

Mobile System Considerations for SDRAM Interface Trends

Andrew B. Kahng
CSE and ECE Depts., UC San Diego
La Jolla, CA 92093, USA
abk@ucsd.edu

Vaishnav Srinivas
ECE Dept., UC San Diego
La Jolla, CA 92093, USA
vaishnav@ucsd.edu

ABSTRACT

A variety of interconnect technologies and standards (DIMMs, MCP, POP, stacked-die and 3D-stack) enable a controller IC to communicate with an external SDRAM, or with multiple SDRAMs over a shared interconnect. Low-power requirements have driven mobile controllers to mobile-SDRAM (LPDDR) memory solutions. However, LPDDR configurations do not scale to match the throughput and capacity requirements of mobile processors, or of emerging tablet products that bring new and divergent tradeoffs among memory subsystem metrics. As a result, identifying the memory configuration best suited to a given mobile application becomes quite challenging.

This paper highlights considerations in choosing a particular memory configuration for a mobile processor based on capacity, throughput, latency, power, cost and thermal concerns. We distinguish various choices according to interconnect implementation and performance, including power and timing in the IO and interconnect. To do this, we apply a three-part framework: (1) driving questions in the form of a *decision tree*, (2) a *calculator* that projects power and timing for mobile IO implementations, and (3) propagated top-down requirements and bottom-up capabilities that distinguish interconnect implementations.

Our framework can support abstraction of timing and power for various interconnect configurations, to feed higher-level tools such as CACTI [19]. We anticipate that it can also be used to project mobile system requirements and memory interconnect capabilities into the future, so as to identify any gaps or bottlenecks in memory product roadmaps.

Keywords

SDRAM, DDR, 3D, Mobile, Tablet, memory configuration, IO, LPDDR, serial memory, Wide IO.

1. INTRODUCTION

SDRAM (Synchronous Dynamic Random Access Memory) is the mainstay main-memory solution for processors today. Many offerings exist in the market [1, 10, 11, 13, 14], spanning multiple packaging and interconnect options. Most SDRAMs available are DDR (Double Data Rate), such as DDR3 and LPDDR2 (Low-Power DDR2). While LPDDR2 and DDR3 memories exist in the market today, others are being discussed in JEDEC [35] or elsewhere as standards for the future, and will probably be available in the market soon based on the traction they receive.¹

¹JEDEC [35] memory standards and nascent standards being discussed in JEDEC <name, IO voltage, pin-width, data rate (Gbps), interconnect, signalling>: (i) <LPDDR2, 1.2V, x16/x32, 1.066, POP/MCP, single-ended LVCMOS output / SSTL input>, (ii) <DDR3, 1.5V, x4/x8/x16, 2.133, DIMM, SSTL>, (iii) <LPDDR3,

POP (Package-on-Package) and MCP (Multi-Chip Package) offer good point-to-point interconnection options, while DIMMs (Dual Inline Memory Module) are required for discrete parts. 3D stacking offers a new way to stack the die through a TSV (Through Silicon Via [4]) and promises to enable high bandwidth memory access through wide interfaces (up to x512 proposed [13]).

Currently, mobile controllers have been driven to LPDDR2 solutions by their need for lower power and point-to-point unterminated interconnect. As requirements on throughput and capacity scale upwards for mobile platforms, alternate solutions are needed since data rates for unterminated signaling do not scale easily beyond 1 Gbps [17]. Section 2 outlines some of the trends in the DRAM industry, including the competing standards available for future mobile memory solutions [13, 14, 26] and ITRS projections on DRAM densities and data rates [6, 7, 8]. It is shown that DRAM densities are expected to double every three years, while DRAM data rates are expected to double every four years. Section 3 highlights the trends for mobile system requirements on the memory [15], including capacity and throughput. It is shown that capacity requirements are expected to scale 3-4x every three years, while throughputs are expected to double every three years [16].

Most of these competing DRAM options differ in their interface implementations. A memory-subsystem framework that includes the SDRAM interface metrics would be useful to compare them and help identify the best memory solution for a particular system. As the first step towards such a framework, we provide a *memory interface calculator*, described in Section 4, for the existing memory interface options available for a mobile system. We also describe how the calculator could be extended for future memory interface options based on the basic categories it supports. For each interface option, the calculator provides a framework to predict interface power and latency for a given throughput and capacity. Integration of this interface calculator into a memory-subsystem calculator such as CACTI [19] is a direction for future work.

