


Parallel Programmable Ethernet Controllers: Performance and Security

Viktor Puš

Brno University of Technology, Faculty of Information Technology
Božetechova 2, 612 00 Brno, CZ
ipus@fit.vutbr.cz



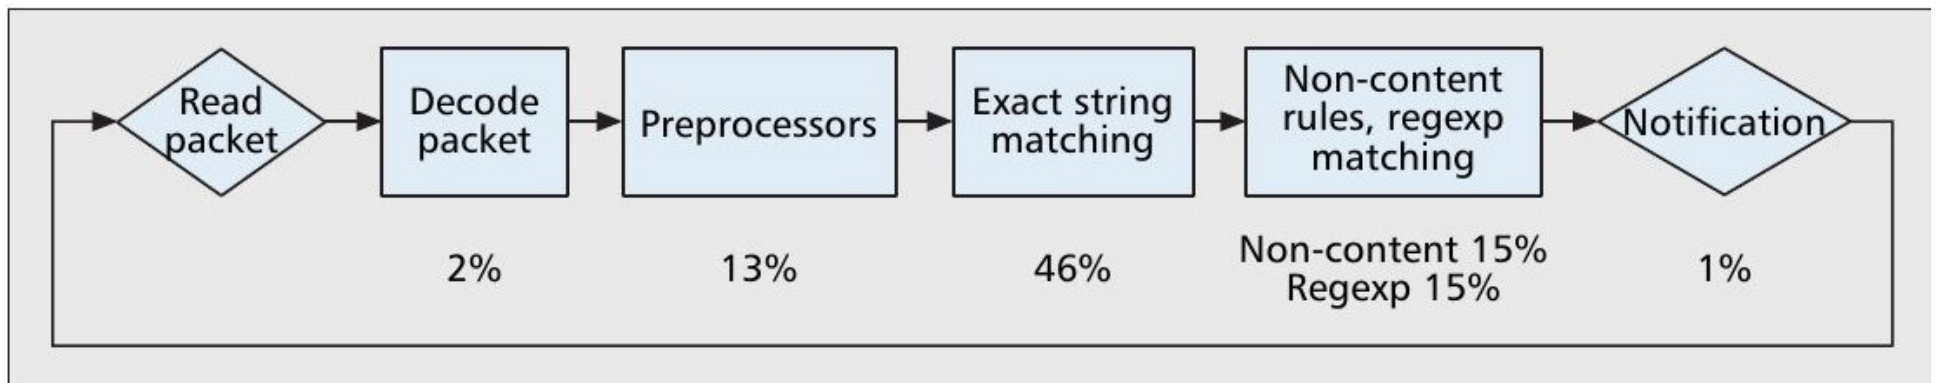
INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ

Motivace 1

- Příjem paketu:
 - 253 instrukcí, 340 B metadat
- Odeslání paketu:
 - 281 instrukcí, 404 B metadat
- 10 Gb/s, pakety 1518 B obousměrně:
 - 435 MIPS, 4.8 Gb/s
- Rychlé procesory → velká spotřeba
- Network procesory nevhodné kvůli latenci DMA

