HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Future Internet From the User Perspective

Jussi Kangasharju
University of Helsinki

Outline

- Has Future Internet forgotten the users?
- How to include users? What do they want and need?
- Application-level fairness
  - Joint work with M. Mu and G. D. Colussi

What Is Internet?

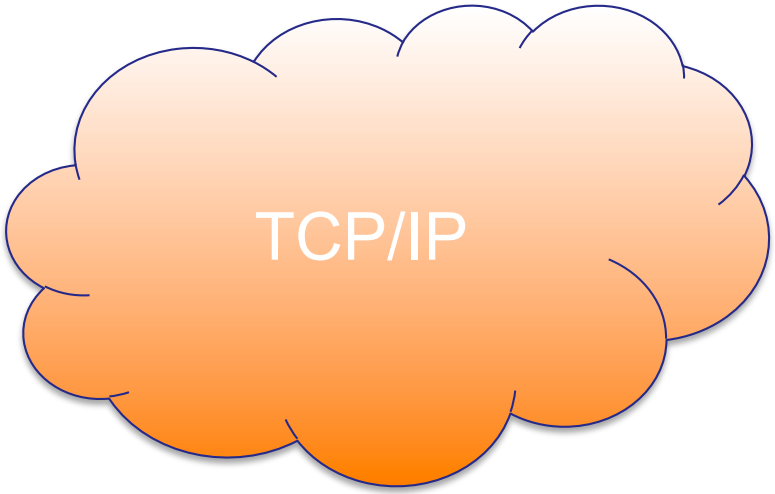

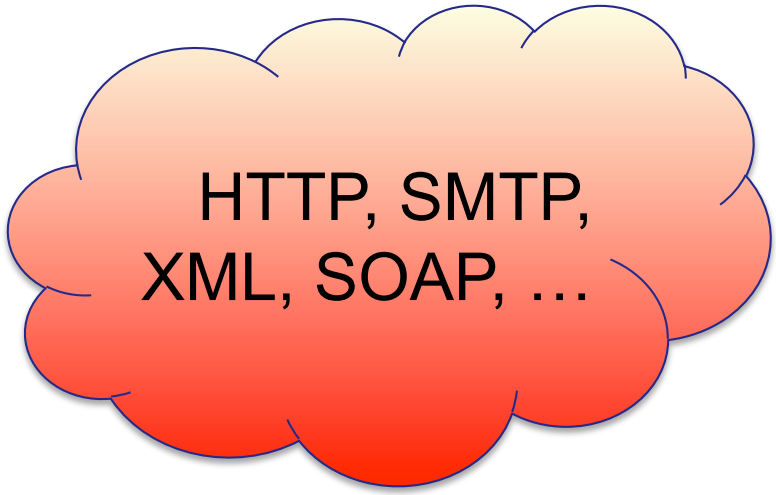

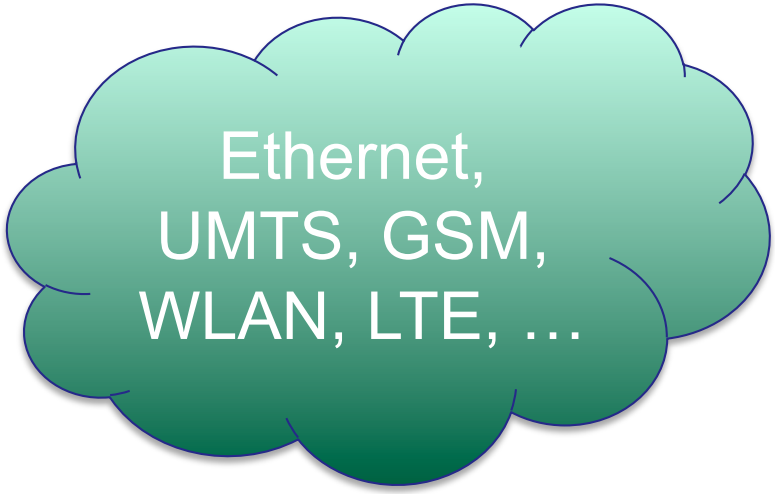

What is your answer?

