

Improving TCP's performance using programmable networks

Belma Turkovic and Fernando Kuipers

RoN++ Meeting, December 16th, 2020

TCP complexity is increasing!

- New protocols and congestion control algorithms are continuously being developed

Extensions

- Deploying TCP extensions is difficult
 - Can take more than a decade (e.g., timestamp, SACK)

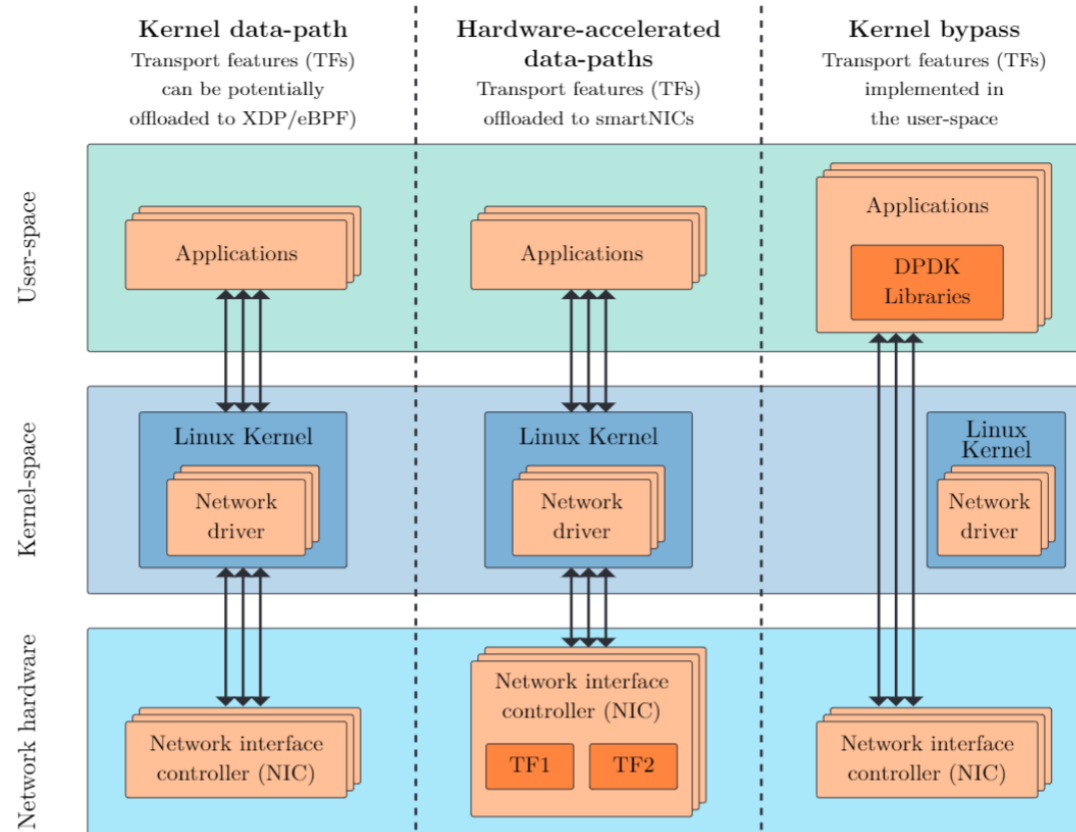
Extensions

- Deploying TCP extensions is difficult
 - Can take more than a decade (e.g., timestamp, SACK)
- Only options:
 - Socket options
 - Modules

