Optimizing 802.11 WLANs

Key Insights for High-Performance MAC Configurations
Leveraging Experimental Studies for Capacity and QoE

Pablo Serrano (pablo@it.uc3m.es), Wireless Broadband Alliance, December 10th, 2025

Roadmap

Three Experimental WLAN Optimization Studies

1. Dynamic Control Theory Adaptation: Implementing two algorithms to achieve **optimal throughput and fairness** in real deployments
2. Static QoS Configuration for EDCA: Maximizing **VoIP capacity** using derived rules for 802.11e parameters
3. MAC Layer Innovation: Introducing VoIPiggy to double capacity, validating **video QoE** using 802.11aa GATS

Adaptive algorithms for optimal throughput and fairness in real deployments

P. Serrano, P. Patras, A. Mannocci, V. Mancuso, Albert Banchs,
Control Theoretic Optimization of 802.11 WLANs: Implementation and Experimental Evaluation, Elsevier Computer Networks, vol. 57, no. 1, January 2013, http://dx.doi.org/10.1016/j.comnet.2012.09.010

Optimal throughput and fairness

Dynamic Control Theory Adaptation

- IEEE 802.11: access scheme whose performance depends on the Contention Window (CW)
- "Slotted ALOHA": Given the number of stations, there exists a CW* that maximizes performance
- Adjust the CW based on conditions
  - A lot of activity: increase the CW
  - Less activity: decrease the CW

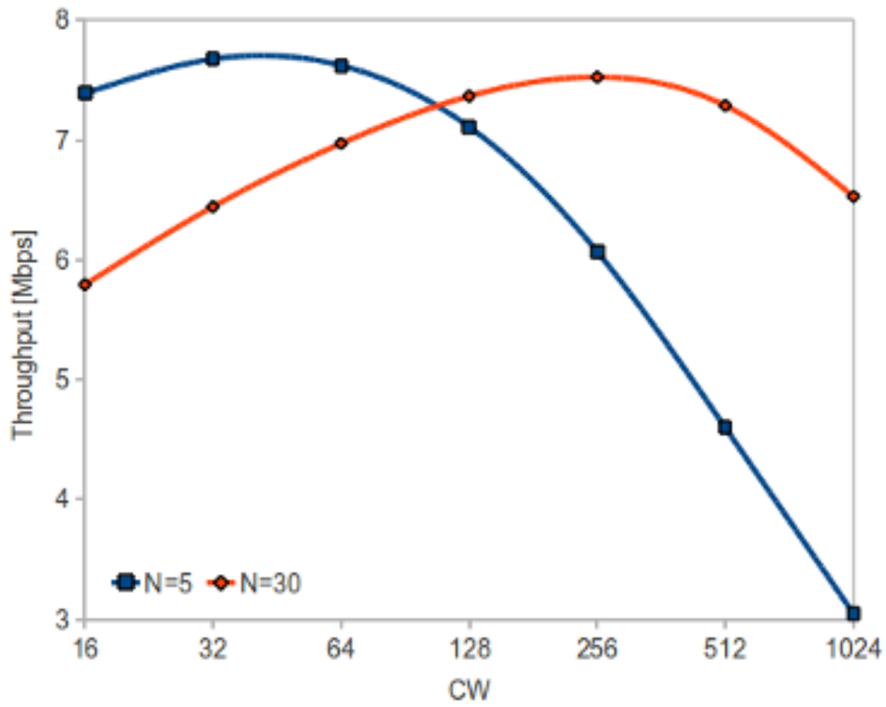

A lot of previous work

Centralized, distributed, etc.

P. Serrano, P. Salvador, V. Mancuso, Y. Grunenberger, Experimenting with Commodity 802.11 Hardware: Overview and Future Directions, IEEE Communications Surveys and Tutorials, vol.17, no.2, pp.671-699, Second quarter 2015, http://dx.doi.org/10.1109/COMST.2015.2417493

- Two types of solutions
  - Centralized approaches: AP computes the configuration and distributes it (now a standard feature)
  - Distributed approaches: stations compute their configuration independently → suitable also for ad-hoc mode
- Limitations
  - Require modifications of the hardware and/or firmware (until EDCA)
  - Their performance has not been assessed in practice
  - Based on heuristics, no analytical foundations

CAC: Centralized Adaptive Control

Centralized approach

- Assuming saturation conditions, optimal transmission probability

$$\tau_{opt} \approx \frac{1}{n}\sqrt{\frac{2T_e}{T_c}}$$

- Optimal collision probability

$$p_{opt} = 1-(1-\tau_{opt})^{n-1} = 1-\left(1-\frac{1}{n}\sqrt{\frac{2T_e}{T_c}}\right)^{n-1} \approx 1-e^{-\sqrt{\frac{2T_e}{T_c}}}$$

CAC: main idea

No need to estimate the number of stations, just the retry bit

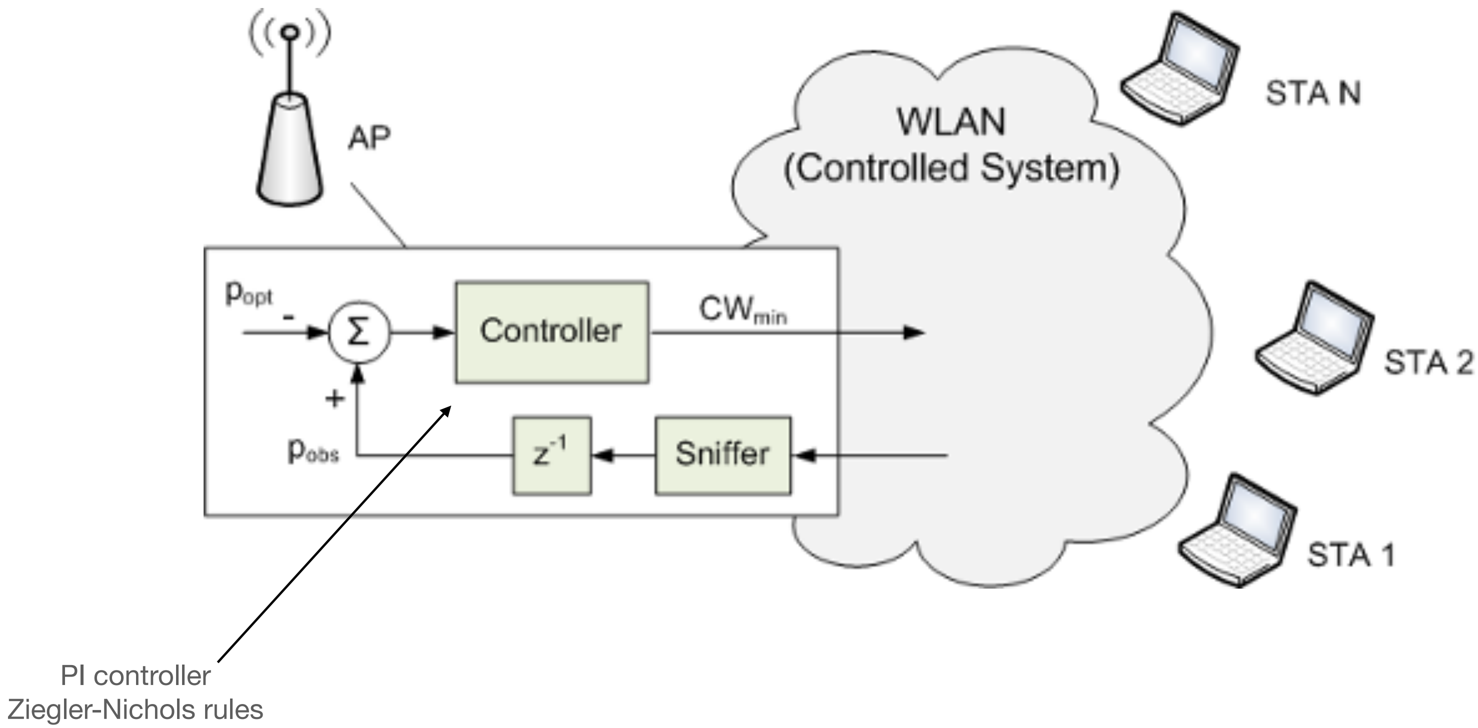

Distributed Adaptive Control

Could we apply the same approach, per station?

- TL; DR: no

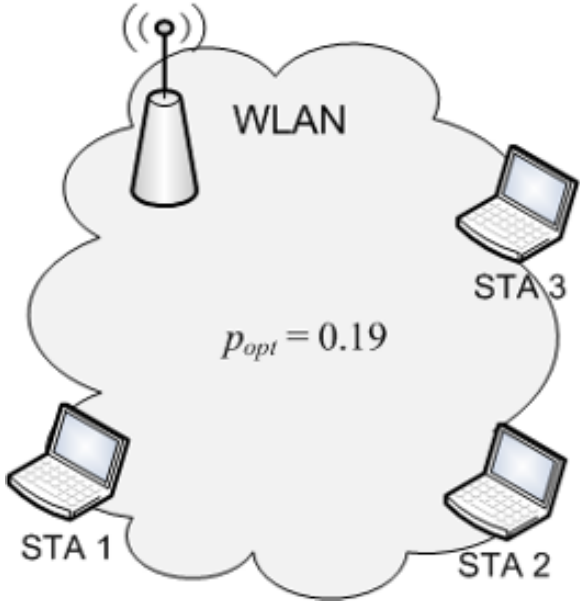

$$p_1 = 1 - (1 - \tau_2)(1 - \tau_3)$$

CW_1	CW_2	CW_3
48	3	28
9	18	42
19	19	19

- Even accounting for this global vs. local vision: capture effect.

