

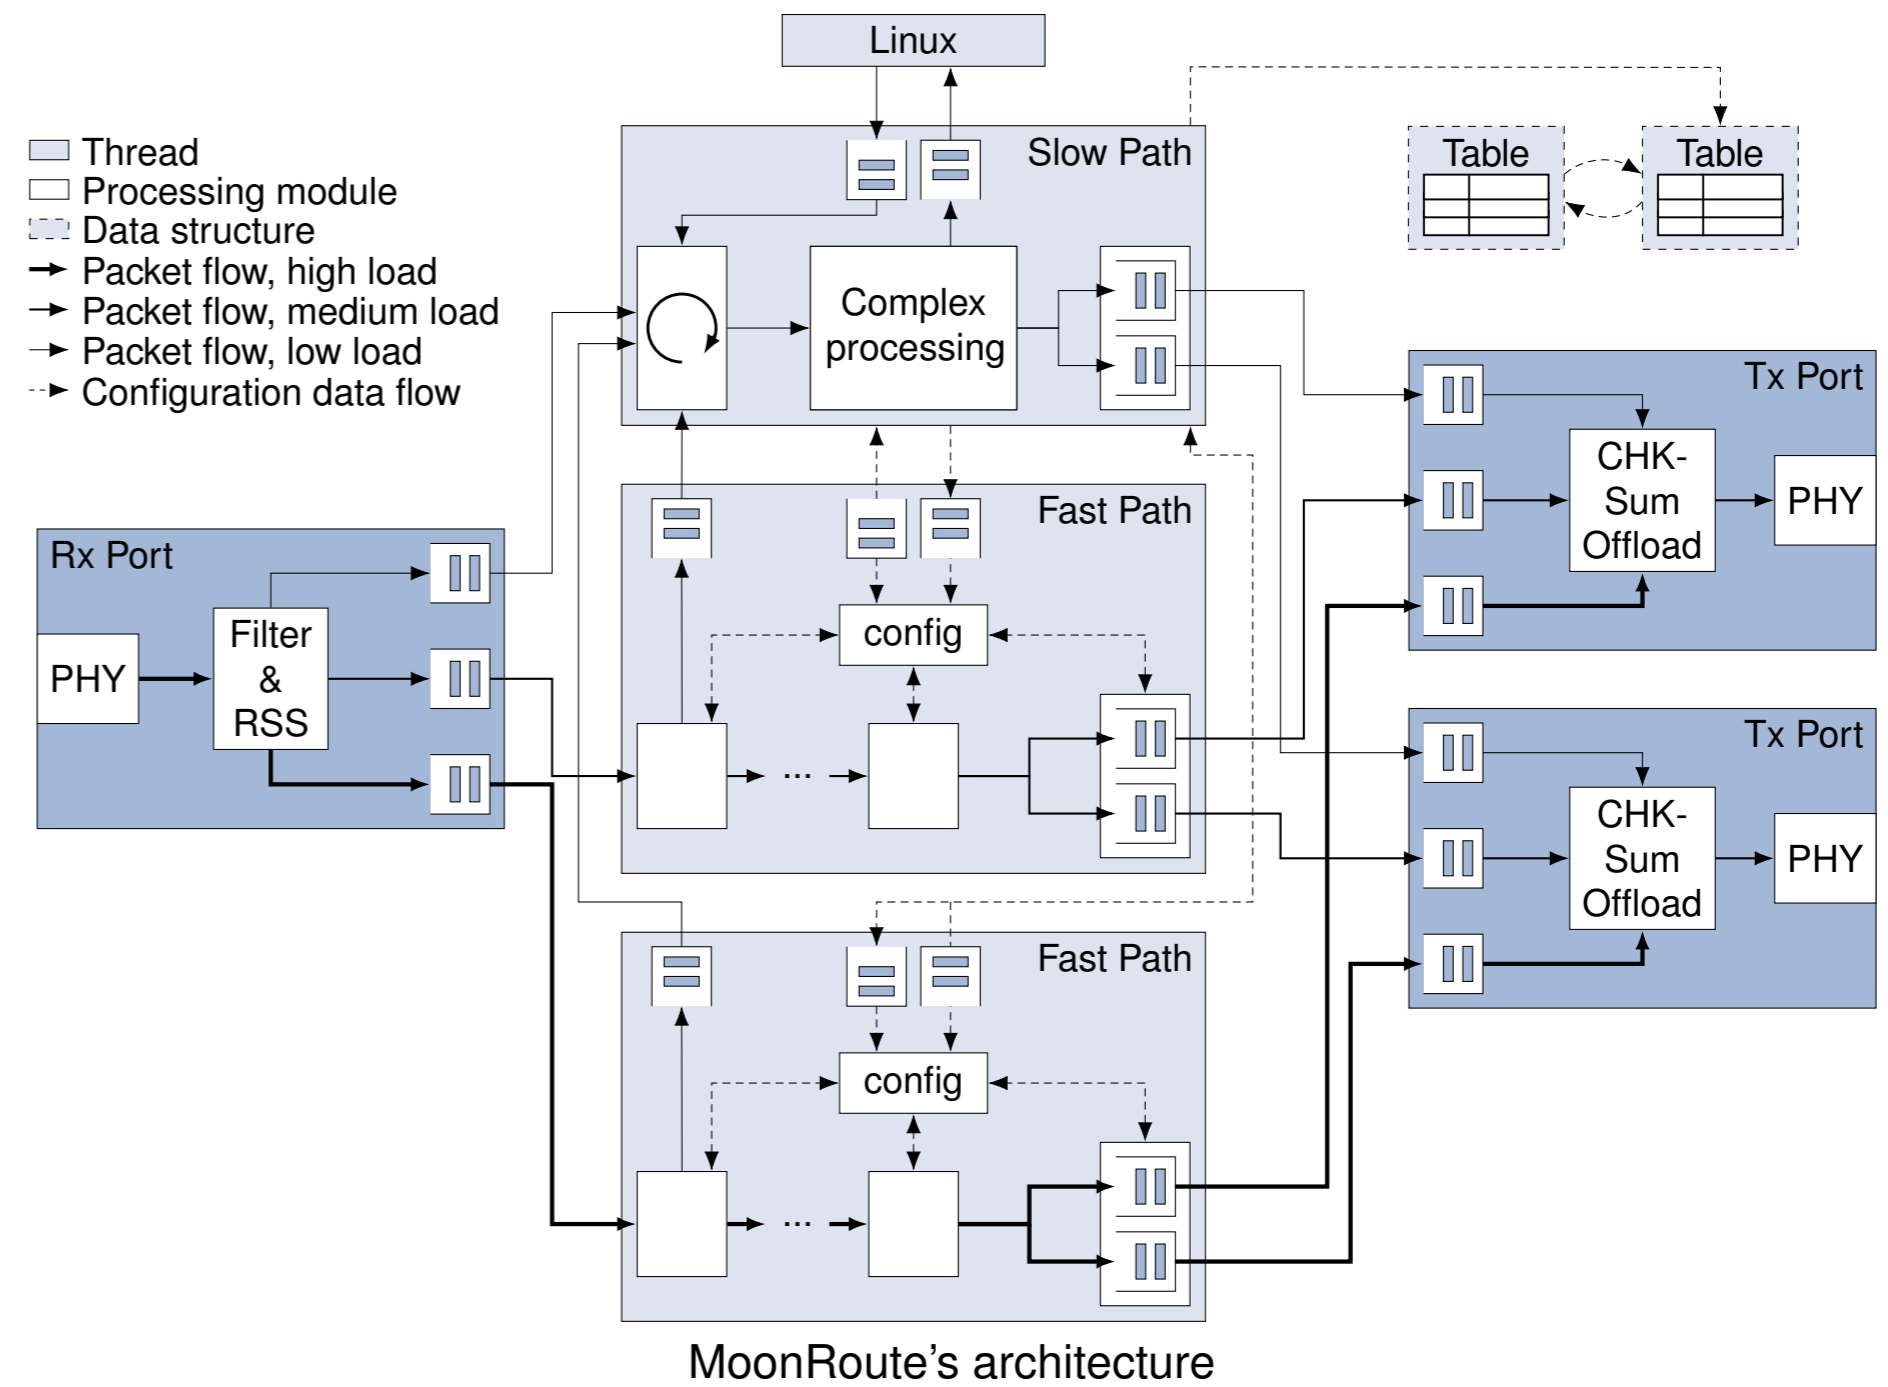
Building Fast but Flexible Software Routers