The key metrics for the main memory are: availability, capacity, cost, throughput, power, latency and thermal considerations. These we call the *primary bounds* on the main memory design space. When choosing between SDRAM options, a priority order of the primary bounds is helpful to identify the best memory solution easily. One such priority order [25] is used to illustrate our framework. The most important bound is availability. Often availability and cost of the memory are tied to volume demand. During early adoption, enough traction is needed for demand to drive down cost and speed up availability. This normally happens through the process

1.2V, x16/x32, 1.6, POP/MCP, being discussed>, (iv) <DDR4, 1.2V, x4/x8/x16, 3.2>, (v) <Mobile-XDR, 1.2V, x16/x32, 4.3, MCP, differential>, (vi) <Wide IO, 1.2V, x512, 0.2-0.3, 3D-stack, LVCMOS>, (vii) <Serial Memory, 1.2V, x16/x32, 4-8, MCP, SPMT or MIPI M-PHY>

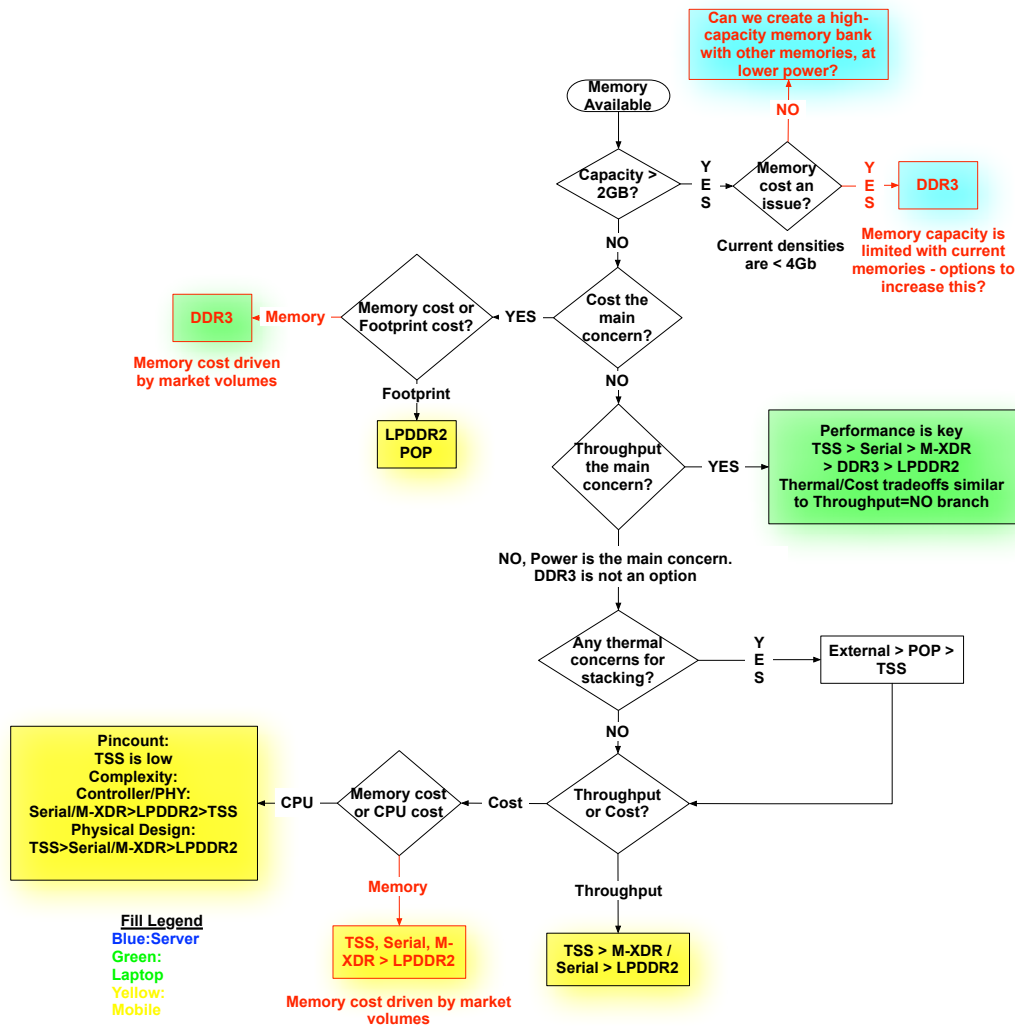


Figure 1: The decision tree with forks based on SDRAM capabilities and regions based on system requirements.

of standardization, where early adopters drive the co-development of the standard. The traction of a standard or solution depends on how it performs with respect to the primary bounds and intangible market forces that are beyond the scope of this paper. What is worthwhile to note, though, is that there is a large transition period from one standard to the next. Often backward compatibility with respect to the interface and interconnect plays an important role in the success of a smooth transition to the new standard, as OEMs (Original Equipment Manufacturer) need time to slowly phase out the old memory type for the new one. This requires the same mobile controller IC to support both the old and new memory types over the transition period.

Shown in Figure 1 is an example *decision tree* with priority ordering of primary bounds to help identify potential memory solutions. The tree is divided into three segments based on the requirements on these bounds – server (blue), laptop (green) and mobile (yellow). The priority of the primary bounds shown in Figure 1 is an example order of decisions by which a memory solution may be identified. We use this order in the paper to illustrate how the calculator aids in determining the choice of memory interface. Other priority orders may be suited to different system requirements, and

could serve just as well to identify a memory solution.

