



# TCP Meets Mobile Code

Parveen Patel  
Jay Lepreau  
*(Univ. of Utah)*

David Wetherall  
Andrew Whitaker  
*(Univ. of Washington)*

1



## The Key Idea

- Transport protocols, such as TCP, need a better upgrade mechanism
- Untrusted mobile code will work!

2

## TCP is a work-in-progress

- A steady stream of TCP extensions and new transport protocols
  - TCP SACK (1996)
  - TCP Connection Migration (2000)
  - ECN and ECN nonce (2001)
  - TCP Nice (2002)
  - TFRC (2000)
  - DCCP (2002)
  - SCTP (2002)
  - ...

3

## Upgrading TCP takes forever

- Research and simulation
- Prototype
- Standards committee
- Implementation in OS 1
- Implementation in OS 2
- ...
- Addition into standard build OS 1
- Addition into standard build OS 2
- ...
- Enable by default
- Enable by default on peer

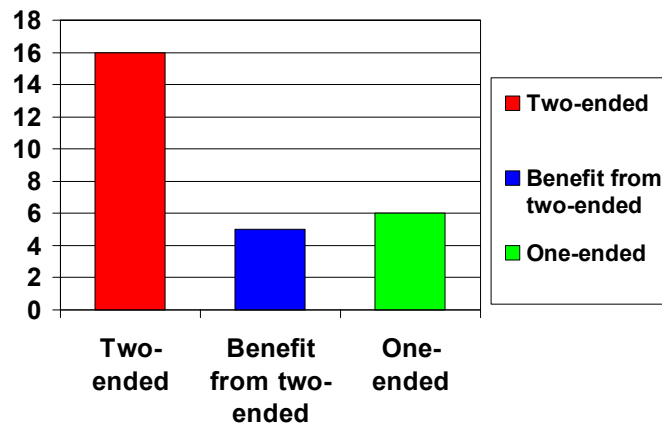
4

## Lousy fallback: one-ended change

- Immediate deployment for self benefit
- Does not always work
  - ◆ Can't exchange new information
- Does not work very well
  - ◆ Lose the benefit of cooperation between both ends

5

## Survey of transport extensions



6

## Our Solution: XTCP

- Connection peers can upgrade each other with new transport protocols using mobile code
- Deployment at one end is all we need !

7

## Upgrading with XTCP is faster

- Research and simulation
- Prototype
- Standards committee
- Implementation to the XTCP API
- Implementation in OS 1
- Implementation in OS 2
- ...
- Addition into standard build OS 1
- Addition into standard build OS 2
- ...
- Enable by default
- Enable by default on peer

8


## XTCP usage scenario #1

- A web server pushes a “high-performance” TCP to its clients the first time they connect

Server



Client



9


## XTCP usage scenario #2

- A mobile client pushes “TCP connection migration” [MobiCom '00] to a server to allow itself to move

Server



Client




10

## XTCP usage scenario #3



- A user installs “TCP nice” [OSDI '02] to support background data transfer

Host A

(With TCP-Nice)



Host B



11

## Will it work ?

- XTCP sounds similar to the challenging domain of active networking
- Challenges can be met in this restricted domain

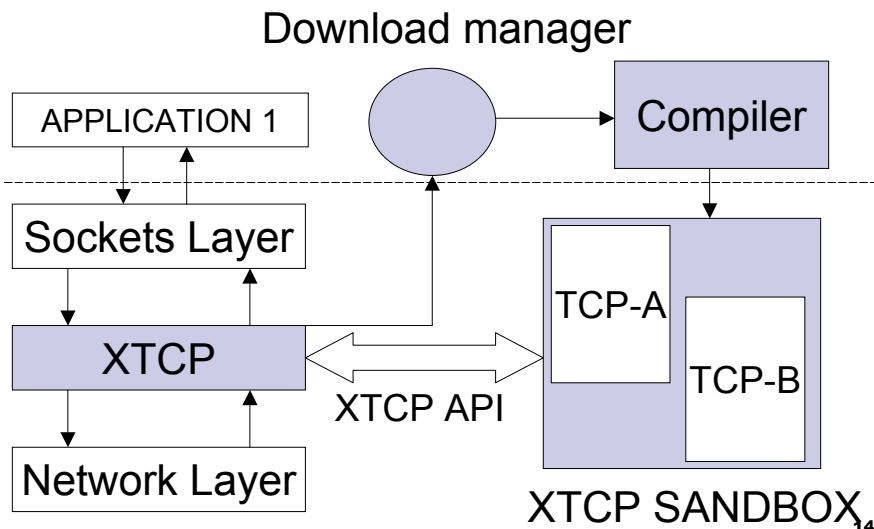
12

# XTCP Challenges

1. Host safety – must isolate and limit resource consumption
2. Network safety – should not compete unfairly or attack other nodes
3. Performance – should not undermine improvement due to extensions

