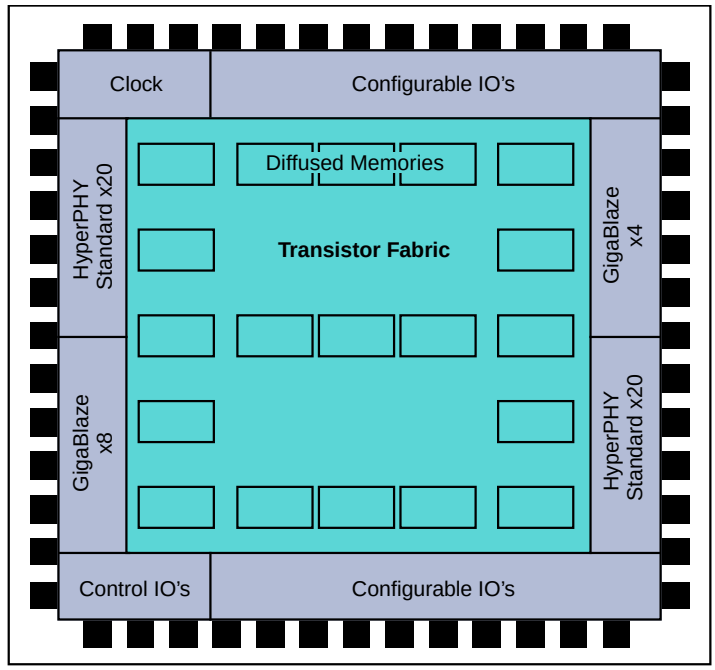

RapidSlices leverage LSI Logic systems and applications expertise in the communications, storage, and consumer markets to meet specific application needs. Custom RapidSlices also can be developed to meet specific customer needs. Each RapidSlice is configured using only metal layers and can address the needs of many different systems and applications. As a result, the up-front, non-recurring engineering (NRE) costs are much lower than in the traditional, standard-cell ASIC model, where a particular IC is used in only one application.

Figure 2 shows a StreamSlice™ layout, a RapidSlice available for the development of devices used in communications applications, such as switches and routers.

