

能数据(主要是 IPC)作为调度的依据,将线程调度到 IPC 高的核上执行.该文通过工作集签名(WSS)表示线程某个执行阶段,WSS 表示线程在一定指令窗口内提交指令的集合表示.线程不同执行阶段的区分通过两个 WSS 的距离(即 WSS 中不同位的个数)计算得到,为了去噪,定义执行阶段变化阈值,如果两个 WSS 之间不同位数除以 WSS 的总位数高于 50%的话,将两个阶段标识为不同,阶段发生变化并且是新出现的阶段,将对线程在各个核上的 IPC 进行统计,并根据 IPC 进行调度.该文设计实现工作签名历史表保存每个阶段的 WSS 及此阶段的 IPC 数据.每个执行阶段的 WSS 计算是通过指令执行地址计数器(PC)的哈希计算映射到 512 个字节的向量表中,向量表中的 1 位表示指令缓存的 64 个字节的大小.当映射到向量表里对应的位的指令已提交,那么相应的位置 1,没有执行的位置 0,向量表各位值的序列就是 WSS 的值.文献[90]提出,计算 WSS 的指令窗口如果选择过小,将会导致线程执行阶段的划分比较多,从而会导致过大的迁移开销,因此,通过实验选取 50K 作为执行阶段识别的指令窗口.

文献[91]提出了以降低功耗为目标的感知程序执行阶段的调度算法,证明当程序行为和特性发生变化的时候进行线程的动态分配可以有效降低功耗.采用与文献[83]相同的程序执行阶段探测算法,采用的度量标准是系统的 EDP(energy-delay product),调度算法将分配线程到运行此线程 EDP 比较小的核上执行.文献[91]假设程序相同执行阶段具有相似的 EDP,当执行阶段发生变化的时候,进行重新调度.线程新的阶段出现的时候,将在每种类型的核上对 EDP 进行采样,根据采样结果进行调度.下次出现相同执行阶段时,根据之前的结果进行调度.

3.4 算法评估

对异构调度算法有效地进行评估需要实验平台更加真实地对异构多核处理器环境进行建模.由于通过 DVFS 调整核的电压和时钟频率无法真实地仿真异构多核处理器,所以,越来越多的研究工作采用真实的异构硬件、模拟器或者通过评估模型来进行算法的效果评估.

3.4.1 真实硬件

文献[66,92]采用(non-uniform memory access,简称 NUMA)服务器作为实验平台,文献[66]提出可以感知 NUMA 架构的虚拟机调度算法,以降低 NUMA 系统的内存局部性问题所带来的影响.文献[92]提出感知 NUMA 架构的线程迁移策略,通过在线监控线程的常驻工作集(WSS)预测线程迁移的代价,如果线程从核 A 迁移核 B,并且 A 和 B 在 NUMA 不同的节点,线程在核 A 的 WSS 最大并且超过了核 B 的 LLC 大小,则预测线程迁移的代价非常大,由于导致很高的 LLC Misses 而不进行迁移,根据此策略进行线程迁移的优化从而提高系统性能.文献[93]采用 ARM big.LITTLE 架构的开发板作为实验平台,主要包括 2 个 Cortex A15 和 3 个 Cortex A7,该文主要提出了基于价格理论的调度优化技术,目的是在满足系统性能的前提下降低功耗,而 ARM big.LITTLE 架构可以表现良好功耗性能的异构特性.

3.4.2 模拟器

现在进入市场的异构产品,比如 NUMA 和 ARM big.LITTLE,针对特定架构和特性,配置相对固定,无法通过灵活的参数配置和定制进行算法的测试.与之相比,采用架构模拟器支持更灵活的配置,可以全面地对算法进行评估.文献[46]使用 sniper 进行异构的配置,大核配置为 4 发射乱序处理核心,小核配置为 4 发射顺序处理核心.文献[71]采用 Simscalar 模拟异构多核处理器,每个核都采用乱序超标量技术,通过在发射带宽、L1D 缓存的大小、分支预测器的大小采用不同的配置来进行异构的模拟.文献[85]采用 Intel Quick IA 作为评估的异构平台,Quick IA 是 Intel 的 FPGA 的 SoC 原型验证平台,该文在平台上对低功耗的 Atom Core 和高性能的 Xeon Core 的异构处理器构建进行仿真.通过模拟器进行不同核的微架构配置相对方便,而且核之间的差异度也更加贴近真实的异构系统,但是模拟器模拟的正确性与真实硬件毕竟存在偏差,因此,为了调度算法评估结果适用于真实的硬件平台,首先需要对模拟器的正确性进行验证和校准.

