A Queue Management Algorithm for Intra-Flow Service Differentiation in the „Best Effort“ Internet

Henning Sanneck

GMD FOKUS, Berlin

sanneck@fokus.gmd.de

ICCCN 99, Natick, MA

October 12, 1999

Overview

- Introduction
  - Motivation: graceful degradation under congestion for real-time multimedia flows in the Internet
- Differential RED algorithm (DiffRED)
  - Properties, Differences to RED
- Evaluation
  - Traffic Model
  - Results
- Conclusions

GMD FOKUS

Motivation

- Loss sensitivity of Internet real-time flows:
  - Video: bursty frame losses (packet losses affecting several frames)
  - Voice: bursty packet losses (dependent on codec)

    ➢ „drop-outs“, high perceptual impact
- Solutions:
  - Reservation (IntServ): complete deployment incl. charging
  - Exploit flow inhomogeneity (Application-layer filtering within the network): payload-specific, large amount of resources needed, might affect security, difficult to apply to voice

GMD FOKUS

Motivation (cont'd)

- Solutions:
  - End-to-End loss recovery (FEC, loss concealment): efficiency also subject to loss patterns, might worsen congestion
- Adequate mapping of applications' requirements (ADU structure, End-to-End loss recovery) to *simple* network mechanisms
- Provision of a *service* which offers control over (transient) loss distribution / correlation
- Bridge the gap between "Best Effort" and full QoS deployment (QoS migration)

GMD FOKUS

Approach

- What basic mechanisms are needed at a gateway to realize such a service ?
- Simple queue management algorithms (RED) already influence loss correlation (gradual adjustment of the drop probability)
- RIO extends RED to provide *inter-flow* service differentiation
- Application of RIO approach to *intra-flow* service differentiation

GMD FOKUS

Differential RED (DiffRED)

