REFERENCES

[1] Joy Arulraj and Andrew Pavlo. 2017. How to Build a Non-Volatile Memory Database Management System. In *Proceedings of the 2017 ACM International Conference on Management of Data - SIGMOD '17*. ACM Press, Chicago, Illinois, USA, 1753–1758. https://doi.org/10.1145/3035918.3054780

[2] Joy Arulraj, Andrew Pavlo, and Subramanya R. Dulloor. 2015. Let's Talk About Storage & Recovery Methods for Non-Volatile Memory Database Systems. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data - SIGMOD '15*. ACM Press, Melbourne, Victoria, Australia, 707–722. https://doi.org/10.1145/2723372.2749441

[3] Carsten Binnig, Andrew Crotty, Alex Galakatos, Tim Kraska, and Erfan Zamanian. 2016. The end of slow networks: it's time for a redesign. *Proceedings of the VLDB Endowment* 9, 7 (March 2016), 528–539. https://doi.org/10.14778/2904483.2904485

[4] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache Flink™: Stream and Batch Processing in a Single Engine. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* 36, 4 (2015), 12.

[5] The Apache Software Foundation. 2019. Apache Pulsar. https://pulsar.apache.org/

[6] Jian Huang, Karsten Schwan, and Moinuddin K. Qureshi. 2014. NVRAM-aware logging in transaction systems. *Proceedings of the VLDB Endowment* 8, 4 (Dec. 2014), 389–400. https://doi.org/10.14778/2735496.2735502

[7] InfiniBand Trade Association. 2019. InfiniBand Roadmap - Advancing InfiniBand. https://www.infinibandta.org/infiniband-roadmap/

[8] Intel Corporation. 2019. Intel® Optane™ DC Persistent Memory Product Brief. https://www.intel.com/content/www/us/en/products/docs/memory-storage/optane-persistent-memory/optane-dc-persistent-memory-brief.html

[9] Jay Kreps, Neha Narkhede, and Jun Rao. 2011. Kafka: A Distributed Messaging System for Log Processing. 1–7.

[10] Steven Pelley, Thomas F. Wenisch, Brian T. Gold, and Bill Bridge. 2013. Storage management in the NVRAM era. *Proceedings of the VLDB Endowment* 7, 2 (Oct. 2013), 121–132. https://doi.org/10.14778/2732228.2732231

[11] Animesh Trivedi, Patrick Stuedi, Jonas Pfefferle, Radu Stoica, Bernard Metzler, Ioannis Koltsidas, and Nikolas Ioannou. 2016. On the [ir] relevance of network performance for data processing. In *8th {USENIX} Workshop on Hot Topics in Cloud Computing (HotCloud 16)*.

[12] Alexander van Renen, Viktor Leis, Alfons Kemper, Thomas Neumann, Takushi Hashida, Kazuichi Oe, Yoshiyasu Doi, Lilian Harada, and Mitsuru Sato. 2018. Managing Non-Volatile Memory in Database Systems. In *Proceedings of the 2018 International Conference on Management of Data - SIGMOD '18*. ACM Press, Houston, TX, USA, 1541–1555. https://doi.org/10.1145/3183713.3196897

[13] Alexander van Renen, Lukas Vogel, Viktor Leis, Thomas Neumann, and Alfons Kemper. 2019. Persistent Memory I/O Primitives. *arXiv:1904.01614 [cs]* (April 2019). http://arxiv.org/abs/1904.01614 arXiv: 1904.01614.

[14] Erfan Zamanian, Carsten Binnig, Tim Harris, and Tim Kraska. 2017. The end of a myth: distributed transactions can scale. *Proceedings of the VLDB Endowment* 10, 6 (Feb. 2017), 685–696. https://doi.org/10.14778/3055330.3055335

[15] Steffen Zeuch, Bonaventura Del Monte, Jeyhun Karimov, Clemens Lutz, Manuel Renz, Jonas Traub, Sebastian Breß, Tilmann Rabl, and Volker Markl. 2019. Analyzing efficient stream processing on modern hardware. *Proceedings of the VLDB Endowment* 12, 5 (Jan. 2019), 516–530. https://doi.org/10.14778/3303753.3303758