# Asymmetry-Aware Work-Stealing Runtimes Errata and Addendum

Christopher Torng, Moyang Wang, and Christopher Batten
School of Electrical and Computer Engineering, Cornell University, Ithaca, NY
{clt67,mw828,cbatten}@cornell.edu

The analytical modeling used in our conference paper *Asymmetry-Aware Work-Stealing Runtimes* [1] assumed that the big core consumed 10% static leakage power (i.e., $\lambda = 0.1$) at nominal voltage and frequency. Refer to the parameters listed at the end of Section II.B to see this. Due to a mistype in our scripts, the first-order estimates plotted in Figures 2–5 correspond instead to 50% static leakage power (i.e., $\lambda = 0.5$ ) for the big core at nominal voltage and frequency. In this errata, we provide both the original and corrected versions of these figures for reference by any future researchers in this area. To first order, the difference in leakage does not change the high-level conclusions in the paper and does not impact the applicability of the proposed AAWS runtime techniques. We also take the opportunity to present a small addendum sweeping the leakage parameter $\lambda$ in order to analyze the first-order impact of leakage on opportunities for balancing marginal utility in a 4B4L system.

Please refer to Section II in the conference paper for descriptions of the analytical model, assumed parameters for numerical analysis, and descriptions of each figure. The scripts for our analytical model and a description of their usage are available online [2].

## I. Corrected Leakage

This errata provides both the original and corrected versions of Figures 2–5 from the conference paper. For each figure, the version on the left is original with $\lambda = 0.5$ (higher leakage), and the version on the right is corrected with $\lambda = 0.1$ (lower leakage). In this subsection, we briefly discuss the impact of leakage on each figure.

Comparing the original and corrected Figure 2, we see that to first order, the system on the left with higher leakage tends to achieve lower energy efficiency across all points, where each point corresponds to a different ($V_{Bi}$, $V_{Lj}$) pair. This result is not surprising because leakage power scales linearly with voltage, while dynamic power scales cubically with voltage. Note that the pareto-optimal isopower point denoted by the open circle does not change, implying that the potential benefit of exploiting marginal utility imbalance using the AAWS techniques is likely to be similar for systems with either lower or higher leakage.

Comparing the original and corrected Figures 3 and 5, we can see how leakage impacts the slope of the power versus performance curves across the DVFS operating range for each core type. Lower leakage results in steeper slopes. Again, this is not surprising due to the scaling of leakage power and dynamic power. In lower leakage systems, dynamic power dominates, which results in a greater relative impact on system power when scaling voltage. In higher leakage systems, static power is more significant, resulting in a shallower curve. In both Figure 3 and Figure 4, note that the highest throughput does not change significantly and still occurs when the marginal utilities curves intersect (i.e., when the marginal utilities are equal). The high-level conclusions of the paper are therefore the same for both low-leakage and high-leakage systems.

Comparing the original and corrected Figure 4, we see that the theoretical speedups do not vary significantly with leakage. Recall that the analytical modeling in the conference paper assumes $\beta = 2$ and $\alpha = 3$. For both lower and higher leakage, the theoretical optimum speedup is roughly the same (i.e., between 1.100 and 1.125) for the 4B4L system. Similarly, for both lower and higher leakage, the feasible speedup within $V_{min}$ and $V_{max}$ is roughly the same (i.e., about 1.100).

## II. Leakage Sweep

We take the opportunity to present a small addendum sweeping the leakage parameter $\lambda$ in order to analyze the first-order impact of leakage on opportunities for balancing marginal utility in a 4B4L system. Figure 6 and Figure 7 are new plots that do not appear in the conference paper. Each row in Figure 6 corresponds to Figure 3 with a specific value of $\lambda$, and similarly, each row in Figure 7 corresponds to Figure 4 with a specific value of $\lambda$. Starting from the first row with zero leakage and moving down to the fourth row with 50% leakage, we see that the power versus performance curves slowly become shallower for both core types, just as we saw earlier in our analysis of the original versus corrected versions of Figure 3 and Figure 4.