3.4.3 评估模型

然而,由于模拟器本身速度受限,无法支持大量混合工作负载(比如:数百个并行程序)在合理时间内执行完成.有些工作建立了分析模型以进行异构处理环境下的程序性能(或功耗)的评估,分析模型不需要像模拟器那

样对异构多核处理器的微架构进行详细模拟,通过模型和算法对复杂工作负载场景下的程序性能进行评估,相比于模拟器来说,速度更快,可扩展性更好,但是分析模型的正确性在算法评估时变得尤为重要。

文献[94]提出分析模型 MPPM(multi-program performance model)来评估并发多程序的性能.由于多个程序并发执行时,多核之间的资源竞争会影响程序的进度,反过来,每个程序执行的进度也会影响资源竞争的程度.MPPM 的输入是每个程序独立地在每个核运行的性能数据,包括总的 CPI、访存 CPI 和访问 Cache 地址的记录,因此需要提前对这些性能数据进行测试和统计.模型算法的初始状态是假设所有的程序都以单核执行的状态开始,通过对 Cache 访问竞争和访存 CPI 的监控分析,评估多核间资源竞争对于程序 CPI 的影响,也就是说,由于资源竞争导致的 CPI 下降,算法将调整 CPI 为考虑资源竞争影响的 CPI(文献[94]中称其为多核 CPI),然后算法不断迭代直到退出.通过 MPPM 可以对并发多程序工作负载下每个程序的 CPI 进行评估,从而根据 CPU 计算系统的吞吐量(STP)和每个程序的平均执行时间(ANTT).通过与时钟精确的架构模拟器的结果进行对比,在 2 核、4 核和 8 核的情况下,MPPM 模型估计的 STP 平均误差是 1.4%、1.6%和 1.7%,ANTT 的平均误差为 1.5%、1.9%和 2.1%.

MPPM 同时支持同构和异构多核处理器的性能,在异构多核处理器环境下,算法将首先对所有程序在所有核上独立执行的性能数据进行收集分析,然后随机选择测试的 benchmark 程序分配到 N 个 Core 上,通过 MPPM 进行异构多核处理器环境下的程序 CPI 的估计.由于性能评估的结果与测试程序调度的算法密切相关,模型也可被用作调度算法效果的评估.

4 总结与展望

伴随计算机系统的日趋复杂和应用需求的多样化,异构多核处理系统逐渐成为主流.调度技术作为充分管理和利用异构多核处理器的主要手段变得尤为重要^[95],然而传统操作系统中的调度技术主要针对同构多核处理器结构设计,没有充分考虑程序的行为特性、资源需求和 CPU 微架构的差异性,无法对硬件架构和工作负载的差异性进行有效感知,影响了调度决策的准确性.本文针对性能非对称性异构多核处理环境下的调度优化技术的挑战及已有研究工作进行了系统的总结.

异构调度的目标是根据负载特性和核之间微架构的差异,将程序分配到最适合的核上运行.最主要的问题是如何感知微架构和程序特性的差异,并根据优化目标采用相对最优的任务分配和迁移策略.因此,本文从优化目标、分析模型、调度决策和算法评估这 4 个方面对异构调度技术已有工作进行了详述.由于异构环境中比如大核和小核的流水线设计差异、缓存架构差异等为各种优化目标的定义和实现提供了更加复杂的场景,优化目标需要在定义和满足优化目标的度量标准时,感知程序在不同微架构运行的差异.本文分别从满足性能、能效、公平性、并发程序瓶颈优化及特定应用领域优化等目标来描述异构调度的工作,比如满足性能主要通过加速比、IPC 比率、关键代码大核执行等的衡量标准将程序分配到更加受益的核上.分析模型主要在程序资源需求分析的时候建立程序特性与不同微架构之间的关联,作为调度的主要依据.本文主要总结了架构无关分析、CPI 模型、经验模型方法,由于分析方法各有利弊,在近期的工作中综合这几种分析方法提出了混合分析模型.由于程序资源需求伴随执行阶段而发生变化,为了实现细粒度的异构调度,调度决策主要通过对程序阶段变化的感知来进行任务迁移的决策,主要的方法包括训练和预测模型,预测决策不需要对每个执行阶段进行训练,效率比训练决策会高,但是预测模型的正确性会直接影响调度的效果.最后,总结了对调度算法进行评估的方法,在使用真实的异构硬件和模拟器的基础上,描述了评估模型相关的工作,在异构硬件环境有限的前提下,分析模型不需要像模拟器那样对异构多核处理器的微架构进行详细模拟,通过模型和算法对复杂工作负载场景下的程序性能进行评估,相比于模拟器来说,速度更快,可扩展性更好,但是评估模型的准确性就变得尤为重要.

