GreenABR: Energy-Aware Adaptive Bitrate Streaming with Deep Reinforcement Learning

Bekir Oguzhan Turkkan[*], Ting Dai[+], Adithya Raman[*], Tevfik Kosar[*], Changyou Chen[*], Muhammed Fatih Bulut[+], Jaroslaw Zola[*], Daby Sow[+]

[*]*University at Buffalo* , [+]*IBM Research*

University at Buffalo
The State University of New York

Introduction

Global Traffic
60%

Mobile Traffic
79%

- Diverse clients and streaming conditions
- HTTP Streaming – Standardized as DASH streaming
  - Adaptive Bitrate Algorithms (ABR)

⇧ Quality ⇩ Stalling ⇩ Oscillations

QoE

- Energy consumption is an important constraint for mobile users

State of the Art

QoE Targeted

- Rule-based Models
  - BOLA [INFOCOM '16][1]
  - Festive [CoNEXT '12][2]
  - MPC [SIGCOMM '15][3]
  - BOLAE, Dynamic ABR [MMSys '18][4]
- Learning-based Models
  - Pensieve [SIGCOMM '17][5]
  - Oboe [SIGCOMM '18][6]
  - Fugu [NSDI '20][7]
  - Comyco [MM '19][8]

Energy Aware

- Context aware
  - Chen et al. [ICDCS '19][9]
- Additional layer
  - ef-HAS [IEEE '18][10], e-DASH [WoWMoM '17][11]
- Using other codecs
  - Vitto et al. [ICMEW '18][12]
- Progressive download
  - EnDASH [IFIP '20][13], Breitbach et al. [CCNC '18][14], Meng et al. [ATC '21][15]
- Data/quality budget
  - QUAD [MMSys '19][16], DataPlanner [MMSys '21][17]

Limitations of Existing Studies

- Training stability and efficiency
  - Pensieve, Comyco, Oboe
- Do not consider the perceptual quality
  - Bola, Festive, MPC, BolaE, DynamicABR, Pensieve, Oboe, ef-HAS, e-DASH, EnDASH
- Either do not consider energy consumption
  - Bola, Festive, MPC, BolaE, DynamicABR, Pensieve, Oboe, Fugu, Comyco
- Or sacrifice QoE to save power consumption
  - ef-HAS, e-DASH, EnDASH

To Address Limitations

stability&efficiency → DQN

perceptual quality → VMAF

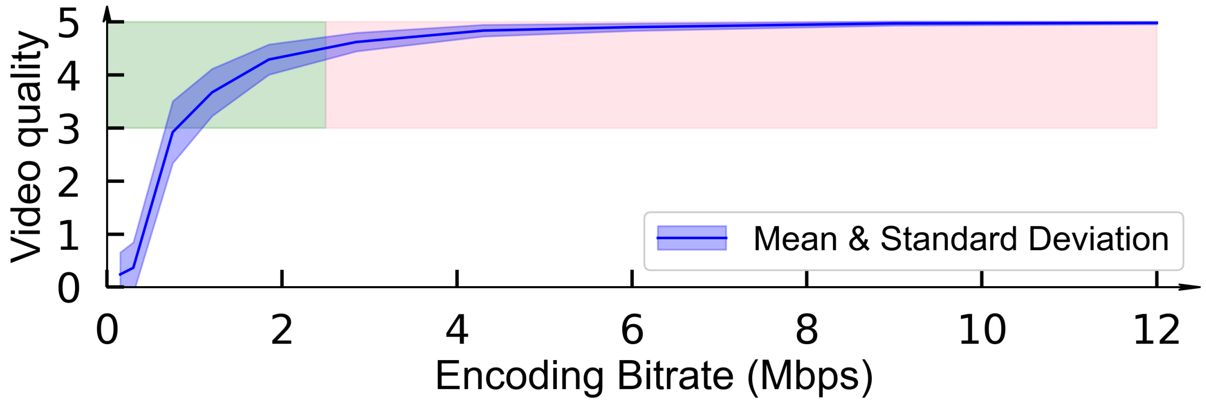

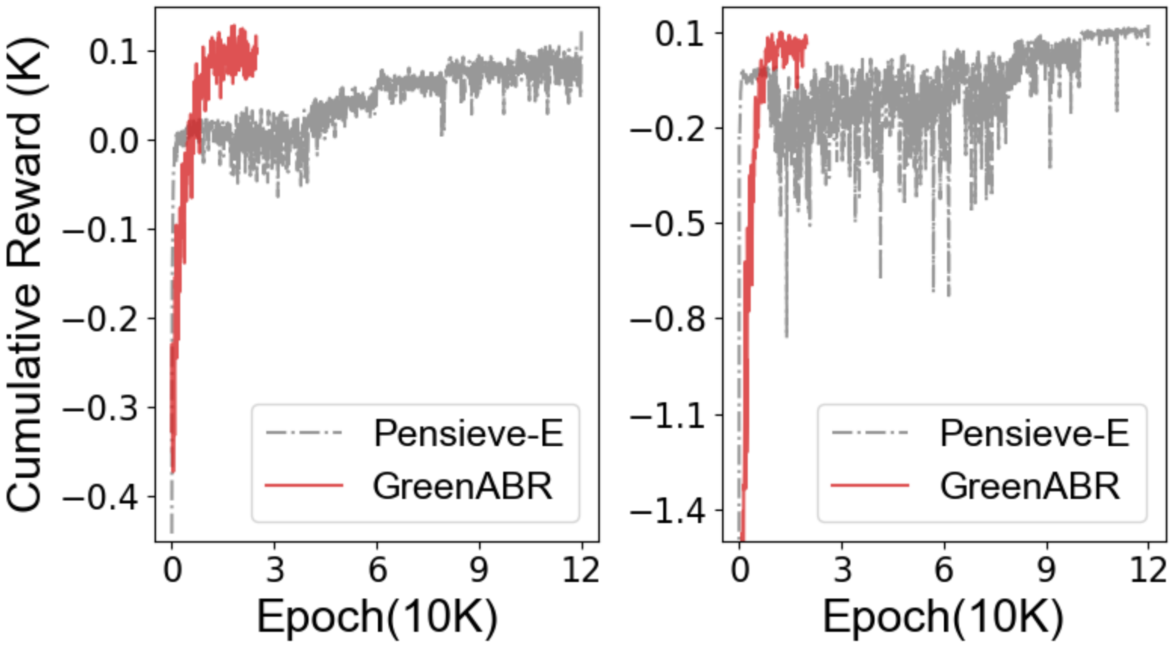

energy consumption, sacrifice QoE → **Optimize for both**

GreenABR: Energy-Aware Adaptive Bitrate Streaming with Deep Reinforcement Learning

Optimizes QoE and energy together

- Targets high perceptual quality
- Minimizes rebuffering events and quality oscillations
- Minimizes capacity violations
- Reduces the energy and data consumption
- Designed for general video streaming scenarios
- Does not require additional hardware support (e.g., HEVC support)

GreenABR Design

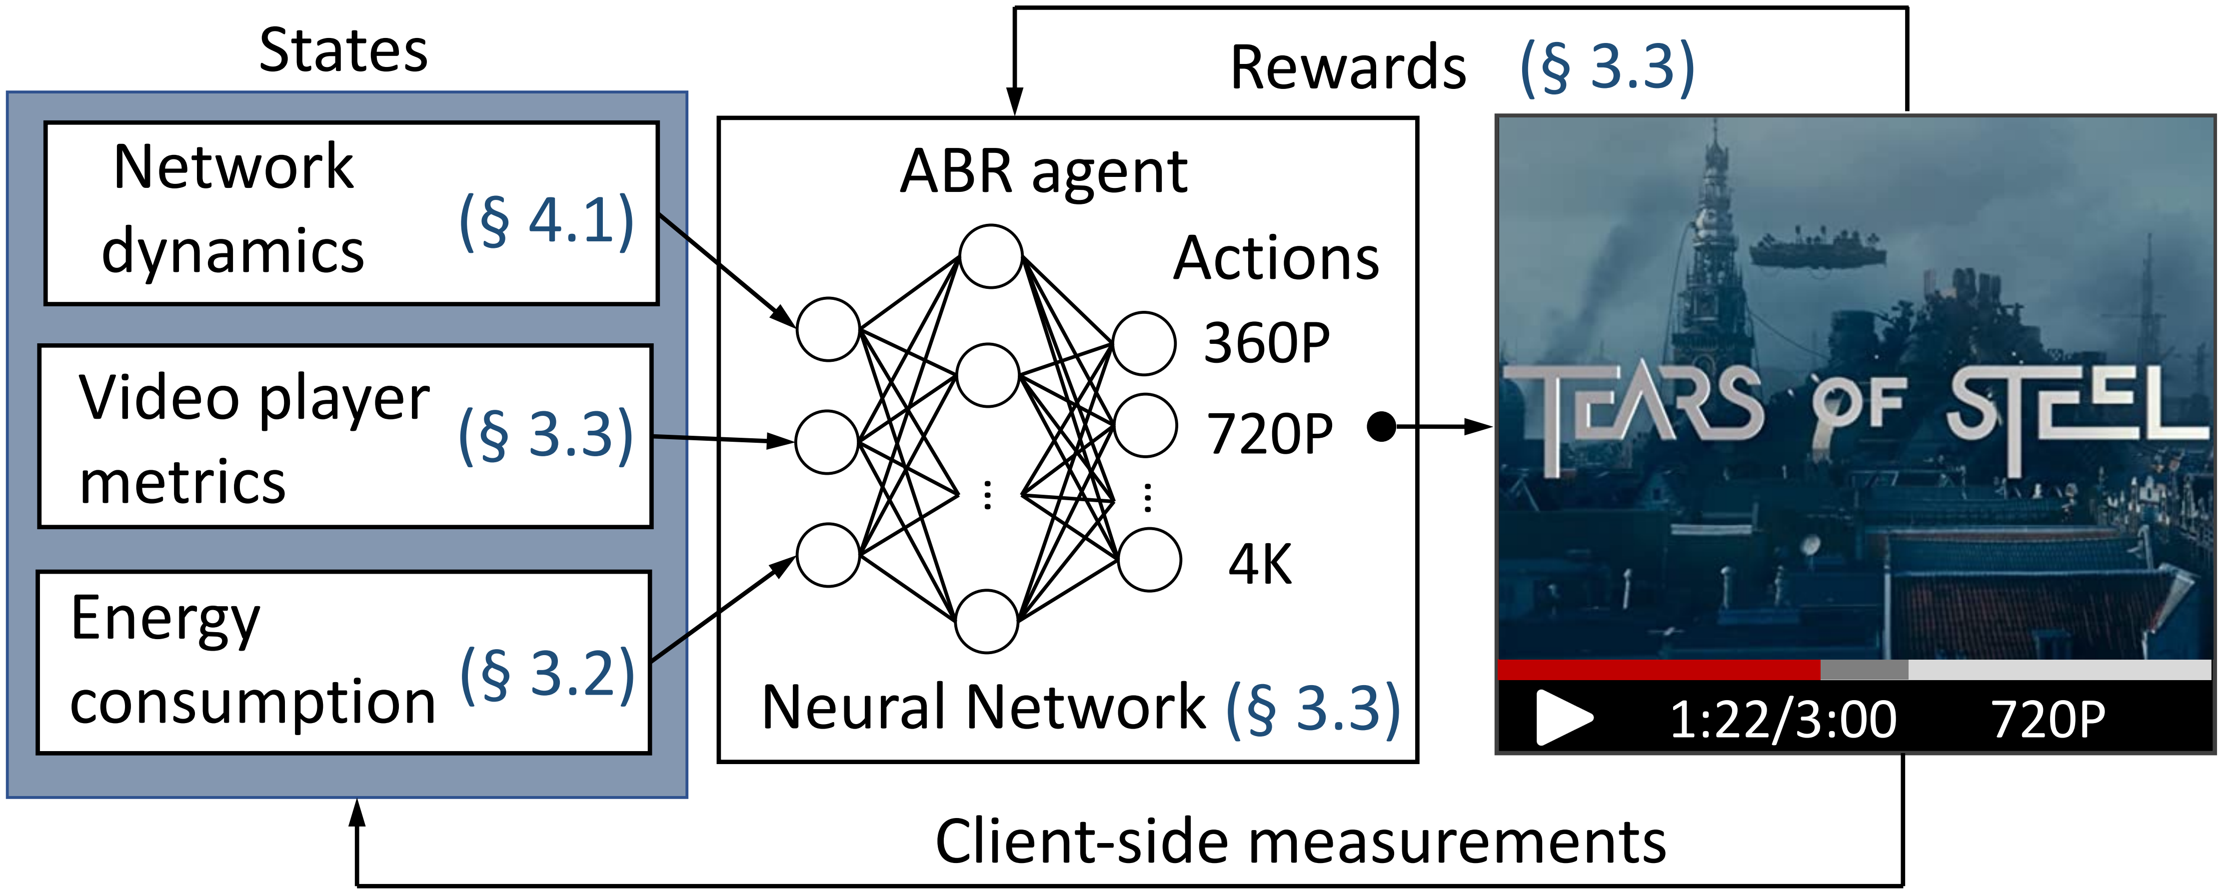

Modelling Power Consumption

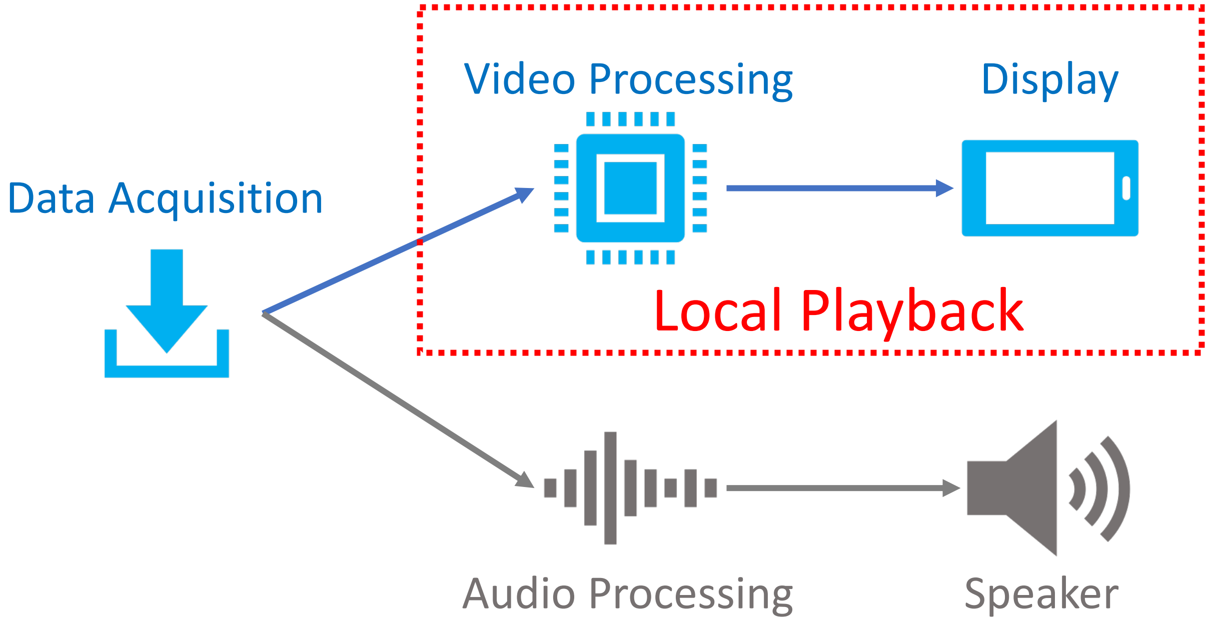

Limitations of existing models

- Designed based on encoding bitrate
- Device heterogeneity

Our approach

- Considers ABR-related components
- Targets to capture pattern
- Minimizes device specific impacts
  - Excludes base power consumption