Each endpoint in the tree identifies one or more particular memory options based on the requirements, and taking into account market bottlenecks (availability, cost) or fundamental technical bottlenecks. The forks identified in red are those currently identified as market-driven. Currently, a high-capacity requirement (>2GB) would make DDR3 memories the only viable option. However, if throughput is not a key primary bound for a particular system, then other options may become viable based on MCP solutions (this is the focus of recent investigations elsewhere [22]). Such alternatives, if capable of increasing demand for a memory type across market segments, could further help reduce cost and improve availability. It is useful to note that while the memory options of the decision tree change over time due to the capabilities of the SDRAM options (both market-driven and fundamental) described in Section 2, the location of the market segments on the decision tree change based on the system requirements described in Section 3. The decision tree shown reflects the capabilities and requirements in today's market.

Once availability and cost have been established, the remaining driving questions (bounds) are capacity, throughput, power, latency

and thermal considerations. The rest of the paper focuses on the first four of these primary bounds and ignores availability and cost issues.

Section 4 describes how the interface considerations for each of the primary bounds impact the circuit and physical design of the mobile controller's IO logic and interconnect. This is shown through a set of equations that contain the requirements and capabilities that address each primary bound. The equations listed are part of the memory interface calculator.

Section 5 summarizes the paper, while providing the motivation for future work, especially:

(1) The need to integrate power and latency for IO and interconnect configurations of various memory types into a higher-level abstraction such as CACTI [19]. This will help provide memory subsystem metrics of latency and energy for these different interconnect options [20, 21].

(2) Predict bottlenecks or gaps in the memory offerings that will provide clues to novel IO and interconnect solutions [2, 22] or help motivate a new memory standard for the future.

2. SDRAM CAPABILITY TRENDS

ITRS projections for DRAM density and DRAM data rate [6, 7, 8] are shown in Figures 2 and 3 respectively. The projections are shown from those made in 1999 till those made in 2009, each year being revised based on new ITRS data available. The data rate projections shown are for the DDR2/DDR3/DDR4 family.




Figure 2: DRAM density projections [8].

Based on the latest projections (2009), it can be observed that DRAM densities are expected to double every three years, while DRAM data rates are expected to double every four years. Also interesting to note, are the blue arrows in Figures 2 and 3, which indicate that over the years, projections for DRAM densities have been revised downwards while DRAM data rates have been revised upwards.

Some of the standards being discussed in JEDEC today include DDR4, DDR3-ULV, LPDDR3, Wide IO, Serial-MIPI and Serial-SPMT [13, 14, 27]. Mobile-XDR and XDR2 are signaling standards driven by Rambus [26]. Recent nascent standards also highlight the trend to either increase pin-width (using TSVs) or use differential signaling for higher throughputs on existing pin-limited interconnects. This is because data rates using single-ended signaling are hard to scale with good power efficiency and signal-integrity [2, 3, 23].

DRAM I/O Rate




Figure 3: DRAM data rate projections [8].

3. MOBILE REQUIREMENTS TRENDS

Mobile system requirements for peak throughput and capacity are shown in Figure 4 [16] and Table 1 respectively. Figure 4 shows throughput requirements of products based on volume production timelines. Capacity requirements are expected to scale 3-4x every three years, while throughputs required are expected to double every three years. Other more aggressive throughput requirements for mobile application processors [32] suggest close to a doubling every year.




Figure 4: Peak throughput projections for mobile platforms (in GB/s) [16].

Table 1: Capacity projections for mobile and PC markets (Source: IDC Market Survey).

Market	2010	2011	2012	2013	2014
Desktop	3.0	4.2	5.6	7.4	10.2
Laptop	2.0	3.3	4.6	6.3	8.0
Mobile	0.3	0.5	0.8	1.0	1.3

4. INTERFACE CONSIDERATIONS

Sections 4.1 through 4.4 address the SDRAM interface considerations of a mobile system for capacity, throughput, power and latency as they appear in the decision tree. Each of these four sections highlights the inferences it passes onto subsequent sections. The considerations are highlighted Eqs. (1) - (31), which show the requirements and capabilities surrounding the circuit design stage of the mobile processor's design hierarchy. The memory interface calculator is based on these equations.

4.1 Capacity

Capacity is the total memory available, typically in GB. It is equal to the density of each memory die multiplied by the number of dies. The dies can be organized into multiple channels, each channel with multiple ranks comprised of x4, x8, x16 or x32 memories to support a required bus width of the mobile controller.

The capacity determines the number of memory dies needed for given density availability. This in turn determines the interconnect lengths and number of multi-drop loads on the bus. The interconnect structure is a key input for the signal-integrity analysis that determines data rate capability. For mobile systems with lower capacities (<2GB), this is typically one of three options: a stacked die, POP or a short trace to an MCP off-chip. Trace lengths within an inch or two are common for the SDRAM interface, keeping it unterminated. However, higher capacity requirements, especially for tablet products, could see the need to support DDR3 DIMMs or multi-drop structures.

