# CISTER

## Research Center in
## Real-Time & Embedded
## Computing Systems

# Technical Report

## System and Circuit Level Power Modeling of Energy-Efficient 3D-Stacked Wide I/O DRAMs

**Karthik Chandrasekar**

**Christian Weis**

**Benny Åkesson**

**Norbert Wehn**

**and Kees Goossens**

# System and Circuit Level Power Modeling of Energy-Efficient 3D-Stacked Wide I/O DRAMs

Karthik Chandrasekar, Christian Weis, Benny Åkesson, Norbert Wehn, and Kees Goossens

CISTER Research Unit

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8340509

E-mail:

http://www.cister.isep.ipp.pt

## Abstract

JEDEC recently introduced its new standard for 3D-stacked Wide I/O DRAM memories, which defines their architecture, design, features and timing behavior. With improved performance/power trade-offs over previous generation DRAMs, Wide I/O DRAMs provide an extremely energy-efficient green memory solution required for next-generation embedded and high-performance computing systems. With both industry and academia pushing to evaluate and employ these highly anticipated memories, there is an urgent need for an accurate power model targeting Wide I/O DRAMs that enables their efficient integration and energy management in DRAM stacked SoC architectures.

In this paper, we present the first system-level power model of 3D-stacked Wide I/O DRAM memories that is almost as accurate as detailed circuit-level power models of 3D-DRAMs. To verify its accuracy, we experimentally compare its power and energy estimates for different memory workloads and operations against those of a circuit-level 3D-DRAM power model and show less than 2% difference between the two sets of estimates.

# System and Circuit Level Power Modeling of Energy-Efficient 3D-Stacked Wide I/O DRAMs

Karthik Chandrasekar[1], Christian Weis[2], Benny Akesson[3], Norbert Wehn[2], Kees Goossens[4]

[1]Computer Engineering, TU Delft, The Netherlands
[2]Microelectronic Systems Design, TU Kaiserslautern, Germany
[3]CISTER-ISEP Research Centre, Polytechnic Institute of Porto, Portugal
[4]Electronic Systems Group, TU Eindhoven, The Netherlands

*Abstract*—**JEDEC recently introduced its new standard for 3D-stacked Wide I/O DRAM memories, which defines their architecture, design, features and timing behavior. With improved performance/power trade-offs over previous generation DRAMs, Wide I/O DRAMs provide an extremely energy-efficient green memory solution required for next-generation embedded and high-performance computing systems. With both industry and academia pushing to evaluate and employ these highly anticipated memories, there is an urgent need for an accurate power model targeting Wide I/O DRAMs that enables their efficient integration and energy management in DRAM stacked SoC architectures.**

**In this paper, we present the first system-level power model of 3D-stacked Wide I/O DRAM memories that is almost as accurate as detailed circuit-level power models of 3D-DRAMs. To verify its accuracy, we experimentally compare its power and energy estimates for different memory workloads and operations against those of a circuit-level 3D-DRAM power model and show less than 2% difference between the two sets of estimates.**

## I. INTRODUCTION

In modern embedded SoC architectures [1], [2] and high-end servers and data centers [3], DRAM memories contribute significantly to the overall system power and energy consumption. With the industry pushing for both high-performance and green computing solutions, the demand for higher memory bandwidth has increased, albeit under tight power and energy budgets. Such contrasting needs have driven JEDEC and DRAM vendors to continuously improve DRAM architectures in terms of bandwidth and power efficiency, leading to the introduction of low voltage LVDDR3 and DDR4 memories for servers and desktops and LPDDR3 memories for mobile/embedded platforms. Although overall power efficiency has improved in these DRAM generations, power consumption during data transfer continues to be high, due to their power-hungry I/O circuits and high capacitance of the packaged/off-chip PCB interconnects between the DRAM memories and processors (between 8pF and 20pF for packaged (PoP) interconnects in LPDDR2 memories [4]).

To overcome this issue, JEDEC proposed a new standard for Wide I/O DRAM memories [5] that enables 3D stacking of the DRAM dies directly on top of processors to reduce the distance between the processor and memory to a few micrometers. The wider I/O in these memories increases the peak memory bandwidth, while the 3D stacking drastically brings down I/O power consumption, due to the low-capacitance (around 2pF [4]) Through Silicon Via interconnects (TSVs) used to accomplish the vertical stacking. The introduction of the Wide I/O DRAM standard now provides a platform for integrated processor and 3D-stacked memory design-space exploration to derive future high-performance and extremely energy-efficient embedded SoCs and server systems, helping meet both the green [6] and exascale computing goals [7]. However, the

key missing link required to facilitate the exploration of these opportunities is *an accurate system-level power model targeting Wide I/O 3D-DRAMs* that is: (a) easily integrable with system-level SoC design flows and (b) enables design-time 3D-DRAM energy estimation in future DRAM-stacked SoC architectures.

In this paper, we present the *first system-level power model of 3D-stacked Wide I/O DRAM memories* and verify its accuracy against a circuit-level 3D-DRAM power model considering JEDEC-specified Wide I/O DRAM configurations [5]. Towards this, we first describe the adaptations made to a baseline circuit-level DRAM architecture model to support 3D-stacked DRAM memories (also used in [8]–[10]) in Section III. We then propose our system-level power model for 3D-stacked Wide I/O DRAMs in Section IV. Finally, in Section V, we experimentally compare the power and energy estimates of the proposed system-level power model for different memory operations against those of the circuit-level power model and show their near equivalence.