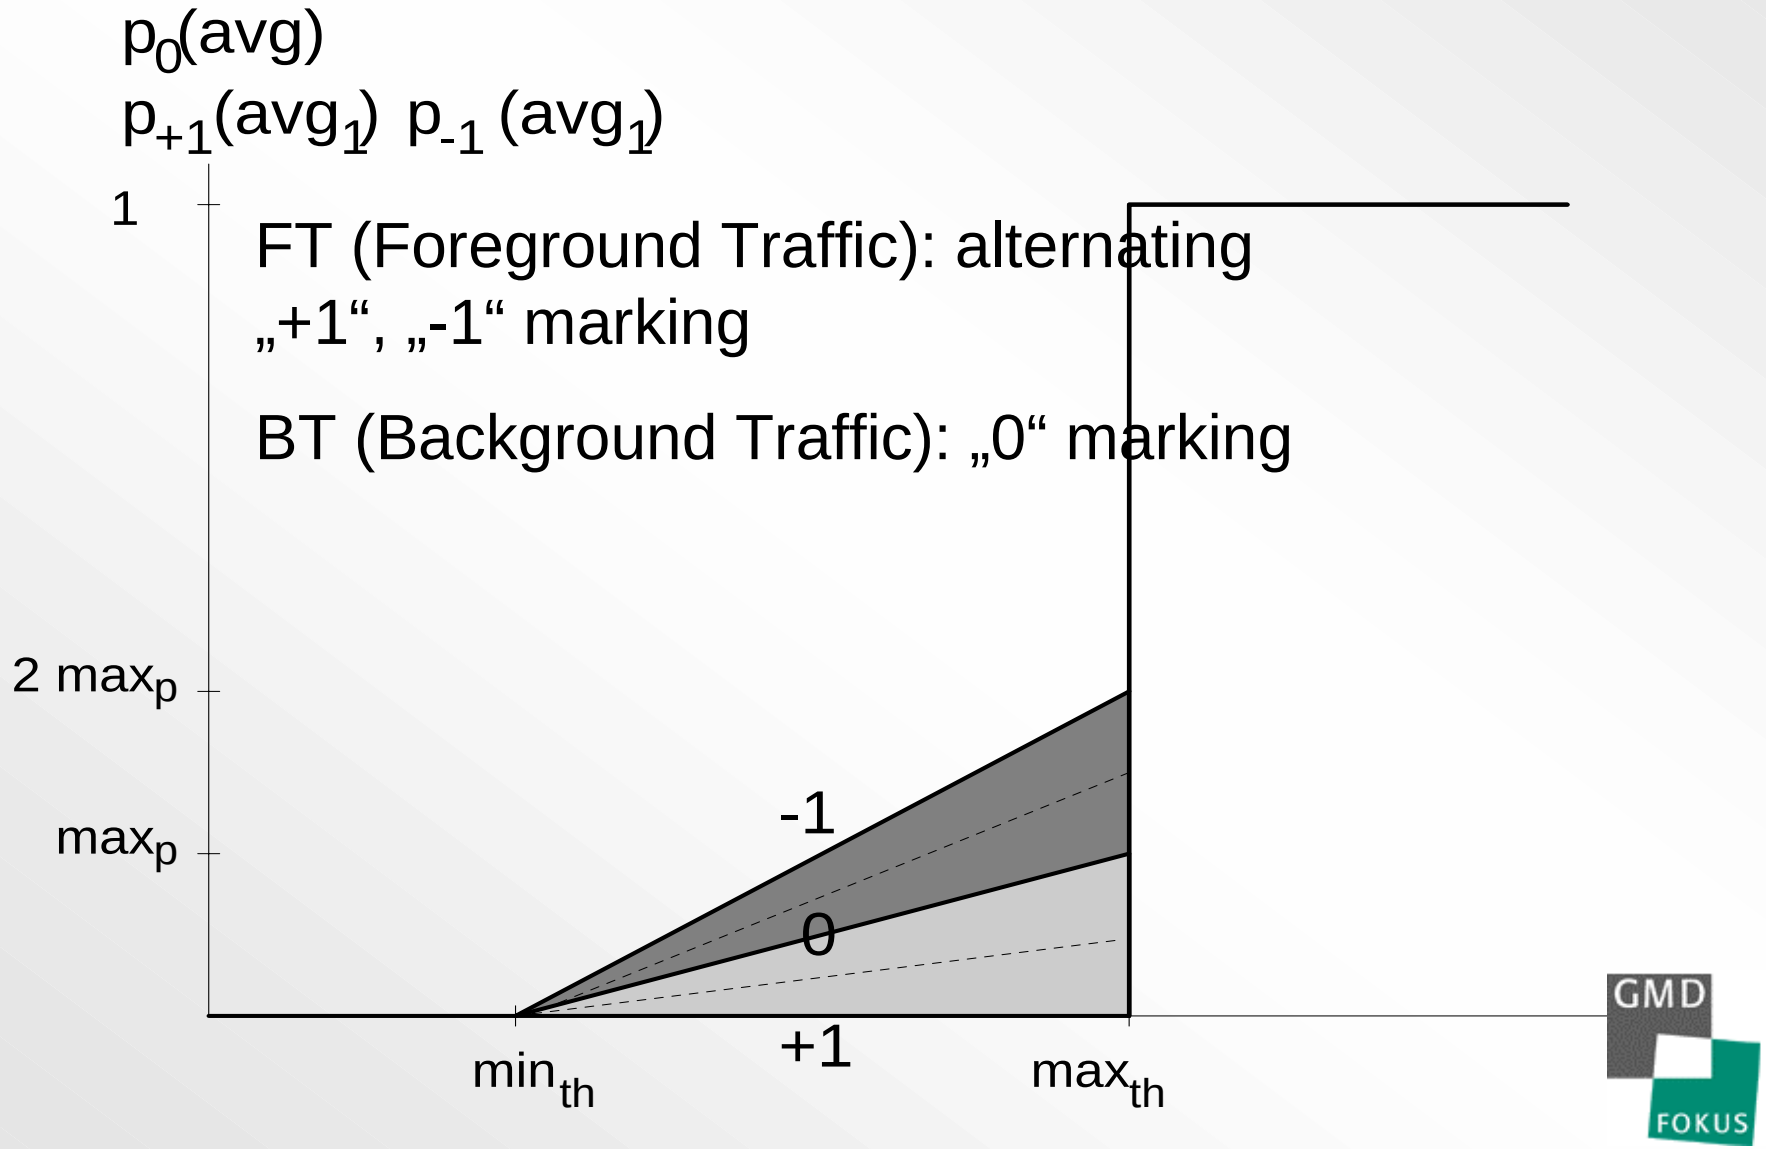

DiffRED: Issues (1)