In Figure 6 and Figure 7, notice that the optimal operating point (denoted by the star) occurs at slightly higher $V_L$ as leakage increases. For example, the optimum $V_L$ in Figure 7 with zero leakage is 1.62 V, but with $\lambda = 0.5$, the optimum $V_L$ is 1.69 V. With higher leakage, the little cores must travel further along the power versus performance curve to balance marginal utility with the big core. Notice that the *feasible* operating point does not change, since the little core quickly hits $V_{max}$ in both figures. Finally, notice that the highest throughput still occurs where the marginal utility curves intersect.

## III. Conclusions

In this errata, we present side-by-side original and corrected figures for a mistyped leakage value in our conference paper. We also briefly survey the impact of leakage by sweeping $\lambda$ in our analytical model. We find that the high-level conclusions of the paper do not change. We hope that this errata and addendum will help make the case that holistically combining static asymmetry, dynamic asymmetry, and work-stealing runtimes can improve performance and energy efficiency in future multicore systems.

[1] C. Torng, M. Wang, and C. Batten. "Asymmetry-Aware Work-Stealing Runtimes." 43rd ACM/IEEE Int'l Symp. on Computer Architecture (ISCA-43), June 2016.

[2] AAWS Scripts: https://github.com/cornell-brg/torng-aaws-scripts-isca2016

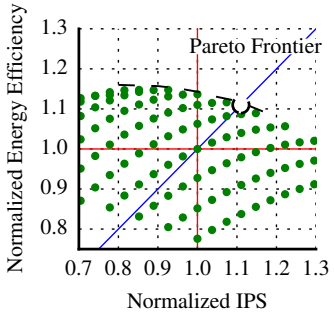


**Figure 2 (Original) with $\lambda$ = 0.5.** Pareto-Optimal Frontier for 4B4L System – Projected energy efficiency vs. performance of a busy 4B4L system across different ($V_{Bi}$, $V_{Lj}$) pairs. Points normalized to (1.0 V, 1.0 V) system. Diagonal line is isopower. Open circle = pareto-optimal isopower system. $\alpha = 3$, $\beta = 2$.

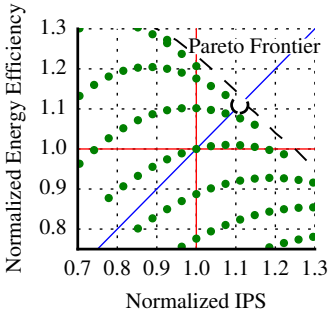


**Figure 2 (Corrected) with $\lambda$ = 0.1.** Pareto-Optimal Frontier for 4B4L System – Projected energy efficiency vs. performance of a busy 4B4L system across different ($V_{Bi}$, $V_{Lj}$) pairs. Points normalized to (1.0 V, 1.0 V) system. Diagonal line is isopower. Open circle = pareto-optimal isopower system. $\alpha = 3$, $\beta = 2$.