Power Model

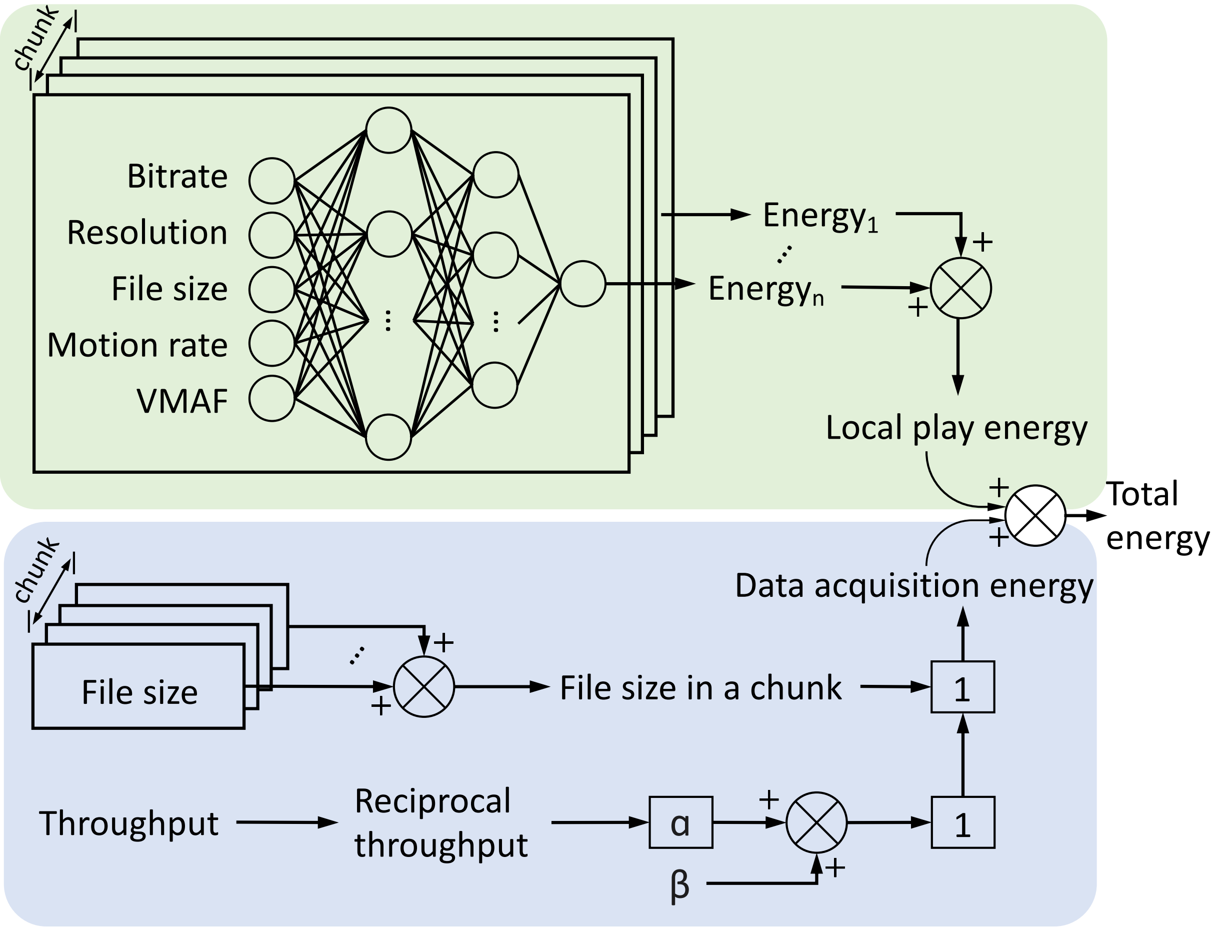

Local Play Energy

- Video processing + Display
- Parameters affect non-linearly
- Regression model
  - RMSE < 0.01 <7% estimation error
- Evaluations on unseen device
  - RMSE < 0.036 <12% estimation error

Data Acquisition Energy

- Adopt an existing model [18]
- Compatible throughput range
- Accurate with > 90%

DQN architecture in GreenABR

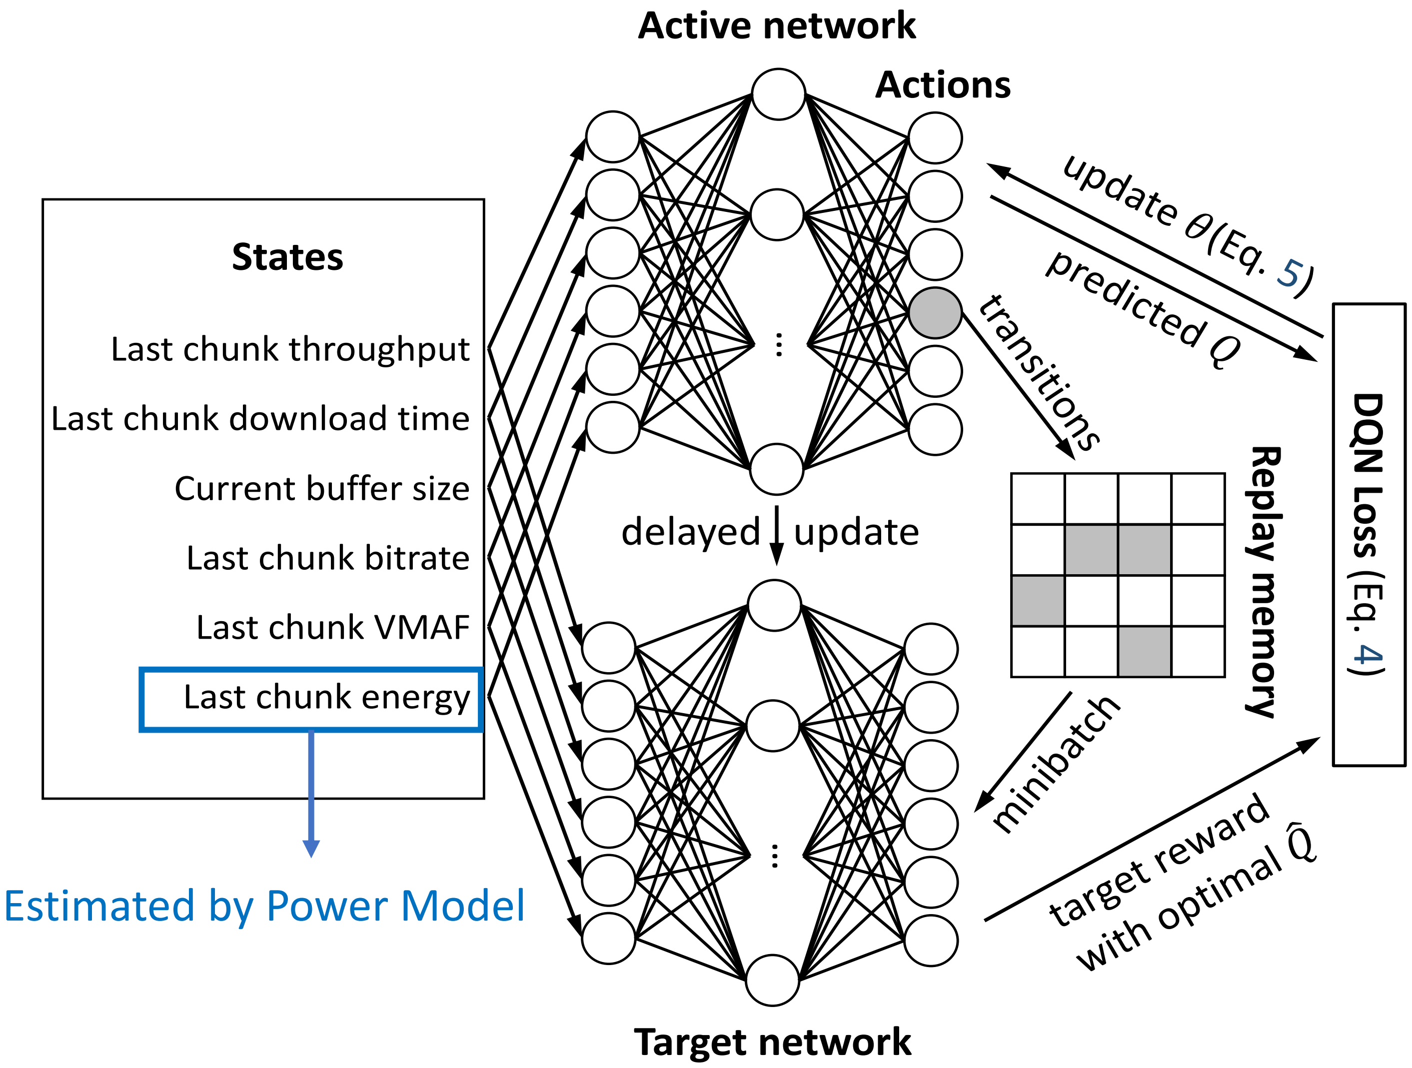

Reward Model

+ quality(VMAF)

- rebuffering(time)

- oscillations(VMAF)

- energy usage

Evaluation Setup

Videos

- Tears of steel (sci-fi), Nature(documentary), BBB(cartoon)
- 6-reps (0.3-4.3 Mbps), 10-reps (0.1-12Mbps), 10-HD-reps (2K&4K)

Energy Data Collection

- Monsoon Power meter collects at each 200 microseconds
- Results are aggregated for each chunk