Testbed

Under raised floor, different link qualities

- 1 AP and 17 stations

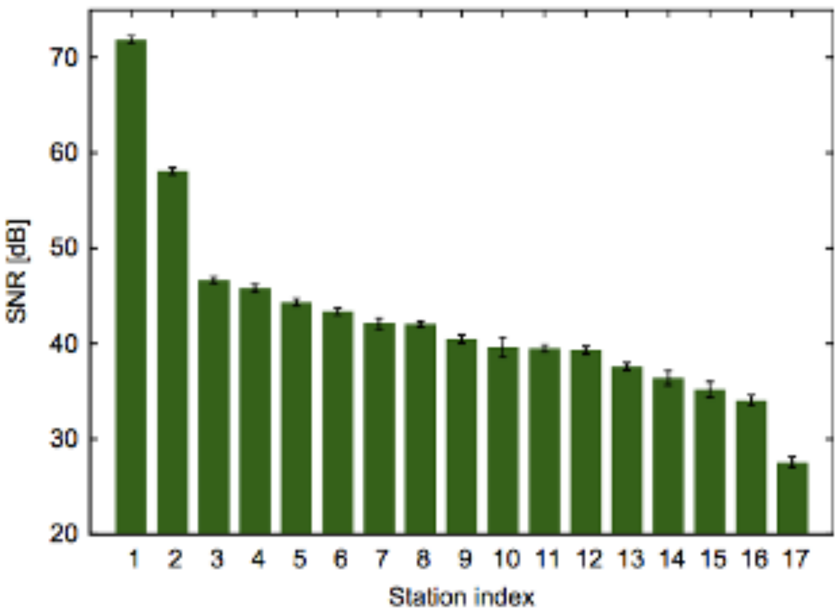

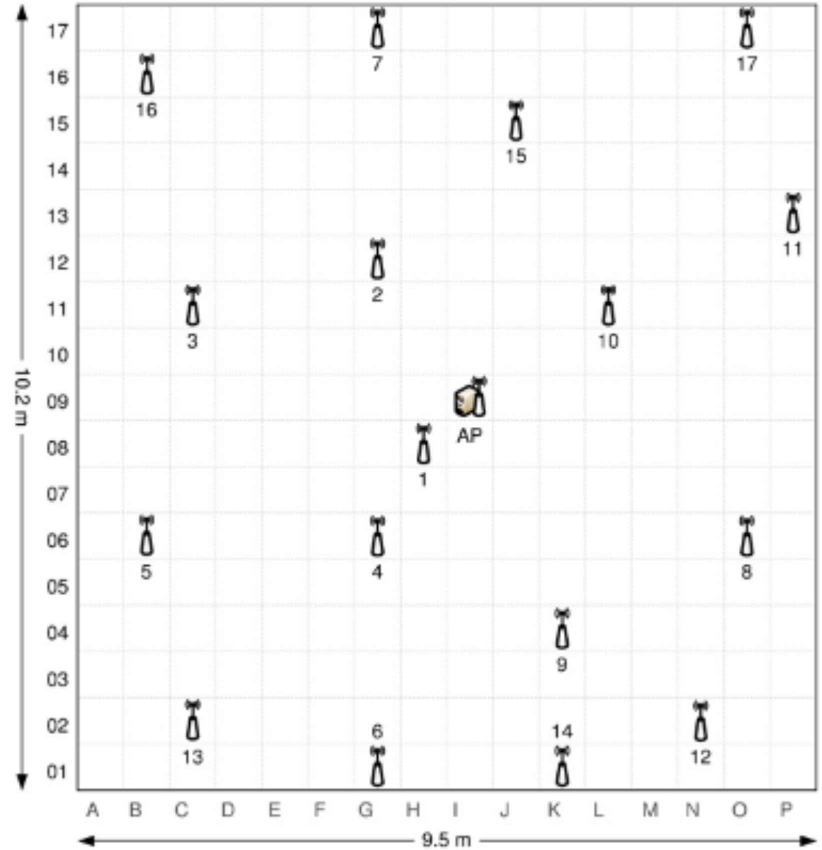

Results: CAC & DAC

Point of operation

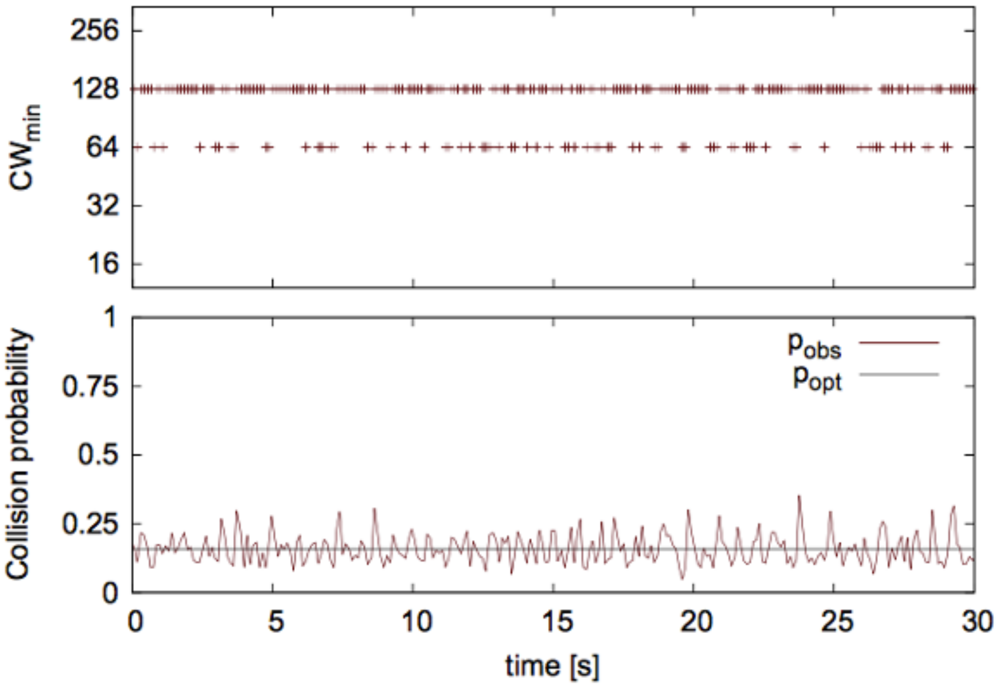

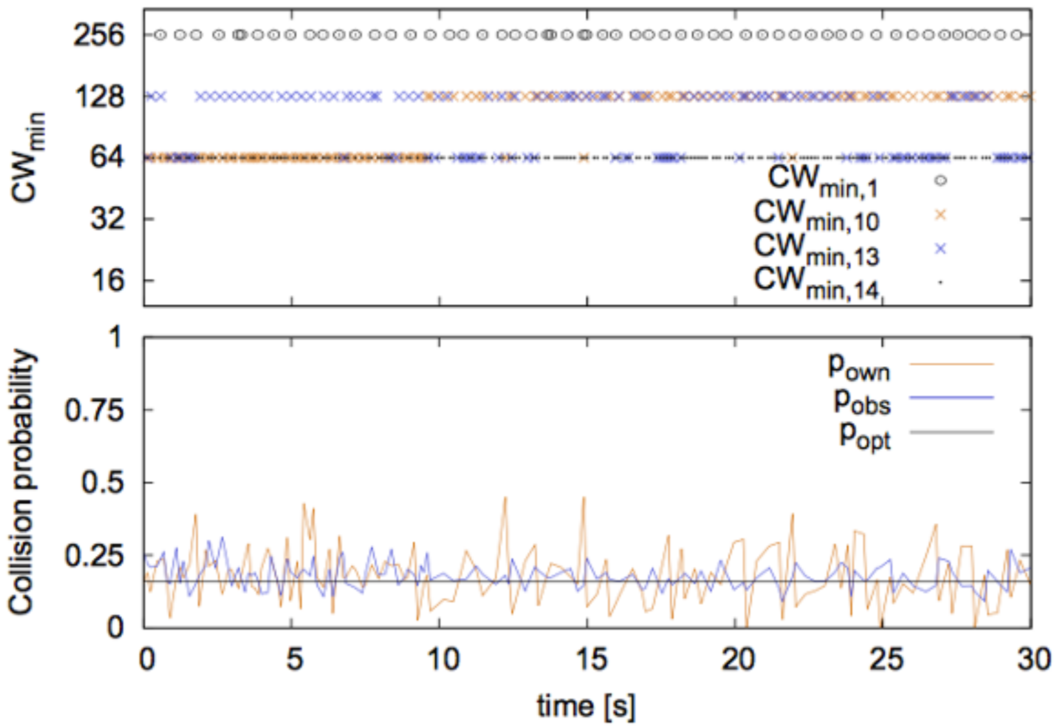

Results

Total and per-station throughput

Mechanism	JFI
EDCA	0,787
CAC	0,996
DAC	0,692

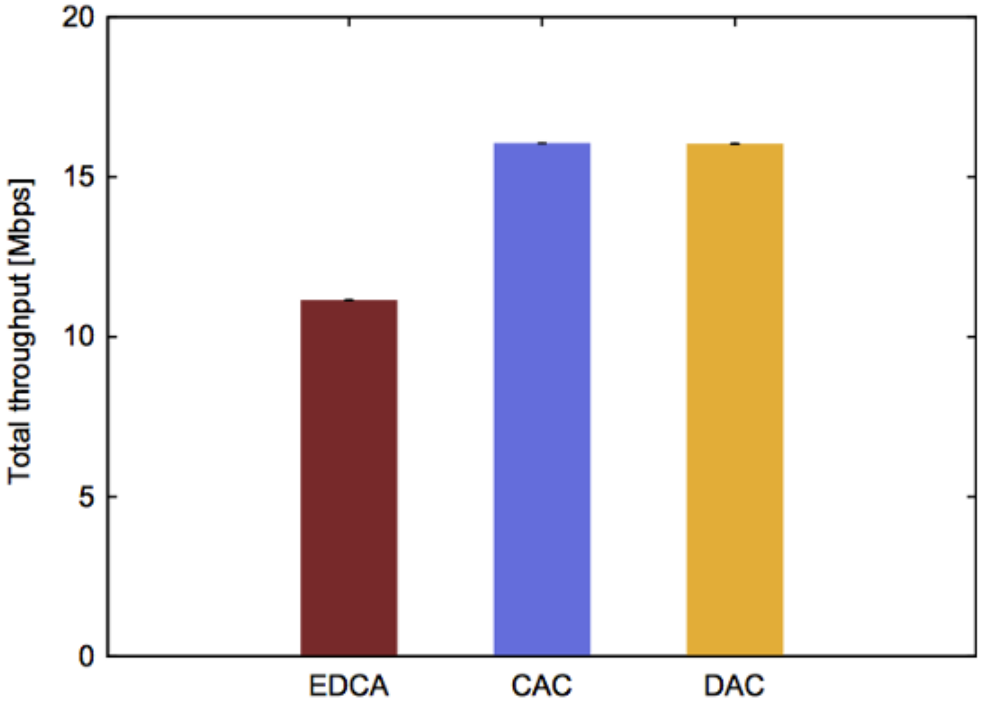

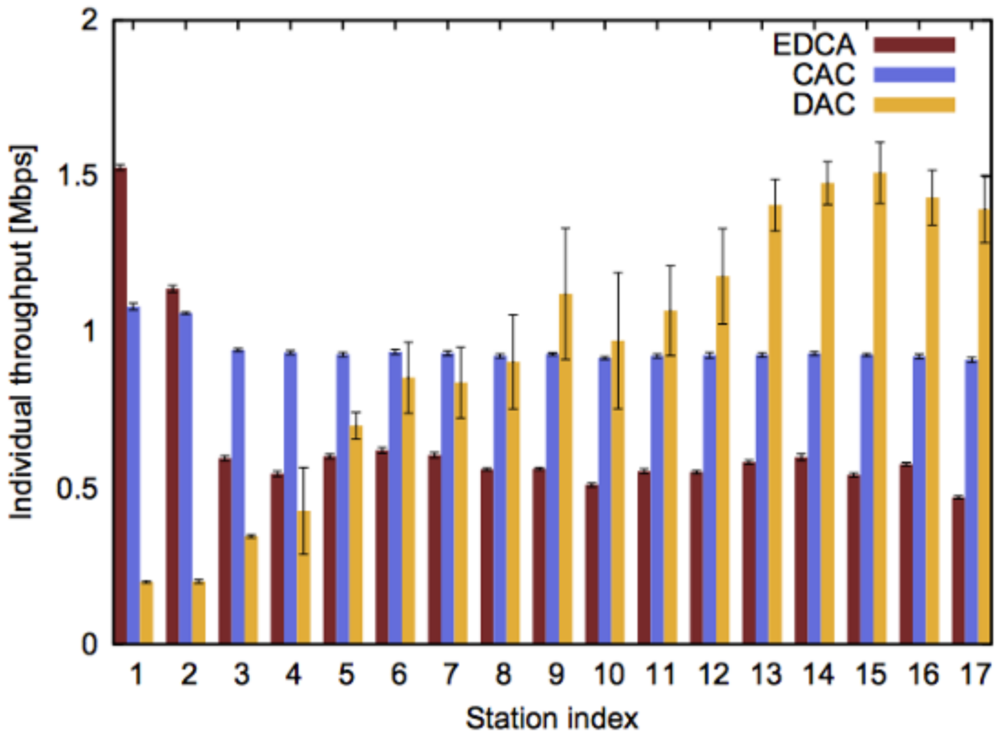

Results

Non-saturation conditions

P. Serrano, P. Patras, A. Mannocci, V. Mancuso, Albert Banchs, Control Theoretic Optimization of 802.11 WLANs: Implementation and Experimental Evaluation, Elsevier Computer Networks, vol. 57, no. 1, January 2013, http://dx.doi.org/10.1016/j.comnet.2012.09.010