然而,当前异构计算机处理系统变得越来越复杂,核的数量和种类不断增加,同时,更多关注微架构设计的持续优化与新软件技术的深度融合,出现了 AI 芯片和深度学习处理器并开始应用.以内存为例,在关于摩尔定律未来的讨论中,有专家表明,传统的 DRAM 访问延时如果从 200ns 降低为 150ns,整体的数据中心的功耗将下降 15%.片上内存的架构设计成为一个有潜力的改进方向,文献[96]证明,片上内存有大于 10 倍的延时缩短和功

耗降低.2016年,Google公司公布了张量处理器——TPU(tensor processing unit),将深度学习等新型软件技术应用在芯片设计中,将处理复杂应用的难度从软件系统降低到硬件架构的层面来实现,在未来还可能会将频繁使用的软件处理算法,比如大数据的排序、查找等直接固化在片上内存中.同时,3D集成电路^[97,98](3DIC)等新技术在芯片设计中的应用,将额外或者共享的资源(比如重排序队列、缓存系统等)拓展到第三维度,达到资源的细粒度共享和相对较低的延迟.这些复杂的异构环境为调度优化目标的建立和满足提出了更为复杂的环境和需求,需要考虑更多的因素.异构感知的调度技术与硬件及微架构的深度融合还有很多工作可以探索.

(1) 伴随异构处理器的应用尤其是在移动终端市场的普及,在满足性能和其他服务质量(QoS)指标的前提下,通过调度来降低功耗是被关注的工作,处理器从硬件层面提供 DVFS(dynamic voltage and frequency)、DPM(dynamic power management)等电源管理技术,软件层面的调度技术如何根据用户 QoS 需求与硬件提供的电源管理技术深度融合以降低功耗是值得研究的方向.同时,目前的硬件暂没有提供度量程序功耗的接口,软件一般也没有提供程序行为变化的接口,因此,在满足异构调度优化目标尤其是功耗目标的工作中,软硬件深度融合的协同工作将是未来主要的方向.

(2) 分析模型是异构感知调度重要的工作,伴随着异构环境的日趋复杂,指令和硬件类型的种类越来越多,差异度也越来越大,比如存在多种指令集(ARM、X86),存在普通内存(DRAM)和非易失性内存(non-volatile memory,简称 NVM)等多种内存,而且由于程序执行过程中竞争或者执行异常和错误会造成动态的异构变化,这就给分析模型在不同微架构配置核上的性能评估带来了困难.如果处理器硬件能够提供合理的辅助程序分析模块,就可以在系统层面提供相应的接口供调度做出更为准确的决策.比如文献[99]设计硬件模块监控线程的行为,并进行其他核上的性能预测.硬件辅助分析的方式虽然提高了调度决策的准确性,速度也会更快,但是增加了硬件设计和实现的开销,需要软硬件设计更好的协同,最好可以在操作系统层面提供接口,为调度提供直接的支持.

(3) 在调度决策方面,现在主要还是在当前指令窗口中采用边训练边决策的方式进行调度,在复杂的异构环境下,如何能够快速、准确地预测程序执行阶段的变化,并快速进行决策值得继续探索.同时,任务迁移也需要程序设计、编译器和系统架构给出优化的思路,如果两个核之间具有不同的指令集架构就又增加了实现的难度,目前在这方面还需要做更多的研究.

(4) 近年出现软件技术如近似计算(approximate computing)、近阈值计算(near-threshold computing)、内存系统的数据压缩技术也将推进更多样化的多核异构系统进入市场,这些计算对于能效提出了更高的要求,例如:百万兆的计算目标在 20MW 的功耗下 1s 完成 10^{18} 次计算^[100],通过异构调度来满足优化目标还有很多工作需要探索.

总之,随着硬件芯片技术和软硬件融合技术的发展,为了更好地适配和管理异构硬件,不仅仅是调度算法,与之相关的操作系统及相关软件都有很多工作要做.