Equations governing capacity are as follows.

$$Capacity = N_{channels} \cdot N_{ranks} \cdot \left(\frac{N_{bus_width}}{N_{memory_width}} \right) \cdot Density \quad (1)$$

$$N_{memories} = \left(\frac{Capacity}{Density} \right) \quad (2)$$

$$N_{memories/channel} = \left(\frac{N_{memories}}{N_{channels}} \right) \quad (3)$$

$$N_{ranks} = N_{memories/channel} \cdot \left(\frac{N_{memory_width}}{N_{bus_width}} \right) \quad (4)$$

$$C_{interconnect} = C_{RD_L} + C_{TSV} + C_{PKG} + C_{PCB} \quad (5)$$

$$C_{ext_DQ} = C_{interconnect} + C_{IO_DRAM} \cdot N_{ranks} \quad (6)$$

$$C_{ext_CA} = C_{interconnect} + C_{I_DRAM} \cdot N_{memories/channel} \quad (7)$$

$$C_{Total} = C_{ext} + C_{self} \quad (8)$$

$$Delay_{interconnect} = f(N_{memories}) \quad (9)$$

4.2 Throughput

Throughput in GB/s is the total number of bytes of data transferred over the memory interface. The data rate is a multiple of the clock frequency; the multiple is 1 for SDR, 2 for DDR, and higher than 2 for multi-data rate.

Equations governing throughput are as follows.

$$Throughput = \left(\frac{2 \cdot N_{bus_width}}{T_{ck}} \right) \quad (10)$$

DQ-DQS WRITE:

$$\left(\frac{T_{ck}}{4} \right) - T_{error} - T_{jitter_hold} - T_{skew_hold} > T_{DH} \quad (11)$$

$$\left(\frac{T_{ck}}{4} \right) - T_{error} - T_{jitter_setup} - T_{skew_setup} > T_{DS} \quad (12)$$

CA-CLK (DDR for LPDDR2/3):

$$\left(\frac{T_{ck}}{4} \right) - T_{error} - T_{jitter_hold} - T_{skew_hold} > T_{IH} \quad (13)$$

$$\left(\frac{T_{ck}}{4} \right) - T_{error} - T_{jitter_setup} - T_{skew_setup} > T_{IS} \quad (14)$$

For DDR3 the CA interface is SDR, and the above timing is relaxed to a half-cycle as opposed to a quarter-cycle. DQS-CLK:

$$T_{jitter_hold} + T_{skew_hold} < \left(\frac{T_{ck}}{2} \right) - T_{DSH} \quad (15)$$

$$T_{jitter_setup} + T_{skew_setup} < \left(\frac{T_{ck}}{2} \right) - T_{DSS} \quad (16)$$

DQ-DQS READ:

$$T_{QSH/QSL} - T_{error} - T_{jitter_hold} - T_{skew_hold} - T_{QHS} > T_{SOC_hold} \quad (17)$$

$$\left(\frac{T_{ck}}{4} \right) - T_{error} - T_{jitter_setup} - T_{skew_setup} - T_{DQSQ} > T_{SOC_setup} \quad (18)$$

Voltage Noise:

$$V_{noise}(crosstalk, SSN, reflection - ringbacks) < V_{NM} \quad (19)$$

$$V_{overshoot} < V_{MAX}(NBTI, HCI) \quad (20)$$

The data rate determined based on the throughput requirement forms a key input for the interface design. Wider interfaces, such as Wide IO, have low data rates for equivalent throughputs, making timing closure easier. On the other hand, serialization increases the data rate and requires higher power to meet timing (due to power in the termination and the analog blocks that help recover and retime the signals).

The key consideration for throughput is the link timing budget based on the data rate. This is based on the implementation of the PHY and IO and the signal-integrity analysis of the link. Each of these factors reduces the data-valid window for the READ, WRITE and CA operations and limits the maximum achievable frequency.

For mobile systems, throughput considerations also include meeting timing with minimum power, and the ability to reduce power at low throughputs. This constrains the termination scheme (unterminated being preferred at lower throughputs) and the choice of retiming scheme. As mobile memories are designed with similar considerations, the power they consume is lower at the expense of timing. For example, LPDDR2 memories do not have a DLL, so although their power is lower than DDR3 memories, they have worse timing parameters [10, 11, 12] and thus the link timing budget is harder to close (LPDDR2 DRAM accounts for more than 33% of the Unit Interval (UI)). This puts considerable burden on the mobile controller to close timing on the SDRAM interface. Traditional link design for the SDRAM interface divides the UI into three equal (33%) buckets for the controller, channel and DRAM. Shorter interconnects (owing to smaller capacity requirements highlighted in Section 4.1) help keep the channel's contribution to the link timing budget to less than 33% at lower data rates. But as data rates get higher, the crosstalk and ISI increase, and lack of terminations for LPDDR2 interfaces impose greater challenges on the channel jitter. Mobile-XDR, Wide IO and serial memories offer different

ways to solve these challenges while keeping the power efficiency low. While Wide IO proposes to use lower data rates to ease timing without increasing pin-count, Mobile-XDR and serial memories propose to use differential signaling to improve signal quality and reduce signal swing [2, 24, 30]. Figure 5 shows throughput projections for the various standards.




Figure 5: Throughput (GB/s) projections based on the memory interface calculator.

The support of multiple memory standards, especially during transitions from one standard to the next, poses challenges for the mobile processor’s IO as the IO will most likely need to support multiple voltages with a single transistor device.