---

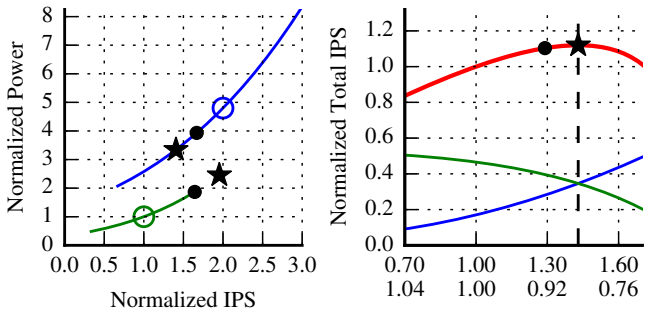


(a) Power vs. IPS (b) Total IPS vs. $V_{Bi}$ and $V_{Lj}$

**Figure 3 (Original) with $\lambda$ = 0.5.** 4B4L System w/ All Cores Active – (a) Power vs. performance curves across the DVFS operating points for each core type, green = little, blue = big, circle = nominal; (b) blue = $\partial P_{BAi}/\partial IPS_{BAi}$ (axis not shown), green = $\partial P_{LAj}/\partial IPS_{LAj}$ (axis not shown), red = $IPS_{tot}$ (axis on left) assuming $V_{Lj}$ and $V_{Bi}$ shown on x-axis ($V_{Lj}$ on top, $V_{Bi}$ on bottom) with constant $P_{target}$. (a–b) star = optimal operating point, dot = feasible operating point, $\alpha = 3$, $\beta = 2$.

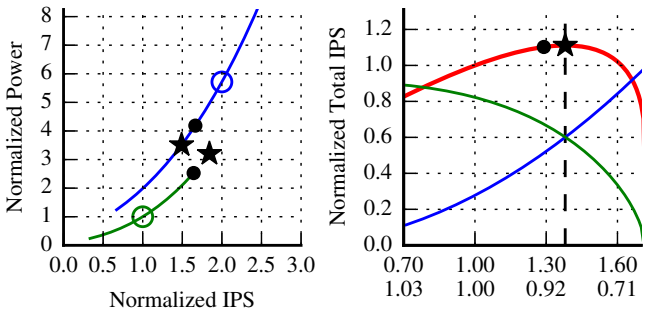


(a) Power vs. IPS (b) Total IPS vs. $V_{Bi}$ and $V_{Lj}$

**Figure 3 (Corrected) with $\lambda$ = 0.1.** 4B4L System w/ All Cores Active – (a) Power vs. performance curves across the DVFS operating points for each core type, green = little, blue = big, circle = nominal; (b) blue = $\partial P_{BAi}/\partial IPS_{BAi}$ (axis not shown), green = $\partial P_{LAj}/\partial IPS_{LAj}$ (axis not shown), red = $IPS_{tot}$ (axis on left) assuming $V_{Lj}$ and $V_{Bi}$ shown on x-axis ($V_{Lj}$ on top, $V_{Bi}$ on bottom) with constant $P_{target}$. (a–b) star = optimal operating point, dot = feasible operating point, $\alpha = 3$, $\beta = 2$.

---

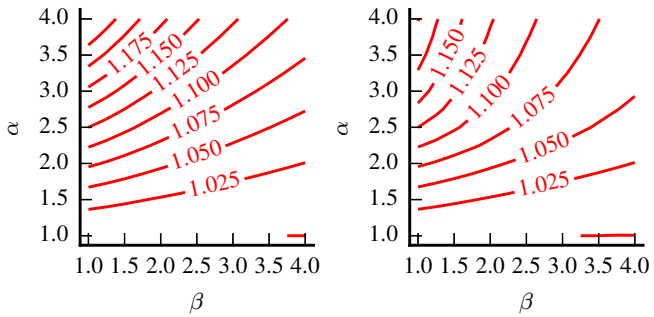


(a) Optimum Speedup (b) Feasible Speedup

**Figure 4 (Original) $\lambda$ = 0.5.** Theoretical Speedup for 4B4L System vs. $\alpha$ and $\beta$ – (a) optimum speedup ignoring $V_{min}$ and $V_{max}$; (b) = feasible speedup within $V_{min}$ and $V_{max}$. Speedups relative to all cores running at $V_N$.

