== Paper summary ==

The growth of applications and processors has caused stringent requirements in latency, bandwidth, capacity, predictability and energy. On the other hand, DRAM faced a fundamental device scaling challenge, while more than 40% of its time will be spent on refreshing cells in the near future. The memory system needs a fundamental overhaul to overcome these challenges.

This book chapter summarized all known promising research directions and preliminary results in threefold.
1) Overcome DRAM scaling challenge by reducing refreshing impact, or co-designing system and DRAM.
2) Enable emerging resistive memory technology by taking the tradeoff between DRAM, PCM, and STT-MRAM.
3) Design a predictable memory system in performance and QoS.

== Strengths of paper and mechanisms (bullet points) ==

* The challenges stated in this chapter are important. Memory system, known as system performance bottleneck for a long time, indeed faces fundamental scaling challenges, both from manufacturer requirements and technology.

* This chapter summarized the state-of-art research on main memory systematically. It is a good tutorial for anyone who wants to understand or conduct memory related research.

* The outlined future research directions highlight great potentials, such as heterogeneous main memory, system/memory co-design, DRAM translation layer, end-to-end performance guarantees.

== Weaknesses of paper and mechanisms (bullet points) ==

* As an introductory chapter, it’s a bit difficult for readers without DRAM background to understand the problem of scaling, bandwidth, and latency.

* The "challenge one" is about DRAM scaling but the discussion covered different aspect of DRAM challenges (e.g., parallelism, reducing waste, in-DRAM copy). The statement of the challenge may need some modification to make the flow of this section better.

* The future direction was separated across different paragraphs. It will be better to have a summary of future direction for each challenge, which would help the readers refresh the key ideas to go forward.

== Detailed comments (expand on the summary as necessary) ==

The challenge of memory system has been extremely tough during these years. The evolution of applications and processors, both in terms of speed and size, are much faster than memory system. This problem gets deteriorated while DRAM, the most widely used main memory technology, faced fundamental scaling problem.

There have been plenty of researches in different system layers, from system software, processors, caches, memory controllers to memory devices. However it is not easy to get the whole picture by reading individual papers. This book chapter provided a good summary on where we are and the likely directions.

The coverage of this chapter was quite broad. It would be a bit difficult for novice readers without DRAM background or system architecture. For instance, why will refresh take much more time when memory density increases? What is a bank/subarray? While it is impossible to address all the necessary background in a chapter, some brief introduction would be very helpful.

The future research directions provided in this book chapter are encouraging. Specifically, it is intriguing to tackle the challenges from architecture’s point of view. System/memory co-design would make system much more powerful in the years to come.

Overall, it’s a very good summary on the problem and research for future memory

system. It is a starting point for researchers going toward this field.

== Ideas for improvement - Can you do better? ==

* Include some background of DRAM to help readers to understand the topics

* Section 6.2 and 6.3 use different points of view (trend and requirement) to state the problem, but some of them overlap with each other. Another approach can be explaining the problem in one section (bandwidth, latency, energy, QoS, scaling) and introducing the emerging technology (resistive memory) in another section

* Summarize the future direction in a separate section to let readers catch the main ideas more easily.

== Lessons learned ==

By reading this chapter, I got better understanding of the whole picture of memory research. During the years in industry, we did feel the memory system is a bottleneck and there must be a scaling limit at some point. With this reading, all these ideas are crystallized.

While there are already some encouraging results, there is still a long way to go. For example, system/memory co-design is still at a very early phase. How to design a memory to fit the system or make memory smarter are big questions. How to get most out of resistive memory is another big topic. There are also quite a few interesting topics on memory scheduling, parallelism, and scaling.