User is HERE! →

Who takes care of this?

Most Future Internet projects are HERE →

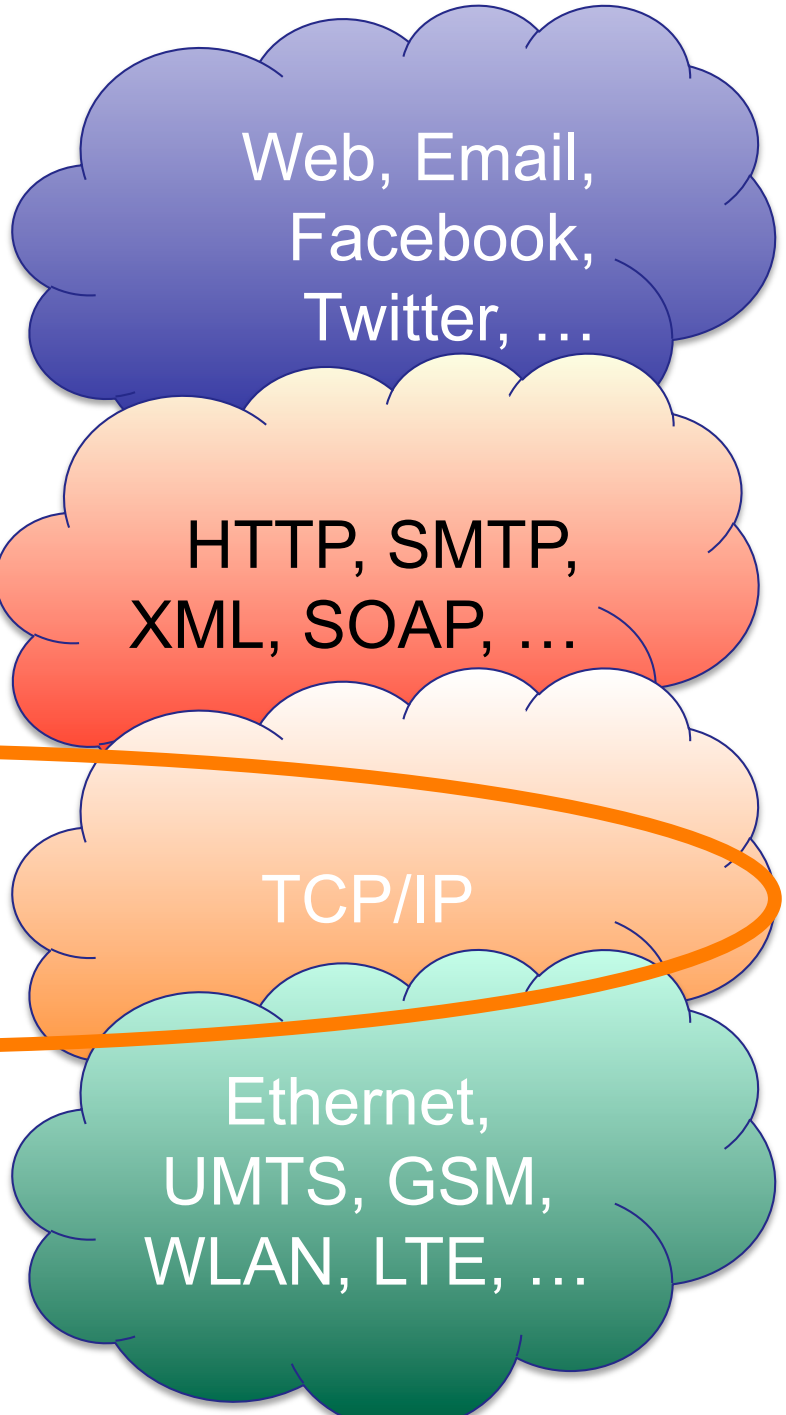

Problem and Solution?