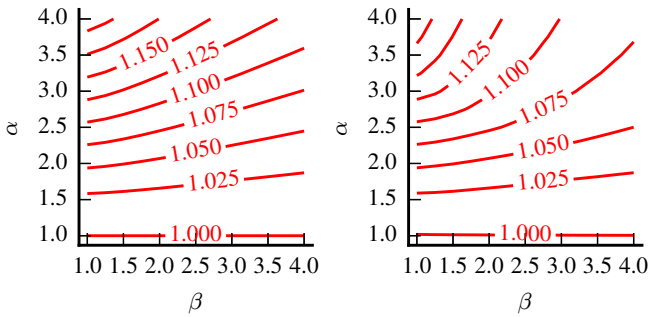


(a) Optimum Speedup (b) Feasible Speedup

**Figure 4 (Corrected) $\lambda$ = 0.1.** Theoretical Speedup for 4B4L System vs. $\alpha$ and $\beta$ – (a) optimum speedup ignoring $V_{min}$ and $V_{max}$; (b) = feasible speedup within $V_{min}$ and $V_{max}$. Speedups relative to all cores running at $V_N$.

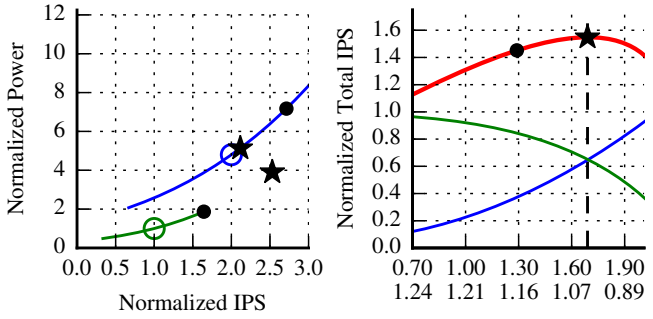


(a) Power vs. IPS (b) Total IPS vs. $V_{Bi}$ and $V_{Lj}$

**Figure 5 (Original) with $\lambda = 0.5$**. 4B4L System w/ 2B2L Active – Assume we rest inactive cores at $V_{min}$. See Figure 3 for legend. $\alpha = 3$, $\beta = 2$.

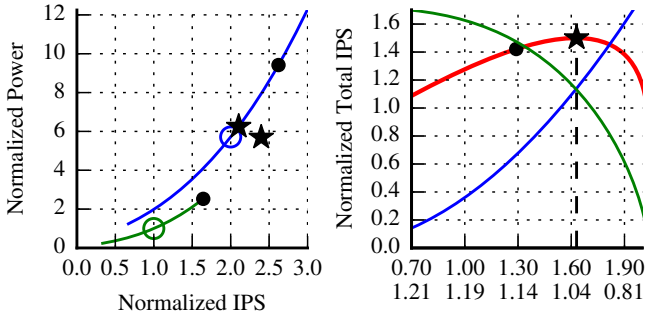


(a) Power vs. IPS (b) Total IPS vs. $V_{Bi}$ and $V_{Lj}$

**Figure 5 (Corrected) $\lambda = 0.1$**. 4B4L System w/ 2B2L Active – Assume we rest inactive cores at $V_{min}$. See Figure 3 for legend. $\alpha = 3$, $\beta = 2$.