The *four major contributions* of this work include:

(a) We propose the first system-level power-model of 3D-stacked Wide I/O DRAM memories.
(b) We describe the adaptations made to the circuit-level DRAM power model employed in [8]–[10] to address 3D-stacked Wide I/O DRAMs.
(c) We derive estimates for JEDEC current measures for different 3D-DRAM configurations using the circuit-level model, in place of the *as yet unavailable datasheets*.
(d) This system-level power model has been released online at [11] as an open-source 3D-DRAM power estimation tool.

## II. RELATED WORK

Many system-level DRAM memory power models have been proposed in the recent past, of which Micron's DRAM power model [12] is the most widely used. However, it was found to be inaccurate by Schmidt et al., in [13], who empirically measured power consumption of a DRAM device and showed that Micron's power model approximated power measures and over-estimated the actual savings of the Self-Refresh mode. Also, Micron's model did not employ details of memory command scheduling and hence could not report accurate power/energy consumption numbers. These issues were later fixed by [14], which used actual command scheduling information, accounted for power consumption during state transitions, and performed cycle-accurate analysis. However, both these system-level power models [12] and [14], target off-chip DDR2/DDR3 DRAMs, and have not yet been verified against independent detailed circuit-level DRAM power models such as, [8]–[10], [15] or [16].

When it comes to circuit-level power modeling of DRAMs, Thoziyoor et al., provided support for analysis of power and

timings of DRAMs in CACTI 5.1 [16], however, its architectural and circuit assumptions could not be employed for 3D-DRAMs. Hence, CACTI published another power model aimed at 3D-DRAMs in [17] (CACTI-3DD), but has not yet released its source code. Facchini et al., in [18] employed their internal circuit-level 3D-DRAM model for power estimation, but did not disclose much details about it. Rambus proposed a detailed circuit-level DRAM power model in [15] and calculated overall power consumption by modeling each DRAM memory component in detail. However, it targeted DDR2/DDR3 devices and did not address 3D-DRAMs, neither did it provide details on how it could be adapted to represent them. Weis et al., in [8]–[10] on the other hand, employed a similar circuit-level architecture and power model that was adapted to perform speculative design-space exploration of 3D-stacked Wide I/O DRAM memories. However, the details of the power model were not published due to the non-availability of a standard for 3D-DRAM architectures and their design and timings.

Although the circuit-level DRAM power models perform detailed and accurate power analysis, they employ complex device-level architecture details and technology specifications making it difficult to integrate them into existing system-level SoC design flows. Furthermore, DRAM vendors only reveal abstract JEDEC-specified worst-case current and voltage information in datasheets, and one needs to have a complete circuit-level understanding of DRAM architectures to adapt them to be used with the circuit-level models. Hence, for system-level SoC designers planning to employ 3D-stacked Wide I/O DRAM memories, there is a need for a system-level power model based on datasheet current and voltage specifications that is: (1) easily integrable into existing SoC design flows, (2) enables fast design-time DRAM energy estimation, and (3) reports power and energy estimates as accurate as the circuit-level models.

In this paper, we first present the adaptable circuit-level DRAM architecture and power model used in [8]–[10] and describe the adaptations made to it to support 3D-stacked Wide I/O DRAM memories [5]. We then present the system-level 3D-DRAM power model that addresses all of the issues discussed above and verify it against the circuit-level model.

## III. CIRCUIT-LEVEL POWER MODELING OF 3D-DRAMs

In this section, we first describe the baseline DRAM architecture model used in [8]–[10] in Section III-A. We then detail the adaptations made to it to target 3D-stacked JEDEC Wide I/O DRAM configurations [5], such as the introduction of TSVs and increase in I/O width, in Sections III-B and III-C. This circuit-level model is developed in SPICE, which provides details on the circuit behavior, such as device and wiring delays, current consumption of different circuit components etc., during different DRAM operations. These together with architectural parameters (such as operating frequency and capacity) and electrical data (such as different voltage sources) are employed as inputs to calculate timing and power consumption (using the lumped element model) of a target 3D-DRAM memory.

### A. Baseline Circuit-Level DRAM Architecture Model

DRAMs are organized as a set of memory banks that include memory elements arranged in arrays of rows and columns. The memory arrays are organized in a hierarchical structure of memory sub-arrays for efficient wiring and reduced power consumption. Each memory cell is modeled as a transistor-capacitor (1T1C) pair and the data is stored in the capacitor as a charge. The individual cells in each sub-array connect to

local wordlines and local bitlines. To read data from the memory, a *Precharge* is issued to prepare the local bitlines to a halfway voltage level and an *Activate* is issued to drive the local wordline high and transfer the charge between the memory cells and the connected local bitlines. This transfer of charge (data) is sensed by the primary sense amplifiers (row buffer), where they are latched. Then, *Read* commands can be issued to read out the specific columns of data (using column select lines) from the row buffer. The data is then switched from the row buffer via local datalines to master datalines and then to the secondary sense amplifiers, which interact with the I/Os. Once finished, the wordlines can be switched off, the cell capacitors disconnected, and the local bitlines can be precharged again.