- Poisson traffic
- 10 MB TCP tranfers
- Every 30 s (average)

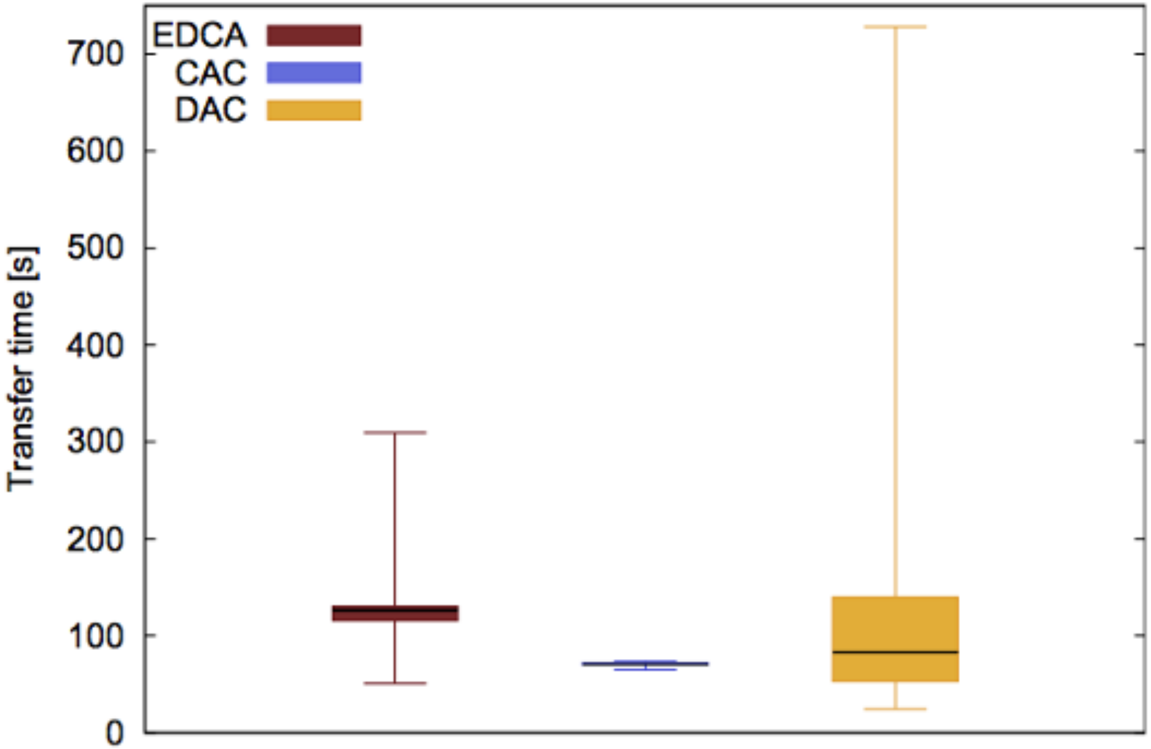

(a) λ^{-1} = 30 [s]

Maximizing VoIP capacity using derived experimental rules for 802.11e

P. Serrano, A. Banchs, J. F. Kukielka, G. D'Agostino, S. Murphy, Configuration of IEEE 802.11e EDCA for Voice and Data traffic: An Experimental Study, ICT-MobileSummit 2008, 10-12 June 2008, Stockholm, Sweden

Motivation

Voice ≠ Data

- EDCA to improve traffic differentiation
- Set of recommended values, but these can be changed
- Voice has different requirements than data
    - Delivery on time
- Develop a methodology to derive easy configuration rules

Setup
Phy: IEEE 802.11b

- Atheros card + MADWIFI driver, iPerf (80 bytes / 10 ms, no silence)
- Parameters: CWmin, CWmax, AIFS (SIFS+A∂), TWOP (us)

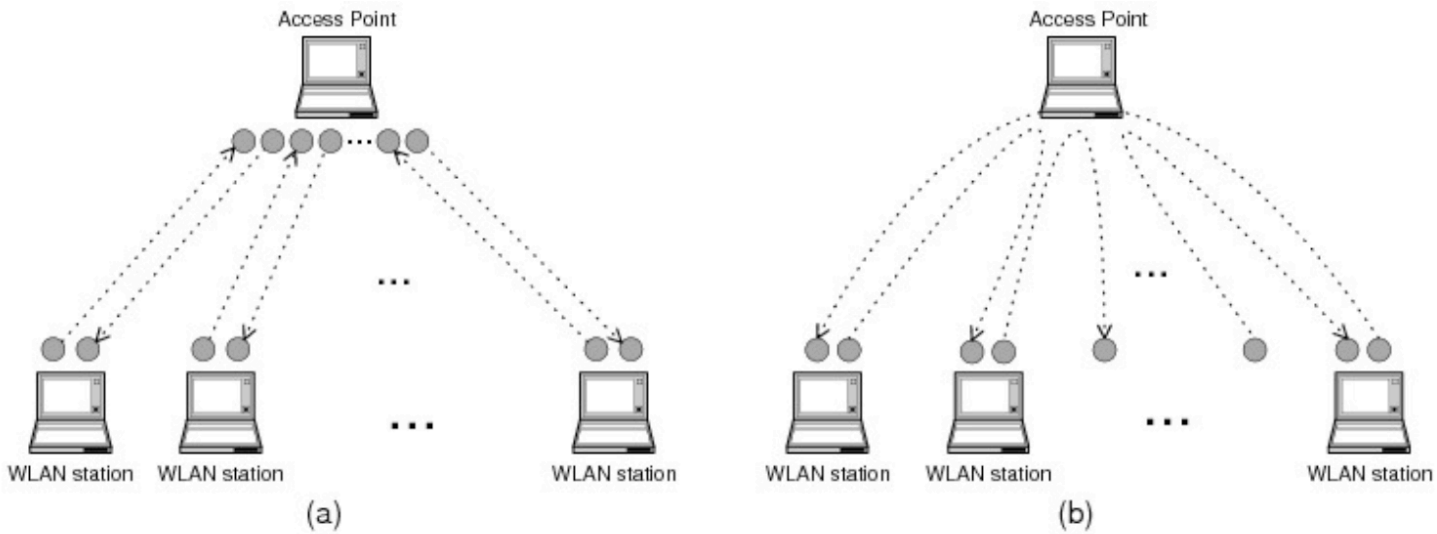

Benchmark
Default DCF configuration

- ITU-T G.114 specifies ≤ 150 ms with < 5% loss
- We assume 3 hops, WiFi + Wired + WiFi
- Criterion: delay ≤ 50 ms w.p. 98%
- Methodology: 5 runs, 60 s/run
- **DCF: up to 8 calls**

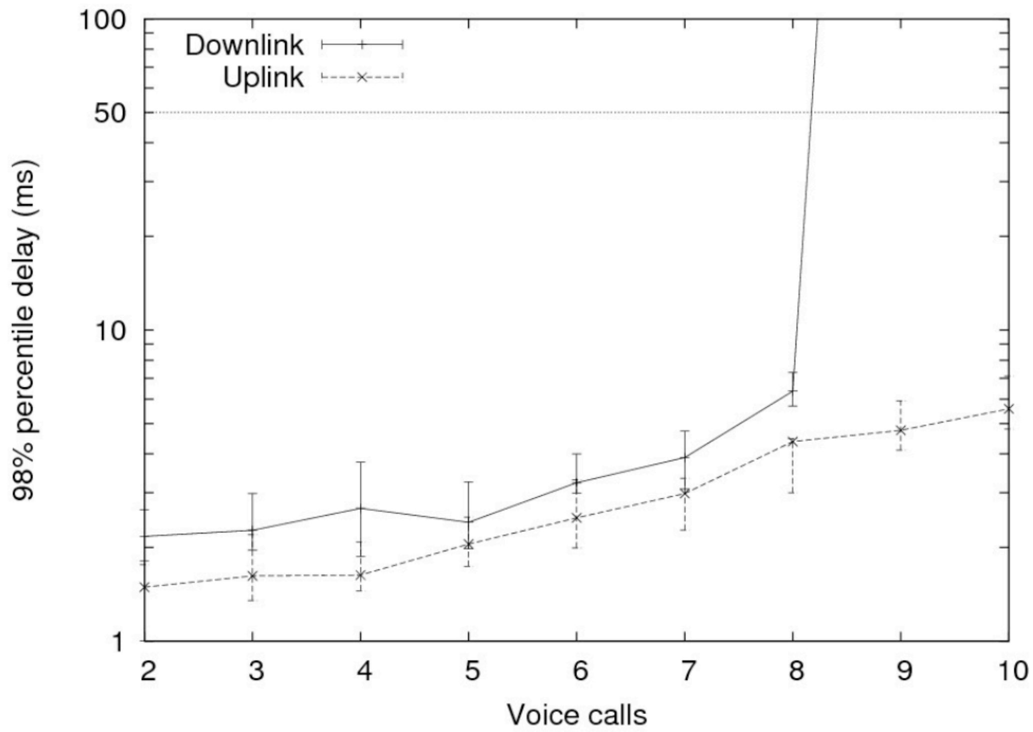

Real Time (RT) and Non RT configuration
Approach

- A combination of reasoning and experiments
- Objective: add more voice stations, then maximize throughput (for non RT)
- Parameters
  - AIFS
  - TXOP
  - CWmin and CWmax
- For Real Time 1st, and then Non Real Time

Real Time conf.: AIFS and TXOP

Set by reasoning

- AIFS: priority to access the channel
  - Minimum value: 2

Rule 1: $A_{voice} = 2$

- TXOP: increase it to empty the queue
  - Up to 2 more stations (from 8 to 10)

