JEDEC

Solid State Technology Association
2500 Wilson Boulevard
Arlington, Virginia 22201-3834
TEL: (703) 907-7560
FAX: (703) 907-7583
http://www.jedec.org

Electronic Industries Alliance

COMMITTEE LETTER BALLOT

COMMITTEE: JC-40.4

Committee Item Number: 142.35

Subject: LRDIMM DDR3 Memory Initialization Chapter Proposal

Background: At September '09 JEDEC meeting in San Jose, "motion by Intel and seconded by Montage to authorize the TG to issue one or more ballots on the items listed on pages 3 and 4 of the presentation that have item numbers, have been shown in the committee and for which the TG has reached a consensus".

The motion passed by acclamation.

This ballot proposal is to specify LRDIMM DDR3 Memory Initialization, and has been approved by DDR3 MB TG on 11/09/2009.

Keywords: DDR3, MB, Memory Buffer, LRDIMM, Initialization

(This page is intentionally left blank)

6 Initialization

6.1 Initialization Overview

The sequence of steps below is an overview of the power-up initialization of the Memory Buffer. Figure 1 shows a graphical representation of the same steps. Detailed requirements for each step are provided in the following sections. Note that the sequence below does not distinguish between host HW and host SW (i.e. BIOS).

1. Power-Up (Same requirements as 882)
   a. Ramp of voltage rails
   b. Generation of stable clock
   c. De-assertion of RESET#
2. Determination of DIMM configuration via SPD read out
3. CA Clock to CMD Training[1]
   a. Sets the optimal phase of the clock to CA and Ctrl signals.
4. Memory Buffer Initialization
   a. Host initializes buffer control words
5. DRAM Initialization
   a. Host initializes DRAM MRS registers
6. DRAM ZQ Calibration
   a. Host issues ZQCAL command to DRAM
7. Memory Buffer to DRAM Interface Training
   a. Done by the buffer, triggered by host RCW write
   b. Write Leveling
   c. Read Enable training
   d. Read/Write DQ/DQS training[2]

Host can wait for a per-physical rank time-out of tCAL=10ms or periodically poll the buffer via SMBus CSR read or wait for ERROUT#.

8. Host to Memory Buffer Interface Training[3]
   a. Done by host
9. Normal Operation

PROPOSED

1.This step is optional for the MB but may be done by some platforms. No special support by the MB is required.
2.This step is optional for the MB but may be required for stable DIMM operation.
3.This step is optional for the MB but may be done by some platforms.