We modeled a memory sub-array to consist of 256k cells connecting up to 512 cells per local bitline and per local wordline. We then connected 256 memory sub-arrays (organized as 16x16) to form 64Mb memory array macros, with master wordlines and column select lines (CSLs) extending over all the sub-arrays. 16 local wordlines across 16 horizontally organized memory sub-arrays connect to one master wordline per memory row and 8 local bitlines (8 memory columns) across 16 vertically organized memory sub-arrays connect to one CSL. This hierarchical organization of the DRAM model is shown in Figure 1.
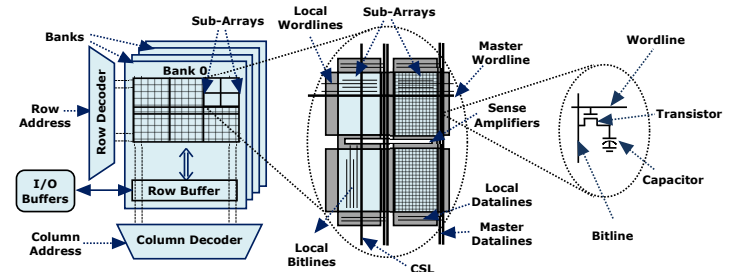


Fig. 1. Baseline DRAM Architecture - Circuit-Level Description and Modeling

The row and column decoders, the master wordline drivers, and the secondary sense amplifiers are placed per memory array. The data buffers, control signals, voltage regulators, charge pumps and other peripherals are shared between different banks.

### B. Extending to 3D-stacked Wide I/O DRAMs

When moving from the baseline DRAM architecture to 3D-stacked Wide I/O DRAM memories, the three biggest changes to be modeled include: (1) enabling three dimensional (3D) stacking of DRAM dies with the help of TSV interconnects, (2) supporting four independent memory channels, and (3) extending I/O interfaces to x128 bits per channel. 3D-stacked DRAMs offer increased memory bandwidth and improved energy efficiency, due to the increased I/O interface width and reduced I/O power consumption. The latter is a result of 3D-stacking DRAM dies with the help of low capacitance TSVs, compared to the traditional horizontal organization on one plane.

Figure 2 depicts the top view of the 3D-stacked multi-channel DRAM memory, with the four channels organized on four quadrants and the four banks of each channel on the top-most layer. In the figure, each quadrant contains the memory cell arrays, the bitline and wordline drivers, the control logic and the sense amplifiers. The power network, test pads, the charge pumps for the wordline high voltage, voltage generators and peripheral circuits are shared between the channels. The TSVs are all restricted to the marked area in each layer. This is compliant with the JEDEC-specified Wide I/O DRAM architecture [5].
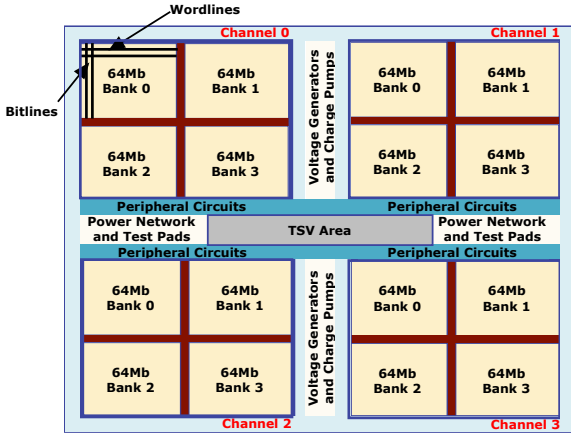
Fig. 2. Top view of 3D-Stacked 4-Channel Wide I/O DRAM with TSV area

In comparison to LPDDR2/3 memories, 3D-stacked DRAMs move away from package-on-package (PoP) interconnects, towards low-capacitance TSV interconnects between the memory and processor [4]. Additionally, the On-Die Termination (ODT) feature, which was re-introduced in LPDDR3 memories due to their higher frequencies, has been completely omitted in 3D-DRAMs due to further reduction in operating frequencies and lower I/O load due to 3D stacking, which further brings down I/O power consumption. As in LPDDR2/3 memories, the Delay-locked Loop circuit (DLLs) have been substituted by a programmable delay to align the data bus to the clock, to keep DRAM latencies and power consumption down.

3D-stacked DRAMs employ four external voltage supplies. The $V_{DD1}$ source at 1.8V serves as the supply voltage for the 2-stage charge pumps with improved efficiency to generate the wordline high voltage (around 2.8V). The $V_{DDCA}$ voltage source (1.2V) is used to drive the command and address buses. The $V_{DD2}$ source (1.2V) corresponds to the core voltage and is supplied to the control logic and parts of the peripheral circuitry in the DRAM device. The interface signaling voltage $V_{DDQ}$, which was absent in off-chip (DDR2/3) memories and only connected to the I/O buffers in mobile (LPDDR2/3) DRAMs, now reflects the entire I/O circuitry in 3D-DRAMs. This includes the I/O pins, I/O pads, I/O drivers, data TSVs and micro-bumps that connect the DRAM and controller directly and is also tied to 1.2V. Other circuit-level modeling details include:

(1) Design of the memory cell architecture of $6F^2$ area using 50nm technology [15].
(2) Use of high-k dielectric gate oxide for better sub-threshold behavior and reduced gate leakage [8], [19].
(3) Design of efficient voltage regulators, charge pumps, sense amplifiers, and buffers according to [15], [20].
(4) Use of appropriate TSV interconnect capacitances (between 2pF and 3pF [4]).
(5) Accurate dimensioning of transistor gate length and width in decoders, buffers, drivers and sense amplifiers [20].
(6) Modeling of appropriate local and global wiring for power distribution, data buses and control signals using TSVs.

Besides modifying or adding these features, electrical modeling of TSVs plays a significant role in accurate circuit-level modeling of 3D-DRAMs and is presented next.