Traces

- 3G, 4G-LTE, Broadband, Synthetic

Perceptual Quality Calculation

- QoE of the streaming session

video quality, rebuffering time and frequency, oscillations, quality switches

- Linear regression model on SQoE-III[19] dataset

Results - 6 representations

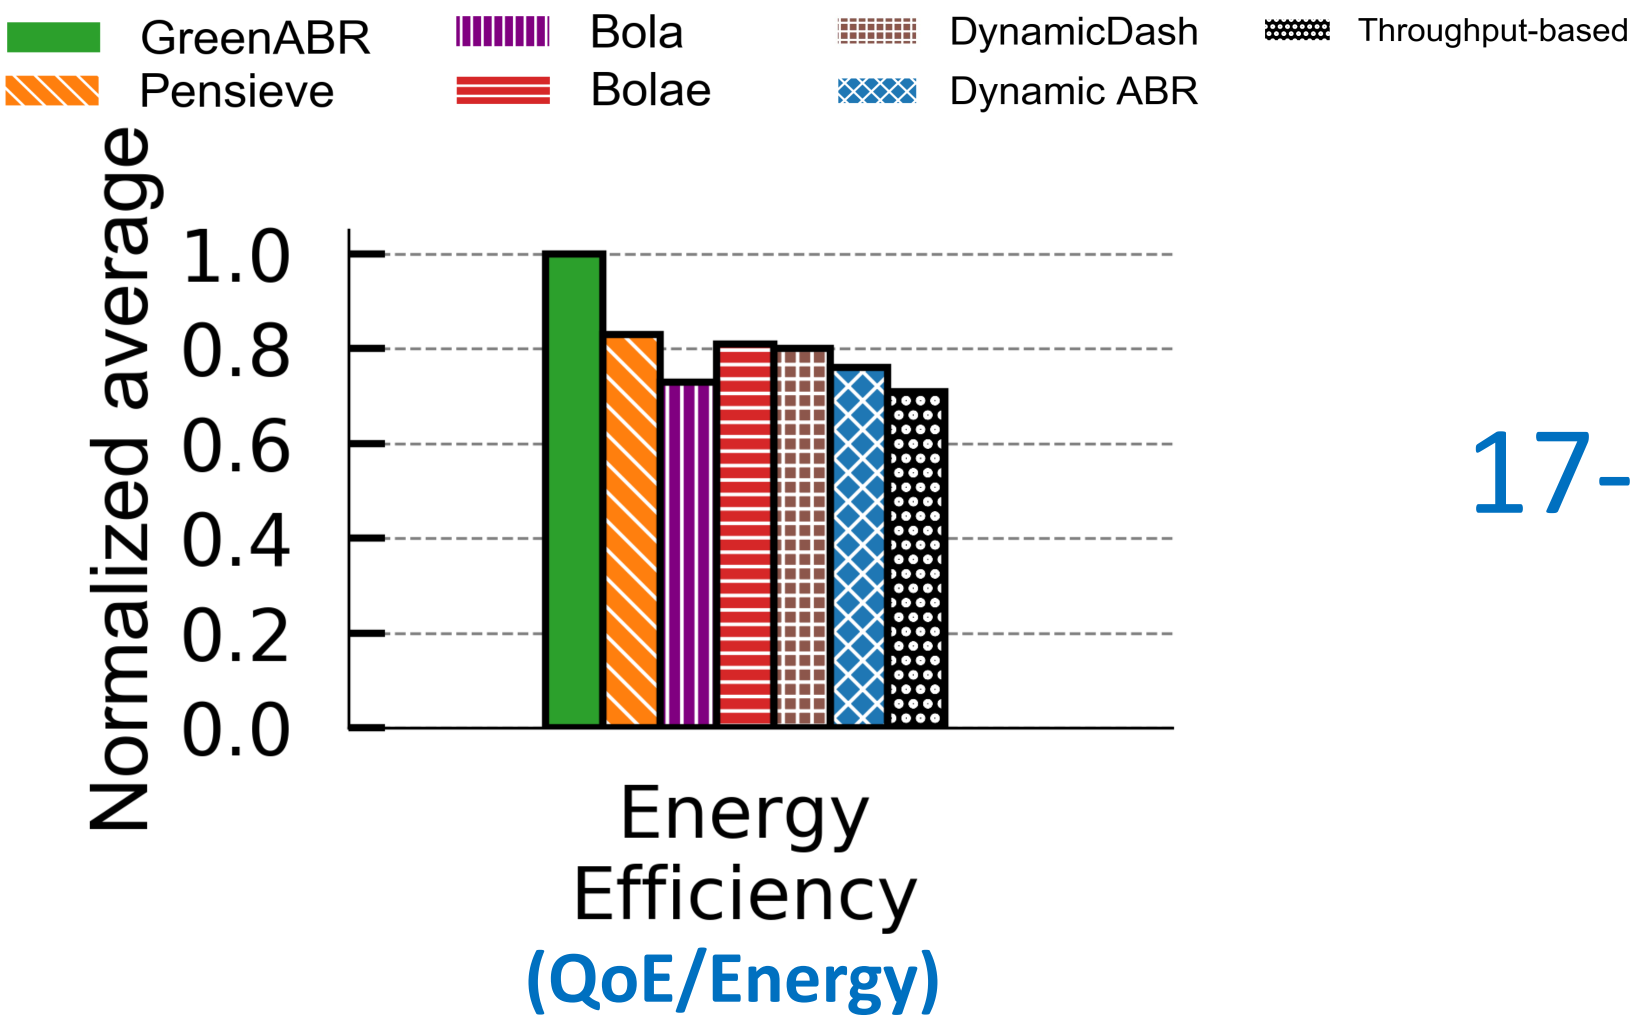

17-29% better