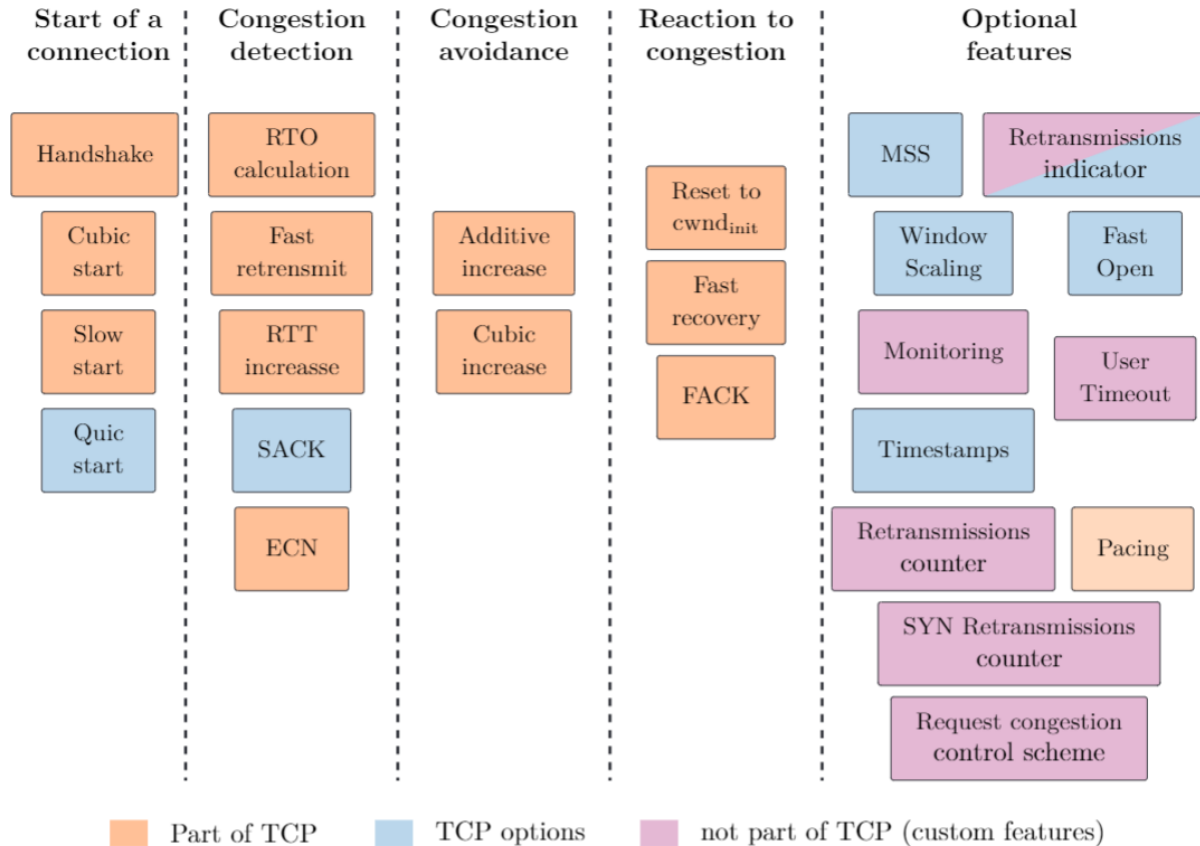
Goal

- Investigate how different data-plane techniques can be used to introduce programmability in the end-hosts