## C. Electrical Modeling of a TSV

The circuit-level power model calculates accurate values for resistance and intrinsic capacitance for a Tungsten TSV, by em-

ploying an electrical model of a TSV similar to [21]. This model considers the TSV through a silicon substrate and oxidation layer as a co-axial wire, and estimates its intrinsic capacitance with respect to the oxide layer and the depletion region in the silicon substrate, besides calculating its resistance. It also considers at both ends of each TSV, I/O buffers used to drive the signal through the TSVs. Also included are the horizontal wires connecting the buffers to the TSV and their capacitances. Figure 3 shows the vertical cross-section and top view of a TSV.
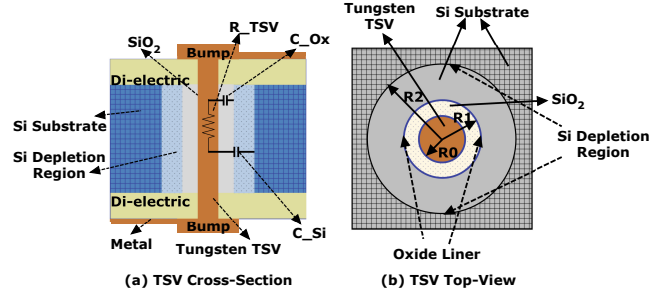


Fig. 3. Vertical and Horizontal (Top-view) Cross-Section of a TSV interconnect

Tungsten TSVs are employed in our model instead of Copper TSVs, because Tungsten has a relatively low thermal impact and high resistance to electro-migration. It also has relatively low resistivity and can be used to fill the very narrow contact structures. Samsung has also used Tungsten for its fabricated TSVs in [22]. In contrast to aluminum wires of $0.8\mu m$ pitch and $0.4\mu m$ width, which have a capacitance of 350fF/mm, the Tungsten TSVs that we employed with a diameter of $7.5\mu m$, pitch of $40\mu m$ and length $50\mu m$, have an intrinsic capacitance of 47.04fF. This is similar to the one reported by Samsung in [22]. However, our calculated resistance numbers of the TSV were different compared to Samsung's reported numbers. Our TSV's resistance value evaluated to $0.0896\Omega$, whereas Samsung reported $0.22\Omega$ in [22]. This difference is due to the additional resistance introduced by the manufacturing process itself. The numbers reported by Samsung are an indication of the partial filling of the tungsten inside the trenches, since via resistance is determined by the thickness of the tungsten layer inside the cavity (pitch) and not by the size of the cavity itself [23]. In our simulations, we decided to employ Samsung's TSV resistance numbers, since they correspond to post-manufacturing values. The I/O buffers driving through the TSVs have an output resistance of $100\Omega$ and a lumped capacitance of 100fF, including the wiring that connects the TSVs and buffers (similar to [8]).

Having presented the detailed circuit-level 3D-DRAM power model, the next section presents the proposed system-level power model for 3D-stacked Wide I/O DRAM memories before validating the accuracy of the same in Section V.

## IV. SYSTEM-LEVEL POWER MODELING OF 3D-DRAMs

As stated before, the complexity and level of detail employed by circuit-level DRAM power models, besides the non-availability of device-level technology specifications, makes it difficult to integrate them with existing system-level SoC design flows. Hence, the most viable method for estimating power consumption of 3D-DRAMs is to use system-level power models that use JEDEC-specified current and voltage values from memory datasheets, which are based on real hardware measurements. However, it should be kept in mind that *the accuracy of the DRAM power model using these datasheet measures, defines the accuracy of the DRAM power and energy estimates.*

To ensure accuracy, a system-level DRAM power model must satisfy all three requirements defined below:
(1) It should consider all memory activities in every clock cycle and keep track of the bank, channel and memory states that vary depending on the memory activity, instead of merely employing the minimum timing constraints given by memory vendors. It will also help obtain a temporal view on the power consumption of the DRAM memory.
(2) It should identify and account for memory operations that are enforced as a result of usage of certain memory features, such as all-bank precharging, refreshing and powering-up when transitioning into/out of the power-saving modes.
(3) Its power and energy estimates should be very similar to those of the circuit-level model for any memory operation, of any granularity (request size), for any variations in memory load or power-down/self-refresh durations.

The system-level power model proposed in this work, satisfies all the criteria mentioned above and is devised to target different 3D-Stacked Wide I/O DRAM memories. In the following subsections, we first discuss equations estimating average power consumption of the basic memory operations in Section IV-A (applicable to mobile DRAMs as well). We then present equations for accurate modeling of Power-Down and Self-Refresh modes and their related transitions, specific to 3D-DRAMs, in Section IV-B. Note that all the equations presented in this section correspond to a single channel of 3D-DRAM memories, since all channels in 3D-DRAMs can be independently analyzed [5].

In comparison to the power equations presented in [12], [14] for off-chip DRAMs, the proposed system-level 3D-DRAM power model has the following features:
(1) It explicitly considers the multiple voltage sources in 3D-DRAMs for different parts of the memory device.
(2) It reflects the changes in DRAM timing parameters due to removal of DLLs. This applies especially to the power-down and self-refresh power-saving modes.
(3) It calculates the I/O power consumption directly from datasheets using $V_{DDQ}$ domain current estimates, since the DRAM has moved 'on-chip'. This was previously not possible for DRAMs and was only recently addressed by [24] using circuit-level I/O description.
(4) It has better accuracy compared to [12], [14], since it performs deeper and exhaustive analysis of power consumption during state transitions, as shown in Section IV B.