Results - 6 representations

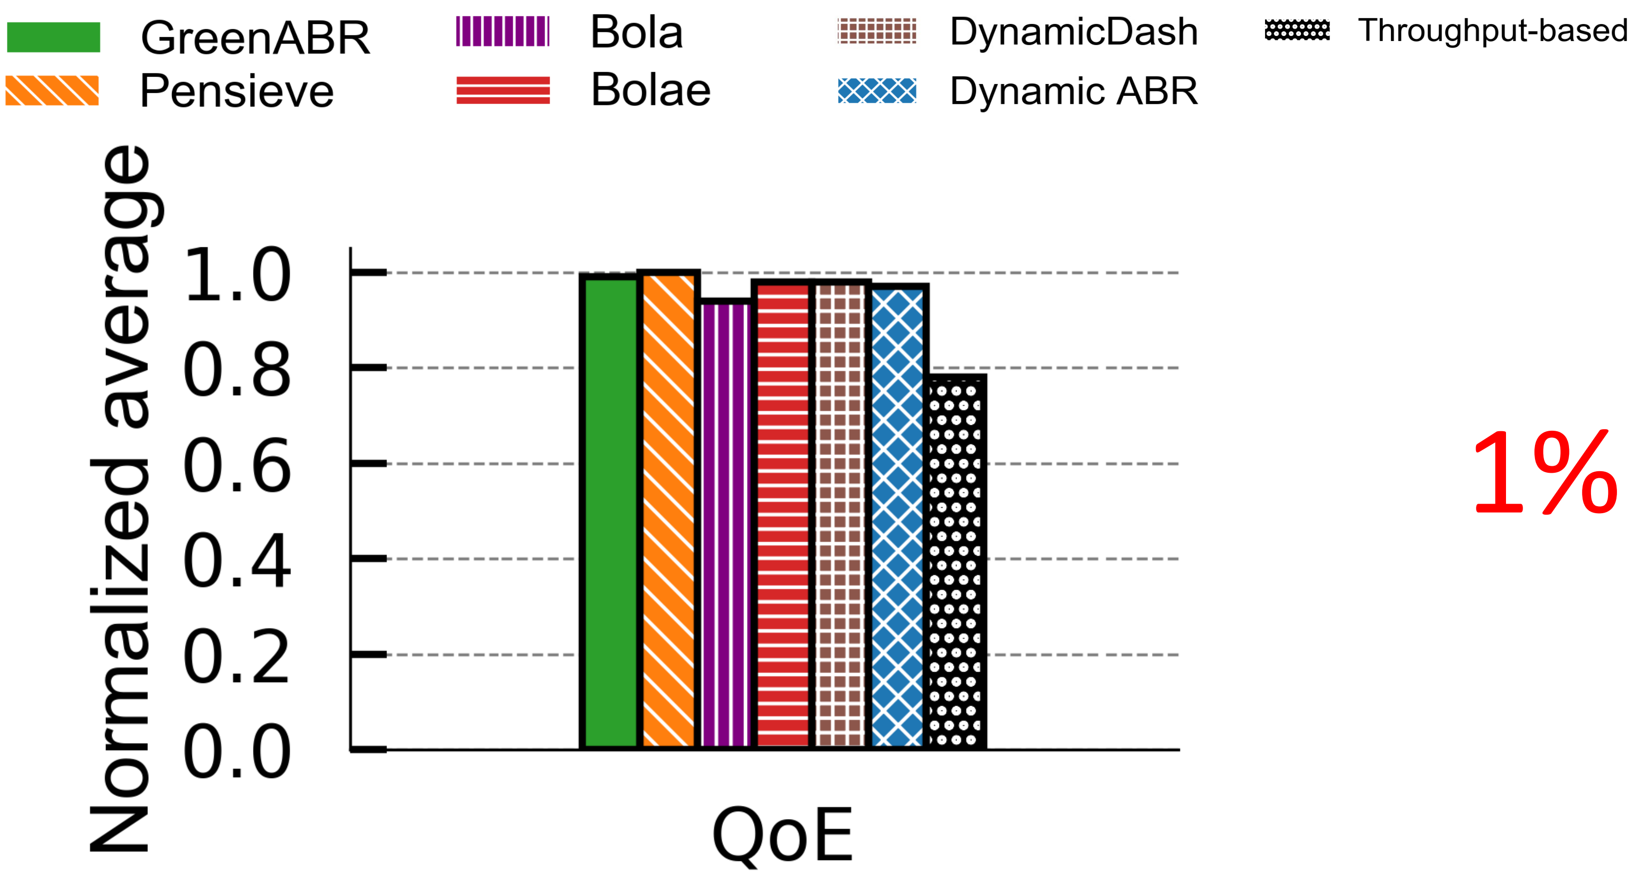

1% degradation

Results - 6 representations

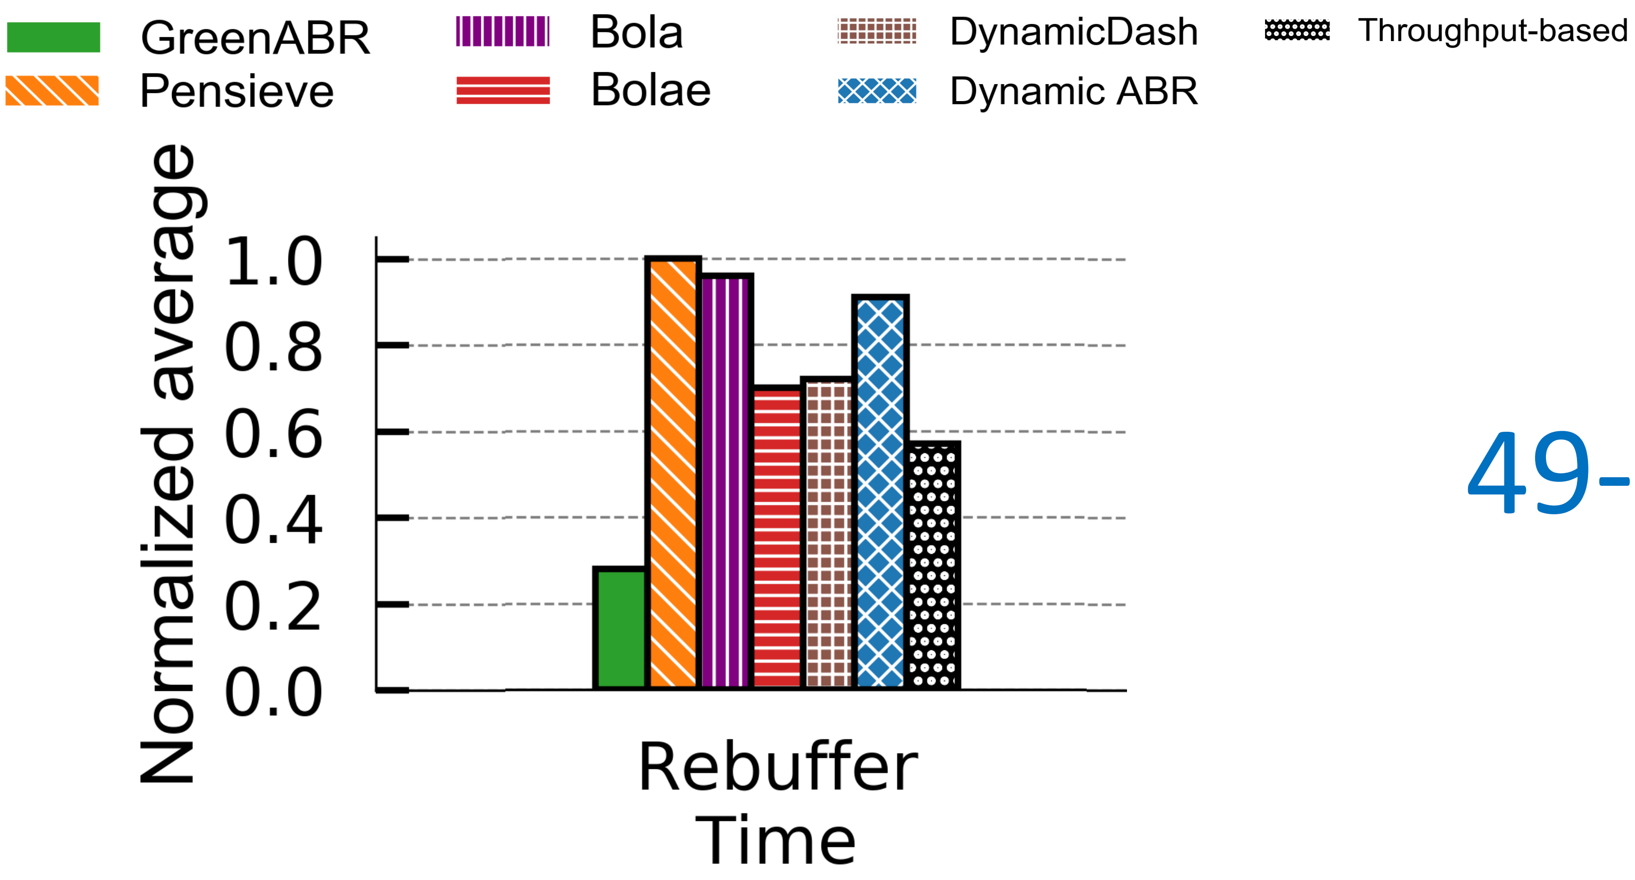

49-72% lower

Results - 6 representations

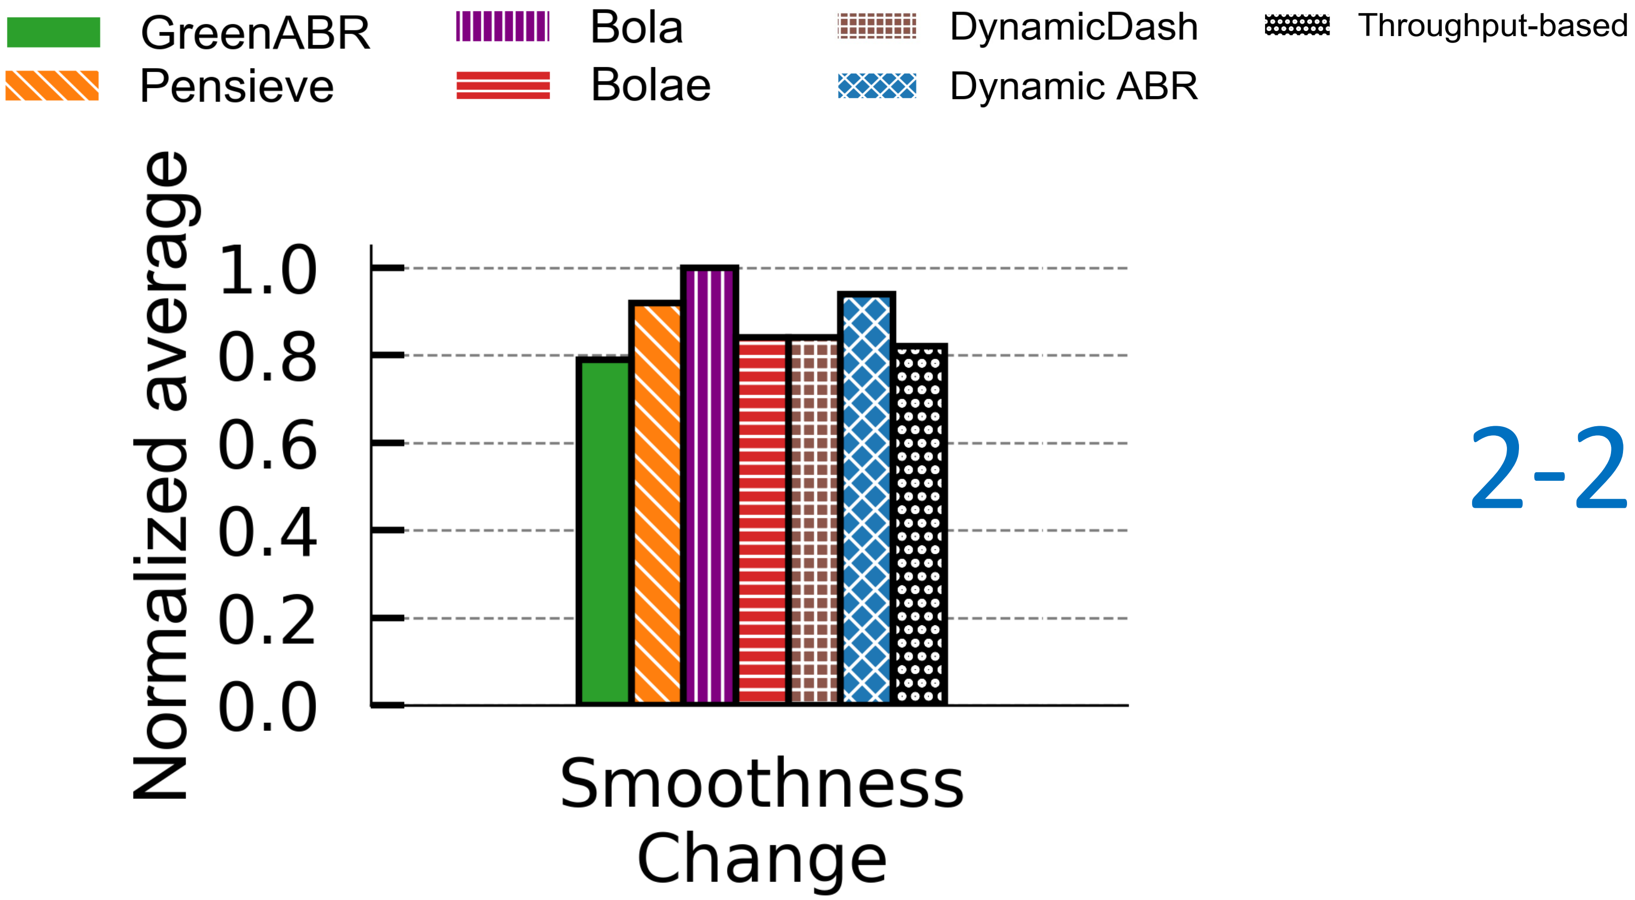

2-21% lower

Results - 6 representations

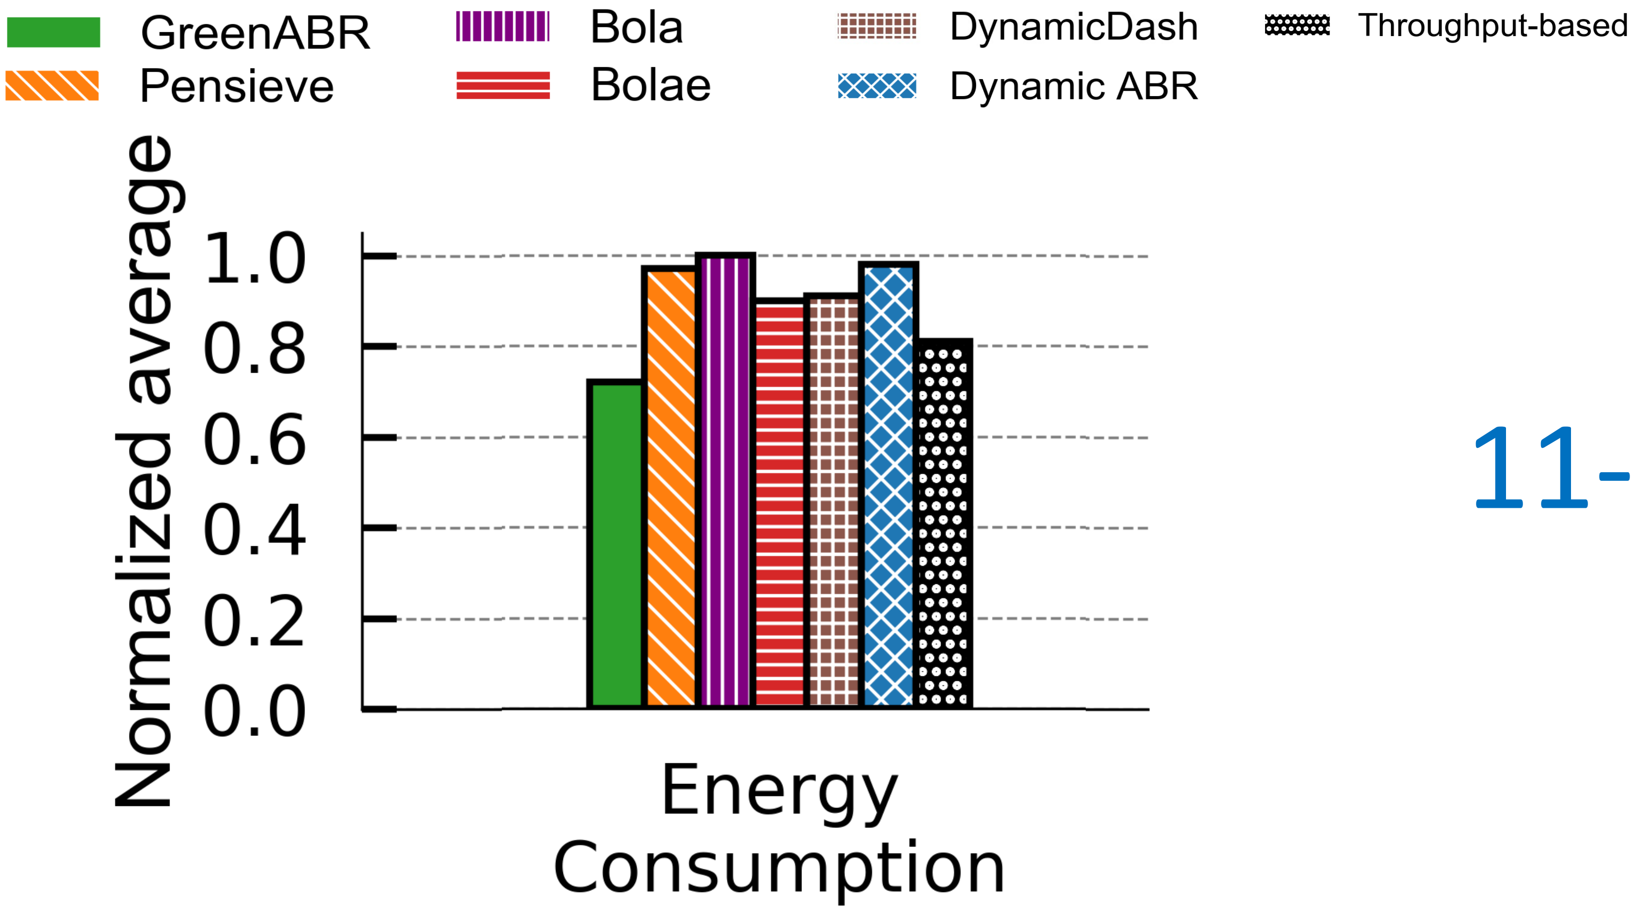

