
A survey of architectural techniques for DRAM power management

Sparsh Mittal

Electrical and Computer Engineering,
Iowa State University,
Iowa, 50014, USA
E-mail: sparsh@iastate.edu

Abstract: Recent trends of CMOS technology scaling and wide-spread use of multicore processors have dramatically increased the power consumption of main memory. It has been estimated that modern data-centres spend more than 30% of their total power consumption in main memory alone. This excessive power dissipation has created the problem of ‘memory power wall’ which has emerged as a major design constraint inhibiting further performance scaling. Recently, several techniques have been proposed to address this issue. The focus of this paper is to survey several architectural techniques designed for improving power efficiency of main memory systems, specifically DRAM systems. To help the reader in gaining insights into the similarities and differences between the techniques, this paper also presents a classification of the techniques on the basis of their characteristics. The aim of this paper is to equip the engineers and architects with knowledge of the state of the art DRAM power saving techniques and motivate them to design novel solutions for addressing the challenges presented by the memory power wall problem.

Keywords: architectural techniques; power efficiency; energy saving; DRAM; main memory; survey; review; classification; system architecture.

Reference to this paper should be made as follows: Mittal, S. (2012) ‘A survey of architectural techniques for DRAM power management’, *Int. J. High Performance Systems Architecture*, Vol. 4, No. 2, pp.110–119.

Biographical notes: Sparsh Mittal received his BTech in Electronics and Communications Engineering from the Indian Institute of Technology, Roorkee, India. He was the graduating topper of his batch and his major project was awarded institute silver medal. At present, he is pursuing his PhD in Electrical and Computer Engineering at Iowa State University, USA. He has been awarded scholarship and fellowship from IIT Roorkee and ISU. His research interests include memory system power efficiency, cache architectures in multicore systems and real-time systems.

1 Introduction

Recent years have witnessed a dramatic increase in power consumption of computing systems. As an example, in year 2006 alone, the data-centres and servers in US consumed 61 billion kilowatt hours (kWh) of electricity (Brown et al., 2007). Further, the energy consumption of main memory is becoming an increasing fraction of total energy consumption of the system, in processors ranging from low-end to high-end. It has been estimated that as much as 40% of the total power consumed in smartphones and data-centres is attributed to the memory system (Barroso and Hölzle, 2009; Lefurgy et al., 2003; Ware et al., 2010). Studies conducted on real server systems show that memory system consumes as much as 50% more power than the processor cores (Lefurgy et al., 2003).

There are several architectural and technological trends that mandate the use of high amount of main memory resources and thus, contribute to the increase in memory power consumption. Firstly, in modern processors, the

number of cores on a single chip is on rise (Borkar, 2007; Intel, <http://ark.intel.com/products/53575/>) and hence, the pressure on the memory system has been increasing. Secondly, as we move to the exascale era, the requirements of data storage and computation are growing exponentially (Agrawal et al., 2008; Bergman et al., 2008; Khaitan et al., 2010; Raju et al., 2009) and to fully optimise the value of such data, modern processors are using main memory (as opposed to persistent storage) as the primary data storage for critical applications. Thirdly, modern data-centres exhibit low average utilisation (Barroso and Hölzle, 2009), but frequent, brief bursts of activity, and thus, to meet the requirements of service-level-agreements (SLAs), operators are forced to provision high amount of main memory resources, suitable to meet the worst case requirement. To cater to the above mentioned demands, modern processors are using main memory with high bandwidth, frequency and capacity. Finally, CMOS technology scaling has enabled higher transistor packing density, which have further exacerbated the problems of power consumption and

heating and have also inhibited the effectiveness of cooling solutions. Thus, improving the power efficiency of memory systems has become extremely important to continue to scale performance (Bergman et al., 2008) and also achieve the goals of sustainable computing.

Recently, many techniques have been proposed for optimising the power efficiency of memory systems. In this paper, we review several of these techniques. As it is practically infeasible to review all the techniques proposed in the literature, we take the following approach to limit the scope of this paper. We only include techniques proposed for saving energy in DRAM systems and do not discuss other emerging storage technologies (e.g., phase change memory). We consider DRAM, since it has been traditionally used as main memory because of its properties such as high density, high capacity, low cost and device standardisation, etc. (Cooper-Balis and Jacob, 2010). Also, we focus on *architecture-level* techniques which allow runtime power management and do not discuss *circuit-level* innovations for reducing power consumption. Further, although the techniques aimed at improving memory performance (e.g., reducing latency) are also likely to improve memory power efficiency, we only include those techniques that have been *evaluated* for improving memory power efficiency. Finally, since different techniques have been evaluated using different experimentation methodologies, we do not present their quantitative results; rather, we only discuss the key ideas of those techniques.

The remainder of this paper is organised as follows. In Section 2, we briefly discuss the terminology used in DRAM systems and the sources of power consumption in them. Understanding the sources of power consumption also helps in gaining insights into the opportunities available for improving power efficiency. In Section 3, we present a classification of the techniques proposed for managing DRAM power consumption to highlight the similarities and differences among them. A more detailed discussion of these techniques is provided in Section 4. Finally, Section 5 provides the concluding remarks.

2 Background

In this section, we briefly discuss DRAM terminology (Jacob et al., 2007) and the sources of power consumption in DRAM systems, to aid in discussion of DRAM power management techniques discussed in the next sections.

2.1 DRAM terminology

In DRAM terminology, a *column* is the smallest addressable portion of the DRAM device and a *row* is a group of bits in the DRAM-array that are sensed together at the instance of receiving an activate signal. A DRAM *bank* is an array of DRAM cells, which can be active independently and has same data bus width as external output bus width. A DRAM *rank* is a group of DRAM devices which operate together to service requests from the memory controller. A dual in-line memory module (*DIMM*) is a printed circuit board

containing one or more DRAM ranks on it, which provides an interface to the memory bus. A DRAM *channel* is a group of one or more DIMMs of DRAM that handle requests from the memory controller. A typical DRAM can have two channels, two DIMMs per channel, two ranks per DIMM, eight banks per rank, for a total of 64 ($= 2 \times 2 \times 2 \times 8$) banks.

2.2 Sources of power consumption

The power consumption in DRAM memory is broadly classified in three categories, namely activation power, read/write power, background power [for a more detailed analysis, see Cooper-Balis and Jacob (2010) and Vogelsang (2010)]. Activation power refers to the power dissipated in activating a memory array row and in precharging the arrays bitlines. The read/write power refers to the power which is consumed when the data moves either into or out of the memory device. The background power is independent of the DRAM access activity and is due to the transistor leakage, peripheral circuitry, and the data refresh operations. Note that the DRAM memory cells store data using capacitors that lose their charge over time and must be periodically recharged; this is referred to as data refreshing.

