



16K x 68-bit Entry NETWORK SEARCH ENGINE

FEATURES SUMMARY

- 16K ENTRIES IN 68-BIT MODE
- TABLE MAY BE PARTITIONED INTO UP TO FOUR (4) QUADRANTS
(Data entry width in each quadrant is configurable as 34, 68, 136, or 272 bits.)
- UP TO 83 MILLION SUSTAINED SEARCHES PER SECOND IN 68-BIT and 136-BIT CONFIGURATIONS
- UP TO 41.5 MILLION SEARCHES PER SECOND IN 34-BIT and 272-BIT CONFIGURATIONS
- SEARCHES ANY SUB-FIELD IN A SINGLE CYCLE
- OFFERS BIT-BY-BIT and GLOBAL MASKING
- SYNCHRONOUS, PIPELINED OPERATION
- UP TO 31 SEARCH ENGINES CASCADABLE WITHOUT PERFORMANCE DEGRADATION
- WHEN CASCADED, THE DATABASE ENTRIES CAN SCALE FROM 124K to 992K DEPENDING ON THE SIZE OF THE ENTRY
- GLUELESS INTERFACE TO INDUSTRY-STANDARD SRAMS
- SIMPLE HARDWARE INSTRUCTION INTERFACE
- IEEE 1149.1 TEST ACCESS PORT
- OPERATING SUPPLY VOLTAGES INCLUDE:
 V_{DD} (Operating Supply Voltage) = 1.8V
 V_{DDQ} (Operating Supply Voltage for I/O) = 2.5 or 3.3V
- 272 BALL, 27mm x 27mm, CAVITY-UP BGA

Figure 1. 272-ball PBGA Package




TABLE OF CONTENTS

DESCRIPTION 6

- Overview 6
- Performance 6
- Applications 6
- Product Range (Table 1.) 6
- Switch/Router Implementation Using the M7010R (Figure 2.) 6
- Signal Names (Table 2.) 7
- Connections (Figure 3.) 8
- M7010R Block Diagram (Figure 4.) 9

MAXIMUM RATING. 10

- Absolute Maximum Ratings (Table 3.) 10

DC AND AC PARAMETERS. 11

- DC and AC Measurement Conditions (Table 4.) 11
- M7010R 2.5V AC Testing Load (Figure 5.) 12
- M7010R 2.5V Input Waveform (Figure 6.) 12
- M7010R 2.5V Output Load Equiv. (Figure 7.) 12
- M7010R 3.3V AC Testing Load (Figure 8.) 12
- M7010R 3.3V Input Waveform (Figure 9.) 12
- M7010R 3.3V Output Load Equiv. (Figure 10.) 12
- Capacitance (Table 5.) 13
- DC Characteristics (Table 6.) 13
- AC Timing Waveforms with CLK2X (Figure 11.) 14
- AC Timing Parameters with CLK2X (Table 7.) 15

OPERATION 16

- Command Bus and DQ Bus 16
- Database Entry (Data Array and Mask Array) 16
- Arbitration Logic. 16
- Pipeline and SRAM Control. 16
- Full Logic 16
- Connections Descriptions 16

CLOCKS	18
Registers	18
Clocks (Figure 12.)	18
Register Overview (Table 8.)	18
Comparand Registers	18
Comparand Register Selection During SEARCH and LEARN (Figure 13.)	19
Mask Registers	18
Addressing the Global Mask Register (GMR) Array (Figure 14.)	19
SEARCH-Successful Registers	19
SEARCH-Successful Register (SSR) Description (Table 9.)	19
The Command Register	20
Command Register Field Descriptions (Table 10.)	20
SEARCH PROCEDURE FOR 32-BIT WIDE PREFIXES	22
Global Mask Register Patterns (Figure 15.)	22
Storing left half of a Data or Mask Array (Figure 16.)	22
The Information Register	23
Information Register Field Descriptions (Table 11.)	23
The READ Burst Address Register (RBURREG)	23
READ Burst Register Description (Table 12.)	23
The WRITE Burst Address Register (WBURREG)	23
WRITE Burst Register Description (Table 13.)	23
The NFA Register	24
NFA Register (Table 14.)	24
SEARCH ENGINE ARCHITECTURE	24
Data and Mask Addressing	24
M7010R Database Configuration (Figure 17.)	25
Bit Position Match (Table 15.)	25
Multi-width Configuration Example (Figure 18.)	25
M7010R Data and Mask Array Addressing (Figure 19.)	26
COMMAND CODES AND PARAMETERS	27
Command Codes	27
Commands and Command Parameters	27
Command Codes (Table 16.)	27
Command Parameters (Table 17.)	27
READ COMMAND	28
Single Location READ Cycle Timing (Figure 20.)	29
Burst READ of the Data and Mask Arrays (BLEN = 4) (Figure 21.)	29
READ Command Parameters (Table 18.)	30
Data and Mask Array, SRAM READ Address Format (Table 19.)	30
READ Address Format for Internal Registers (Table 20.)	30
READ Address Format for Data and Mask Arrays (Table 21.)	31

WRITE COMMAND	31
Single Location WRITE Cycle Timing (Figure 22.)	32
Burst WRITE of the Data and Mask Arrays (BLEN = 4) (Figure 23.)	32
(Single) WRITE Address Format for Data and Mask Arrays or SRAM (Table 22.)	33
WRITE Address Format for Internal Registers (Table 23.)	33
WRITE Address Format for Data and Mask Array (Burst WRITE) (Table 24.)	33
SEARCH COMMAND	34
68-bit Configuration	34
Hardware Diagram for a Table with a Single Device (68-bit Operation) (Figure 24.)	34
68-Bit Configuration SEARCH Timing Diagram (One Device) (Figure 25.)	35
Right-Shift of 68-bit Signals for TLSZ Values (Table 25.)	36
Shift of SSF and SSV from SADR (for different HLAT Values) (Table 26.)	36
Latency of SEARCH from Instruction to SRAM Access Cycle (68-bit Mode) (Table 27.)	36
68-bit Logical SEARCH	37
x68 Table with One Device (Figure 26.)	37
136-bit Configuration	38
Hardware Diagram for a Table with One Device (136-bit Operation) (Figure 27.)	38
136-Bit Configuration SEARCH Timing Diagram (One Device) (Figure 28.)	39
Right-Shift of 136-bit Signals for TLSZ Values (Table 28.)	40
Shift of SSF and SSV from SADR (for different HLAT values) (Table 29.)	40
Latency of SEARCH from Instruction to SRAM Access Cycle (136-bit Mode) (Table 30.)	40
136-bit Logical SEARCH	41
x136 Table with One Device (Figure 29.)	41
272-bit Configuration	42
Hardware Diagram for a Table with One Device (272-bit Operation) (Figure 30.)	42
272-Bit Configuration SEARCH Timing Diagram (One Device) (Figure 31.)	43
Right-Shift of 272-bit Signals for TLSZ Values (Table 31.)	44
Shift of SSF and SSV from SADR (for different HLAT Values) (Table 32.)	44
Latency of SEARCH from Instruction to SRAM Access Cycle (272-bit Mode) (Table 33.)	44
272-bit Logical SEARCH	45
x272 Table with One Device (Figure 32.)	45
Mixed-sized Searches on Tables Configured with Different Width Using an M7010R Device	46
Multiwidth Configuration Example (Figure 33.)	46
Timing Diagram for Mixed SEARCH (One Device) (Figure 34.)	47
LRAM and LDEV Description	48
LEARN COMMAND	48
LEARN Command Timing Diagram (TLSZ = 00) (Figure 35.)	49
LEARN Timing Diagram (TLSZ = 1, except on Last Device) (Figure 36.)	50
LEARN Timing Diagram on Device Number 7 (TLSZ = 01) (Figure 37.)	51
SRAM WRITE Cycle Latency from Second Cycle of LEARN Instruction (Table 34.)	51

DEPTH-CASCADING	52
Depth-Cascading Up to Eight Devices (One Block)	52
Depth-Cascading Up to 31 Devices (4 Blocks)	52
Depth-Cascading to Generate a “FULL” State for a Block	52
Depth-Cascading to Form a Single Block (8 Devices) (Figure 38.)	53
Four Blocks (31 Devices Cascaded) SEARCH, 68-bit Configured with LDEV = 1 (Figure 39.) ...	54
“FULL” State Generation in a Cascaded Table (Figure 40.)	55
ARBITRATION	56
Timing Diagram for Arbitration Within a Block (Figure 41.)	56
Timing for Arbitration for Two or More Blocks for the Last Device (Figure 42.)	57
SRAM ADDRESSING	58
SRAM PIO Access	58
SRAM READ Access for One M7010R Device (Figure 43.)	59
SRAM WRITE Access for One M7010R Device (Figure 44.)	61
SRAM Bus Address Generation (Table 35.)	61
Right-Shift of SRAM Signals for TLSZ Values (Table 36.)	62
Right-Shift of SRAM Signals for HLAT Values (Table 37.)	62
JTAG (1149.1) TESTING	62
Test Access Port Controller Instructions (Table 38.)	62
TAP Device ID Register (Table 39.)	62
POWER DISTRIBUTION GUIDELINE	63
Network Search Engine Power Distribution (Figure 45.)	63
PART NUMBERING	64
PACKAGE MECHANICAL INFORMATION	65
REVISION HISTORY	66

M7010R

DESCRIPTION

Overview

The M7010R is a feature-rich, TCAM-based hardware search engine optimized for networking and communications applications. It incorporates leading-edge Associative Processing Technology (APT, trademark of Cypress Semiconductor, Inc.) and Advanced Power Management. The data table may be partitioned into up to four (4) quadrants, allowing the user to configure each quadrant with different table entry widths (x34, x68, x136, or x272-bit). It is also programmable to accelerate performance.

Performance

The M7010R outperforms competitive solutions using software sequential search algorithms in conjunction with SRAMs or ASICs, or hardware implementation with ASICs and CAMs. The latter solution, while faster than a software-based solu-

tion, still suffers from performance degradation when depth-cascaded and is unable to scale to next-generation requirements. The M7010R-based solutions overcome all of these drawbacks.

Applications

The performance and features of the M7010R makes it ideal in applications such as enterprise LAN switches, broadband switching and routing equipment, supporting multiple data rates from OC-48 and beyond.

Figure 2 illustrates how a search engine subsystem can be optimized using a host bridge ASIC, the M7010R, and synchronous or non-synchronous SRAMs. It also illustrates how this system fits into a switch-router implementation.

Table 1. Product Range

Part Number	Operating Supply Voltage	Operating I/O Voltage	Speed
M7010R-083ZA1	1.8V	2.5 or 3.3V	83MHz
M7010R-066ZA1	1.8V	2.5 or 3.3V	66MHz

Figure 2. Switch/Router Implementation Using the M7010R




Table 2. Signal Names

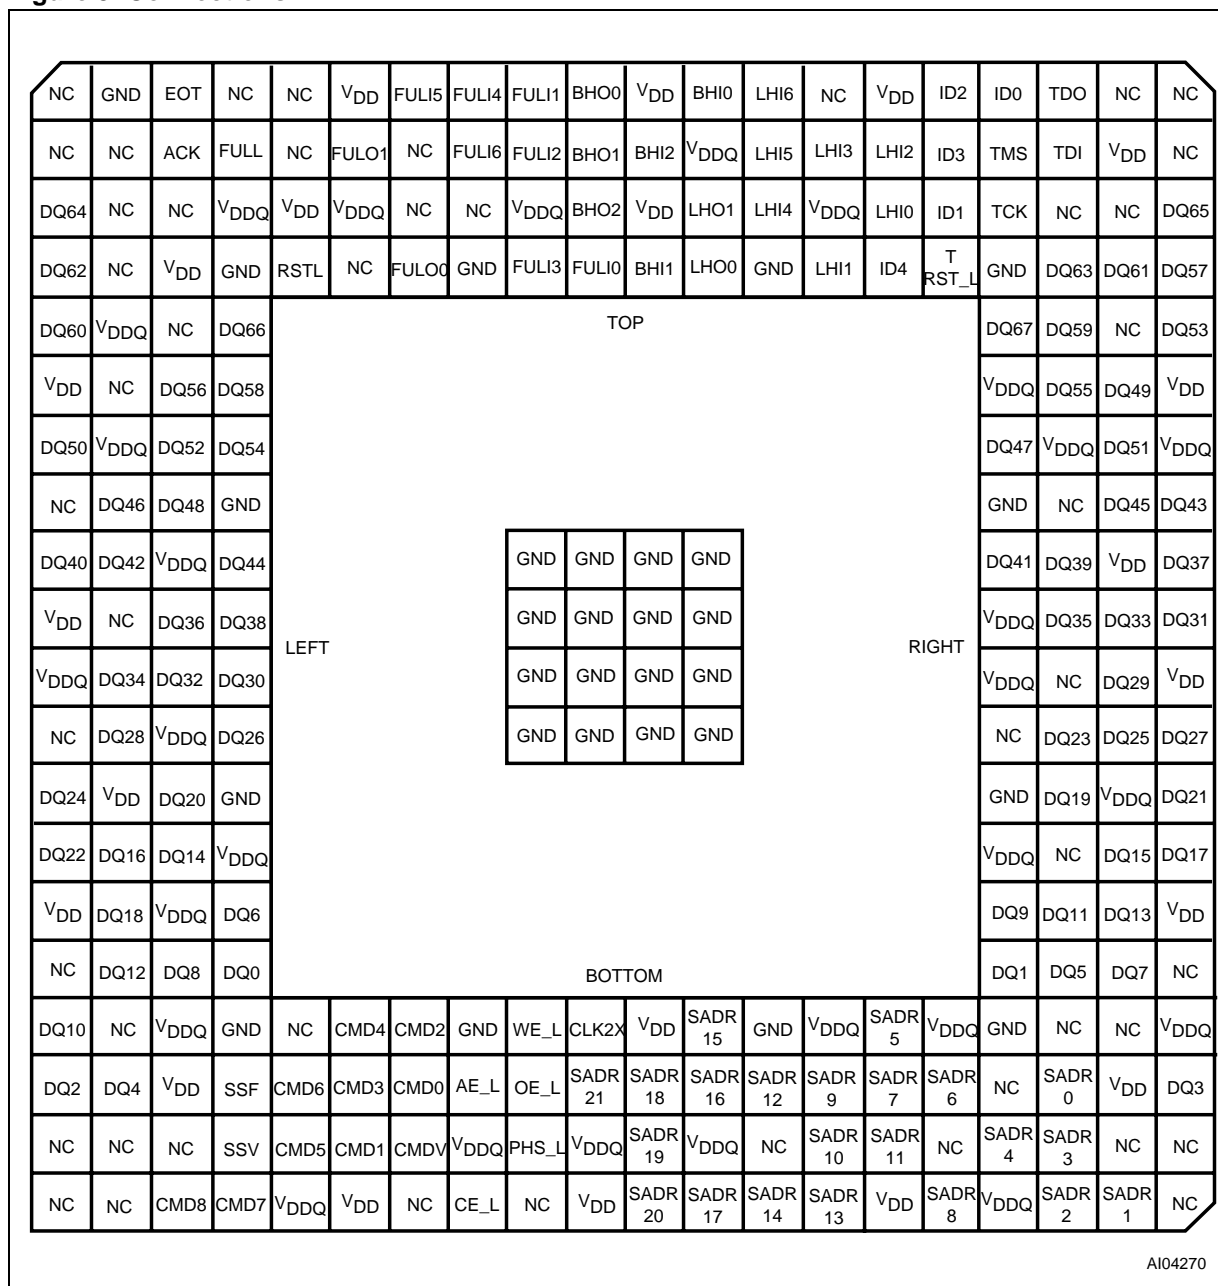
Symbol	Type	Connection Name
Clocks and Reset		
CLK2X	I	Master Clock
PHS_L	I	Phase
RST_L	I	Reset
Command and DQ Bus		
CMD[8:0]	I	Command Bus
CMDV	I	Command Valid
DQ[67:0]	I/O	Address/Data Bus
ACK ⁽¹⁾	T	READ Acknowledge
EOT ⁽¹⁾	T	End of Transfer
SSF	T	SEARCH Successful Flag
SSV	T	SEARCH Successful Flag Valid
SADR[21:0]	T	SRAM Address
CE_L	T	SRAM Chip Enable
WE_L	T	SRAM WRITE Enable
OE_L	T	SRAM Output Enable
ALE_L	T	Address Latch Enable

Cascade Interface		
LHI[6:0]	I	Local Hit In
LHO[1:0]	O	Local Hit Out
BHI[2:0]	I	Block Hit In
BHO[2:0]	O	Block Hit Out
FULI[6:0]	I	Full In
FULO[1:0]	O	Full Out
FULL	O	Full Flag
Device Identification		
ID[4:0]	I	Device Identification
Test Access Port		
TDI	I	Test Access Port's Test Data In
TCK	I	Test Access Port's Test Clock
TDO	T	Test Access Port's Test Data Out
TMS	I	Test Access Port's Test Mode Select
TRST_L	I	Test Access Port's Reset

Note: Signal types are: I = Input only; I/O = Input or Output; O = Output; and T = Tristate