Rule 2: $TXOP_{limit\ voice} = 65535$ us

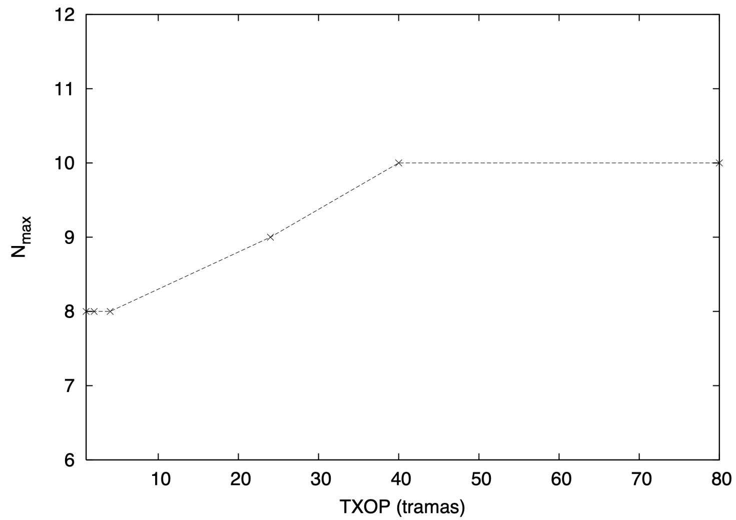

Real Time conf.: CW setting

Exhaustive search

- No Binary Exponential Backoff (would increase delay)

Rule 3: $CW_{max}(voice) = CW_{min}(voice)$

- Search on AP and STA for n=11 stations (maximum)

	CW_{min}(voice, AP)					
CW_{min}(voice, STA)	4	8	16	32	64	128
4	-	-	-	-	-	-
8	287.55	-	-	-	-	-
16	105.79	102.96		-	-	-
32	33.61	45.87	36.3	69.63	-	-
64	23.4	29	25.08	$\boxed{22.56}$	-	-
128	28.08	29.53	29.61	26.94	-	-
256	39.76	44.97	47.77	43.48	42.67	-
512	69.18	68.54	76.43	73.45	-	-

Rule 5: Admit no more than 11 voice conversations

Non Real Time configuration

First rules

- No need to empty queues

$$\text{Rule 6: } TXOP_{limit}(\text{data}) = 0$$

- Binary Exponential Backoff could help

$$\text{Rule 7: } CW_{max}(\text{data}) = CW_{min}(\text{data})\ 2^5$$

Non Real Time configuration
About AIFS and CWmin

- Both parameters regulate the aggressiveness
- 1st experiment: with n=11 voice stations and 1 data station, two approaches
  - Use minimum AIFS, find CWmin to guarantee QoE
  - Use maximum AIFS, find CWmin to guarantee QoE
- Results {AIFS, CWmin} = {(2, 128), (15, 16)}

Non Real Time configuration

AIFS and CWmin, cont'd

- One data station
- Use both configurations for an increasing number of voice calls
- Measure throughput

Rule 8: A_{data} = 15 and CW_{min}(data)=16

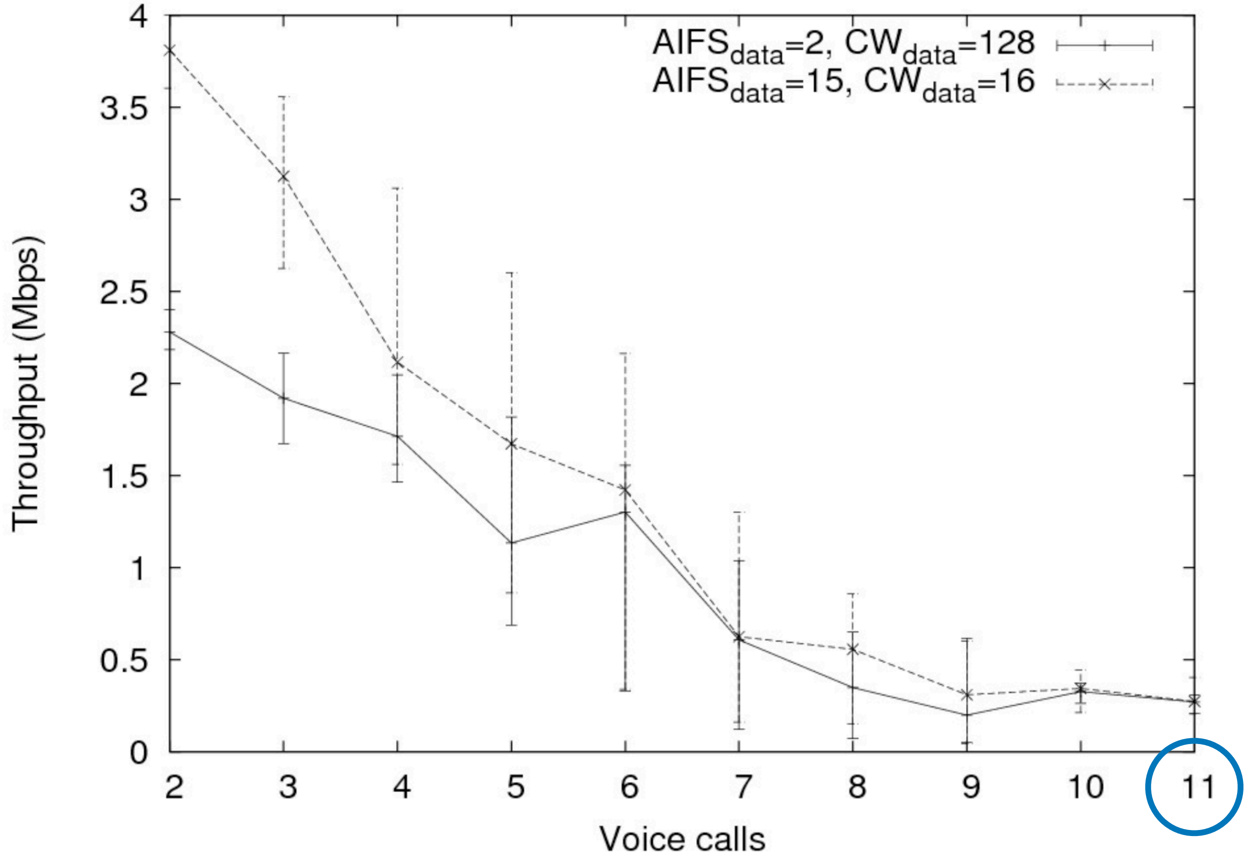

Non Real Time configuration

More than 1 data station

- Doubling the number of stations ~ halving the CWmin
- With n data stations, set the data to CWmin = 16 x n
- Experiment: n=11 calls, increasing number of data stations

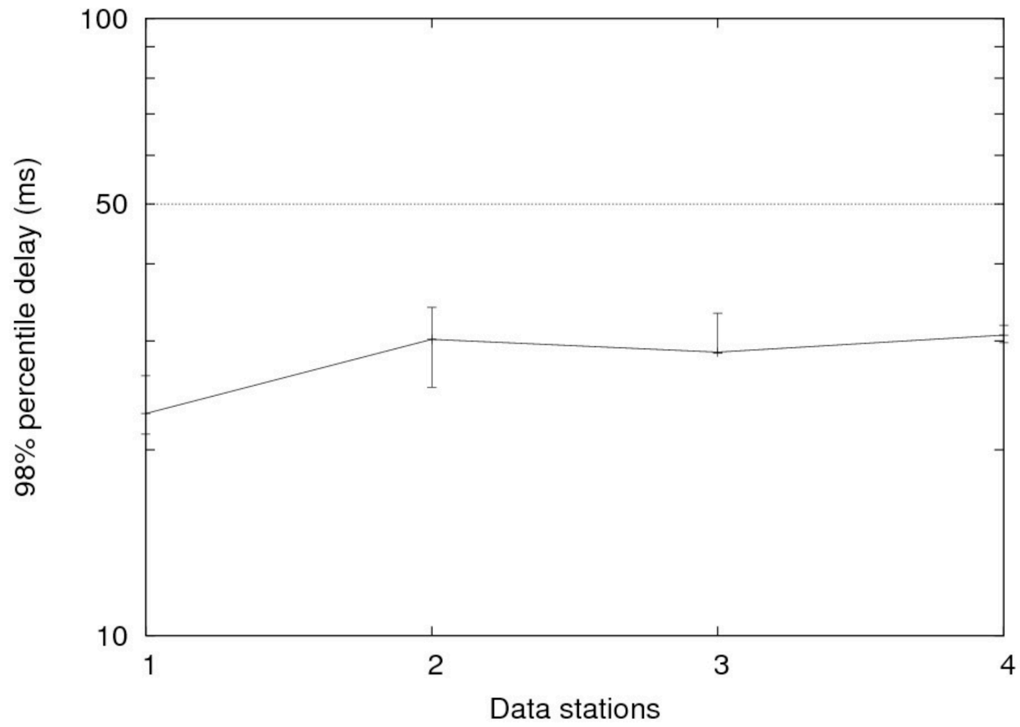

Comparison vs. EDCA

Recommended parameters

- With no data (figure)
  - EDCA, shorter delays until 9 calls (+1 vs DCF)
  - Our proposal: 2 more conversations
- With data (not shown)
  - EDCA: only 7 calls
  - Our proposal: 11 (as before)

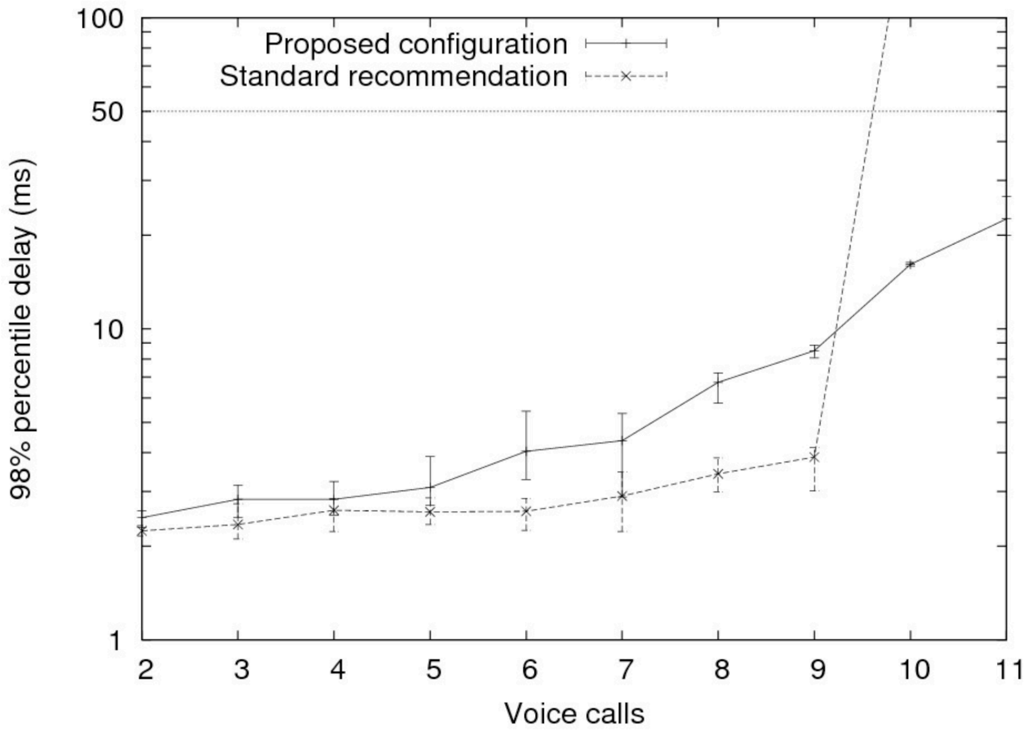

MAC Layer Innovation: Introducing VoIPiggy to double capacity, improving QoE using 802.11aa GATS