3 A classification of DRAM power management techniques

In this section, we present a classification of the DRAM power management techniques based on their characteristics.

As shown by Barroso and Hölzle (2009), modern servers operate most of the time between 10% and 50% of maximum utilisation. Thus, considerable opportunities exist to transition idle inactive memory banks into low-power modes. These power modes could be either state-preserving modes (i.e., the data is retained) or state-destroying modes (i.e., the data is not retained). For this purpose, DRAM chips provision several modes of operation. Each mode is characterised by its power consumption and the time that it takes to transition back to the active mode. This time is referred to as resynchronisation latency or exit latency. Typically, the modes with lower energy consumption also have higher reactivation time and vice versa. Also, a DRAM module may enter a low-power state-preserving mode when it is idle, but must return to the active mode to service a request. A large number of techniques have been proposed which utilise the adaptive power saving capability offered by the modern multi-banked memory systems (Ware et al., 2010; Amin and Chishti, 2010; Anagnostopoulou et al., 2012; Ayoub et al., 2010; Ben Fradj et al., 2006; Bi et al., 2010; Chandrasekar et al., 2012; Chen et al., 2002b, 2011; David et al., 2010; Delaluz et al., 2001, 2002a, 2002b, 2002c; Deng et al., 2011, 2012; Diniz et al., 2007; Fan et al., 2001; Floyd et al., 2007; Huang et al., 2003, 2005; Hur and Lin, 2008; Irani et al., 2003; Kandemir et al., 2001; Khargharia et al., 2007; Koc et al., 2006; Lebeck et al., 2000; Li et al., 2004, 2007; Liu et al., 2011b; Lyuh and

Kim, 2004; Malladi et al., 2012; Meisner et al., 2009; Mukundan and Martinez, 2012; Ozturk et al., 2006; Ozturk and Kandemir, 2008; Pandey et al., 2006; Pisharath et al., 2004; Rodero et al., 2010; Sudan et al., 2012; Udipi et al., 2010, 2012; Wang and Hu, 2005; Zheng et al., 2008; Zheng and Zhu, 2010; Zhou et al., 2004).

Some techniques perform memory access redistribution (also called memory traffic reshaping), which involves changing the address mapping in DRAM, migrating data within DRAM, etc., for increasing the idle period of certain memory banks or increasing memory reference locality (Amin and Chishti, 2010; Ayoub et al., 2010; Delaluz et al., 2002a; Huang et al., 2005; Khargharia et al., 2007; Koc et al., 2006; Lebeck et al., 2000; Ozturk et al., 2006; Ozturk and Kandemir, 2008; Pandey et al., 2006; Kaseridis et al., 2011; Kim and Park, 2001; Li and Zhang, 2012; Sudan et al., 2010; Tolentino et al., 2009).

Several other techniques reduce the power consumed in each DRAM access by accessing only part of DRAM in each memory access (Cooper-Balis and Jacob, 2010; Udipi et al., 2012; Zheng et al., 2008; Ahn et al., 2009; Fang et al., 2010; Seongil et al., 2011; Zhang et al., 2012). Thus, these techniques provision activating a much smaller portion of DRAM circuit component than what is activated in conventional DRAMs.

Many techniques for reducing memory power consumption use mechanisms based on memory access scheduling (Amin and Chishti, 2010; Floyd et al., 2007; Lyuh and Kim, 2004; Hanson and Rajamani, 2012; Lin et al., 2007b, 2008; Liu et al., 2008; Trajkovic et al., 2008; Yoon et al., 2011, 2012; Zheng et al., 2009). This includes mechanisms such as memory access throttling or buffering or coalescing. Also, techniques have been proposed which reduce the number of accesses to memory by smart management of last level caches (LLCs) or programme level transformations, etc. (Amin and Chishti, 2010; Aggarwal et al., 2008; Isen and John, 2009; Mazumdar et al., 2010).

Some techniques use data compression to reduce the memory footprint of the application (Ozturk and Kandemir, 2008; Chen et al., 2002a; Tremaine et al., 2001; Yang et al., 2010). This helps in reducing the number of memory banks occupied by application data. The techniques based on data replication increase idle time duration of memory banks by duplicating their selected read-only data blocks on other active banks (Ozturk and Kandemir, 2008).

Like other CMOS circuits, DRAM power also depends on operating frequency and supply voltage and hence DVFS (dynamic voltage/frequency scaling) mechanism has been used in several techniques to save memory power (Deng et al., 2011; Tolentino et al., 2009; Lin et al., 2008; David et al., 2011). Some techniques reduce the DRAM refresh power, by trading off either performance or reliability (Liu et al., 2011b; Isen and John, 2009; Ghosh and Lee, 2007; Liu et al., 2012; Ohsawa et al., 1998; Patel et al., 2005; Phadke and Narayanasamy, 2011; Stuecheli et al., 2010). Some techniques account for the effect of temperature while optimising for memory energy (Ayoub et al., 2010; Lin et al., 2008; Lin et al., 2007a, 2009; Liu et al., 2011a). Such

techniques are referred to as thermal-aware memory energy saving techniques.

Also, while most of the techniques work to reduce total power consumption of DRAM system, a few techniques work to *limit* their average power consumption (David et al., 2010), while some other techniques work to limit their peak power consumption (Chen et al., 2011; Diniz et al., 2007).

4 DRAM power management techniques

In this section, we review several DRAM power saving techniques. As discussed before, we only present the key ideas of each technique and do not discuss their qualitative results.

Lebeck et al. (2000) propose a technique for turning off DRAM chips in low-power mode. Their technique works by controlling virtual address to physical address mapping such that the physical pages of an application are clustered into a minimum number of DRAM chips and the unused chips are transitioned to low power modes. In addition, their technique also monitors the time period between accesses to a chip as a metric for measuring the frequency of reuse of a chip. When this time is greater than a certain threshold, the chip is transitioned to the low-power mode.

Fan et al. (2001) present an analytical model to approximate the idle time of memory chips. Based on this, they identify the threshold time, after which the memory chip can transitioned to low-power state. They observe that, for their experimentation framework, the simple policy of immediately transitioning a DRAM chip to a low-power mode when it becomes idle, performs better than more sophisticated policies that predict DRAM chip idle time. Li et al. (2004) propose a technique for saving memory energy which adaptively transitions memory module to low-power modes while also providing guarantees on the maximum performance loss.