11-28% savings

Results - 6 representations

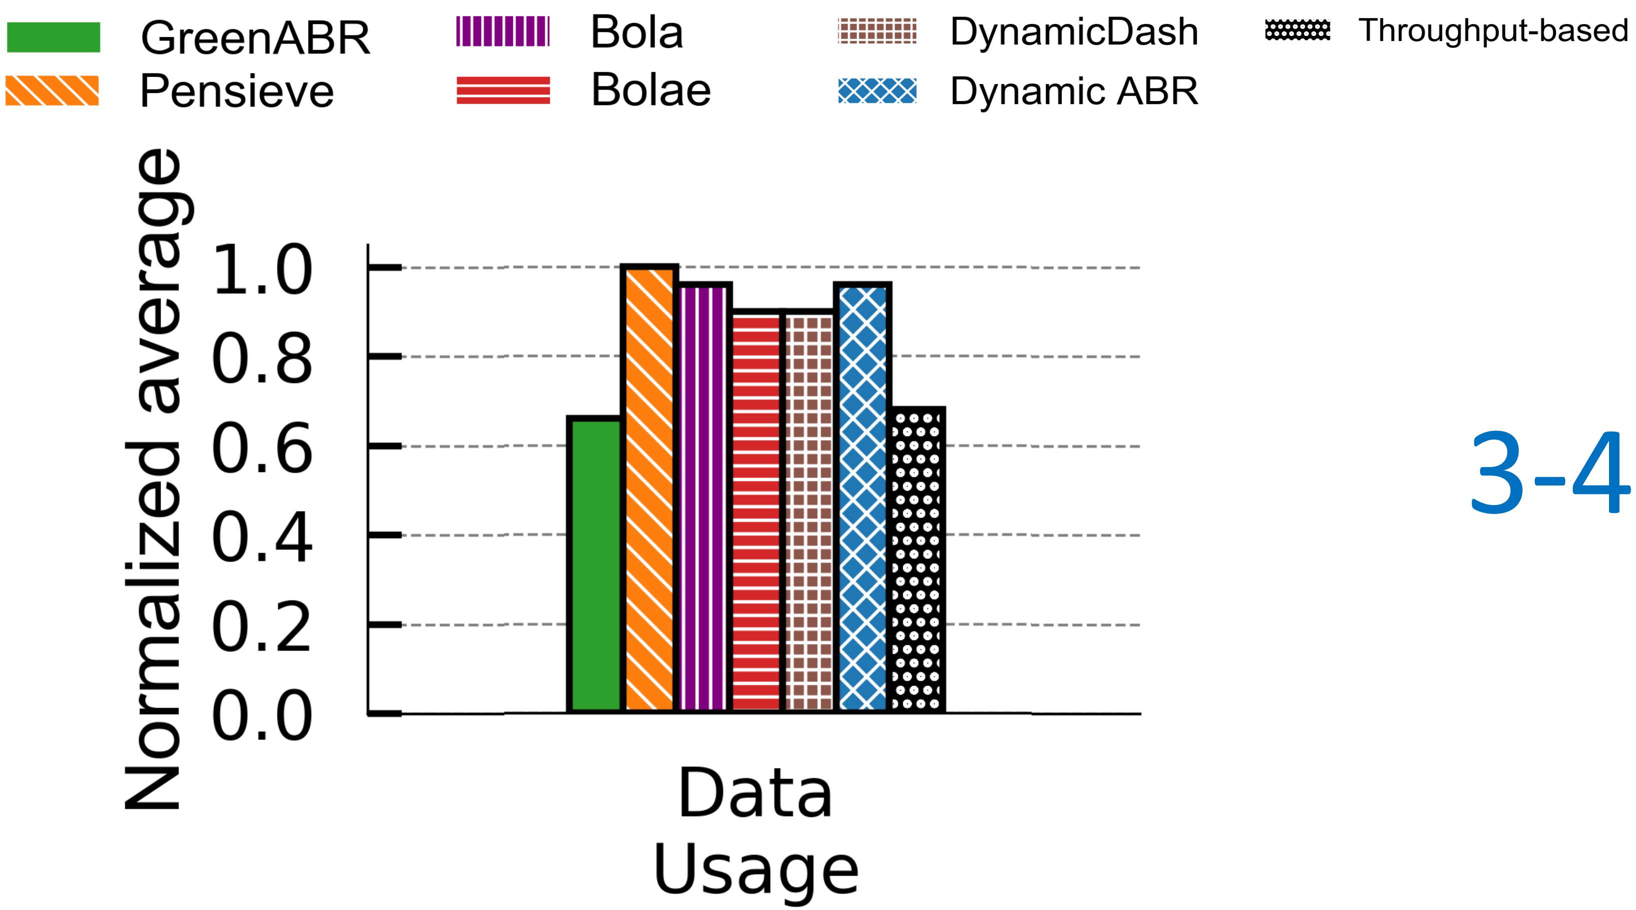

3-44% savings

Results - 10 representations

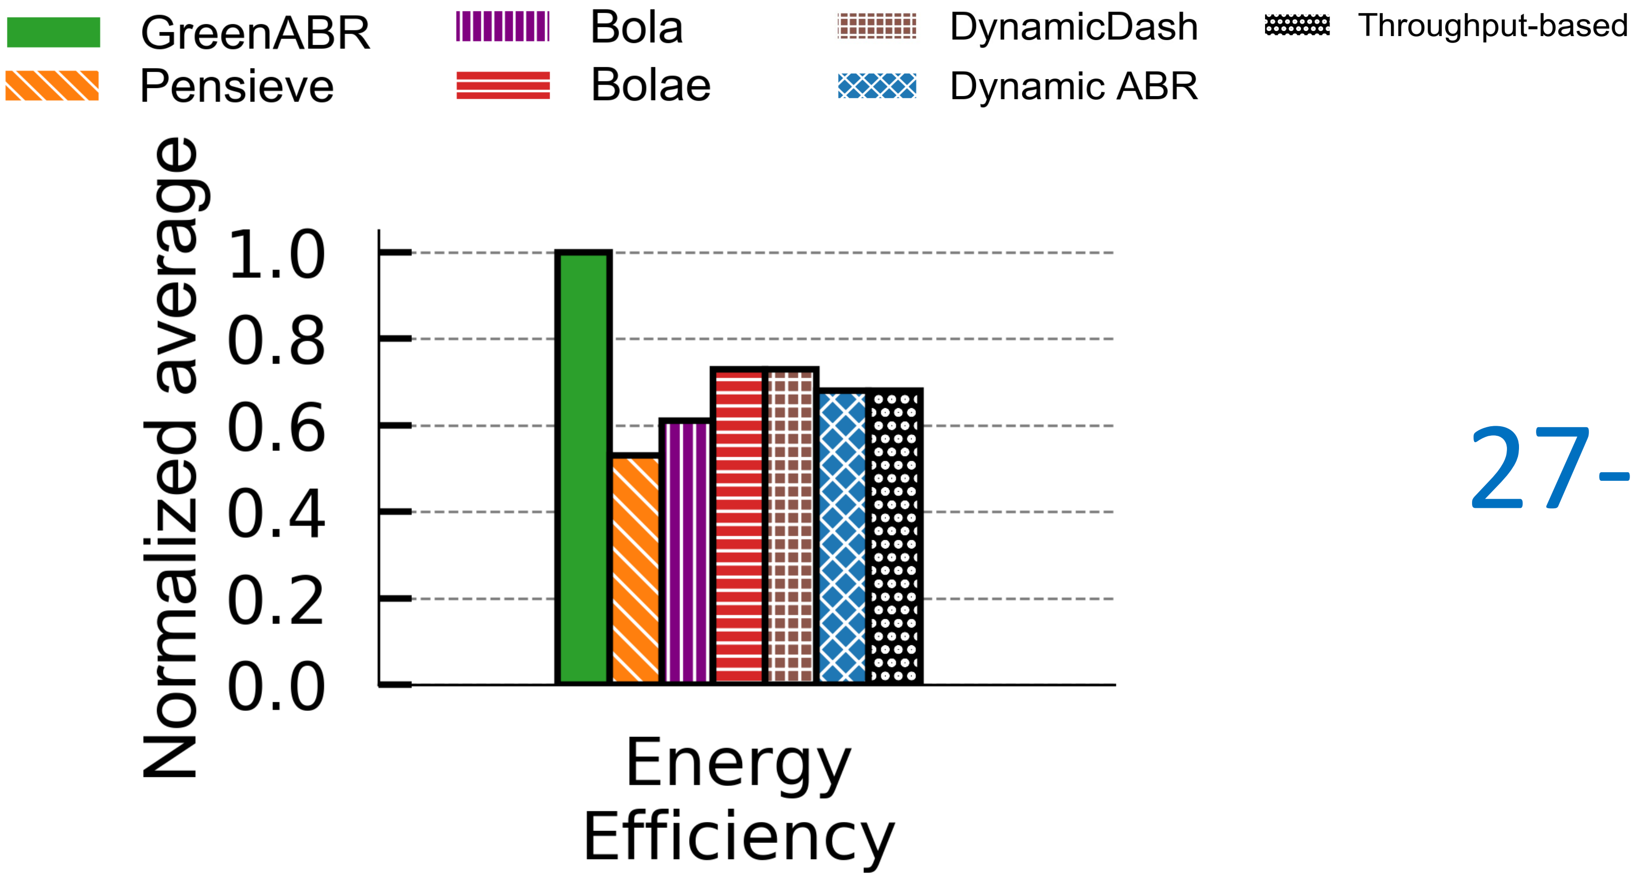

27-47% better

Results - 10 representations

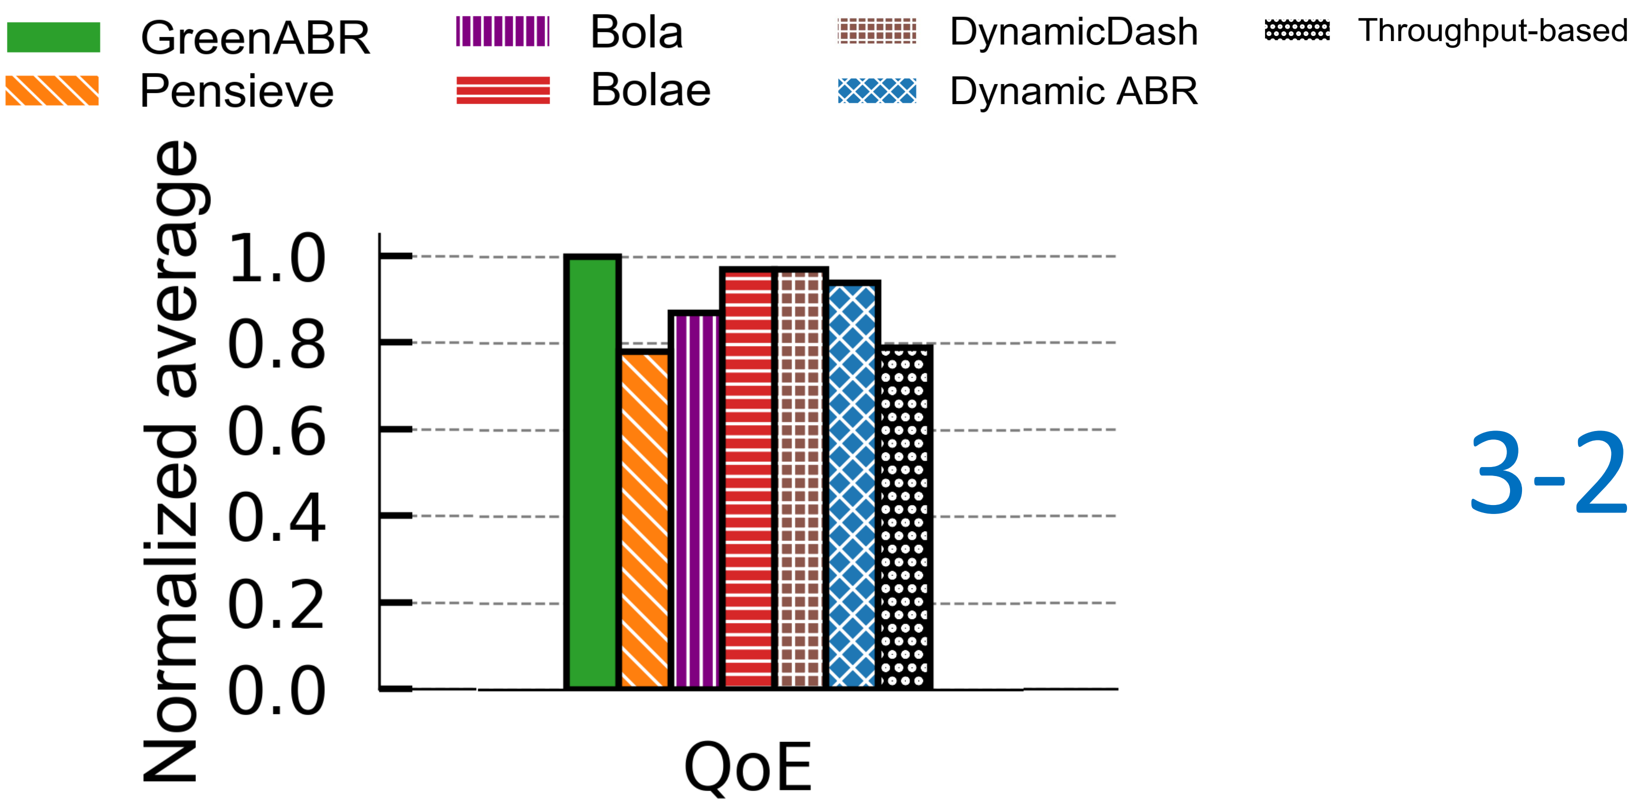

3-22% better

Results - 10 representations

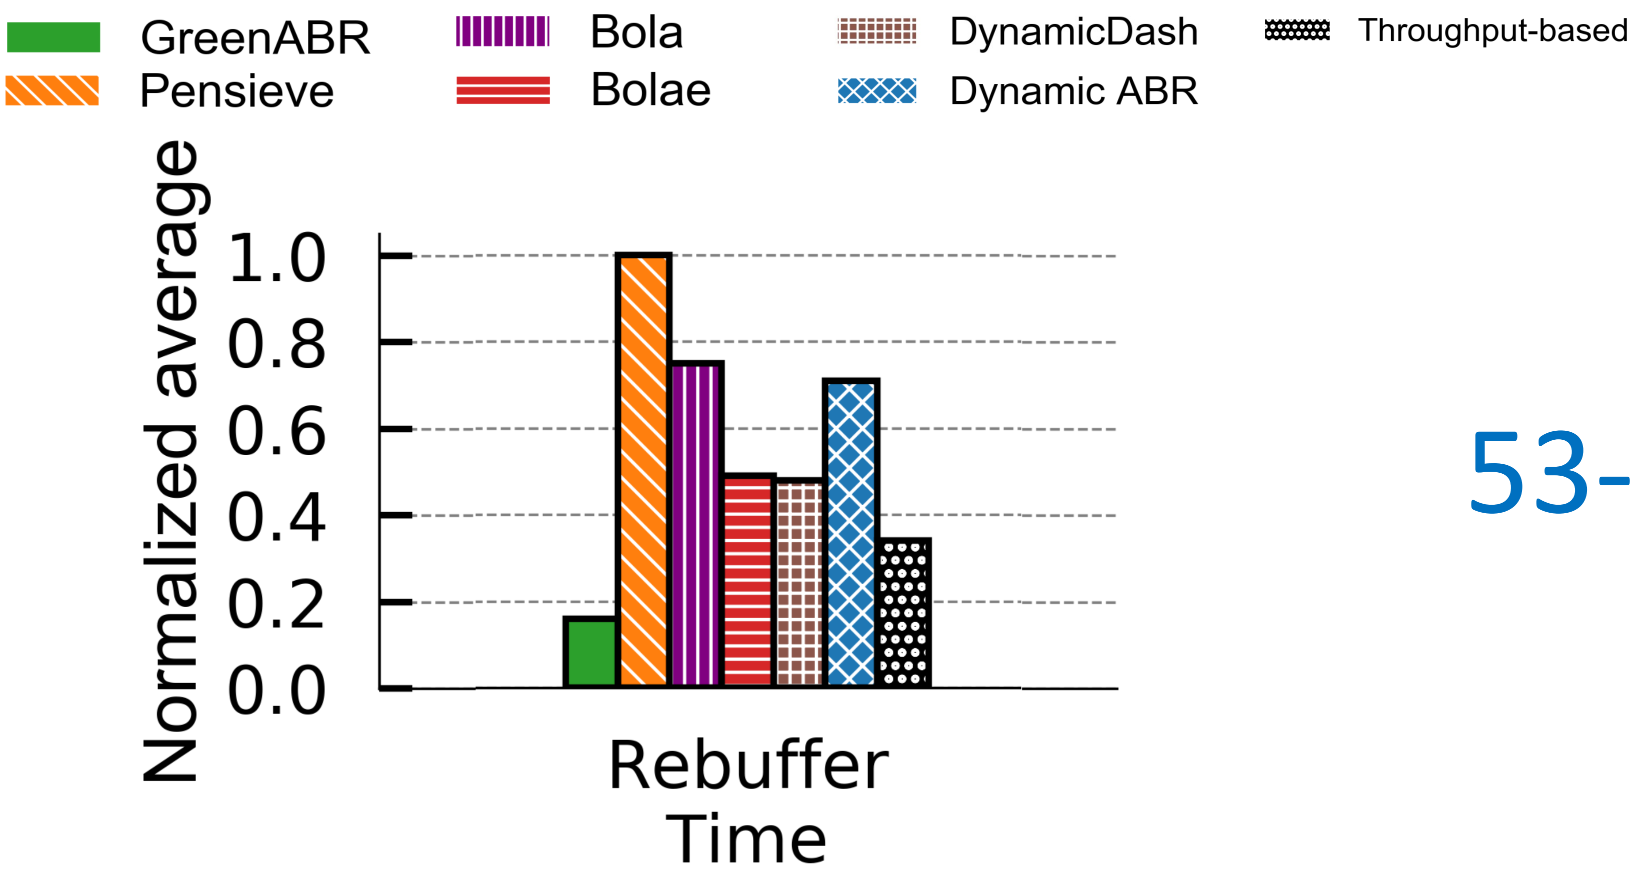

