

Position Summary: A Streaming Interface for Real-Time Interprocess Communication

Jork Löser

Hermann Härtig

Lars Reuther

Dresden University of Technology, Germany

Abstract

Timely transfer of long, continuous data streams and handling data omission are stringent requirements of multimedia applications. To cope with these requirements, we extend well known mechanisms for inter-address-space data transmission, such as zero-copy and fast IPC, by the notion of time. Therefore, we add a time track to data streams and add mechanisms to limit the validity of data. For cases of overload we add notification and revocation techniques.

Inadequacies of current IPC mechanisms Efficient interprocess communication schemes for long data transfers in non real-time applications [2, 5, 6] address the problem of copy avoidance by using shared memory. But they do not cover issues of time, such as data loss due to CPU shortage.

In hard real-time systems data loss does not occur, because the entire system is designed for the worst case regarding resource needs during execution. However, this results in poor resource utilization and is therefore not practical. Designing the system for the average case improves the overall resource utilization [4, 1], at the cost of quality. Resource shortages during execution can happen and lead to data loss then. To cope with this, data loss should be expressed at the communication layer.

Tolerating occasional resource shortages allows multiple applications to share resources, e.g. memory pools for communication buffers. This in turn requires retracting these resources in overload situations and hence must be supported.

The DROPS Streaming Interface The DROPS Streaming Interface (DSI) is our approach to a real-time interprocess communication subsystem. It defines a user-level timed packet-oriented zero-copy transport protocol between real-time components. The data flows are represented by streams with assigned traffic specifications.

For actual data transfer, DSI uses a consumer-producer scheme on a ring buffer containing packet descriptors. The packet descriptors provide an indirection for data access and allow a flexible use of the shared data buffers.

The specifications of streams in DSI base on jitter-constrained periodic streams (JCS) [3]. JCS allow to esti-

mate the resources needed for a given stream, e.g. buffer capacity. On stream creation, DSI uses these estimates when establishing the shared data buffers. If the communication peers behave conforming to their specification, no buffer shortage and no data loss occurs.

In cases of resource shortages, the communication peers cannot always meet their traffic specification. To cope with this at the sender, DSI adds timestamp information to the transferred data packets. To cope with resource shortage at the receiver, DSI limits the validity of data by time. This means, the data packets produced at the sender will expire after a certain time, even if they were not consumed by the receiver. For both techniques, DSI uses virtual time, which is assigned to and stored together with each data packet. The virtual time corresponds to the position of the data in the entire stream. The mapping of virtual time to real-time is the responsibility of the communication partners.

A problem arises when the expiration of data must be enforced. It must not impose any blocking, but the sender must know for sure, that the receiver will not continue to access old data anymore. Thus sending a message to the receiver and waiting for an answer is not an option. To enforce the expiration of data DSI uses virtual memory techniques. For this, the sender can request retraction of shared memory pages from the receiver. When the receiver noticed the retraction, it requests re-establishing the memory mapping. This allows an immediate notification without blocking.

References

- [1] J.-Y. Chung, J. W. S. Liu, and K.-J. Lin. Scheduling periodic jobs that allow imprecise results. *IEEE Transactions on Computers*, 39(9):1156–1173, Sept. 1990.
- [2] P. Druschel and L. L. Peterson. Fbufs: A high-bandwidth cross-domain transfer facility. In *14th ACM Symposium on Operating System Principles (SOSP)*, pages 189–202, Asheville, NC, Dec. 1993.
- [3] C.-J. Hamann. On the quantitative specification of jitter constrained periodic streams. In *MASCOTS*, Haifa, Israel, Jan. 1997.
- [4] H. Härtig, L. Reuther, J. Wolter, M. Borriß, and T. Paul. Cooperating resource managers. In *Fifth IEEE Real-Time Technology and Applications Symposium (RTAS)*, Vancouver, Canada, June 1999.
- [5] F. W. Miller, P. Keleher, and S. K. Tripathi. General data streaming. In *19th IEEE Real-Time Systems Symposium (RTSS)*, Madrid, Spain, Dec. 1998.
- [6] V. S. Pai, P. Druschel, and W. Zwanenpoel. IO-Lite: A unified I/O buffering and caching system. *ACM Transactions on Computer Systems*, 18(1):37–66, Feb. 2000.