

## Errata Slip (Updated 12/23/15)

In the paper “Spartan: A Distributed Array Framework with Smart Tiling” by Chien-Chin Huang, *New York University*; Qi Chen, *Peking University*; Zhaoguo Wang and Russell Power, *New York University*; Jorge Ortiz, *IBM T.J. Watson Research Center*; Jinyang Li, *New York University*; Zhen Xiao, *Peking University* (Wednesday session “Parallel & Distributed Systems,” pp. 1–15 of the Proceedings), the authors corrected the following:

Figure 6(b) original

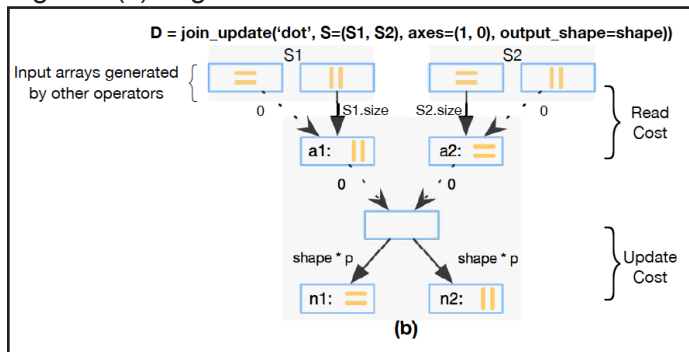
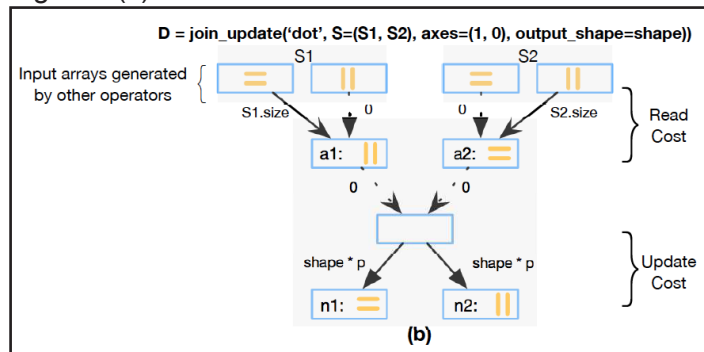


Figure 6(b) corrected



The read costs of figure 6.b are incorrect in the original version. The read cost of S1 in figure 6(b) should be  $S1.size$  if S1 is partitioned by row (originally mistyped as 0); if S1 is partitioned by column, the cost should be 0 (originally mistyped as  $S1.size$ ).

The read cost of S2 in figure 6(b) should be  $S2.size$  if S2 is partitioned by column (originally mistyped as 0); if S2 is partitioned by row, the cost should be 0 (originally mistyped as  $S2.size$ ).

In the paper “LAMA: Optimized Locality-aware Memory Allocation for Key-value Cache” by Xiameng Hu, Xiaolin Wang, Yechen Li, Lan Zhou, and Yingwei Luo, *Peking University*; Chen Ding, *University of Rochester*; Song Jiang, *Wayne State University*; Zhenlin Wang, *Michigan Technological University* (Wednesday session “Cloud Storage,” pp. 57–69 of the proceedings), the following errors occurred:

Page 60, Section 3.2

Replace sentence:

Original Text

Now we can profile the MRC using  $fp$  distribution.

The miss ratio for cache size of  $x$  is the fraction of reuses that have an average footprint smaller than  $x$ :

Corrected Text

Now we can profile the MRC using  $fp$  distribution.

The miss ratio for cache size of  $x$  is the fraction of reuses that have an average footprint larger than  $x$ :

Page 68, References

Replace reference:

Original Text

[19] Jacob Brock, Yechen Li, Chencheng Ye, and Chen Ding. Optimal cache partition-sharing : Dont ever take a fence down until you know why it was put up. robert frost. *In Proceedings of ICPP*, 2015.

Corrected Text

[19] Jacob Brock, Chencheng Ye, Chen Ding, Yechen Li, Xiaolin Wang, and Yingwei Luo. Optimal cache partition-sharing. *In Proceedings of ICPP*, 2015

Continued on next page →

---

In the paper “Latency-Tolerant Software Distributed Shared Memory,” by Jacob Nelson, Brandon Holt, Brandon Myers, Preston Brigg, Luis Ceze, Simon Kahan, and Mark Oskin, *University of Washington* (Thursday session “Memory,” pp. 291–305 of the proceedings), the authors omitted the following:

**Additional Text**

**Acknowledgements**

This work was supported by NSF Grant CCF-1335466, Pacific Northwest National Laboratory, and gifts from NetApp and Oracle.

