

# pktgen RX & Traffic Generation with improved TX stack

Daniel Turull

2014-11-26

# Overview

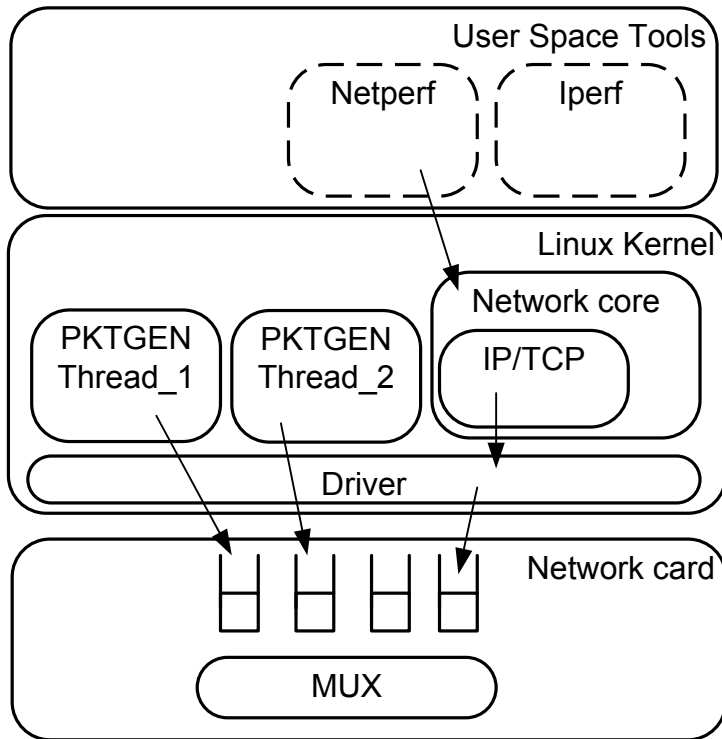
- Introduction
  - Architecture
- Pktgen RX
- Updated results

# Introduction

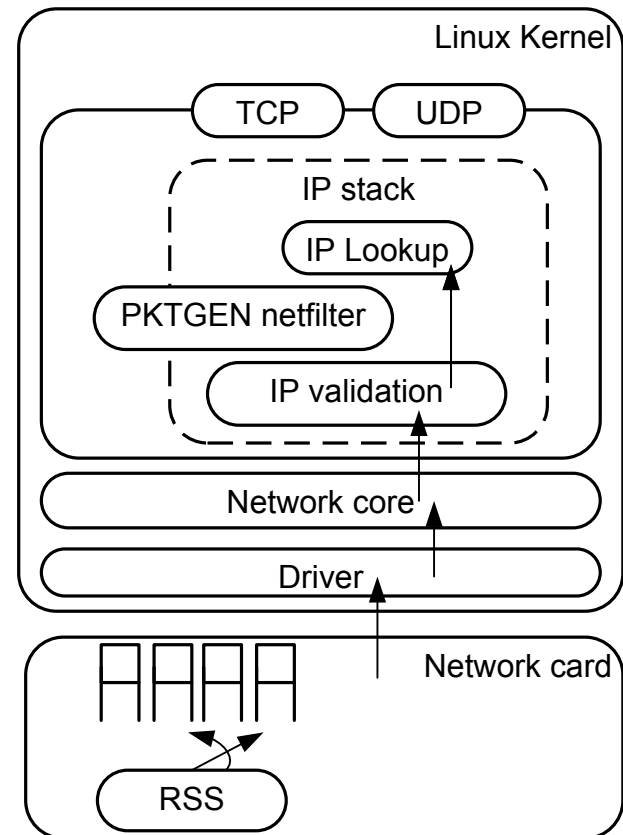
- Packet generator in the Linux Kernel
- Wire speed with small packets
- Out-of-tree receiver patch

# Architecture

## Transmitter



## Receiver



# /proc/net/pktgen/pgrx

- rx [device] to enable the receiver part for a specific device. If it is wrong, all the devices are used.
- rx\_reset: to reset the counters
- rx\_disable: to disable the receiver
- display [human or script]
- statistics [counter, basic, basic6 or time]

