

Scalable Network I/O for the JVM

D. Bromberg — F. Morandat — L. Réveillère

Existing network APIs on multi-core systems are known to have difficulties scaling to high connection rates. Several research efforts have addressed aspects of these performance problems, proposing new techniques that offer valuable performance improvements [Satran et al., 2008, Shalev et al., 2010, Han et al., 2012]. However, they all innovate within the confines of the traditional socket-based networking APIs at the level of the operating system.

On the other hand high-level languages (i.e., Java) raise the level of abstraction and hides OS low level details, allowing programmers to focus on functionalities. Legacy Java applications require a JRE to be executed. This runtime is the keystone of independence between language and system. Even if most parts of JREs are system agnostic, a few primitives needs tight integration within the OS. Currently none of the state of the art JRE addresses efficiently the specific issue of network IOs in the context of multicore architectures. Thus the aim of this project is to do a step towards scalability of network IO on managed runtime environment (MRE) through the example of Java and its virtual machine.

This internship will start by performing a survey of existing approaches for improving the performance of network IOs on top of multicore hardware. The state of the art will mainly include the lasts two years of all major system conferences (Eurosys, Osd, Nsdi, Sosps, Plos, Asplos & Usenix). Then according to one — or some — of these articles, we will patch the network API of a JRE to support a multi-core aware technique. This implementation will be evaluated with gathered benchmarks of popular Java network applications, and compared to a vanilla JVM/JRE. Finally we will design an approach for improving the performance of legacy Java code by leveraging on low-level systems improvement of network connections.

References

- Sangjin Han, Scott Marshall, Byung-Gon Chun, and Sylvia Ratnasamy. Megapipe: a new programming interface for scalable network i/o. In *Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation, OSDI'12*, pages 135–148, Berkeley, CA, USA, 2012. USENIX Association. ISBN 978-931971-96-6.
- Julian Satran, Leah Shalev, Muli Ben-Yehuda, and Zorik Machulsky. Scalable i/o - a well-architected way to do scalable, secure and virtualized i/o. In Muli Ben-Yehuda, Alan L. Cox, and Scott Rixner, editors, *Workshop on I/O Virtualization*. USENIX Association, 2008.
- Leah Shalev, Julian Satran, Eran Borovik, and Muli Ben-Yehuda. Isostack: highly efficient network processing on dedicated cores. In *Proceedings of the 2010 USENIX conference on USENIX annual technical conference, USENIXATC'10*, pages 5–5, Berkeley, CA, USA, 2010. USENIX Association.