VoIPiggy
Motivation

P. Salvador, V. Mancuso, P. Serrano, F. Gringoli, A. Banchs, VoIPiggy: Analysis and Implementation of a Mechanism to Boost Capacity in IEEE 802.11 WLANs Carrying VoIP traffic, IEEE Transactions on Mobile Computing, vol. 13, no. 7, July 2014, pp. 1640–1652., http://dx.doi.org/10.1109/TMC.2013.114

- In case of voice traffic: two flows (UL and DL) of small packets
- This is very inefficient
- Approach
    - After receiving a DL frame
    - Wait for UL frame
    - Piggyback
    - Omit the ACK (since the exchange is after a SIFS)

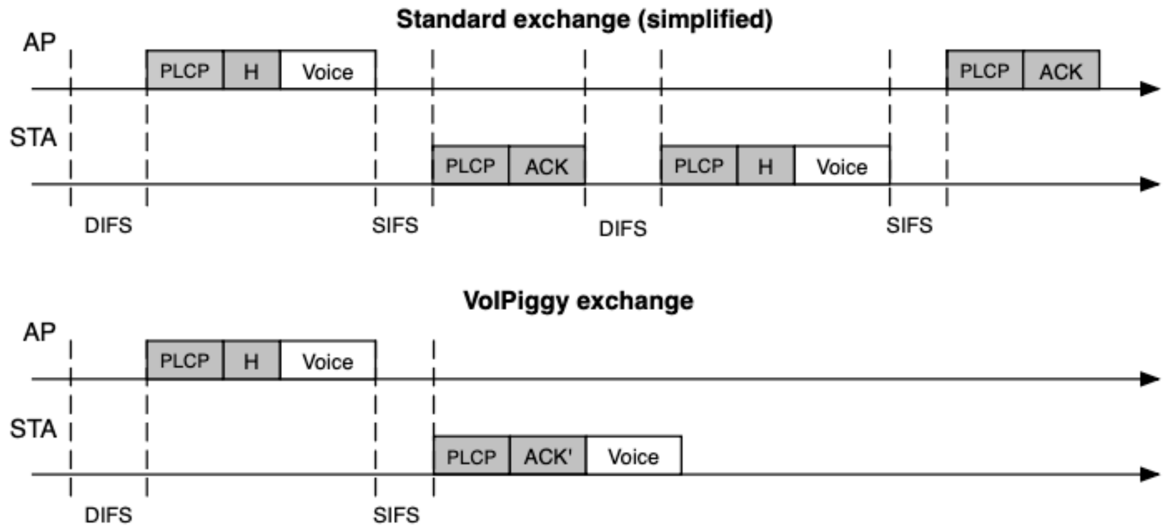

Performance evaluation
Testbed description

- One desktop acting as AP, center of the testbed, 7 dBi
- 30 Alix 2d2 devices acting as clients, 2Bi
- Mgen as traffic emulator, 60 B or 160 B / 20 ms
- PTP for sync

Performance evaluation

Voice-only Scenarios

Model: nD/D/1 queue

Voice codec	MCS (Mbps)	n_v experimental EDCA			n_v model VoIPiggy	
		EDCA	VoIPiggy	Gain	VoIPiggy	Gain
G. 711	2	5	9	80%	9	80%
	5.5	10	18	80%	18	80%
	11	12	26	116%	26	116%
	6	19	29	53%	29	53%
	9	24	≥ 30	-	49	71%
	12	26	≥ 30	-	52	100%
	54	≥ 30	≥ 30	-	125	-
G. 726	2	8	14	75%	14	75%
	5.5	10	26	160%	26	160%
	11	12	≥ 30	-	33	175%
	6	23	≥ 30	-	49	113%
	9	26	≥ 30	-	66	153%
	12	≥ 30	≥ 30	-	79	-
	54	≥ 30	≥ 30	-	154	-

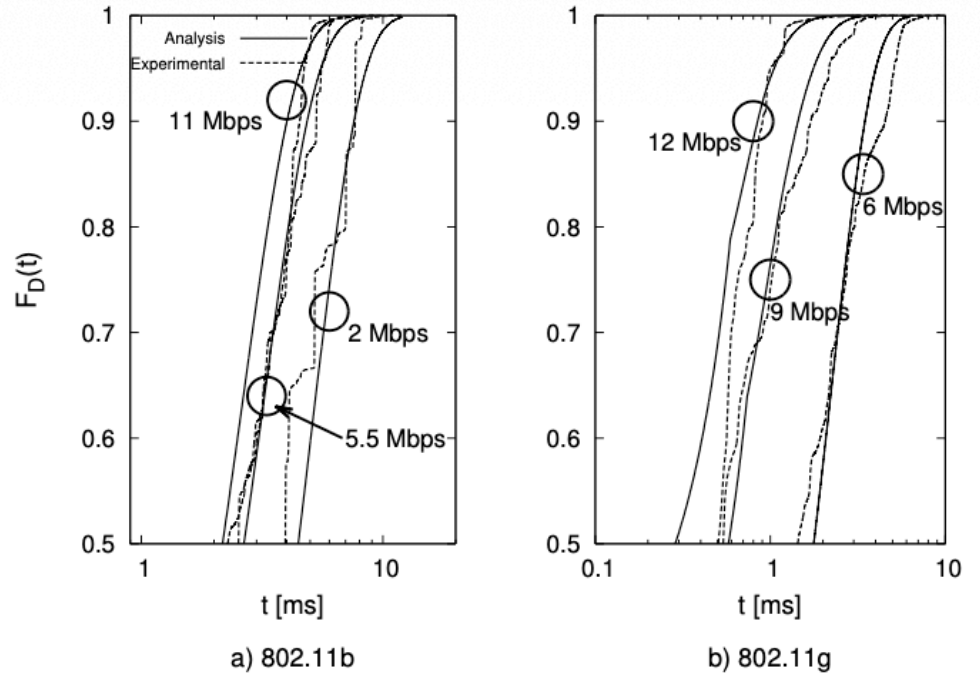

Fig. 7: CDF of the delay for voice-only scenarios.

Performance evaluation

Voice and Data

P. Salvador, V. Mancuso, P. Serrano, F. Gringoli, A. Banchs, VoIPiggy: Analysis and Implementation of a Mechanism to Boost Capacity in IEEE 802.11 WLANs Carrying VoIP traffic, IEEE Transactions on Mobile Computing, vol. 13, no. 7, July 2014, pp. 1640–1652., http://dx.doi.org/10.1109/TMC.2013.114

- One TCP station receiving data from the AP
- Stop when losses > 1% (EDCA)
- VoIPiggy achieves maximum # calls

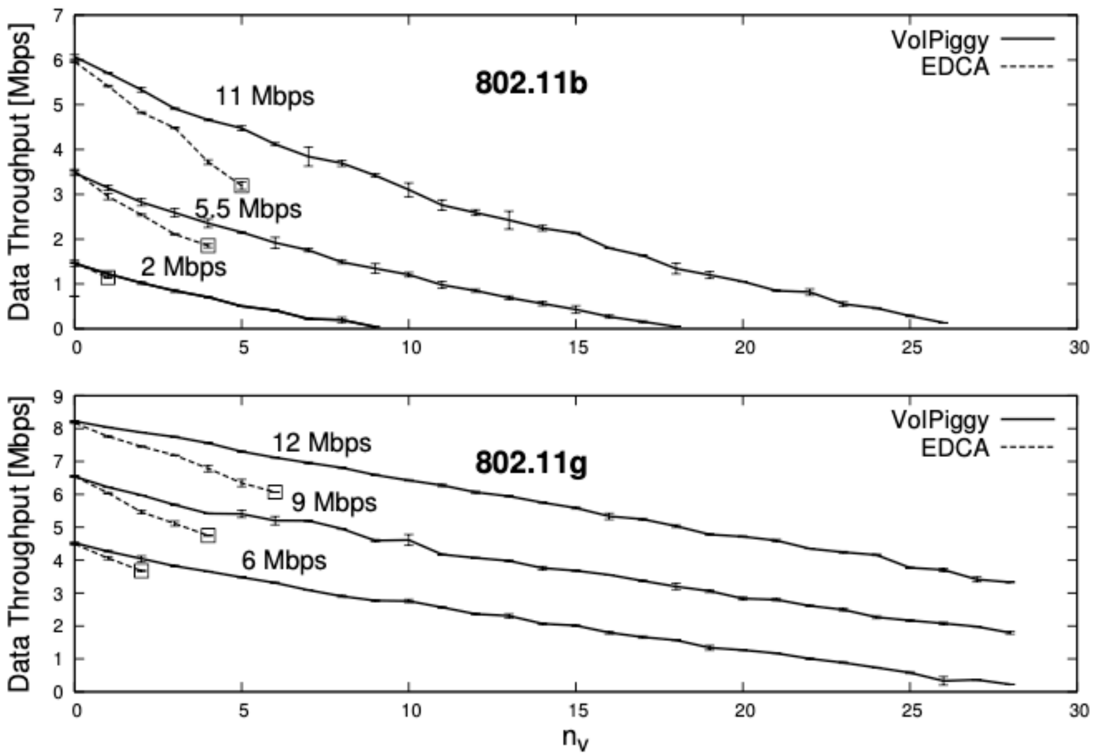

Evaluating three RT multicast schema

The case for RT multicast

- Wireless set-top boxes, Enhanced driving, Crowded scenarios

- Neither legacy Unicast nor Multicast are adequate

F. Gringoli, P. Serrano, I. Úcar, N. Facchi, A. Azcorra, Experimental QoE Evaluation of Multicast Video Delivery over IEEE 802.11aa WLANs, IEEE Transactions on Mobile Computing, Vol. 18, Issue 11, Nov. 2019, http://doi.org/10.1109/TMC.2018.2876000

QoE: emulating Mean Opinion Scores (MOS)

Netflix's VMAF

- Mean Opinion Scores (MOS) are impractical
  - We process 15k videos !
- Image QoE-oriented
  - MSE, PSNR
  - Structural Similarity (SSIM) index
- Netflix's VMAF: state of the art correlation with human perception