Figure 1 — Initialization Overview

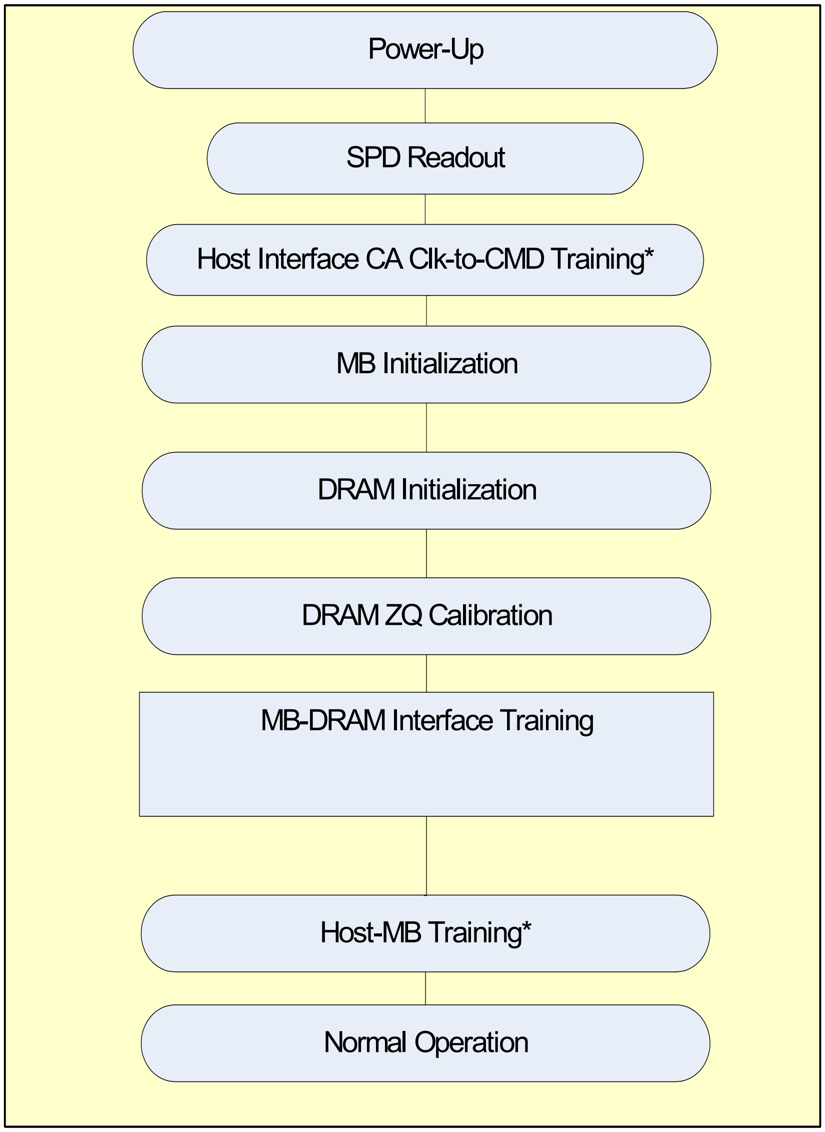

* This step is optional for the MB but may be done by some platforms

6.2 Power-on Initialization

The Memory Buffer can be powered-on at 1.5V or 1.35V. After the voltage ramp, stable power is provided for a minimum of 200 uS with RESET# asserted. When the reset input RESET# is LOW, all input receivers are disabled, and can be left floating. The MB output pin QRST# follows the MB input pin RESET#. Therefore the reference voltage (V_{REF}) doesn't need to be stable. In addition, when RESET# is LOW, all control registers are restored to their default states. The outputs QACKE[3:0] and QBCKE[3:0] must drive LOW during reset. All other outputs must float. As long as the RESET# input is pulled LOW the register is in low power state and input termination is not present. A certain period of time (t_{ACT}) before the RESET# input is pulled HIGH the reference voltage needs to be stable within specification, the clock input signal must be stable, the register inputs DCS[1:0]# must be pulled HIGH to prevent accidental access to the control registers and DCKE[2:0] as well as DCKE3/DODT[1] must be pulled LOW.

6.2.1 Clock Stabilizaion Time t_{STAB}

During PLL stabilization time t_{STAB} the memory buffer is not fully operational. In order to avoid invalid comands being sent to the DRAMs some rules apply to the inputs of the buffer:

- All DCS signals need to be kept high. No DRAM command or control word write[1] may take place.
- All DCKE signals don't change their state.
- DODT[0] or DODT[1:0] signals (depending on F0RC6 setting) are kept at a stable valid logic level.

These rules apply to any instance where stabilization time t_{STAB} is required:

- Exit from Reset
- Exit from clock stop power down
- Changing clocking related registers (F0RC2, F0RC10, F0RC11, F1RC8, F1RC11-F1RC15)[2]
- Changing input clock frequency during boot or run time

Since the buffer has not reached a stable state the termination on the host interface will be undefined before the end of the stabilization time.

After reset and after the PLL stabilization time (t_{STAB}) the device must meet the input setup- and hold specification, as well as accept and transfer input signals to the corresponding outputs. The RESET# input must always be held at a valid logic level once the input clock is present.