13

# XTCP Design



# 1. Host safety

- No shared state between extensions
  - ◆ Easy resource accounting
  - ◆ Easy termination
  
- Memory safety: type-safety of Cyclone
  
- CPU timer-based CPU protection

15

# 2. Network safety

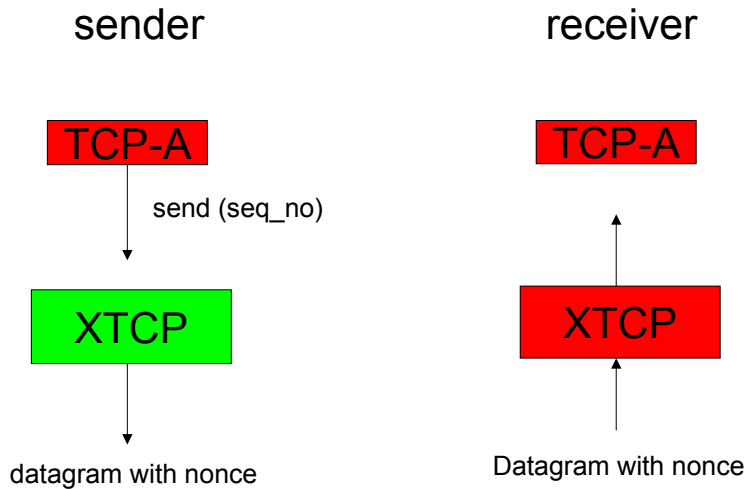
- Well-defined notion of network safety
  - ◆ TCP-friendliness [RFC 2914]
  - ◆ TCP response function is mathematically defined [SIGCOMM '98]
  
- Enforcement without trusting transports
  - ◆ Adapt ECN nonce mechanism is used for validation [ICNP '01]

16



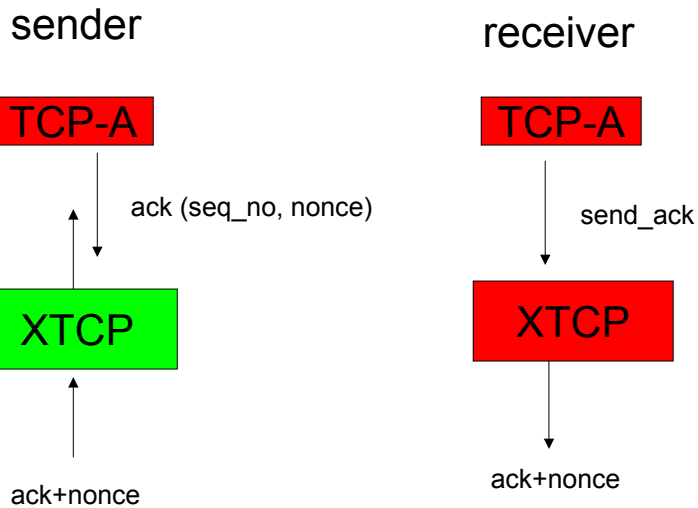
# Nonce Mechanism

[untrusted red, trusted green components]



17

# Nonce Mechanism



18

## 3. Performance

- Connections proceed without delays
  - ◆ Code is downloaded out-of-band
  - ◆ Benefits later connections
- Efficient to share data between the C-based kernel and Cyclone code
  - ◆ No garbage collection
  - ◆ Lightweight runtime

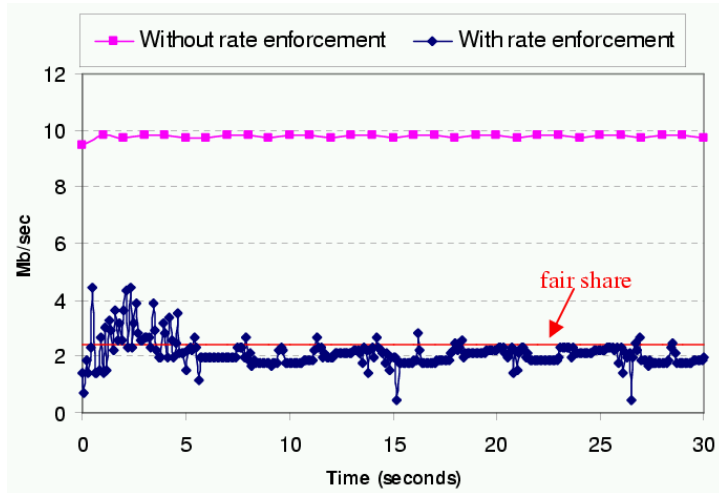
19

## Status

- Prototype in FreeBSD 4.7
- Modest memory and CPU cost
  - ◆ CPU cost is 80% more than base TCP on the sender side, without any optimizations.
- Ported TCP Friendly UDP, TCP NewReno and TCP SACK to the XTCP API
- User-level version in progress