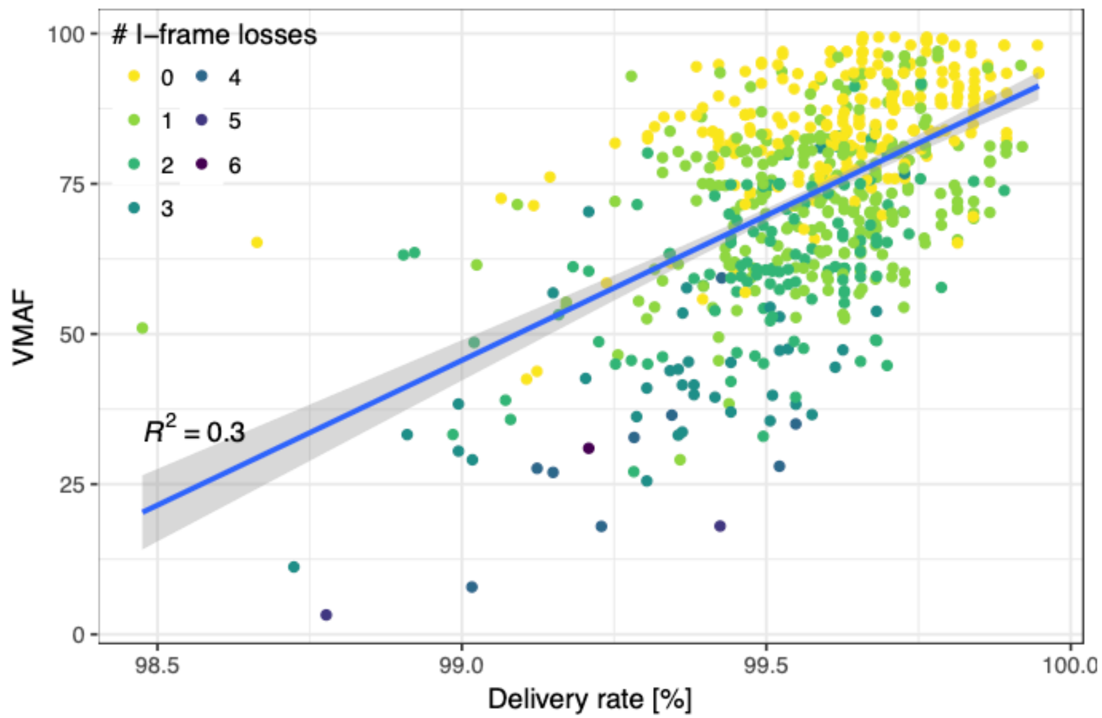

Mechanisms considered

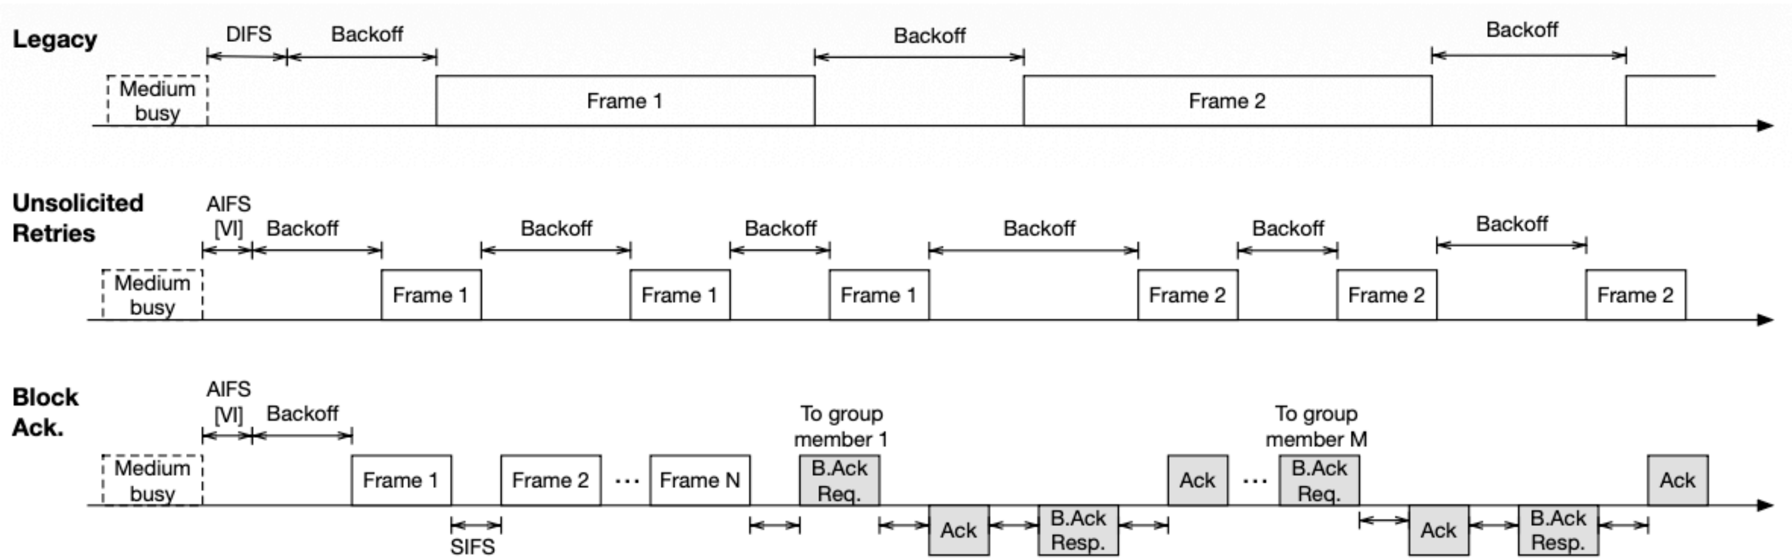

- Parameters: MCS, # retries

Testbed description

Initial assessment of legacy @ 6 Mbps

- 1 desktop (AP), 25 Alix 2d2 (STA), 9 videos from Video Quality Experts Group

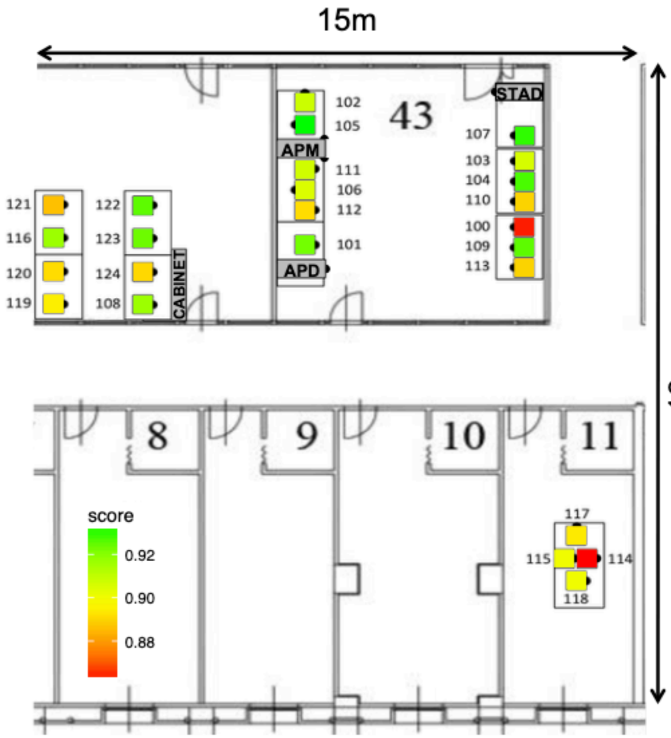

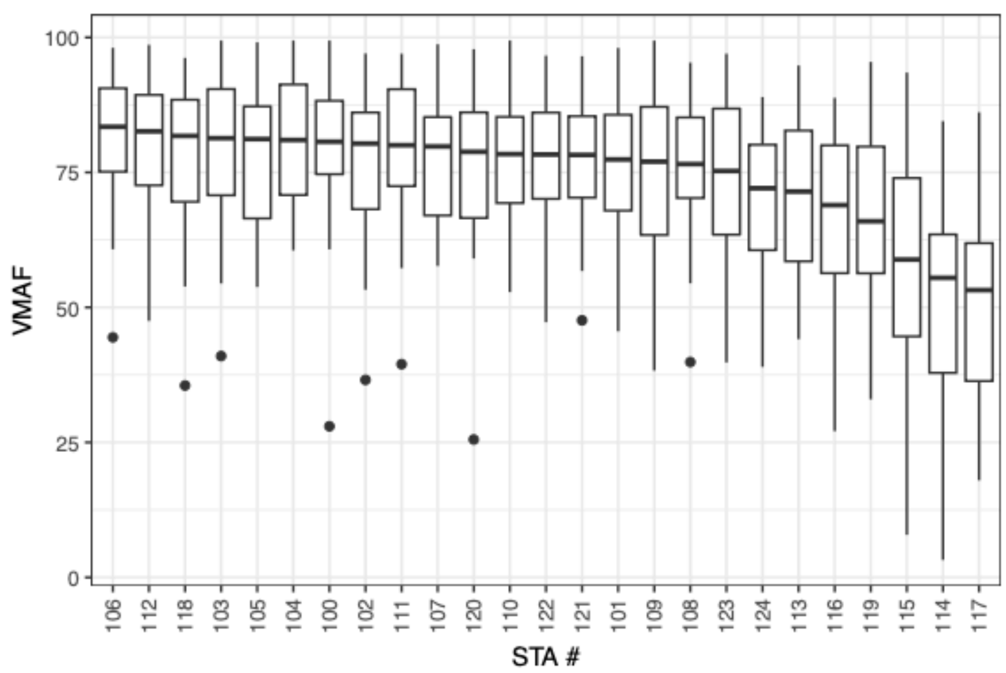

Performance assessment

Ideal conditions (no interference)

