Application Note 001: Initialization of LNI7010/LNI7020 Network Database Search Engines on Power-up

LNI-0401-2*1-A-27

1 INTRODUCTION

Lara Networks, Inc.'s (Lara's) LNI7010/LNI7020 devices incorporate leading edge Associative Processing Technology ™ (APT) and are designed to be high-performance, pipelined, synchronous, 32K-entry Network Database Search Engines (NDSEs). The LNI7010/LNI7020 database entry size is configurable as 34 bits, 68 bits, 136 bits, or 272 bits. In 34-bit entry mode, the size of the database is 32K entries. In 68-bit entry mode, the size of the database is 16K entries. In 136-bit entry mode, database size is 8K, and in 272-bit mode, the database size is 4K. The LNI7010/LNI7020 devices are programmable to support multiple databases with varying entry sizes.

This application note describes how to initialize NDSEs on power-up and loading into LNI7010/LNI7020 devices.

1.1 Functional Description

The following subsections outline the functional description of the NDSEs.

1.2 Initializing the LNI7010/LNI7020 NDSEs

Data is read or written to the devices via the DQ[67:0] bus (see Figure 1). The command (CMD) bus CMD[8:0] determines the operations, including READ, WRITE, SEARCH, and LEARN. During the same cycle the data array, the mask array, the external SRAM, or the register file can be selected by supplying an index using DQ[21:19] for the LNI7010 device and DQ[21:20] for the LNI7020 device. The SEARCH index appears on the SADR[21:0] bus.

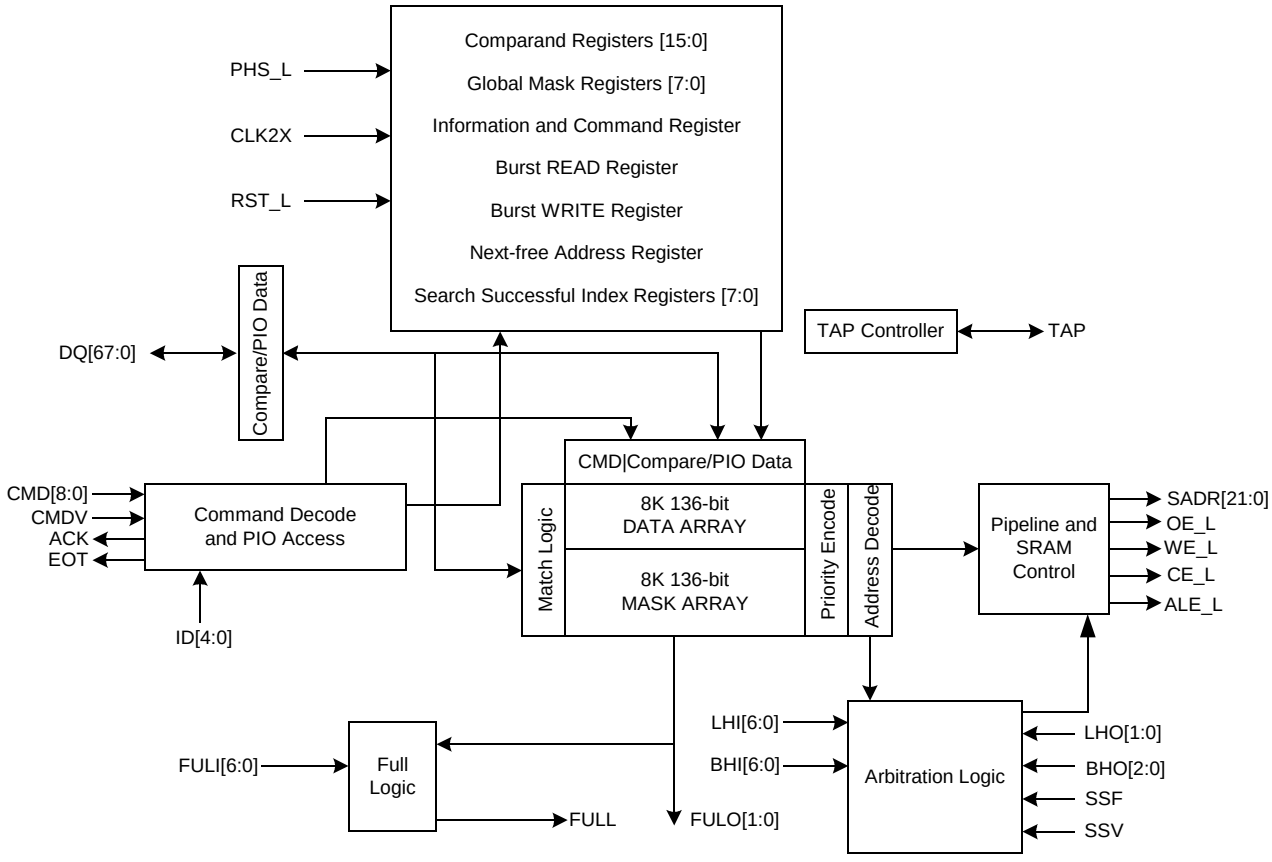

FIGURE 1. LNI7010 BLOCK DIAGRAM

A large register set allows the user to store multiple comparand and SEARCH results on the NDSE instead of at the host. In addition to single-address READ and WRITE operations, the burst READ and burst WRITE registers allow bursts of up to 512 operations.

The example shown in this application note is for 136-bit lookups, but the device can easily be configured and used for 68-bit and 272-bit lookups.