- „Differential“ loss probability curves (compensation of lower/higher FT loss probabilities in the long term)
- Queue state *(avg)* might change substantially between FT arrivals

$$avg \leftarrow (1- w_q)avg + w_q\, q$$

- Possible solutions *:*
  - $w_{q,1} = f(FT, BT\ arrivals) \geqslant avg_1$
  - $q_1 =$ *f(FT arrivals):* sub-sampling of *q* at FT arrivals $\geqslant avg_1$

GMD FOKUS

DiffRED: Sub-sampling

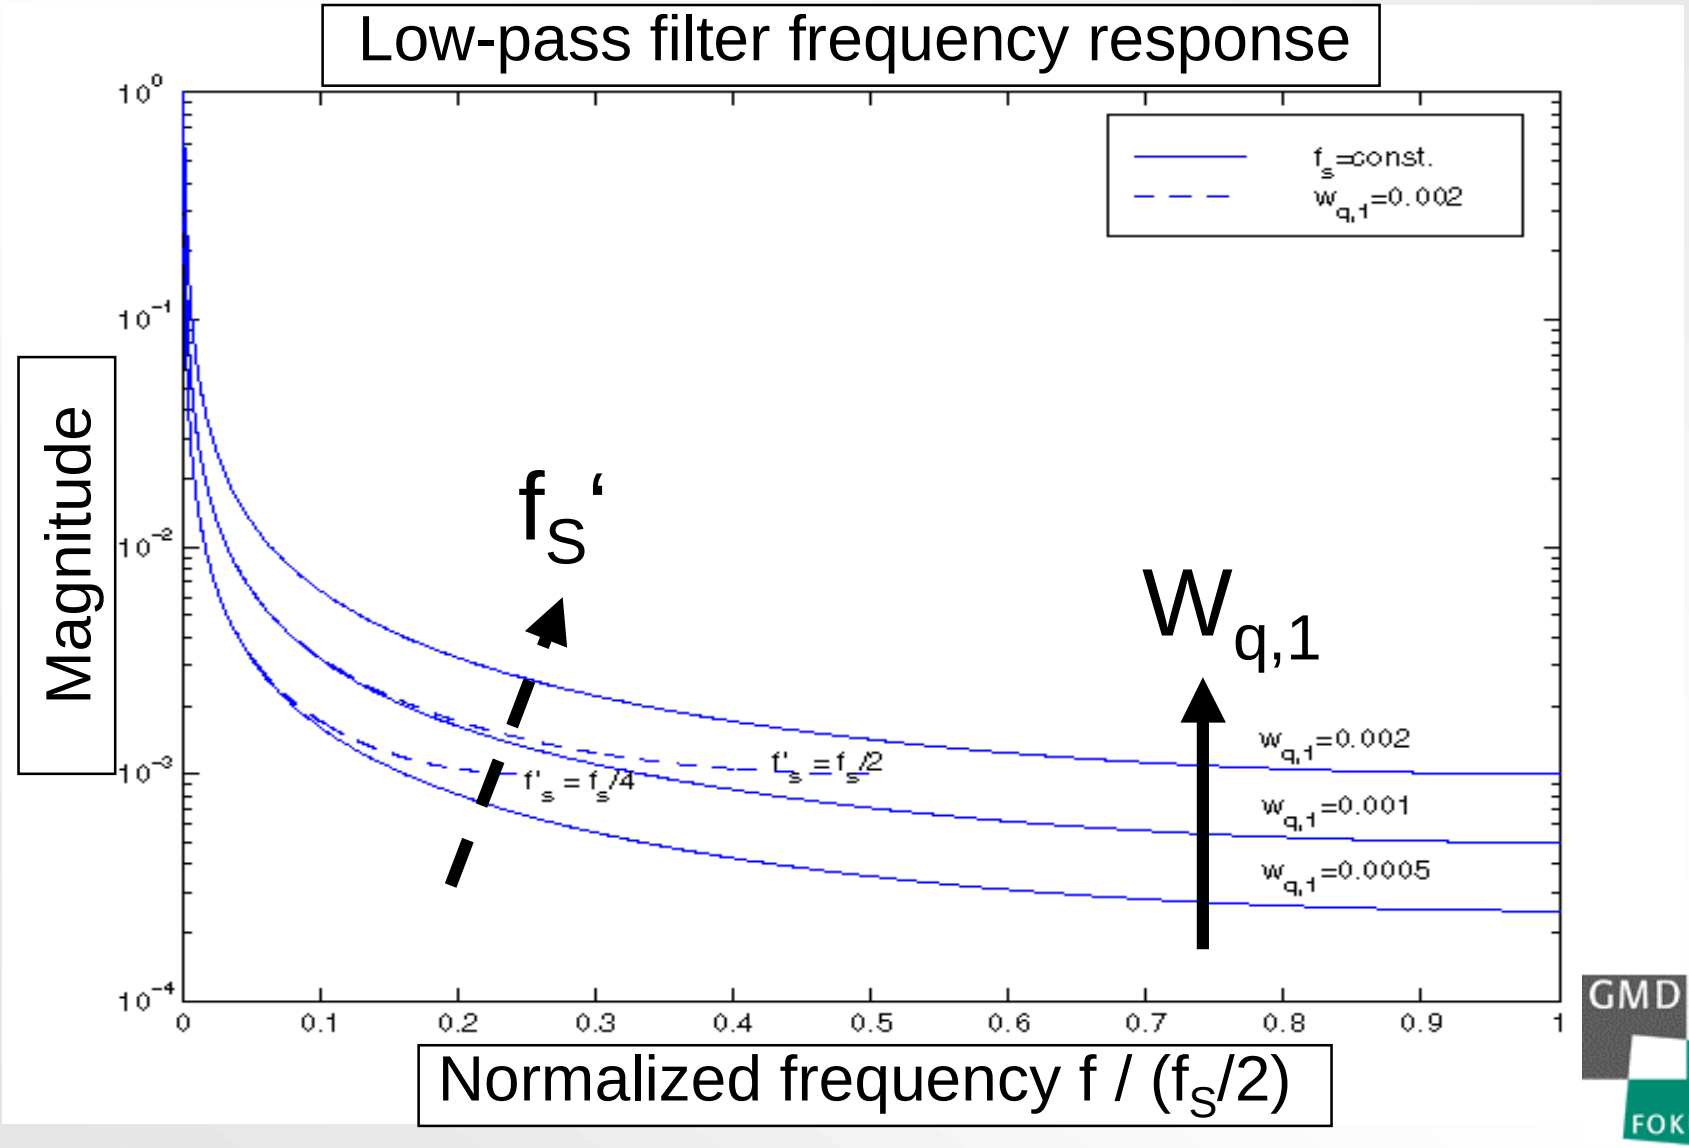

GMD FOKUS

DiffRED: Issues (2)

- Irregular partition of +1/-1 arrivals
- Possible solutions:
  - monitor and penalize misbehaving flows
  - *adjust loss probability curves to ratio of +1/-1*
- Injection of -1 traffic to mark a flow entirely as +1
- Possible solutions:
  - monitor and penalize misbehaving flows
  - *volume-based charging*

GMD FOKUS

DiffRED: Summary

- „Differential“ loss probability curves (compensation of lower/higher probabilities in the long term)
- *avg*: sub-sampling of *q*
- Monitoring of *+1/-1* arrivals

➢ fair loss sharing between FT and BT

GMD FOKUS

Results

- Traffic Model (based on measurement studies)

Traffic type	*H-BT*	*D-BT*	*FT (voice)*
flow share (%) of BT	75	25	-
peak bandwidth ($\frac{kBit}{s}$)	256	30...34	83.2
packet size (bytes)	560	128	208
on/off distribution	Pareto	Expo.	Expo.
shape parameter	1.9	—	—
mean burst (packets)	20	4	18
mean ontime (s)	0.35	0.12...0.14	0.36
mean offtime (s)	0.7	0.12...0.14	0.64

- Experiment: variation of the FT load share λ_{FT}/λ at a fixed traffic intensity $\rho=\lambda/\mu=0.95$, single gateway

