Errata Slip

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

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.