2 CONFIGURATION IN SINGLE-DEVICE MODE

There are three steps for device configuration, as follows:

1. Interconnection
2. Device reset
3. Command register initialization.

2.1 Interconnection

All bits of the LNI and BHI input bus are connected to ground when an NDSE is configured to be in single-device mode (see Figure 2). All bits of FULI input are connected to V_{CC}.

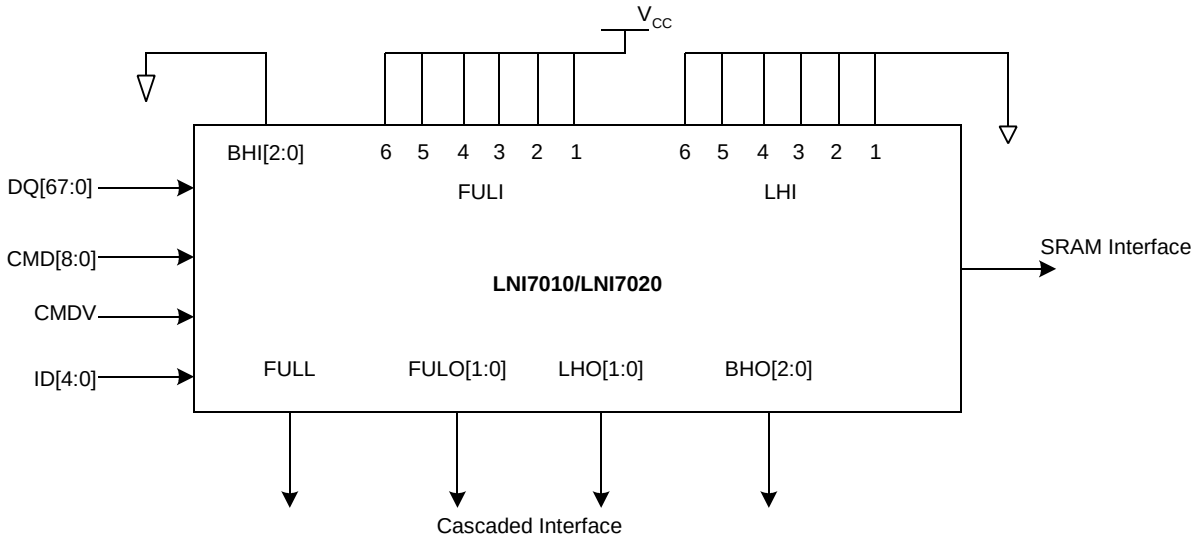

FIGURE 2. LNI7010/LNI7020 IN SINGLE-DEVICE MODE

2.2 Reset the NDSE

There are two ways of initializing the NDSE. The first is a hardware reset using RST_L; the second is a software reset using the command register bit (SRST).

Hardware reset: A hardware reset must be provided for 20 clock cycles once the power is on. Hardware or software resets execute an initialization sequence that is described in the remainder of Section 2.

2.3 Initialize the Command Register

The next step is to initialize the command register with a WRITE command. The host ASIC writes in the NDSE whose ID[4:0] matches DQ[25:21]. If there is no match, the Host writes to all the NDSEs when DQ[25:21] = 11111 (a broadcast address. **Note**. ID[4:0] is a hardware connect in the NDSE.

It takes three operation cycles to WRITE into the command register (as shown in Figure 3). The address is supplied onto the DQ bus in the first cycle, data is supplied in the second cycle, and NOP is the third cycle.

Table 1 describes the commands in the command register. See Table 2 for a typical command register initialization for a 8K x 136 NDSE in single-device mode.

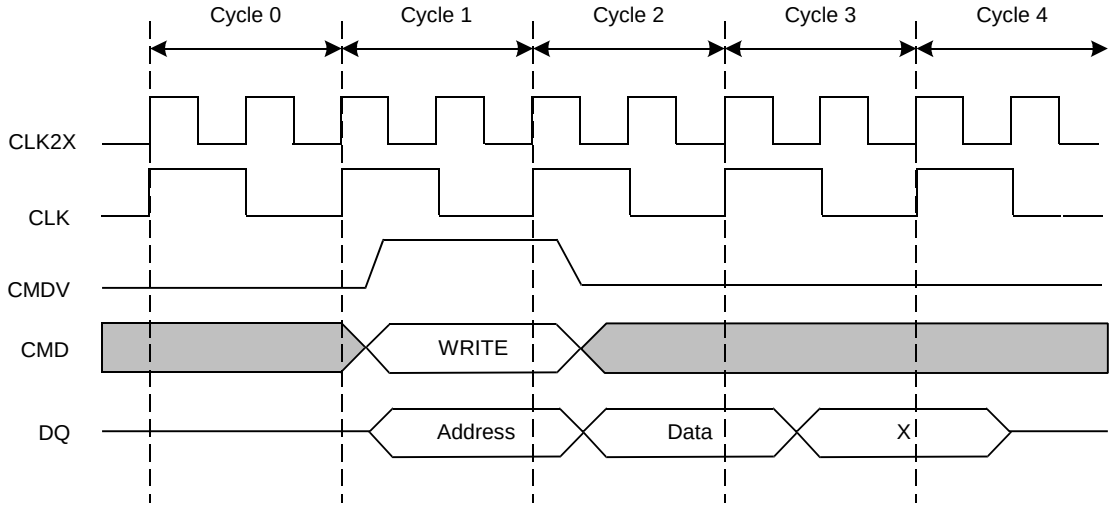

FIGURE 3. WRITE CYCLE TIMING

TABLE 1. COMMAND REGISTER DESCRIPTION FOR DATA CYCLE (CYCLE 2)

		INITIAL VALUE	
SRST	[0]	0	**Software Reset**. A 1 resets the device, with the same effect as the hardware reset. Internally, it generates a reset pulse lasting eight CLK cycles. This bit automatically resets to 0 when 1 is written to it.
DEVE	[1]	0	**Device Enable**. If the device enable reads 0, it does not perform any SEARCH, WRITE, LEARN, or READ operations. It keeps the SRAM bus in a 3-state condition and forces the cascade interface outputs to 0.
TLSZ	[3:2]	01	**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 PIO accesses to the SRAM (SADR[21:0], CE_L, OE_L, WE_L, ALE_L, SSV, SSF, and ACK). Once programmed, the SEARCH latency is constant. Latency #CLK 00: 1 device 4 01: 2-8 devices 5 10: 8-31 devices 6 11: Reserved.