Problem:

- Future Internet == Research into network infrastructure
- Users don’t care for infrastructure

Conjecture/Fact:

- User actions affect even lowest levels of network stack
- Should not (cannot?) design infrastructure in isolation
- Innovation driven by applications, not infrastructure

Solution:

- Include users and applications
- Get “user people” and “infrastructure people” talking

How Can You Design the Future Internet if You Are Not Using the Current Internet?

My Messages

For infrastructure people:

- Use the Internet!
  - Internet = Facebook, Flickr, YouTube, Twitter, web, P2P
  - ssh is a dinosaur
- Understand user needs
  - Me, my stuff & my friends

- For user people:

  ???????

How to Include Users?

- Users + Network infrastructure = ?
- What should we do?

Systems for “Normal” People

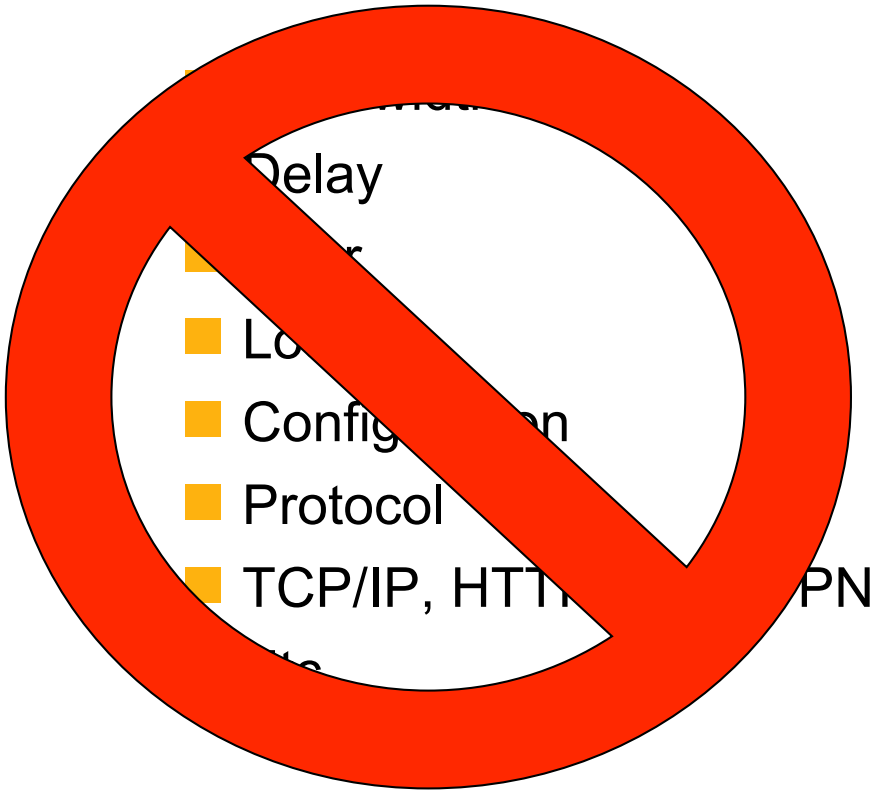

Must speak their language!

Language of "Normal" People

- It's a very simple language:

Am I getting what I want?

- In other words, is the user satisfied with the service?
  - Service meets user's expectations and requirements
  - (User = Human user or another computer program!)

Satisfied ≠ Best possible performance

- It can be, but doesn't have to be

- **Success measured in terms of user satisfaction**

Fairness in Networks

- How to evaluate network performance?
- Currently: Network-centric measures, e.g., utilization
- Better: Does network satisfy user's requirements?
- Problem: How to model and measure user satisfaction?
- Answer: Application-level utility metrics for different network parameters

Network Parameters

Fairness != Fair bandwidth sharing

- Can have fair bandwidth sharing and unfair treatment of applications
  - Actually: Happens very often with TCP
- How should network behave towards applications?