Delaluz et al. (2001) propose software and hardware-based approaches to save memory energy. Their hardware based approach works by estimating the time of next access to a memory bank and then, depending upon the time, switching the bank to a suitable low-power mode. The software-directed approach uses compiler analysis to insert memory module transition instructions in the programme binary. To avoid the time overhead of resynchronisation, they propose bringing the memory bank to active mode before its next use. Depending upon the break-even analysis of length of idle time and power saving opportunity of different power saving modes, a suitable power saving mode is chosen. Delaluz et al. (2002c) describe an OS scheduler-based power mode control scheme for saving DRAM energy. Their scheme tracks memory banks that are being used by different applications and selectively turns on/off these banks at context switch points.

A limitation of the approaches based on power mode control is that most of the idle times between different accesses to memory ranks are shorter than the resynchronisation time between different power modes. To address this issue, Huang et al. (2005) propose a method for

saving memory energy by concentrating the memory access activities to merely few memory ranks, such that rest of the ranks can be switched to low-power modes. Their method migrates the frequently-accessed pages to 'hot' ranks and the infrequently-used and unmapped pages to 'cold' ranks. This also helps in elongating the idle periods of cold ranks.

Delaluz et al. (2002a) propose a technique to save memory energy by dynamically placing the arrays with temporal affinity into the same set of banks. This increases the opportunities of exploiting deeper sleep modes (more energy-saving operating modes) and keeping modules in low-power modes for longer durations of time. Using the same principle, they also propose an array interleaving mechanism (Delaluz et al., 2002b) for clustering multiple arrays, which are accessed simultaneously, into a single common data space. Interleaving enhances spatial locality of the programme and reduces the number of accesses to the off-chip memory. Along with interleaving the arrays, their mechanism also transforms the code accordingly by replacing the original array references and declarations with their transformed equivalents.

Huang et al. (2003) propose a technique for saving memory energy using virtual memory management. Their technique works by using virtual memory remapping to reduce the memory footprint of each application and transitioning the unused memory modules to low-power modes. Zhou et al. (2004) discuss a utility-based memory allocation scheme where different applications are allocated memory in proportion to their utility (i.e., the performance benefit gained by allocation of the memory). After allocation, the rest of the memory is transitioned to low-power modes for saving power. For estimating the utility of allocating memory to different applications, their scheme dynamically tracks page miss-rate curve (MRC) for virtual memory system using either hardware or software methods. Lyuh and Kim (2004) propose a technique which uses analytical model for saving memory power. Their technique selects a suitable low-power mode for a memory bank by synergistically controlling assignment of variables to memory banks and scheduling of memory access operations, such that total memory power consumption is minimised.

Bi et al. (2010) propose methods to hide the latency of resynchronisation of memory ranks to low-power modes by exploiting the knowledge of system input/output (I/O) calls. Their technique works on the observation that a majority of file-I/O accesses are made through system calls, the operating system knows the completion time of these accesses. Thus, using this knowledge, their technique transitions idle memory ranks into low-power modes. Further, to hide the resynchronisation delay, their technique uses prediction mechanism to estimate the most likely rank to be accessed on a system call entry and speculatively turns on that rank. On a correct prediction, the rank transition completes before the memory request arrives and thus, the resynchronisation latency is fully hidden.

Pandey et al. (2006) propose a technique for saving energy in direct memory access (DMA) transfers. Since

DMA transfers are usually larger than the transfers initiated by the processors, they are divided into multiple number of smaller transfer operations. However, due to the availability of only short time gaps between any two DMA-memory requests, the opportunity of transitioning the memory to low-power mode remains small. To address this, Pandey et al. propose temporally aligning DMA requests coming from different I/O buses to the same memory device. For this purpose, their technique delays DMA-memory requests directed to a memory chip which is in low-power mode and tries to gather enough requests from other I/O buses before transitioning that chip to the normal power mode. This helps in elongating the idle time of memory chips and also maximises the utilisation of active time of memory chips.

Koc et al. (2006) discuss a data-recomputation approach for increasing idle time of memory banks to save their energy. When an access to a bank in low-power mode is made, their technique first checks the active banks. If the requested data can be recomputed by using the data stored in already active banks, their technique does not activate the bank in low-power mode. Rather, the data request is fulfilled based on the computations performed on the data obtained from the already active banks.

To reduce the refresh power of the DRAM devices, Ghosh and Lee (2007) propose adaptive refresh method. The conventional refresh mechanism of DRAM periodically refreshes all the memory rows for retaining the data. However, from the standpoint of data retention, an access to a memory row performs an operation equivalent to a regular refresh. The technique proposed by Ghosh and Lee (2007) uses this observation to avoid refreshing a memory row, if it has been recently read out or written to by the processor. To track the recency of memory access operation, they use counters for each row in the memory module. Using this technique, the number of regular row-sweeping refresh operations are greatly reduced, which results in saving of power.

Liu et al. (2012) propose a technique for saving DRAM energy by avoiding unbeneficial refreshes. Their technique works on the observation that in the DRAM, only a small number of cells need to be refreshed at the minimum conservative refresh rate. The rest of the cells can be refreshed at a much higher rate, while still maintaining their charge. Based on this observation, their technique groups DRAM rows in multiple bins and uses different refresh interval for different bins. Thus, by refreshing most of the cells less frequently than the leaky cells, their technique reduces the number of refresh operations required and reduces memory power consumption.

Isen and John (2009) discuss a technique for utilising programme semantics to save memory energy. Their technique uses memory allocation/deallocation information to identify inconsequential data and avoids refreshing them. For example, the regions of memory which are free (unallocated and invalid) or freshly allocated (allocated but invalid) do not store meaningful data and hence, retaining the data of those regions is not important. Thus, their technique saves power by avoiding refreshing such data.

Liu et al. (2011b) propose an application level technique to reduce refresh level power in DRAM memories. They show that many applications are tolerant to errors in the non-critical data, and errors in non-critical data show little or no impact in the application's final result. Based on this observation, their technique works by using programmer supplied information to identify critical and non-critical data in the programmes. Using this information, at runtime, these data are allocated in different modules of the memory. The memory modules containing critical data are refreshed at the regular refresh-rate, while the modules containing non-critical data are refreshed at substantially lower rates. The use of lower refresh rates leads to saving in refresh power, however, it also increases the probability of data corruption. Thus, their technique exercises a trade-off between energy saving and data corruption.