---

In the paper “Hawk: Hybrid Datacenter Scheduling,” by Pamela Delgado and Florin Dinu, *École Polytechnique Fédérale de Lausanne (EPFL)*; Anne-Marie Kermarrec, *Inria*; Willy Zwaenepoel, *École Polytechnique Fédérale de Lausanne (EPFL)* (Friday session, “Scheduling at Large Scale,” pp. 499–510 of the proceedings), the authors corrected the following:

Page 505, Section 4.2

First column, last paragraph: “the percentage of jobs” was modified to “the fraction of jobs”

Page 509, Acknowledgements Section

“Schwartzkopf” was modified to “Schwarzkopf”

---

In the paper “Bolt: Faster Reconfiguration in Operating Systems,” by Sankaralingam Panneerselvam and Michael M. Swift, *University of Wisconsin—Madison* (Friday session, “OS & Hardware,” pp. 511–516 of the proceedings), changes have been made to text references regarding speculated hardware in the abstract, introduction and evaluation sections.

**Original version:**

The operating system is required to wait for 10ms for the purpose of hardware initialization while starting up a x86 processor core. The speculated hardware assumed zero initialization delay.

**Corrected version:**

The speculated hardware is real since the modern x86 multi-core processors need not incur the initialization delay.

---

In the paper “Fence: Protecting Device Availability With Uniform Resource Control,” by Tao Li and Albert Rafetseder, *New York University*; Rodrigo Fonseca, *Brown University*; Justin Cappos, *New York University* (Wednesday session, “Dependability,” pp. 177–191 of the proceedings), the following changes were made:

In the original paper, the affiliation for Tao Li, Albert Rafetseder, and Justin Cappos was listed as Polytechnic Institute of New York Univeristy. This was corrected to say New York University.

Table 1 and Table 2 were reformatted to fit within the column margins.

Continued on next page →

In the paper “Data Sharing or Resource Contention: Toward Performance Transparency on Multicore Systems,” by Sharanyan Srikanthan, Sandhya Dwarkadas, and Kai Shen, *University of Rochester*, (Friday session, “OS & Hardware,” pp. 529–540 of the proceedings), the following changes were made:

1. Figures 6(c) and Figure 8(c) contained incorrect Intra-Socket Coherence counts. This error has been corrected. This errata reflects no change to the technique or the remaining qualitative and quantitative results. The only change to the text to reflect the figure changes is as follows: “Note that for workloads 14 and 15, SAM shows reductions in” has been changed to “Note that for some workloads, SAM shows reductions in.”

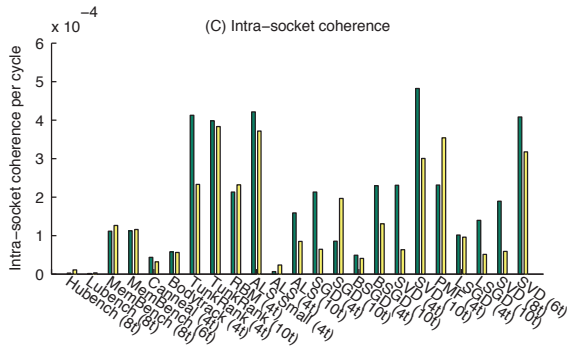


Figure 6(c) original




Figure 6(c) corrected

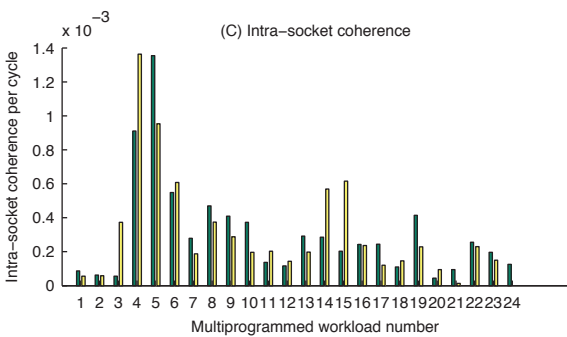


Figure 8(c) original




Figure 8(c) corrected

2. In Table 2, workload 16 characteristics were incorrectly a copy of workload 15. This error has been corrected to “8 SVD, 8 LSGD”.