References:

- [1] Yeric G. Moore's law at 50: Are we planning for retirement. In: Proc. of the IEEE Int'l Electron Devices Meeting. 2015. 1.1.1-1.1.8.
- [2] Venkat A, Tullsen DM. Harnessing ISA diversity: Design of a heterogeneous-ISA chip multiprocessor. ACM SIGARCH Computer Architecture News, 2014, 121-132.
- [3] Pricopi M, Mitra T. Bahurupi: A polymorphic heterogeneous multi-core architecture. ACM Trans. on Architecture and Code Optimization, 2012,8(4):1-21.
- [4] Greenhalgh P. Big.LITTLE processing with ARM Cortex™-A15 & Cortex-A7. In: Proc. of the ARM. 2011. 1-8.
- [5] Becchi M, Crowley P. Dynamic thread assignment on heterogeneous multiprocessor architectures. In: Proc. of the 3rd Conf. on Computing Frontiers. 2006. 29.
- [6] Lugini L, Petrucci V, Mosse D. Online thread assignment for heterogeneous multicore systems. In: Proc. of the 41st Int'l Conf. on Parallel Processing Workshops. 2012. 538-544.

- [7] Rehman M, Asfand-E-Yar M. Scheduling on heterogeneous multi-core processors using stable matching algorithm. *Int'l Journal of Advanced Computer Science and Applications*, 2016,7(6).
- [8] Nie P, Duan Z. Efficient and scalable scheduling for performance heterogeneous multicore systems. *Journal of Parallel and Distributed Computing*, 2012,72(3):353–361.
- [9] Saez JC, Prieto M, Fedorova A, Blagodurov S. *A Comprehensive Scheduler for Asymmetric Multicore Systems*. New York: Association for Computing Machinery, 2010. 139–152.
- [10] Luo YC, Packirisamy V, Hsu W C, Zhai A. Energy efficient speculative threads: Dynamic thread allocation in same-ISA heterogeneous multicore systems. In: *Proc. of the 19th Int'l Conf. on Parallel Architectures and Compilation Techniques*. 2010. 453–464.
- [11] Liu GS, Park J, Marculescu D. Dynamic thread mapping for high-performance, power-efficient heterogeneous many-core systems. In: *Proc. of the 31st IEEE Int'l Conf. on Computer Design (ICCD)*. 2013. 54–61.
- [12] Navada S, Choudhary NK, Wadhavkar SV, Rotenberg E. A unified view of non-monotonic core selection and application steering in heterogeneous chip multiprocessors. In: *Proc. of the 19th Int'l Conf. on Parallel Architectures and Compilation Techniques*. 2013. 133–144.
- [13] Eyerman S, Eeckhout L. The benefit of SMT in the multi-core era: Flexibility towards degrees of thread-level parallelism. In: *Proc. of the 19th Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*. 2014. 591–606.
- [14] Kim J, Lee J, Jeong J. Exploiting asymmetric CPU performance for fast startup of subsystem in mobile smart devices. *IEEE Trans. on Consumer Electronics*, 2015,61(1):103–111.
- [15] Lin FX, Wang Z, Zhong L. K2: A mobile operating system for heterogeneous coherence domains. In: *Proc. of the 19th Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*. 2014. 285–300.
