

Potential future research in computing: Heterogeneous systems, memory subsystems - Process-in-Storage, or not to Process-in-Storage? That is the question

Uri Weiser

Technion - Israel Institute of Technology, Israel

Abstract. The era of Heterogeneous systems and Big Data computing is already here. Handling huge amount of data poses new challenges in data management and in the effective usage of memory, caches, heterogeneous structures and available bandwidth. In addition, computing requirements of Big Data is unique; in many occasions the processing required per storage access is limited (i.e. low Instructions/Byte) which presents a new challenge and opportunities for computer architects. The increasing performance requirement are driving the industry towards parallel execution of threads (or tasks) being executed by energy efficient computing (e.g. heterogeneous) engines.

In this talk we will present some of the research being performed at the Technion, which is related to effective usage of caches and heterogeneous systems in multi-tasking environment and will provide a glimpse of remedy to the situation. In addition we will present the substantial issue of non-temporal-memory-access in Big Data environment and suggest some hopefully stimulating solutions.