		INITIAL VALUE	
HLAT	[6:4]	000	**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 SRAM PIO READ operations. 000: 0 100: 4 001: 1 101: 5 010: 2 110: 6 011: 3 111: 7.
LDEV	[7]	0	**Last NDSE in Table**. This device is the last NDSE in the depth-cascaded table. In the event of a SEARCH failure, the device with this bit set drives the hit signals as follows: SSF = 0; SSV = 0.
LRAM	[8]	0	**Last Device on this SRAM Bus**. This device is the last NDSE on this SRAM bus. In cycles where an NDSE does not drive the SRAM bus, the device with this bit set drives the SRAM bus (SADR, CD_L, and WE_L) in their inactive state. This bit sets a default driver for the SRAM control signals (SADR, CE_L, SE_L, and OE_L). Note. OE_L is always asserted or disserted. Connect the OE_L from the device with the set LRAM bit to the SRAM bus.
CFG	[16:9]	00000000	**Table Configuration**. The device is internally divided into four quadrants of 4K x 68 bits, each of which can be configured as 4K x 68, 2K x 136, or 1K x 272 as follows: 00: 4K x 446 01: 2K x 136 10: 1K x 272 Bits[10:9] apply to the configuration of the first quadrant in the address space. Bits[12:11] apply to the configuration of the second quadrant in the address space. Bits[14:13] apply to the configuration of the third quadrant in the address space. Bits[16:15] apply to the configuration of the fourth quadrant in the address space.
	[67:17]	0	**Reserved**.

TABLE 2. TYPICAL COMMAND REGISTER INITIALIZATION FOR 8K x 136 NDSE IN SINGLE-DEVICE MODE

CYCLE	CMD[8:0]	CMDV	DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
1	xxxxxx001	1	xxxxxxxxxx	ID[4:0]	11	xxxxxxxxxxxxxx	111000
2	xxxxxxxxx	0	00… …00	00000	00	0001010101110	000010
3	xxxxxxxxx	0	xxxxxxxxxx	xxxxx	xx	xxxxxxxxxxxxxx	xxxxxxx

In the first cycle, DQ[5:0] selects command registers whereas DQ[20:19] selects register access. During the second cycle, the following bits must be asserted or disserted correctly to perform initialization.

1. SRST bit must be set to 0 (DQ[0] in second cycle).
2. Set DEVE.
   - Device must be enabled to drive the SRAM bus as well as cascading output flags. The DEVE line must be set to a logic 1 (DQ[1] in the second cycle) for all devices in a block.
   - Table size (TLSZ), which is indicative of single-device depth, is set to 00 (DQ[3:2] in the second cycle).
   - Latency of hit signals (HLAT), which is indicative of 4-cycle latency, is set to 000 (DQ[6:4] in the second cycle).
3. Configure the last NDSE.
   - The last NDSE in the table (LDEV) is indicative of the position of the last NDSE in the depth-cascaded table.
   - The device with LDEV set to 1 (DQ[7] in the second cycle) drives SSF and SSV signals in a cycle during a SEARCH with a miss or a nonSEARCH command.
   - The last RAM of the table (LRAM) is indicative of the last NDSE on the SRAM bus. If LRAM is set to 1 (DQ[8] in the second cycle), the last device on the SRAM bus drives the SADR, OE_L, WE_L, and ALE_L when none of the upstream devices drive.
   - When the LNI7010 and LNI7020 devices are used in a single-device operation, it drives the SRAM bus whether a hit or a miss occurs.
4. Program the layer.
   - When configuration (CFG) bits CFG[16:9] are set to 0101010101 (DQ[16:9] in the second cycle), it is indicative of 136-bit operation in all four quadrants of a single device.
   - The 8K x 136 bit NDSE is partitioned into four internal quadrants. This allows flexibility in performing multilayer lookups (layer 2 and layer 3) in the same device.
   - The third cycle is NOP.

3 CONFIGURATION IN CASCADED MODE

Each NDSE has a unique ID address. The hard-wired ID[4:0] allows the identification of up to 31 NDSEs that can be cascaded to allow 248K entries of 136 bits. Most practical applications use up to eight cascaded devices.

The command register provides information on whether a READ, WRITE, SEARCH, or LEARN instruction will be executed. The device with the lowest address has the highest priority in the cascaded chain. The last device in the chain must be programmed first in order to drive the SRAM bus interface signals.

To configure a device, use the following steps.

1. Interconnection.
2. Reset the device.
3. Initialize the command register.

3.1 Interconnection

The NDSE application can be depth-cascaded into devices of various table-size configurations (68 bits, 136 bits, or 272 bits) by programming the TLSZ field of the command register. The NDSEs perform all the necessary arbitration to decide which device will drive the SRAM bus. The SEARCH rate remains constant while the latency of searches increases as the table size increases.