4.3 Power

The Micron DDR3 power calculator [5] provides a framework to calculate memory interface and core power for DDR3 configurations. Our memory interface calculator models power with a similar framework, but includes calculations for all of the mobile memory types of interest today – LPDDR2/3, Wide IO, Mobile-XDR and Serial Memory.

The interface power is mainly made up of the following sources:

1. Switching (dynamic) power: typically drawn when capacitors are charged, including self-loading.

2. Termination power: The static power drawn through the termination network. The termination power depends on the termination current, which in turn depends on whether the line is at a logic 0 or a logic 1. For DDR3, CA topology uses fly-by termination [10, 12]. If mobile systems use DDR3, sub-optimal termination values are used to save power. The Micron DDR3 power calculator [5] shows how the power would change for dual-rank configurations with different termination values.

3. Bias/static power: analog currents used to bias the IO circuits - transmitter as well as receiver, adaptive circuits, DLLs and other timing circuits. The memory interface calculator provides a look-up for the power consumed for various implementations in literature [2, 3, 31] that relate to reported jitter and skew performance of recovery and retiming blocks along with adaptive schemes for calibrating the IO for the channel and the recovery logic for the eye opening. The ability to reduce such power based on data rate is very important for maintaining good power-efficiency at lower bandwidths [2, 3, 24]. SPMT provides a novel way to run a parallel interface at lower data rates and a serial interface at higher data rates [27], which provides good power efficiency at both ends of the throughput range.

Shown below are the considerations for power. Table 3 lists the power efficiencies of the various interfaces included in the memory interface calculator. Mobile processors use a variety of schemes to reduce power in low-throughput modes that include shutting off

terminations, using low-power retiming schemes at lower throughputs, and power-gating schemes that turn off blocks not critical at lower throughputs. To achieve these power savings, transitions between various power states are enabled.

$$P_{Total} = \sum_{mode_i} (D_{ci} \cdot (P_{Total_Interface_i} + P_{DRAM_Core_i})) \quad (21)$$

Interface Power:

$$P_{Total_Active_Peak} = P_{dynamic} + P_{term} + P_{static/bias} \quad (22)$$

$$P_{Total_Active_Idle} = P_{term} + P_{static/bias} \quad (23)$$

$$P_{Sleep} = P_{leakage} \quad (24)$$

Dynamic Power:

$$P_{dynamic} = N_{pins} D_c \alpha C_{Total} V_{sw} V_{dd} f \quad (25)$$

Single-ended termination:

$$P_{term_oh} = (V_{dd} - V_{TT})(V_{oh} - V_{TT})/R_{TT} \quad (26)$$

$$P_{term_ol} = V_{TT}(V_{TT} - V_{ol})/R_{TT} \quad (27)$$

DDR3 DQ:

$$P_{DQ_Term} = 0.25 \cdot V_{dd}^2 \cdot \left(\frac{1}{R_{TT}} + \frac{1}{R_{on} + R_{stub} + R_{TT}} \right) \quad (28)$$

DDR3 CA:

$$P_{CA_Term} = 0.25 \cdot V_{dd}^2 \cdot \left(\frac{1}{R_{TT}} + \frac{1}{50 + R_{TT}} \right) \quad (29)$$

Differential Termination:

$$P_{diff_term} = V_{dd} V_{sw} / R_{TT} \quad (30)$$

4.4 Latency

Latency from the memory controller to main memory is a portion of the cache miss penalty during a load instruction from the CPU. Most of the mobile memories have similar access times (T_{RCD} , T_{RL} , T_{DQSK}) as they depend on the DRAM core that they share. The key latency differentiator is the implementation of the PHY and IO on the mobile processor. Latency is not currently included in the decision tree as the various memory options have latencies of the same order (50-70ns from controller to DRAM and back), but this range could well be a differentiator in performance. Future work will examine the architecture impact of different latencies in this range.

The memory interface calculator calculates latency during a load cache miss (a memory READ operation) from the controller CA path back to the DQ READ path. This includes the pipeline stages in the controller and PHY for CA and READ, the combinational output and input delay in the IO, the passive delay in the interconnect, T_{RCD} and the read latency (T_{RL} and T_{DQSK}).

Apart from the latency during the load miss, the latency of waking up from sleep to full activity is an important consideration for mobile systems, since frequent mode changes are needed to enable power savings, and wakeup times impact performance of the CPU. DLL/PLL lock times and wakeup times for any of the adaptive schemes involved will impact the latency of the IO during wakeup from sleep modes. DLLs on DDR3 DRAMs take over 500 clock cycles to lock [10, 12], whereas LPDDR2 memories do not have DLLs and hence can exit self-refresh mode without a latency penalty. Periodic calibration for adaptive schemes, such as impedance and delay calibration, require periodic stalls to update

the calibrated values. The stall frequency for such updates, and the stall time needed for propagating the updated settings, are other latency considerations. Eq. (31) shows the calculation for latency.

$$Latency = N_{pipes} \cdot T_{ck} + T_{del_int} + T_{del_comb} + T_{RCD} + T_{RL} + T_{DQSCK} \quad (31)$$

4.5 Examples

To further illustrate the calculator and the decision tree, we provide two examples that go through the steps described in Sections 4.1 through 4.4.