- [16] Lustig D, Trippel C, Pellauer M, Martonosi M. ArMOR: Defending against memory consistency model mismatches in heterogeneous architectures. In: *Proc. of the IEEE Int'l Symp. on Computer Architecture*. 2015. 388–400.
- [17] Rodrigues R, Koren I, Kundu S. Performance and power benefits of sharing execution units between a high performance core and a low power core. In: *Proc. of the 27th Int'l Conf. on VLSI Design*. 2014. 204–209.
- [18] Deng Y, Cheng XH. A heterogeneous multiprocessor task scheduling algorithm based on SFLA. In: *World Automation Congress*. 2016. 1–5.
- [19] Zhou N, Qi D, Wang X, Zheng Z, Lin W-W. A list scheduling algorithm for heterogeneous systems based on a critical node cost table and pessimistic cost table. *Concurrency and Computation: Practice and Experience*, 2017,29(5):e3944.
- [20] Petrucci V, Loques O, Mossé D, *et al.* Energy-efficient thread assignment optimization for heterogeneous multicore systems. *ACM Trans. on Embedded Computing Systems*, 2015,14(1):1–26.
- [21] Petrucci V, Loques O, Mossé D. Lucky scheduling for energy-efficient heterogeneous multi-core systems. In: *Proc. of the 2012 USENIX Conf. on Power-Aware Computing and Systems*. 2012. 7.
- [22] Muthukaruppan TS, Pathania A, Mitra T. Price theory based power management for heterogeneous multi-cores. In: *Proc. of the 19th Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*. 2014. 161–176.
- [23] Winter JA, Albonesi DH. Scheduling algorithms for unpredictably heterogeneous CMP architectures. In: *Proc. of the Int'l Conf. on Dependable Systems & Networks*. 2008. 42–51.
- [24] Gutierrez A, Dreslinski RG, Mudge T. Evaluating private vs. shared last-level caches for energy efficiency in asymmetric multi-cores. In: *Proc. of the Int'l Conf. on Embedded Computer Systems: Architecture, Modeling, and Simulation*. 2014. 191–198.
- [25] Srinivasan S, Kurella N, Koren I, *et al.* Exploring heterogeneity within a core for improved power efficiency. *IEEE Trans. on Parallel and Distributed Systems*, 2016,27(4):1057–1069.
- [26] Muthukaruppan TS, Pricopi M, Venkataramani V, Mitra T, Vishin S. Hierarchical power management for asymmetric multicore in dark silicon era. In: *Proc. of the Design Automation Conf*. 2013. 1–9.
- [27] Colin A, Kandhalu A, Rajkumar R, *et al.* Energy-efficient allocation of real-time applications onto heterogeneous processors. In: *Proc. of the IEEE Int'l Conf. on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. 2014,84(1):1–10.
- [28] Cheng T, Muthukaruppan TS, Mitra T, Lei J. Approximation-aware scheduling on heterogeneous multi-core architectures. In: *Proc. of the Design Automation Conf*. 2015. 618–623.