Data-plane techniques



Transport features




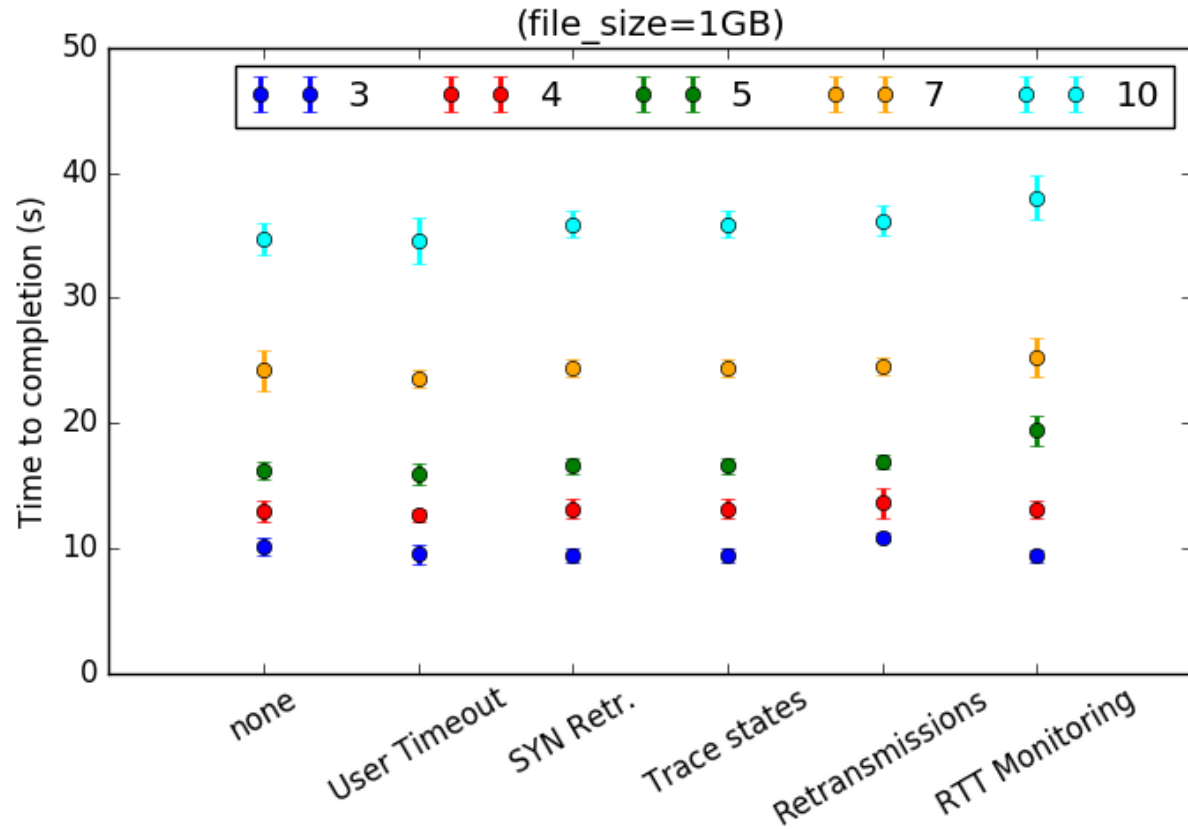
eBPF

- Lightweight VM
- Enables running sandboxed programs in the Linux kernel
- No need to change kernel source-code or load kernel modules

eBPF evolution

- BPF
- eBPF
- Exposure to user-space
- TCP-BPF
- Support for user-defined TCP options







- + Easy to implement
- + Small overhead



- Limited program size
- Limited hooks
- Uses the “slow” kernel
- Needs a custom kernel
(for custom TCP options, TCP-BPF supported since 4.15)

DPDK



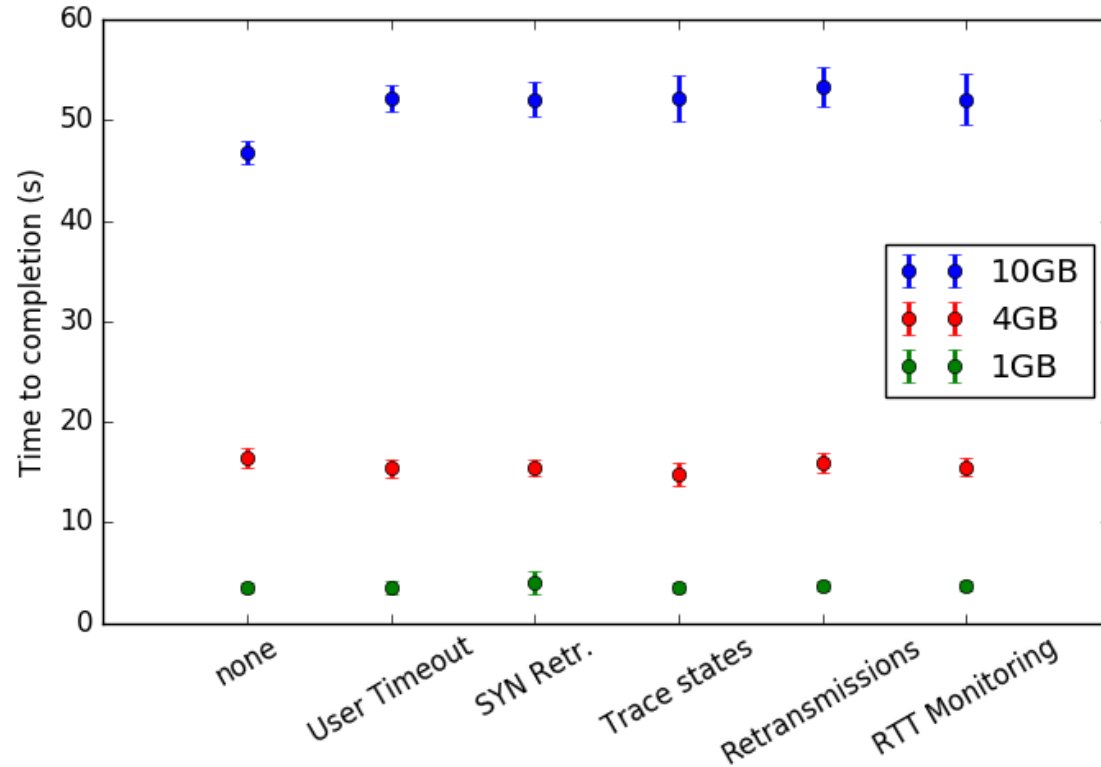
- Kernel bypass technique
- Set of libraries to accelerate packet processing

