

# Accelerator Level Parallelism

Mark D. Hill  
Computer Sciences Department  
University of Wisconsin-Madison

Vijay Janapa Reddi  
School of Engineering and Applied Sciences  
Harvard University

VIEWPOINT August 14, 2020

*Future applications demand more performance, but technology advances have been faltering. A promising approach to further improve computer system performance under energy constraints is to employ hardware **accelerators**. Already today, mobile systems concurrently employ multiple accelerators in what we call **accelerator-level parallelism (ALP)**. To spread the benefits of ALP more broadly, we charge computer scientists to develop the science needed to best achieve the performance and cost goals of ALP hardware and software.*

## Introduction

While past information technology (IT) advances have transformed society, future advances hold great additional promise. For example, we have only just begun to reap the changes from artificial intelligence--especially machine learning--with profound advances expected in medicine, science, education, commerce, and government. All too often forgotten, underlying the IT impact are the dramatic improvements in the programmable hardware. Hardware improvements deliver performance that unlocks new capabilities. However, unlike in the 1990s and early 2000s, tomorrow's performance aspirations must be achieved with much less technological advancement (Moore's Law and Dennard Scaling). How then does one deliver AR/VR, self-driving vehicles, and health wearables at costs that enable great customer value?

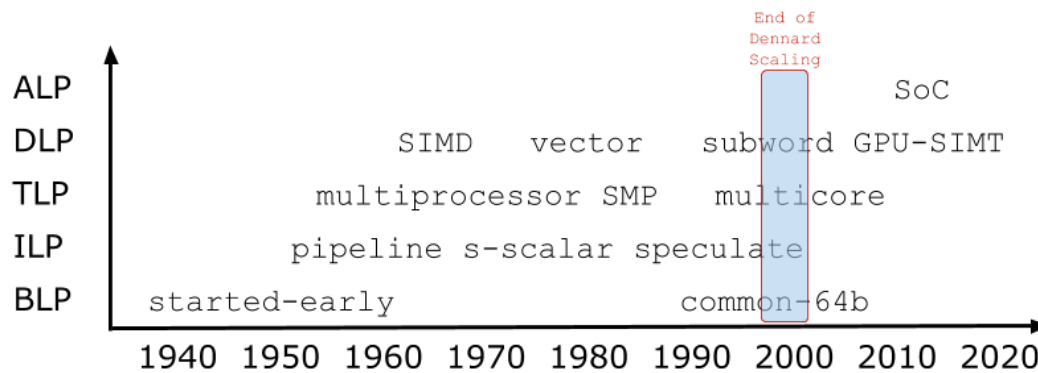
One approach that has emerged is to use **accelerators**: hardware components that execute a targeted computation class faster and usually with much less energy. An accelerator's flexibility can vary from high (GP-GPU) to low (fixed-function block). Recent work tends to focus on targeting specific application domains, such as graphics (before GPUs generalized), deep machine learning, physics simulations, and genomics. Moreover, most work on accelerators, including in CACM [J+18,NGS19,DTH20], has focused on CPUs using a single accelerator, with one early forecast of multiple accelerator use [BC11].

In our view, many future computing systems obtain greater efficiency by employing **multiple accelerators** where each accelerator efficiently targets an aspect of the ongoing computation, much as a Swiss Army knife has specific tools for specific tasks. Smartphones foreshadow this future by employing many accelerators concurrently, but unlike a Swiss Army knife these accelerators often operate in parallel using separately-developed software stacks.

We assert that there is as yet no "science" for debating and systematically answering basic questions for how to best facilitate broad, flexible and effective use of multiple accelerators. Herein we expose this opportunity (the **what**), but charge our readers with determining **how** best to address it. We first review past computer system improvements exploiting levels of parallelism, introduce **Accelerator-Level Parallelism (ALP)** as a way to frame new challenges, and examine the "point" success of smartphone ALP.

## Past, Present and Future Parallelism

As technology scaling provided more and smaller transistors, computer processor architects transformed the transistor bounty into faster processing by using the transistors in parallel. Effectively using repeated transistor doubling required **new levels of transistor parallelism**. Figure 1 looks at the past and present, and depicts the different levels of parallelism (y-axis) that have emerged as computing evolved over the decades (x-axis).