Figure 2 Example of a StreamSlice Layout

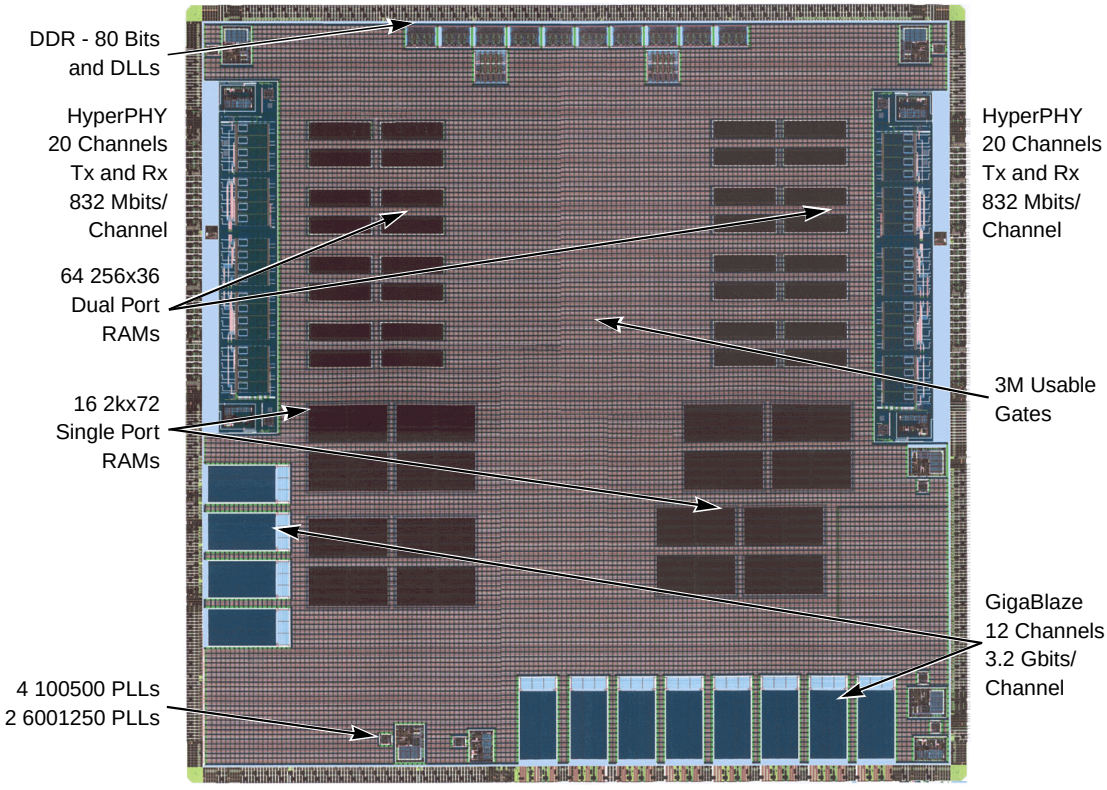

RapidChip Platform Use of the CoreWare Library

A key benefit of the RapidChip platform is that it provides convenient access to the vast LSI Logic CoreWare library. Depending on the specifications, an IP block is delivered in one of three forms: Hard, Soft, or Firm.

- Hard IP is diffused at optimal locations within a RapidSlice using cell-based elements for maximum performance and density. CoreWare examples of Hard IP include:
  - Embedded processors: ARM®, MIPS, and ZSP™ (LSI Logic Open Standard DSP)

  - HyperPHY® transceivers (600 – 3125 Mb/s): enabling standard interfaces such as SPI4.2 and SPI5, as well as backplane SerDes functionality
  - GigaBlaze® transceivers (1.0 – 4.25 Gb/s): enabling standard interfaces such as 1Gb/s and 10Gb/s Ethernet Fibre Channel, RapidIO, Infiniband, PCI Express, S-ATA, and SAS
  - Interconnect Transceivers: USB, PCI-X, HyperTransport, SGMII
- Soft IP is incorporated into the RapidChip platform technology as a functional block and implemented in the transistor fabric like any other block in the design, with specific timing criteria to ensure its functionality.
  - Microprocessor peripherals to ensure OS-readiness (interrupt controllers, UARTs, timers, GPIO, Ethernet interfaces, bus structure)
  - Link Layers for SPI 4.2, Ethernet, etc.
- Firm IP allows fully routed and characterized high-performance blocks to be incorporated into the RapidChip platform design. The blocks can be located anywhere within the transistor fabric. These are R-cell based, but can be run at a frequency exceeding the R-cell core frequency. CoreWare examples of Firm IP include:
  - MAC Link layers to support SPI 4.2, 10/100/1000 MAC
  - ARM processors
  - 10Gbit Ethernet MAC

RapidChip Platform Use of Embedded Memory Blocks

Single-port and dual-port, high-density, high-performance RAM blocks are diffused into each RapidSlice, defining a memory space that can be configured to meet a particular application need. The customization of the memory space is done using the RapidChip platform tool set. Memory blocks can be combined to create larger memories. For example, controllers can be built in the transistor fabric to create multi-port memory structures. Some of the key features that can be configured are:

- 1RW, 1R/1W port, 2RW port memories available
- Tile multiple memory blocks into wider and/or deeper memories

- Reorganize a single memory block, e.g., 512x18 from a 256x36
- Membist automatically inserted

Certain networking application-specific RapidSlice also include content addressable memory (CAM) blocks.

Configurable I/O

Each RapidSlice includes I/Os that meet certain interface standard requirements, such as SPI-4, XAUI, Fibre Channel, USB, etc. I/Os not dedicated to specific standards are configurable and support the bus types shown in Table 1.

Table 1 Standards Supported by Configurable I/Os

I/O Type	Types	Voltage	Pull-Up/-Down	Controlled Impedance
Bidirectional LVCMOS LVTTL	2, 4, 8, 12 mA	1.8, 2.5, 3.3 V	Yes	No
Bidirectional SSTL_2	Class 1, 2	2.5 V, 1.8 V	N/A	Yes
Bidirectional HSTL	Class 1, 2, 3 Single End	1.5 V, 1.8 V	N/A	Yes
LVDS	Differential Input / Output	2.5 V	N/A	No

Packaging

Each RapidSlice has an associated custom package to ensure all electrical and mechanical design requirements are met based on worst-case conditions. This package also meets the requirements for very high-speed I/Os. Depending on the number of I/Os, die size, and power consumption of a particular RapidSlice, more than one package might be offered.

Transistor Fabric and R-Cells

The transistor fabric provides the basis for the implementation of the user's logic. An R-cell is the basic unit within the transistor fabric; it is

made up of specially sized “N” and “P” type transistors for maximum flexibility and performance. R-cells are diffused in a regular pattern throughout the RapidSlice. They can be configured via metal to create the full range of logic functions available within the RapidChip platform logic cell library. The library contains over 200 cells, with a range of drive strengths. R-cells also can be efficiently configured as small memory blocks, further adding flexibility to a designer’s memory implementation. Transistors within the fabric are activated only when they are part of the implementation of a function used in the design, ensuring the most power-efficient solution.