53-84% lower

Results - 10 representations

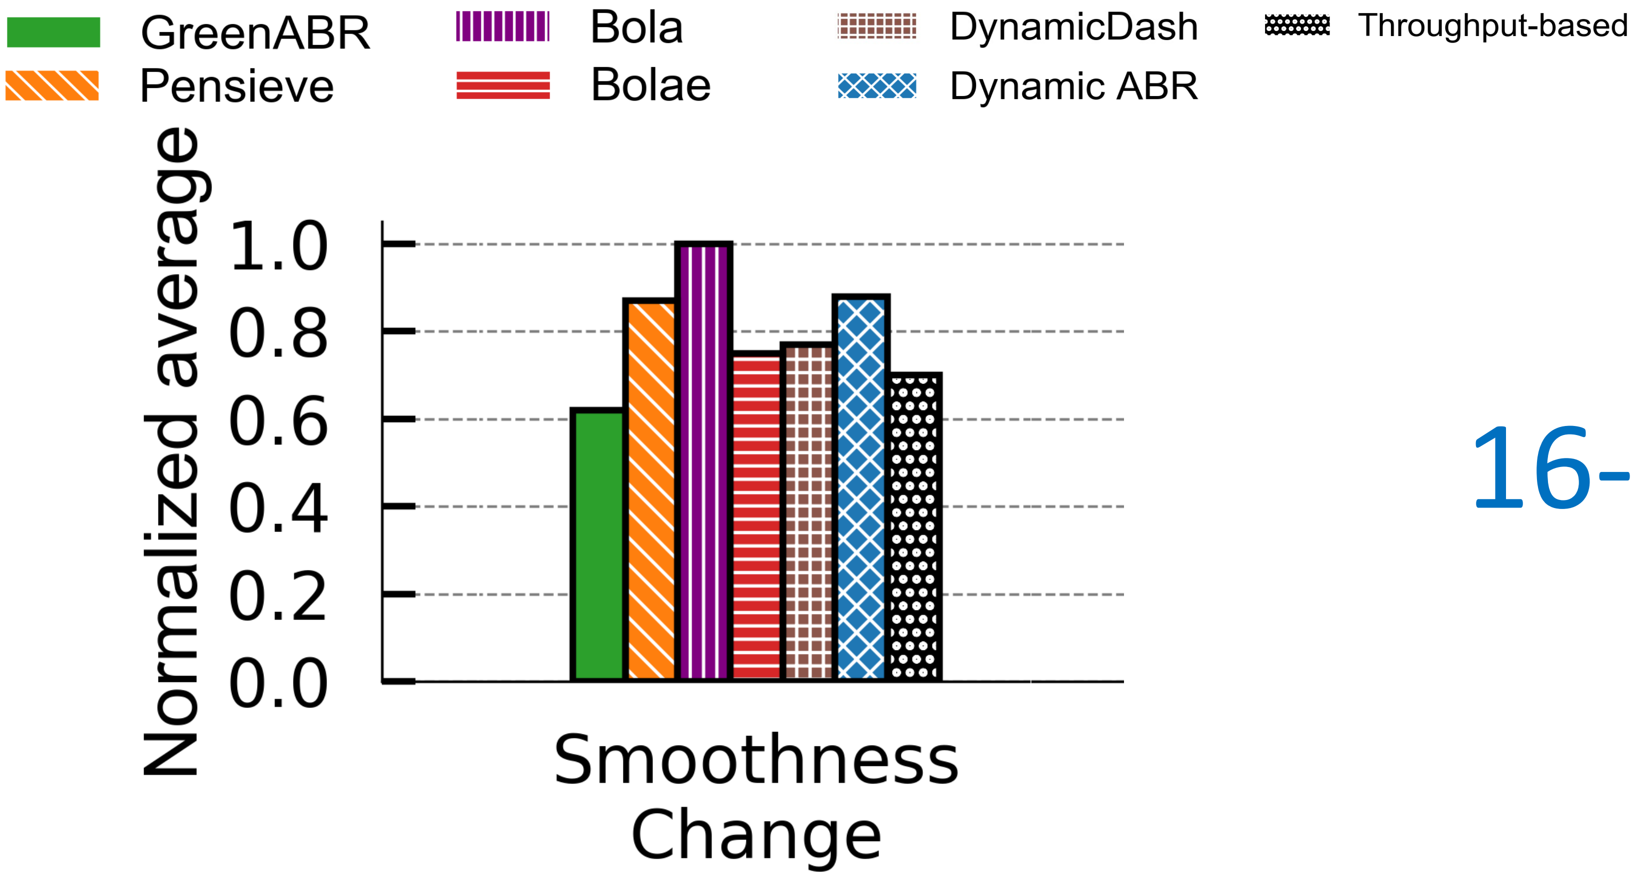

16-31% lower

Results - 10 representations

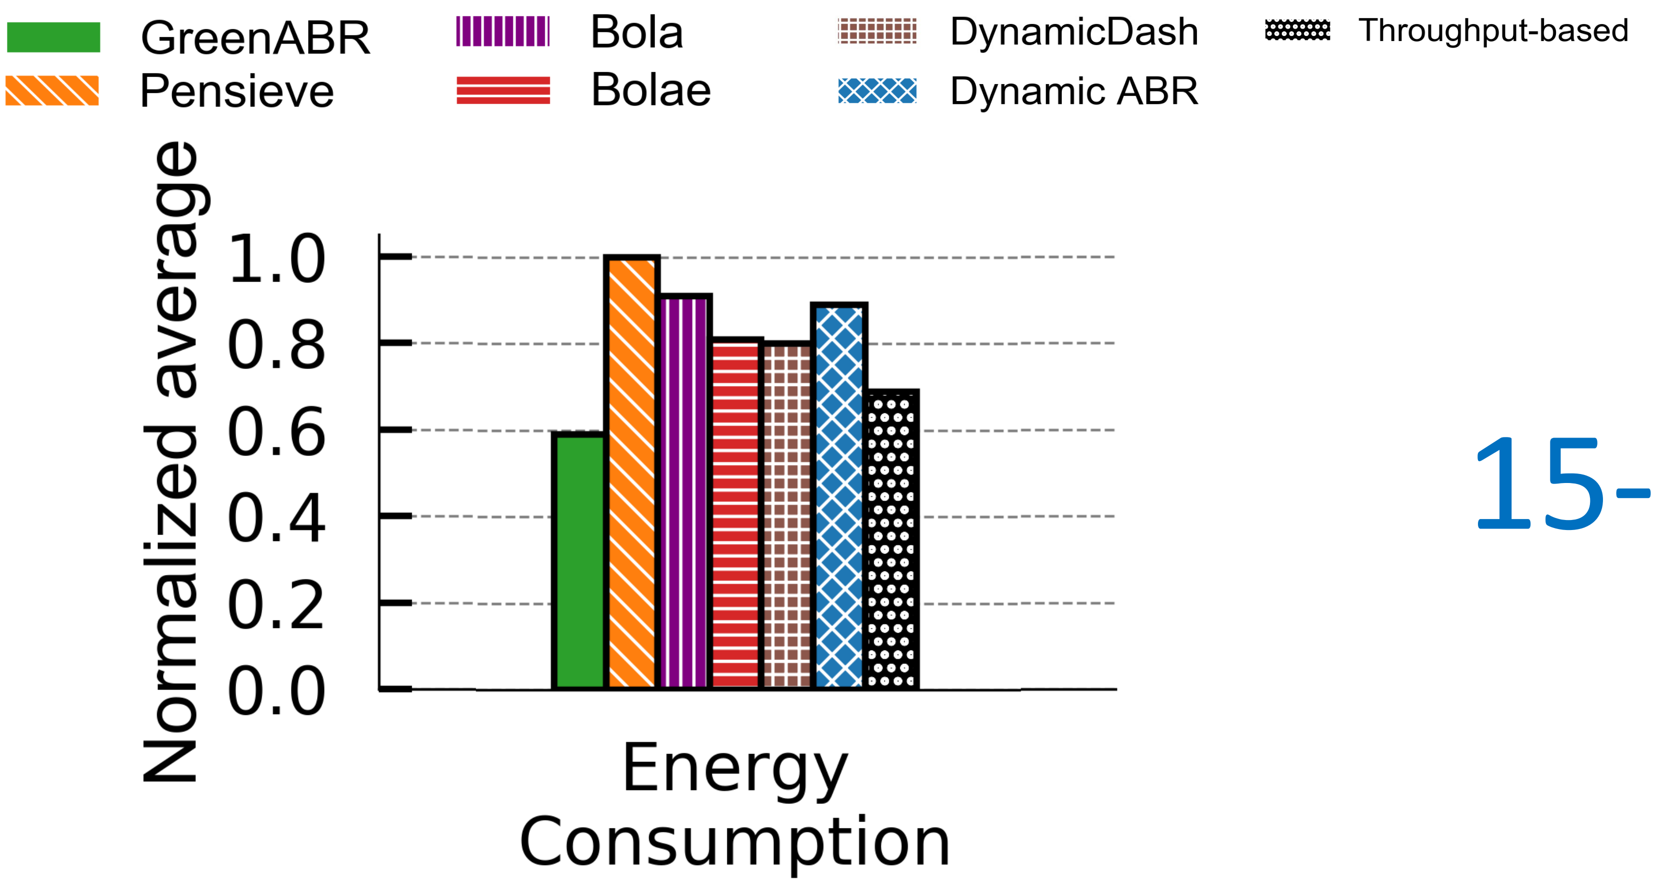

15-41% savings

Results - 10 representations

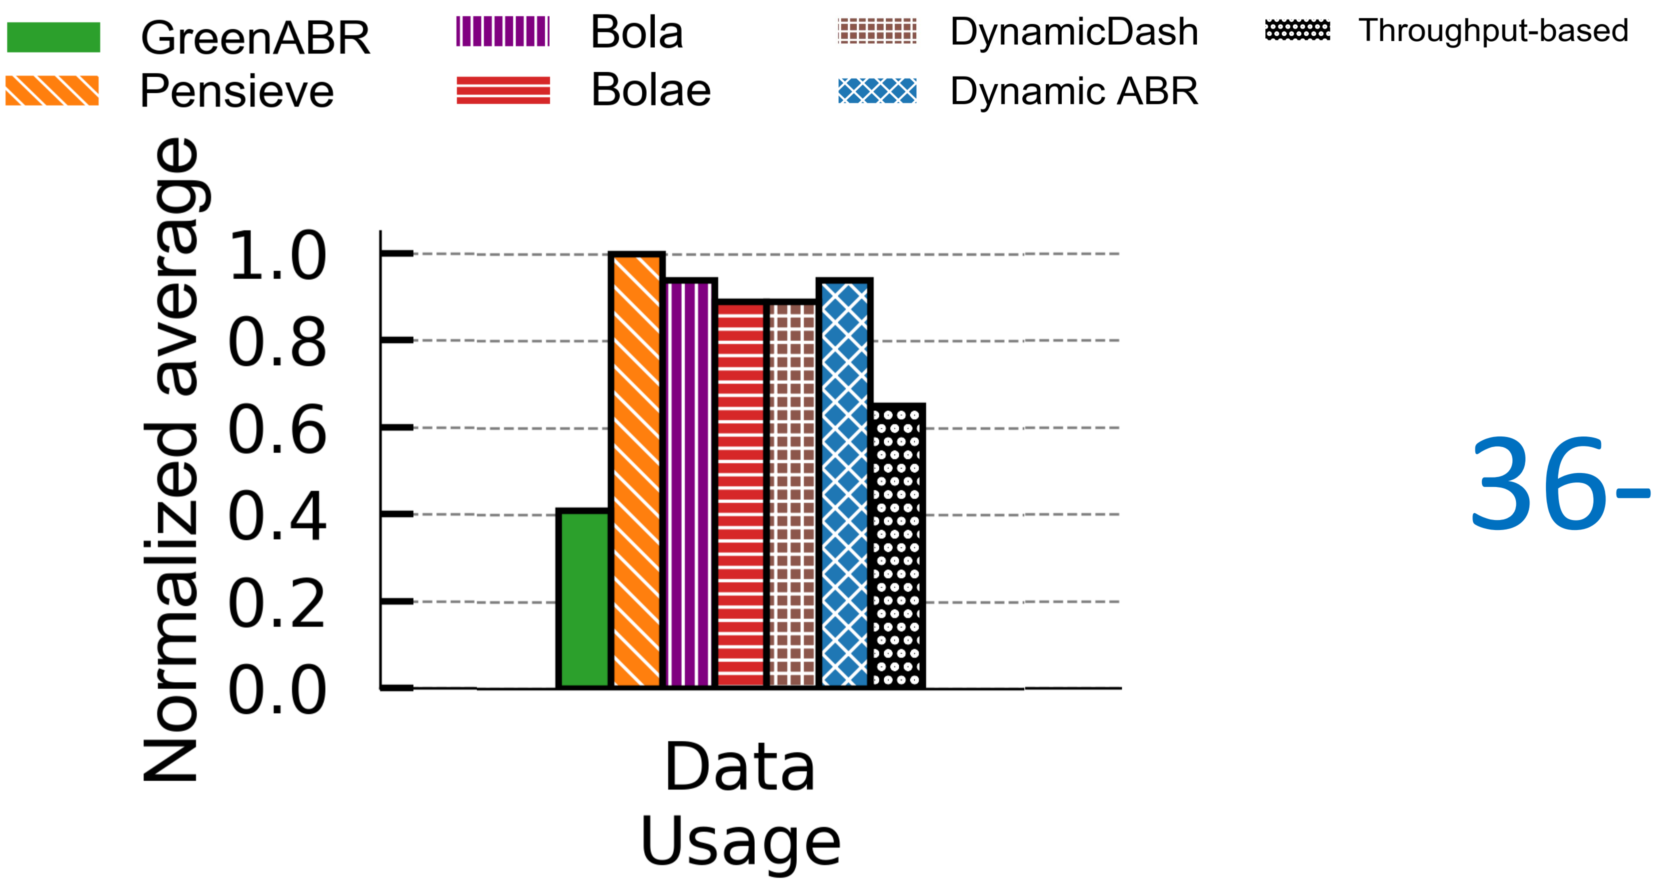

36-59% savings

Results – 10 representations (2K-4K)