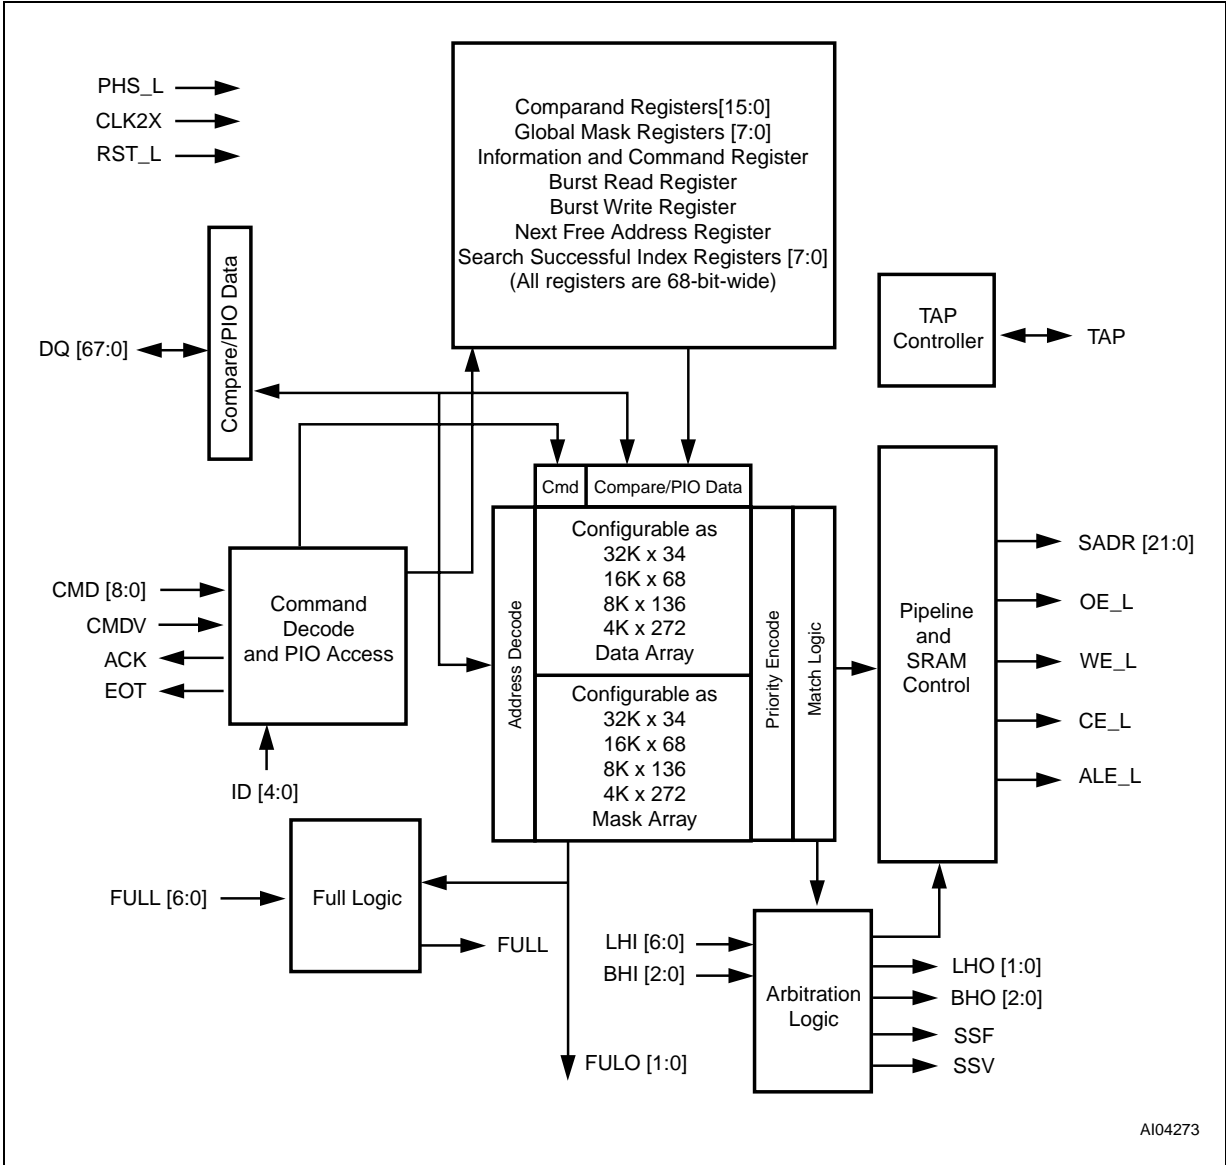
1. ACK and EOT Signals require a pull-down resistor of 47 ohms.

Figure 3. Connections



Note: This diagram is TOP VIEW perspective (view through package).

Figure 4. M7010R Block Diagram



MAXIMUM RATING

Stressing the device above the rating listed in the “Absolute Maximum Ratings” table may cause permanent damage to the device. These are stress ratings only and operation of the device at these or any other conditions above those indicated in the Operating sections of this specification is

not implied. Exposure to Absolute Maximum Rating conditions for extended periods may affect device reliability. Refer also to the STMicroelectronics SURE Program and other relevant quality documents.

Table 3. Absolute Maximum Ratings

Symbol	Parameter	Value	Unit
T _{STG}	Storage Temperature (V _{DD} Off)	–0 to 70	°C
T _{SLD} ⁽¹⁾	Lead Solder Temperature for 10 seconds	235	°C
V _{DDQ}	Input or Output Voltages	3.3	V
V _{DD}	Supply Voltage	–0.4 to 2.7	V
I _O	Output Current	100	mA
P _D	Power Dissipation	< 3	W

Note: 1. Soldering temperature not to exceed 260°C for 10 seconds (total thermal budget not to exceed 150°C for longer than 30 seconds).

DC AND AC PARAMETERS

This section summarizes the operating and measurement conditions, as well as the DC and AC characteristics of the device. The parameters in the following DC and AC Characteristic tables are derived from tests performed under the Measure-

ment Conditions listed in the relevant tables. Designers should check that the operating conditions in their projects match the measurement conditions when using the quoted parameters.

Table 4. DC and AC Measurement Conditions

Sym	Parameter	M7010R 2.5V	M7010R 3.3V	Units
V _{DD}	V _{DD} Operating Supply Voltage	1.7 to 1.9	1.7 to 1.9	V
V _{DDQ}	V _{DDQ} Voltage for I/O	2.4 to 2.6	3.1 to 3.5	V
t _A	Ambient Operating Temperature	0 to 70	0 to 70	°C
C _L	Load Capacitance	6	6	pF
V _{IH}	Input High Voltage ⁽¹⁾	1.7 to V _{DDQ} + 0.3	2.0 to V _{DDQ} + 0.3	V
V _{IL}	Input Low Voltage ⁽²⁾	-0.3 to 0.7	-0.3 to 0.8	V
	Supply Voltage Tolerance	±5	±5	%
t _R , t _F	Input Rise and Fall Times (at 0.3V and 2.7V)	≤ 2 (see Figure 6, page 12)	≤ 2 (see Figure 9, page 12)	ns
	Input Timing Reference Levels	1.25	1.5	V
	Output Timing Reference Levels	1.25	1.5	V
	Input Pulse Voltages	GND to 2.5	GND to 3.3	V
	Input and Output Timing Ref. Voltages	(see Figure 7, page 12)	(see Figure 10, page 12)	V

Note: 1. Maximum allowable applies to overshoot only (V_{DDQ} is 3.3V supply).

2. Minimum allowable applies to undershoot only.

M7010R

Figure 5. M7010R 2.5V AC Testing Load




Figure 8. M7010R 3.3V AC Testing Load




Figure 6. M7010R 2.5V Input Waveform

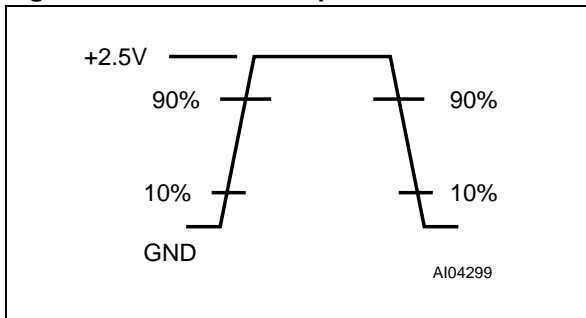


Figure 9. M7010R 3.3V Input Waveform




Figure 7. M7010R 2.5V Output Load Equiv.




Figure 10. M7010R 3.3V Output Load Equiv.




Table 5. Capacitance

Symbol	Parameter	Test Condition	Min	Max	Unit
C _{IN}	Input Capacitance	V _{IN} = 0V		6	pF
C _{IO} ⁽¹⁾	Input / Output Capacitance	V _{OUT} = 0V		6	pF

Note: Effective capacitance measured with power supply. Sampled only, not 100% tested.

1. Outputs deselected.

Table 6. DC Characteristics

Symb	Parameter	Test Condition ⁽¹⁾	Min	Max	Unit	
I _{LI}	Input Leakage Current	V _{DDQ} = V _{DDQ} (max) 0V ≤ V _{IN} ≤ V _{DDQMAX}	-10	+10	μA	
I _{LO}	Output Leakage Current	V _{DDQ} = V _{DDQ} (max) 0V ≤ V _{OUT} ≤ V _{DDQMAX}	-10	+10	μA	
I _{DD1}	1.8V Supply Current @ V _{DDMAX}	M7010R	I _{OUT} = 0mA, 83MHz Search		1250	mA
		M7010R	I _{OUT} = 0mA, 66MHz Search		1000	mA
I _{DD2}	2.5V Supply Current @ V _{DDMAX}	M7010R	I _{OUT} = 0mA, 83MHz Search		180	mA
		M7010R	I _{OUT} = 0mA, 66MHz Search		150	mA
I _{DD3}	3.3V Supply Current @ V _{DDMAX}	M7010R	I _{OUT} = 0mA, 83MHz Search		300	mA
		M7010R	I _{OUT} = 0mA, 66MHz Search		240	mA
V _{IL}	Input Low Voltage		-0.3	0.8	V	
V _{IH}	Input High Voltage		2.0	V _{DDQ} + 0.3	V	
V _{OL}	Output Low Voltage	V _{DDQ} = V _{DDQ} (min) I _{OL} = 8mA		0.4	V	
V _{OH}	Output High Voltage	V _{DDQ} = V _{DDQ} (min) I _{OH} = 8mA	2.4		V	

Note: 1. Valid for Ambient Operating Temperature: T_A = 0 to 70°C; V_{DD} = 1.8V.

Figure 11. AC Timing Waveforms with CLK2X

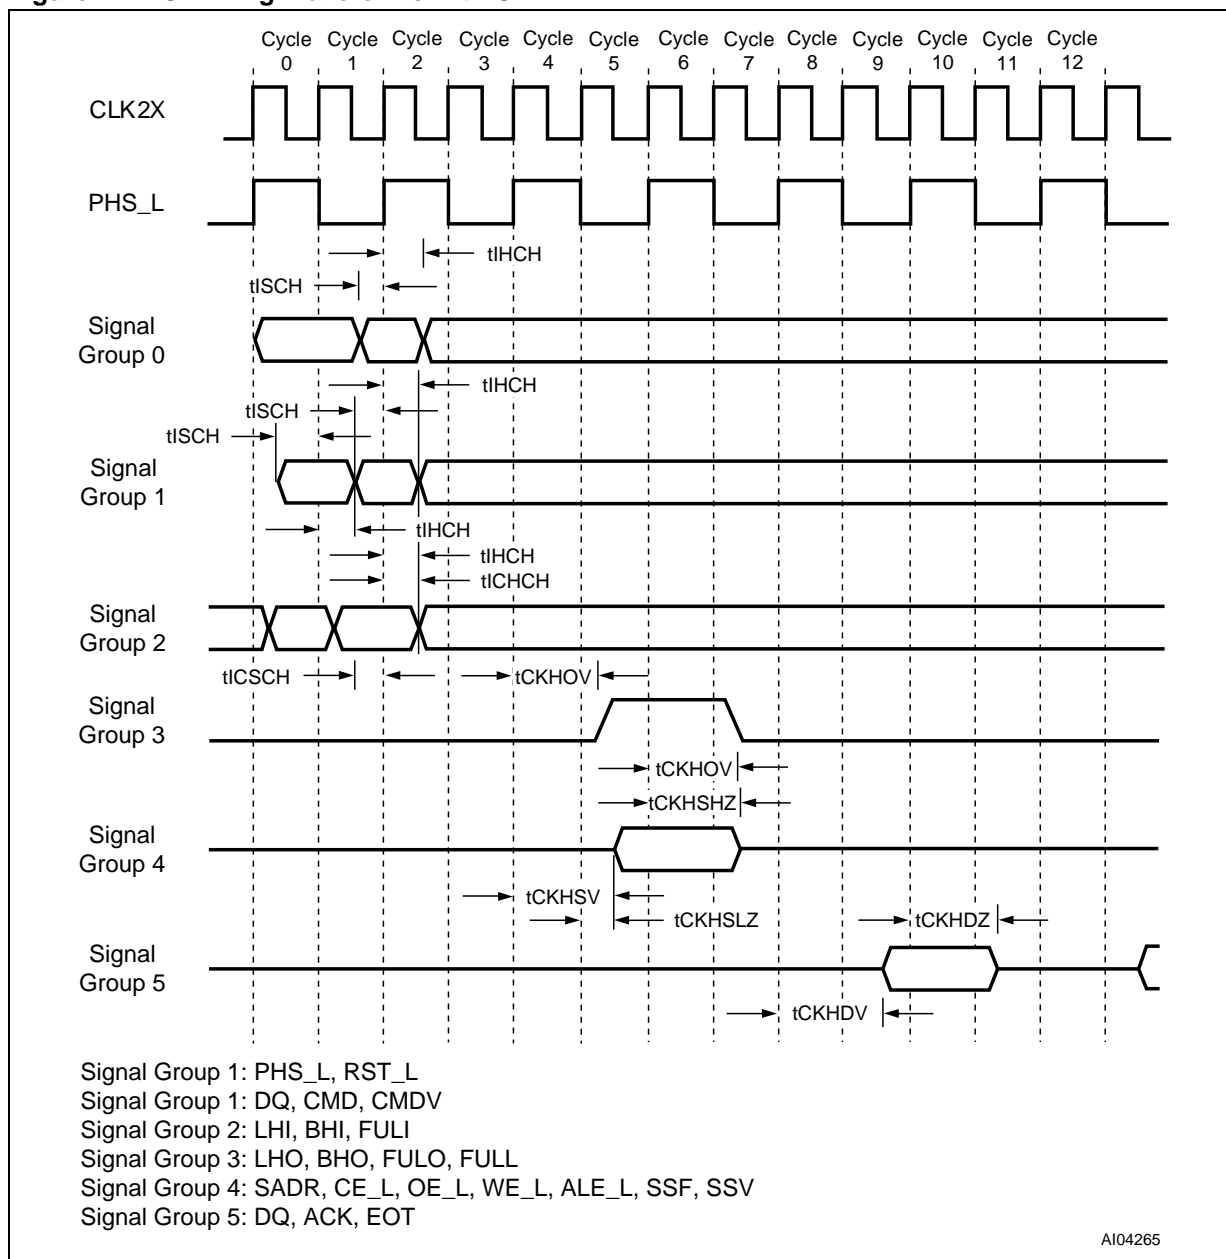


Table 7. AC Timing Parameters with CLK2X

Row	Symbol	M7010R-066		M7010R-083		Unit	Description ⁽¹⁾
		Min	Max	Min	Max		
1	f _{CLK}		133		166	MHz	CLK2X frequency
2	t _{CLK}	7.5		6.0		ns	CLK2X period
3	t _{CKHI}	3.0		2.4		ns	CLK2X high pulse ⁽²⁾
4	t _{CKLO}	3.0		2.4		ns	CLK2X low pulse ⁽²⁾
5	t _{ISCH}	2.5		1.8		ns	Input Setup Time to CLK2X rising edge ⁽²⁾
6	t _{IHCH}	0.6		0.6		ns	Input Hold Time to CLK2X rising edge ⁽²⁾
7	t _{ICSCH}	4.2		3.5		ns	Cascaded Input Setup Time to CLK2X rising edge ⁽²⁾
8	t _{ICHCH}	0.6		0.6		ns	Cascaded Input Hold Time to CLK2X rising edge ⁽²⁾
9	t _{CKHOV}		8.5		7.0	ns	Rising edge of CLK2X to LHO, FULO, BHO, FULL valid ⁽³⁾
10	t _{CKHDV}		9.0		7.5	ns	Rising edge of CLK2X to DQ valid ⁽⁴⁾
11	t _{CKHDZ}		8.5		7.0	ns	Rising edge of CLK2X to DQ high-Z ⁽⁵⁾
12	t _{CKHSV}		9.0		7.5	ns	Rising edge of CLK2X to SRAM Bus valid ⁽⁴⁾
13	t _{CKHSHZ}		6.5		6.0	ns	Rising edge of CLK2X to SRAM Bus high-Z ^(4,5)
14	t _{CKHSLZ}	7.0		6.5		ns	Rising edge of CLK2X to SRAM Bus low-Z ^(4,5)

Note: 1. Valid for Ambient Operating Temperature: T_A = 0 to 70°C; V_{DD} = 1.8V.

2. Values are based on 50% signal levels.

3. Based on an AC load of CL = 50pF (see Figure 5, page 12 and Figure 8, page 12).

4. Unless otherwise noted, all values are based on AC load of CL = 50pF (see Figure 5, page 12 and Figure 8, page 12).

5. These parameters are sampled and not 100% tested.

OPERATION

Command Bus and DQ Bus

CMD[8:0] carries the command and its associated parameter. DQ[67:0] is used for data transfer to, and from the data base entries. The database entries are comprised of a data field and a mask field which are organized as a data array and a mask array. The DQ Bus carries the SEARCH data during the SEARCH command as well as the address and data during Pipelined I/O (PIO) READ/WRITE operations, of the data array, mask array, and internal registers. The DQ Bus also can carry the address information for the PIO accesses to the SRAM.

Database Entry (Data Array and Mask Array)

Each database entry comprises a data field and a mask field. The resultant value of the entry is a logical AND of the corresponding data and mask bits and can take logical values of '1,' '0' and 'X' (don't care), depending on the value in the mask bit. The on-chip priority encoder selects the first matching entry in the database which is nearest to location 0.