Note that the $I_{DD}$ measures in the datasheets of 3D-DRAMs already reflect the impact of TSVs used in internal wiring and I/O, hence, there is no need to separately account for them in the system-level power model. Using TSVs reduces the I/O power per data bit transferred during a read operation to 0.7mW instead of 2.3mW in LPDDR2 memories [22] (PoP interconnects) and 4.6mW in DDR3 memories [12] (off-chip interconnects), leading to savings of 75% and 85%, respectively. By employing current measures and architecture and timing information from datasheets, the system-level power model can be easily integrated into existing system-level SoC design flows, without any complex changes or additions.

### A. Modeling Basic 3D-DRAM Operations

When it comes to basic memory operations, such as, Activate (ACT), Precharge (PRE), Read (RD), Write (WR) and Refresh (REF), 3D-DRAMs are not very different compared to off-chip and mobile DRAM generations, except for the use of multiple voltage sources and the computation of I/O power consumption.

Hence, we propose a generic power estimation model in Equation (1) for all basic DRAM operations and memory states that takes into account the different voltage sources, including $V_{DD1}$, $V_{DD2}$, $V_{DDCA}$ and $V_{DDQ}$. As can be noticed from the equation, it adds up the corresponding power estimates for all the voltage sources (calculated using the associated current measures) for the relevant memory operations. In the equation, $i$ is used to represent the $V_{DD1}$ and $V_{DD2}$ voltage domains. Note that the current measures corresponding to the $V_{DDCA}$ and $V_{DD2}$ sources have been added up and represented by $V_{DD2}$ (in Equation (1) and Table I), since they are both tied to 1.2V supply.

$$P(OP) = \sum_{n=1}^{t_{OP}} \Big( \sum_{i=1}^{2} \big( I_{DDi} \times V_{DDi} \big) + \big( I_{DDQ} \times V_{DDQ} \big) \Big) / t_{TOT} \quad (1)$$

Table I gives the values of currents (in mA) and timings (in ns) for the respective memory operations that should be substituted in this generic power equation. Accurate scaling of the power estimates for the basic memory operations, has been presented and described in [14]. The table also lists background currents consumed when the memory is in the active or precharged states. The I/O current numbers ($I_{DDQ}$) reported for the read/write operations corresponding to the $V_{DDQ}$ source, account for the I/O power consumption in the generic power model in Equation (1).

TABLE I
AVERAGE POWER CONSUMPTION OF BASIC MEMORY OPERATIONS

| Operation | $I_{DD1}$ | $I_{DD2}$ | $I_{DDQ}$ | $t_{OP}$ (ns) |
|---|---|---|---|---|
| ACT | $I_{DD0\_1} - I_{DD3N\_1}$ | $I_{DD0\_2} - I_{DD3N\_2}$ | - | $t_{RAS}$ |
| PRE | $I_{DD0\_1} - I_{DD2N\_1}$ | $I_{DD0\_2} - I_{DD2N\_2}$ | - | $t_{RP}$ |
| RD | $I_{DD4R\_1} - I_{DD3N\_1}$ | $I_{DD4R\_2} - I_{DD3N\_2}$ | $I_{DD4R\_Q}$ | $t_{RD}$ |
| WR | $I_{DD4W\_1} - I_{DD3N\_1}$ | $I_{DD4W\_2} - I_{DD3N\_2}$ | $I_{DD4W\_Q}$ | $t_{WR}$ |
| REF | $I_{DD5\_1} - I_{DD3N\_1}$ | $I_{DD5\_2} - I_{DD3N\_2}$ | - | $t_{RFC}$ |
| Active | $I_{DD3N\_1}$ | $I_{DD3N\_2}$ | - | $t_{act}$ |
| Precharged | $I_{DD2N\_1}$ | $I_{DD2N\_2}$ | - | $t_{pre}$ |

In Equation (1) and Table I, $t_{OP}$ corresponds to the period for which the corresponding operation must be active. For instance, $t_{OP}$ for a read and a write command is given by $t_{RD}$ and $t_{WR}$, respectively, which correspond to the period of data transfer during the respective read and write operations. However, $t_{OP}$ equates to $t_{RAS}$, $t_{RP}$ and $t_{RFC}$ for ACT, PRE and REF, commands respectively, which are JEDEC-specified minimum timing constraints to be satisfied for these operations to finish [20]. If these operations continue to be active beyond these minimum timing constraints, appropriate scaling of power numbers must be employed as shown in [14]. The $t_{act}$ and $t_{pre}$ timings correspond to the total time period spent in the active and precharged modes, respectively, when performing the basic DRAM operations. These are employed to estimate the background power consumption during these operations. $t_{TOT}$ refers to the total operation time window considered when estimating power for the particular operation. It is equal to $t_{OP}$ for all operations except activate and precharge commands, for which it is at least equal to the $t_{RC}$ timing constraint [20] (and may be longer [14]). Note that for accurate power and energy estimation, the actual command timings from the given memory trace must be employed instead of the minimum timing constraints, and the average power numbers must be appropriately scaled [14].

### B. Modeling Power-Saving Modes

When modeling average power consumption of the power-saving modes in 3D-DRAMs, the power model must take into account the memory operations and transitions that are enforced as a result of using these modes. For instance, when employing the self-refresh mode, all banks must be precharged and an explicit auto-refresh must be issued when entering the self-refresh

mode, and the energy consumption due to powering-up must also be accounted for. The power model must employ the appropriate currents associated with these resultant operations and transitions for their relevant time intervals, which are different in off-chip DRAMs due to the presence of DLLs. Below, we present the power equations (specific to 3D-DRAMs) for two of the power-saving modes: (1) Power-Down and (2) Self-Refresh. In all the equations, $i$ is used to represent the $V_{DD1}$ and $V_{DD2}$ domains.