RapidChip Platform Design Flow

The design flow for creating a device using the RapidChip platform consists of:

1. Selecting a RapidSlice that most closely meets the end-system requirements. If the required RapidSlice does not exist, LSI Logic can work with the customer to develop a new RapidSlice.
2. Selecting intellectual property (IP) that meets the end-system requirements. Customer-specific IP can also be developed if the available RapidSlices do not meet those requirements.
3. Completing the design by adding customer-specific functionality.

This design flow is supported by the methodology and tools incorporated in the RapidChip platform; these accelerate design times using a predictable design process. This incorporates best-in-class third- party tools and automation tools for the generation of design structures. Rules and constraints are also provided to guide designers through the RapidChip platform design process.

RapidChip Platform Tool Flows

Initially, LSI Logic is offering a single RapidChip platform design flow that uses industry standard ASIC design tools. This ensures that end-users efficiently move through the design process. This tool flow is shown in Figure 3, below.

A second tool flow, which LSI Logic will introduce later in 2003, dramatically reduces the cost of the EDA tools required to design a

RapidChip platform device. The design tools included in this flow will be extremely easy to use and will be capable of implementing highly complex custom RapidChip platform devices. However, they will be available at a fraction of the cost of ASIC design tools available today, further reducing the total cost of designing complex ICs.

Figure 3 RapidChip Platform Tool Flow Diagram

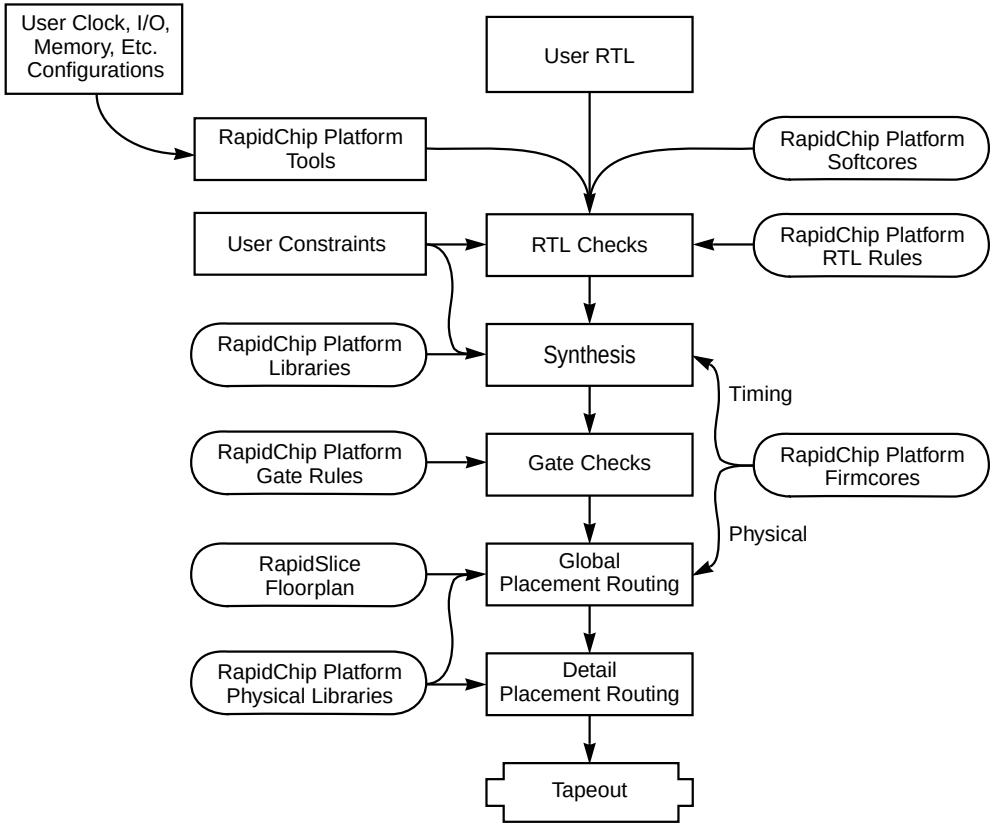

RTL Generation Tools

The RapidChip platform technology automatically generates clock, memory, test, and I/O structures, based on user inputs. This relieves the design team of tedious design tasks and dramatically reduces the overall resources required to design custom, high-performance ICs. Also, these structures are “correct by construction” and optimized to ensure ease of implementation during the physical design phase.