Arbitration Logic

When multiple (Silicon) Search Engines are cascaded to create large databases, the data being searched is presented to all Search processors simultaneously in the cascaded system. When more than one device has duplicate entries, the arbitration logic on the Search Engine with the matching entry which is closest to address 0 of the cascaded database, will be selected to drive the SRAM Bus.

Pipeline and SRAM Control

Pipeline latency is added to give enough time to the arbitration logic in a cascaded system to determine the index with the highest priority. The pipeline logic adds latency to the SRAM access cycles and the SSF and SSV signals to align them to the host ASIC receiving the associated data. Refer to Table 27, page 36 for details.

Full Logic

Bit[0] in each of the 68-bit entries has a special purpose for the LEARN command (0 = empty, 1 = full). When all the data entries have Bit[0] set to '1,' the database asserts the FULL flag, indicating that all the Search Engines in the depth-cascaded array are full.

Connections Descriptions

Master Clock (CLK2X). The M7010R samples all of the control and data signals on the positive edge of CLK2X when PHS_L is low.

Phase (PHS_L). This signal runs at half the frequency of CLK2X and generates an internal clock from CLK2X (see Figure 12, page 18).

Reset (RST_L). Driving RST low initializes the device to a known state.

Command Bus (CMD[8:0]). [1:0] specifies the command; [8:2] contains the command parameters. The descriptions of individual commands explains the details of the parameters. The encoding of commands based on the [1:0] field are:

- 00: PIO READ
- 01: PIO WRITE
- 10: SEARCH
- 11: LEARN

Command Valid (CMDV). Qualifies the CMD bus as follows:

- 0: No Command
- 1: Command

Address/Data Bus (DQ[67:0]). Carries the READ and WRITE address as well as the data during register, data, and mask array operations. It carries the compare data during SEARCH operations. It also carries the SRAM address during SRAM PIO accesses.

READ Acknowledge (ACK). Indicates that valid data is available on the DQ Bus during register, data, and mask array READ operations, or the data is available on the SRAM data bus during SRAM READ operations.

Note: ACK Signals require a pull-down resistor of 47Ω.

End of Transfer (EOT). Indicates the end of burst transfer during READ or WRITE burst operations.

Note: EOT Signals require a pull-down resistor of 47 ohms.

SEARCH Successful Flag (SSF). When asserted, this signal indicates that the device is the global winner in a SEARCH operation.

SEARCH Successful Flag Valid (SSV). When asserted, this signal qualifies the SSF signal.

SRAM Address (SADR[21:0]). This bus contains address lines to access off-chip SRAMs that contain associative data. See Table 35, page 61 for the details of the generated SRAM address.

SRAM Chip Enable (CE_L). This is Chip Enable control for external SRAMs. When more than one device is cascaded, CE_L of all devices must be connected.

SRAM WRITE Enable (WE_L). This is WRITE Enable control for external SRAMs. When more than one device is cascaded, WE_L of all devices must be connected.

SRAM Output Enable (OE_L). This is Output Enable control for external SRAMs. Only the last device drives this signal (with the LRAM Bit set).

Address Latch Enable (ALE_L). When this signal is low, the addresses on the SRAM address bus have been validated. When more than one device is cascaded, the ALE_L of all devices must be connected.

Local Hit In (LHI[6:0]). These pins depth-cascade the device to form a larger table size. One signal of this bus is connected to the LHO[1] or LHO[0] of each of the upstream devices in a block. Connect all unused LHI pins to a logic '0.' (For more information, see DEPTH-CASCADING, page 52.)

Local Hit Out (LHO[1:0]). LHO[1] and LHO[0] are the same logical signal. LHO[1] or LHO[0] is connected to one input of the LHI bus of up to four downstream devices (in a block that contains up to eight devices; for more information, see DEPTH-CASCADING, page 52.)

Block Hit In (BHI[2:0]). Inputs from the previous BHO[2:0] are tied to the BHI[2:0] of the current device (see DEPTH-CASCADING, page 52). In a four-block system, the last block can contain only seven devices because the ID code 11111 is used for broadcast access.

Block Hit Out (BHO[2:0]). Outputs from the current device are connected to the BHI[2:0] of the next device (see DEPTH-CASCADING, page 52).

Full In (FULI[6:0]). Each signal in this bus is connected to FULO[0] or FULO[1] of an upstream device to generate the FULL flag for the depth-

cascaded block. For more information, see DEPTH-CASCADING, page 52 to Generate Full for a Block Section.

Full Out (FULO[1:0]). FULO[1] and FULO[0] are the same logical signal. One of these two signals must be connected to the FULI of up to four downstream devices in a depth-cascaded table. Bit [0] in the data array indicates if the entry is full (1) or empty (0). This signal is asserted if all of the bits in the data array are '1s.' Refer to Depth-Cascading to Generate a "FULL" State for a Block, page 52.

Full Flag (FULL). When asserted, this signal indicates that the table consisting of many depth-cascaded devices is full.

Device Identification (ID[4:0]). The binary-encoded device ID for a depth-cascaded system starts at 00000 and goes up to 11110. 11111 is reserved for a special broadcast address that selects all cascaded (silicon) Search Engines in the system. On a broadcast read-only, the device with the LDEV Bit set to '1' responds.

Test Data In (TDI). This is the Test Access Port's Test Data In.

Test Clock (TCK). This is the Test Access Port's Test Clock.

Test Data Out (TDO). This is the Test Access Port's Test Data Out.

Test Mode Select (TMS). This is the Test Access Port's Test Mode Select.

Test Reset (TRST_L). This is the Test Access

CLOCKS

The M7010R receives a Clock (CLK2X) signal and Phase (PHS_L) signal. The Phase (PHS_L) divides the CLK2X signal to generate the Internal Clock (CLK), as shown in Figure 12. The CLK2X and CLK signals are used for internal operations.

Registers

All the M7010R registers are 68 bits wide. The M7010R contains 32 comparand storage registers, 16 global mask registers, 8 SEARCH-successful index registers, command, information, burst READ, burst WRITE, and next free address registers. Table 8 provides a register overview of all the registers. The registers are ordered in ascending address order.

Comparand Registers

The device contains thirty-two 68-bit comparand registers dynamically selected in every SEARCH operation to store the comparand presented on the DQ Bus. The LEARN command will also use these registers when it is executed. The M7010R stores the SEARCH command's "Cycle A" comparand in the even-number register and the "Cycle B" comparand in the odd-numbered register, as shown in Figure 13, page 19.


Mask Registers

The device contains sixteen (8 pairs) 68-bit global mask registers dynamically selected in every SEARCH operation to select the SEARCH sub-field (see Figure 14, page 19). The three-bit GMR Index supplied on the CMD bus applies eight pairs of global masks during the SEARCH and WRITE operations, also shown in Figure 14.

Note: In 68-bit SEARCH and WRITE operations, the host ASIC must program the even and odd mask register with the same values, and the M7010R uses even-numbered mask registers as global masks.

Each mask bit in the global mask registers is used during SEARCH and WRITE operations. In SEARCH operations, setting the Mask Bit to '1' enables compares; setting the Mask Bit to '0' disables compares (forced match) at the current bit position. In WRITE operations to the data or mask array, setting the Mask Bit to '1' enables WRITES; setting the Mask Bit to '0' disables WRITES at the corresponding bit position.

Figure 12. Clocks



Note: Any reference to "CLK Cycles" means 2 cycles of the signal, "CLK2X."
 1. "CLK" is an internal signal. The period for this clock is specified in Table 7, page 15.

Table 8. Register Overview

Address	Abbreviation	Type	Name
0–31	COMP0–31	R	32 Comparand Registers. Stores comparands from the DQ Bus for learning later.
32–47	MASKS	RW	16 Global Mask Registers Array.
48–55	SSR0–7	R	8 SEARCH Successful Index Registers.
56	COMMAND	RW	Command Register.
57	INFO	R	Information Register.
58	RBURREG	RW	Burst READ Register.
59	WBURREG	RW	Burst WRITE Register.
60	NFA	R	Next Free Address Register.
61–63	–	–	Reserved

Figure 13. Comparand Register Selection During SEARCH and LEARN





Figure 14. Addressing the Global Mask Register (GMR) Array



SEARCH-Successful Registers

The device contains eight SEARCH-successful registers (SSRs) to hold the index of the location where a successful search occurred. The format of each register is described in Table 9. The SEARCH command specifies which SSR stores the index of a specific SEARCH command in "Cycle B" of the SEARCH Instruction.

After the index location is specified, the host ASIC can use this register to access that data array, mask array, or external SRAM using the index as part of the address (see SRAM ADDRESSING, page 58). The device with a valid bit set performs a READ or WRITE operation. All other devices suppress the operation.

Table 9. SEARCH-Successful Register (SSR) Description

Field	Range	Initial Value	Description
INDEX	[13:0]	X	Index. This is the address of the 68-bit entry where a successful search occurs. The device updates this field if it has a successful search. In 136-bit, the LSB is '0;' in a 272-bit configuration, the two LSBs are '00.' The index updates if the device is either a local or global winner in a SEARCH operation.
–	[30:14]	0	Reserved.
VALID	[31]	0	Valid. The device sets this bit to '1' if it is a global winner (first device downstream with a hit) in a SEARCH operation, in a depth-cascaded configuration.
–	[67:32]	0	Reserved.

The Command Register

Table 10. Command Register Field Descriptions

Field	Range	Initial Value	Description										
SRST	[0]	0	Software Reset. If '1,' this bit resets the device, with the same effect as the hardware reset. Internally, it generates a reset pulse lasting for eight CLK cycles. This bit automatically resets to a '0' during the reset cycle.										
DEVE	[1]	0	Device Enable. If '0,' it keeps the SRAM bus (SADR, WE_L, CE_L, OE_L, and ALE_L), SSF, and SSV signals in a tri-state condition and forces the cascade interface output signals LHO[1:0] and BHO[2:0] to '0.' It also keeps the DQ Bus in Input mode. The purpose of this bit is to make sure that there is no bus contention when the devices power-up in the system.										
TLSZ	[3:2]	01	<p>Table Size. The host ASIC must program this field to configure the chips into a table of a certain size. This field affects the pipeline latency of the SEARCH and LEARN operations as well as the READ and WRITE accesses to the SRAM (SADR[21:0], CE_L, OE_L, WE_L, ALE_L, SSV, SSF, and ACK). Once programmed, the SEARCH latency stays constant.</p> <table border="0"> <thead> <tr> <th></th> <th>Latency # CLK Cycles</th> </tr> </thead> <tbody> <tr> <td>00: 1 device</td> <td>4</td> </tr> <tr> <td>01: 2-8 devices</td> <td>5</td> </tr> <tr> <td>10: 9-31 devices</td> <td>6</td> </tr> <tr> <td>11: Reserved</td> <td></td> </tr> </tbody> </table>		Latency # CLK Cycles	00: 1 device	4	01: 2-8 devices	5	10: 9-31 devices	6	11: Reserved	
	Latency # CLK Cycles												
00: 1 device	4												
01: 2-8 devices	5												
10: 9-31 devices	6												
11: Reserved													
HLAT	[6:4]	000	<p>Latency of Hit Signals. This field adds latency to the SSF, SSV, and ACK signals by the following number of CLK cycles during SEARCH and ACK during an SRAM READ access.</p> <table border="0"> <tbody> <tr> <td>000: 0</td> <td>100: 4</td> </tr> <tr> <td>001: 1</td> <td>101: 5</td> </tr> <tr> <td>010: 2</td> <td>110: 6</td> </tr> <tr> <td>011: 3</td> <td>111: 7</td> </tr> </tbody> </table>	000: 0	100: 4	001: 1	101: 5	010: 2	110: 6	011: 3	111: 7		
000: 0	100: 4												
001: 1	101: 5												
010: 2	110: 6												
011: 3	111: 7												
LDEV	[7]	0	<p>Last device in the cascade. When set, this device is the last device in the depth-cascaded table and is the default driver for the SSF and SSV signals. In the event of a SEARCH failure, the device with this bit set drives the hit signals as follows: SSF = 0, SSV = 1</p> <p>During non-search cycles, the device with this bit set drives the signals as follows: SSF = 0, SSV = 0</p>										
LRAM	[8]	0	<p>Last device on this SRAM Bus. When set, this device is the last device on the SRAM bus in the depth-cascaded table and is the default driver for the SADR, CE_L, WE_L, and ALE_L signals. In cycles where no M7010R device (in a depth-cascaded table) drives these signals, the signals are driven as follows: SADR = 22'h3FFFFFF, CE_L = 1, WE_L = 1, and ALE_L = 1. OE_L is always driven by the device for which this bit is set.</p>										

Field	Range	Initial Value	Description
CFG	[16:9]	0000 0000	<p>Database Configuration. The device is internally divided into four quadrants of 8K x 68, each of which can be configured as 4K x 68, 2K x 136, or 1K x 272 as follows:</p> <p>00: 4K x 68 01: 2K x 136 10: 1K x 272 11: Reserved</p> <p>Bits [10:9] apply to configuring the 1st quadrant in the address space. Bits [12:11] apply to configuring the 2nd quadrant in the address space. Bits [14:13] apply to configuring the 3rd quadrant in the address space. Bits [16:15] apply to configuring the 4th quadrant in the address space.</p>
	[67:17]	0	Reserved.

SEARCH PROCEDURE FOR 32-BIT WIDE PREFIXES

The Global Mask Register is used for 32-bit wide data paths as follows:

Writing a '1' in the Global Mask Register allows data to be written into the M7010R. A '0' in the Global Mask Register disallows data modification. Information is written into the left half of the 68-bit word Search Engine as long as space for 34 bits of data is available and then into the right half of the Search Engine. 32-bit data can be entered in two cycles.

The first step is to write into two of the eight Global Mask Registers with the patterns shown in Figure 15. Writing this data using Global Mask Register 1 allows the left half of the data array to be completely filled.

Figure 16 shows Bits 67 through 36 in the left section of the data array representing 32-bits of data. Bits 35 and 34 shown separately can be defined by the user for table management. In this application 34-bit operation occurs in each half-section of the Data and Mask arrays of the Search Engine. The left half is filled first, then the right. Not all locations have to be filled.

SEARCH operations are performed twice, once on the left half and then on the right half. Note that a '1' in the Global Mask register enables a compare during a SEARCH operation and a '0' forces a match condition regardless of the state of the data bit.

The SEARCH throughput for 34-bit operations is half of the 68-bit operations. A search is performed by using the Global Mask Register "0" for the left half of the 68-bit, then another search is performed using Global Mask Register 1 for the right half of the 68-bit word. The order is important, as the left half has a higher priority than the right half.

For example, if a search on the left half produces a match and a search on the right half also produc-

es a match, then in that case, the left half is a higher priority. So if only one unique match exists in a particular system, then a match on the left side may alleviate the need to do a search on the right half of the Data array.

Figure 15. Global Mask Register Patterns

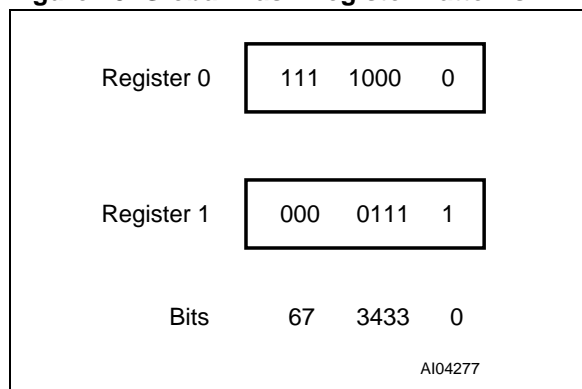



Figure 16. Storing left half of a Data or Mask Array



The Information Register

Table 11. Information Register Field Descriptions

Field	Range	Initial Value	Description
Revision	[3:0]	0001	Revision Number. This is the current device revision number. Numbers start from one and increment by one for each revision of the device.
Implementation	[6:4]	000	This is the M7010R implementation number.
Reserved	[7]	0	Reserved.
Device ID	[15:8]	00000001	This is the Device Identification Number.
MFID	[31:16]	1101_1100_ 0111_1111	Manufacturer ID. This field is the same as the manufacturer ID and continuation bits.
	[67:32]		Reserved.

The READ Burst Address Register (RBURREG)

These READ burst address register fields must be programmed before burst READ (see Table 12).

The WRITE Burst Address Register (WBURREG)

These WRITE burst address register fields must be programmed before burst WRITE (see Table 13).

Table 12. READ Burst Register Description

Field	Range	Initial Value	Description
AADR	[13:0]	0	Address. This is the starting address of the data array or mask array during a burst READ operation. It automatically increments by 1 for each successive read of the data array or mask array. Once the operation is complete, the contents of this field must be reinitialized for the next operation.
	[18:14]		Reserved.
BLEN	[27:19]	0	Length of Burst Access. The device provides the capability to read from 4 up to 511 locations in a single burst. The BLEN decrements automatically. Once the operation is complete, the contents of this field must be reinitialized for the next operation.
	[67:28]		Reserved.

Table 13. WRITE Burst Register Description

Field	Range	Initial Value	Description
AADR	[13:0]	0	Address. This is the starting address of the data array or mask array during a burst WRITE operation. It automatically increments by 1 for each successive write of the data array or mask array. It increments by 1 for each successive read of the data array or mask array. Once the operation is complete, the contents of this field must be reinitialized for the next operation.
	[18:14]		Reserved.
BLEN	[27:19]	0	Length of Burst Access. The device provides the capability to write from 4 up to 511 locations in a single burst. The BLEN decrements automatically. Once the operation is complete, the contents of this field must be reinitialized for the next operation.
	[67:28]		Reserved.

