

TCP over SONIC

Xinxiang Lao
Xuke Fang

SoNIC

Software-defined Network Interface Card

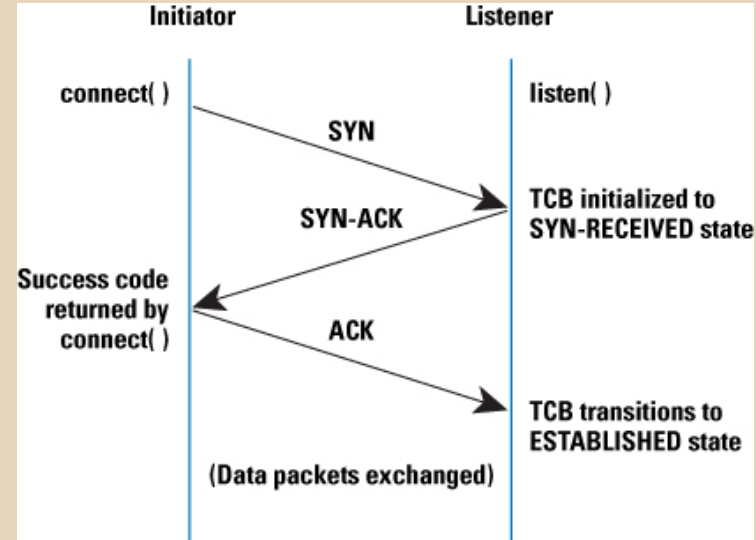
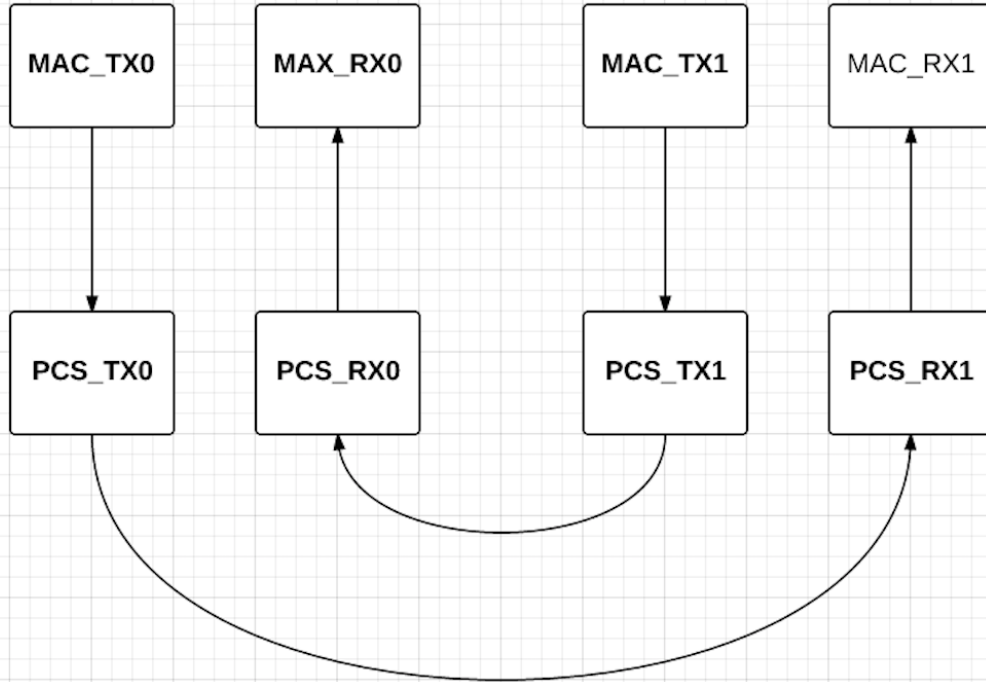
Provides access to the physical and data link layers in software by implementing them in software.

TCP over SoNIC

Design, Implement, Test
Improve performance

Motivation: Reliable data transmission over
SoNIC

Architecture



TCP Throughput Limitation

$$\text{Throughput} \leq \frac{\text{MSS}}{\text{RTT} \sqrt{P_{\text{loss}}}}$$


where MSS is the maximum segment size and P_{loss} is the probability of packet loss. If packet loss is so rare that the TCP window becomes regularly fully extended, this formula doesn't apply.

