CSR: Core Surprise Removal in Commodity Operating Systems

Noam Shalev, Eran Harpaz, Hagar Porat, Idit Keidar, Yaron Weinsberg* | Technion, Israel Institute of Technology | *IBM Research, Israel

Motivation

- **Current operating system crash in face of any hardware fault.**

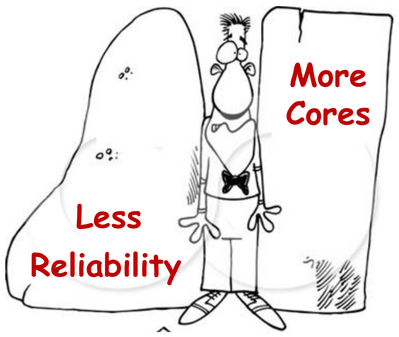

Contributions

- ***Core Surprise Removal*** **strategy.**
- **HTM in kernel code for reliability.**
- **Implement CSR, using HTM, in Linux.**
- **Evaluation on real system.**

Fault Model

- ***Fail-stop*** **model:**

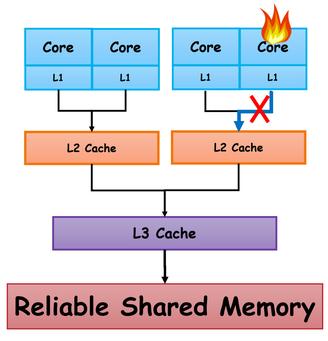

Recovery Strategy

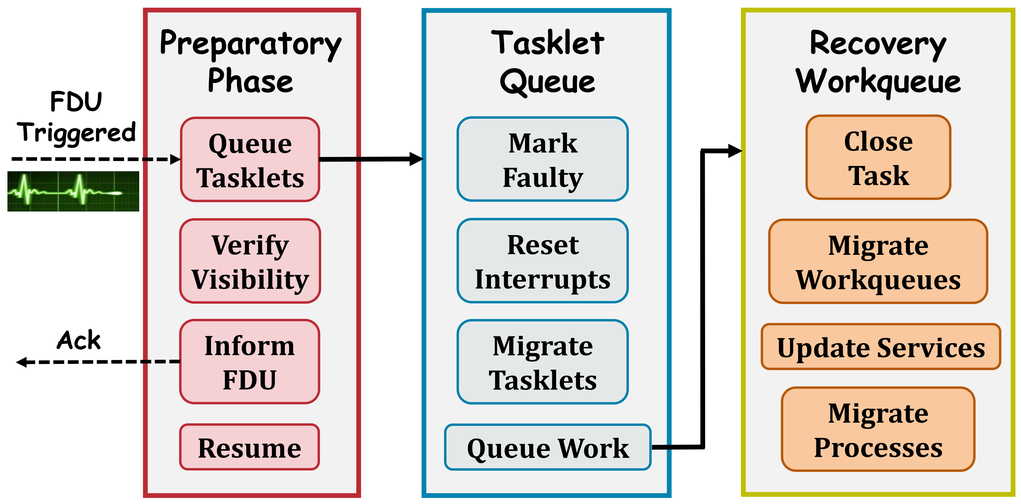

Evaluation on Virtualized Environment

- **User mode: 100% Success**
- **Idle mode: 100% Success**
- **Kernel mode: It's complicated**
  - **Crashes are due to held locks!**

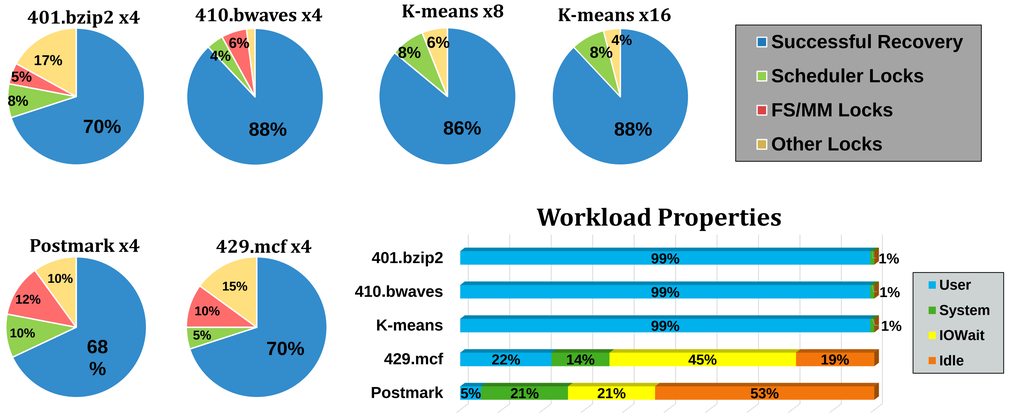

Hardware Transactional Memory

Solution: Replace OS locks for transactions

- **Execute atomically**
- **Does not use locks**

Implementation using Intel TSX®:


```
MY_CRITICAL_SECTION:
if (_xbegin() == _XBEGIN_STARTED) {
   if (raw_spin_is_locked(&rq->lock)){
      _xabort(1);
   }
 /* Critical Section Body */
   _xend();
} else { //fallback
   if (retries < MAX_RETRIES){//retry
      goto MY_CRITICAL_SECTION;
   }
   raw_spin_lock(&rq->lock);
   /* Critical Section Body */
   raw_spin_unlock(&rq->lock);
}
```

Evaluation on a Real System

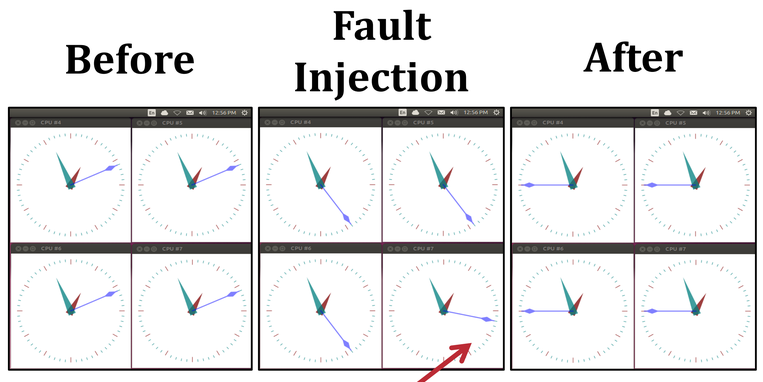

Workload	Commit Rate	Performance Gain	Energy Saving
Idle	100%	-	4%
16-threads	99.9%	0%	1%
32-threads	99.9%	3%	3%
64-threads	99.8%	4%	2%

TECHNION Israel Institute of Technology

IBM