

Userspace RCU Overview

Userspace RCU: New APIs and
Data Structures, 2013 Update

E-mail:
mathieu.desnoyers@efficios.com

> Userspace RCU

- Userspace API and data structures for *scalability, real-time* response, and *speed*,
- API semantic similar to the Linux kernel,
- LGPLv2.1 license,
- Multi-architecture, POSIX OSes,
- Optimized on Linux, fallbacks for other OS.

> URCU 0.7.x APIs

- RCU read-side critical sections, grace periods,
- ***call_rcu()*** helper threads,
- RCU lock-free ***queue***, wait-free ***stack***, RCU ***lists***,
- RCU lock-free automatically resized ***hash table***,
- Atomic operations, memory barriers, reference counting.

> URCU New APIs, Ongoing Work

- URCU 0.8.x
 - *Wait-free concurrent queue*
 - Lock-free stack
- Ongoing development:
 - *RCU Judy Array*
 - git://git.urcu.so/urcu.git branch: `urcu/rcuja-range`
 - Linked lists with multi-direction traversals
 - git://git.urcu.so/urcu.git branch: `rcudq`

> Wait-Free Concurrent Queue

- Wait-free enqueue
 - Extremely fast, also great for real-time
- Queue head and tail 2 separate objects
 - User controls cache-line layout
- Supports ***splice*** operation
 - Constant time $O(1)$ merge of N elements

> RCU Judy Arrays

- Trie: lookup $O(\text{key length})$,
- RCU lookups and traversals,
- Permits ordered key traversals,
- Scalable updates,
- Cache-efficient,
- Reasonably fast updates.

> Ongoing RCU Judy Array Implementation

- Complete, needs testing and review:
 - Add, removal, lookup, traversal, bit-distribution selection for pool internal nodes (densely populated).
- Work in progress:
 - Storing *range*, looking up points and intersections, traversal,

> Questions ?

- Userspace RCU library available at: <http://urcu.so>

- lttng-dev@lists.lttng.org



*Effici*OS

- <http://www.efficios.com>