3.1.1. Depth Cascading up to Eight Devices (One Block)

Figure 4 and Figure 5. Block Cascading for LHO/BHO Generation show how up to eight devices can be cascaded to form a table of 128K x 68 bits, 64K x 136 bits, or 32K x 272 bits within a block. It also shows the interconnection between the devices for depth cascading. The host ASIC must program the TLSZ field to 01. For each SEARCH, if a device determines a local hit within the device, it asserts the LHI[1] and LHI[0] signals. Each NDSE has it's own device ID[4:0] pins. If the device is full it asserts the FULO[1] and FULO[0] for each LEARN or WRITE into the data array. If all the upstream devices are full, the last device asserts a FULL signal in addition to FULO[1:0].

The NDSEs are informed of the SEARCH result within the device after four cycles from the issuance of a SEARCH command. They then drive their LHO signals. At the next cycle, all downstream devices know the outcome of the SEARCH from all the upstream devices.

If any of the upstream devices has a hit, all the subsequent devices defer the driving of the SRAM bus. If a miss occurs, the device 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.

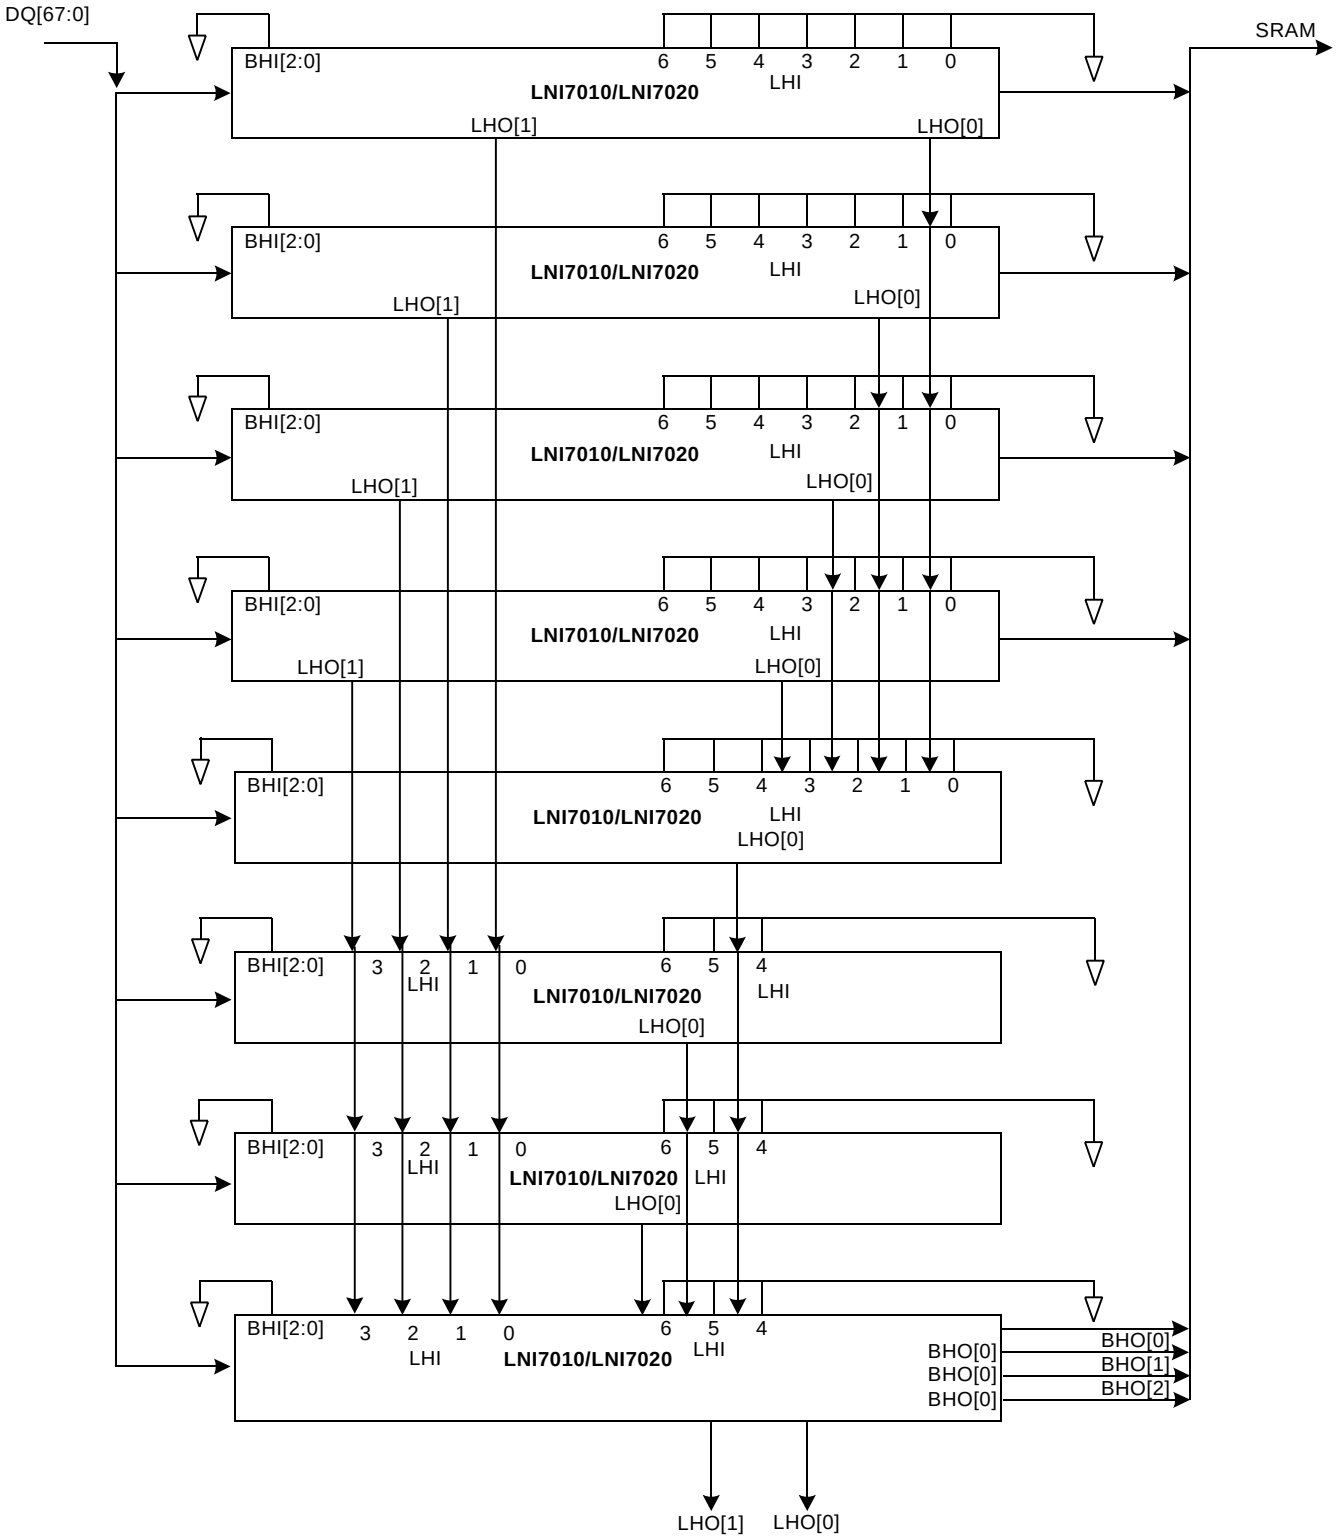

FIGURE 4. BLOCK CASCADING FOR LHO/BHO GENERATION