The NFA Register

Bit [0] of each 68-bit data entry is a special bit designated for use in the operation of the LEARN command. In 68-bit configurations, the Bit[0] indicates whether a location is full (bit set to '1') or empty (bit set to '0'). Every WRITE/LEARN command loads the address of first 68-bit location that contains a "0" in the entry's Bit[0]. This is stored in the NFA register. If all the bits in a device are set to '1,' the M7010R asserts FULO[1:0] to '1.'

In a 136-bit configuration, the LSB of this register is always set to '0.' The host ASIC must set Bit 0

and Bit 68 in a 136-bit word to either '0' or '1' to indicate full/empty status for a 136-bit entry.

Note: Both Bits 0 and Bit 68 must be set to '0' or '1' (e.g., '10' or '01' settings are invalid).

Table 14. NFA Register

Address	67 - 14	13 - 0
60	Reserved	Index

SEARCH ENGINE ARCHITECTURE

The M7010R consists of 16k x 68-bit storage cells referred to as "data bits." There is a mask cell corresponding to each data cell. Figure 17 shows the three organizations of the device based on the value of CFG bits in the COMMAND register.

During a SEARCH operation, the SEARCH Data Bit (S), Data Array Bit (D), Mask Array Bit (M) and the Global Mask Bit (G) are used in the following manner to generate a match at that bit position (see Table 15, page 25).

The entry with all matched bit positions results in a successful search in the M7010R. In order for a successful SEARCH to make the device the local winner in the SEARCH operation, all 68-bit positions within a device must generate a match for a 68-bit entry in 68-bit-configured quadrants, or all 136-bit positions must generate a match for two consecutive even and odd 68-bit entries in quadrants configured as 136 bits, or all 272-bit positions must generate a match for four consecutive

entries aligned to four entry-page boundaries of 68-bit entries in quadrants configured as 272 bits.

An arbitration mechanism using a cascade bus determines the global winning device among the local winning devices in a SEARCH cycle. The global winning device drives the SRAM bus, SSV, and The SSF signals. In the case of a SEARCH failure, the device(s) with LDEV and LRAM bits set drive the SRAM bus, SSF, and SSV signals.

The M7010R may be partitioned into up to four (4) quadrants of different widths (e.g., 34, 68, 136, or 272 bits), even within the same chip (see Application Notes AN1338 and AN1339). Figure 18 shows a sample configuration of different widths.

Data and Mask Addressing

Figure 19, page 26 shows the M7010R data array and mask array addressing procedure. The data array and mask array addresses differ only in one bit in the address cycle of the READ and WRITE commands.

Figure 17. M7010R Database Configuration




Table 15. Bit Position Match

G	M	S	D	Match
0	x	x	x	1
1	0	x	x	1
1	1	0	0	1
1	1	0	1	0
1	1	1	0	0
1	1	1	1	1

Figure 18. Multi-width Configuration Example





Figure 19. M7010R Data and Mask Array Addressing



COMMAND CODES AND PARAMETERS

A master device, such as an ASIC controller, issues commands to the M7010R using the CMDV signal and the CMD Bus. The following subsections describe the functions of the commands.

Command Codes

The M7010R implements four basic commands shown in Table 16. The Command code must be presented to CMD[1:0] while keeping the command valid (CMDV) signal high for two CLK2X cycles.

These two CLK2X cycles are designated as "Cycle A" and "Cycle B." The CMD[8:2] field passes the parameters of the command in CLK2X Cycles A and B. The controller ASIC must align the instructions with the CLK2X signal.

Commands and Command Parameters

Table 17 lists the CMD bus fields that contain the M7010R command parameters as well as their respective cycles.

Table 16. Command Codes

CMD Code	Command	Description
00	READ	Reads one of the following: data array, mask array, device registers, or external SRAM.
01	WRITE	Writes one of the following: data array, mask array, device registers, or external SRAM.
10	SEARCH	Searches the data array for a desired pattern using the specified register from the global mask register array and local mask associated with each data cell.
11	LEARN	The device has internal storage for up to 16 comparands that it can learn. The device controller can insert these entries at the next free address (as specified by the NFA register) using the LEARN Instruction.

Table 17. Command Parameters

Cmd	Cyc	8	7	6	5	4	3	2	1	0
READ	A	SADR[21]	SADR[20]	SADR[19]	0	0	0	0 = Single 1 = Burst	0	0
	B	0	0	0	0	0	0	0 = Single 1 = Burst	0	0
WRITE	A	SADR[21]	SADR[20]	SADR[19]	GMR Index[2:0]			0 = Single 1 = Burst	0	1
	B	0	0	0	GMR Index[2:0]			0 = Single 1 = Burst	0	1
SEARCH	A	SADR[21]	SADR[20]	SADR[19]	GMR Index[2:0]			68-bit or 136-bit: 0 272-bit: 1 in 1st Cycle 0 in 2nd Cycle	1	0
	B	Successful SEARCH Register Index[2:0]			Comparand Register Index				1	0
LEARN ⁽¹⁾	A	SADR[21]	SADR[20]	SADR[19]	Comparand Register Index				1	1
	B	0	0	Mode 0: 68-bit 1: 136-bit	Comparand Register Index				1	1

Note: The SRAM Address Bit SADR [19] in the command bit C6 will not be passed to the SRAM (see Table 28).

1. The 272-bit configuration does not support the LEARN Instruction.

READ COMMAND

The READ can be a single read of a data array, a mask array, an SRAM, or a register location (CMD[2] = 0). It can be a burst READ (CMD[2] = 1) using an internal auto-incrementing address register (RBURADR) of the data or mask array locations (see Table 18, page 30 and Table 19, page 30 for formats).

A single-location READ operation takes six cycles, as shown in Figure 20, page 29. The burst READ adds two cycles for each successive read. The SADR[21:19] bits supplied in the READ Instruction Cycle A drives SADR[21:19] signals during the PIO READ of an SRAM location.

The single READ operation takes six CLK cycles, in the following sequence:

- **Cycle 1:** The host ASIC applies the READ Instruction on the CMD[1:0] (CMD[2] = 0), using CMDV = 1, and the DQ Bus supplies the address, as shown in Table 19, page 30 and Table 20, page 30. The host ASIC selects the device for which ID[4:0] matches the DQ[25:21] lines. If DQ[25:21] = 11111, the host ASIC selects the M7010R with the LDEV Bit set. The host ASIC also supplies SADR[21:19] on CMD[8:6] in Cycle A of the READ Instruction if the READ is directed to the external SRAM.
- **Cycle 2:** The host ASIC releases the DQ[67:0] bus to a tri-state condition.
- **Cycle 3:** The host ASIC keeps DQ[67:0] bus in a tri-state condition.
- **Cycle 4:** The selected device starts to drive the DQ[67:0] bus and drives the ACK signal from Z to low.
- **Cycle 5:** The selected device drives the READ data from the addressed location on the DQ[67:0] bus and drives the ACK signal high.
- **Cycle 6:** The selected device floats the DQ[67:0] bus and drives the ACK signal low.

At the termination of Cycle 6, the selected device releases the ACK line to a tri-state condition. The

READ Instruction is complete, and a new operation can begin.

The burst READ operation lasts $4 + 2n$ CLK-cycles (where “n” stands for the number of accesses in the burst specified by the BLEN field of the RBURREG) in the sequence shown in Figure 21, page 29. This operation assumes that the host ASIC has programmed the RBURREG with the starting address (ADDR) and the length of transfer (BLEN) before initiating the burst READ command.

- **Cycle 1:** The host ASIC applies the READ Instruction on the CMD[1:0] (CMD[2] = 1), using CMDV=1 and the address supplied on the DQ Bus, as shown in Table 21, page 31. The host ASIC selects the device for which ID[4:0] matches the DQ[25:21] lines. If DQ[25:21] = 11111, the host ASIC selects the M7010R with the LDEV Bit set.
 - **Cycle 2:** The host ASIC floats DQ[67:0] to a tri-state condition.
 - **Cycle 3:** The host ASIC keeps DQ[67:0] bus in a tri-state condition.
 - **Cycle 4:** The selected device starts to drive the DQ[67:0] bus and drives ACK, and EOT from Z to low.
 - **Cycle 5:** The selected device drives the READ data from the addressed location on the DQ[67:0] bus and drives the ACK signal high.
- Note:** Cycles four and five repeat for each additional access until all the accesses specified in the burst length (BLEN) field of RBURREG are complete. On the last transfer, the M7010R drives the EOT signal high.
- **Cycle (4 + 2n):** The selected device drives the DQ[67:0] to 3-state condition and drives the ACK and the EOT signals low.

At the termination of Cycle $4 + 2n$, the selected device floats the ACK line to 3-state condition. The burst READ Instruction is complete, and a new operation can begin (see Table 21, page 31 for burst READ address formats).

Figure 20. Single Location READ Cycle Timing

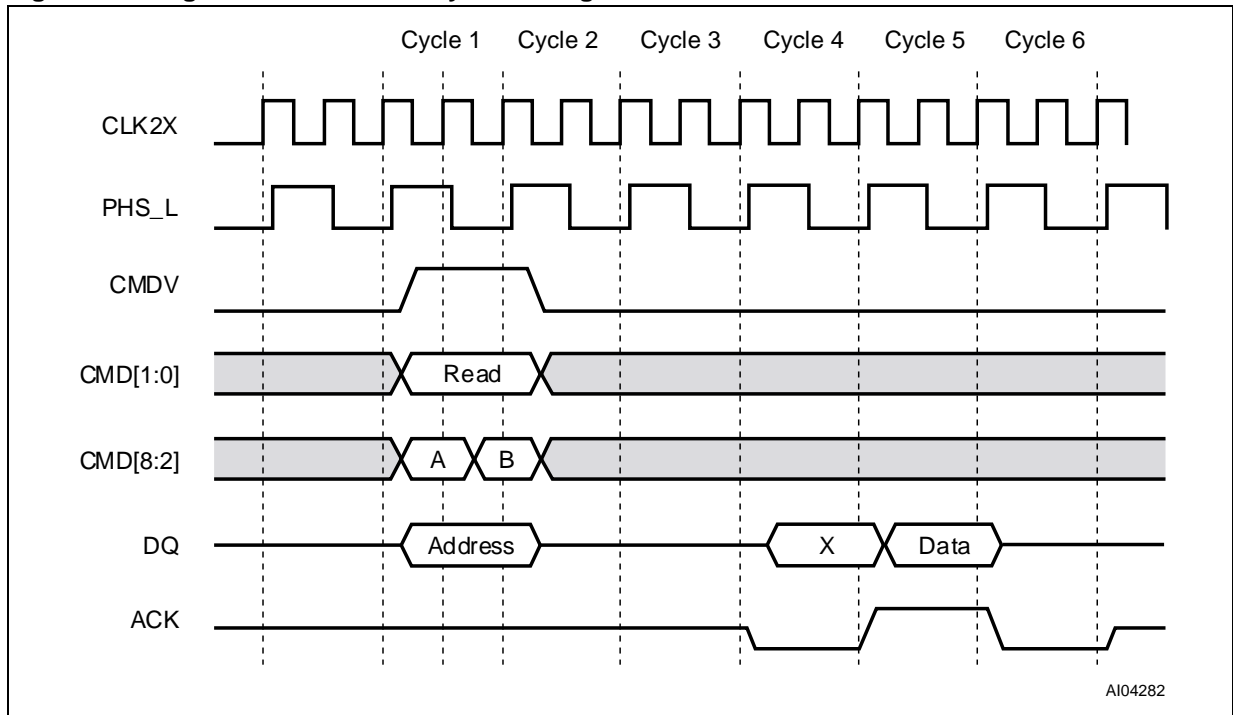


Figure 21. Burst READ of the Data and Mask Arrays (BLEN = 4)

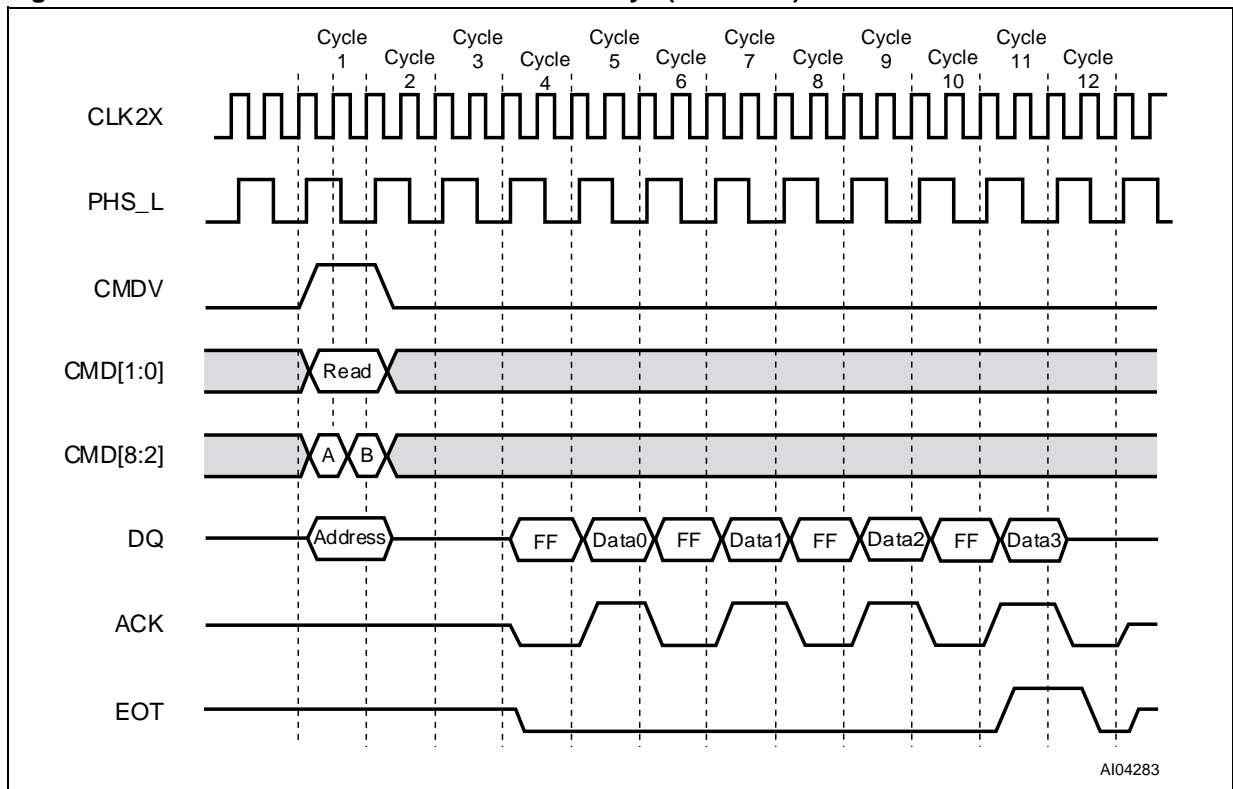


Table 18. READ Command Parameters

CMD Parameter CMD[2]	Read Command	Description
0	Single Read	Reads a single location of the data array, mask array, external SRAM, or device registers. All access information is applied on the DQ Bus.
1	Burst Read	Reads a block of locations from the data array or mask array as a burst. The internal register (R BURADR) specifies the starting address and the length of the data transfer from the data array or mask array, and it auto-increments the address for each access. All other access information is applied on the DQ Bus. Note: The device registers and external SRAM can only be read in single-read mode.

Table 19. Data and Mask Array, SRAM READ Address Format

DQ [67:30]	DQ [29]	DQ [28:26]	DQ [25:21]	DQ [20:19]	DQ [18:14]	DQ [13:0]
Reserved	0: Direct 1: Indirect	Successful SEARCH Register Index (Applicable if DQ[29] is indirect)	ID	00: Data Array	Reserved	If DQ[29] is '0,' this field carries address of data array location. If DQ[29] is '1,' the successful SEARCH Register specified on DQ[28:26] supplies the address of the data array location: {SSR[13:2], SSR[1] DQ[1], SSR[0] DQ[0]} ⁽¹⁾
Reserved	0: Direct 1: Indirect	Successful SEARCH Register Index (Applicable if DQ[29] is indirect)	ID	01: Mask Array	Reserved	If DQ[29] is '0,' this field carries address of mask array location. If DQ[29] is '1,' the successful SEARCH Register specified on DQ[28:26] supplies the address of the mask array location: {SSR[13:2], SSR[1] DQ[1], SSR[0] DQ[0]} ⁽¹⁾
Reserved	0: Direct 1: Indirect	Successful SEARCH Register Index (Applicable if DQ[29] is indirect)	ID	10: External SRAM	Reserved	If DQ[29] is '0,' this field carries address of SRAM location. If DQ[29] is '1,' the successful SEARCH Register specified on DQ[28:26] supplies the address of the SRAM location.

Note: 1. “|” stands for logical OR operation, and “{ }” stands for concatenation operator.

Table 20. READ Address Format for Internal Registers

DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
Reserved	ID	11: Register	Reserved	Register Address



Table 21. READ Address Format for Data and Mask Arrays

DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:14]	DQ[13:0]
Reserved	ID	00: Data Array	Reserved	Do not care. These 14 bits come from the internal register (RBURADR) which increments for each access.
Reserved	ID	01: Mask Array	Reserved	Do not care. These 14 bits come from the internal register (RBURADR) which increments for each access.

WRITE COMMAND

The WRITE can be a single write of a data array, mask array, register, or external SRAM location (CMD[2] = 0). It can also be a burst WRITE (CMD[2] = 1) using an internal auto-incrementing address register (WBURADR) of the data array or mask array locations (see Table 23, page 33 for format). A single-location WRITE is a three-cycle operation, shown in Figure 22, page 32. The burst WRITE adds one extra cycle for each successive location write.