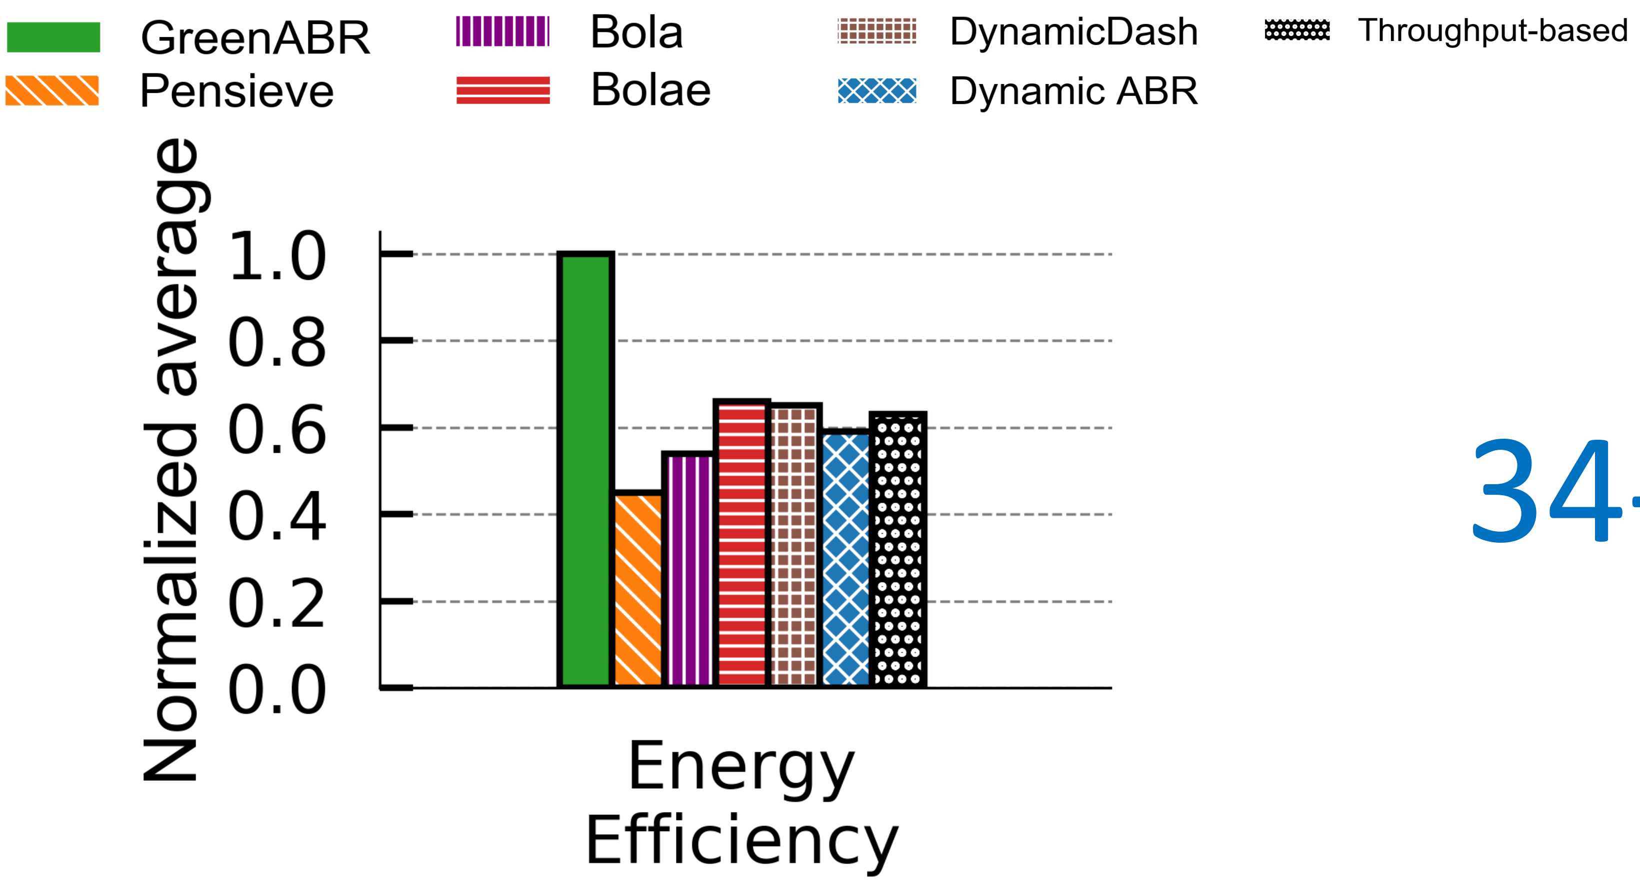

34-55% better

Results - 10 representations (2K-4K)

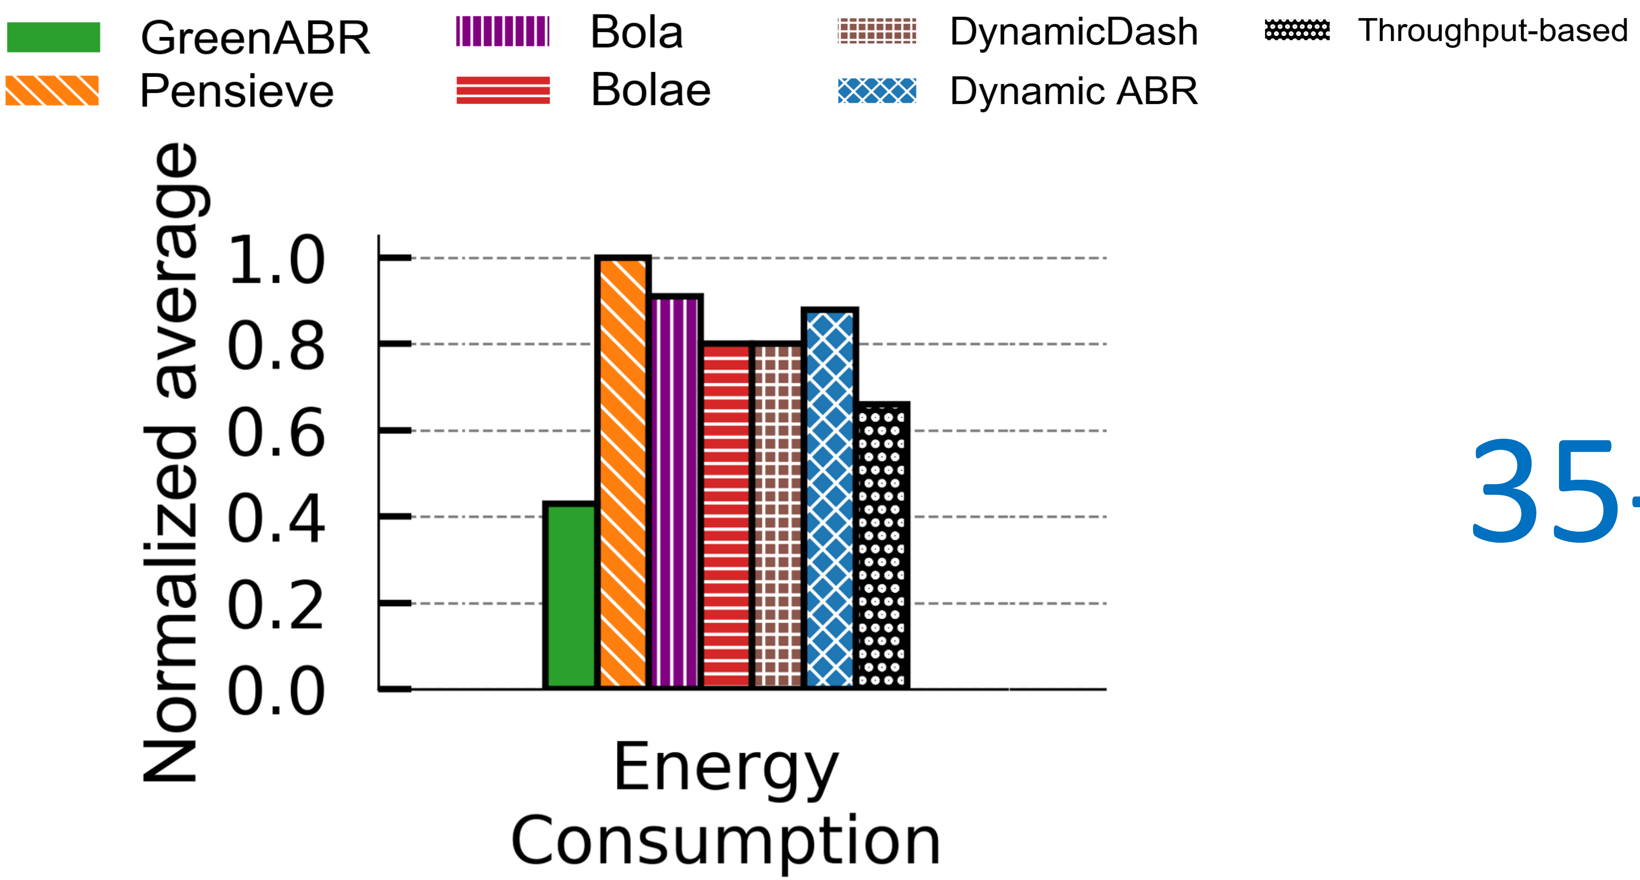

35-57% savings

Energy Consumption & Capacity Violations

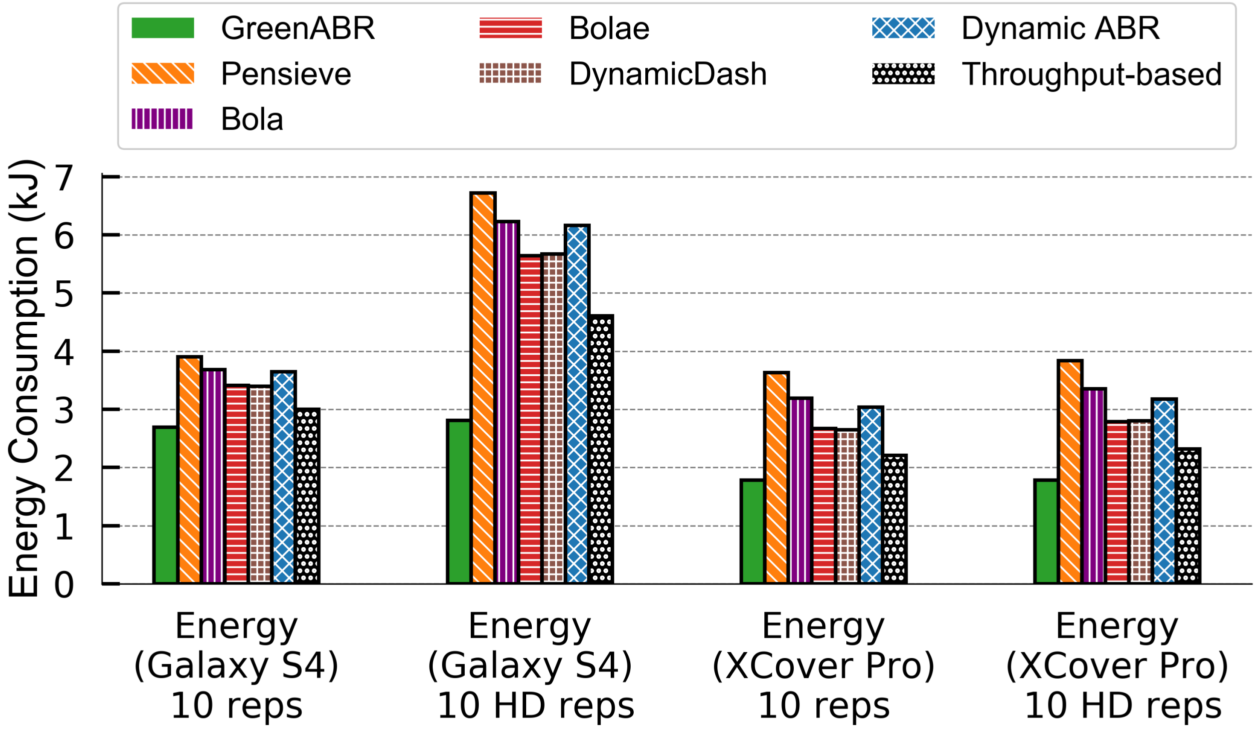

Energy Consumption Analysis

- Slight increase for Xcover Pro
- GreenABR preserves its efficiency
- Other ABRs **doubles** their energy consumption

Pixel Rate > Decoder Capacity

- Software decoder takes over
- Drops frames
- Boosts Energy Consumption

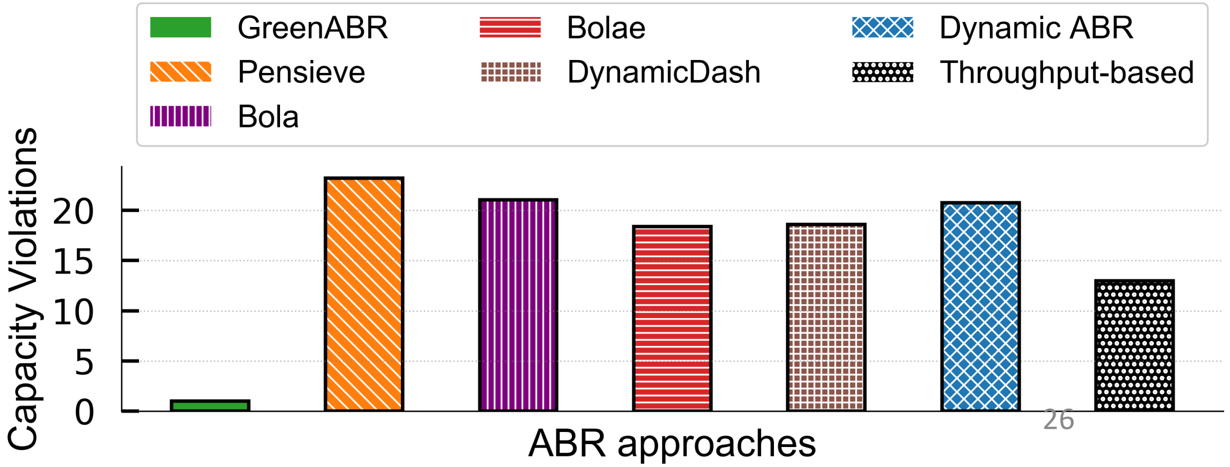

Summary

- RL models improve ABRs
  - Requires training efficiency and stability
- Perceptual quality metrics (e.g.,VMAF) improve ABRs
  - Bitrate greedy approaches
    - Fail for unstable network conditions
    - Incur additional power and data usage
- Sustainable solutions needed for mobile users
- Standardized evaluations needed to avoid unfair comparisons
- GreenABR outperforms SOTA up to
  - 57% energy savings
  - 22% QoE
  - 60% data savings
  - 84% less rebuffering

Future Work

- Generalized GreenABR for different encoding ladders
- GreenABR for live streaming scenarios
- GreenABR on server side

Questions

References

1. K. Spiteri, R. Urgaonkar, and R. K. Sitaraman. 2016. BOLA: Near-optimal bitrate adaptation for online videos. In IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications. 1–9
2. Junchen Jiang, Vyas Sekar, and Hui Zhang. 2012. Improving Fairness, Efficiency, and Stability in HTTP-Based Adaptive Video Streaming with FESTIVE (CoNEXT ’12). New York, NY, USA, 97–108.
3. Yasir Zaki, Thomas Pötsch, Jay Chen, Lakshminarayanan Subramanian, and Carmelita Görg. 2015. Adaptive Congestion Control for Unpredictable Cellular Networks. In Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (London, United Kingdom) (SIGCOMM ’15).
4. Kevin Spiteri, Ramesh Sitaraman, and Daniel Sparacio. 2018. From Theory to Prac- tice: Improving Bitrate Adaptation in the DASH Reference Player. In Proceedings of the 9th ACM Multimedia Systems Conference (Amsterdam, Netherlands) (MMSys ’18).
5. Hongzi Mao, Ravi Netravali, and Mohammad Alizadeh. 2017. Neural Adaptive Video Streaming with Pensieve. In Proceedings of the Conference of the ACM Special Interest Group on Data Communication (Los Angeles, CA, USA) (SIGCOMM ’17).
6. Zahaib Akhtar, Yun Seong Nam, Ramesh Govindan, Sanjay Rao, Jessica Chen, Ethan Katz-Bassett, Bruno Ribeiro, Jibin Zhan, and Hui Zhang. 2018. Oboe: Auto- Tuning Video ABR Algorithms to Network Conditions. In Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (Budapest, Hungary) (SIGCOMM ’18).
7. FrancisY.Yan,HudsonAyers,ChenzhiZhu,SadjadFouladi,JamesHong,Keyi Zhang, Philip Levis, and Keith Winstein. 2020. Learning in situ: a randomized experiment in video streaming. In 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20).
8. Tianchi Huang, Chao Zhou, Rui-Xiao Zhang, Chenglei Wu, Xin Yao, and Lifeng Sun. 2019. Comyco: Quality-Aware Adaptive Video Streaming via Imitation Learning. CoRR abs/1908.02270 (2019).
9. B. Varghese, G. Jourjon, K. Thilakarathne, and A. Seneviratne. 2017. e-DASH: Modelling an energy-aware DASH player. In 2017 IEEE 18th International Sym- posium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM).