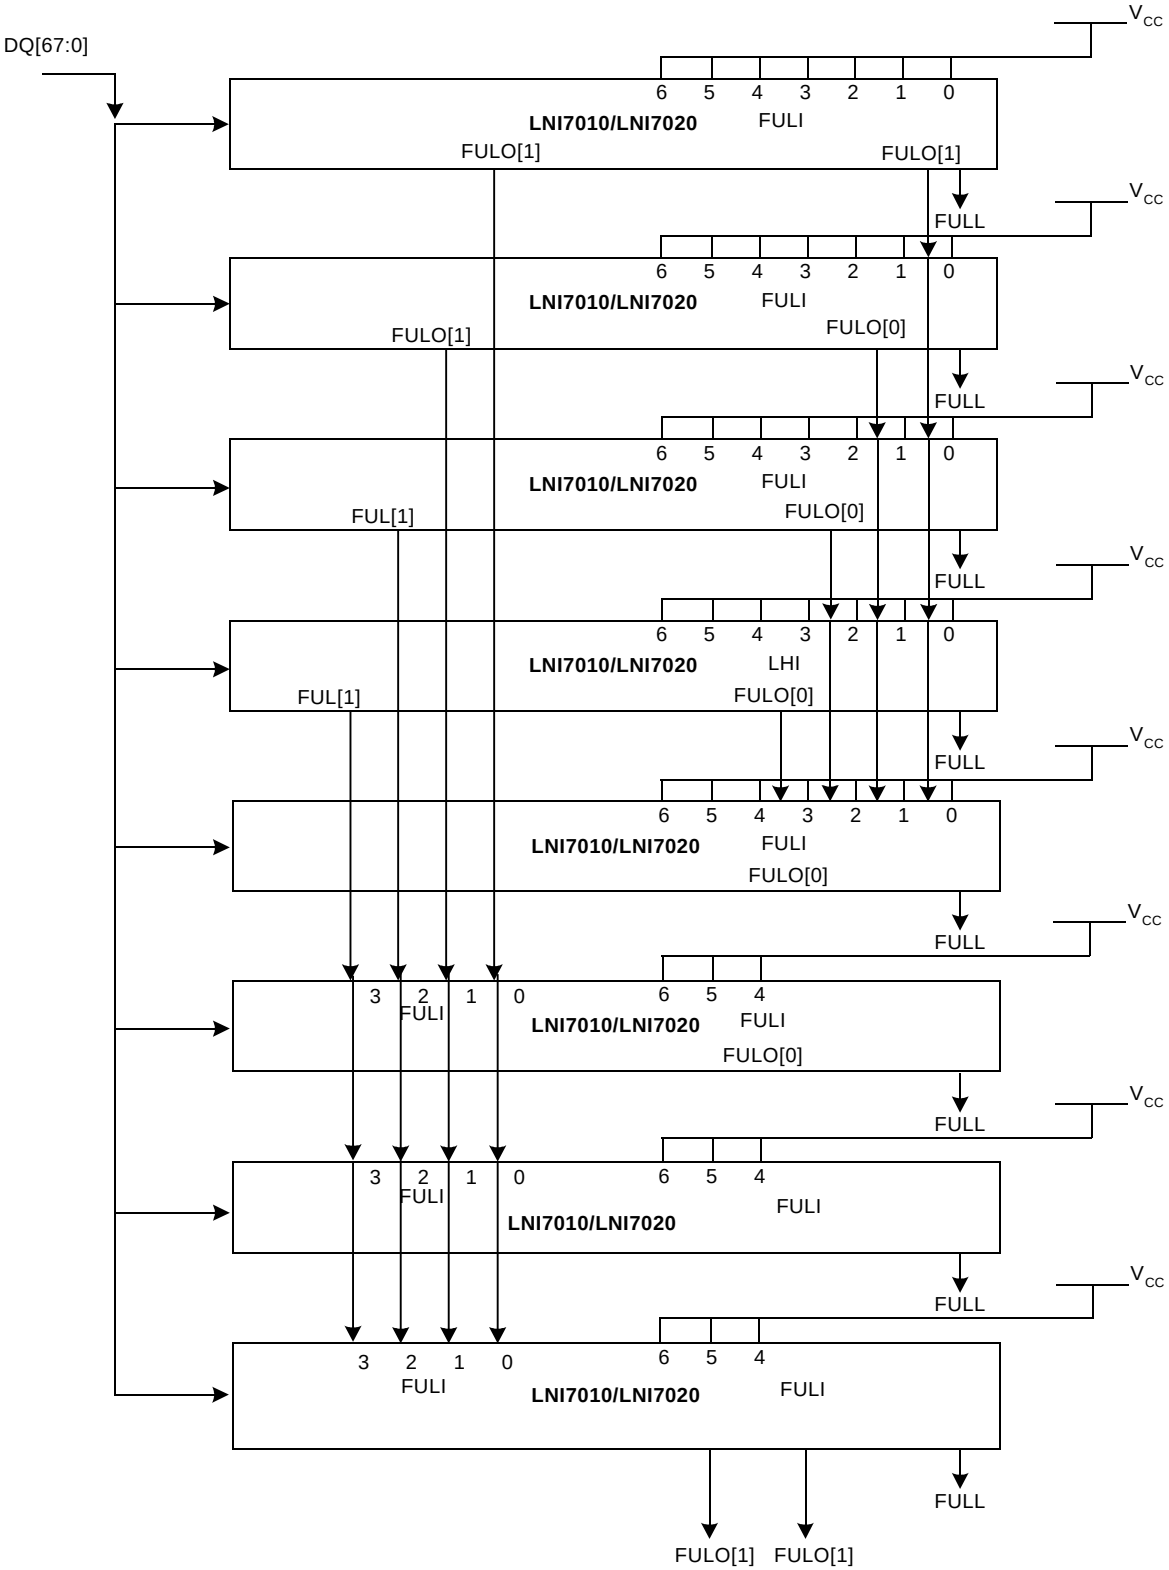

FIGURE 5. BLOCK CASCADING FOR LHO/BHO GENERATION

3.1.2. Depth Cascading up to 31 Devices (Four Blocks)

Figure 6 shows how to cascade up to four blocks. Each block except the last contains up to eight NDSEs to form a 496K x 68-bit, a 248K x 136-bit, or a 124K x 272-bit table. **Note**. See the interconnection between blocks required for depth cascading.

The host ASIC must program the TLSZ field to 10 for cascading between eight and 31 devices (in up to four blocks). A block asserts BHO[2], BHO[1], and BHO[0] for each SEARCH if there is a hit in one of the eight devices within that block.