Observation

- What are effects of congestion on applications?
- Traffic is ***affected*** by congestion in network
  - Increased delay and loss
- ***Impact*** of congestion on application is application-specific
- Users ***experience*** the impact of congestion
- Must study all three aspects! → **Application-level fairness**

Modeling Applications

- Bandwidth is **positive**
  - Increase in bandwidth makes life better
- Delay, jitter, and loss are **negative**
  - Increase in these makes life worse
  - Also called damaging parameters
- Generalization of ITU's E-Model

Bandwidth Utility

- Elastic and real-time applications

- Use logarithmic utility function
  - Similar to work of F. Kelly

$$u(x) = C\log(1+x)$$

- C normalizes utility to 1 when user is satisfied

Utility for Damaging Parameters

- Application dependent bounds for delay, jitter, and loss
  - Below a threshold not visible to the user
  - Above another threshold, becomes “unusable”

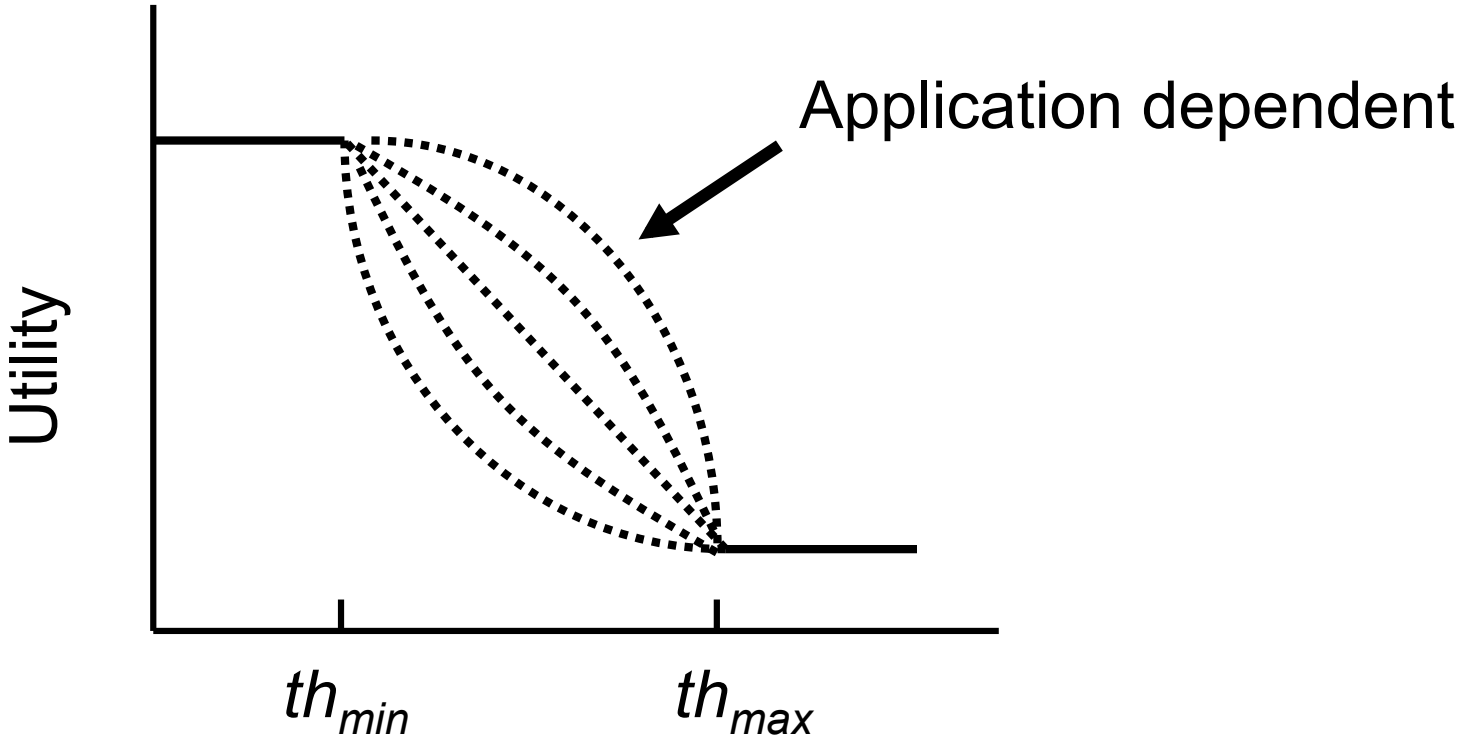

Utility for Damaging Parameters

- Damage utility function:

$$u_{\tau,\phi,p}(z) = \begin{cases} 1 & if \quad z \le th_{min} \\ F(z) & if \quad th_{min} < z < th_{max} \\ u_{min} & if \quad z \ge th_{max} \end{cases}$$

- Parameters th_{min}, th_{max}, $F(z)$, and u_{min} application-specific

- General form applies to any application