Sudan et al. (2010) propose a technique for saving memory power by using OS management approach. Their technique works by controlling the address mapping of OS pages to DRAM devices such that the clusters of cache blocks from different OS pages, which have similar access counts are co-located in a row-buffer. This improves the hit rate of the row-buffer and thus leads to saving of memory power. For co-locating pages, Sudan et al. propose two techniques. One of their technique reduces OS page size such that the frequently accessed blocks are clustered together in the new, reduced size page (called a 'micro-page'). Then, the hot micro-pages are migrated in the same row-buffer. The second technique proposed by them uses a hardware scheme. This scheme introduces a layer of translation between physical addresses assigned by the OS and those used by the memory controller to access the DRAM devices. By taking advantage of this layer of mapping, hot pages are migrated in the same row-buffer.

Trajkovic et al. (2008) propose a buffering-based technique for reducing memory power consumption. Their technique works on the observation that if in a synchronous DRAM, two memory access (i.e., read/write) operation are done in a same activate-precharge cycle, the cost of activation and precharging can be avoided. This is because, the DRAMs allow the row to be left 'on' after a memory access. Based on this observation, on read accesses, their technique prefetches additional cache blocks. Similarly, for write accesses, combines multiple blocks which are to be written to the same DRAM row. To store the extra prefetched lines, their technique uses a small storage structure in the memory controller. Similarly, to buffer the writes to the same DRAM row also, a small storage structure is used. By adapting the above mentioned prefetching and write-combining scheme for each application, their technique achieves reduction in memory power consumption.

Zheng et al. (2008) propose a technique for saving memory power consumption by reducing the number of memory chips involved in each memory access. This is referred to as 'rank-subsetting' approach. Their technique adds a small buffer called 'mini-rank buffer' between each DIMM and the memory bus. Using this, a DRAM rank,

which normally provides 64-bit datapath, can be internally designed using either eight 8-bit ranks, or four 16-bit ranks or two 32-bit ranks, which are termed as mini-rank. With this support, on any memory access, only a single mini-rank is activated and the other mini-ranks can be transitioned to low-power modes.

Fang et al. (2010) extend mini-rank approach to heterogeneous mini-rank design which adapts the number of mini-ranks according to the memory access behaviour and memory bandwidth requirement of each workload. Based on this information, for a latency-sensitive application, their technique uses a mini-rank configuration which does not degrade application performance; while for a latency-insensitive application, their technique uses a mini-rank configuration which achieves memory power saving.

Yoon et al. (2012) propose a technique for saving memory power consumption by intelligently utilising low-power mobile DRAM components. Their technique uses buffering mechanism to aggregate the data outputs from multiple ranks of low frequency mobile DRAM devices (e.g., 400 MHz LPDDR2), to collectively provide high bandwidth and high storage capacity equal to server-class DRAM devices (e.g., 1,600 MHz DDR3).

Yoon et al. (2011) propose a technique for saving memory power by dynamically changing the granularity of data transferred in each DRAM access. Their technique works by managing virtual memory such that a specific access granularity can be used for each page based on the spatial locality present in each application. For applications with high spatial locality, their technique uses coarse-grained data accesses, while for applications with low spatial locality their technique uses fine-grained data accesses.

Several researchers have proposed techniques which use DVFS mechanism to save memory energy. Deng et al. (2011, 2012) use memory DVFS to save memory energy. At the time of low memory activity, their technique lowers the frequency of DRAM-devices, memory channels and memory controllers such that the performance loss is minimum. This leads to saving of memory power consumption. They also extend their technique for coordinating DVFS across multiple memory controllers, memory channels, and memory devices to minimise the overall system power consumption.

Diniz et al. (2007) propose a technique to limit the *instantaneous* (peak) power consumption of main memory to a pre-specified power budget. Their technique uses knapsack and greedy algorithms to decide the timings at which memory devices should be transitioned to suitable low-power modes such that the instantaneous power of memory is always within the power budget. David et al. (2010) present a scheme for limiting the average power consumption of memory, by suitably transitioning memory devices into low-power modes. Chen et al. (2011) propose a method for limiting the peak power consumption of the server (which includes power consumption of processor and main memory) system using control theoretic approach.

Amin and Chishti (2010) propose a replacement policy for last-level cache, which tries to increase the idle time of certain pre-chosen DRAM ranks, called 'prioritised ranks'. This replacement policy tries to prevent the replacement of blocks belonging to the prioritised ranks. This reduces the conflict misses and writebacks to the prioritised rank, increasing the idle period between accesses made to those ranks. Amin and Chishti also propose a technique which buffers writeback requests sent to DRAM, to increase the idle period of the DRAM ranks. The requests are buffered as long as the target ranks remain idle or the buffer remains full. When the targeted ranks become active (due to the demand misses), the buffered requests are sent to them.

Ozturk et al. (2006) present a bank-aware cache miss clustering approach for saving DRAM energy. Their technique uses compiler analysis to restructure the code such that the cache misses from the last-level cache are clustered together. Clustering of the cache misses also leads to the clustering of cache hits. Thus, the memory accesses and memory idle cycles are also clustered. This increases the memory access activities in certain banks and the other banks become idle for a long time. By taking advantage of this, idle memory banks are transitioned to low-power modes.

As the computational requirements of state-of-the-art applications is increasing (Khaitan et al., 2009), the pressure on memory systems is also on rise and to mitigate this pressure, researchers have proposed techniques to intelligently manage the LLCs in the processors. Mazumdar et al. (2010) propose a technique for reducing the number of memory accesses in multicore systems by cache aggregation approach. Their technique works on the observation that due to the availability of high-bandwidth point-to-point interconnects between sockets, a read from the LLC of a connected chip consumes less time and energy than an access to DRAM. Based on this, their technique uses the LLC of an idle processor in a connected socket for holding the evicted data from the active processor. This reduces the number of accesses to DRAM and thus reduces the power consumption of DRAM.

Phadke and Narayanasamy (2011) propose a heterogeneous main memory architecture which comprises of three different memory modules. Each memory module is optimised for latency, bandwidth, power consumption, respectively, at the expense of the other two. Their technique works by using offline analysis to characterise an application based on its LLC miss rate and memory level parallelism. Using this information, at runtime, the operating system allocates the pages of an application in one of the three memory modules that satisfies its memory requirements. Thus, their approach saves memory energy and also improves performance of the system.