# Output basic

```
cat /proc/net/pktgen/pgrx
      RECEPTION STATISTICS
      PER-CPU Stats
CPU 0: Rx packets: 0      Rx bytes: 0
CPU 1: Rx packets: 0      Rx bytes: 0
CPU 2: Rx packets: 9385641      Rx bytes: 563138460
      Work time 13483820 us
      Rate: 696066pps 334Mb/sec (334112119bps)
CPU 3: Rx packets: 9490524      Rx bytes: 569431440
      Work time 13483381 us
      Rate: 703868pps 337Mb/sec (337856767bps)
CPU 4: Rx packets: 8949667      Rx bytes: 536980020
      Work time 13483488 us
      Rate: 663750pps 318Mb/sec (318600065bps)
CPU 5: Rx packets: 9030598      Rx bytes: 541835880
      Work time 13483509 us
      Rate: 669751pps 321Mb/sec (321480635bps)
CPU 6: Rx packets: 9085660      Rx bytes: 545139600
      Work time 13483696 us
      Rate: 673825pps 323Mb/sec (323436304bps)
CPU 7: Rx packets: 9009976      Rx bytes: 540598560
      Work time 13483359 us
      Rate: 668229pps 320Mb/sec (320750080bps)

      Global Statistics
Packets Rx: 54952066      Bytes Rx: 3297123960
Work time 13483828 us
4075405pps 1956Mb/sec (1956194611bps)
```