  - Many studies confirm by deriving parameter values
- Feasible to derive parameters for application classes

Combining Utilities

- Intuitive properties of combination function
  - If all damaging utilities are 1, then U = u(x)
  - If any damaging utility is < 1, then U < u(x)
  - If any damaging utility is 0, then U = 0
- We use product of individual utilities as combination
  - Same used in E-Model
- Choice of right combination function still an open question

Thresholds for Real Applications

- Examples, see more in paper

Application	Bandwidth		Delay		Jitter		Loss		Source
	$th_{\_min}$	$th_{\_max}$	$th_{\_min}$	$th_{\_max}$	$th_{\_min}$	$th_{\_max}$	$th_{\_min}$	$th_{\_max}$	
VoIP	C	64kbps	100ms	150ms	40ms	75ms	1%	3%	[9-12]
Video phone	16	384kbps	150ms	400ms	50ms	80ms		1%	[13]
Web	Elastic		2s	4s	N/A		N/A		[13,16-18]
Xbox Halo	Framerate		50ms	200ms			1.5%	3.5%	[20]
Bulk data	Elastic		N/A		N/A		N/A		[13]

Analysis and Evaluation

- Show that fair bandwidth sharing is not enough

- Several TCP and VoIP flows over same link
- VoIP flows get their bandwidth, TCP shares the rest
  - Model analytically with RED
  - RED hard to tune, but easy to model
  - Cover all “sensible” scenarios
- Two cases:
  - Vary number of flows, keep propagation delay fixed
  - Vary also delay
- Bandwidth always shared fairly, utilities NOT fair