DPDK




- Tested two different frameworks:
 - F-stack
 - TLDK





- Tested two different frameworks:
 - F-stack
 - TLDK





 Fast performance

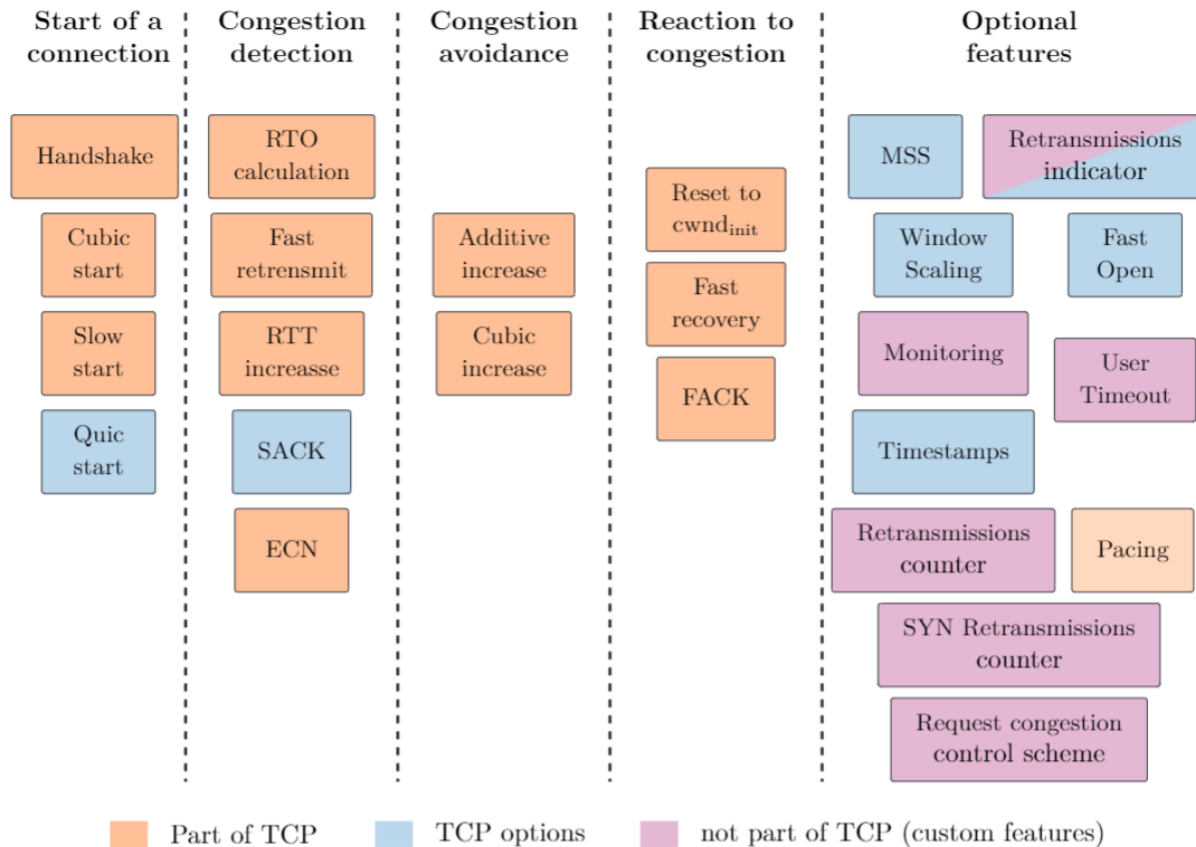