- [29] Tseng PH, Hsiu PC, Pan CC, Kuo TW. User-centric energy-efficient scheduling on multi-core mobile devices. In: Proc. of the 51st Annual Design Automation Conf. 2014. 1–6.
- [30] Haque ME, He Y, Elnikety S, *et al.* Exploiting heterogeneity for tail latency and energy efficiency. In: Proc. of the IEEE/ACM Int'l Symp. 2017. 625–638.
- [31] Kim YG, Kim M, Chung SW. Enhancing energy efficiency of multimedia applications in heterogeneous mobile multicore processors. *IEEE Trans. on Computers*, 2017,P(99):1.
- [32] Yuhao Z, Reddi VJ. High-performance and energy-efficient mobile Web browsing on big/little systems. In: Proc. of the IEEE Int'l Symp. on High Performance Computer Architecture. 2013. 13–24.
- [33] Sang JN, Kim YG, Chung SW. An energy-efficient task scheduler for mobile web browsing. In: Proc. of the IEEE Int'l Conf. on Consumer Electronics. 2017. 188–189.
- [34] Annamalai A, Rodrigues R, Koren I, Kundu S. An opportunistic prediction-based thread scheduling to maximize throughput/watt in AMPs. In: Proc. of the Int'l Conf. on Parallel Architectures & Compilation Techniques. 2013.
- [35] Homayoun H, Kontorinis V, Shayan A, *et al.* Dynamically heterogeneous cores through 3D resource pooling. In: Proc. of the 18th Int'l Symp. on High Performance Computer Architecture. 2012. 1–12.
- [36] Khubaib K, Suleman MA, Hashemi M, *et al.* MorphCore: An energy-efficient microarchitecture for high performance ILP and high throughput TLP. In: Proc. of the IEEE Int'l Symp. on Microarchitecture. 2012. 305–316.
- [37] Kim M, Kim K, Geraci JR, Hong S. Utilization-aware load balancing for the energy efficient operation of the big.LITTLE processor. In: Proc. of the Design, Automation & Test in Europe Conf. & Exhibition. 2014.
- [38] Ko B-M, Lee J, Jo H. AMP aware core allocation scheme for mobile devices. In: Proc. of the Spring World Congress on Engineering and Technology. 2012. 1–4.
- [39] Lin FX, Wang Z, Likamwa R, *et al.* Reflex: Using low-power processors in smartphones without knowing them. *ACM SIGARCH Computer Architecture News*, 2012,40(1):13.
- [40] Lukefahr A, Padmanabha S, Das R, *et al.* Heterogeneous microarchitectures trump voltage scaling for low-power cores. In: Proc. of the Int'l Conf. on Parallel Architectures & Compilation Techniques. 2014. 237–250.
- [41] Mühlbauer T, Rödiger W, Seilbeck R, *et al.* Heterogeneity-conscious parallel query execution: Getting a better mileage while driving faster. In: Proc. of the Heterogeneity-conscious Parallel Query Execution. 2014. 1–10.
- [42] Nishtala R, Mosse D, Petrucci V. Energy-aware thread co-location in heterogeneous multicore processors. In: Proc. of the Int'l Conf. on Embedded Software. 2013. 1–9.
- [43] Liu Y, Li Y, Zhao Y, *et al.* A scheduling algorithm in the randomly heterogeneous multi-core processor. In: Proc. of the 12th Int'l Conf. on Natural Computation. 2016. 2140–2146.
- [44] Singh P, Hailu N. Energy-aware online non-clairvoyant multiprocessor scheduling: Multiprocessor priority round robin. *IET Computers & Digital Techniques*, 2017,11(1):16–23.
- [45] Luo T, Ma S, Lee R, *et al.* Fairness-aware scheduling on single-ISA heterogeneous multi-cores. In: Proc. of the Int'l Conf. on Parallel Architectures & Compilation Techniques. 2013. 177–188.
- [46] Fedorova A, Seltzer M, Smith MD. Cache-fair thread scheduling for multicore processors. Harvard University, 2009. https://www.researchgate.net/publication/248502496_Cache-Fair_Thread_Scheduling_for_Multicore_Processors
- [47] Munkres J. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 1957,5(1):32–38.
- [48] Saez JC, Pousa A, Castro F, *et al.* ACFS: A completely fair scheduler for asymmetric single-isa multicore systems. In: Proc. of the 30th Annual ACM Symp. on Applied Computing, Vols. I and II. 2015. 2027–2032.
- [49] Kwon Y, Kim C, Maeng S, *et al.* Virtualizing performance asymmetric multi-core systems. In: Proc. of the Int'l Symp. on Computer Architecture. 2011. 45.
- [50] Sun T, An H, Wang T, *et al.* CRQ-based fair scheduling on composable multicore architectures. In: Proc. of the 26th ACM Int'l Conf. on Supercomputing. 2012. 173.
- [51] Markovic N, Nemirovsky D, Unsal O, *et al.* Thread lock section-aware scheduling on asymmetric single-isa multi-core. *IEEE Computer Architecture Letters*, 2015,14(2):160–163.


