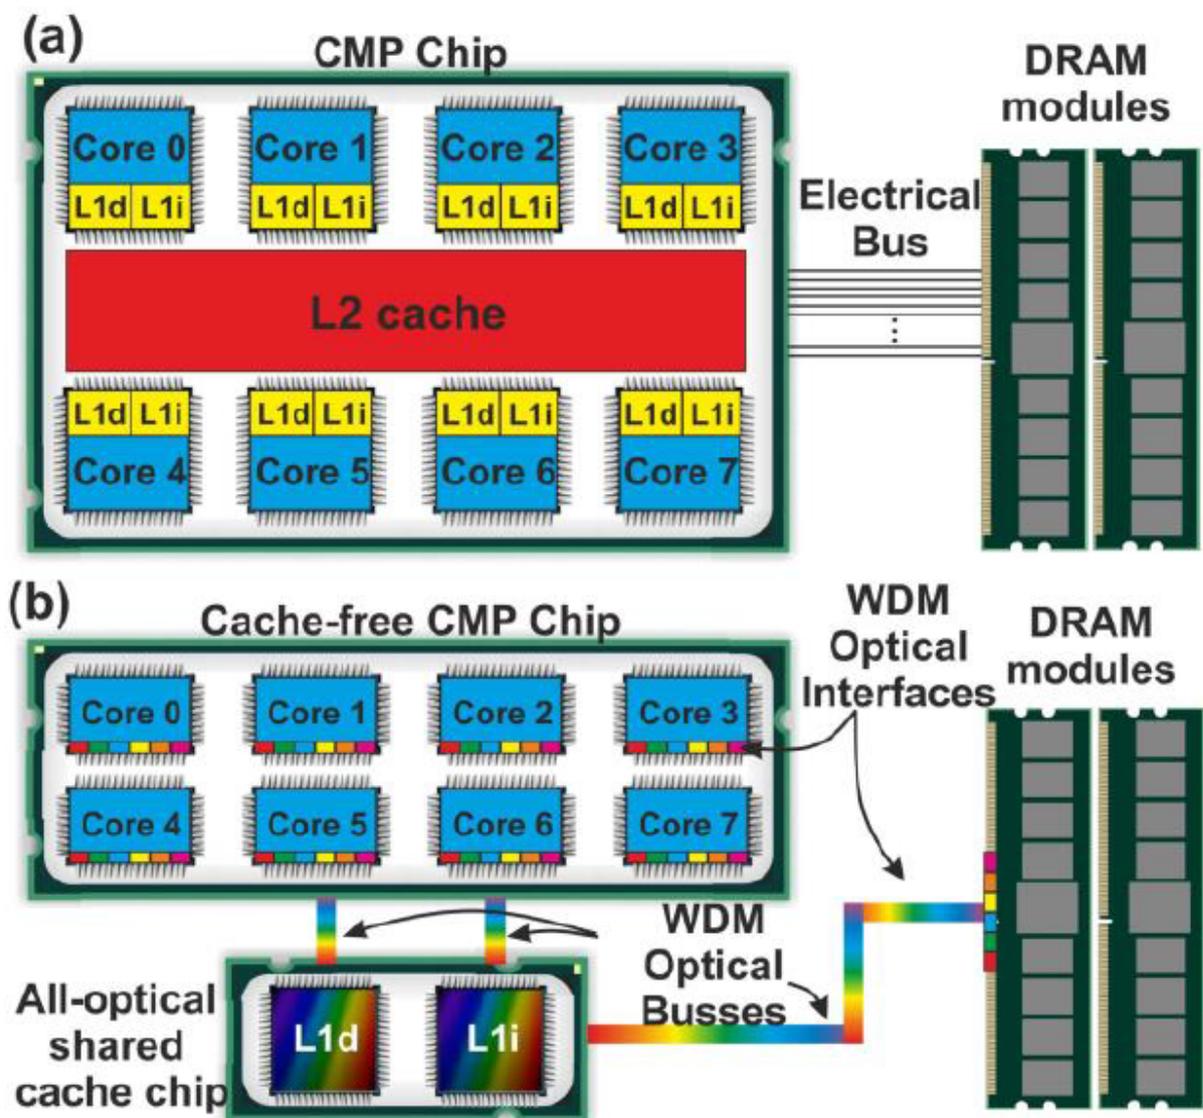


Optics in Computing: from Photonic Network-on-Chip to Chip-to-Chip Interconnects and Disintegrated Architectures