$$\text{Guaranteed VMAF} = Q_5(\mathbb{S}_j)$$

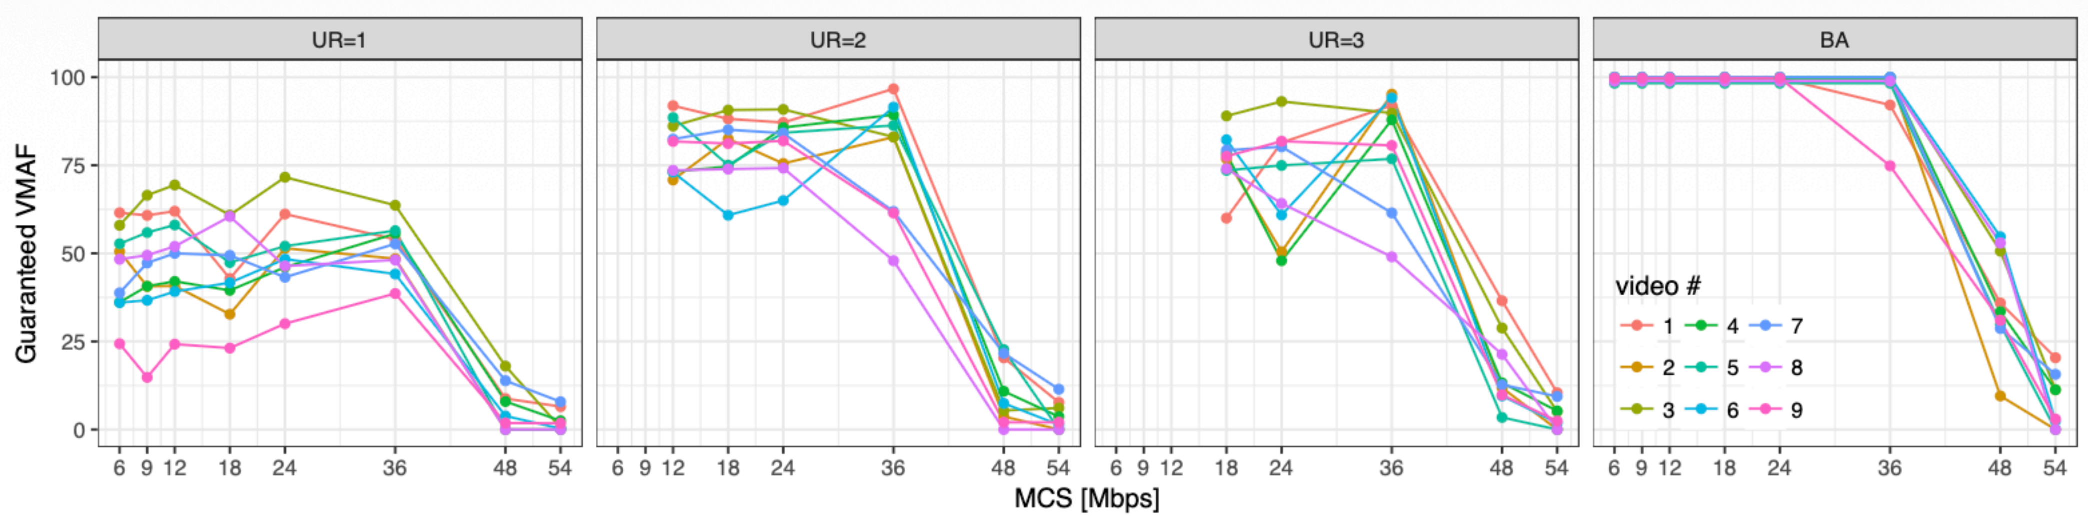

Performance assessment

Controlled interference - Unsolicited Retries

- One STA sending 1500 B UDP traffic using iperf at 54 Mbps to a different AP

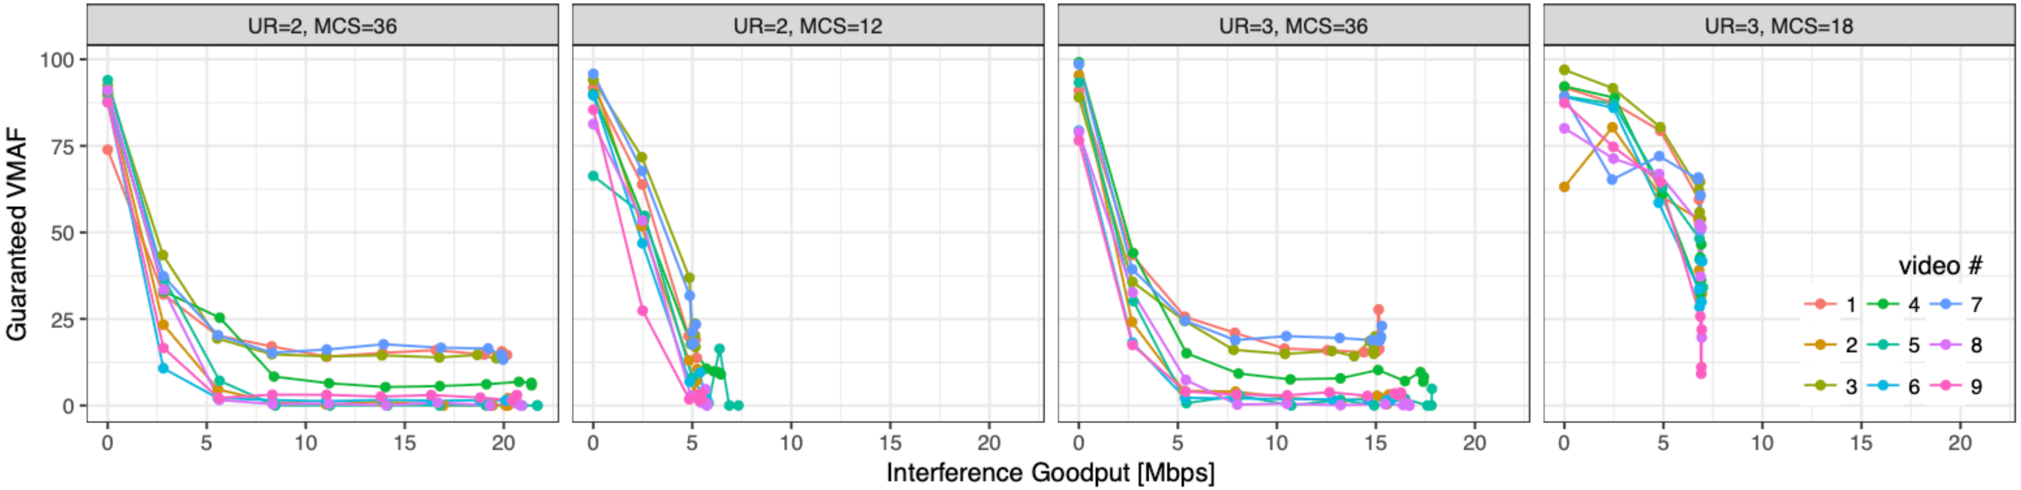

Performance assessment

Controlled interference - Block acknowledgment

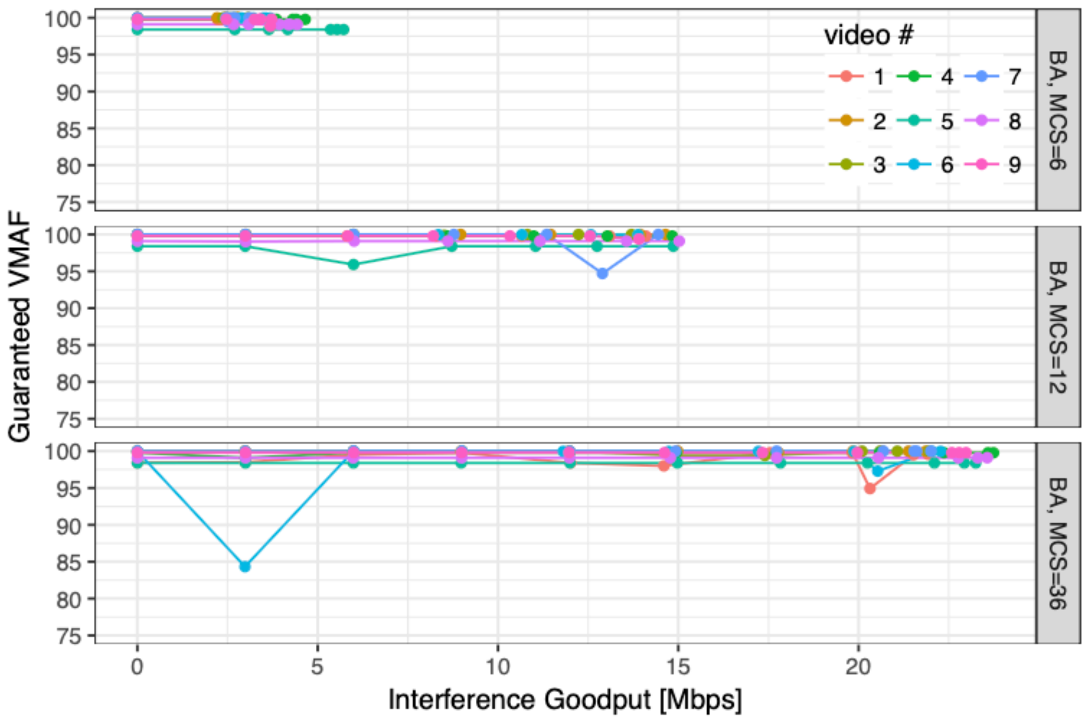

Performance anomaly

Performance assessment

TCP interference

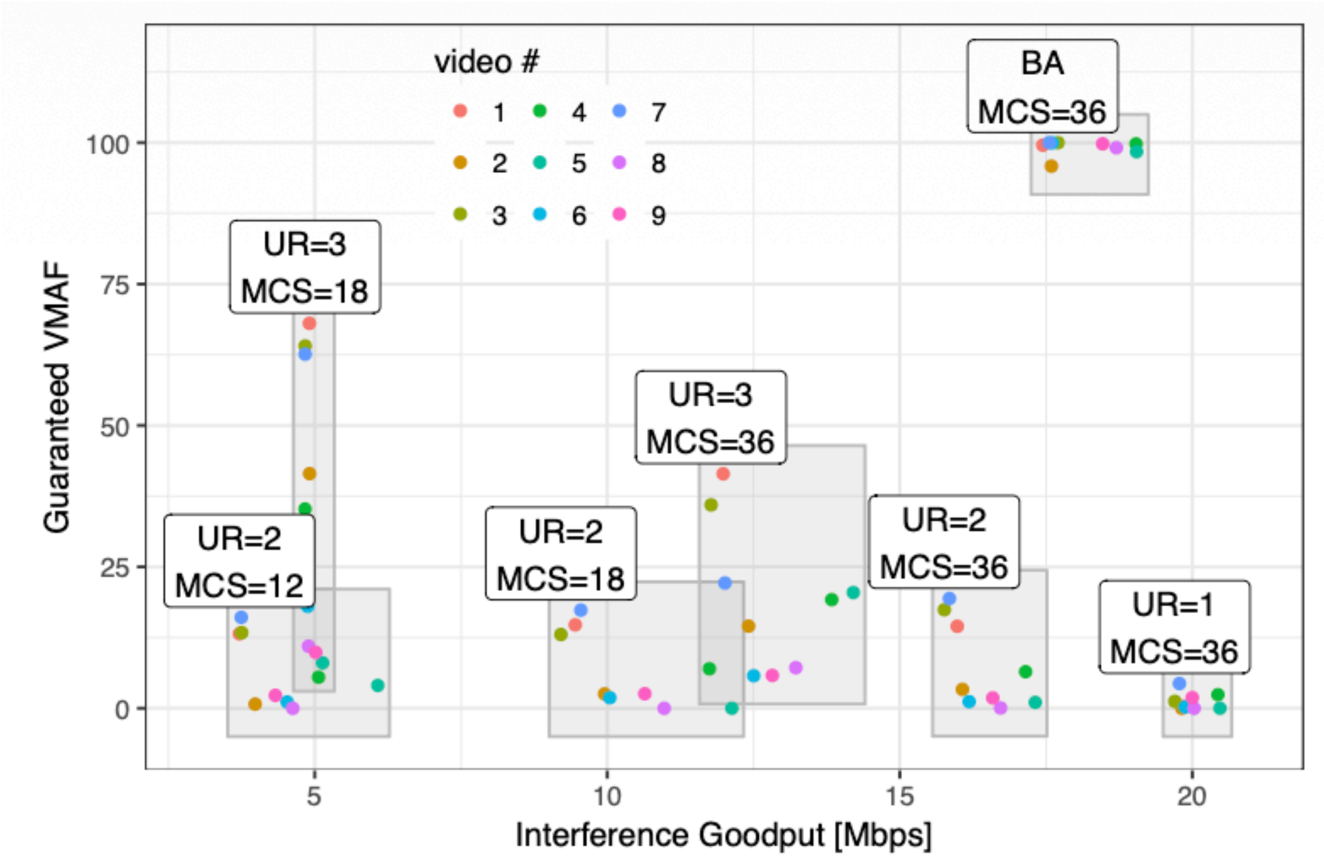

Performance asessment

Performance in the wild

- Chose a video at random
- Chose a channel in round robin
  - Channel = {1, 6, 11, 14}
  - Send video, compute VMAF
- Repeat during 24 h, working day