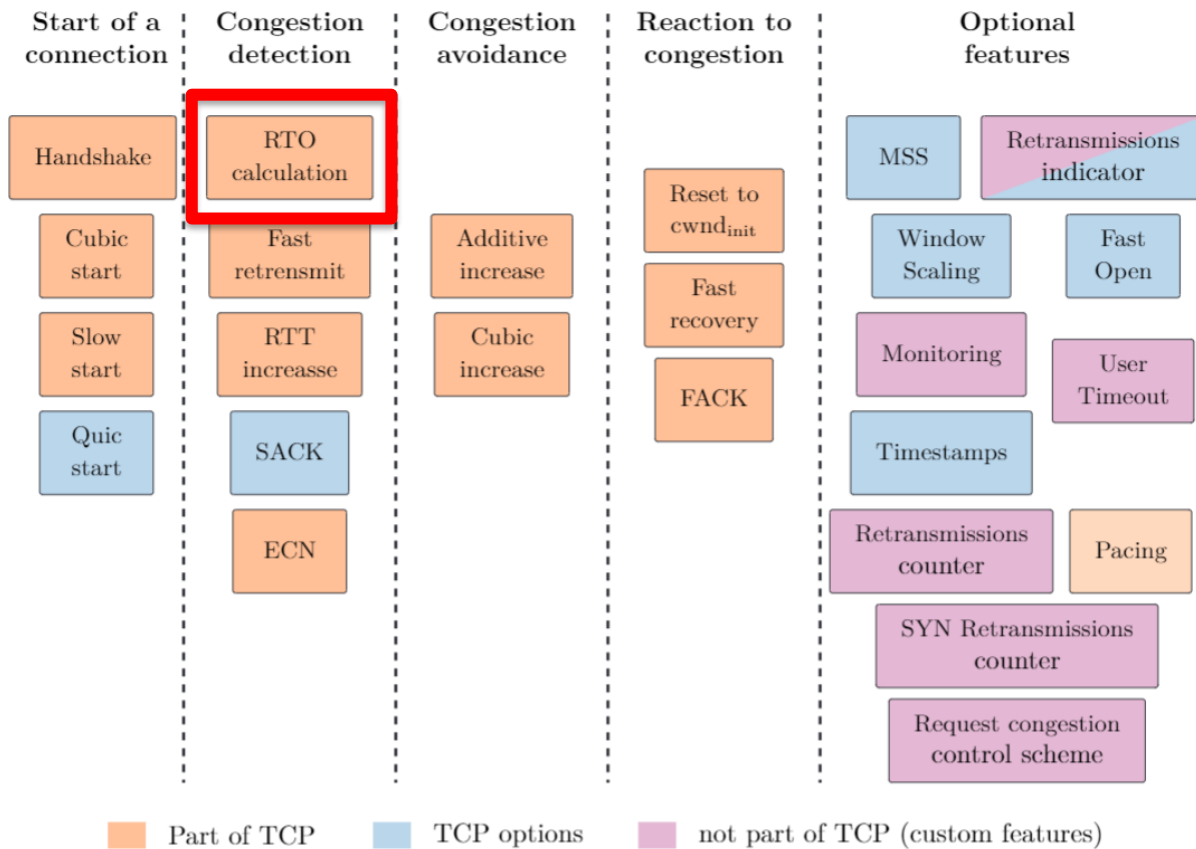
-  Does not have TCP/IP stack out of the box
-  Debugging is difficult (Few tools)
-  Specialized hardware
-  Steep learning curve