Following a decade of radical advances in the areas of integrated photonics and computing architectures, we discuss the use of optics in the current computing landscape attempting to re-define and refine their role based on the progress in both research fields. We present the current set of critical challenges faced by the computing industry and provide a thorough review of photonic Network-on-Chip (pNoC) architectures and experimental demonstrations, concluding to the main obstacles that still impede the materialization of these concepts. We propose the employment of optics in chip-to-chip (C2C) computing architectures rather than on-chip layouts towards reaping their benefits while avoiding technology limitations on the way to manycore set-ups. We identify multsocket boards as the most prominent application area and present recent advances in optically enabled multsocket boards, revealing successful 40Gb/s transceiver and routing capabilities via integrated photonics. These results indicate the potential to bring energy consumption down by more than 60% compared to current QuickPath Interconnect (QPI) protocol, while turning multsocket architectures into a single-hop low-latency setup for even more than 4 interconnected sockets, which form currently the electronic baseline. We go one step further and demonstrate how optically-enabled 8-socket boards can be combined via a 256×256 Hipoλaos Optical Packet Switch into a powerful 256-node disaggregated system with less than 335nsec latency,

forming a highly promising solution for the latency-critical rack-scale memory disaggregation era. Finally, we discuss the perspective for disintegrated computing via optical technologies as a means to increase the number of synergized high-performance cores overcoming die area constraints, introducing also the concept of cache disintegration via the use of future off-die ultra-fast optical cache memory chiplets.



(a) Conventional CMP architecture with on-chip Cache Memories and Electrical Bus for CPU-MM communication (b) The proposed CMP architecture with off-chip optical

Assuming, for example, an optical CMP-to-cache bus speed and

optical cache operational speed of 16GHz, as has been modelled in [164], with a reasonable processing core clock speed of 2GHz, the cache access system performs 8x faster than the processing cores. This indicates that the optical cache can serve all 8 processing cores within a single 2GHz cycle. Regarding latency, every core has 8 cache clock cycles available to complete its request within a single core clock cycle, including of course optoelectronic conversion at the CMP interface, propagation in the optical bus and cache accessing. Assuming a bus length of 1cm, which can be considered as a reasonable value within a macrochip System-in-Package, the time-of-flight is just 50psec for a waveguide-based bus refractive index of 1.5. With optoelectronic conversion taking place at the bus clock speed and at the Memory Address and Memory Buffer Register (MAR and MBR, respectively) interfaces, ultra-fast cache access latency can be obviously easily retained. For detailed timing diagrams that present the optical cache circuitry operation at various stages for both Read and Write operations and the TDM-based access scheme followed in the proposed system of Fig. (b).

This has been extensively analyzed, where also the performance of the system depicted in Fig. was thoroughly investigated via detailed simulations using the gem5 simulation engine and the PARSEC benchmark suite. The main findings when comparing the system of Fig.10(a) with the system of Fig.(b) for the same amount of total cache capacity can be summarized as follows:

- The use of a shared L1 cache yields an important reduction in the cache miss rate of more than 75%, especially when executing parallel programs with high data sharing and exchange needs among their threads; the high volumes of data exchange increase the traffic and consequently the miss rate among the dedicated L1d caches in typical architectures with dedicated L1 caching.
- The shared L1 cache negates the need for cache coherency updates and cache coherency protocols, simplifying the program execution and contributing significantly in cache miss ratio reduction by cancelling all cache coherency misses.

- Cache miss ratio reduction and concurrent multiple core service translate to important execution time speed-up factors that were shown to range between 10% and 20% for computational settings that employed cache capacities equal to the Sparc T5 processor and IBM's Power7 processor, respectively.

Extending this concept into a macrochip layout with multiple core and optical cache chiplets can bring additional benefits, since caching will be rather utilized as a pool of resources that will facilitate time and energy savings. Moreover, it can transform computing from a rigid into a versatile and flexible environment, where caching and processing resources can be exploited on demand depending on the workload requests, allowing eventually also for cache and processing power upgrades similar to the way that DRAM upgrades are currently being performed.

For more information: DOI [10.1109/JLT.2018.2875995](https://doi.org/10.1109/JLT.2018.2875995)