The WRITE operation sequence is as follows:

- **Cycle 1A:** The host ASIC applies the WRITE Instruction to CMD[1:0] (CMD[2] = 0), using CMDV=1 and the address supplied on the DQ Bus, as shown in Table 22, page 33. The host ASIC also supplies the index to the global mask register (GMR) to mask the WRITE to the data array or mask array location in CMD[5:3]. For SRAM writes, the host ASIC must supply SADR[21:19] on CMD[8:6].
- **Cycle 1B:** The host ASIC continues to apply the WRITE Instruction to CMD[1:0] (CMD[2] = 0) using CMDV = 1 and the address supplied on the DQ Bus. The host ASIC continues to supply the GMR Index to mask the WRITE to the data or mask array locations in CMD[5:3]. The host ASIC selects the device where ID[4:0] matches the DQ[25:21] = 11111.
- **Cycle 2:** The host ASIC drives the DQ[67:0] with the data to be written to the data array, mask array, external SRAM, or register location of the selected device.
- **Cycle 3:** Idle cycle. At the termination of this cycle, another operation can begin.

The burst WRITE operation lasts for (n + 2) CLK cycles, where “n” signifies the number of accesses in the burst as specified in the BLEN field of the WBURREG register (see Figure 23, page 32).

This operation assumes that the host ASIC has programmed the WBURREG with the starting address (ADDR) and the length of transfer (BLEN) before initiating the burst WRITE command (see

Table 24, page 33 for format). The sequence is as follows:

- **Cycle 1A:** The host ASIC applies the WRITE Instruction on the CMD[1:0] (CMD[2] = 1), using CMDV = 1 and the address supplied on the DQ Bus, as shown in Table 23, page 33. The host ASIC also supplies the index to the global mask register to mask the WRITE to the data or mask array locations in CMD[5:3].
- **Cycle 1B:** The host ASIC continues to apply the WRITE Instruction to CMD[1:0] (CMD[2] = 0) using CMDV = 1 and the address supplied on the DQ Bus. The host ASIC continues to supply the GMR Index to mask the WRITE to the data or mask array locations in CMD[5:3]. The host ASIC selects the device where ID[4:0] matches the DQ[25:21] = 11111.
- **Cycle 2:** The host ASIC drives the DQ[67:0] with the data to be written to the data array or mask array location of the selected device. The host ASIC writes the data on the DQ[67:0] bus only to the subfield that has the corresponding mask bit set to '1' in the global mask register specified by the index CMD[5:3] and supplied in Cycle 1.
- **Cycles 3 to n + 1:** The host ASIC drives DQ[67:0] with the data to be written to the next data array or mask array location (addressed by the auto-increment AADR field of the WBURREG register) of the selected device.
The host ASIC writes the data on the DQ[67:0] bus only to the subfield that has the corresponding mask bit set to '1' in the global mask register specified by the index CMD[5:3] and supplied in Cycle 1. The M7010R drives the EOT signal low from Cycle 3 to Cycle n; the M7010R drives the EOT signal high in Cycle n + 1 (n is specified in the BLEN field of the WBURREG).
- **Cycle n + 2:** The M7010R drives the EOT signal low. At the termination of the Cycle n + 2, the M7010R floats the EOT signal to a 3-state, and a new instruction can begin.

Figure 22. Single Location WRITE Cycle Timing




Figure 23. Burst WRITE of the Data and Mask Arrays (BLEN = 4)




Table 22. (Single) WRITE Address Format for Data and Mask Arrays or SRAM

DQ [67:30]	DQ [29]	DQ [28:26]	DQ [25:21]	DQ [20:19]	DQ [18:14]	DQ [13:0]
Reserved	0: Direct 1: Indirect	Successful SEARCH Register Index (Applicable if DQ[29] is indirect)	ID	00: Data Array	Reserved	If DQ[29] is '0,' this field carries the address of the data array location. If DQ[29] is '1,' the SSR specified on DQ[28:26] is used to generate the address of the data array location: {SSR[13:2], SSR[1] DQ[1], SSR[0] DQ[0]}. ⁽¹⁾
Reserved	0: Direct 1: Indirect	Successful SEARCH Register Index (Applicable if DQ[29] is indirect)	ID	01: Mask Array	Reserved	If DQ[29] is '0,' this field carries address of the mask array location. If DQ[29] is '1,' the SSR specified on DQ[28:26] is used to generate the address of the data array location: {SSR[13:2], SSR[1] DQ[1], SSR[0] DQ[0]}. ⁽¹⁾
Reserved	0: Direct 1: Indirect	Successful SEARCH Register Index (Applicable if DQ[29] is indirect)	ID	10: External SRAM	Reserved	If DQ[29] is '0,' this field carries address of the data SRAM location. If DQ[29] is '1,' the SSR specified on DQ[28:26] is used to generate the address of the data array location: {SSR[13:2], SSR[1] DQ[1], SSR[0] DQ[0]}. ⁽¹⁾

Note: 1. "||" stands for logical OR operation, and "{" }" stands for concatenation operator.

Table 23. WRITE Address Format for Internal Registers

DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
Reserved	ID	11: Register	Reserved	Register address

Table 24. WRITE Address Format for Data and Mask Array (Burst WRITE)

DQ [67:26]	DQ [25:21]	DQ [20:19]	DQ [18:14]	DQ [13:0]
Reserved	ID	00: Data array	Reserved	Don't care. These 14 bits come from the internal register (WBURADR), which increments with each access.
Reserved	ID	01: Mask array	Reserved	Don't care. These 14 bits come from the internal register (WBURADR), which increments with each access.

SEARCH COMMAND

The M7010R Search Engine can be configured in three ways:

1. 68-bit
2. 136-bit
3. 272-bit
4. Mixed-sized SEARCHES on tables configured with different widths

68-bit Configuration

Figure 25, page 35 shows the timing diagram for a SEARCH operation in the 68-bit-configured table (one device only). This illustration assumes that the host ASIC has programmed TLSZ to '00,' HLAT to '000,' LRAM to '1,' and LDEV to '1' in the command register. The hardware diagram for this search subsystem is shown in Figure 24.

- **Cycle A:** The host ASIC drives CMDV high and applies the SEARCH command code (10) on CMD[1:0]. CMD[5:3] must be driven by the index to the global mask register pair for use in the SEARCH operation. CMD[8:6] signals must be driven by the same bits that will be driven on SADR[21:19] by this device if it has a hit. DQ[67:0] must be driven with the data to be compared. CMD[2] signal must be driven to logic '0.'
- **Cycle B:** The host ASIC continues to drive CMDV high and to apply the SEARCH command (10) on CMD[1:0]. CMD[5:2] must be driven by the index of the comparand register pair for storing the 136-bit word presented on the DQ Bus during Cycles A and B. CMD[8:6] signals must be driven with the index of the SSR that will be used for storing the address of the matching entry and the hit flag. The DQ[67:0] continues to carry the 68-bit data to be compared.

Note: In the 68-bit configuration, the host ASIC must supply the same data on DQ[67:0] during cycles A and B. The even and odd GMR pairs selected for the compare must be programmed with the same value.

The SEARCH command is a pipelined operation and executes a SEARCH at half their rate of frequency of CLK2X for 68-bit searches in x68-configured tables. The latency of SADR, CE_L, ALE_L, WE_L, SSV, and SSF from 68-bit SEARCH Command cycle (= two CLK2X cycles) is shown in Table 27, page 36.

The timing diagram for all SRAM interface signals, SSV, and SSF shift to the right for different values of TLSZ, as specified in Table 25, page 36 and Table 26, page 36.

In addition, SSV and SSF shift to the right for different values of HLAT, as specified in Table 26, page 36.

68-bit Configuration with LDEV = 1. The device is configured to be the last in the depth-cascaded table by setting LDEV to '1' in the Command Register. The device with LDEV set to '1' drives the SSF and SSV signals in cycles when all upstream devices do not drive these signals. The M7010R with its LDEV Bit set drives SSF and SSV during a search with a miss or with non-search commands (see the LDEV Bit definition in Table 10, page 20).

68-bit Configuration with LRAM = 1. Setting LRAM to '1' in the Command Register configures the device to be the last on the SRAM Bus. In a cycle where the upstream device does not drive the SRAM Bus, the last device of the SRAM Bus (with LRAM = 1) drives the bus (SADR, CE_L, WE_L, ALE_L) when they are active. When set to '1,' the LRAM Bit sets the default driver for the SRAM control signals (SADR, CE_L, WE_L, and ALE_L).

Figure 24. Hardware Diagram for a Table with a Single Device (68-bit Operation)

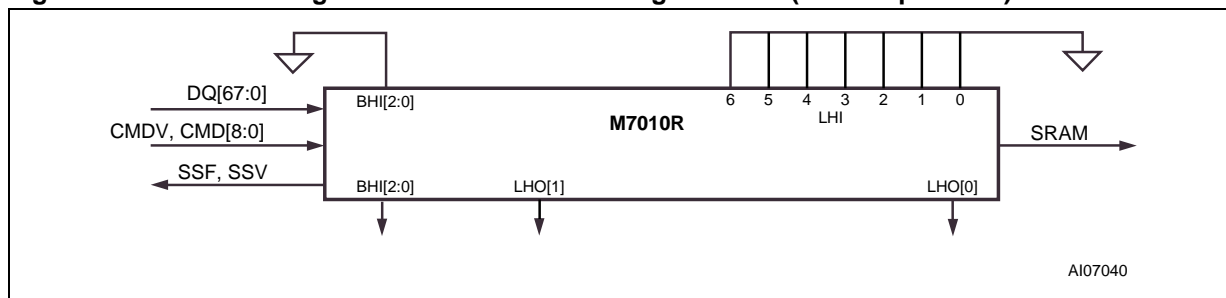


Figure 25. 68-Bit Configuration SEARCH Timing Diagram (One Device)




Table 25. Right-Shift of 68-bit Signals for TLSZ Values

TLSZ	Number of CLK Cycles
00	0
01	1
10	2

Table 26. Shift of SSF and SSV from SADR (for different HLAT Values)

HLAT	Number of CLK Cycles
000	0
001	1
010	2
011	3
100	4
101	5
110	6
111	7

Table 27. Latency of SEARCH from Instruction to SRAM Access Cycle (68-bit Mode)

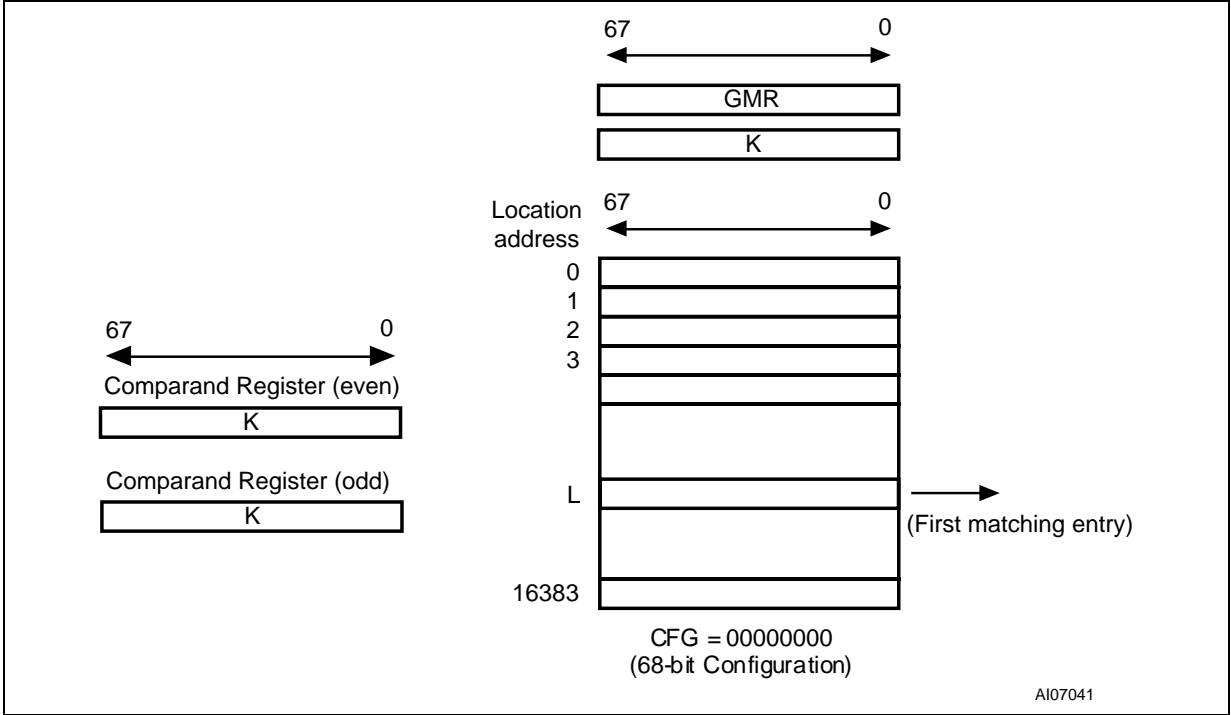
# of devices	Max Table Size	Latency in CLK Cycles
1 (TLSZ = 00)	16K x 68-bit	4
2–8 (TLSZ = 01)	128K x 68-bit	5
9–31 (TLSZ = 10)	496K x 68-bit	6

68-bit Logical SEARCH

The logical, 68-bit SEARCH operation is shown in Figure 26. The entire table of 68-bit entries is compared to a 68-bit word K (presented on the DQ Bus in both Cycles A and B of the command) using the GMR and the local mask bits. The effective GMR is the 68-bit word specified by the identical value in both even and odd GMR pairs selected by the GMR Index in the command's Cycle A. The 68-bit word K (presented on the DQ Bus in both Cycles A and B of the command) is also stored in both even and odd comparand register pairs selected

by the Comparand Register Index in the command's Cycle B. In a x68 configuration, only the even comparand register can subsequently be used by the LEARN command. The word K (presented on the DQ Bus in both Cycles A and B of the command) is compared with each entry in the table, starting at location "0." The first matching entry's location, address "L," is the winning address that is driven as part of the SRAM address on the SADR[21:0] lines.

Figure 26. x68 Table with One Device



136-bit Configuration

Figure 28, page 39 shows the timing diagram for the SEARCH operation in the 136-bit table (CFG = 01010101) consisting of a single device for one set of parameters: TLSZ = 00, HLAT = 001, LRAM = 1, and LDEV = 1. The hardware diagram for the search subsystem is shown in Figure 27.

The following is the operation sequence for a single, 136-bit SEARCH command.

- **Cycle A:** The host ASIC drives the CMDV high and applies the SEARCH command code (10) to CMD[1:0] signals. CMD[5:3] signals must be driven with the index to the GMR pair for use in this SEARCH operation. CMD[8:6] signals must be driven with the same bits that will be driven on SADR[21:19] by this device if it has a hit. DQ[67:0] must be driven with the 68-bit data ([135:68]) to be compared against all even locations. The CMD[2] signal must be driven to logic '0.'
- **Cycle B:** The host ASIC continues to drive the CMDV high and applies SEARCH command code (10) on CMD[1:0]. CMD[5:2] must be driven by the index to the comparand register pair for storing the 136-bit word presented on the DQ Bus during Cycles A and B. CMD[8:6] signals must be driven by the index of the SSR that will be used for storing the address of the matching entry and hit flag. The DQ[67:0] is driven with

68-bit data ([67:0]), compared to all odd locations.

Note: For 136-bit searches, the host ASIC must supply two distinct, 68-bit data words on DQ[67:0] during Cycles A and B. The even-numbered GMR of the pair specified by the GMR Index is used for masking the word in Cycle A. The odd-numbered GMR of the pair specified by the GMR Index is used for masking the word in Cycle B.

The SEARCH command is a pipelined operation that executes searches at half the rate of the frequency of CLK2X for 136-bit searches in x136-bit-configured tables. The latency of SADR, CE_L, ALE_L, WE_L, SSV, and SSF from the 136-bit SEARCH command cycle (two CLK2X cycles) is shown in Table 30, page 40.

The timing diagram for all SRAM interface signals, SSV, and SSF shift to the right for different values of TLSZ, as specified in Table 28, page 40.

In addition, SSV and SSF shift to the right for different values of HLAT, as specified in Table 29, page 40.

The result of the SEARCH operation appears as an SRAM READ Cycle with a pipelined latency. It is specified as shown in Table 30, page 40.