---

# Output time

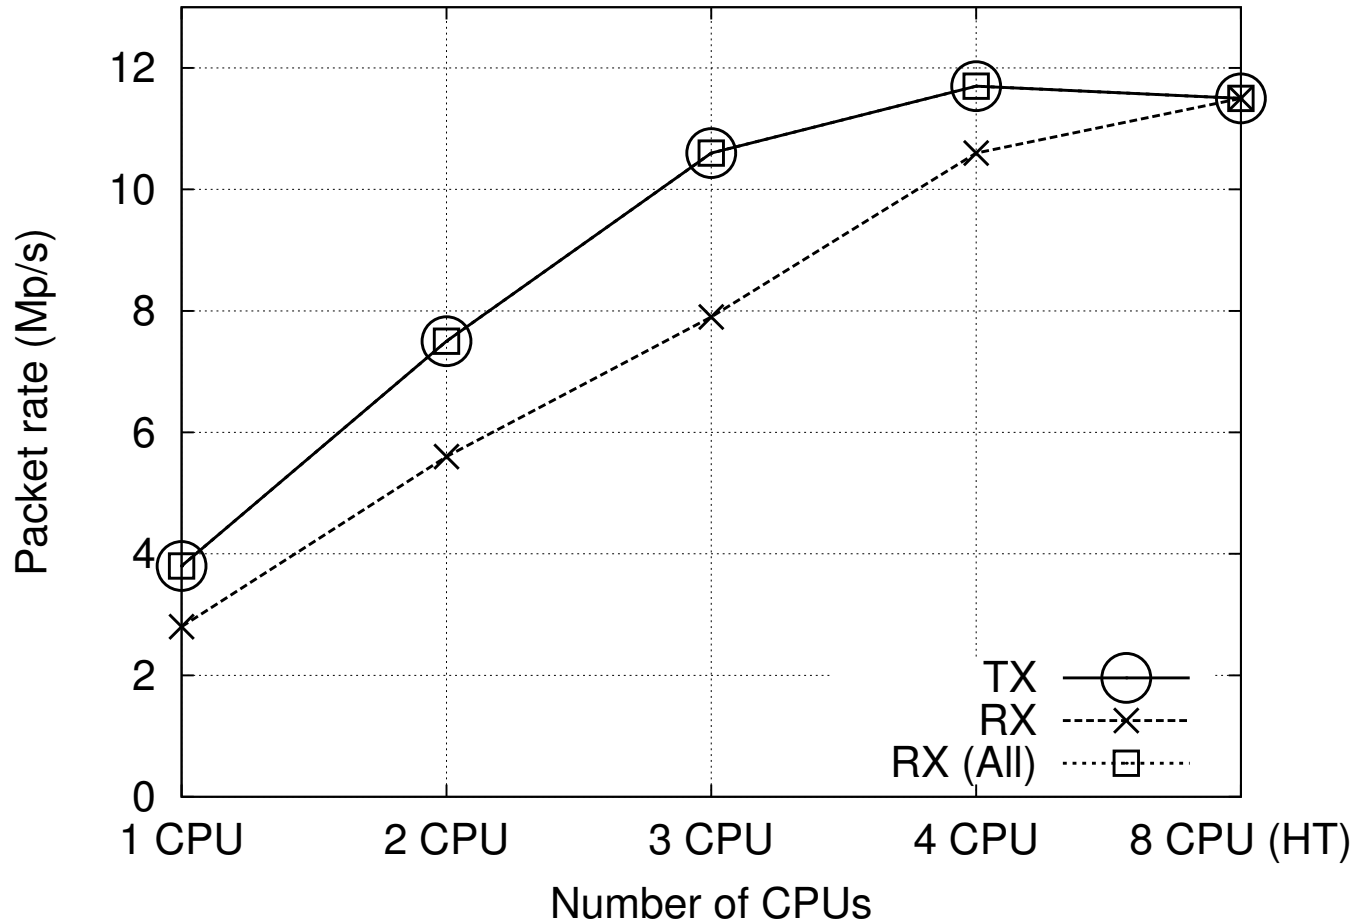
```
RECEPTION STATISTICS
PER-CPU Stats
CPU 0: Rx packets: 0      Rx bytes: 0
CPU 1: Rx packets: 0      Rx bytes: 0
CPU 2: Rx packets: 10000000      Rx bytes: 600000000
Work time 6719014 us
Rate: 1488313pps 714Mb/sec (714390534bps)
Inter-arrival
    Average: 671 ns Variance 381245 ns2
    Max: 10108 ns Min:: 404 ns
    Samples: 9999999
Jitter
    Average: 360 ns Variance 665018 ns2
    Max: 9404 ns Min:: 0 ns
    Samples: 9999998
Latency
    Average: 470578 ns Variance 216731083 ns2
    Max: 545217 ns Min:: 60345 ns
    Samples: 1000000
```

# Considerations

- Fix CPU frequency
- Disable CPU C states
- Interrupt affinity, counters per cpu