- [52] Pricopi M, Mitra T. Task scheduling on adaptive multi-core. *IEEE Trans. on Computers*, 2014,63(10):2590–2603.
- [53] Torng C, Wang M, Batten C. Asymmetry-aware work-stealing runtimes. In: *Proc. of the ACM/IEEE Int'l Symp. on Computer Architecture*. 2016,44(3):40–52.
- [54] Chen Q, Guo M. Adaptive workload-aware task scheduling for single-ISA asymmetric multicore architectures. *ACM Trans. on Architecture and Code Optimization*, 2014,11(1):1–25.
- [55] Lakshminarayana NB, Lee J, Kim H. Age based scheduling for asymmetric multiprocessors. In: *Proc. of the Conf. on High Performance Computing Networking, Storage and Analysis*. 2009. 25.
- [56] Suleman MA, Mutlu O, Qureshi MK, *et al.* Accelerating critical section execution with asymmetric multicore architectures. *IEEE Micro*, 2010,30(1):60–70.
- [57] Lakshminarayana N, Rao S, Kim H. Asymmetry aware scheduling algorithms for asymmetric multiprocessors. *Cc.gatech.edu*, 2008.
- [58] Joao JA, Suleman MA, Mutlu O, *et al.* Bottleneck identification and scheduling in multithreaded applications. *ACM SIGARCH Computer Architecture News*, 2012,40(1):223.
- [59] Joao JA, Suleman MA, Mutlu O, *et al.* Utility-based acceleration of multithreaded applications on asymmetric CMPs. In: *Proc. of the 40th Annual Int'l Symp. on Computer Architecture*. 2013. 154–165.
- [60] Hwang W. HMMSched: Hybrid main memory-aware task scheduling on multicore systems. In: *Proc. of the 5th Int'l Conf. on Future Computational Technologies and Applications*. 2013. 39–48.
- [61] Merkel A, Bellosa F. Memory-aware scheduling for energy efficiency on multicore processors. In: *Proc. of the Workshop on Power Aware Computing & Systems*. 2008. 1.
- [62] Boyd-Wickizer S, Morris R, Kaashoek MF. Reinventing scheduling for multicore systems. In: *Proc. of the Workshop on Hot Topics in Operating Systems*. Monte Verità, 2009.
- [63] Ge R, Feng X, Cameron KW. Modeling and evaluating energy-performance efficiency of parallel processing on multicore based power aware systems. In: *Proc. of the IEEE Int'l Symp. on Parallel & Distributed Processing*. 2009. 1–8.
- [64] Merkel A, Stoess J, Bellosa F. Resource-conscious scheduling for energy efficiency on multicore processors. In: *Proc. of the European Conf. on European Computer Systems*. 2010. 153–166.
- [65] Eyerhan S, Eeckhout L. Modeling critical sections in Amdahl's law and its implications for multicore design. *ACM SIGARCH Computer Architecture News*, 2010,38(3):362.
- [66] Jia R, Kun W, Xiaobo Z, *et al.* Optimizing virtual machine scheduling in NUMA multicore systems. In: *Proc. of the IEEE Int'l Symp. on High Performance Computer Architecture*. 2013. 306–317.
- [67] Takouna I, Dawoud W, Meinel C. Efficient virtual machine scheduling-policy for virtualized heterogeneous multicore systems. 2011. https://www.researchgate.net/publication/228459497_Efficient_Virtual_Machine_Scheduling-policy_for_Virtualized_Heterogeneous_Multicore_Systems
- [68] Kazempour V, Kamali A, Fedorova A. AASH: An asymmetry-aware scheduler for hypervisors. *ACM SIGPLAN Notices*, 2010, 45(7):85–96.
- [69] Wang Y, Wang X, Chen Y. Energy-efficient virtual machine scheduling in performance-asymmetric multi-core architectures. In: *Proc. of the Int'l Conf. on Network & Service Management*. 2012. 288–294.
- [70] Wang G, Wang Y, Liu H, *et al.* HSIP: A novel task scheduling algorithm for heterogeneous computing. *Scientific Programming*, 2016,2016:1–11.
- [71] Jian C, John LK. Energy-aware application scheduling on a heterogeneous multi-core system. In: *Proc. of the IEEE Int'l Symp. on Workload Characterization*. 2008. 5–13.
- [72] Chen J, John LK. Efficient program scheduling for heterogeneous multi-core processors. In: *Proc. of the 46th Annual Design Automation Conf*. 2009. 927–930.
- [73] Shelepov D, Saez Alcaide JC, Jeffery S, *et al.* HASS: A scheduler for heterogeneous multicore systems. *ACM SIGOPS Operating Systems Review*, 2009,43(2):66–75.
- [74] Saez JC, Shelepov D, Fedorova A, *et al.* Leveraging workload diversity through OS scheduling to maximize performance on single-ISA heterogeneous multicore systems. *Journal of Parallel and Distributed Computing*, 2011,71(1):114–131.

