Proceedings of the
18th USENIX Symposium on Networked Systems Design and Implementation (NSDI '21)

Errata Slip #1

In the paper "Segcache: a memory-efficient and scalable in-memory key-value cache for small objects" by Juncheng Yang, *Carnegie Mellon University;* Yao Yue, *Twitter;* Rashmi Vinayak, *Carnegie Mellon University* (Tuesday session, "System Performance and Programmability," pp. 503–518 of the Proceedings), the authors wish to make the following clarification:

Page 509, section 3.6.3, first paragraph:

Original text:
Under the independent reference model (IRM), a popular model used for cache workloads, an object with a higher frequency is more likely to be re-accessed. Moreover, it has been shown in theory that under IRM and for fix-sized objects, the least frequently used (LFU) is *k*-competitive and the best policy [27, 39, 61, 64].

Corrected text:
Under the independent reference model (IRM), a popular model used for cache workloads, an object with a higher frequency is more likely to be re-accessed. And it has been shown to be effective by many works [27, 39, 61, 64].