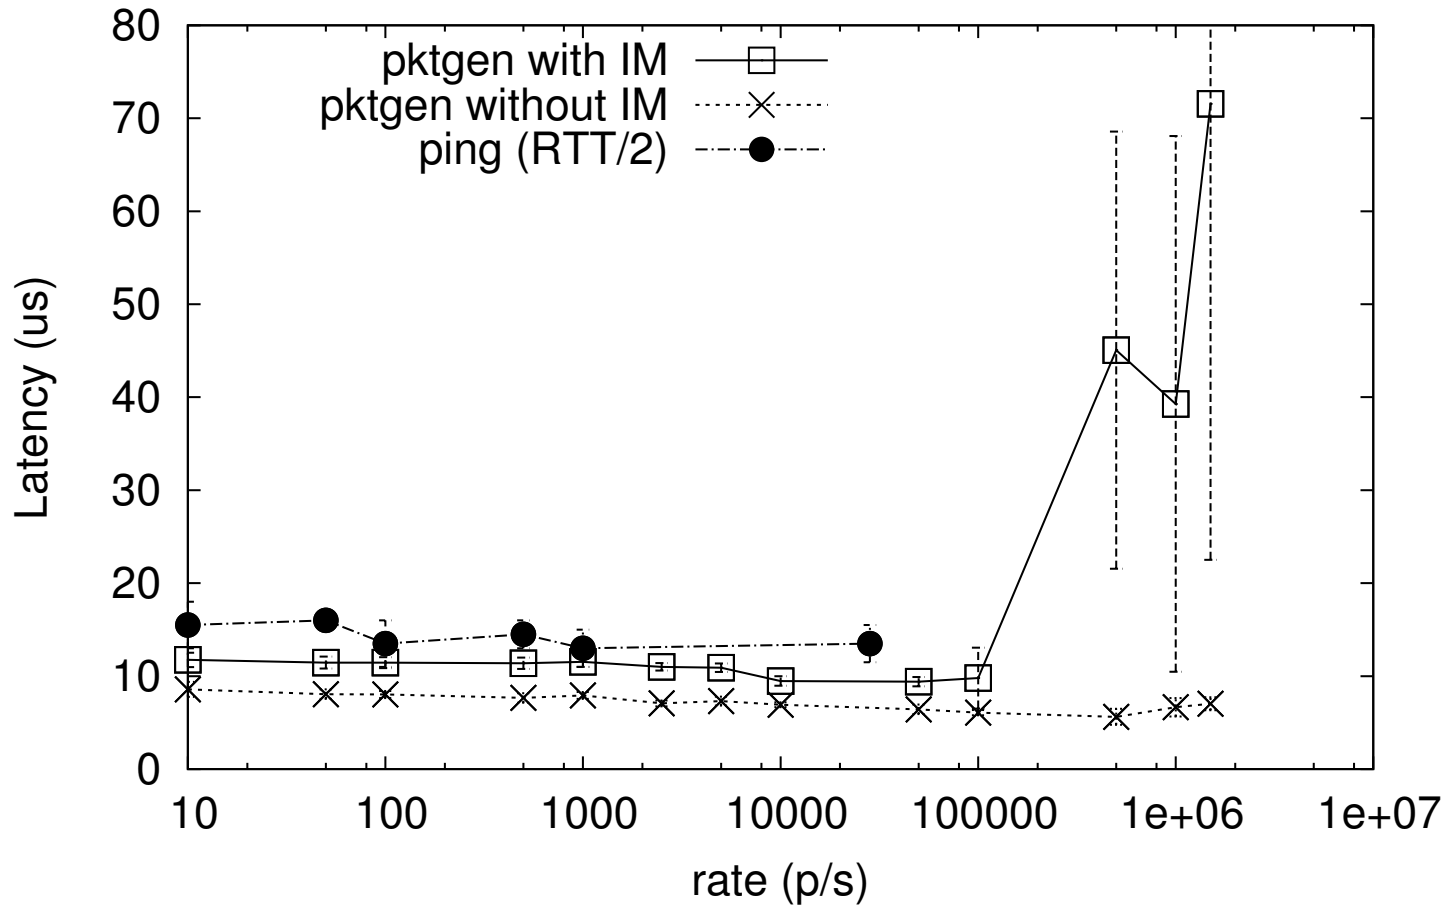


# RX performance



# Latency

Statistics Latency (Packet size: 64Bytes)



# Max Throughput performance with 64 Byte packet

Tool	TX Rate (kp/s)	RX Rate (kp/s)	Kernel driver
Iperf3	294	294	original
Netperf	713	713	Original
Pktgen *	3700	2800	Original
Netmap	14 880	11 200	customized
DPDK	14 880	14 880	customized

\*old version. See next slide

# Updated results (Sep 2014)

- 10Gbit/s full TX wirespeed smallest packet size on a single CPU core (14.8Mpps)
- The SKB xmit\_more API
- Challenge: bulk without added latency
- need to adjust cleanup interval
  - `ethtool -C ethX rx-usecs 30`
- Source (Jesper Dangaard):  
<http://netoptimizer.blogspot.se/2014/10/unlocked-10gbps-tx-wirespeed-smallest.html>

# Sources

- <https://github.com/danieltt/pktgen>
- In Bifrost Distrution Kernel
- Old material from my Master thesis
  - <https://people.kth.se/~danieltt/pktgen/>

Thanks