Figure 2 — Timing of clock and data during power-on initialization sequence

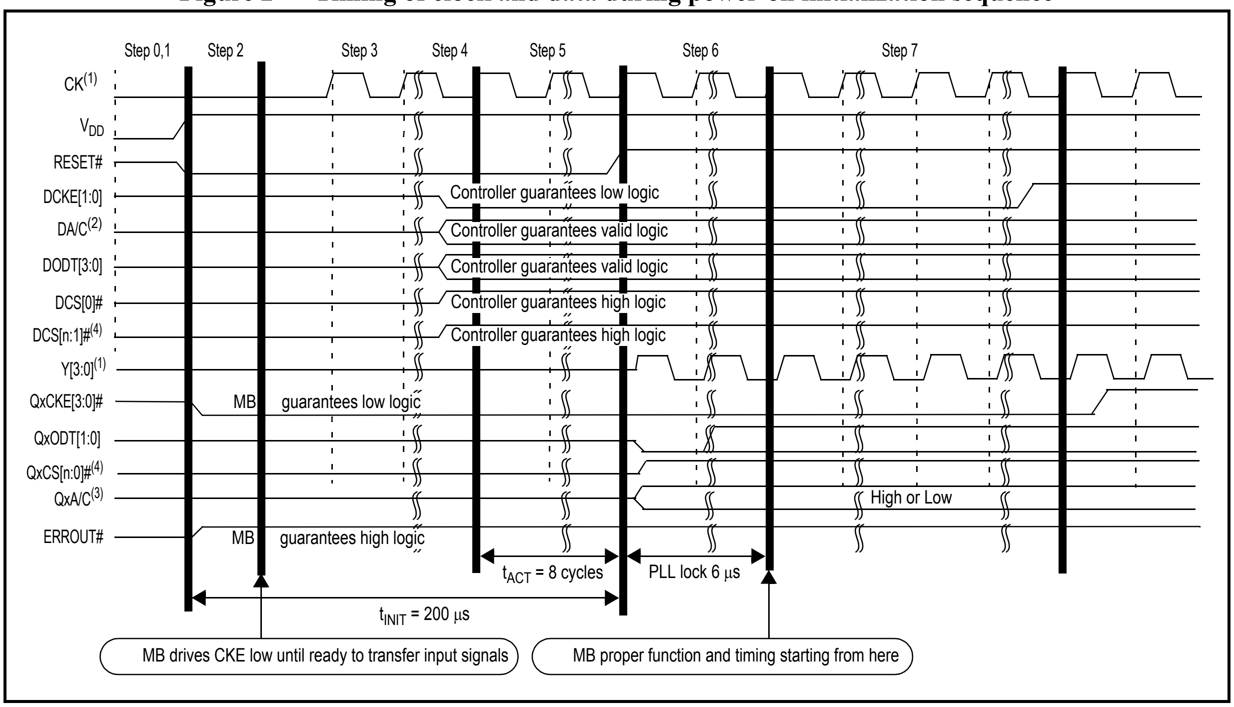

1. No control word writes are allowed during t_{STAB} after reset/frequency change/clock stop. Control word writes are allowed during t_{STAB} after clocking related register write.
2. These requirements mean that we have to wait two t_{STAB} times during every initialization, one after RESET# de-assertion and one after all the clocking related control words have been written.

(1) CK# and Y[3:0]# left out for better visibility
(2) DCKE[2:0], DCKE[3]/DODT[1], DODT[0] and DCS0[3:0]# are not included in this range
(3) QxCKE[3:0], QxODT[1:0], QxCS[7:0]# are not included in this range
(4) n = JEDEC standard DIMMs, n = 7 for non-JEDEC applications
(5) n = 3 for dual or quad rank DIMMs, n = 7 for octal rank DIMMs

From a device perspective, the initialization sequence must be as shown in Table 1

Table 1 — MB Device Initialization Sequence[a]

Step	Power	Inputs: Signals provided by the controller								Outputs: Signals provided by the device					
	VDD, AVDD, PVDD	RESET#	Vref	DCS# [n:0][b]	DODT [1:0]	DCKE [3:0]	DA/C	PAR_IN	CK CK#	QCS# [n:0][c]	QODT [1:0]	QCKE [3:0]	QxA/C	ERR OUT#	Y[3:0] Y#[3:0]
0	0V	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	Z	Z	Z	Z	Z	Z
1	0-->V_{DD}	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	L	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z
2[d]	V_{DD} 1.5V-->1.35V 1.35V-->1.5V	L	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	L	Z	Z	L[e]	Z	H[4]	Z
3	V_{DD}	L	X or Z	X or Z	X or Z	X or Z	X or Z	X or Z	running	Z	Z	L	Z	H	Z
4	V_{DD}	L	X or Z	H	X or Z	L	X or Z	X or Z	running	Z	Z	L	Z	H	Z
5	V_{DD}	L	stable voltage	H	X	L	X	X	running	Z	Z	L	Z	H	Z
6	V_{DD}	H	stable voltage	H	X	L	X	X	running	H	L[f]	L	X	H	running
7[g]	V_{DD}	H	stable voltage	H	X	X	X	X	running	After Step 6 (Step 7 and beyond), the device outputs are as defined in the device Function Tables (see Table 12, Table 14 and Table 16).					