Case 1: Fixed Delay

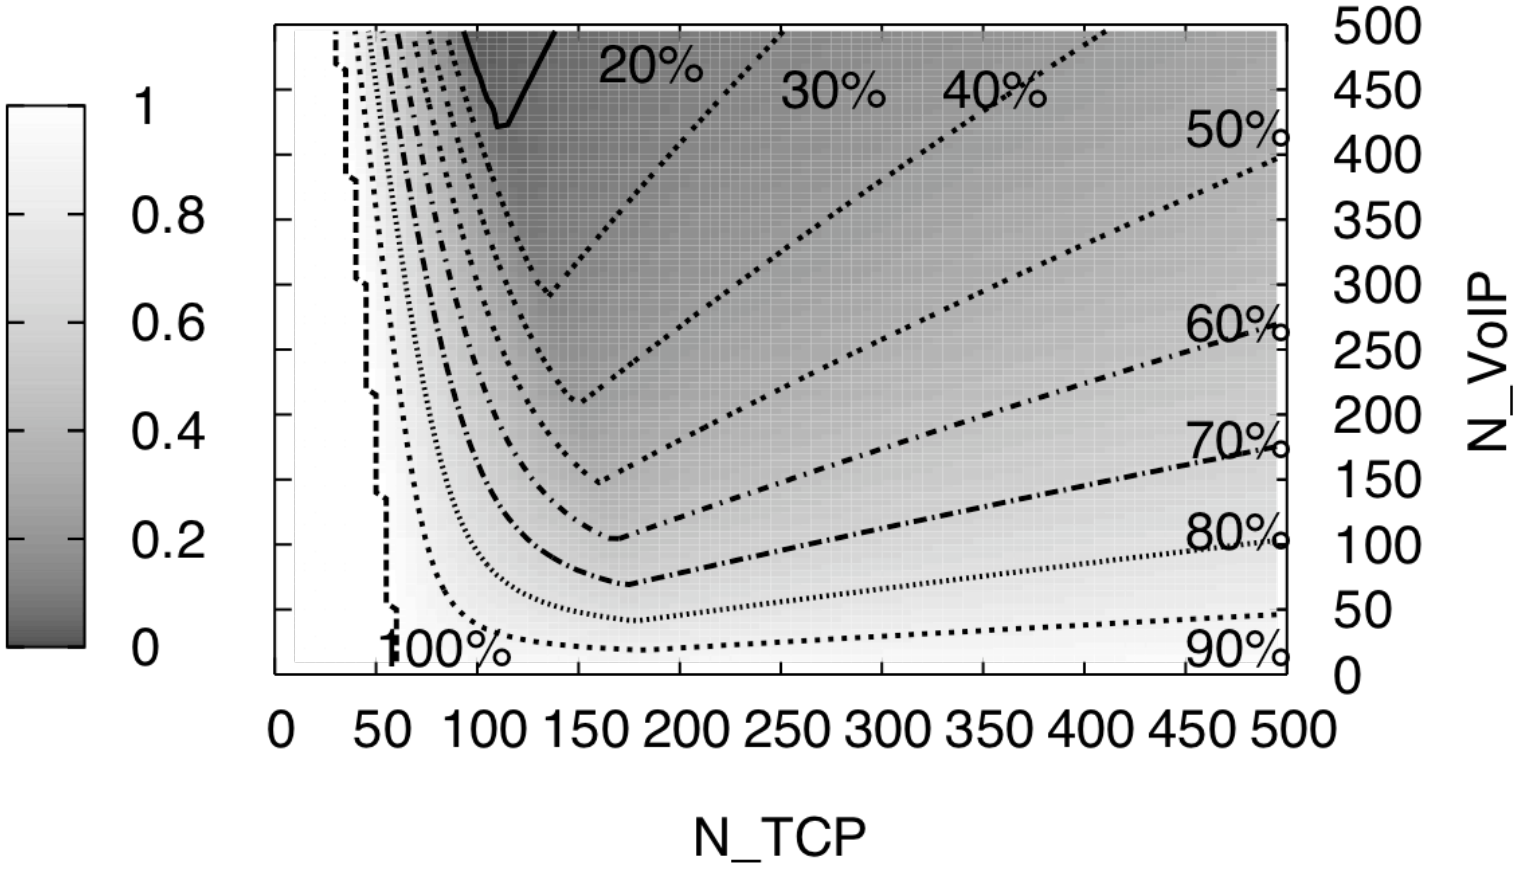

- Utility averaged over all flows in system

Case 1: VoIP vs. TCP

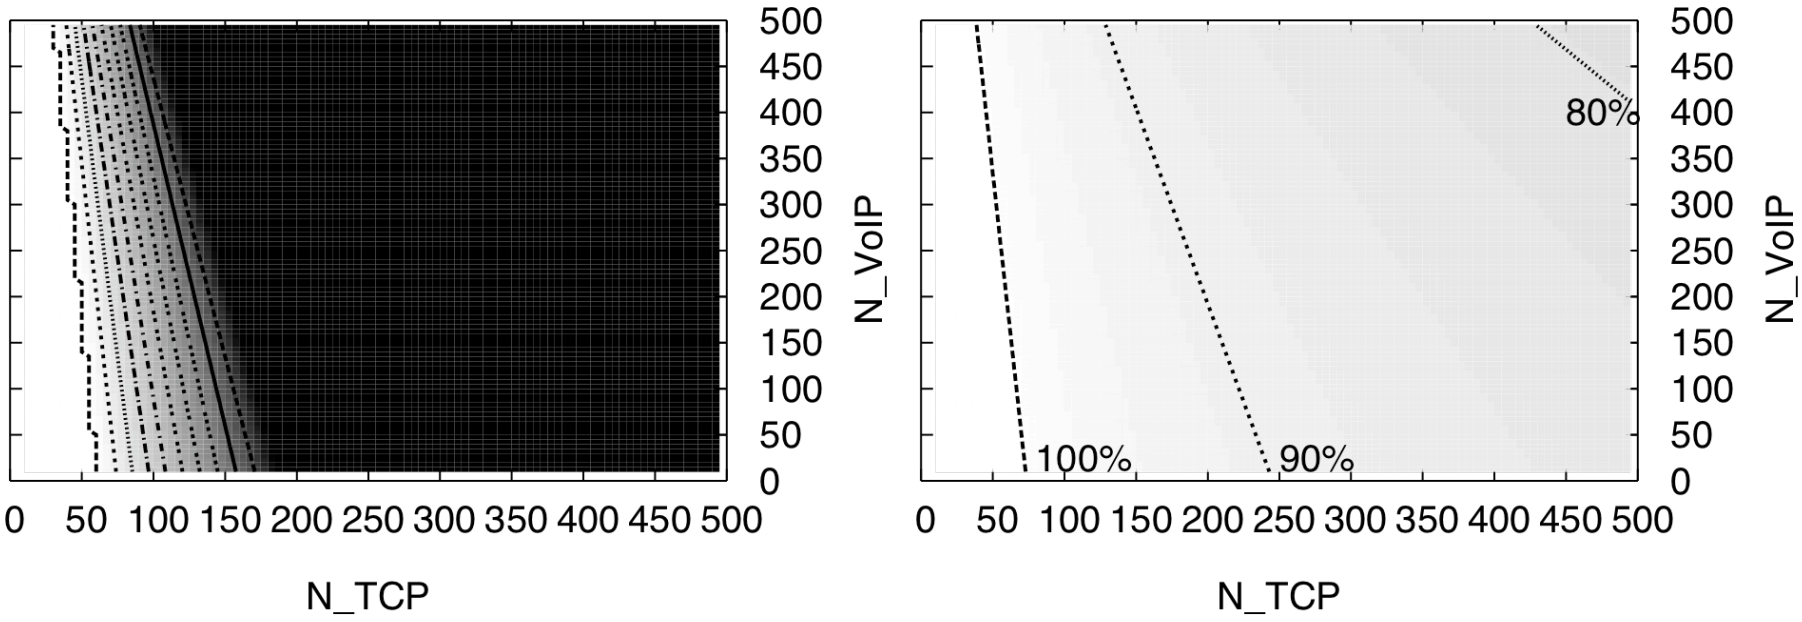

VoIP TCP

- VoIP suffers greatly, TCP does not suffer
- VoIP **NOT** treated fairly, even though bandwidth is fairly shared

Case 2: Vary Flows and Delay

- Similar results apply

- Small delay → High loss → Low utility
- Large delay → Low utility

- Reason for problems:
  Combined effect of damaging parameters has only a small range where VoIP can deliver useful service

Summary and Conclusion

- We need to consider application-level effects in congestion control
- Fair sharing of bandwidth alone does not give fairness
- Must use a wider range of parameters
  - Parameters already exist for many application classes
- Analytical evaluation to show actual effects
- Clear need for future research

Thank You!

Email: Jussi.Kangasharju@cs.helsinki.fi