Yang et al. (2010) discuss a software-based RAM compression technique for saving power in embedded systems. Their technique uses memory compression only for those applications which may gain performance or energy benefits from compression. For such applications, their technique performs compression of memory data and

swapped-out pages in online manner, thus dynamically adjusting the size of the compressed RAM area. Thus, their technique saves power by using compression to increase the effective size of the memory.

Ozturk and Kandemir (2008) integrate different approaches such as dynamic data migration, data compression, and data replication, etc. to effectively transition a large number of memory banks into low-power modes. They formulate DRAM energy minimisation problem as an integer linear programming (ILP) problem and solve it using an ILP solver. Using ILP formulation, they find the (non-uniform) bank architecture and accompanying data mapping strategy which best suits the application-data access patterns. Similarly, they use ILP formulation to find best possible data replication scheme which increases idle time of certain banks by duplicating their selected read-only data blocks on other active banks. They also use ILP formulation to find the best time to compress and/or migrate the data between banks.

Several researchers have used domain-specific optimisations to save DRAM power. Kim and Park (2001) and Li and Zhang (2012) propose techniques for reducing DRAM power consumption in video processing domain. Video processing applications are characterised by abundant spatial and temporal image data correlations, and unbalanced accesses to frames (e.g., certain image frames are accessed much more frequently than other image frames). Hence, to take advantage of these properties, their techniques map image data in DRAM in a way which minimises the number of row-activations. Thus, the power consumption of DRAM is reduced.

Chen et al. (2002b) propose a technique for tuning the garbage collector (GC) in Java to reduce memory power consumption. GC is a tool used in Java virtual machine (JVM) for automatic reclamation of unused memory. Chen et al. propose using GC to turn off the memory banks that do not hold live data. They also observe that the pattern of object allocation and the number of memory banks available in the DRAM architecture crucially influence the effectiveness of GC in optimising energy.

Pisharath et al. (2004) propose an approach to reduce memory power consumption in memory-resident database management systems (DBMS). One of their techniques uses hardware monitors to detect the frequency of use of memory banks during query execution and based on this, switches the idle banks into low-power mode. Another technique uses a software approach. For DBMS, when the query is submitted, it is first parsed and then sent to the query optimiser which uses query tree to find the best suited plan for execution of the query (Pisharath et al., 2004). At this point, query optimiser finds the database tables which will be accessed to answer the query. Based on this information, their technique changes the table-to-bank mapping such that memory accesses can be clustered. Also, the queries presented to the database are augmented with explicit bank turn off or turn on instructions. Using this support, at runtime, the memory banks are dynamically transitioned into low-power mode. Since leakage (static) power varies

exponentially with the temperature, the dissipation of power in DRAM leads to increase of device temperature, which further increases the leakage power dissipation. This may lead to thermal emergencies. Also, many of the above mentioned approaches move or map frequently accessed pages to merely a few active memory modules. This is also likely to increase the temperature of the active modules. To address this, Ayoub et al. (2010) propose a technique which monitors the temperature of the active modules. When the temperature reaches a threshold, it selectively migrates a small number of memory pages between active and dormant memory modules and transitions the active modules in the self-refresh mode. Since this approach spreads out the memory accesses to multiple modules, it reduces the power density of the active modules and thus avoids thermal emergencies.

Lin et al. (2009) propose a technique for addressing memory thermal issues which works by orchestrating thread scheduling and page allocation. Their technique groups the programme threads in multiple groups such that all the threads in a group can be active simultaneously. Then each group is mapped to certain DIMMs and at any time, only one group and its corresponding DIMMs remain active and the rest of the DIMMs are inactivated to reduce their temperature. Similarly, Lin et al. (2007a, 2008) propose techniques to mitigate overheating in the memory system by adjusting memory throughput to stay below the emergency level.

5 Conclusions

Recent advances in CMOS fabrication and chip design have greatly increased the power consumption of main memory in modern computing systems. To provide a solution to this problem, several research efforts have been directed towards managing the power consumption of main memory. In this paper, we surveyed several architectural techniques which are designed for improving DRAM memory power efficiency. We also presented a classification of the proposed techniques across several parameters, to highlight their similarities and differences. We believe that this survey will help researchers and designers to understand the state of the art in approaches pursued for reducing memory power consumption. At the same time, it will also encourage them to design innovative solutions for memory systems of future green computing infrastructure.

References

- Aggarwal, N., Cantin, J., Lipasti, M. and Smith, J. (2008) 'Power-efficient DRAM speculation', in *HPCA*, IEEE, pp.317–328.
- Agrawal, A. et al. (2008) 'A new heuristic for multiple sequence alignment', in *IEEE EIT*, pp.215–217, DOI: <http://dx.doi.org/10.1109/EIT.2008.4554299>.
- Ahn, J., Jouppi, N., Kozyrakis, C., Leverich, J. and Schreiber, R. (2009) 'Future scaling of processor-memory interfaces', in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, ACM, p.42.
- Amin, A. and Chishti, Z. (2010) 'Rank-aware cache replacement and write buffering to improve DRAM energy efficiency', in *Proceedings of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design*, ACM, pp.383–388.
- Anagnostopoulou, V., Biswas, S., Saadeldeen, H., Savage, A., Bianchini, R., Yang, T., Franklin, D. and Chong, F.T. (2012) 'Barely alive memory servers: Keeping data active in a low-power state', in *ACM Journal on Emerging Technologies in Computing Systems, Special issue on Sustainable and Green Computing Systems*, April, in press.
- Ayoub, R., Indukuri, K. and Rosing, T. (2010) 'Energy efficient proactive thermal management in memory subsystem', in *International Symposium on Low-Power Electronics and Design (ISLPED)*, IEEE, pp.195–200.
- Barroso, L. and Hözl, U. (2009) 'The datacenter as a computer: an introduction to the design of warehouse-scale machines', *Synthesis Lectures on Computer Architecture*, Vol. 4, No. 1, pp.1–108.
- Ben Fradj, H., Belleudy, C. and Auguin, M. (2006) 'System level multi-bank main memory configuration for energy reduction', in Vounckx, J., Azemard, N. and Maurine, P. (Eds.): *Integrated Circuit and System Design. Power and Timing Modeling, Optimization and Simulation, Lecture Notes in Computer Science*, Vol. 4148, pp.84–94, Springer Berlin/Heidelberg, DOI: http://dx.doi.org/10.1007/11847083_9.
- Bergman, K. et al. (2008) 'Exascale computing study: technology challenges in achieving exascale systems', Tech. rep., DARPA Technical Report.
- Bi, M., Duan, R. and Gniady, C. (2010) 'Delay-hiding energy management mechanisms for DRAM', in *HPCA*, IEEE, pp.1–10.
- Borkar, S. (2007) 'Thousand core chips: a technology perspective', in *Proceedings of the 44th Annual DAC*, ACM, pp.746–749.
- Brown, R., Masanet, E., Nordman, B., Tschudi, B., Shehabi, A., Stanley, J., Koomey, J., Sartor, D., Chan, P., Loper, J. et al. (2007) 'Report to congress on server and data center energy efficiency', *Public Law*, 109–431.
- Chandrasekar, K., Akesson, B. and Goossens, K. (2012) 'Runtime power-down strategies for real-time SDRAM memory controllers', in *Proceedings of the 49th Annual DAC*, ACM, pp.988–993.
- Chen, G., Kandemir, M., Vijaykrishnan, N., Irwin, M.J. and Wolf, W. (2002a) 'Energy savings through compression in embedded java environments', in *Proceedings of the Tenth International Symposium on Hardware/software Codesign, CODES '02*, ACM, pp.163–168.
- Chen, G., Shetty, R., Kandemir, M., Vijaykrishnan, N., Irwin, M.J. and Wolczko, M. (2002b) 'Tuning garbage collection for reducing memory system energy in an embedded java environment', *ACM Trans. Embed. Comput. Syst.*, November, Vol. 1, No. 1, pp.27–55.
- Chen, M., Wang, X. and Li, X. (2011) 'Coordinating processor and main memory for efficient server power control', in *International Conference on Supercomputing, ICS '11*, ACM, New York, NY, USA, pp.130–140.