Figure 27. Hardware Diagram for a Table with One Device (136-bit Operation)




Figure 28. 136-Bit Configuration SEARCH Timing Diagram (One Device)

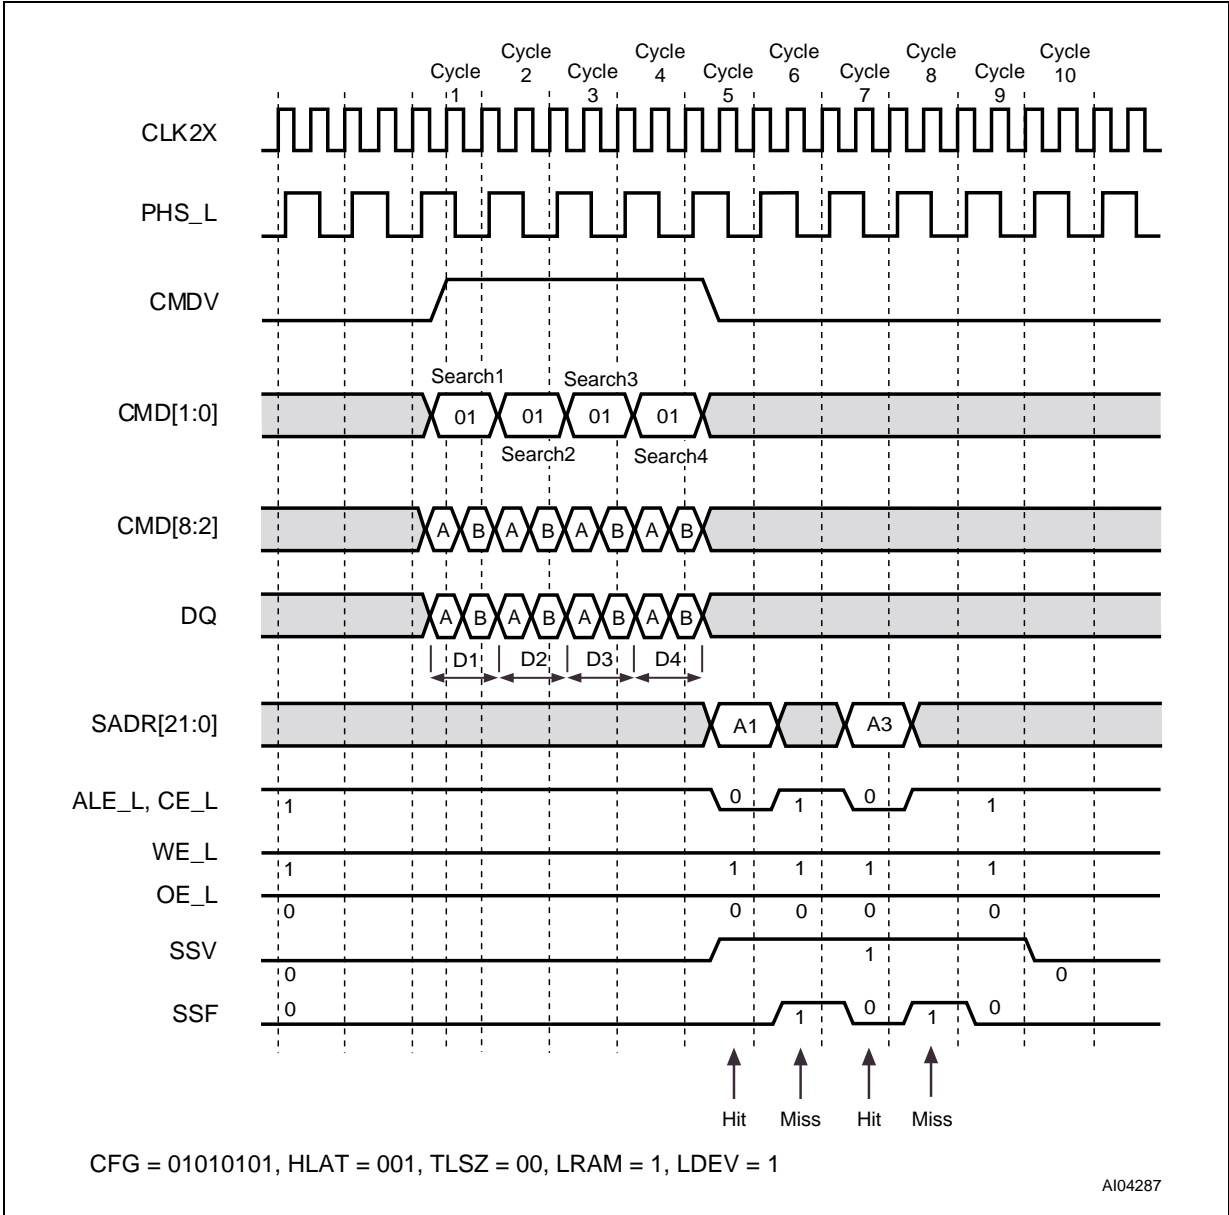


Table 28. Right-Shift of 136-bit Signals for TLSZ Values

TLSZ	Number of CLK Cycles
00	0
01	1
10	2

Table 29. Shift of SSF and SSV from SADR (for different HLAT values)

HLAT	Number of CLK Cycles
000	0
001	1
010	2
011	3
100	4
101	5
110	6
111	7

Table 30. Latency of SEARCH from Instruction to SRAM Access Cycle (136-bit Mode)

# of devices	Max Table Size	Latency in CLK Cycles
1 (TLSZ = 00)	8K x 136-bit	4
2–8 (TLSZ = 01)	64K x 136-bit	5
9–31 (TLSZ = 10)	248K x 136-bit	6


136-bit Logical SEARCH

The logical, 136-bit SEARCH operation is shown in Figure 29. The entire table of 136-bit entries is compared to a 136-bit word K (presented on the DQ Bus in both Cycles A and B of the command) using the GMR and the local mask bits. The GMR is the 136-bit word specified by the even and odd global mask pair selected by GMR Index in the command's Cycle A. The 136-bit word K (presented on the DQ Bus in both Cycles A and B of the command) is also stored in both even and odd comparand register pairs selected by the Comparand Register Index in the command's Cycle B. The two comparand registers can subsequently be used by the LEARN command with the even-

numbered comparand register stored in an even-numbered location, and the odd-numbered comparand register stored in an adjacent, odd-numbered location. The word K (presented on the DQ Bus in Cycles A and B of the command) is compared with each entry in the table starting at location "0." The first matching entry's location, address "L," is the winning address that is driven as part of the SRAM address on the SADR[21:0] lines.

Note: The matching address is always going to be an even-numbered address for a 136-bit SEARCH.

Figure 29. x136 Table with One Device



272-bit Configuration

Figure 31, page 43 shows the timing diagrams for a SEARCH operation in the 272-bit-configured table (CFG = 10101010) consisting of a single device for one set of parameters: TLSZ = 00, HLAT = 001, LRAM = 1, and LDEV = 1. The hardware diagram for this search subsystem is shown in Figure 30.

– **Cycle A:** The host ASIC drives the CMDV high and applies the SEARCH command code (10) to CMD[1:0] signals. CMD[5:3] signals must be driven with the index to the GMR pair for bits [271:136] of the data being searched. DQ[67:0] must be driven with the 68-bit data ([271:204]) to be compared to all locations “0” in the four 68-bits-word page. The CMD[2] signal must be driven to logic '1.'

Note: CMD[2] = 1 signals that the search is a x272-bit search. CMD[8:3] is ignored.

– **Cycle B:** The host ASIC continues to drive CMDV high and continues to apply SEARCH command code (10) on CMD[1:0]. The DQ[67:0] is driven with 68-bit data ([203:136]) to be compared to all locations “1” in the 68-bits-word page.

– **Cycle C:** The host ASIC drives the CMDV high and applies the SEARCH command code (10) to CMD[1:0] signals. CMD[5:3] signals must be driven with the index to the GMR pair for bits [135:0] of the data being searched. CMD[8:6] signals must be driven with the bits that will be driven on SADR[21:19] by this device if it has a hit. DQ[67:0] must be driven with the 68-bit data ([135:68]) to be compared to all locations “2” in the four 68-bits-word page. The CMD[2] signal must be driven to logic '0.'

– **Cycle D:** The host ASIC continues to drive CMDV high and applies SEARCH command

code (10) on CMD[1:0]. CMD[8:6] signals must be driven with the index of the SSR that will be used for storing the address of the matching entry and hit flag (see Table 9, page 19 for a description of SSR[0:7]). The DQ[67:0] is driven with the 68-bit data ([67:0]) to be compared to all locations “3” in the four 68-bits-word page. CMD[5:2] is ignored because the LEARN Instruction is not supported for x272 tables.

Note: For 272-bit searches, the host ASIC must supply four distinct 68-bit data words on DQ[67:0] during Cycles A, B, C, and D. The GMR Index in Cycle A selects a pair of GMRs that apply to DQ data in Cycles A and B. The GMR Index in Cycle C selects a pair of GMRs that apply to DQ data in Cycles C and D.

The SEARCH command is a pipelined operation that executes searches at one-fourth the rate of the frequency of CLK2X for 272-bit searches in x272-bit-configured tables. The latency of SADR, CE_L, ALE_L, WE_L, SSV, and SSF from the 272-bit SEARCH command (measured in CLK cycles) from the CLK2X cycle that contains the C and D cycles is shown in Table 33, page 44.

The timing diagram for all SRAM interface signals, SSV, and SSF shift to the right for different values of TLSZ, as specified in Table 31, page 44.

In addition, SSV and SSF shift to the right for different values of HLAT, as specified in Table 32, page 44.

In the 272-bit configuration, SEARCH takes two CLK cycles. The result of the SEARCH operation appears as an SRAM READ Cycle with a pipelined latency measured from the second cycle of the command, as specified in Table 33, page 44.

Figure 30. Hardware Diagram for a Table with One Device (272-bit Operation)




Figure 31. 272-Bit Configuration SEARCH Timing Diagram (One Device)




Table 31. Right-Shift of 272-bit Signals for TLSZ Values

TLSZ	Number of CLK Cycles
00	0
01	1
10	2

Table 32. Shift of SSF and SSV from SADR (for different HLAT Values)

HLAT	Number of CLK Cycles
000	0
001	1
010	2
011	3
100	4
101	5
110	6
111	7

Table 33. Latency of SEARCH from Instruction to SRAM Access Cycle (272-bit Mode)

# of devices	Max Table Size	Latency in CLK Cycles
1 (TLSZ = 00)	4K x 272-bit	4
2–8 (TLSZ = 01)	32K x 272-bit	5
9–31 (TLSZ = 10)	124K x 272-bit	6

272-bit Logical SEARCH

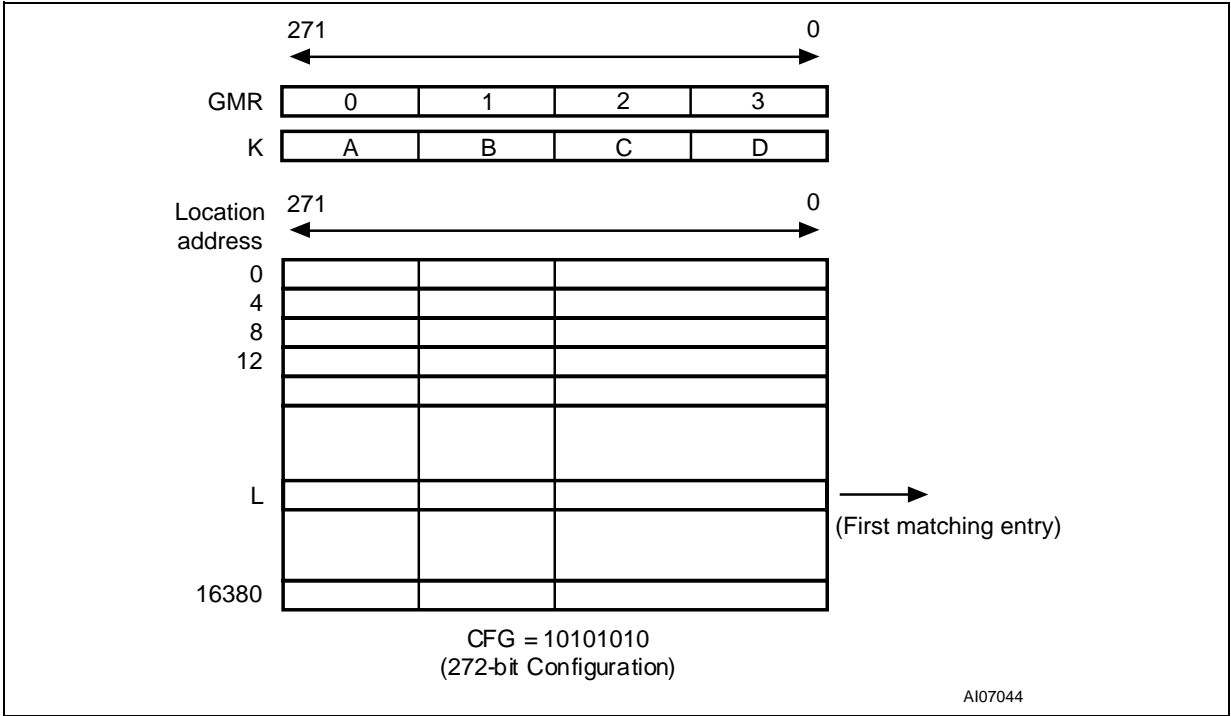
The logical 272-bit SEARCH operation is shown in Figure 32. The entire table of 272-bit entries is compared to a 272-bit word K (presented on the DQ Bus in both Cycles A, B, C, and D of the command) using the GMR and the local mask bits. The GMR is the 272-bit word specified by the two pairs of GMRs selected by GMR Indexes in the command's Cycles A and C. The 272-bit word K (presented on the DQ Bus in Cycles A, B, C, and D of

the command) is compared with each entry in the table starting at location "0."

The first matching entry's location, address "L," is the winning address that is driven as part of the SRAM address on the SADR[21:0] lines.

Note: The matching address is always going to be location "0" in a four-entry page for a 272-bit SEARCH (two LSBs of the matching index will be "00").

Figure 32. x272 Table with One Device



Mixed-sized Searches on Tables Configured with Different Width Using an M7010R Device

This subsection will cover mixed searches (x68, x136, and x272) with tables of different widths (x68, x136, and x272). The sample operation shown is for a single device with CFG = 10010000, containing three tables of x68, x136, and x272 widths. The operation can be generalized to a block of 8 to 31 devices using four blocks; the timing and the pipeline operation is the same as described previously for fixed searches on a table of one-width-size.

Figure 34, page 47 shows three sequential searches; first, a 68-bit SEARCH on a table configured as x68, then a 136-bit SEARCH on a table configured as x136, and finally a 272-bit SEARCH on the table configured as x272. Each results in a hit.

Note: The DQ[67:66] will be “00” in each of the two A and B Cycles of the x68-bit SEARCH (Search1). DQ[67:66] is “01” in each of the A and B Cycles of the x136-bit SEARCH (Search2). DQ[67:66] is “10” in each of the A, B, C, and D Cycles of the x272-bit SEARCH (Search3). By having table designation bits, the M7010R device enables the creation of many tables of different widths in a bank of search engines.

Figure 33 shows the sample table. Two bits in each 68-bit entry need to be designated as table number bits. One example choice might be: the “00” values for the table configured as x68, “01” values for tables configured as x136, and “10” values for tables configured as x272. For the above explanation, it is further assumed that bits for DQ[67:66] for each entry will be designed as such table designation bits.

Figure 33. Multiwidth Configuration Example





Figure 34. Timing Diagram for Mixed SEARCH (One Device)



LRAM and LDEV Description

When search engines are cascaded using multiple M7010R devices, the SADR, CE_L, and WE_L (tri-state signals) are all tied together. To eliminate external pull-up and pull-downs, one device in a bank is designated as the default driver. For non-SEARCH or non-LEARN cycles (see LEARN COMMAND, page 48) or SEARCH cycles with a global miss, the SADR, CE_L, and WE_L signals are driven by the device with the LRAM Bit set. It is important that only one device in a bank of cascaded search engines have this bit set. Failure to do so will cause contention on SADR, CE_L, and WE_L, and can potentially cause damage to the device(s).

LEARN COMMAND

Bit [0] of each 68-bit data location specifies whether an entry in the database is occupied. If all the entries in a device are occupied, the device asserts FULO signal to inform the downstream devices that it is full.

The result of this communication between depth-cascaded devices determines the global FULL signal for the entire table. On a miss by the SEARCH (signalled to the ASIC through the SSV and SSF signals [SSV = 1, SSF = 0]), the host ASIC can apply the LEARN command to learn the entry from a comparand register to the next-free location (see The NFA Register, page 24). The NFA updates to the next-free location following each WRITE or LEARN command.

In a depth-cascaded table, only a single device will learn the entry through the application of a LEARN Instruction. The determination of the LEARN device is based on the FULL and FULO signalling between the devices. The first non-full device learns the entry by storing the contents of the specified comparand registers to the location(s) pointed to by the NFA.

In a x68-configured table, the LEARN command writes a single 68-bit location. In a 136-bit-configured table, the LEARN command writes the next even and odd 68-bit locations. In 136-bit mode, Bit[0] of the even and odd 68-bit locations is '0,' indicating that they are cascaded empty, or '1,' which indicates that they are occupied.

The global FULL signal indicates to the table controller (the host ASIC) that all entries within a block are occupied and that no more entries can be learned. The M7010R device updates the signal to

Similarly, when search engines using multiple M7010R devices are cascaded, SSF and SSV (also tri-state signals) are tied together. To eliminate external pull-up and pull-downs, one device in a bank is designated as the default driver. For non-SEARCH or SEARCH cycles with a global miss, the SSF and SSV signals are driven by the device with the LRAM Bit set. It is important that only one device in a bank of cascaded search engines have this bit set. Failure to do so will cause contention on SSF and SSV, and can potentially cause damage to the device(s).

a data array after each WRITE or LEARN command. Also using the NFA Register as part of the SRAM address, the LEARN command generates a WRITE cycle to the external SRAM.

The LEARN command is supported on a single block containing up to eight devices if the table is configured as either a x68 or a x136. The LEARN command is not supported for x272-configured tables.

The LEARN operation lasts two CLK cycles. The sequence of this operation is as follows:

- **Cycle 1A:** The host ASIC applies the LEARN Instruction on CMD[1:0] using CMDV = 1. The CMD[5:2] field specifies the index of the comparand register pair that will be written to the data array in the 136-bit-configured table. For a LEARN in a 68-bit-configured table, the even-numbered comparand specified by this index will be written. CMD[8:6] carries the bits that will be driven on SADR[21:19] in the SRAM WRITE cycle.
- **Cycle 1B:** The host ASIC continues to drive CMDV to '1,' CMD[1:0] to '11,' and CMD[5:2] with the comparand pair index. CMD[6] must be set to '0' if the LEARN is being performed on a 68-bit-configured table, and to '1' if the LEARN is being performed on a 136-bit-configured table.
- **Cycle 2:** The host ASIC drives CMDV to '0.'