a. X = Logic LOW or logic HIGH. Z = floating.
b. n = 3 for JEDEC standard DIMM, n = 7 for non-JEDEC application
c. n = 3 for dual or quad rank DIMMs, n = 7 for octal rank DIMMs.
d. The system may power up using either 1.5V or 1.35V. The BIOS reads the SPD and adjusts the voltage if needed from 1.35V to 1.5V or from 1.5V to 1.35V. After the voltage transition, stable power is provided for a minimum of 200 uS with RESET# asserted.
e. QxCKE[3:0] and ERROUT# will be driven to these logic states by the register after RESET# is driven LOW and VDD is 1.5V or 1.35V (nominal).
f. This indicates the state of QxODT[1:0] after RESET# switches from LOW-to-HIGH and before the rising CK edge (falling CK# edge).
g. Step 7 is a typical usage example and is not a MB requirement.

To ensure defined outputs from the register before a stable clock has been supplied, the memory buffer must enter the reset state during power-up. It may leave this state only after a LOW to HIGH transition on RESET# while a stable clock signal is present on CK and CK#. In the DDR3 LRDIMM application, RESET# is specified to be completely asynchronous with respect to CK and CK#. Therefore, no timing relationship can be guaranteed between the two.

As part of the initialization all control words are reset to their default state which is “0”, except when explicitly defined otherwise. After initialization, the host only needs to write to those control registers whose contents need to be changed.

6.3 Initialization with Stable Power (Soft Reset)

The timing diagram in Figure 3 depicts the initialization sequence with stable power and clock. This will apply to the situation when we have a soft reset in the system. RESET# will be asserted for minimum 100ns. This RESET# timing is based on DDR3 DRAM Reset Initialization with Stable Power requirement, and is a minimum requirement. Actual RESET# timing can vary base on specific system requirement, but it cannot be less than 100ns as required by JESD79-3 Specification.

Figure 3 — Timing of clock and data during initialization sequence with stable power