The RapidChip platform technology uses RTL generation tools to create design structures for the implementation of a custom IC design. A brief description of the generation tools is provided below:

- *GenMem* implements the memory on a RapidSlice based on user requirements using transistor diffused and/or fabric memories. GenMem also creates the MemBIST test structures of all the memories in the design.
- *GenIO* implements the I/Os on a RapidSlice based on the user requirements. GenIO also adds the JTAG structures for boundary test and provides SSO analysis.
- *GenClock* creates clocks based on the user requirements. Clock circuits are the most timing-critical, layout-sensitive structures within any complex IC design. Automated clock design also takes into account the requirements for controlling the clock during the testing of the device.
- *GenTest* implements the test structures for the logic design and connects the MemBIST, JTAG, and clocking structures to the TAP controller.

RapidChip Platform Rules and Constraints

To ensure the fastest design turn-around-time possible, the physical implementation of a RapidChip platform device must be accomplished in a single pass. Design iterations requiring RTL or gate level redesign due to timing issues or routing congestion found during physical design is not acceptable. Accomplishing first-pass physical design success requires that the inputs to the physical design process—the design netlist and timing constraints—must meet a set of criteria. These criteria guarantee that the required design performance is met and the design structures can be implemented within the available physical resource of the selected RapidChip platform.

As the customer RTL is being written, the RapidChip platform technology makes extensive use of rule-based and structure-based feedback and checking, isolating design issues early in the flow. RTL analysis is a key component in this strategy. Constraints are provided to the chip designer detailing RTL design rules. RTL analysis tools analyze and identify “implementation unfriendly” structures that might correctly describe the

required functionality, but which cause problems in the physical design process. Identifying these structures early in the design process allows changes to be made to the RTL, eliminating the need for iterations during the physical design phase.

Figure 4 below shows the flow used for RTL analysis.

Figure 4 RTL Analysis Flow

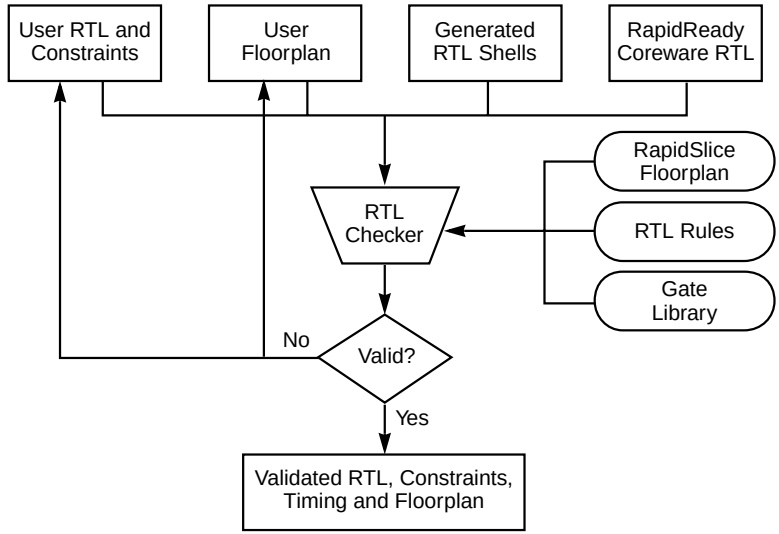

Examples of restrictions and checks are:

- Register all signals into, and out of, all major synthesis blocks. This ensures timing requirements between major functional blocks can be easily met.
- Avoid multicyle paths. When this rule is not observed, verifying the timing constraints for physical design can become a significant effort.
- Avoid flip flops clocked using the negative-edge of the clock (commonly found in memory controllers, etc.).
- Prohibit pulse generation in the design. Any signal going into the “Clock Input” of a storage element must originate from a GenClock-generated clock.
- Prohibit large multiplexing structures, such as those required for bus switching. Employ distributed multiplexing strategies.

- Asynchronous loops are not allowed.

Guidelines for synthesizing the RTL into gates also ensure the synthesis tool makes use of the most optimal RapidChip platform logic cells. Once the gate level netlist is available, the structure can be analyzed using a tool that lets engineers accurately and quickly analyze and modify the customer netlist, explore the hierarchies, and identify potential design issues. This tool runs design rule checks, generates design statistics, and lets the engineer explore design variations.

Figure 5 below shows the flow used for validated RTL to placed gates.