*1) Power-Down Mode:* When an active/precharged power-down is issued, the DRAM must be in power-down mode for a time period of $t_{PD}$, which may vary from a minimum of $t_{CKE}$ to a maximum of $9 \times t_{REFI}$ (ns). When exiting the power-down mode, a time period of $t_{XP}$ is needed to restart regular operations to the DRAM (instead of $t_{XPDLL}$ in off-chip DRAMs [12]). When employing the power-down mode in the precharged state, the memory consumes $I_{DD2N}$ current when exiting from the power-down state and $I_{DD2P}$ current in the power-down state, as shown in Equation (2). If the power-down mode in the employed in the active state, $I_{DD3N}$ and $I_{DD3P}$ currents must be used instead.

$$P(PD) = \sum_{i=1}^{2} \left( \left( \sum_{n=1}^{t_{PD}} I_{DD2P\_i} + \sum_{n=1}^{t_{XP}} I_{DD2N\_i} \right) \times V_{DDi} \right) / (t_{PD} + t_{XP}) \quad (2)$$

Before entering the power-down mode, care should be taken that the last initiated memory operation is completed and the power consumption during this transition is accurately modeled. For instance, if a read (with or without auto-precharge) was issued before the power-down entry ($RD_{PDEN}$), $I_{DD3N}$ current is consumed during the Read 'Operation Time' ($t_{OP}$) and $I_{DD4R}$ during the cycles of data transfer (defined by Burst Length (BL)), as shown in Equation (3). For a write operation, $I_{DD4W}$ is consumed during the data transfer. The $t_{OP\_RD}$ for a Read is given by the sum of Read Latency ($t_{RL}$), data alignment time ($t_{DQSCK}$), Burst Length (BL) and 1 cycle for the auto-precharge (if any) to register ($t_{OP\_RD} = t_{RL} + BL + t_{DQSCK} + 1$). The $t_{OP\_WR}$ for a Write is given by the sum of Write Latency ($t_{WL}$), write to precharge time ($t_{WR}$), Burst Length (BL) and 1 cycle for the auto-precharge (if any) to register ($t_{OP\_WR} = t_{WL} + BL + t_{WR} + 1$).

For other basic memory operations preceding a power-down, such as ACT, PRE, and REF, one clock cycle must be spent in the active, precharged and active modes, respectively, before entering power-down. Also the ACT, PRE, and REF operation power must be considered using currents in Table I. The $I_{DDQ}$ measures do not apply here, since there is no data transfer.

$$P(RD_{PDEN}) = \left( \sum_{i=1}^{2} \left( \left( \sum_{n=1}^{t_{OP\_RD}} I_{DD3N\_i} + \sum_{n=1}^{BL} I_{DD4R\_i} \right) \times V_{DDi} \right) + \right.$$
$$\left. \left( \sum_{n=1}^{BL} I_{DD4R\_Q} \times V_{DDQ} \right) \right) / t_{OP\_RD} \quad (3)$$

*2) Self-Refresh Mode:* The self-refresh mode is used to retain data even when the clock is stopped (not just gated). When in self-refresh mode, the memory internally performs refreshes to maintain its contents without an external clock. When entering self-refresh, all banks must have been precharged and an explicit auto-refresh must be issued at the start of the self-refresh period.

The $I_{DD6}$ current is consumed for the time period spent in the self-refresh mode ($t_{SR}$), which excludes the time spent in finishing the explicit auto-refresh. The auto-refresh consumes $I_{DD5} - I_{DD3N}$ over one refresh period ($t_{RFC}$) from the start of the

self-refresh. $I_{DD2N}$ current is consumed when exiting the self-refresh state for the $t_{XSR}$ exit period (instead of $t_{XSDLL}$ in off-chip DRAMs [12]). If the the auto-refresh finishes before the self-refresh exit begins, during these auto-refresh cycles ($t_{SR\_REF}$), $I_{DD2P0}$ current is consumed in the background, instead of the $I_{DD6}$ self-refresh current. However, if the self-refresh exit begins before the end of the explicit auto-refresh, the remaining cycles of the auto-refresh operation ($t_{EX\_REF}$) carry forward to the self-refresh exit period. In this case, the $I_{DD3N}$ current is consumed in the background during these remaining cycles, instead of the $I_{DD2N}$ self-refresh exit current. This accurate modeling of transitions (in contrast to [12], [14]) is shown in Equation (4).

$$P(SR) = \sum_{i=1}^{2} \left( \left( \sum_{n=1}^{tSR} I_{DD6\_i} + \sum_{n=1}^{tXSR} I_{DD2N\_i} + \sum_{n=1}^{tRFC} (I_{DD5\_i} - I_{DD3N\_i}) + \right.\right.$$
$$\left.\left. \sum_{n=1}^{tSR\_REF} I_{DD2P0\_i} + \sum_{n=1}^{tEX\_REF} I_{DD3N\_i} \right) \times V_{DDi} \right) / (t_{SR} + t_{XSR} + t_{RFC}) \quad (4)$$

Having presented the circuit-level and system-level power models in Sections III and IV, respectively, the next section evaluates the latter against the former by comparing the energy estimates of the two, for different memory operations.