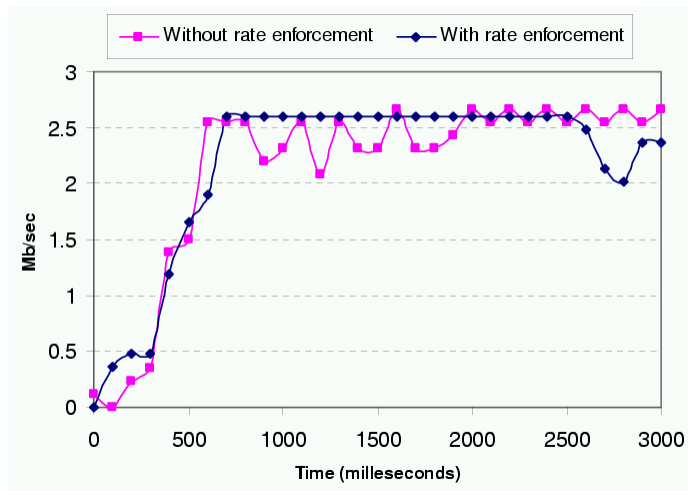
20

# It works! UDP



21

# It works! TCP



22

## Open research issues

- TCP rate policing function
  - ◆ Quickly detect unresponsive extensions
  - ◆ Admit all responsive extensions
  
- XTCP API
  - ◆ Must be sufficient and portable

23

## Conclusions

- Transport protocols need self upgrade mechanism
  
- Mobile code works !
  - ◆ Constrained domain and recent advances
    - Mathematical definition of TCP response function (1998)
    - Cyclone (2002)

24



END OF TALK

....

BACKUP/DETAIL SLIDES

25



## Policies

- Applications can use socket options
- Administrators can set policies
- Policy daemons can collaborate

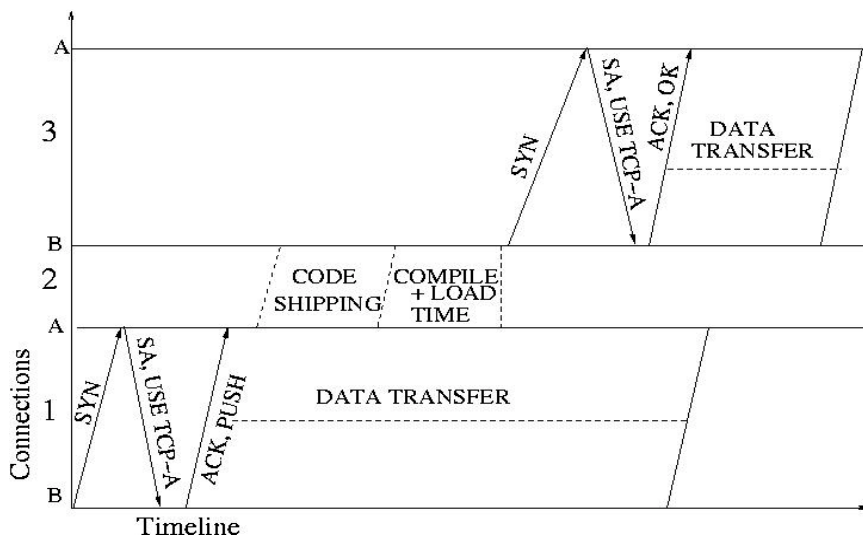
26

## Granularity of extensions

- Transport protocols are written to the XTCP API
- Complete transport protocols are transferred over the network
  - ◆ Retains the way protocols are written now
  - ◆ Maximum flexibility
  - ◆ Maximum simplicity
  - ◆ Code is not large: 85K compressed source

27

## Connection setup and code loading



## Network safety

- TCP friendly network access API
  - ◆ *xtcp\_net\_send* (seq\_no)
  - ◆ *xtcp\_net\_resend* (seq\_no)
  - ◆ *xtcp\_net\_ack* (seq\_no, nonce)
  - ◆ *xtcp\_net\_acksum* (seq\_no, nonce\_sum)

29

## A Fourth Challenge: Deployment of XTCP framework

- Benefits self
- Can only harm self
- Deployment only needed at end points
- TCP-friendliness is non-threatening

30