GMD FOKUS

Results

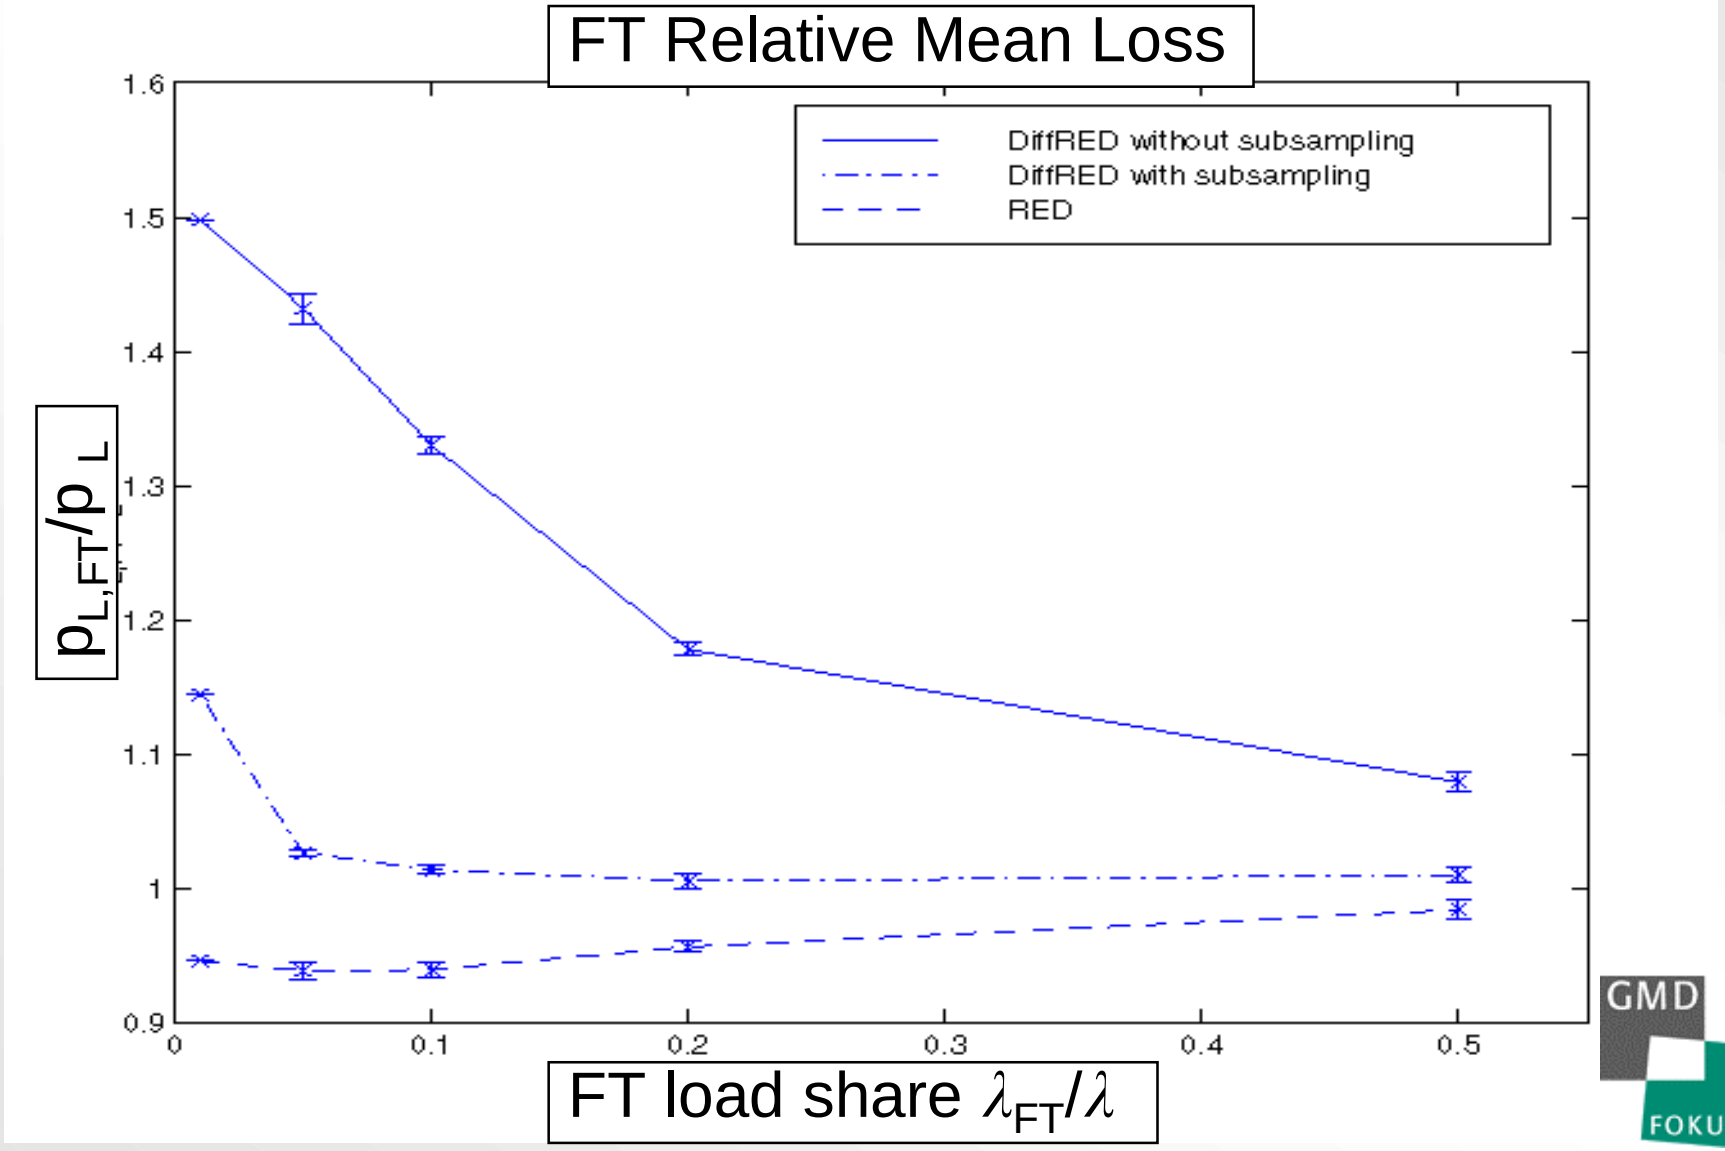

GMD FOKUS

Results

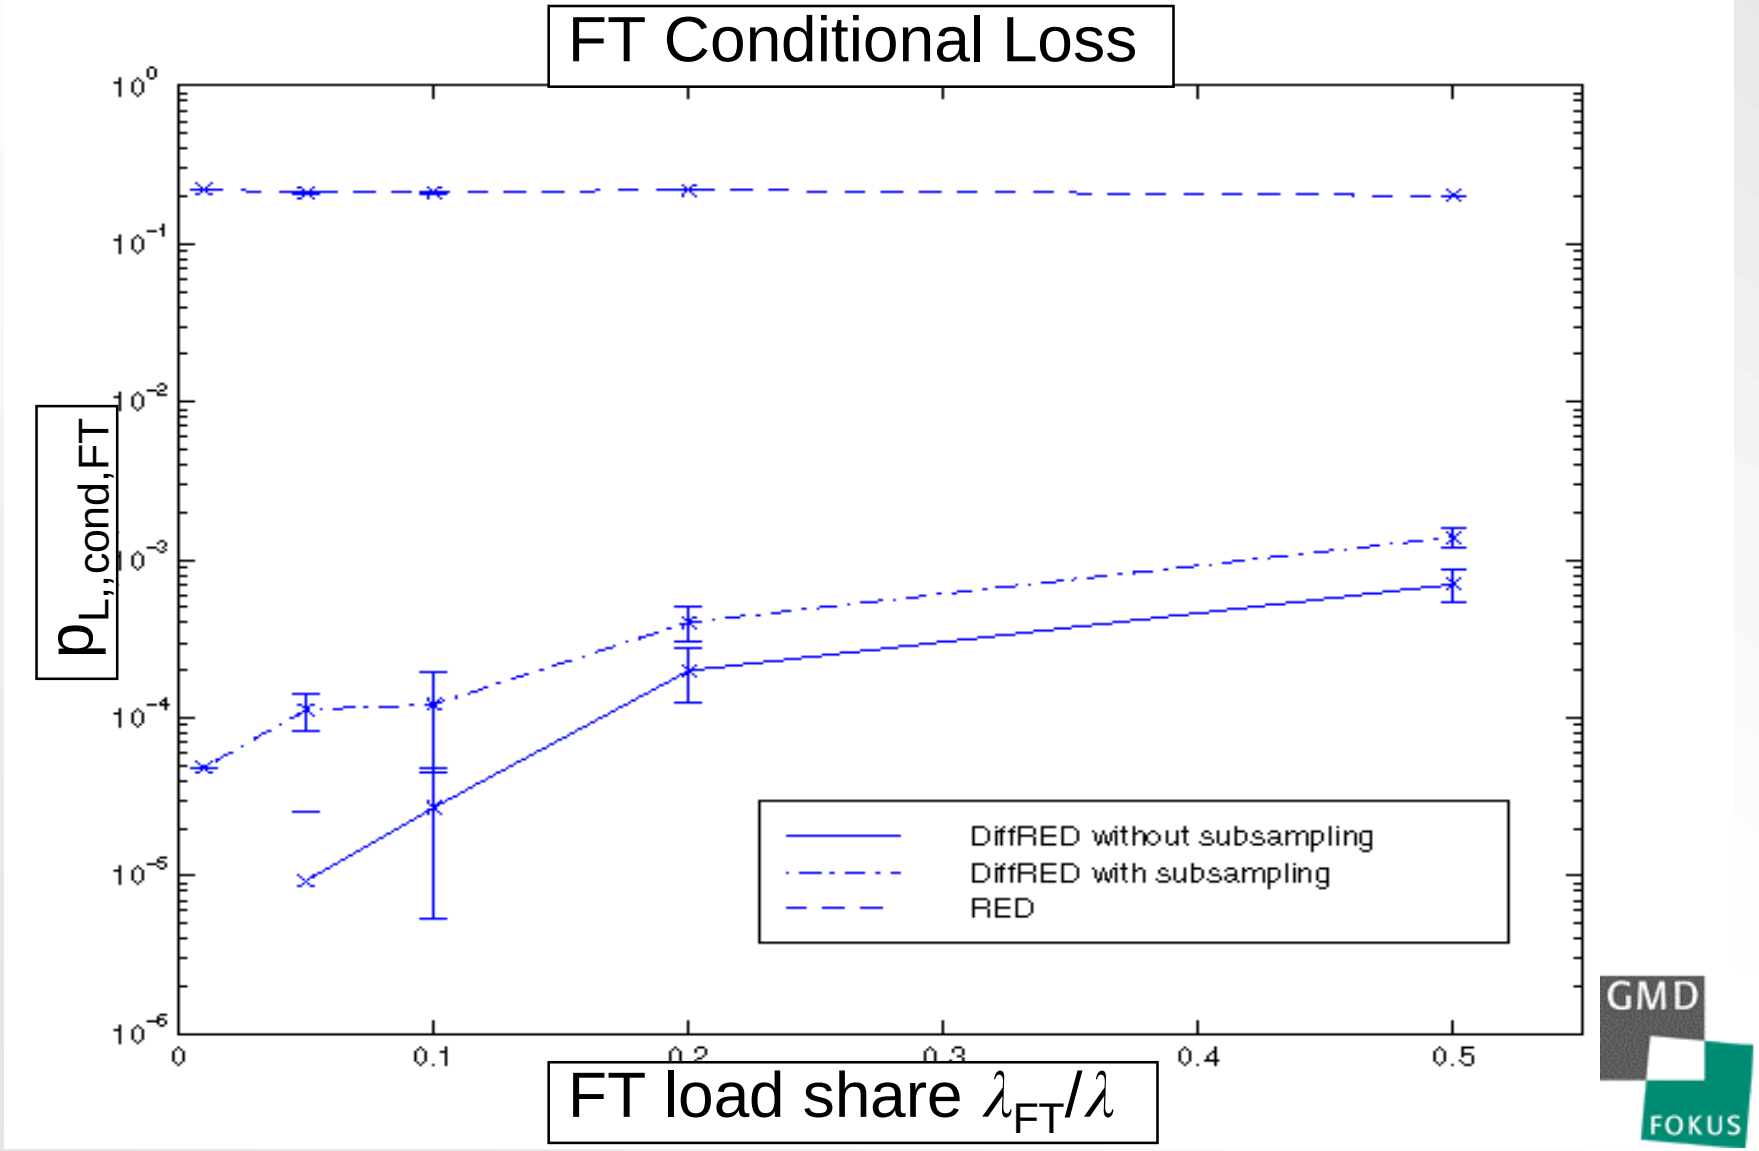

GMD FOKUS

Conclusions

- Control loss distribution for certain flows while maintaining RED properties in the long term
  ➢ applications‘ requirements (ADU, e2e loss recovery) can be mapped on a simple network mechanism
- No complete QoS architecture needed, partial deployment beneficial,
  suitable framework: DiffServ AF (three drop precedences within a class)
- Further study for other (bursty) traffic types needed

GMD FOKUS