## V. RESULTS AND ANALYSIS

In this section, we present experiments to verify the accuracy of the system-level power model by comparing its power and energy estimates against those of the circuit-level power model. In these experiments, we employed four randomly selected MediaBench applications [25], mapped to the four channels of two JEDEC-specified 3D-DRAM configurations viz., 200 MHz and 266 MHz [5] including: (1) H.263 Encoder, (2) EPIC Encoder, (3) JPEG Encoder and (4) MPEG2 Decoder. These applications were independently executed on SimpleScalar simulator [26] with a 16KB L1 D-cache, 16KB L1 I-cache, 128KB shared L2 cache and 64-byte cache line configuration. We filtered out the L2 cache misses meant for the different channels of the 3D-DRAM memory and forwarded them through four trace players to a DRAM controller [27], which generated the memory commands for the different channels. Since 3D-DRAM datasheets are currently unavailable, we first derive the expected values of JEDEC current measures using the circuit-level power model and JEDEC test loops and conditions [5].

### A. Circuit-Level current estimates for Wide I/O DRAMs

In order to account for manufacturing process variations [28]–[30] and to avoid large yield losses, DRAM vendors provide worst-case current measures in datasheets. Hence, to be compliant with datasheet values, we also account for the expected variations and report worst-case current numbers by performing Monte-Carlo analysis on our circuit-level SPICE model. We further compared our worst-case estimates for 2Gb LPDDR2 memories against Micron's datasheets [31] and observed less than 1% difference between the two. In Table II, we present the generated worst-case current values (in place of the as yet unavailable datasheets) for the two JEDEC-specified 3D-DRAM configurations for a single channel viz., 200 and 266 MHz.

Although the current measures are generated using the circuit-level model, *the accuracy of the system-level power model defines the accuracy of the energy estimates*. Hence, if these were used by [12], [14], the accuracy of the power estimates would be worse. Note that the actual current measures in datasheets (when available) will be vendor-specific and can be different.

TABLE II
CURRENT MEASURES FOR WIDE I/O DRAM MEMORIES (ONE CHANNEL)

| Current | 3D-200 $V_{DD1}$ (mA) | 3D-200 $V_{DD2}$ (mA) | 3D-200 $V_{DDQ}$ (mA) | 3D-266 $V_{DD1}$ (mA) | 3D-266 $V_{DD2}$ (mA) | 3D-266 $V_{DDQ}$ (mA) |
|---|---|---|---|---|---|---|
| $I_{DD0}$ | 5.88 | 21.18 | - | 6.06 | 21.82 | - |
| $I_{DD2N}$ | 0.13 | 4.04 | - | 0.16 | 4.76 | - |
| $I_{DD2P}$ | 0.05 | 0.17 | - | 0.05 | 0.17 | - |
| $I_{DD3N}$ | 0.52 | 6.55 | - | 0.58 | 7.24 | - |
| $I_{DD3P}$ | 0.25 | 1.49 | - | 0.25 | 1.49 | - |
| $I_{DD4R}$ | 1.41 | 70.27 | 15.46 | 1.82 | 91.16 | 20.06 |
| $I_{DD4W}$ | 1.42 | 56.71 | 4.08 | 1.82 | 72.76 | 5.24 |
| $I_{DD5}$ | 6.26 | 28.17 | - | 6.39 | 28.74 | - |
| $I_{DD6}$ | 0.07 | 0.27 | - | 0.07 | 0.27 | - |

## B. Energy Comparison for Different Memory Operations

In our second experiment, we compare the energy estimates reported by the two models for different synthetically generated memory operations on a single DRAM channel for the two memory configurations. Table III presents the energy estimates for: (1) read and write operations of different granularities (request sizes) and (2) power-down and self-refresh operations with periods of different lengths. From the results, it is clear that despite variations, the system-level model deviates by less than 2% from the circuit-level model for different operations.

TABLE III
IMPACT OF VARIATION IN GRANULARITY ON ACCURACY (ONE CHANNEL)

| Memory Operation | Circuit 3D-200 Energy ($\mu J$) | System 3D-200 Energy ($\mu J$) | Diff % | Circuit 3D-266 Energy ($\mu J$) | System 3D-266 Energy ($\mu J$) | Diff % |
|---|---|---|---|---|---|---|
| RD (64 B) | 4.113 | 4.159 | 1.13 | 4.146 | 4.153 | 0.17 |
| RD (256 B) | 10.31 | 10.44 | 1.3 | 10.126 | 10.197 | 0.7 |
| WR (64 B) | 3.676 | 3.65 | -0.72 | 3.653 | 3.592 | -1.67 |
| WR (256 B) | 8.116 | 8.18 | 0.78 | 7.945 | 7.951 | 0.07 |
| PD (200 cc) | 0.341 | 0.346 | 1.46 | 0.285 | 0.289 | 1.4 |
| PD (1000 cc) | 1.493 | 1.522 | 1.94 | 1.149 | 1.171 | 1.91 |
| SR (200 cc) | 4.549 | 4.538 | -0.24 | 4.538 | 4.525 | -0.28 |
| SR (1000 cc) | 6.301 | 6.338 | 0.58 | 5.852 | 5.875 | 0.39 |

## C. Energy Comparison for Different Memory Loads

In our third experiment, we compare the energy estimates reported by the two power models for different workloads on all four channels of the memory. For this analysis, we employed the four memory traces obtained using the four MediaBench applications and employed either the power-down mode or the self-refresh mode for the idle periods [32] in all of them. We then increased the trace player frequency in steps, thereby varying the rate of traffic injection to the memory. Here as well, we observed less than 2% difference between the two estimates for both memories for all variations in traffic (depicted in Figure 4).
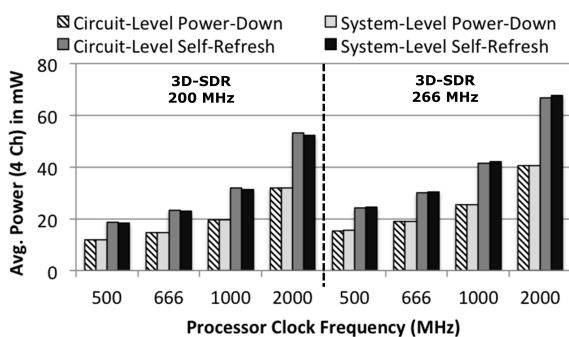


Fig. 4. Impact of Memory Load Variation on Accuracy (Four Channels)

In both these experiments the observed difference is due to the approximation of currents consumed during state transitions to JEDEC-specified currents, in the system-level power model.