Example 1:

A particular system requires 4GB of DRAM and expects a peak throughput of 10 GB/s.

The 4GB capacity requirement, using 4Gb DRAM density, translates to 8 DRAM dies ($N_{memories} = 8$). This rules out an MCP or stacked solution, and an external interconnect through a DIMM. Hence, DDR3 is the memory of choice within the decision tree. Eight x8 DRAMs connected to a x64 bus width provides the needed single-rank memory configuration. The loading on the DQ pin is a single DRAM along with the interconnect. The loading on the CA pin is eight DRAMs along with the interconnect.

A throughput of 10 GB/s on a x64 bus translates to a data rate of 1.25 Gbps, or a clock frequency of 625MHz for the DDR3 DRAM. Use of a 667 MHz DDR3 speed bin is required. The interface timing parameters of interest for this DRAM are listed in Table 2. Plugging into the timing equations of Section 4.2 shows that the jitter and skews required over a DIMM interconnect require the use of terminations, and a fly-by channel for the CA bus. Based on this, outputs of the power and latency calculations from the calculator are summarized in Table 2.

Example 2:

Another system requires 1GB of DRAM and expects a peak throughput of 4 GB/s.

The 1GB capacity requirement, using 4Gb DRAM density, translates to 2 DRAM dies ($N_{memories} = 2$). POP, MCP or stacked solutions are possible. Given the lower throughput requirement, LPDDR2 POP is the preferred choice based on the decision tree (to reduce power and footprint cost). A dual-rank x32 POP is the suitable memory configuration. The loading on the DQ pin is two DRAMs along with a short POP interconnect. The loading on the CA pin is also two DRAMs along with the short POP interconnect.

A throughput of 4 GB/s on a x32 bus translates to a data rate of 1 Gbps, or a clock frequency of 500MHz for the LPDDR2 DRAM. Use of a 533 MHz DDR3 speed bin is required. The interface timing parameters of interest for this DRAM are also listed in Table 2. Plugging into the timing equations of Section 4.2 shows that the jitter and skews required over a POP interconnect require some careful retiming as the LPDDR2 interface does not support terminations. Based on this, the power and latency calculations are as summarized in Table 2.

5. SUMMARY AND FUTURE WORK

We have described a framework to compare mobile memory interfaces that includes a decision tree, a memory interface calculator and requirements/capabilities for the primary bounds. Considerations for a mobile system that concern the memory interface can be elaborated with the help of requirements and capabilities concerning each primary bound, and applied according to a priority order embodied in a decision tree. Our memory interface calculator cur-

Table 2: Inputs and outputs of the calculator for two design examples.

Parameter	Example 1	Example 2
Capacity	4GB	1GB
Throughput	10 GB/s	4 GB/s
$N_{memories}$	8	2
Configuration	x64 DDR3	x32 LPDDR2
N_{ranks}	1	2
Clockrate	625 MHz	500 MHz
T_{DH}	240 ps	210 ps
T_{DS}	205 ps	210 ps
T_{IH}	315 ps	220 ps
T_{IS}	240 ps	220 ps
T_{DSH}	300 ps	380 ps
T_{DSS}	300 ps	380 ps
$T_{QSH/QSL}$	570 ps	713 ps
T_{QHS}	180 ps	230 ps
T_{DQSQ}	125 ps	200 ps
$P_{Total_Active_Peak}$	~1.5 W	~170 mW
$P_{Total_Active_Idle}$	~500 mW	~5-10 mW
Latency	~50 ns	~50 ns

rently implements Eqs. (1) - (31) in Excel to project power and latency for various interface options that meet a given capacity and throughput.

A summary of the outputs of the memory interface calculator for the competing mobile memory standards is shown in Table 3 below. This includes the power and latency of the mobile controller's IO logic. Table 3 clearly highlights that LPDDR2 needs to scale beyond 533 MHz if it is to maintain its place in the mobile market. LPDDR3, currently projected to support 800MHz would be a suitable choice until 12 GB/s. Given the large power consumption of DDR3, a suitable alternative is required for throughputs higher than 12 GB/s. Wide IO, Mobile-XDR and Serial Memory (SPMT, MIPI) provide such alternatives. Wide IO is clearly the most power-efficient, but uses TSS. Mobile-XDR and serial memories provide traditional interconnect options with differential signaling. Efficiency of serial schemes at lower throughputs is a concern. Larger (>2GB) capacity requirements need to be addressed for all the mobile memory types, but Wide IO provides an option of a multi-die stack using TSS.

Going forward, interface timing and power can be added into a memory-subsystem calculator such as CACTI [20, 34]. By doing this, we can quantify impact of memory interface choices on memory subsystem performance and power. This is an important next step in projecting capabilities against requirements of the memory subsystem.