- Cooper-Balis, E. and Jacob, B. (2010) 'Fine-grained activation for power reduction in DRAM', *Micro*, Vol. 30, No. 3, pp.34–47, IEEE.
- David, H., Fallin, C., Gorbato, E., Hanebutte, U.R. and Mutlu, O. (2011) 'Memory power management via dynamic voltage/frequency scaling', in *Proceedings of the 8th ACM International Conference on Autonomic Computing, ICAC '11*, ACM, New York, NY, USA, pp.31–40.
- David, H., Gorbato, E., Hanebutte, U.R., Khanna, R. and Le, C. (2010) 'RAPL: memory power estimation and capping', in *International Symposium on Low-Power Electronics and Design (ISLPED)*, August, pp.189–194.
- Delaluz, V., Kandemir, M. and Kolcu, I. (2002a) 'Automatic data migration for reducing energy consumption in multi-bank memory systems', in *DAC*, IEEE, pp.213–218.
- Delaluz, V., Kandemir, M., Vijaykrishnan, N., Irwin, M., Sivasubramaniam, A. and Kolcu, I. (2002b) 'Compiler-directed array interleaving for reducing energy in multi-bank memories', in *ASPAC*, IEEE, pp.288–293.
- Delaluz, V., Sivasubramaniam, A., Kandemir, M., Vijaykrishnan, N. and Irwin, M. (2002c) 'Scheduler-based DRAM energy management', in *DAC*, ACM, pp.697–702.
- Delaluz, V., Kandemir, M., Vijaykrishnan, N., Sivasubramaniam, A. and Irwin, M. (2001) 'DRAM energy management using software and hardware directed power mode control', in *HPCA*, IEEE, pp.159–169.
- Deng, Q., Meisner, D., Bhattacharjee, A., Wenisch, T.F. and Bianchini, R. (2012) 'MultiScale: memory system dvfs with multiple memory controllers', in *International Symposium on Low Power Electronics and Design (ISLPED)*, July.
- Deng, Q., Meisner, D., Ramos, L., Wenisch, T. and Bianchini, R. (2011) 'Memscale: active low-power modes for main memory', *ACM SIGPLAN Notices*, Vol. 46, No. 3, pp.225–238.
- Diniz, B., Guedes, D., Meira, W., Jr. and Bianchini, R. (2007) 'Limiting the power consumption of main memory', in *ACM SIGARCH Computer Architecture News*, Vol. 35, No. 2, pp.290–301, ACM.
- Fan, X., Ellis, C. and Lebeck, A. (2001) 'Memory controller policies for DRAM power management', in *Proceedings of the 2001 International Symposium on Low Power Electronics and Design*, ACM, pp.129–134.
- Fang, K., Zheng, H. and Zhu, Z. (2010) 'Heterogeneous minirank: Adaptive, power-efficient memory architecture', in *39th International Conference on Parallel Processing (ICPP)*, IEEE, pp.21–29.
- Floyd, M., Ghiasi, S., Keller, T., Rajamani, K., Rawson, F., Rubio, J. and Ware, M. (2007) 'System power management support in the IBM POWER6 microprocessor', *IBM Journal of Research and Development*, Vol. 51, No. 6, pp.733–746.
- Ghosh, M. and Lee, H. (2007) 'Smart refresh: an enhanced memory controller design for reducing energy in conventional and 3D die-stacked DRAMs', in *MICRO-40*, pp.134–145, IEEE Computer Society, DOI: <http://dx.doi.org/10.1109/MICRO.2007.38>.
- Hanson, H. and Rajamani, K. (2012) 'What computer architects need to know about memory throttling', in *Computer Architecture*, pp.233–242, Springer, DOI: http://dx.doi.org/10.1007/978-3-642-24322-6_20.
- Huang, H., Pillai, P. and Shin, K. (2003) 'Design and implementation of power-aware virtual memory', *USENIX Annual Technical Conference*, pp.57–70.
- Huang, H., Shin, K., Lefurgy, C. and Keller, T. (2005) 'Improving energy efficiency by making DRAM less randomly accessed', in *Proceedings of the 2005 International Symposium on Low Power Electronics and Design*, ACM, pp.393–398.
- Hur, I. and Lin, C. (2008) 'A comprehensive approach to DRAM power management', in *14th International Symposium on High Performance Computer Architecture, 2008, HPCA*, IEEE, pp.305–316.
- Intel [online] <http://ark.intel.com/products/53575/> (accessed 26 September 2012).
- Irani, S., Shukla, S. and Gupta, R. (2003) 'Online strategies for dynamic power management in systems with multiple power-saving states', *ACM Transactions on Embedded Computing Systems (TECS)*, Vol. 2, No. 3, pp.325–346.
- Isen, C. and John, L. (2009) 'Eskimo-energy savings using semantic knowledge of inconsequential memory occupancy for DRAM subsystem', in *MICRO-42*, pp.337–346, IEEE, DOI: <http://dx.doi.org/10.1145/1669112.1669156>.
- Jacob, B., Ng, S. and Wang, D. (2007) *Memory Systems: Cache, DRAM, Disk*, Morgan Kaufmann Publication, Burlington, MA, USA.
- Kandemir, M., Sezer, U. and Delaluz, V. (2001) 'Improving memory energy using access pattern classification', in *International Conference on Computer-Aided Design*, IEEE Press, pp.201–206.
- Kaseridis, D., Stuecheli, J. and John, L.K. (2011) 'Minimalist open-page: a DRAM page-mode scheduling policy for the many-core era', in *MICRO, MICRO-44 '11*, New York, NY, USA, ACM, pp.24–35.
- Khaitan, S. et al. (2009) 'Fast parallelized algorithms for on-line extended-term dynamic cascading analysis', in *IEEE/PES PSC*, pp.1–7.
- Khaitan, S., McCalley, J. and Raju, M. (2010) 'Numerical methods for on-line power system load flow analysis', *Energy Systems*, Vol. 1, No. 3, pp.273–289.
- Khargharia, B., Hariri, S. and Yousif, M.S. (2007) 'Self-optimization of performance-per-watt for interleaved memory systems', in *14th International Conference on High Performance Computing, HiPC'07*, Springer-Verlag, Berlin, Heidelberg, pp.368–380.
- Kim, H. and Park, I-C. (2001) 'High-performance and low-power memory-interface architecture for video processing applications', *IEEE Transactions on Circuits and Systems for Video Technology*, November, Vol. 11, No. 11, pp.1160–1170.
- Koc, H., Ozturk, O., Kandemir, M. and Ercanli, E. (2006) 'Minimizing energy consumption of banked memories using data recomputation', in *International Symposium on Low Power Electronics and Design*, pp.358–361.
- Lebeck, A., Fan, X., Zeng, H. and Ellis, C. (2000) 'Power aware page allocation', *ACM SIGPLAN Notices*, Vol. 35, No. 11, pp.105–116.
- Lefurgy, C., Rajamani, K., Rawson, F., Felter, W., Kistler, M. and Keller, T. (2003) 'Energy management for commercial servers', *Computer*, Vol. 36, No. 12, pp.39–48.
- Li, X., Gupta, R., Adve, S. and Zhou, Y. (2007) 'Cross-component energy management: joint adaptation of processor and memory', *ACM Transactions on Architecture and Code Optimization (TACO)*, Vol. 4, No. 3, p.14.
- Li, X., Li, Z., David, F., Zhou, P., Zhou, Y., Adve, S. and Kumar, S. (2004) 'Performance directed energy management for main memory and disks', in *ACM SIGARCH Computer Architecture News*, Vol. 32, No. 5, pp.271–283, ACM.