**Figure 1: A snapshot of parallelism over the years, showing how the various forms of parallelism were exploited through different types of architectural mechanisms.**


In Figure 1, Bit-level parallelism (BLP) refers to performing basic operations (arithmetic, etc.) in parallel. It was common in early computers and was later enhanced with larger word sizes in commodity systems. Instruction-level parallelism (ILP) is the execution of logically sequential instructions concurrently with pipelining, superscalar, and increasing speculation. Thread-level parallelism (TLP) is the use of multiple processor cores, which initially started with discrete processors and were later integrated as on-chip cores. Data-level parallelism (DLP) pertains to performing similar operations on multiple data operands via arrays and pipelines before broad success via general-purpose graphics processing units (GP-GPUs).

In this viewpoint and in Figure 1, we assert that another major parallelism level is emerging: **Accelerator-Level Parallelism (ALP)**. We define ALP as *the parallelism of workload components concurrently executing on multiple accelerators*. A goal of ALP is to unlock many accelerators at the same time in a manner analogous to how ILP concurrently employs multiple functional units. ALP does not replace other parallelism levels but builds upon them, as most accelerators internally employ one or more of BLP, ILP, TLP, and DLP. Moreover, much like ILP that has been exploited at different levels of the stack, ranging from superscalar and out-of-order execution at the microarchitecture level up to instruction scheduling at the compiler level, ALP opens up many degrees of freedom for novel hardware and software design and optimization. It also opens up possibilities for new runtime resource management, which is analogous to heterogeneous scheduling across CPUs and GPUs, but with the added complexity of scheduling tasks in real time across a sea of hardware accelerators.

ALP is emerging today, especially in Systems on a Chip (SoCs) for mobile, edge and cloud computing that concurrently employ multiple accelerators. We next present a case study of ALP in mobile SoCs to understand how ALP is currently used, albeit in somewhat of a limited form, and then lay a foundation for future work that can exploit ALP more generally.


## Mobile SoCs as Harbingers of Multiple Accelerators Using ALP

Driven by the need for extreme energy efficiency, mobile SoCs are the very early adopters of ALP. Figure 2 shows SoCs from four major vendors: Apple, Qualcomm, Samsung, and Huawei. Across all vendors, in each SoC, much less than 50% of the die is dedicated to the CPUs (white boxes). The majority of the area is dedicated to specialized accelerators, such as a Digital Signal Processor, Image Signal Processor, GPU, Neural Processing Unit, and Video Encoder/Decoder, as well as I/O interfaces for audio, networking, video, etc.



**Figure 2: Apple, Qualcomm, Samsung, & Huawei SoCs (top-left & clockwise) [G18]**

It is common in smartphone SoCs for workloads to **exhibit ALP with multiple accelerators in concurrent--not exclusive--use**. Figure 3 shows 4K, 60 frame-per-second video capture usecase with two paths. One path goes to the display, rendering real-time content to the end-user, and the other path goes to flash storage to save the content for offline viewing. In this example, data traverses accelerators with both parallelism (two paths) and pipelining, all choreographed by CPUs (not shown). In other usecases like an interactive multi-party video conferencing application, data flow and CPU choreographing can be even more dynamic and complex. Nevertheless, we expect accelerators to increasingly handle “data plane” computation while CPUs retain the “control plane” tasks. Doing so will enable richer computation from a fixed power budget, valuable from smartphones to cars to the cloud.



**Figure 3: ALP in action in a 4K video capture usecase on a smartphone [RYK18]**

Mobile SoCs are clearly relying on ALP for low-power and efficient execution. However, they are not yet exploiting the full potential of ALP, which we see as needed for recouping the flexibility that the CPU delivered for decades. For instance, in the above example, the dataflow and the binding between the application tasks and accelerators is fixed. The ISP cannot be programmatically repurposed for tasks aside from processing camera image inputs. To this end, we believe we need better science and engineering toward ALP utilization.

### **Toward a Science for Multiple-Accelerator Systems using ALP**

John Hennessy and David Patterson asserted in their 2018 Turing Award Lecture that we are upon a *new golden age for computer architecture* [HP19].

We assert that the challenge put forth by Hennessy and Patterson ought to be generalized to a *new golden age for computer science and engineering* and that employing multiple accelerators with ALP is an opportunity that opens up new vistas for research as accelerators are integrated into complex SoCs. We do not know all of the possibilities, but we discuss some ideas here to seed research directions.