Even if there is no packet loss in the network, window size can limit throughput.

$$\text{Throughput} \leq \frac{\text{RWIN}}{\text{RTT}}$$

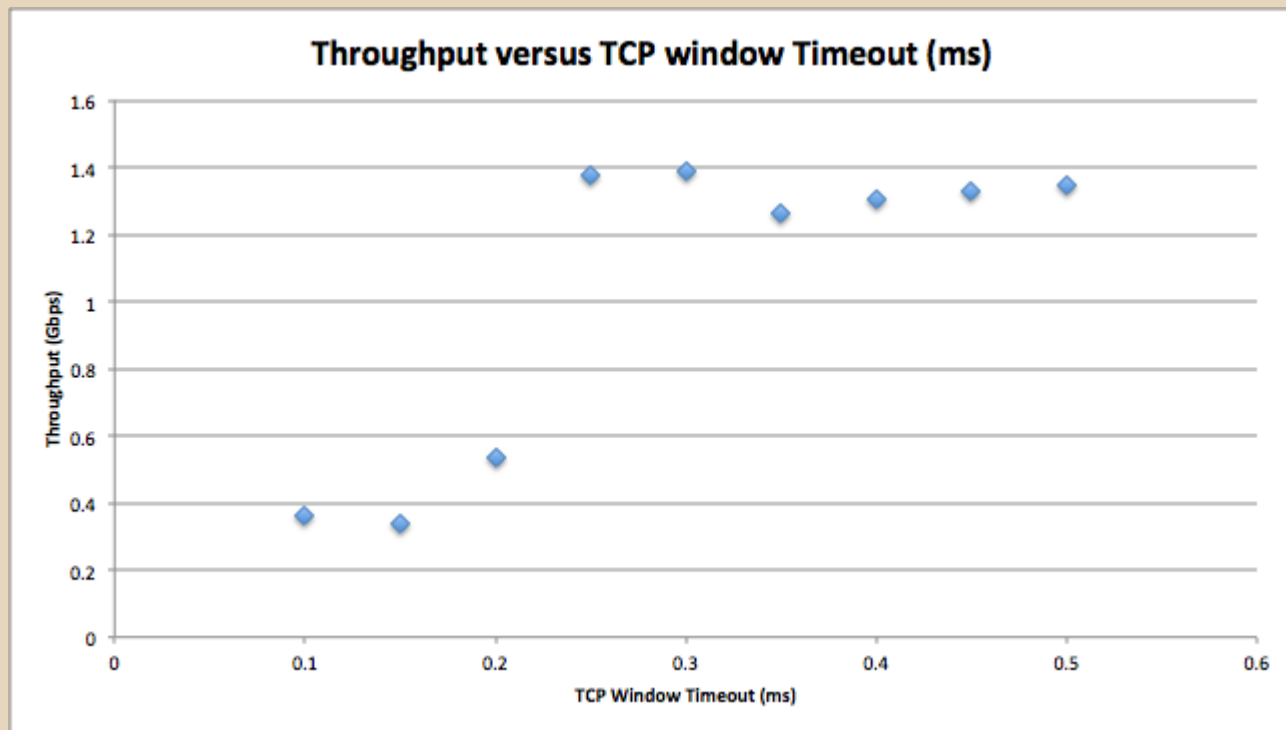
where RWIN is the TCP Receive Window and RTT is the round-trip time for the path.

Evaluation



Configuration[Timeout: 0.2 ms; 100000 pkts;]

Evaluation



Evaluation

```
SONIC: [sonic_print_port_stat] ---Port [0] stats---
SONIC: [sonic_print_pcs_stat] [tx_pcs0] thru= 14.303579 pkts= 1003568 blks= 1950283646 err_states= 0 err_blks= 0 fifo= 0 dma= 0
SONIC: [sonic_print_mac_stat] [tx_mac0] thru= 1.361825 pkts= 1003569 bytes= 1531446294 err_crc= 0 err_len= 0 fifo= 0
SONIC: [sonic_print_pcs_stat] [rx_pcs0] thru= 13.599506 pkts= 963315 blks= 1853897216 err_states= 0 err_blks= 0 fifo= 0 dma= 0
SONIC: [sonic_print_mac_stat] [rx_mac0] thru= 1.306807 pkts= 963315 bytes= 1470018690 err_crc= 0 err_len= 0 fifo= 0
SONIC: [sonic_print_app_stat] [app0] thru=-nan time= 0 pkts= 0 bytes= 0 fifo= 0
SONIC: [sonic_print_port_stat] ---Port [1] stats---
SONIC: [sonic_print_pcs_stat] [tx_pcs1] thru= 13.596026 pkts= 963315 blks= 1853913722 err_states= 0 err_blks= 0 fifo= 0 dma= 0
SONIC: [sonic_print_mac_stat] [tx_mac1] thru= 1.307224 pkts= 963316 bytes= 1470020216 err_crc= 0 err_len= 0 fifo= 0
SONIC: [sonic_print_pcs_stat] [rx_pcs1] thru= 14.306804 pkts= 1003566 blks= 1950272000 err_states= 0 err_blks= 0 fifo= 0 dma= 0
SONIC: [sonic_print_mac_stat] [rx_mac1] thru= 1.361329 pkts= 1003566 bytes= 1531441716 err_crc= 0 err_len= 0 fifo= 0
SONIC: [sonic_print_app_stat] [app1] thru=-nan time= 0 pkts= 0 bytes= 0 fifo= 0
SONIC: [__sonic_exit]
```

Configuration: [Timeout 0.3 ms; 100000 pkts]

Future Work

Improve performance(Theoretically 10 Gbs)

Measure performance with tcp flows of different characteristics

Question?

Thank You All!