MoonRoute a high-performance software router build on top of libmoon/DPDK

Features & Architecture

MoonRoute is a flexible framework for implementing high-performance software routers based on libmoon/DPDK [2].

- **Flexibility:** possibility to add/change Lua script code (almost) anywhere in the packet processing queue
- **Code reuse:** existing C libraries can be mixed with Lua script code using the foreign function interface of LuaJit
- **Performance:** a single server is able to saturate multiple 10 GbE ports with minimum-sized packets
- **Scalability:** the framework incorporates the benefits and challenges modern multi/many-core architectures employing RSS, multi-queue & hardware filters
- **Multiple processing paths:** Packets can be processed in a high-performance fast path, a feature-rich slow path, or handled by the host OS minimizing the overall implementation effort



Design principles

MoonRoute's design is optimized for **maximum scalability**:

- Lock free data structures are used (queues)
- Shared data is kept to a minimum
- Immutable data structures are used for sharing

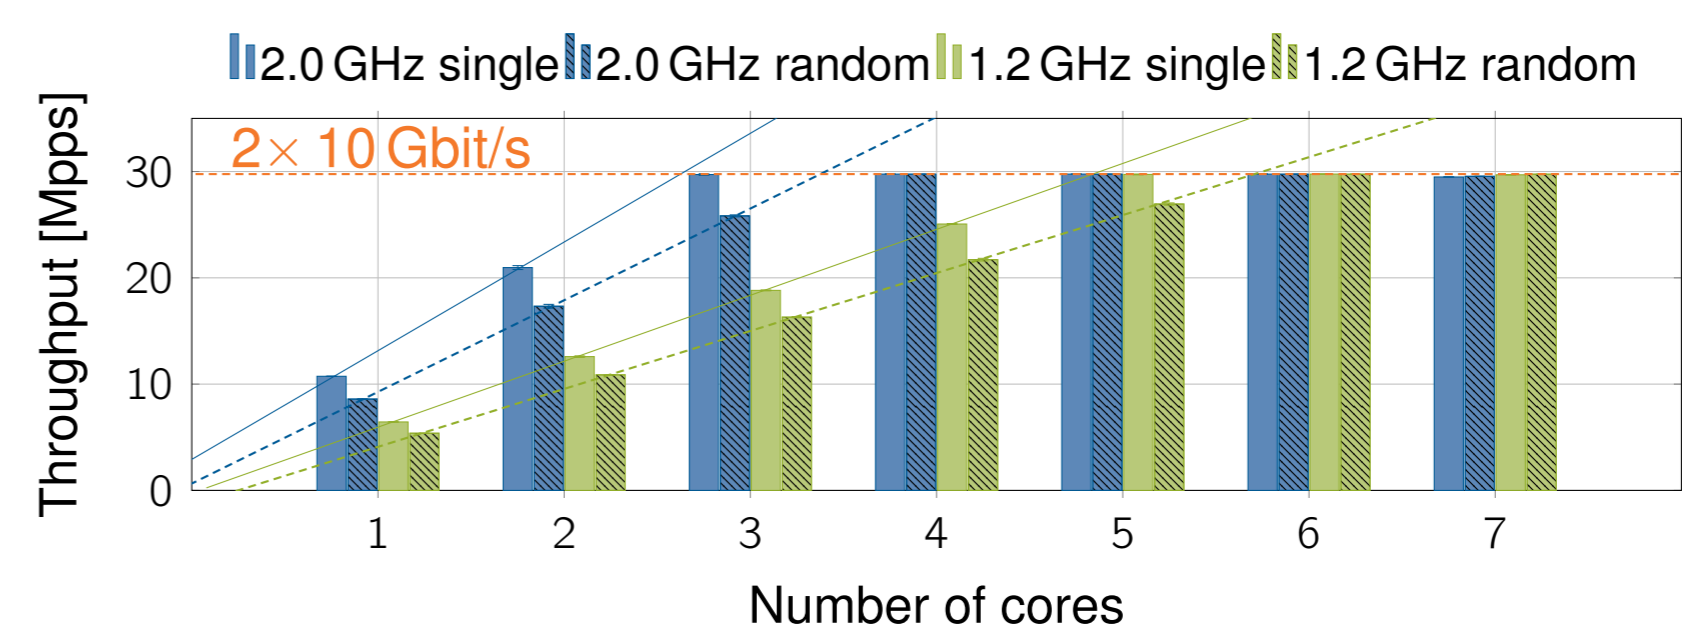
MoonRoute's routing table is **easily replaceable**:

- Currently DPDK's routing table is used (DIR-24-8)
- **Double buffering** is used for routing table updates:
 - One read-only routing table is used for routing
 - Another routing table is updated
 - Routing tables are switched on update