At the end of Cycle 2, a new instruction can begin. SRAM WRITE latency is the same as the SEARCH to the SRAM READ cycle measured from the second cycle of the LEARN Instruction.

LEARN is a pipelined operation and last for two CLK cycles where TLSZ = 00, as shown in Figure 35, page 49, and TLSZ = 01, as shown in Figure 36, page 50 and Figure 37, page 51. Figure 36 and Figure 37 assume that the device performing the LEARN operation is not the last device in the table and has its LRAM Bit set to '0.'

Note: The OE_L for the device with the LRAM Bit set goes high for two cycles for each LEARN (one during the SRAM WRITE cycle, and one during the cycle before it). The latency of the SRAM WRITE cycle from the second cycle of the instruction is shown in Table 34, page 51.

Figure 35. LEARN Command Timing Diagram (TLSZ = 00)




Figure 36. LEARN Timing Diagram (TLSZ = 1, except on Last Device)

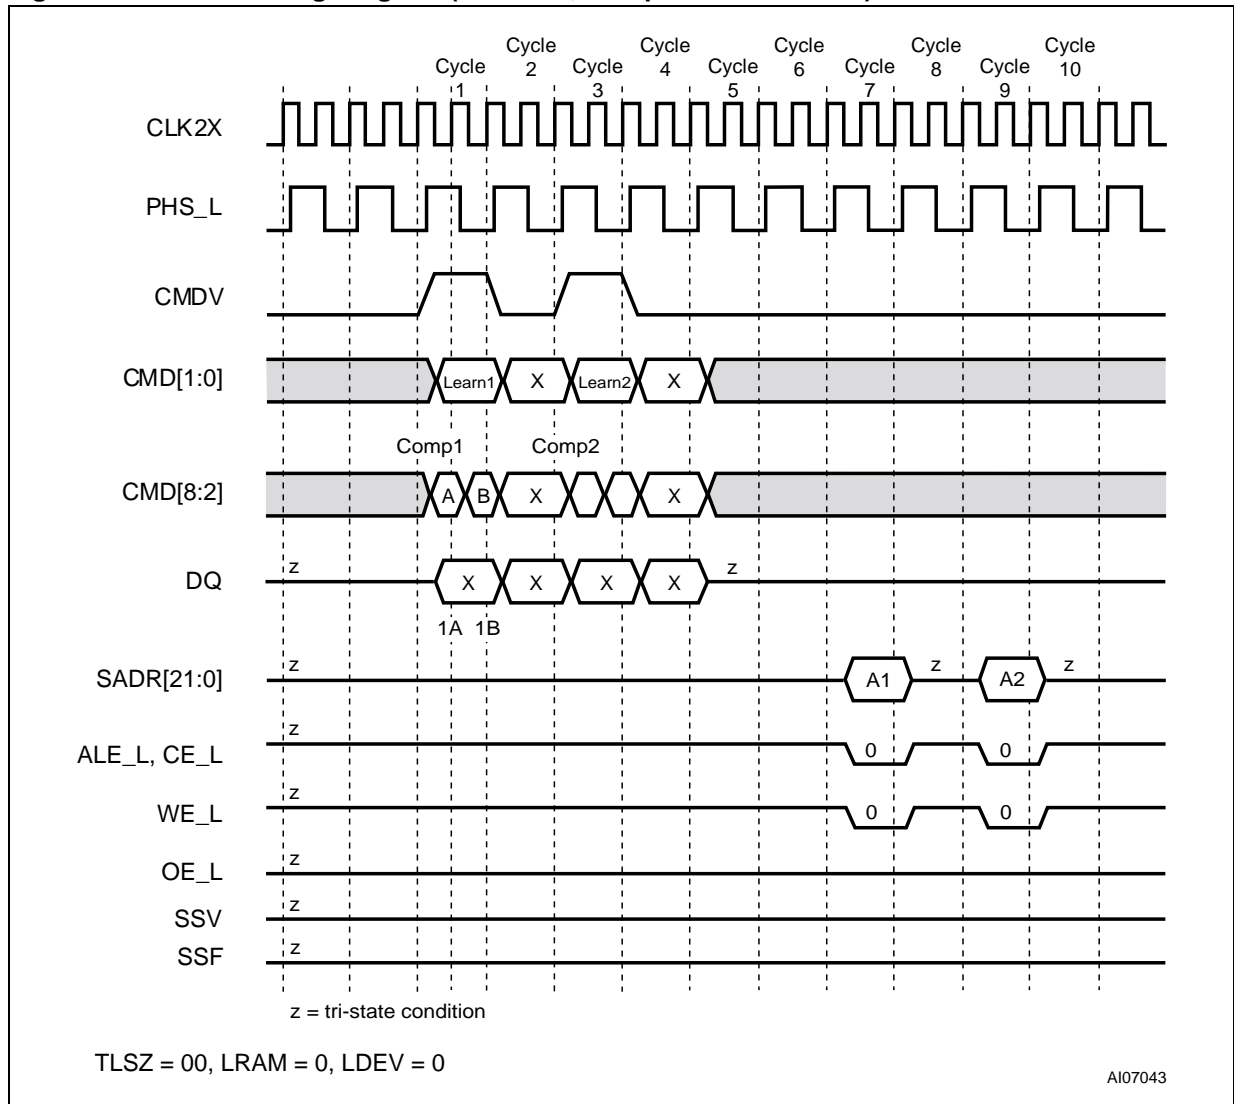


Figure 37. LEARN Timing Diagram on Device Number 7 (TLSZ = 01)

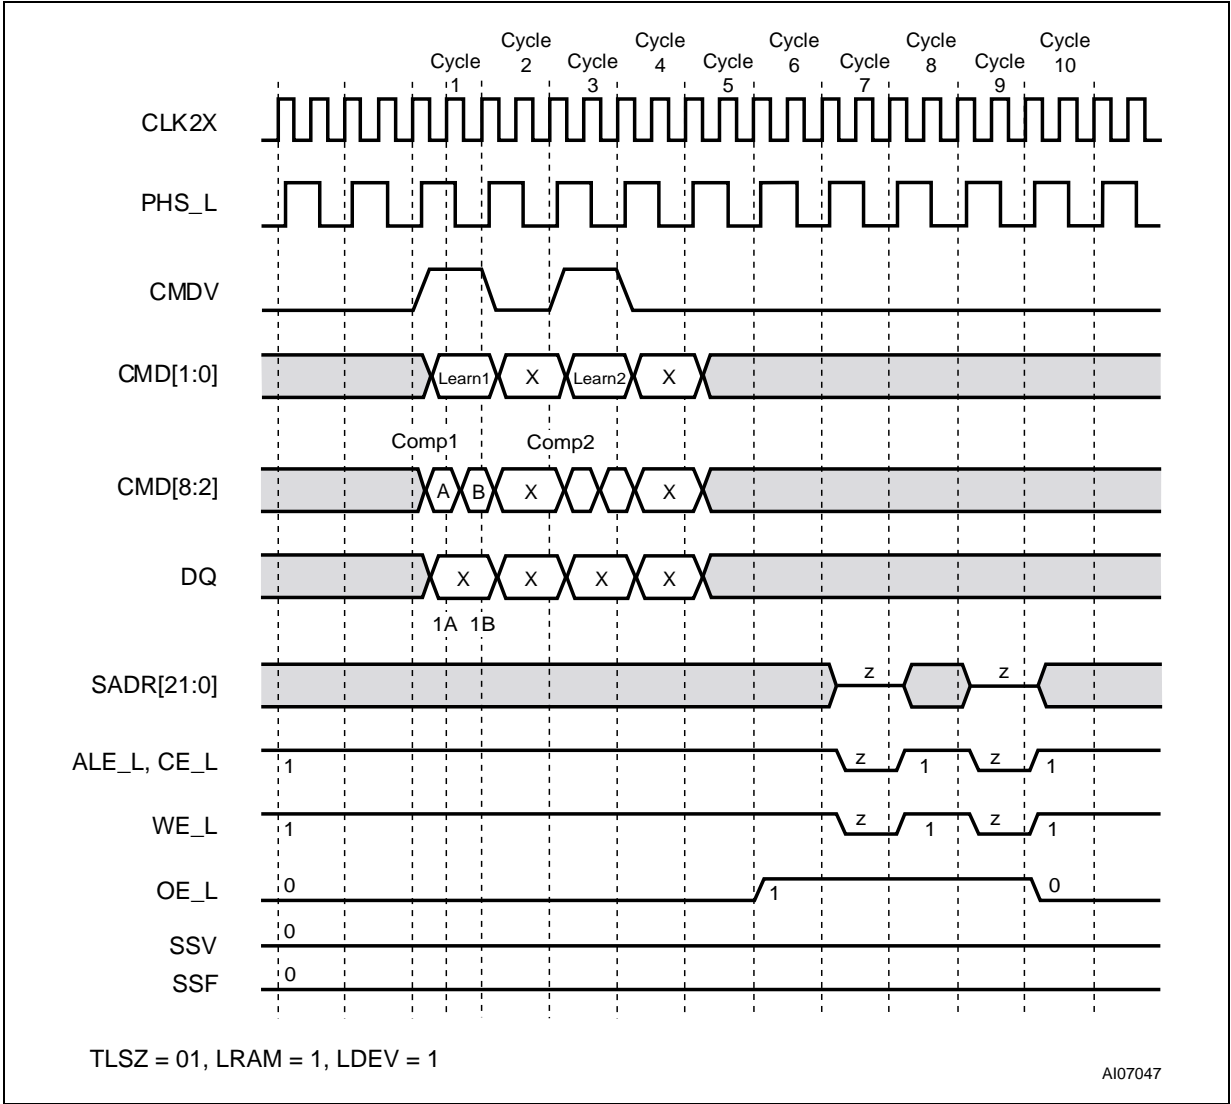


Table 34. SRAM WRITE Cycle Latency from Second Cycle of LEARN Instruction

Number of Devices	Latency in CLK Cycles
1 (TLSZ = 00)	4
1-8 (TLSZ = 01)	5
1-31 (TLSZ = 10)	6

DEPTH-CASCADING

The Search Engine application can depth-cascade the device to various table sizes in 68-bit, 136-bit, and 272-bit configurations by programming the table size (TLSZ) field of the Command Register. The devices perform all the necessary arbitration to decide which device drives the SRAM Bus. The latency of the searches increases as the table size increases while the search rate remains constant.

Depth-Cascading Up to Eight Devices (One Block)

Figure 38, page 53 shows how up to eight devices can cascade to form a 128K x68-bit, 64K x136-bit, or 32K x272-bit table. It also shows the interconnection between the devices for depth-cascading. The host ASIC must program the table size (TLSZ) field to '01.' Each Search Engine asserts the LHO[1] and LHO[0] signals to inform downstream devices in the cascade of its results. The LHI[6:0] signals for any device are connected to the LHO signals of the upstream device. A single device alone drives the SRAM bus in any given cycle.

Depth-Cascading Up to 31 Devices (4 Blocks)

Figure 39, page 54 shows how to cascade up to four blocks. Each block contains up to eight M7010Rs (except the last block, which contains 7 devices), to form a 496K x68, 248K x136, or 124K x272 table. Note the interconnection between blocks for depth-cascading. The host ASIC must program the table size (TLSZ) field to 10 for cascading 8 to 31 devices (in up to four blocks). For each search, a block asserts BHO[2], BHO[1], and BHO[0]. The BHO[2:0] signals for a block are only taken from the last device in the block. See Figure 41, page 56 for the arbitration cycle between

blocks to determine which device drives the SRAM Bus.

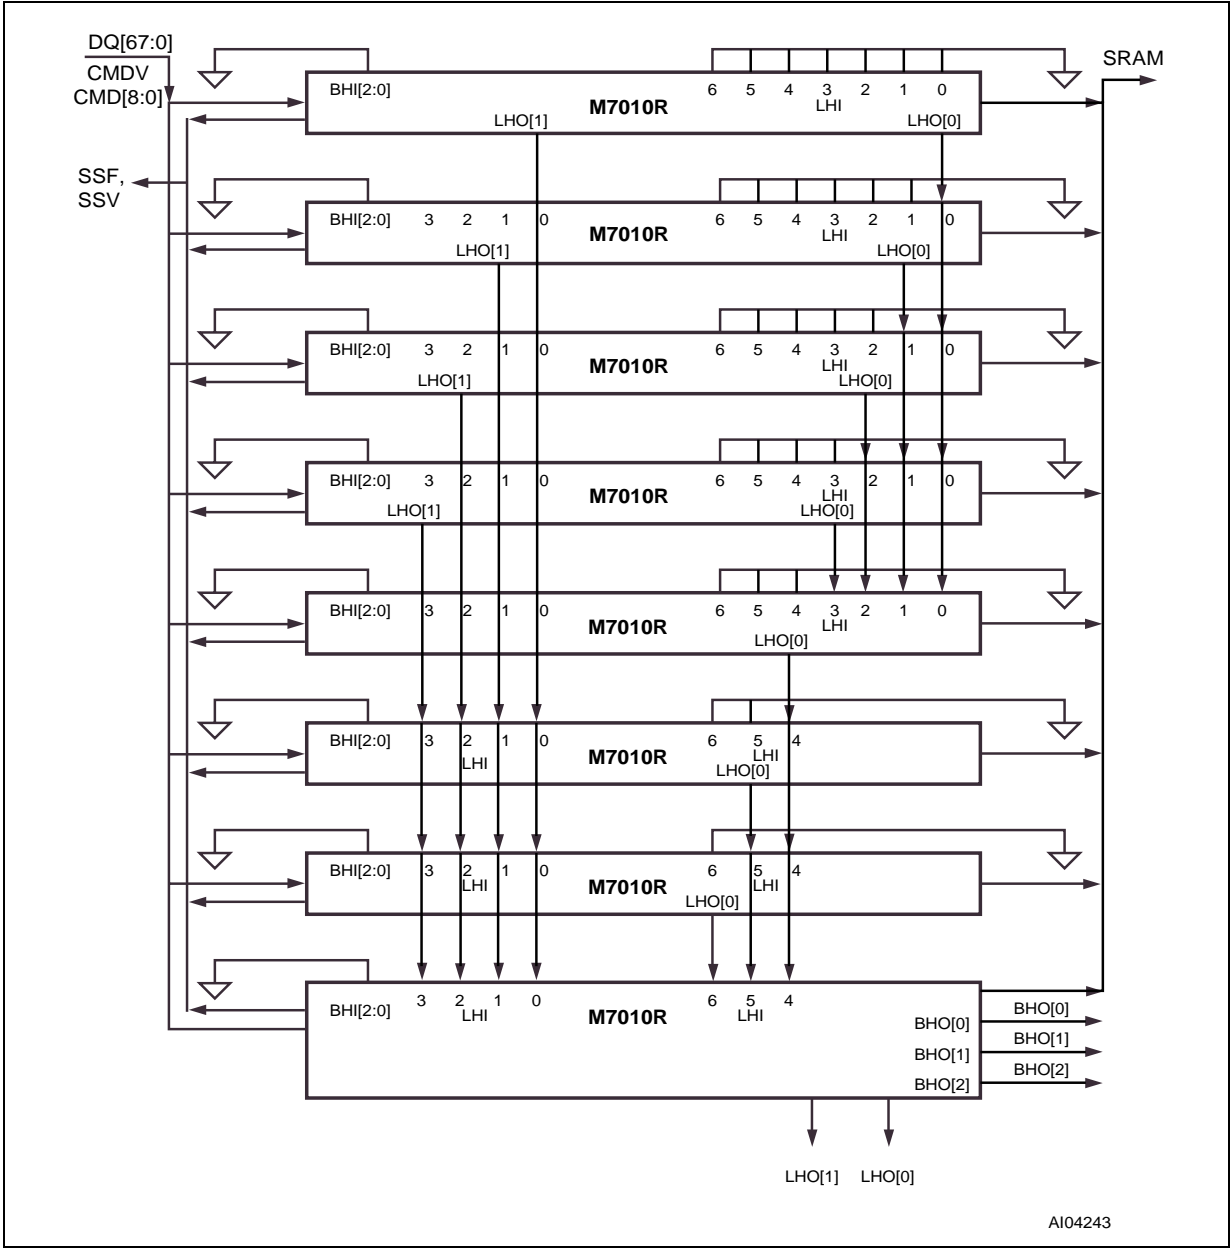
The device is configured to be the last in the depth-cascaded table by setting LDEV to 1 in the Command Register. The device with LDEV set to 1 drives the SSF and SSV signals in cycles when all upstream devices do not drive these signals. The M7010R with its LDEV Bit set drives SSF and SSV during a search with a miss or with non-search commands. See the LDEV Bit definition in Table 10, page 20.

Depth-Cascading to Generate a "FULL" State for a Block

Bit[0] of each of the 68-bit entries is designated as a special bit (1 = FULL; 0 = Empty). For each LEARN or PIO WRITE to the data array, each device asserts FULO[1] and FULO[0] if it does not have any empty locations (see Figure 40, page 55). Each device combines the FULO signals from the devices above it with its own full status to generate a FULL signal, which will then give a "full" status of the table up to the device asserting the FULL signal. Figure 40, page 55 shows the hardware connection diagram for generating the FULL signal that goes back to the ASIC. In a depth-cascaded block of up to eight devices, the FULL signal from the last device should be fed back to the ASIC controller to indicate the fullness of the table. The FULL signal of the other devices should be left open.

Note: The LEARN Instruction is supported for up to eight devices, whereas FULL cascading is allowed for one block in tables containing more than eight devices. In tables for which a LEARN Instruction will not be used, the Bit[0] of each 68-bit entry should always be set to '1.'