A key challenge is developing abstractions and implementations to enable programmers to target the whole SoC and implementers to holistically design its software and hardware. We take inspiration from the Single Instruction Multiple Thread (SIMT) model that effectively abstracts GPU hardware’s cornucopia of parallelism and scheduling mechanisms. SIMT both enabled GPUs to expand from graphics workloads to general-purpose DLP use and enabled software-hardware implementation improvements beneath the abstraction.

As ALP emerges, we expect new paradigms must be invented to flexibly and effectively exploit its potential. This is not the case today. In contrast to a SIMT-like holistic view, today's SoCs only exploit ALP in limited niches with each accelerator acting as a "silo" with its own programming model, and often its own (domain-specific) language, runtime, software development kit (SDK), and driver interface. While employing multiple accelerators with no abstraction can work in restricted situations (e.g., for 10-20 phone usecases), it is unlikely to make ALP generally useful. How can we transcend per-accelerator software silos of different languages, SDKs, etc? What are abstractions and mechanisms for scheduling/sequencing accelerators or partitioning/virtualizing them (perhaps stream data flow)? What belongs in runtimes versus above/below the OS hardware abstraction layer?

Even more than previously parallel levels, ALP exploitation will likely require software-hardware co-design due to the heterogeneous nature of accelerators and ALP. Moreover, this is also likely to incentivize computer-aided design tool chain innovations to facilitate the rapid exploration of heterogeneous design spaces. ALP implementations should aspire toward globally optimal software-hardware systems, whereas much good work today focuses on making each accelerator "locally" optimal. While good accelerators are essential, a collection of locally-optimal accelerators is unlikely to be globally optimal. For this reason, we need better methods for holistically designing SoCs from accelerator, memory, and interconnect components, more like how processor cores are crafted from ALUs, register files, and buses. Analysis in both cases centers on parallel operation: ALP for SoCs and ILP for cores.

In more detail, there are many ALP questions that need better answers and better methods for systematically determining answers. For instance, from a compute perspective, we lack the fundamental science on how we must select, size, make efficient, and sometimes combine similar accelerators? Similarly, from a memory perspective, when should on-chip memory be private to accelerators or shared? When should this memory be a software-visible scratchpad or software-transparent cache? From an integration perspective, how do we best communicate data (shared memory or queues) and control (polling, interrupts, other) among accelerators? From an operational perspective, once an SoC is deployed, can we schedule heterogeneous parallel resources with (non-convex) optimization or must heuristics suffice? In sum, a more systematic approach is needed to design many accelerators as blocks to create holistic ALP systems that excel at performance and cost goals.

**Conclusion.** This viewpoint has argued that employing multiple accelerators with ALP has much promise for enhancing future computing efficiency, that we don't yet know how to do it well beyond niches, and that we can work together to make this happen. We have identified **what** the opportunity is, but leave to our readers **how** best to solve it.

## References

- [BC11] S. Borkar and A. A. Chien, "The Future of Microprocessors," CACM, May 2011, pp. 67-77, 10.1145/1941487.1941507.
- [DTH20] W. J. Dally, Y. Turakhia, and S. Han, "Domain-Specific Hardware Accelerators," CACM, July 2020, pp. 48-57, 10.1145/3361682.
- [G18] GSMARENA, "Comparing the sizes of flagship chipsets: Qualcomm, Samsung, Huawei and Apple," [https://www.gsmarena.com/comparing\\_the\\_sizes\\_of\\_flagship\\_chipsets\\_qualcomm\\_samsung\\_huawei\\_adn\\_apple-news-30240.php](https://www.gsmarena.com/comparing_the_sizes_of_flagship_chipsets_qualcomm_samsung_huawei_adn_apple-news-30240.php), March 2018.

- [J+18] N. P. Jouppi, C. Young, N. Patil, and D. Patterson, "A Domain-Specific Architecture for Deep Neural Networks," CACM, September 2018, pp. 50-59, 10.1145/3154484
- [HP19] J. L. Hennessy and D. A. Patterson, "A New Golden Age for Computer Architecture," CACM, February 2019, pp. 48-60, 10.1145/3282307.
- [NGS19] T. Nowatzki, V. Gangadhar, and K. Sankaralingam, "Heterogeneous Von Neumann/Dataflow Microprocessors," CACM, June 2019, pp. 83-91, 10.1145/3323923
- [RYK18] V. J. Reddi, H. Yoon, and A. Knies, "Two Billion Devices and Counting," IEEE Micro, January-February 2018, pp. 6–21.