Projecting both capabilities and requirements can also help identify gaps in memory interface offerings or predict bottlenecks in capabilities. Shown below in Figure 6 is the commodity memory space for the mobile market, overlaid with the mobile requirements from Figure 4 but leaving out the LPDDR3 curve. The figure shows a clear gap in the commodity offerings between LPDDR2 (until 8 GB/s) and the future mobile offerings of serial and Wide IO (above 12 GB/s). LPDDR3 is being discussed in JEDEC, but was not on the roadmap a year ago. It is not expected to scale beyond 12.8 GB/s, which establishes the need for either Wide IO or serial memory beyond 2014 for the mobile market.

Table 3: Summary of parameters of existing memory interfaces from the calculator.

Bound	LPDDR2	TSS-Wide IO	DDR3	Serial	Mobile-XDR
Clock Speed (MHz)	300-533, DDR	200-333, SDR	400-800, DDR	4-8 GHz, Serial	400-533, Octal
Throughput (GB/s)	3-4.3	12-24	6-13	12-17	12-17
Peak IO Power Efficiency (mW/Gbps)	~40	~10	~120	~60	~20
Peak Core Power Efficiency (mW/Gbps)	~50	~35	~100	~50	~50
Total Peak Power Efficiency (mW/Gbps)	~90	~45	~220	~110	~70
Active Idle IO Power (mW)	6-10	2-4	~500-600	~450	~200
Active Idle Core Power (mW)	~20	~20	~150	~20	~20
Capacity (GB) (Current trends)	0.5-1 for x32 dual rank	0.5-2 through multi-die stacking	2-8 for dual-rank DIMM	0.5-1 0.5-1 GB	0.5-1 for x32 dual rank
Latency from MC-DRAM-MC	~50ns	~40ns	~45ns, but penalty if DLL is off (~512 Tck)	~65ns, PLL lock penalty if off	~60ns, DLL penalty if off
Data width	x32	x512	x64	x16-x32	x32
Pin count	60	640 TSVs (No ext.)	130	60-100	60
Form-factor	POP/MCP	Stacked TSS	DIMM	MCP	MCP/POP




Figure 6: Gap in throughput (GB/s) before LPDDR3.

While our calculator currently focuses on the considerations elaborated in Sections 4.1 through 4.4, additional considerations listed below are planned.

- **Cost:** The cost of a mobile processor is affected by many factors beyond memory cost (Section 1). Die area depends on IO/PHY area and PDN/decap area; complex PHY retiming logic and adaptive schemes are significant area adders, being higher for serial links than for lower-speed parallel links. Package pincount is an issue for pin-limited design; serialization relieves pressure on this parameter, while 3D stacking eliminates package pins entirely. Technology integration cost, including cost of testing and yield, is higher for newer technologies such as 3D stacking. Footprint cost is relevant for small form-factor systems (POP and stacking enable smaller footprints than MCP and DIMM configurations). Finally, complexity and time-to-market are also important considerations when comparing memory interfaces.
- **Temperature:** Total power is directly related to peak temperature, and thermal gradients of the environment, including the package, can be a consideration for a mobile system. 3D stacking is especially sensitive to thermal considerations [9]. At the same time, memory interface power is only one component of the mobile processor power.

- **Tablet products:** The tablet market, being an evolution of the laptop and mobile markets, has options that sit on either side of the laptop and mobile segments in the decision tree. Design solutions for the tablet market are being offered as either high-end mobile devices or low-power laptop devices. The memory configurations that a tablet can support are thus limited by the memory configurations supported by its CPU IC, e.g., a high-end mobile processor (supporting LPDDR2) may not support the same memory configurations as a low-end laptop processor (supporting DDR3). Furthermore, tailored ICs for tablets may need to support a wider range of memory configurations to give the OEM the flexibility of choosing between them based on their specific system.
- **IO Programmability:** The need for higher capacity and throughput with larger power budgets changes the requirements for the memory interface on tablets. Support for more memory types and configurations means more programmability and complexity in the PHY and IO, e.g., supporting multiple voltages (1.5V DDR3, 1.2V LPDDR2) needs a programmable IO, and the need to support different power/throughput budgets demands programmable retiming and adaptive features on the IO and PHY. Optimizing pinout of the mobile processor to support different memory types is also a challenge, as the order and location of pins is not the same for POP, MCP and DIMM configurations. This challenge can be addressed by programmability (at the expense of timing and area) in the mobile controller and IO to support the different memory configurations.
- Parameters such as N_{pins} , T_{ck} , termination scheme, retiming scheme, etc. allow exploration outside the confines of existing interface standards. This enables identification of gaps in current commodity offerings (e.g., our previous discussion has highlighted the LPDDR3 gap between 8.6 GB/s and 12.8 GB/s, as well as the lack of a low-power, high-capacity alternative – LPDDR2 MCP [22], 3D stack within a DRAM package, or even a DIMM with non-DDR3 memory type – when performance is not a key requirement). It also straightforwardly enables projection of primary bounds for explo-

ration of potential new interface standards such as Wide IO DDR at 400MHz, or LPDDR3 with on-die termination or pseudo-open-drain signaling.

6. ACKNOWLEDGMENTS

We thank Prof. Mark Horowitz and Prof. Dean Tullsen for illuminating discussions, Roger Barth for DRAM trends used by the ITRS Test ITWG, and Anand Srinivasan, Matt Nowak, Riko Radocic, Michael Suh, Hari Rao and Jose Corleto for valuable inputs on mobile processor memory requirements.