## VI. CONCLUSION

In this work, we proposed the first system-level power model addressing 3D-Stacked Wide I/O DRAM memories and verified its accuracy using a circuit-level 3D-DRAM architecture and power model. We performed experiments for different JEDEC-specified 3D-DRAM configurations by varying memory operations, applications, memory load and power-down and self-refresh durations and showed less than 2% difference between the estimates of the two power models in all cases. This model has been released as an open-source 3D-DRAM power estimation tool and can be easily integrated with existing system-level SoC design flows for early design-time DRAM power and energy estimation in future 3D-DRAM stacked SoCs.

## REFERENCES

[1] O.Vargas, Infineon Technologies AG, *Achieve minimum power consumption in mobile memory subsystems*, Technical Report, 2006.
[2] C.H.Berkel, *Multi-core for mobile phones*, In Proc. DATE 2009.
[3] L.Minas et al., *Energy Efficiency for Information Technology: How to Reduce Power Consumption in Servers and Data Centers*, Intel Press, 2009.
[4] S.Dumas, *Mobile Market Trends drive Low Power DRAM*, JC42.6, 2011.
[5] JEDEC SST Association, *Wide I/O SDR DRAM Standard*, JESD229, 2012.
[6] DARPA-BAA-10-37, *Ubiquitous High Performance Computing*, 2010.
[7] S.Borkar, *The Exascale challenge*, In Proc. VLSI-DAT 2010.
[8] C.Weis et al., *Design Space Exploration of 3D-stacked DRAMs*, In Proc. DATE 2011.
[9] C.Weis et al., *An Energy Efficient DRAM Subsystem for 3D integrated SoCs*, In Proc. DATE 2012.
[10] M.D.Gomony et al., *DRAM Selection and Configuration for Real-Time Mobile Systems*, In Proc. DATE 2012.
[11] K.Chandrasekar et al., *DRAMPower: Open-source DRAM power & energy estimation tool*, www.es.ele.tue.nl/drampower, 2012.
[12] Micron Technology Inc., *TN-41-01: Calculating Memory System Power for DDR3*, Technical Report, 2007.
[13] D.Schmidt et al., *DRAM Power Management and Energy Consumption: A Critical Assessment*, In Proc. SBCCI 2009.
[14] K.Chandrasekar et al., *Improved Power Modeling of DDR SDRAMs*, In Proc. DSD 2011.
[15] T.Vogelsang, Rambus Inc., *Understanding the Energy Consumption of Dynamic Random Access Memories*, In Proc. MICRO 2010.
[16] S.Thoziyoor et al., *CACTI 5.1*, Technical Report, HP Labs, 2008.
[17] K.Chen et al., *CACTI-3DD: Architecture-level modeling for 3D die-stacked DRAM main memory*, In Proc. DATE 2012.
[18] M.Facchini et al., *System-level power/performance evaluation of 3D stacked DRAMs for mobile applications*, In Proc. DATE 2009.
[19] ITRS, Design, Report and Tables, Edition 2011.
[20] B.Jacob et al., *Memory Systems: Cache, DRAM, Disk*, Morgan Kaufmann Publishers, 2008.
[21] T.Bandhyopadhyay et al., *Electrical modeling of Through Silicon and Package Vias*, In Proc. 3DIC 2009.
[22] J.Kim et al., *A 1.2V 12.8GB/s 2Gb mobile Wide-I/O DRAM with 4×128 I/Os using TSV-based stacking*, IEEE JSSC, Vol.47, No.1, 2012.
[23] G.Pares et al., *Mid-process through silicon via technology using tungsten metallization: Process optimization and electrical results*, In ETPC 2009.
[24] N.Jouppi et al., *CACTI-IO: CACTI With Off-Chip Power-Area-Timing Models*, In Proc. ICCAD 2012.
[25] C.Lee et al., *MediaBench: a tool for evaluating and synthesizing multimedia and communications systems*, In Proc. MICRO 1997.
[26] T.Austin et al., *SimpleScalar: An infrastructure for computer system modeling*, IEEE Computer, Vol. 35, No. 2, 2002.
[27] B.Akesson et al., *Architectures and Modeling of Predictable Memory Controllers for Improved System Integration*, In Proc. DATE 2011.
[28] H.David et al,. *RAPL: memory power estimation and capping*, In Proc. ISLPED 2010.
[29] M.Gottscho et al., *Power Variability in Contemporary DRAMs*, IEEE Embedded Systems Letters, Vol. 4, No. 2, 2012.
[30] Intel, *Memory 3 Sigma Power Analysis*, White Paper, 2008.
[31] Micron Technology Inc., *LPDDR2 SDRAM 2Gb Data Sheet*, 2011.
[32] K.Chandrasekar et al., *Run-Time Power-Down Strategies for Real-Time SDRAM Memory Controllers*, In Proc. DAC 2012.