Motivace 2

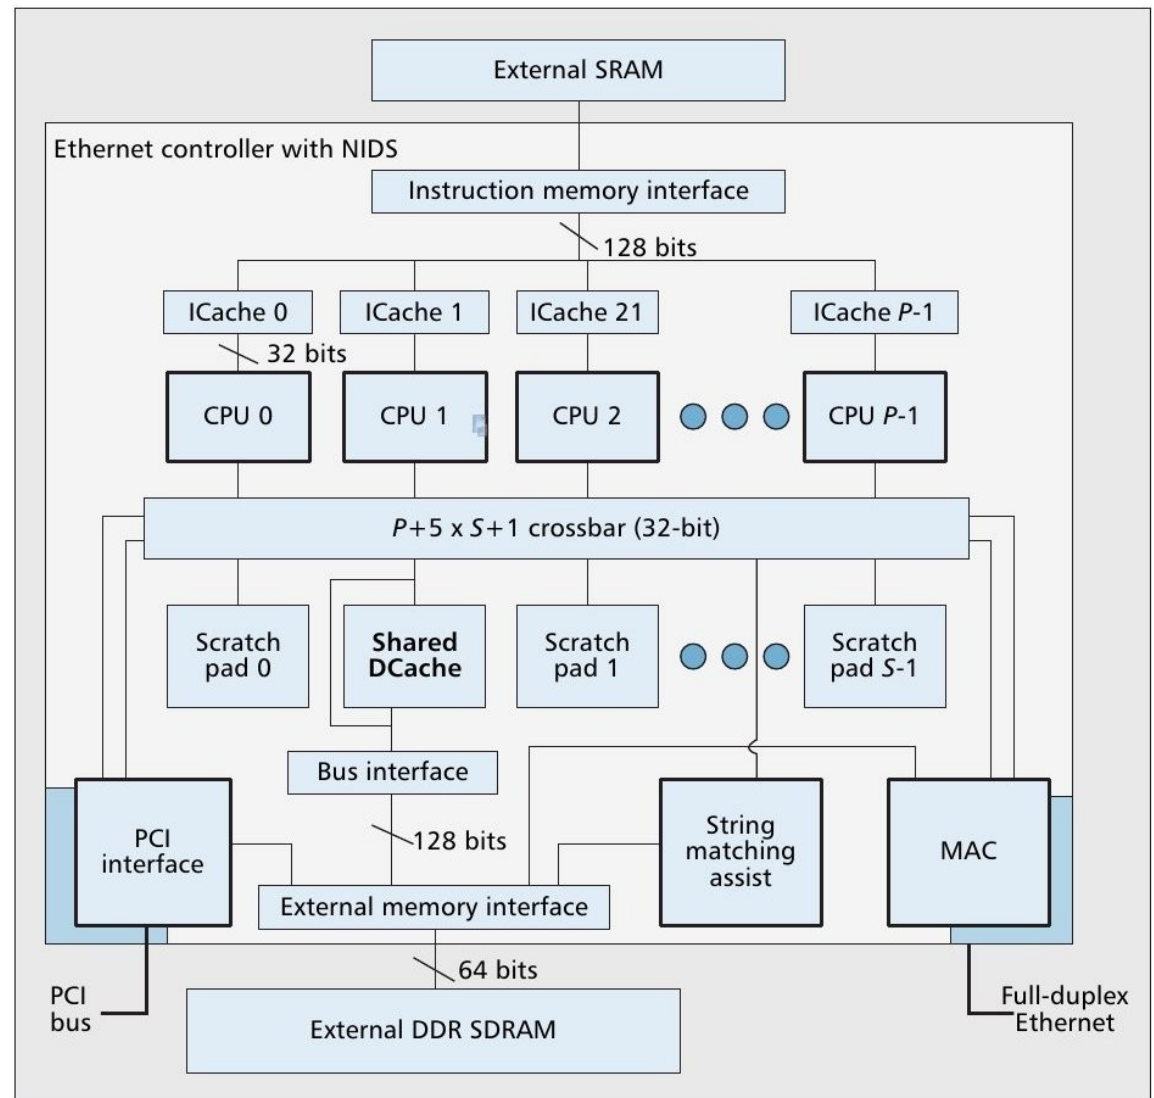
- Snort
 - Nevhodné počítat v SW na serveru
 - Snaha o implementaci na interface



■ Figure 1. *The Snort packet processing loop with percentage of time spent in each phase.*

Návrh

- MIPS
 - Ethernet
 - Non-string Snort
- Memory Transfer Assists
- String Matching Assist
- Paměti
- Crossbar



■ Figure 2. Block diagram of proposed Ethernet controller architecture.