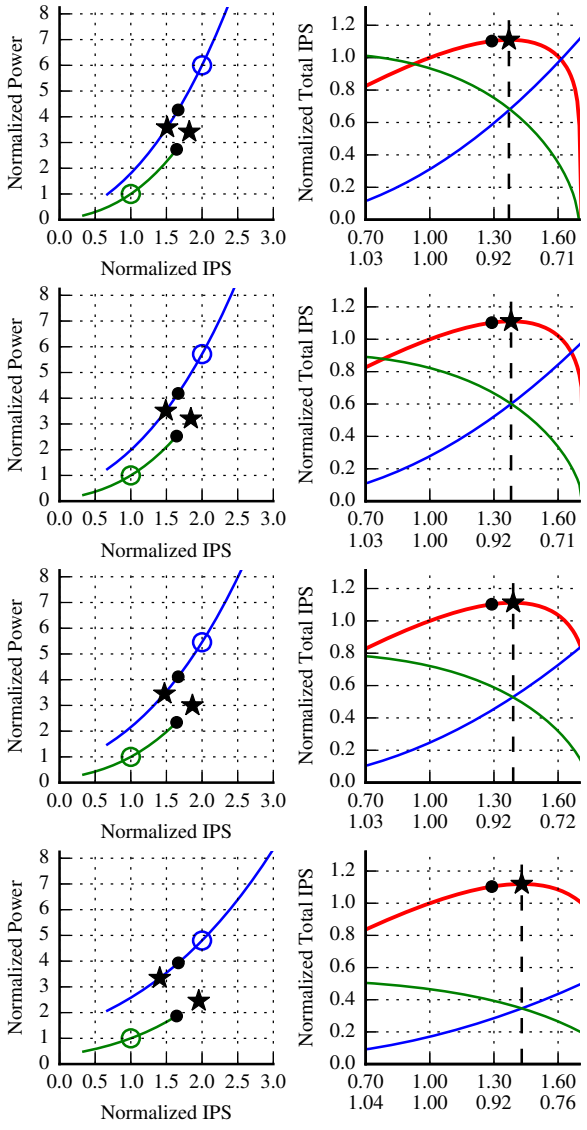


**Figure 6**. Leakage Sweep for 4B4L System w/ All Cores Active – Each set of plots is identical to Figure 3, but each row corresponds to a different value of $\lambda$. *First row* sets $\lambda = 0.0$; *second row* sets $\lambda = 0.1$; *third row* sets $\lambda = 0.2$; *fourth row* sets $\lambda = 0.5$.

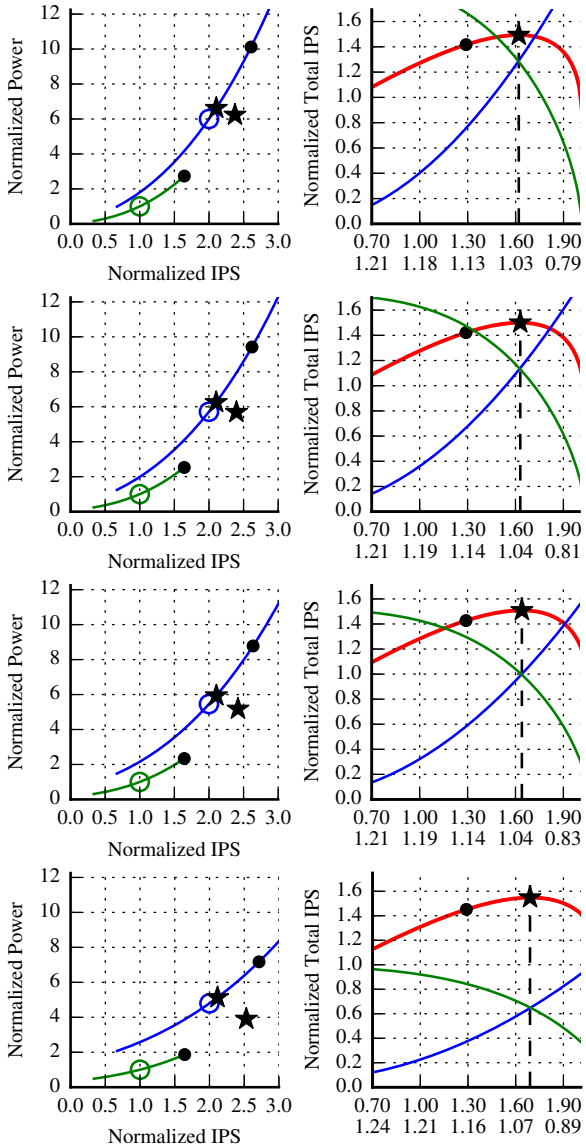


**Figure 7**. Leakage Sweep for 4B4L System w/ 2B2L Active – Each set of plots is identical to Figure 5, but each row corresponds to a different value of $\lambda$. *First row* sets $\lambda = 0.0$; *second row* sets $\lambda = 0.1$; *third row* sets $\lambda = 0.2$; *fourth row* sets $\lambda = 0.5$.