Figure 38. Depth-Cascading to Form a Single Block (8 Devices)



A104243

Figure 39. Four Blocks (31 Devices Cascaded) SEARCH, 68-bit Configured with LDEV = 1


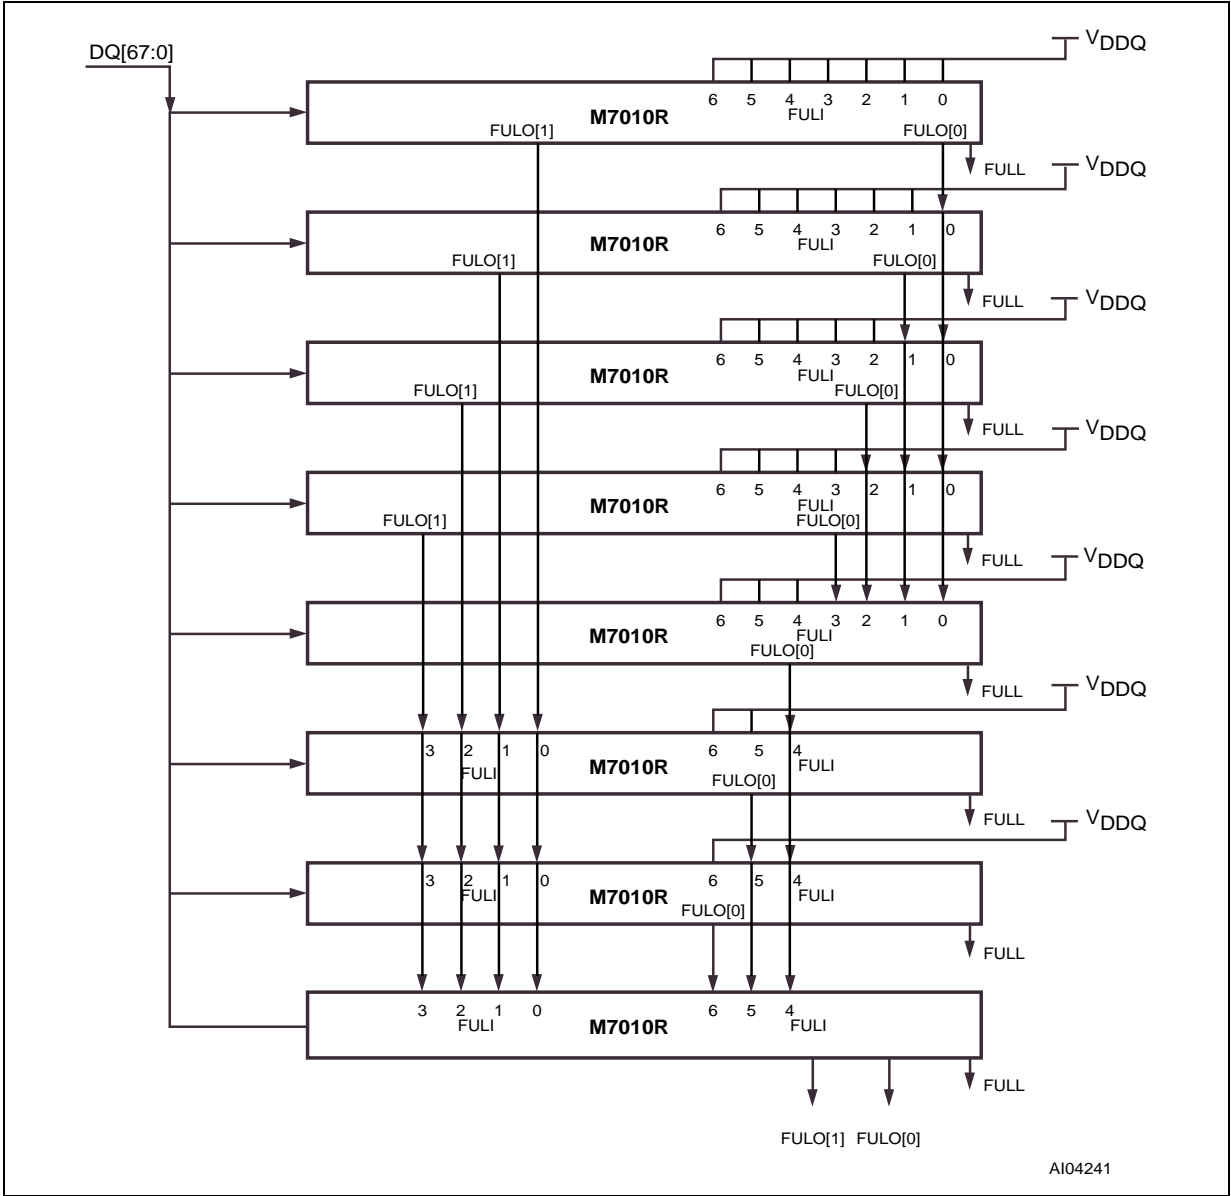


Figure 40. "FULL" State Generation in a Cascaded Table



ARBITRATION

Figure 41, page 56 shows an example of the arbitration cycle for determining which device drives the SRAM Bus in a single block, up to eight M7010Rs with 136-bit configuration settings.

Four cycles from the SEARCH command, all M7010Rs are informed of the SEARCH result within the device and drive their LHO signals. At the next cycle, all downstream devices know the outcome of the SEARCH in all the upstream devices.

If any of the upstream devices has a hit, all the subsequent devices defer driving the SRAM Bus. If a SEARCH failure occurs, the M7010R with the LRAM Bit set (the last in the chain) drives the SRAM Bus signals. The device with LDEV set to '1' is the default driver of the SSV and SSF signals. Figure 42, page 57 shows how an M7010R arbitrates accesses to the SRAM.

Figure 41. Timing Diagram for Arbitration Within a Block





Figure 42. Timing for Arbitration for Two or More Blocks for the Last Device



SRAM ADDRESSING

Table 35, page 61 lists and describes the commands used to generate addresses on the SRAM address bus. The Index[13:0] field contains the address of a 68-bit entry that results in a hit in 68-bit configured quadrant. It is the address of the 68-bit entry that lies at the 136-bit page and 272-bit page boundaries in 136-bit and 272-bit configured quadrants, respectively.

The register section of this specification describes the NFA and SSR registers. Adr[13:0] contains the address supplied on the DQ Bus during PIO access to the M7010R. Command Bits 8 and 7, CMD[8:6] are passed from the command to the SRAM address bus. See COMMAND CODES AND PARAMETERS, page 27 for more information.

SRAM PIO Access

SRAM READ. Enables READ access to the off-chip SRAM that contains associative data. The latency from the issuance of the READ Instruction to the address appearing on the SRAM bus is the same as the latency of the SEARCH Instruction, and will depend on the value programmed for the TLSZ parameter in the device configuration register. The latency of the ACK from the READ Instruction is the same as the latency of the SEARCH Instruction to the SRAM address plus the HLAT programmed into the configuration register.

Note: SRAM READ is a blocking operation - no new instruction can begin until the ACK is returned by the selected device performing the access.


The following explains the SRAM READ operation in a table with only one device and having the following parameters: TLSZ = 00, HLAT = 000, LRAM = 1, and LDEV = 1. Figure 43, page 59 shows the associated timing diagram. For the fol-

lowing description, the selected device refers only to the device in the table because it is the only device to be accessed.

- **Cycle 1A:** The host ASIC applies the READ Instruction on CMD[1:0] using CMDV = 1. The DQ Bus supplies the address, with DQ[20:19] set to “10,” to select the SRAM address. The host ASIC selects the device for which the ID[4:0] matches the DQ[25:21] lines. During this cycle, the host ASIC also supplies SADR[21:19] on CMD[8:6].
- **Cycle 1B:** The host ASIC continues to apply the READ Instruction on CMD[1:0] using CMDV = 1. The DQ Bus supplies the address with DQ[20:19] set to “10” to select the SRAM address.
- **Cycle 2:** The host ASIC floats DQ[67:0] to a tri-state condition.
- **Cycle 3:** The host ASIC keeps DQ[67:0] in a tri-state condition.
- **Cycle 4:** The selected device starts to drive DQ[67:0] and drives ACK from High-Z to LOW.
- **Cycle 5:** The selected device drives the READ address on SADR[21:0]; it also drives ACK HIGH, CE_L LOW, and ALE_L LOW.
- **Cycle 6:** The selected device drives CE_L HIGH, ALE_L HIGH, the SADR Bus and DQ Bus in a tri-state condition, and ACK LOW.

At the end of Cycle 6, the selected device floats ACK in a tri-state condition, and a new command can begin. Table 36, page 62 shows by how many cycles SRAM signals shift to the right for various TLSZ values. Table 37, page 62 shows by how many cycles SRAM signals shift to the right for various HLAT values.

Figure 43. SRAM READ Access for One M7010R Device



SRAM WRITE. Enables WRITE access to the off-chip SRAM containing associative data. The latency from the second cycle of the WRITE Instruction to the address appearing on the SRAM bus is the same as the latency of the SEARCH Instruction, and will depend on the TLSZ value parameter programmed into the device configuration register.

Note: SRAM WRITE is a pipelined operation - new instruction can begin right after the previous command has ended. The following explains the SRAM WRITE operation accomplished through a table of only one device with the following parameters: TLSZ = 00, HLAT = 000, LRAM = 1, and LDEV = 1. Figure 44, page 61 shows the timing diagram. For the following description, the selected device refers to the only device in the table as this is the only device that will be accessed.

– **Cycle 1A:** The host ASIC applies the WRITE Instruction on CMD[1:0] using CMDV = 1. The DQ Bus supplies the address, with DQ[20:19] set to “10,” to select the SRAM address. The host ASIC selects the device for which the ID[4:0] matches the DQ[25:21] lines. The host ASIC also supplies SADR[21:19] on CMD[8:6] in this cycle.

Note: CMD[2] must be set to '0' for SRAM WRITE, because burst WRITES into the SRAM are not supported.

– **Cycle 1B:** The host ASIC continues to apply the WRITE Instruction on CMD[1:0] using CMDV = 1. The DQ Bus supplies the address with DQ[20:19] set to “10” to select the SRAM address.

Note: CMD[2] must be set to '0' for SRAM WRITE, because burst WRITES into the SRAM are not supported.

– **Cycle 2:** The host ASIC continues to drive DQ[67:0]. The data in this cycle is not used by the M7010R.

– **Cycle 3:** The host ASIC continues to drive DQ[67:0]. The data in this cycle is not used by the M7010R.

At the end of Cycle 3, a new command can begin. The WRITE is a pipelined operation; however, the WRITE cycle appears at the SRAM bus with the same latency as the SEARCH Instruction (as measured from the second cycle of the WRITE command).

Figure 44. SRAM WRITE Access for One M7010R Device




Table 35. SRAM Bus Address Generation

Command	SRAM Operation	21	20	19	[18:15]	[14:0]
SEARCH	Read	C8	C7	C6	ID[4:0]	Index[14:0]
LEARN	Write	C8	C7	C6	ID[4:0]	NFA[14:0]
PIO READ	Read	C8	C7	C6	ID[4:0]	Adr[14:0]
PIO WRITE	Write	C8	C7	C6	ID[4:0]	Adr[14:0]
Indirect Access	Write/Read	C8	C7	C6	ID[4:0]	SSR[14:0]

Table 36. Right-Shift of SRAM Signals for TLSZ Values

TLSZ	Number of CLK Cycles
00	0
01	1
10	2

Table 37. Right-Shift of SRAM Signals for HLAT Values

HLAT	Number of CLK Cycles
000	0
001	1
010	2
011	3
100	4
101	5
110	6
111	7

JTAG (1149.1) TESTING

The M7010R supports the Test Access Port (TAP) and Boundary Scan Architecture as specified in the IEEE JTAG standard 1149.1. The pin interface to the chip consists of five signals with the standard definitions: TCK, TMS, TDI, TDO, and TRST_L. Table 38, page 62 describes the opera-

tions that the test access port controller supports. Table 39 shows the TAP Device ID Register.

Note: To disable JTAG functionality, connect the TCK, TMS, and TDI pins to V_{DDQ} through a pull-up, and the TRST_L to ground through a pull-down.

Table 38. Test Access Port Controller Instructions

Instruction	Type	Description
SAMPLE/PRELOAD	Mandatory	Sample/Preload. Loads the values of signals going to and from IO pins into the boundary scan shift register to provide a snapshot of the normal functional operation.
EXTEST	Mandatory	External Test. Uses boundary scan values shifted in from TAP to test connectivity external to the device.
INTEST	Optional	Internal Test. Allows slow-speed, functional testing of the device using the boundary scan register to provide the I/O values.

Table 39. TAP Device ID Register

Field	Range	Initial Value	Description
Revision	[31:28]	0001	Revision Number. This is the current device revision number. Numbers start from one and increment by one for each revision of the device.
Part #	[27:12]	0000 0000 0000 0001	This is the part number for this device.
MFID	[11:1]	000_1101_1100	Manufacturer ID. This field is the same as the manufacturer ID used in the TAP controller.
LSB	[0:0]	1	Least Significant Bit

POWER DISTRIBUTION GUIDELINE

In order to prevent voltage supply sags that can potentially degrade device performance, large bypass capacitors are often recommended. Since the bulk storage not only contains an effective series resistance, but also a fairly high inductance, the large bypass capacitors should be assisted by other capacitors that have a lower inductance (but typically less capacitance). These high frequency capacitors control the switching transients and hold-over the power planes during an average load change until the higher inductance capacitors can react. High frequency bypass capacitors are used having values of 0.01 μ F and 0.1 μ F.

For a single Search Engine application, a recommended power plane and ground plane may be laid out as follows:


- A 1000 μ F bulk capacitor is recommended for the 1.8V V_{DD} source supply.

- A 100 μ F bulk capacitor is recommended for the 3.3V V_{DDQ} source supply.
- Four sets of 0.1 μ F and 0.01 μ F high frequency capacitors are recommended between V_{DD} , V_{DDQ} and ground.

Multiple bulk and high frequency capacitors may also be required. Users can determine the values of such capacitors after computing, based on their system and power supply environment. The device should achieve the search performance as specified with the bypass capacitors.

This application note is a general guideline for Search Engine Design. For more detailed information, please refer to Intel Website for Appnote AP-912, Pentium III Xeon Processor Power Distribution Guidelines. (<http://support.intel.com/design/pentiumiii/xeon/aplnots/245095.htm>).

Figure 45. Network Search Engine Power Distribution



PART NUMBERING

Table 40. Ordering Information Scheme


Example:	M70	10	R	-083	ZA	1	T
Device Type	M70 Search Engine						
Density	10 = 1Mb (16K x 68-bit Table Entries)						
Operating Supply Voltage	R = $V_{DD} = 1.8V$						
Speed	-083 = 83 Million Searches per Second -066 = 66 Million Searches per Second						
Package	ZA = PBGA, 272-count, 27mm x 27mm ⁽¹⁾						
Temperature Range	1 = 0 to 70 °C						
Shipping Option	Tape & Reel Packing = T						

Note: 1. Where "Z" is the symbol for BGA packages and "A" denotes 1.27mm ball pitch

For a list of available options (e.g., Speed, Package) or for further information on any aspect of this device, please contact the ST Sales Office nearest to you.

PACKAGE MECHANICAL INFORMATION

Figure 46. PBGA-Z00 – 272-ball Plastic Ball Grid Array Package Outline



Note: Drawing is not to scale.

Table 41. PBGA-Z00 – 272-ball Plastic Ball Grid Array Package Mechanical Data

Symb	mm			inches		
	Typ	Min	Max	Typ	Min	Max
A ⁽⁴⁾	27.00	26.80	27.20 ⁽¹⁾	1.102	1.094	1.110 ⁽¹⁾
A1 ^(2,3)	0.60	0.50	0.70	0.024	0.020	0.029
A2		1.63	1.90		0.067	0.078
B ⁽⁴⁾	27.00	26.80	27.20	1.102	1.094	1.110
b	0.75	0.60	0.90	0.031	0.024	0.037
D	27.00	26.80	27.20	1.102	1.094	1.110
D1	24.13			0.985		
D2	24.00			0.980		
E	27.00	26.80	27.20	1.102	1.094	1.110
E1	24.13			0.985		
E2	24.00			0.980		
e	1.27			0.052		
ddd			0.20			0.008

- Note: 1. Maximum mounted height is 2.45mm based on a 0.65mm ball pad diameter. Solder paste is 0.15mm thickness and 0.65mm in diameter.
 2. The terminal A1 corner must be identified on the top surface by using a corner chamfer, ink, or metallized markings, or other feature of package body or integral heatslug.
 3. A distinguished feature is allowable on the bottom surface of the package to identify the terminal A1 corner.
 4. Exact shape of each corner is optional.

REVISION HISTORY

Table 42. Document Revision History

Date	Rev. #	Revision Details
January 2002	1.0	First Issue
07/23/02	1.1	Changes after extensive review (Figures 3, 8, 11, 12, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 43, 44; Tables 6, 7, 9, 10, 12, 17, 19, 22, 26, 27, 29, 30, 32, 33, 34, 35)

Information furnished is believed to be accurate and reliable. However, STMicroelectronics assumes no responsibility for the consequences of use of such information nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent or patent rights of STMicroelectronics. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. STMicroelectronics products are not authorized for use as critical components in life support devices or systems without express written approval of STMicroelectronics.

The ST logo is registered trademark of STMicroelectronics
All other names are the property of their respective owners.

© 2002 STMicroelectronics - All Rights Reserved

STMicroelectronics GROUP OF COMPANIES
Australia - Brazil - Canada - China - Finland - France - Germany - Hong Kong - India - Israel - Italy - Japan - Malaysia -
Malta - Morocco - Singapore - Spain - Sweden - Switzerland - United Kingdom - U.S.A.

www.st.com