Trocha historie

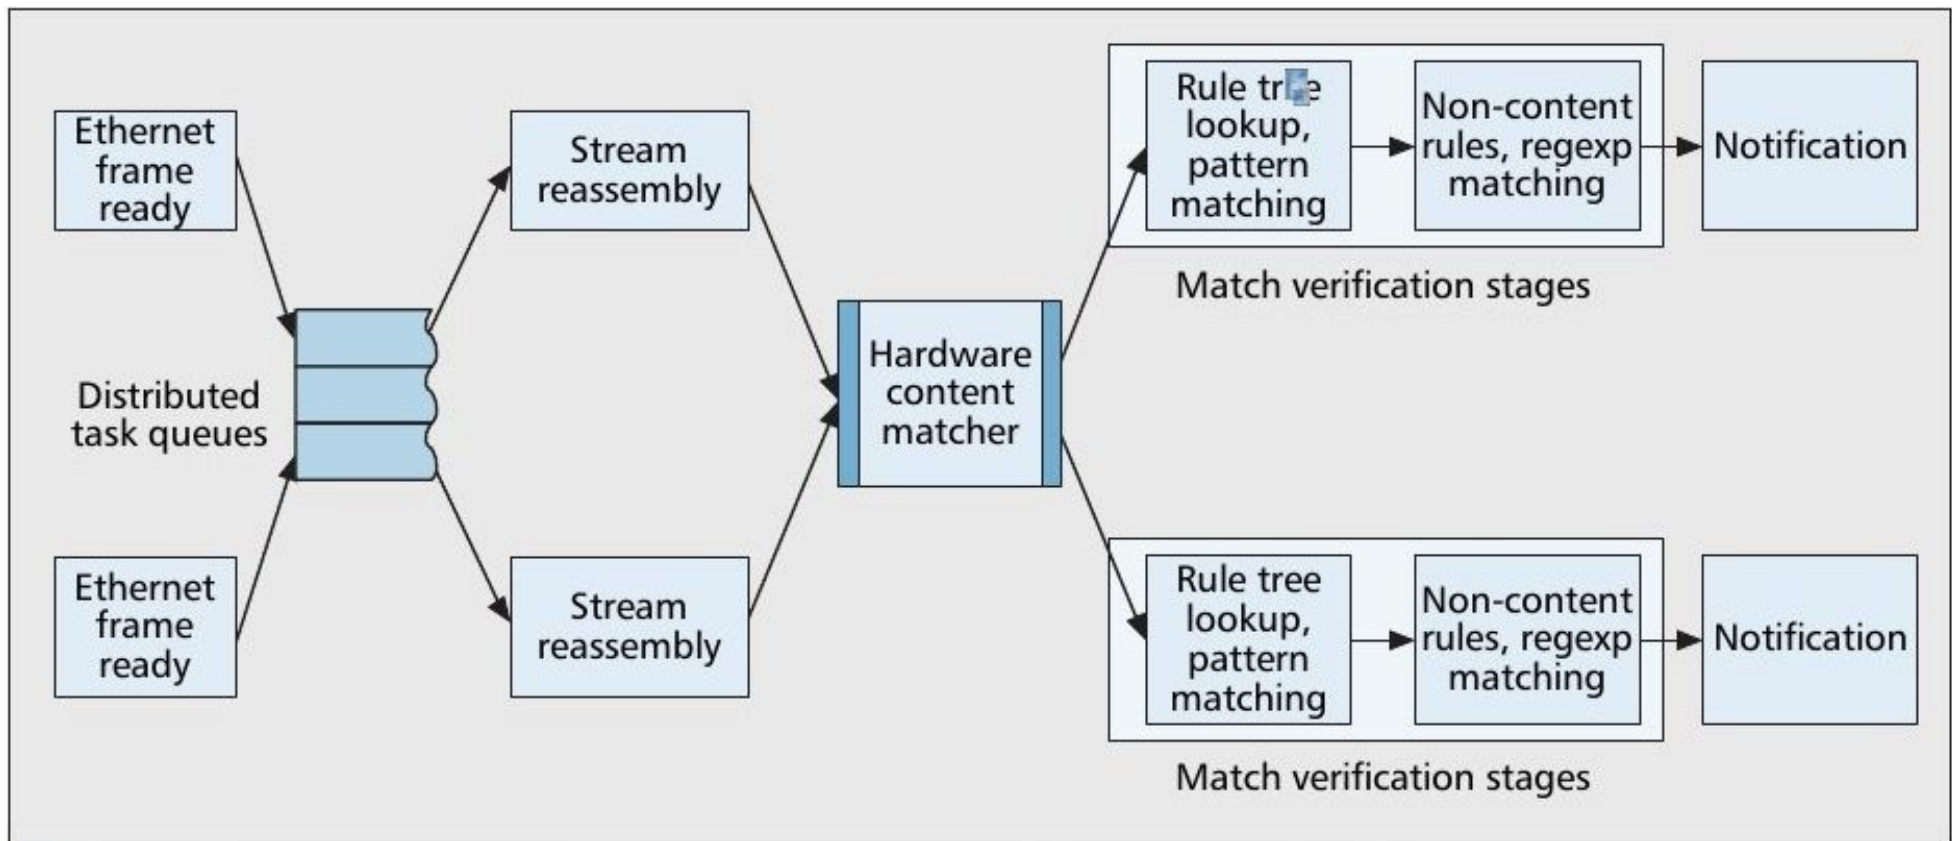
- Alteon Websystems (1996)
 - ACEswitch 180: První tri-speed switch
 - ACEnic: První Jumbo rámce
 - Tigon: První 1G síťovka
 - Dnes Broadcom, 3Com, ...
- Nortel Networks (2000)
- Radware (2009)
- (bez afilací)