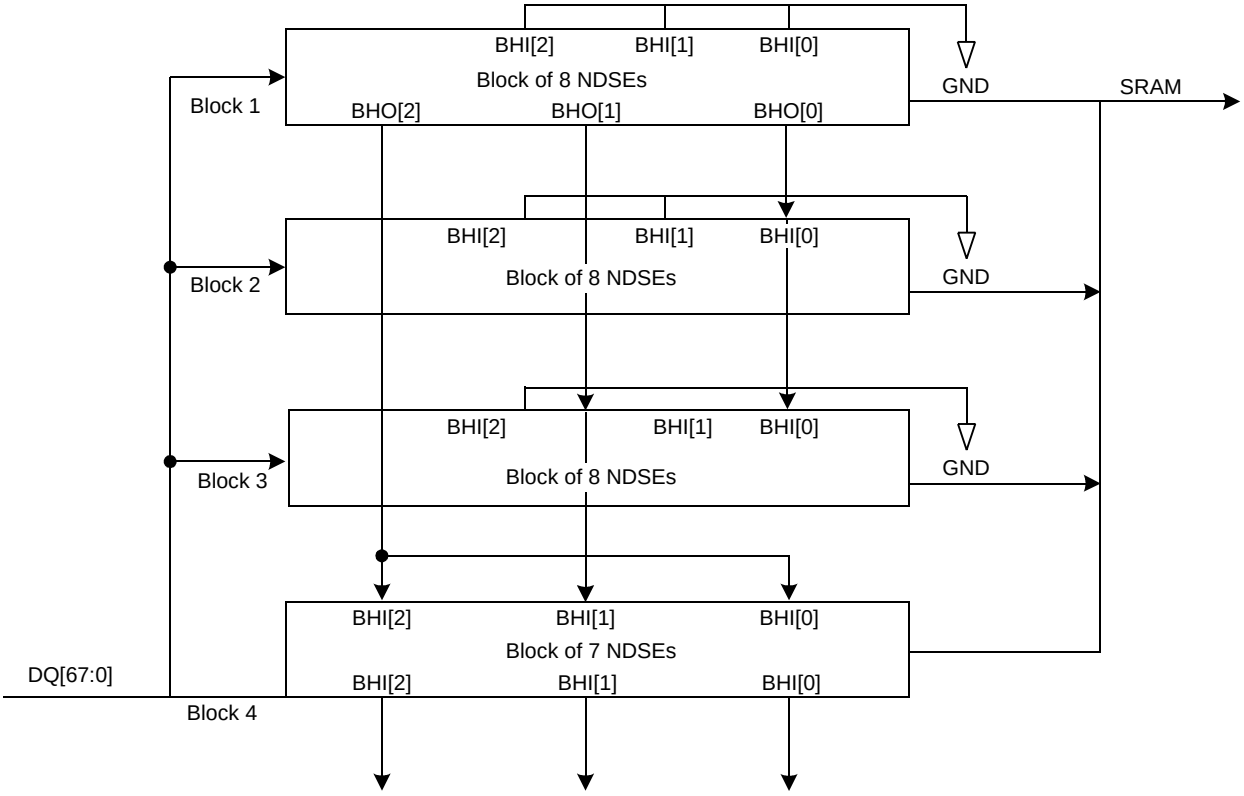

FIGURE 6. BLOCK CASCADING OF TWO OR MORE BLOCKS

3.2 Reset NDSEs

There are two ways of initializing a device. They are as follows.

1. Hardware reset using RST_L.
2. Software reset using SRST.

Hardware Reset. A hardware reset must be provided for 20 clock cycles once the power is on. Hardware or software resets execute the initialization sequence found in Step 3: Initialize the command register (see page 2).

3.3 Initialize the Command Register

The following items are key to the initialization of the command register as described in Table 2 (see page 4).

- The command register is initialized by a WRITE command.
- The host ASIC WRITEs in the NDSE whose ID[4:0] matches the DQ[25:21], or in all NDSEs when DQ[25:21] = 11111.
- Writing into the command register as shown below in Table 3 is a three-cycle operation.
- The last device in the chain must be programmed first in order to drive the SRAM bus interface signals as follows.

TABLE 3. TYPICAL COMMAND REGISTER INITIALIZATION FOR 8K X 136-BIT NDSE FOR THE LAST DEVICE IN SINGLE-DEVICE MODE

CYCLE	CMD[8:0]	CMDV	DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
1	xxxxxx001	1	Reserved	ID[4:0]	11	Reserved	111000
2	xxxxxxxxx	0	00… …00	00000	00	0001010101110	000010
3	xxxxxxxxx	0	xxxxxxxxxx	xxxxx	xx	xxxxxxxxxxxxxx	xxxxxxx

Writing into the command register is a three-cycle operation (as shown above in Figure 3). The address is supplied onto the DQ bus in the first cycle, the data is supplied in the second cycle, and NOP is supplied in the third cycle. During the first cycle, DQ[5:0] selects command registers, and DQ[20:19] selects register access. During the second cycle, the following bits must be asserted or disserted correctly to perform initialization.

1. SRST bits must be set to 0 (DQ[0] in the second cycle).
2. Set DEVE. The device must be enabled to drive the SRAM bus as well as cascading output flags. The DEVE line must be set to a logic 1 (DQ[1] in the second cycle) for all the devices in a block.
   - TLSZ, which is indicative of one device depth, is set to 00 (DQ[3:2] in the second cycle).
   - HLAT is set to 000 (DQ[6:4] in the second cycle).
3. Configure the last NDSE. The last NDSE in the table (LDEV) is indicative of the position of the last NDSE in the depth-cascaded table.
   - The device with LDEV set to 1 (DQ[7] in the second cycle) drives the SSF and SSV signals during a SEARCH with a miss or a nonSEARCH cycle (when no upstream devices drive these signals).
   - The last RAM of the table (LRAM) is indicative of the last NDSE on the SRAM bus.
   - If the LRAM is set to 1 (DQ[8] in the second cycle), the last device on the SRAM bus drives the SADR bus, OE_L, WE_L, ALE_L (when no upstream devices drive these signals).
4. Program the layer. When configuration (CFG) bits [16:9] are set to 0101010101 (DQ[16:9] in the second cycle), it is indicative of 136-bit operations in all four quadrants of a single device.
   - **Note**. The 8K x 136-bit NDSE is partitioned into four quadrants internally. This allows flexibility in performing multilayer lookups (layers 2 and 3) in the same device.

The third cycle is NOP. Programming of all the other devices in the block except the last NDSE is performed similarly to that of the last NDSE, above. **Note**. All the other devices except the last one in the chain must be programmed such that LDEV = 0, LRAM = 0, and DEVE = 1, as shown in Table 4. Therefore to initialize eight NDSEs would take 24 clock (CLK) cycles.

TABLE 4. TYPICAL COMMAND REGISTER INITIALIZATION FOR 8K X 136-BIT NDSE FOR ALL DEVICES EXCEPT THE LAST

CYCLE	CMD[8:0]	CMDV	DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
1	xxxxxx001	1	Reserved	ID[4:0]	11	Reserved	111000
2	xxxxxxxxx	0	00… …00	00000	00	0001010101000	000010
3	xxxxxxxxx	0	xxxxxxxxxx	xxxxx	xx	xxxxxxxxxxxxxxx	xxxxxxx

4 INITIALIZING AND PROGRAMMING THE GLOBAL MASK REGISTER

The global mask register (GMR) must be programmed for NDSE data writes, because one of the 16 mask registers is always used during data array WRITEs, mask array WRITEs, or SEARCH operations. Initially, one mask register must be programmed with all bits set to 1. Selecting that bit during WRITE operations allows the NDSE data to be written.

Table 5 shows the initializing of GMR[0], in order to enable a WRITE into the data array.