Figure 5 Flow – Validated RTL to Placed Gates

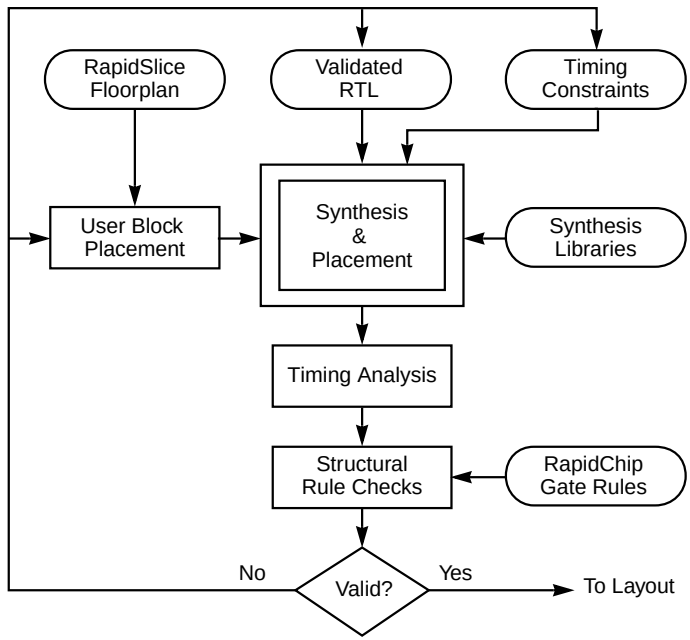

Correct-By-Construction Physical Design Methodology

The RapidChip platform single-pass design flow ensures the fastest turn around time possible. The RapidChip methodology takes into account the following deep submicron electrical phenomena that can affect the

operation of an IC through a Correct-by-Construction approach during the physical design phase:

- Electromigration (EM)

  The cell library contains no high-current drivers, eliminating the possibility of EM. EM effects caused by power grid connections are avoided by using power grid structures that have been designed to take into account the maximum possible current draw.

- Crosstalk

  The RapidChip platform physical design tools follow strict rules to automatically control slew rates, limit the net length of all signals, and provide timing margins, thus eliminating crosstalk effects.

- Static Power Droop

  Each RapidSlice design includes a fixed power grid structure designed to handle worst-case design conditions (frequency, gate count, memory size, etc.).

- Dynamic Power Droop

  Each RapidSlice includes a fixed power grid structure with integrated decoupling capacitors, guaranteeing that dynamic power droop is not a concern. Also, the clustering of higher power cells is limited during placement avoiding any possible hot spot issues.

Simple Migration Path to Cell-Based ASIC

LSI Logic can easily and quickly migrate a RapidChip platform design to a cell-based implementation with only limited support from the customer design team. This is because:

- The diffused Hard IP, PLLs, and memory structures used in the RapidChip platform design are same as those uses in an ASIC implementation. No technology requalification is required.
- The RapidChip platform design methodology is inherently more restrictive compared to what is allowed in a cell-based design methodology. Thus, any complex, high-performance design structure within a RapidChip platform design can be easily re-implemented in a cell-based ASIC with additional margin.

- The deliverables required to implement a RapidChip platform design (netlist and complete set of final timing constraints) are the same as those required to implement a cell-based ASIC.

Summary

Cell-based ASIC solutions continue to be used where the highest performance, highest density, lowest power solution is required. In these applications, end-product margins are high enough to justify the high cost of the cell-based IC development. Also, cell-based ASICs will continue to be used in high-volume/cost-sensitive applications where the high development costs are justified by a large savings in overall component cost over the life of the end-product.

FPGA solutions continue to make sense in a number of applications. One example might be where low volumes are required and production costs are not a major concern. Another example might be where time-to-market is critical, and where limited complexity, low performance and high power consumption are not critical factors in the implementation of the end system.

However, for medium-volume, complex designs with strict TTM demands, complex IP content, and low-power and low-cost development requirements, the RapidChip platform is clearly superior.

The RapidChip platform approach delivers ASIC-like performance, density, and power consumption using a design methodology that ensures FPGA-like design time at a very low development cost. The RapidChip platform enables companies to design IC-based high-performance systems at one-quarter of the development cost of a cell-based ASIC solution and one tenth the production price of an equivalent FPGA solution.

DB10-000017-00 April 2003

Headquarters

LSI Logic Corporation
North American Headquarters
Milpitas CA
Tel: 408.433.8000

LSI Logic Europe Ltd
European Headquarters
Bracknell England
Tel: 44.1344.426544
Fax: 44.1344.481039

LSI Logic K.K.
Headquarters
Tokyo Japan
Tel: 81.3.5463.7821
Fax: 81.3.5463.7820

To receive product literature, visit us at http://www.lsilogic.com.
For a current list of our distributors, sales offices, and design resource centers, view our web page located at http://www.lsilogic.com/contacts/index.html.

ISO 9000 Certified

AI

Order No.
Doc. No. DB10-000017-00