- Li, Y. and Zhang, T. (2012) 'Reducing DRAM image data access energy consumption in video processing', *IEEE Transactions on Multimedia*, April, Vol. 14, No. 2, pp.303–313.
- Lin, C., Yang, C. and King, K. (2009) 'PPT: joint performance/power/thermal management of DRAM memory for multi-core systems', in *Proceedings of the 14th ACM/IEEE International Symposium on Low Power Electronics and Design*, ACM, pp.93–98.
- Lin, J., Zheng, H., Zhu, Z., David, H. and Zhang, Z. (2007a) 'Thermal modeling and management of DRAM memory systems', in *ISCA*, ACM, Vol. 35.
- Lin, J., Zheng, H., Zhu, Z., Zhang, Z. and David, H. (2007b) 'DRAM-level prefetching for fully-buffered DIMM: design, performance and power saving', in *ISPASS*, IEEE, pp.94–104.
- Lin, J., Zheng, H., Zhu, Z., Gorbatoev, E., David, H. and Zhang, Z. (2008) 'Software thermal management of DRAM memory for multicore systems', in *ACM SIGMETRICS Performance Evaluation Review*, Vol. 36, No. 1, pp.337–348, ACM.
- Liu, J., Jaiyen, B., Veras, R. and Mutlu, O. (2012) 'RAIDR: retention-aware intelligent DRAM refresh', in *ISCA*, June, pp.1–12.
- Liu, S. et al. (2011a) 'Hardware/software techniques for DRAM thermal management', in *HPCA*, pp.515–525.
- Liu, S., Pattabiraman, K., Moscibroda, T. and Zorn, B. (2011b) 'Flicker: saving DRAM refresh-power through critical data partitioning', *ACM SIGPLAN Notices*, Vol. 46, No. 3, pp.213–224.
- Liu, S., Memik, S., Zhang, Y. and Memik, G. (2008) 'A power and temperature aware DRAM architecture', in *DAC*, IEEE, pp.878–883.
- Lyu, C. and Kim, T. (2004) 'Memory access scheduling and binding considering energy minimization in multi-bank memory systems', in *Proceedings of the 41st Annual DAC*, ACM, pp.81–86.
- Malladi, K.T., Nothhaft, F.A., Periyathambi, K., Lee, B.C., Kozyrakis, C. and Horowitz, M. (2012) 'Towards energy-proportional datacenter memory with mobile DRAM', in *ISCA*, June, pp.37–48.
- Mazumdar, S., Tullsen, D. and Song, J. (2010) 'Inter-socket victim cacheing for platform power reduction', in *2010 IEEE International Conference on Computer Design (ICCD)*, IEEE, pp.509–514.
- Meisner, D., Gold, B. and Wensch, T. (2009) 'PowerNap: eliminating server idle power', *ACM Sigplan Notices*, Vol. 44, No. 3, pp.205–216.
- Mukundan, J. and Martinez, J.F. (2012) 'MORSE: multiobjective reconfigurable self-optimizing memory scheduler', *HPCA*, pp.1–12.
- Ohsawa, T., Kai, K. and Murakami, K. (1998) 'Optimizing the DRAM refresh count for merged DRAM/logic LSIs', in *Proceedings of the 1998 International Symposium on Low Power Electronics and Design*, ACM, pp.82–87.
- Ozturk, O. and Kandemir, M. (2008) 'ILP-based energy minimization techniques for banked memories', *ACM Trans. Des. Autom. Electron. Syst.*, July, Vol. 13, No. 3, pp.50:1–50:40.
- Ozturk, O., Chen, G., Kandemir, M. and Karakoy, M. (2006) 'Cache miss clustering for banked memory systems', in *IEEE/ACM International Conference on Computer-Aided Design, ICCAD '06*, ACM, pp.244–250.
- Pandey, V., Jiang, W., Zhou, Y. and Bianchini, R. (2006) 'DMA-aware memory energy management', in *HPCA*, February, pp.133–144.
- Patel, K., Benini, L., Macii, E. and Poncino, M. (2005) 'Energy-efficient value-based selective refresh for embedded DRAMs', *Integrated Circuit and System Design. Power and Timing Modeling, Optimization and Simulation*, Vol. 3728, pp.909–909, DOI: http://dx.doi.org/10.1007/11556930_48.
- Phadke, S. and Narayanasamy, S. (2011) 'MLP aware heterogeneous memory system', in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, IEEE, pp.1–6.
- Pisharath, J., Choudhary, A. and Kandemir, M. (2004) 'Reducing energy consumption of queries in memory-resident database systems', in *Proceedings of the 2004 International Conference on Compilers, Architecture, and Synthesis for Embedded Systems*, ACM, pp.35–45.
- Raju, M. et al. (2009) 'Domain decomposition based high performance parallel computing', *International Journal of Computer Science Issues*, Vol. 5, pp.27–32.
- Rodero, I., Chandra, S., Parashar, M., Muralidhar, R., Seshadri, H. and Poole, S. (2010) 'Investigating the potential of application-centric aggressive power management for HPC workloads', in *International Conference on High Performance Computing (HiPC)*, December, pp.1–10.
- Seongil, O., Choo, S. and Ahn, J.H. (2011) 'Exploring energy-efficient DRAM array organizations', in *54th International Midwest Symposium on Circuits and Systems (MWSCAS)*, August, pp.1–4.
- Stuecheli, J., Kaseridis, D., Hunter, H. and John, L. (2010) 'Elastic refresh: techniques to mitigate refresh penalties in high density memory', in *MICRO-43*, pp.375–384, IEEE, DOI: <http://dx.doi.org/10.1109/MICRO.2010.22>.
- Sudan, K., Chatterjee, N., Nellans, D., Awasthi, M., Balasubramonian, R. and Davis, A. (2010) 'Micro-pages: increasing DRAM efficiency with locality-aware data placement', in *ACM SIGARCH Computer Architecture News*, Vol. 38, No. 1, pp.219–230, ACM.
- Sudan, K., Rajamani, K., Huang, W. and Carter, J. (2012) 'Tiered memory: an iso-power memory architecture to address the memory power wall', *IEEE Transactions on Computers*, in press.
- Tolentino, M.E., Turner, J. and Cameron, K.W. (2009) 'Memory MISER: improving main memory energy efficiency in servers', *IEEE Trans. Comput.*, March, Vol. 58, No. 3, pp.336–350.
- Trajkovic, J., Veidenbaum, A. and Kejariwal, A. (2008) 'Improving SDRAM access energy efficiency for low-power embedded systems', *ACM Transactions on Embedded Computing Systems (TECS)*, Vol. 7, No. 3, p.24.
- Tremaine, R., Franaszek, P., Robinson, J., Schulz, C., Smith, T., Wazlowski, M. and Bland, P. (2001) 'IBM memory expansion technology (MXT)', *IBM Journal of Research and Development*, Vol. 45, No. 2, pp.271–285.
- Udipi, A., Muralimanohar, N., Balasubramonian, R., Davis, A. and Jouppi, N. (2012) 'LOT-ECC: localized and tiered reliability mechanisms for commodity memory systems', in *Proceedings of ISCA*.
- Udipi, A., Muralimanohar, N., Chatterjee, N., Balasubramonian, R., Davis, A. and Jouppi, N. (2010) 'Rethinking DRAM design and organization for energy-constrained multi-cores', in *ACM SIGARCH Computer Architecture News*, Vol. 38, No. 3, pp.175–186, ACM.
- Vogelsang, T. (2010) 'Understanding the energy consumption of dynamic random access memories', in *MICRO-43*, pp.363–374, DOI: <http://dx.doi.org/10.1109/MICRO.2010.42>.