TABLE 5. TYPICAL GMR[0] INITIALIZATION FOR 8K X 136-BIT NDSEs TO ENABLE WRITES

CYCLE	CMD[8:0]	CMDV	DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
1	xxxxxx001	1	Reserved	ID[4:0]	11	Reserved	100000
2	xxxxxxxxx	0	11… …11	11111	11	1111111111111	111111
3	xxxxxxxxx	0	xxxxxxxxxx	xxxxx	xx	xxxxxxxxxxxxxxx	xxxxxxx

Writing into the GMR is a three-cycle operation, as indicated in Table 6, which shows bit assignments for the GMR. The address is supplied onto the DQ bus in the first data cycle, data is supplied in the second cycle, and the third cycle is NOP. In the first cycle, DQ[5:0] selects the GMRs. DQ[20:9] selects register access.

TABLE 6. GLOBAL MASK REGISTER FOR WRITE

DATA	GLOBAL MASK	WRITE
x	0	Do not modify
x	1	Modify

In the second cycle, the GMR is written with all bits as 1, via the DQ bus. This enables writing or comparison to all locations. The third cycle is NOP.

The LNI7010/LNI7020 device must be initialized with all 0s in the NDSE's data, and with all 1s for the NDSE's mask before the user loads any useful information.

The easiest way to initialize all NDSE data is by burst WRITE into all locations. Burst WRITE registers must be programmed each time before any burst WRITE operations.

4.1 Programming the Burst WRITE Register

Writing into a burst WRITE register is a three-cycle operation (as shown in Table 5). The address is supplied onto the DQ bus in the first cycle, the data is supplied in the second, and NOP is supplied in the third.

In the first cycle, DQ[5:0] selects burst WRITE registers and DQ[20:19] selects register access. In the second cycle, the start address is supplied by DQ[13:0] and BLEN[27:19] is supplied by DQ[27:19] (as shown in Table 7).

The WRITE burst address register allows a block WRITE of up to 512 locations at a time. Initializing 8K locations is done by setting burst registers to WRITE 512 locations at a time and executing burst WRITEs 16 times.

TABLE 7. TYPICAL BURST ADDRESS REGISTER INITIALIZATION FOR 8K X 136-BIT NDSE TO BURST WRITE

CYCLE	CMD[8:0]	CMDV	DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
1	xxxxxx001	1	Reserved	ID[4:0]	11	Reserved	111011
2	xxxxxxxxx	0	xx... .xx11	11111	11	xxxxx00000000	000000
3	xxxxxxxxx	0	xxxxxxxxxx	xxxxx	xx	xxxxxxxxxxxxxxx	xxxxxxx

For 8K x 136-bit mode, the user can supply 0 as the data and data mask for the NDSE, and use the WRITE burst register to store 0s in the NDSE data and mask register.

5 BURST WRITE INTO THE NDSE DATA ARRAY

NDSE data and masks are written by the burst WRITE command, as shown in Table 8.

TABLE 8. TYPICAL DATA ARRAY INITIALIZATION (ADDRESS 0-512) FOR 8K Z 136-BIT NDSE FOR BURST WRITES

CYCLE	CMD[8:0]	CMDV	DQ[67:26]	DQ[25:21]	DQ[20:19]	DQ[18:6]	DQ[5:0]
1	xxxxxx101	1	Reserved	ID[4:0]	11	Reserved	111011
2	xxxxxxxxx	0	xx... .xx11	11111	11	xxxxx00000000	000000
Repeat							
… … … ..							
Last	xxxxxxxxx	0	xx... .xx11	11111	11	xxxxx00000000	000000

Writing into the NDSE data array is a three-cycle operation, as shown in Figure 3. The address is supplied onto the DQ bus in the first cycle, data is supplied in the second cycle, and NOP is supplied in the third cycle.

In the first cycle, DQ[13:0] selects the starting address of the location and DQ[20:19] selects the NDSE data array. In the second cycle, data is written to selected addresses via the DQ bus with all bits as 1 or 0, since all bits are user-defined. It is recommended that bit[0] be reserved as a valid bit to be set to 1 when writing valid data. From the third cycle onward, data is written on an address that is incremented from the starting address up to the BLEN.

As an example bit[0], when a valid bit is stored, that data should be set to a 1 when the NDSE is updated with an entry. This is comparable to a valid bit provided by the user so tat it can be differentiated from an entry that has not been updated. As all bits are user-defined, it is recommended that bit[0] be reserved as a valid bit.

If binary operations are desired, the entire mask array is loaded with 1 after power-up. **Note**. For ternary operations, a match is done on the entire 136-bit word, and the occurrence of a match or no-match is provided.

The NDSE is now ready for SEARCH operations. SEARCH operations always use one of 16 GMRs, so they must be programmed before a SEARCH operation (see Table 9 for an example). Initially, programming one mask register with all bits to 1, and selecting it during a SEARCH, allows NDSE device data to be compared.

TABLE 9. GMR FOR SEARCH OPERATIONS

DATA	GLOBAL MASK	SEARCH
x	0	Disable comparison (forced match).
x	1	Enable comparison.

The SEARCH data should also be stored in one of the 16 pairs of comparand registers for the LEARN operation.

In a SEARCH operation for 8K x 136-bit mode, SEARCH results appear on the SRAM bus interface as well as being stored in one of eight SEARCH-successful registers.

Conclusions

As explained in this application note, it is very important to properly initialize all registers in the NDSE. After proper initialization, both data and mask arrays can be programmed with the appropriate data for SEARCH operations.

CONTACT

Lara Networks, Inc.
110 Nortech Parkway
San Jose, CA 95134
www.laranetworks.com

Tel: 1-408-942-2000
Fax: 1-408-942-2099