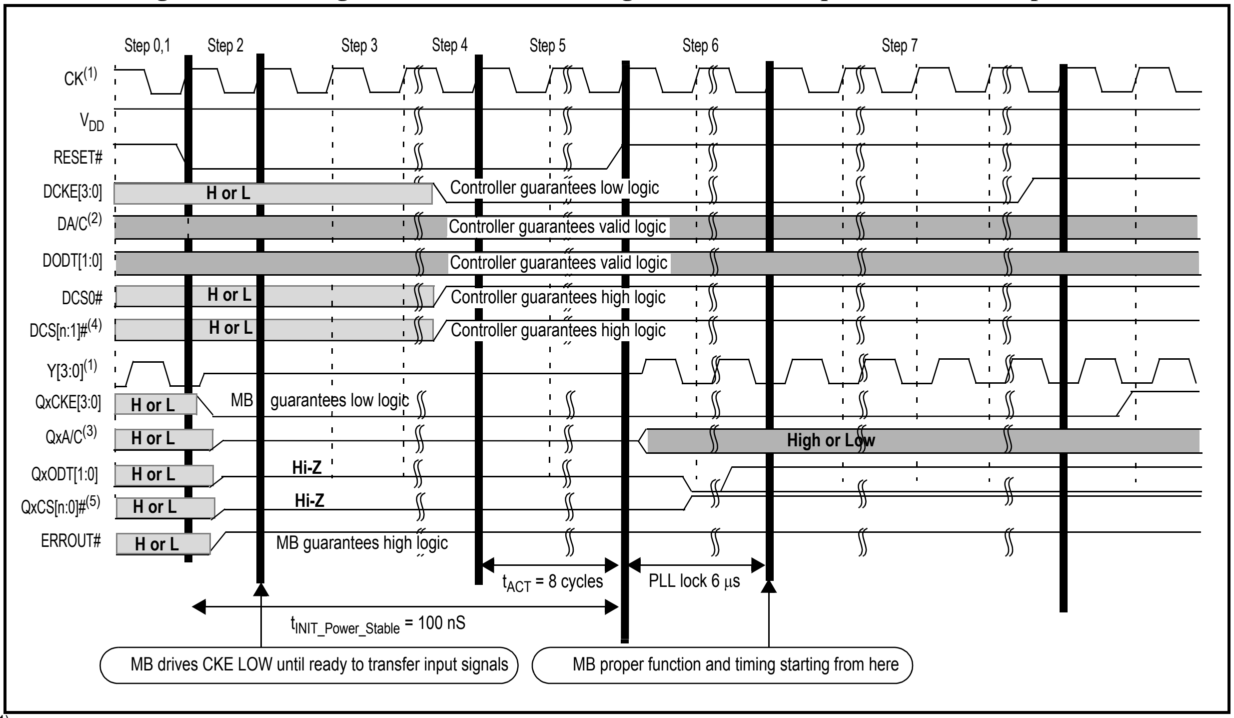

(1) CK# and Y[3:0]# left out for better visibility

(2) DCKE[2:0], DCKE[3]/DODT[1], DODT[0] and DCS0[3:0]# are not included in this range

(3) QxCKE[3:0], QxODT[1:0], QxCS[7:0]# are not included in this range

(4) n = JEDEC standard DIMMs, n = 7 for non-JEDEC applications

(5) n = 3 for dual or quad rank DIMMs, n = 7 for octal rank DIMMs.

Table 2 — MB Device Initialization Sequence[a] when Power and Clock are Stable

Step	Power	Inputs: Signals provided by the controller								Outputs: Signals provided by the device					
	VDD, AVDD, PVDD	RESET#	Vref	DCS# [n:0][b]	DODT [1:0]	DCKE [3:0]	DA/C	PAR_IN	CK CK#	QCS# [n:0][c]	QODT [1:0]	QCKE [3:0]	QxA/C	ERR OUT#	Y[3:0] Y#[3:0]
0	V_{DD}	H	stable voltage	X	X	X	X	X	running	X	X	X	X	X	running
1	V_{DD}	H	stable voltage	X	X	X	X	X	running	X	X	X	X	X	running
2	V_{DD}	L	stable voltage	X	X	X	X	X	running	Z	Z	L[d]	Z	H[4]	Z
3	V_{DD}	L	stable voltage	X	X	X	X	X	running	Z	Z	L	Z	H	Z
4	V_{DD}	L	stable voltage	H	X	L	X	X	running	Z	Z	L	Z	H	Z
5	V_{DD}	L	stable voltage	H	X	L	X	X	running	Z	Z	L	Z	H	Z
6	V_{DD}	H	stable voltage	H	X	L	X	X	running	H	L[e]	L	X	H	running
7	V_{DD}	H	stable voltage	H	X	X	X	X	running	After Step 6 (Step 7 and beyond), the device outputs are as defined in the device Function Tables (see Table 12, Table 14 and Table 16).					

a. X = Logic LOW or logic HIGH. Z = floating.
b. n = 3 for JEDEC standard DIMM, n = 7 for non-JEDEC application
c. n = 3 for dual or quad rank DIMMs, n = 7 for octal rank DIMMs.

d. QxCKE[3:0] and ERROUT# will be driven to these logic states by the register after RESET# is driven LOW and V_{DD} is nominal

e. This indicates the state of QxODT[1:0] after RESET# switches from LOW-to-HIGH and before the rising CK edge (falling CK# edge)

6.4 Host RCW to Configure MB

After the MB is ready to receive commands and addresses without error, the host needs to configure the MB with register control word writes and/or SMBUS writes (step 4). At reset all RCWs had been reset to their default state (which is '0' except when explicitly noted otherwise). Therefore the host only needs to write those control words whose contents need to be changed.