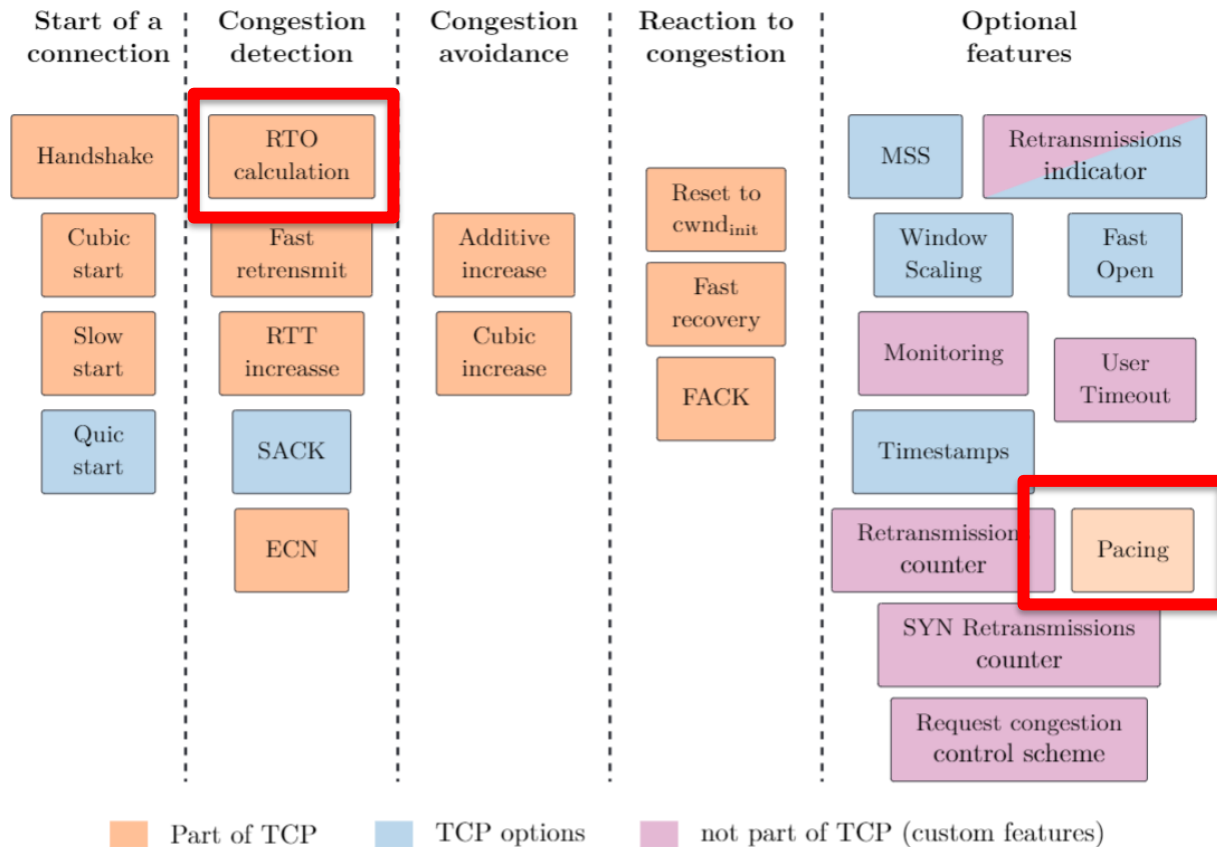


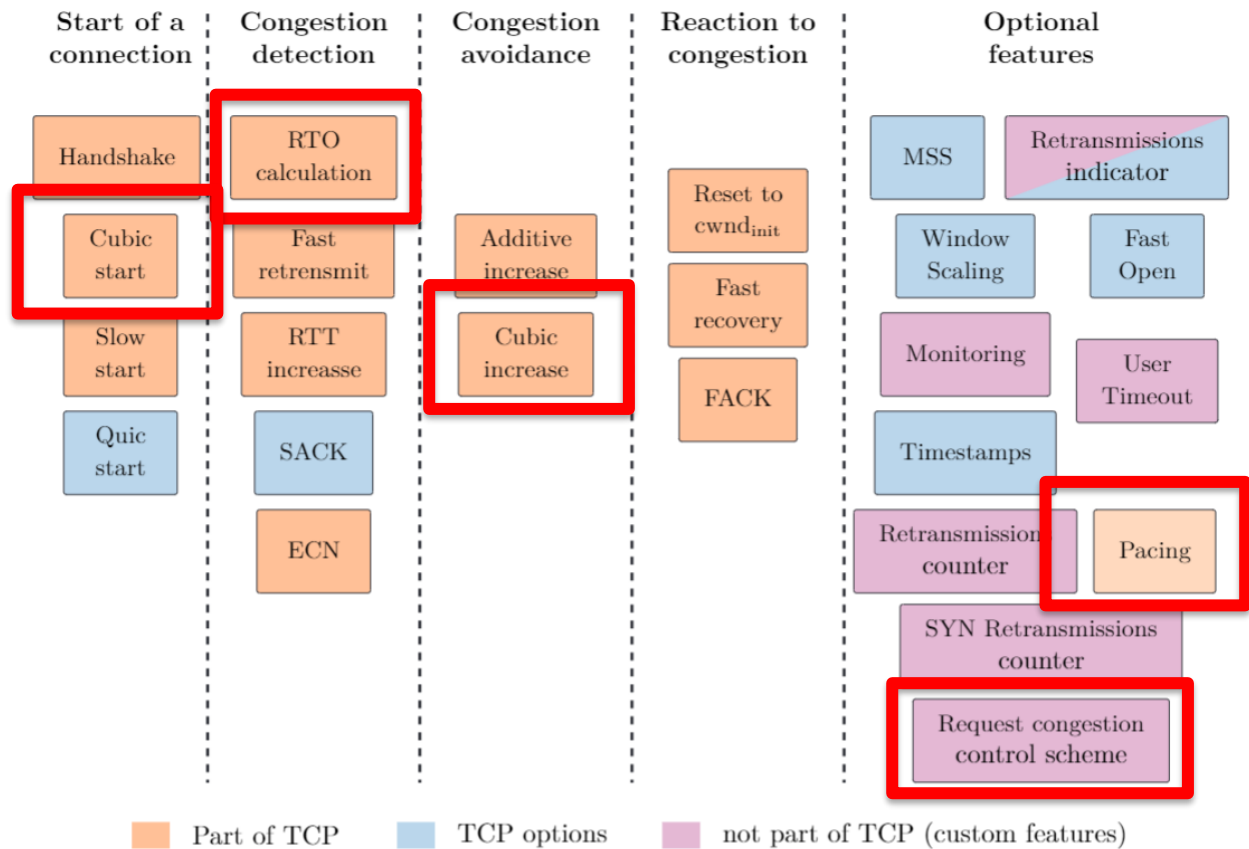
P4

- High-level language
- Defines how packet processors are configured and operated on network packets











- + Easy to implement
- + Fast performance
- + Can be combined with the kernel stack (Only offload TCP options)



- Limited functionality (No floating-point arithmetic, limited accesses to registers)
- Specialized hardware
- Can have high memory utilization
- Can only react to packets

Conclusion

- Three techniques to improve the flexibility of the TCP stack
 - Small overhead
 - Increased flexibility

P4air

TCP complexity is increasing!

- New protocols and congestion control algorithms are continuously being developed
- It is impossible to take their interactions with other protocols and algorithms into account

Goal

Improve fairness between all flows present on a switch by grouping them based on their congestion control algorithm

Goal

Improve fairness between all flows present on a switch by grouping them based on their congestion control algorithm

- From within the data-plane

Goal

Improve fairness between all flows present on a switch by grouping them based on their congestion control algorithm

- From within the data-plane
 - and by taking into account limitations on actions and/or memory accesses

P4air

- P4air: Increasing Fairness among Competing Congestion Control Algorithms
- Conference: IEEE ICNP 2020
- YouTube video:
<https://youtu.be/udXrPi6GVtk>