Funkce firmware 1

- Vychází z Tigon-II
- Frame-level paralelism
- Work discovery
- Snort:
 - SRC+DST IP hash (hash nalezena, nebo je paket přiřazen CPU s nejkratší frontou)
 - Stream reassembly do scratchpadu
 - String Matching Assist
- In-order dokončení

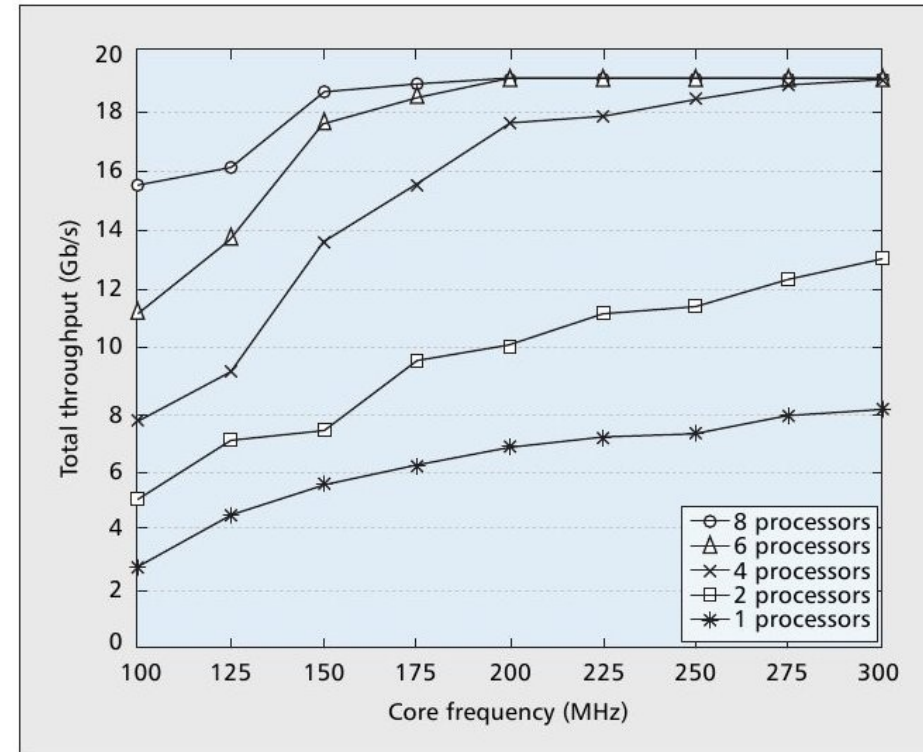
Funkce firmware 2



■ Figure 3. Firmware parallelization strategy used for NIC-embedded Snort.

Výsledky: Bez Snortu


- Spinach – cycle-accurate toolkit for simulating programmable network interfaces
- Graf:
 - 1518 B UDP pakety
 - 10G: 6 jader na 200 Mhz



■ Figure 4. Scaling core frequency and the number of processors for 10 Gb/s Ethernet.

Výsledky: Snort

- Graf: 500 MHz, reálný provoz, reálná pravidla Snortu
- 1G: 6 jader na 300 MHz + String Matching Assist



■ Figure 5. Snort throughput results achieved with 2–8 CPUs for mail and Web rulesets with LL1 and LL2 input traces: a) inspection throughput; b) normalized speedup relative to 2 CPUs.

Diskuze