7. REFERENCES

- [1] Y. Choi, H. Jeong and H. Kim, "Future Evolution of Memory Subsystem in Mobile Applications," *IEEE International Memory Workshop (IMW)*, 2010.
- [2] R. Palmer, J. Poulton, A. Fuller, J. Chen and J. Zerbe, "Design Considerations for Low-Power High-Performance Mobile Logic and Memory Interfaces," *Proc. IEEE ASSCC*, 2008, pp. 205-208.
- [3] M. Lee, W. Dally and P. Chiang, "Low-Power Area-Efficient High-Speed I/O Circuit Techniques," *IEEE JSSC* 35(11) (2000), pp. 1591-1599.
- [4] J. Dukovic et al., "Through-Silicon-Via Technology for 3D Integration," *IEEE International Memory Workshop (IMW)*, 2010.
- [5] Micron DRAM System Power Calculators: http://www.micron.com/support/dram/power_calc.html
- [6] ITRS PIDS Chapter: http://www.itrs.net/Links/2009ITRS/2009Chapters_2009Tables/2009_PIDS.pdf
- [7] ITRS PIDS Tables: http://www.itrs.net/Links/2010ITRS/2010Update/ToPost/2010Tables_PIDS_FOCUS_C_ITRS.xls
- [8] R. Barth, *personal communication*, February 2011.
- [9] "System-Level Impact of 3D", *GSRC Workshop* 2010.
- [10] JEDEC DDR3 Specification JESD79-3B.
- [11] JEDEC LPDDR2 Specification JESD209-2C.
- [12] Micron Data Sheets for DDR3 and LPDDR2.
- [13] Future-Mobile JEDEC Draft Wide IO Specification.
- [14] Future-Mobile JEDEC Draft Serial Memory Specification.
- [15] B. Jacob, "Trends in Memory Systems," *Sun-DARPA UNIC Architecture Workshop*, 2010.
- [16] QCT Memory Analysis, Qualcomm, Inc.
- [17] S. Woo, "DRAM and Memory System Trends," Rambus, Inc.
- [18] W. Dally and J. Poulton, *Digital Systems Engineering*, Cambridge University Press, 1998.
- [19] CACTI, <http://www.hpl.hp.com/research/cacti/>
- [20] S. Thoziyoor, N. Muralimanohar, J. Ahn and N. P. Jouppi, "CACTI 5.1. Technical Report," HP Laboratories, 2008.
- [21] Y. F. Tsai, Y. Xie, V Narayanan and M. J. Irwin, "Three-Dimensional Cache Design Exploration Using 3D Cacti," *Proc. IEEE ICCD*, Oct. 2005, pp. 519-524.
- [22] K. T. Malladi and M. A. Horowitz, *personal communication*, 2011.
- [23] M. A. Horowitz, C-K. K. Yang and S. Sidiropoulos, "High-Speed Electrical Signaling: Overview and Limitations," *IEEE Micro* 18 (1998), pp. 12-24.
- [24] B. Leibowitz et al., "A 4.3 GB/s Mobile Memory Interface with Power-Efficient Bandwidth Scaling," *IEEE JSSC* 45 (2010), pp. 136-137.
- [25] M. Floman, "Memories Are Made for This: Mobile Device Applications," *JEDEC Flash Storage Summit*, 2010.
- [26] Mobile-XDR technical brief http://www.rambus.com/assets/documents/products/mobile_xdr_brief.pdf
- [27] A. Ruberg, "Serial Port Memory: Bandwidth is Not Just for PCs Anymore," *Denali Memcon*, 2010.
- [28] J. Ellis, "Overcoming Obstacles for Closing Timing for DDR3-1600 and Beyond," *Denali Memcon*, 2010.
- [29] A. Vaidyanath, "Challenges and Solutions for GHz DDR3 Memory Interface Design," *Denali Memcon*, 2010.
- [30] D. Oh et al., "Design and Characterization of a 12.8 GB/s Low Power Differential Memory System for Mobile Applications," *Proc. IEEE EPEP*, 2009, pp. 33-36.
- [31] H. Hatamkhani, F. Lambrecht, V. Stojanovic and C-K. K. Yang, "Power-Centric Design of High-Speed I/Os," *Proc. ACM/IEEE DAC*, 2006, pp. 867-872.
- [32] S. Cho, "Technical Challenges in TSV Integration," *Keynote, RTI 3-D Conference*, 2010.
- [33] G. Balamurugan et al., "A Scalable 5-15 Gbps, 14-75 mW Low Power I/O Transceiver in 65nm CMOS," *IEEE JSSC* 43 (2008), pp. 1010-1019.
- [34] S. Thoziyoor, J. Ahn, M. Monchiero, J. B. Brockman and N. P. Jouppi, "A Comprehensive Memory Modeling Tool and its Application to the Design and Analysis of Future Memory Hierarchies," *Proc. IEEE ISCA*, 2008, pp. 51-62.
- [35] JEDEC: <http://www.jedec.org>