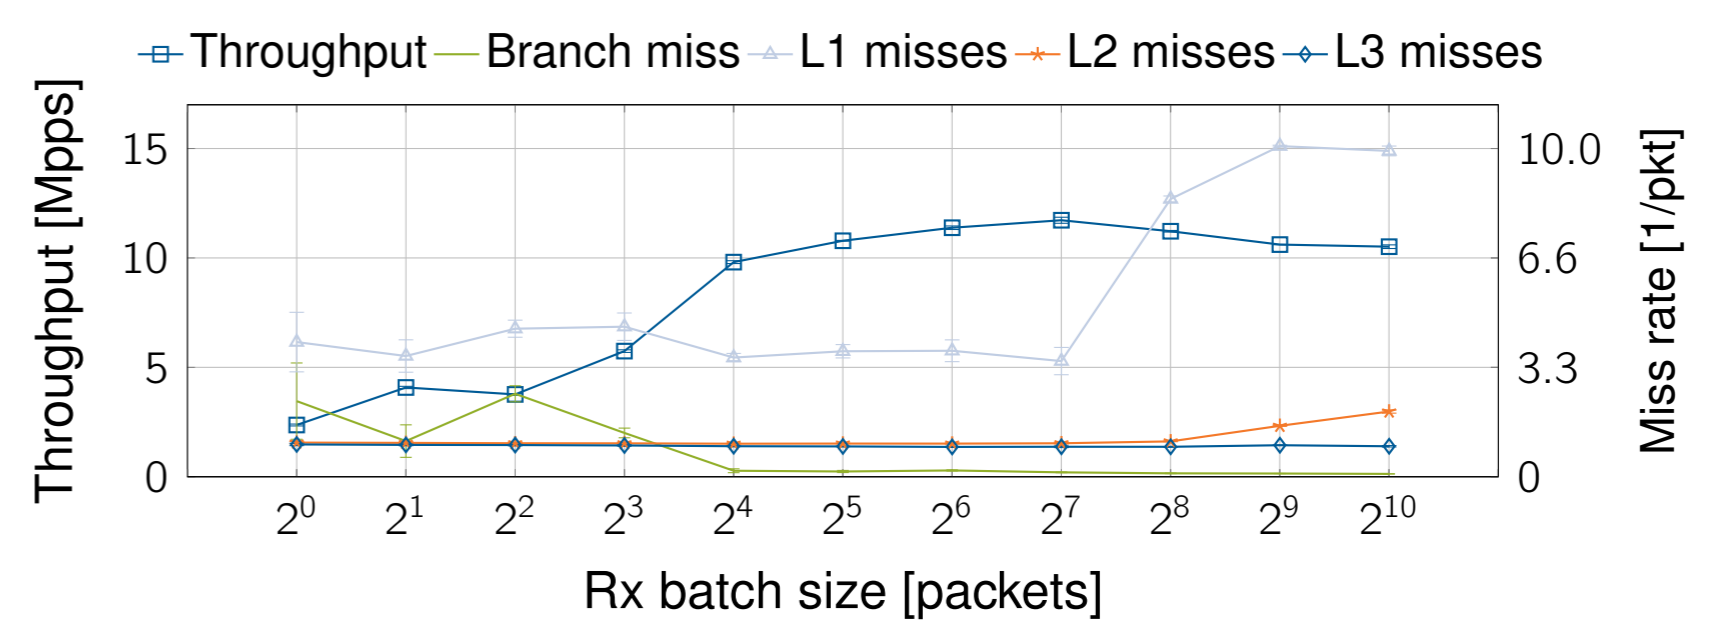
MoonRoute supports **multiple batching strategies**:

- Batching on high-performance paths:
 - Flagging avoids expensive rebatching
- Creating new batches for low-performance paths:
 - Reduces load on slow path

Scalability and Influence of Batching



- MoonRoute's sample router scales perfectly, using a single/randomly chosen flows for fully-populated routing table



- Optimal performance for batch size of 128 packets

Comparison

- MoonRoute achieves superior performance to a number of different software routers (tested with a single routing entry on the same hardware):


Router	Mpps	Relative
MoonRoute	14.6	100%
FastClick (DPDK 2.2) [1]	10.4	72%
Click (DPDK 2.2) [3]	4.3	29%
Linux 3.7	1.5	10%

Availability

- A full paper, the MoonRoute framework, and a sample router implementation are available at:



Full paper



MoonRoute repository

[1] T. Barbette, C. Soldani, and L. Mathy. Fast Userspace Packet Processing. In *Proceedings of the Eleventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, pages 5–16. IEEE Computer Society, 2015.
 [2] P. Emmerich, S. Gallenmüller, D. Raumer, F. Wohlfart, and G. Carle. MoonGen: A Scriptable High-Speed Packet Generator. In *Internet Measurement Conference 2015 (IMC'15)*, Tokyo, Japan, Oct. 2015.
 [3] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click Modular Router. *ACM Transactions on Computer Systems (TOCS)*, 18(3):263–297, 2000.