- Wang, Z. and Hu, X. (2005) 'Energy-aware variable partitioning and instruction scheduling for multibank memory architectures', *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, Vol. 10, No. 2, pp.369–388.
- Ware, M., Rajamani, K., Floyd, M., Brock, B., Rubio, J., Rawson, F. and Carter, J. (2010) 'Architecting for power management: the IBM® POWER7 approach', in *HPCA*, IEEE, pp.1–11.
- Yang, L., Dick, R.P., Lekatsas, H. and Chakradhar, S. (2010) 'Online memory compression for embedded systems', *ACM Trans. Embed. Comput. Syst.*, March, Vol. 9, No. 3, pp.27:1–27:30.
- Yoon, D., Jeong, M. and Erez, M. (2011) 'Adaptive granularity memory systems: a trade-off between storage efficiency and throughput', in *ISCA*, ACM, pp.295–306.
- Yoon, D.H., Chang, J., Muralimanohar, N. and Ranganathan, P. (2012) 'BOOM: enabling mobile memory based low-power server DIMMs', in *ISCA*, June, pp.25–36.
- Zhang, G., Wang, H., Chen, X., Huang, S. and Li, P. (2012) 'Heterogeneous multi-channel: fine-grained DRAM control for both system performance and power efficiency', in *Proceedings of the 49th Annual DAC*, ACM, pp.876–881.
- Zheng, H. and Zhu, Z. (2010) 'Power and performance trade-offs in contemporary DRAM system designs for multicore processors', *IEEE Transactions on Computers*, Vol. 59, No. 8, pp.1033–1046.
- Zheng, H., Lin, J., Zhang, Z. and Zhu, Z. (2009) 'Decoupled DIMM: building high-bandwidth memory system using low-speed DRAM devices', in *ISCA*, ACM, pp.255–266.
- Zheng, H., Lin, J., Zhang, Z., Gorbatoev, E., David, H. and Zhu, Z. (2008) 'Mini-rank: adaptive DRAM architecture for improving memory power efficiency', in *MICRO*, pp.210–221, IEEE.
- Zhou, P., Pandey, V., Sundaresan, J., Raghuraman, A., Zhou, Y. and Kumar, S. (2004) 'Dynamic tracking of page miss ratio curve for memory management', in *ACM SIGOPS Operating Systems Review*, Vol. 38, No. 5, pp.177–188, ACM.