6.5 Host MRS to Configure DRAM

The host sends MRS commands to the DRAM mode registers to configure proper DRAM operation for all ranks behind the MB. The host can utilize either the MB broadcast mode to send MRS commands to all physical ranks associated with a logical rank or the MB physical rank mode to target MRS commands to specific physical ranks by programming the corresponding mode to the F0RC14 DRAM MRS Control bits.

6.6 Host MRS to DRAM ZQ Calibration

The host sends ZQCL (ZQ Calibration Long) commands to all ranks behind the MB. The MB broadcasts ZQ calibration commands (both ZQCL and ZQCS) to all physical ranks associated with a logical rank. The host may issue ZQ calibration commands sequentially to each logical rank (i.e. wait for tZQinit, tZQoper or tZQCS after each ZQ calibration command before issuing the next one) or it may overlap the ZQ calibration commands to all ranks.

The MB may or may not perform any calibration for its own I/O circuits on receipt of ZQCL or ZQCS calibration commands.

6.7 MB-DRAM Training

After all the DRAMs are fully operational, the host triggers the DRAM interface training by setting the control bit DA4 in F0RC12. Once enabled, the MB assumes autonomous control of the Command/Address, Control and Data/Strobe signals to the DRAMs, without any further assistance from the host.

The MB performs 'Write Leveling' to the DRAMs to ensure successful writes, and 'Read Enable Training' to ensure it can capture read data from the DRAMs correctly, as part of the DRAM interface training.

No DRAM commands or control word writes (either over the Command/Address and Control buses or via SMBus) can be issued to the MB until the MB-DRAM Interface training is complete and the DODTn inputs must be kept low. The MB responds to SMBus CSR read accesses whether or not DRAM interface training is active. To determine DRAM interface training completion the host may either wait for a time 'tCAL * number of physical ranks' (which is the maximum amount of time that the MB-DRAM Interface training is allowed to take) or it may periodically poll the content of the MB CSR 'Training Completion' (see 'Configuration Registers' chapter) over the SMBus.

In addition, training completion can be signaled by the assertion of ERROUT#. This is not the power-up default but can be enabled by setting the DBA1 bit in the F3RC14 Training Completion Control Word.

6.8 Host-MB Training

Now the host can train the MB host interface. The host sets the 'Connector DQ interface write leveling' bits in F0RC12 (=0bx001) in order to perform host interface write leveling. Optionally a MB may support RDIMM backward compatible host interface write leveling by intercepting ‘Write leveling enable’ (MR1 A7) and ‘Qoff’ (MR1 A12) commands to the DRAMs. The MB asynchronously feeds back CK/CK#, sampled with the rising edge of DQS/DQS#, through the DQ bus (in the same way as a DDR3 DRAM does). Since write leveling requests from the host terminate at the MB, multiple write leveling requests to the MB (i.e. one per rank) will give the same result and are optional. It is sufficient to perform host interface write leveling only once regardless of how many logical or physical ranks are supported by the MB.

MB DRAM interface training must be completed before any MB host interface read training to reflect the correct read round trip delay to and from an LRDIMM. This delay consists of

(1) Host to MB delay for command

(2) Command delay through the MB[1]

(3) MB command to DRAM delay

(4) DRAM latency

(5) DRAM data/strobe return time to MB

(6) Data/strobe delay through the MB

(7) MB to host delay for data/strobe

MB DRAM interface training must be completed before any MB host interface read or write training to ensure that the MB DRAM interface is configured optimally for DRAM reads and writes. The host performs normal writes to DRAM for write DQ/DQS margining, with failure occurring through incorrect strobing of write data through the MB. The host performs read DQ/DQS margining by using normal reads from DRAM or by using the DRAM multi purpose register (MPR) with failure occurring through invalid DQ/DQS alignment at the host I/O pins. No special support by the MB is required for host interface read training.

PROPOSED

1.Note that the data/strobe delay through the MB during writes can be different than the command delay through the MB.