- Figures: VMAF per station (dots), average SSIM (lines)

Performance evaluation

Legacy multicast

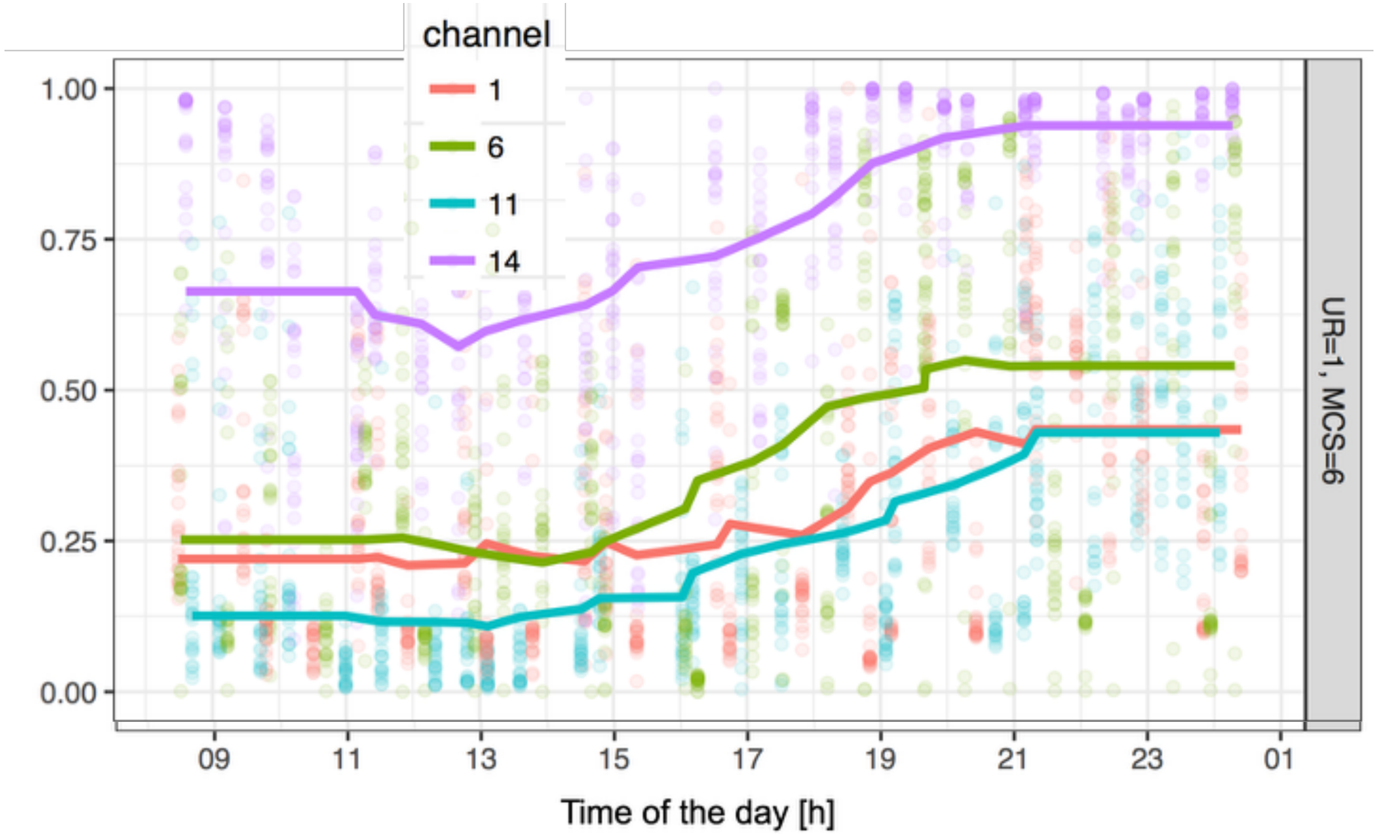

Performance evaluation

UR 2, 3, and BA

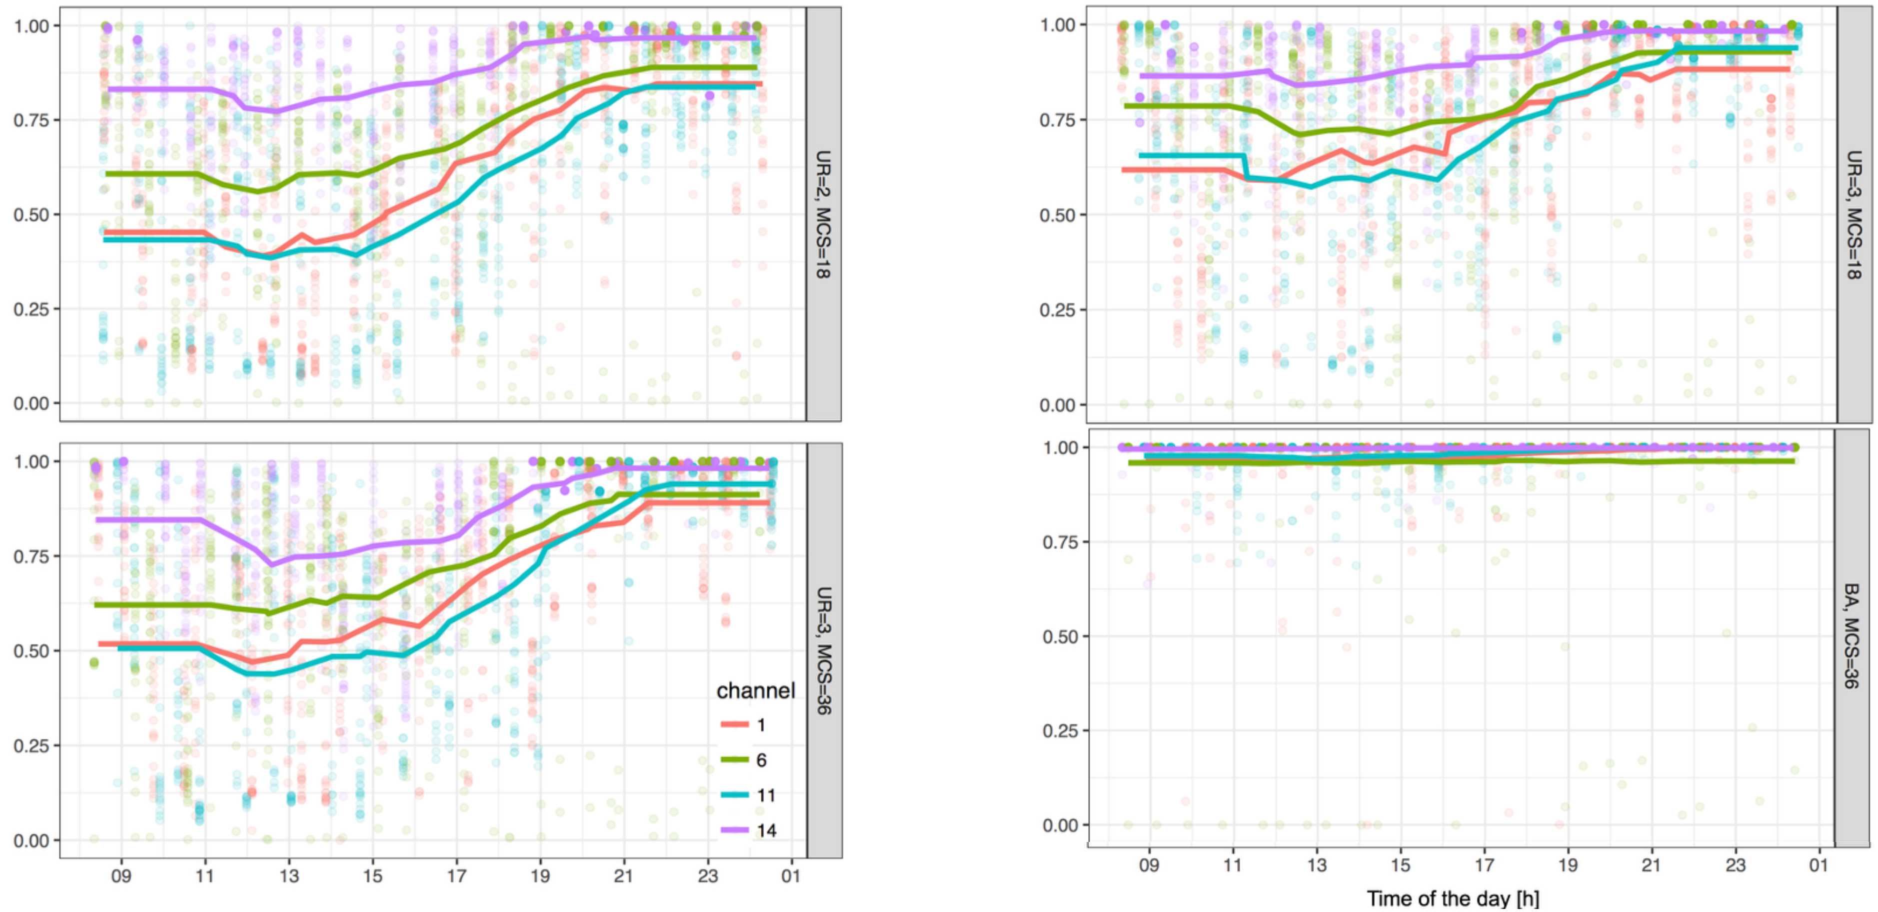

Optimizing 802.11 WLANs

Key Insights for High-Performance MAC Configurations
Leveraging Experimental Studies for Capacity and QoE

Pablo Serrano (pablo@it.uc3m.es), December 10th, 2025