- [75] Kiran DC, Radheshyam B, Gurunaryanan S, *et al.* Compiler assisted dynamic scheduling for multicore processors. In: Proc. of the Int'l Conf. on Process Automation. 2011. 1–6.
- [76] Sondag T, Rajan H. Phase-guided thread-to-core assignment for improved utilization of performance-asymmetric multi-core processors. In: Proc. of the Int'l Workshop on Materials Science and Engineering. 2009. 73–80.
- [77] Sondag T, Krishnamurthy V, Rajan H. Predictive thread-to-core assignment on a heterogeneous multi-core processor. In: Proc. of the 4th Workshop on Programming Languages and Operating Systems. 2007. 1.
- [78] Koufaty D, Reddy D, Hahn S. Bias scheduling in heterogeneous multi-core architectures. In: Proc. of the 5th European Conf. on Computer Systems. 2010. 125–138.
- [79] Chen J, Nair AA, John LK. Predictive heterogeneity-aware application scheduling for chip multiprocessors. *IEEE Trans. on Computers*, 2014,63(2):435–447.
- [80] Van Craeynest K, Jaleel A, Eeckhout L, *et al.* Scheduling heterogeneous multi-cores through performance impact estimation (PIE). *ACM SIGARCH Computer Architecture News*, 2012, 213–224.
- [81] Mattson RL, Gecsei J, Slutz DR, *et al.* Evaluation techniques for storage hierarchies. *IBM Systems Journal*, 1970,9(2):78–117.
- [82] Saez JC, Fedorova A, Koufaty D, *et al.* Leveraging core specialization via OS scheduling to improve performance on asymmetric multicore systems. *ACM Trans. on Computer Systems*, 2012,30(2):1–38.
- [83] Saez JC, Prieto M, Fedorova A, *et al.* A Comprehensive Scheduler for Asymmetric Multicore Systems. New York: Association for Computing Machinery, 2010. 139–152.
- [84] Khan O, Kundu S. A self-adaptive scheduler for asymmetric multi-cores. In: Proc. of the Symp. on Great Lakes VLSI. 2010. 397.
- [85] Cong J, Yuan B. Energy-efficient scheduling on heterogeneous multi-core architectures. In: Proc. of the IEEE Int'l Symp. on Low Power Electronics & Design. 2012. 345.
- [86] Pricopi M, Muthukaruppan TS, Venkataramani V, *et al.* Power-performance modeling on asymmetric multi-cores. In: Proc. of the Int'l Conf. on Compilers. 2013. 1–10.
- [87] Eyerman S, Hoste K, Eeckhout L. Mechanistic-empirical processor performance modeling for constructing CPI stacks on real hardware. In: Proc. of the IEEE Int'l Symp. on Performance Analysis of Systems and Software. 2011. 216–226.
- [88] Jiang X, Mishra A, Zhao L, *et al.* ACCESS: Smart scheduling for asymmetric cache CMPs. In: Proc. of the IEEE Int'l Symp. on High Performance. 2011. 527–538.
- [89] Jooya AZ, Baniyadi A, Analoui M. History-aware, resource-based dynamic scheduling for heterogeneous multi-core processors. *IET Computers & Digital Techniques*, 2011.
- [90] Sawalha L, Wolff S, Tull MP, *et al.* Phase-guided scheduling on single-ISA heterogeneous multicore processors. In: Proc. of the 14th Euromicro Conf. on Digital System Design. 2011. 736–745.
- [91] Sawalha L, Barnes RD. Energy-efficient phase-aware scheduling for heterogeneous multicore processors. In: Proc. of the Green Technologies Conf. 2012. 1–6.
- [92] Li T, Baumberger D, Koufaty DA, *et al.* Efficient operating system scheduling for performance-asymmetric multi-core architectures. In: Proc. of the 2007 ACM/IEEE Conf. on Supercomputing. 2007. 53.
- [93] Muthukaruppan TS, Pathania A, Mitra T. Price theory based power management for heterogeneous multi-cores. In: Proc. of the Int'l Conf. on Architectural Support for Programming Languages & Operating Systems. 2014. 161–176.
- [94] Van Craeynest K, Eeckhout L. Understanding fundamental design choices in single-ISA heterogeneous multicore architectures. *ACM Trans. on Architecture and Code Optimization*, 2013,9(4):1–23.
- [95] Mittal S. A survey of techniques for architecting and managing asymmetric multicore processors. *ACM Computing Surveys*, 2016, 48(3):1–38.
- [96] Shiu E, Prakash S. System challenges and hardware requirements for future consumer devices: From wearable to ChromeBooks and devices in-between. In: Proc. of the 2015 Symp. on VLSI Technology. 2015. 1–5.
- [97] Zhou X, Yang J, Xu Y, *et al.* Thermal-aware task scheduling for 3D multicore processors. *IEEE Trans. on Parallel and Distributed Systems*, 2010,21(1):60–71.
- [98] Xu TC, Liljeberg P, Plosila J, *et al.* Exploration of heuristic scheduling algorithms for 3D multicore processors. In: Proc. of the Map2MPSoc/SCOPEs. 2012. 22–31.

- [99] Srinivasan S, Iyer R, Zhao L, *et al.* HeteroScouts: Hardware assist for OS scheduling in heterogeneous CMPs. In: Proc. of the ACM SIGMETRICS Joint Int'l Conf. on Measurement & Modeling of Computer Systems. 2011. 149.
- [100] Vetter JS, Mittal S. Opportunities for nonvolatile memory systems in extreme-scale high-performance computing. *Computing in Science & Engineering*, 2015,17(2):73–82.



赵姗(1982—),女,河南商丘人,高级工程师,CCF 学生会员,主要研究领域为操作系统,软硬件深度融合.



杨秋松(1977—),男,博士,研究员,博士生导师,CCF 专业会员,主要研究领域为软件工程,形式化方法,操作系统.



李明树(1966—),男,博士,研究员,博士生导师,CCF 会士,主要研究领域为操作系统深度设计(包括安全操作系统、数据操作系统等),可信软件过程,基础软硬件核心技术与应用.

www.jos.org.cn

www.jos.org.cn