References

10. S. Kim, H. Oh, and C. Kim. 2018. eff-HAS: Achieve higher efficiency in data and energy usage on dynamic adaptive streaming. Journal of Communications and Networks 20, 3 (2018), 325–342.
11. M. Uitto and M. Forsell. 2018. Towards Energy-Efficient Adaptive Mpeg-Dash Streaming Using Hevc. In 2018 IEEE International Conference on Multimedia Expo Workshops (ICMEW). 1–6.
12. X. Chen, T. Tan, and G. Cao. 2019. Energy-Aware and Context-Aware Video Streaming on Smartphones. In 2019 IEEE 39th International Conference on Dis- tributed Computing Systems (ICDCS). 861–870.
13. A. Mondal, B. Palit, S. Khandelia, N. Pal, J. Jayatheerthan, K. Paul, N. Ganguly, and S. Chakraborty. 2020. EnDASH - A Mobility Adapted Energy Efficient ABR Video Streaming for Cellular Networks. In 2020 IFIP Networking Conference (Networking). 127–135.
14. T. Breitbach, P. Sanders, and D. Schultes. 2018. Optimizing energy consumption and user experience in a mobile video streaming scenario. In 2018 15th IEEE Annual Consumer Communications Networking Conference (CCNC). 1–9.
15. Jiayi Meng, Qiang Xu, and Y. Charlie Hu. 2021. Proactive Energy-Aware Adaptive Video Streaming on Mobile Devices. In 2021 USENIX Annual Technical Conference (USENIX ATC 21).
16. Yanyuan Qin, Shuai Hao, Krishna R. Pattipati, Feng Qian, Subhabrata Sen, Bing Wang, and Chaoqun Yue. 2019. Quality-Aware Strategies for Optimiz- ing ABR Video Streaming QoE and Reducing Data Usage. In Proceedings of the 10th ACM Multimedia Systems Conference (Amherst, Massachusetts) (MM- Sys ’19).
17. Yanyuan Qin, Chinmaey Shende, Cheonjin Park, Subhabrata Sen, and Bing Wang. 2021. DataPlanner: Data-Budget Driven Approach to Resource-Efficient ABR Streaming.
18. Li Sun, Ramanujan K. Sheshadri, Wei Zheng, and Dimitrios Koutsonikolas. 2014. Modeling WiFi Active Power/Energy Consumption in Smartphones. In 2014 IEEE 34th International Conference on Distributed Computing Systems.
19. Zhengfang Duanmu, Abdul Rehman, and Zhou Wang. 2018. A Quality-of- Experience Database for Adaptive Video Streaming. IEEE Transactions on Broad- casting 64, 2 (2018), 474–487.

Supplementary Slides

Modelling Power Consumption

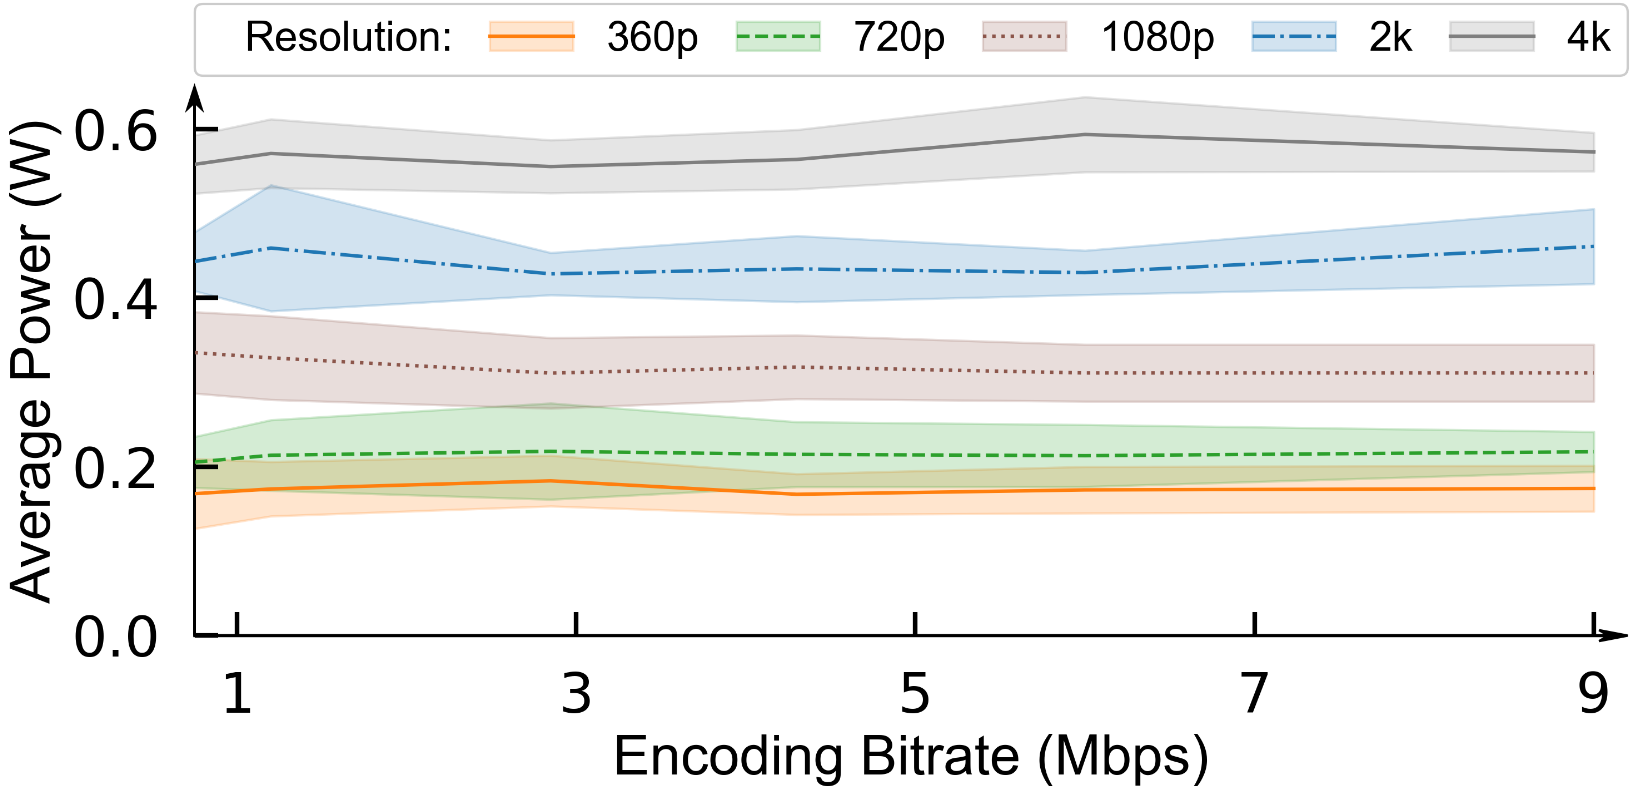

Limitations of existing models

- Designed based on encoding bitrate
- Device heterogeneity

Our approach

- Considers ABR-related components
- Targets to capture pattern
- Minimizes device specific impacts
  - Excludes base power consumption

Evaluation Setup

Videos

- Three-minute videos
- Tears of steel (sci-fi), Nature(documentary), BBB(cartoon)

Energy Data Collection

- Videos are played for each representation
- Monsoon Power meter collects at each 200 microseconds
- Results are aggregated for each chunk

Trace Type	Avg. Band. (Mbps)	Distribution %
3G	1.29 ± 0.77	29%
4G	31.58 ± 13.59	14%
Broadband	3.91 ± 2.37	28%
Synthetic	4.41 ± 5.03	29%

Model	SRCC
SQoE-III sample	0.6534
Pensieve	0.6563
Comyco	0.7517
Proposed Model	**0.7845**

Perceptual Quality Calculation

- QoE of the streaming session

video quality, rebuffering time and frequency, oscillations, quality switches

- Linear regression model on SQoE-III[19] dataset
- VMAF is used as quality measure
- 34 users 450 videos 13 net. conditions

Results - 6 representations

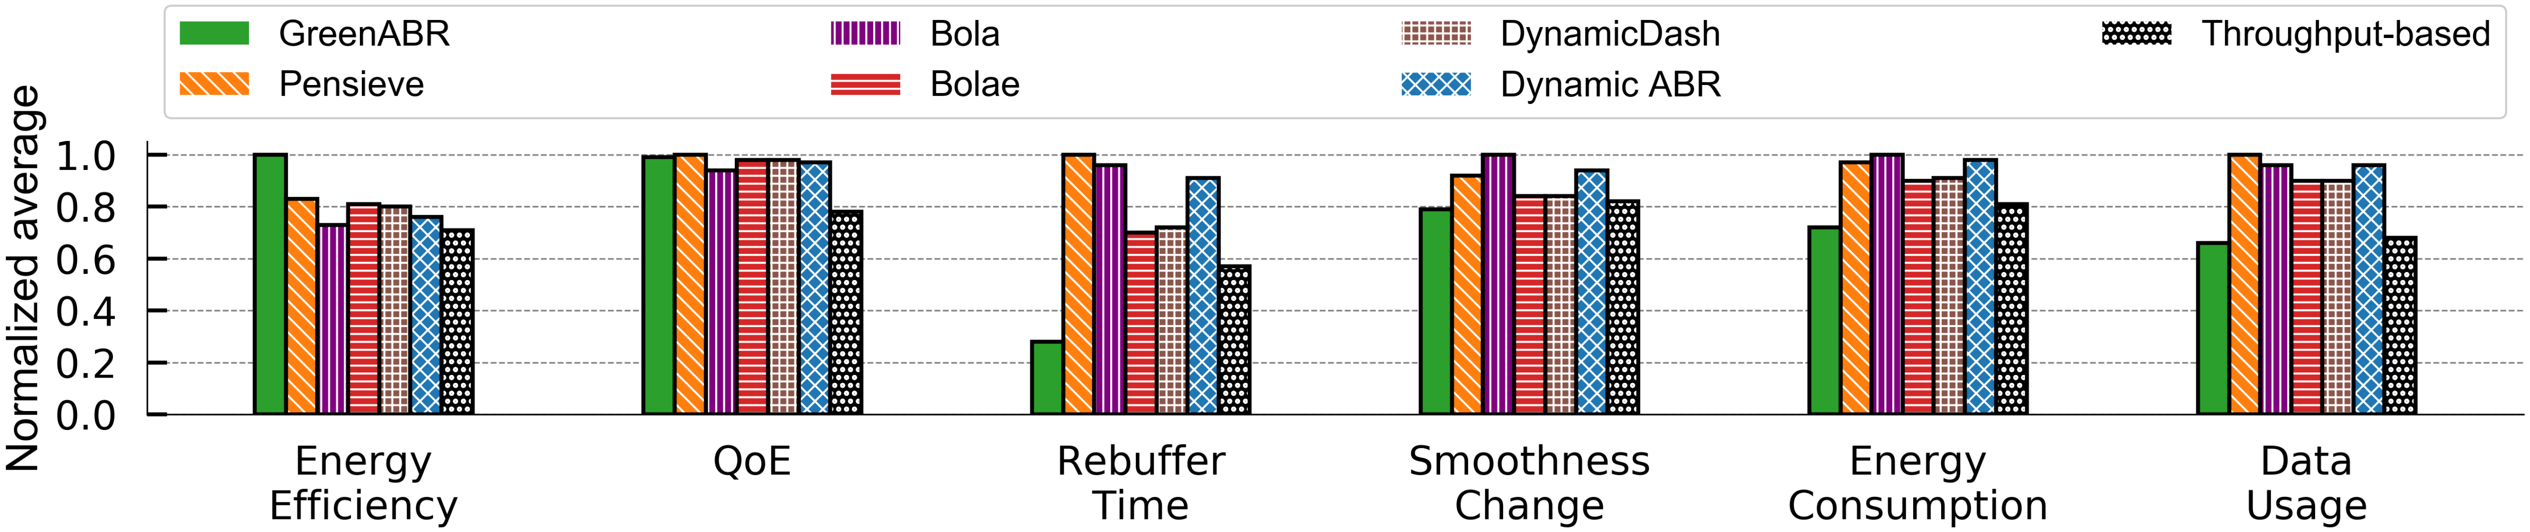

Resolution	Bitrate(kbps)
320 x 180	300
640 x 360	750
768 x 432	1200
1024 x 576	1850
1280 x 720	2850
1920 x 1080	4300

- Average results of 3 videos for all traces
- Energy Efficiency (QoE/Energy) => 17-29 % better
- QoE => 1% degradation
- Rebuffering => 49-72 % lower
- Smoothness => 2-21 % lower
- Energy => 11-28 % savings
- Data => 3-44 % savings

Results - 10 representations

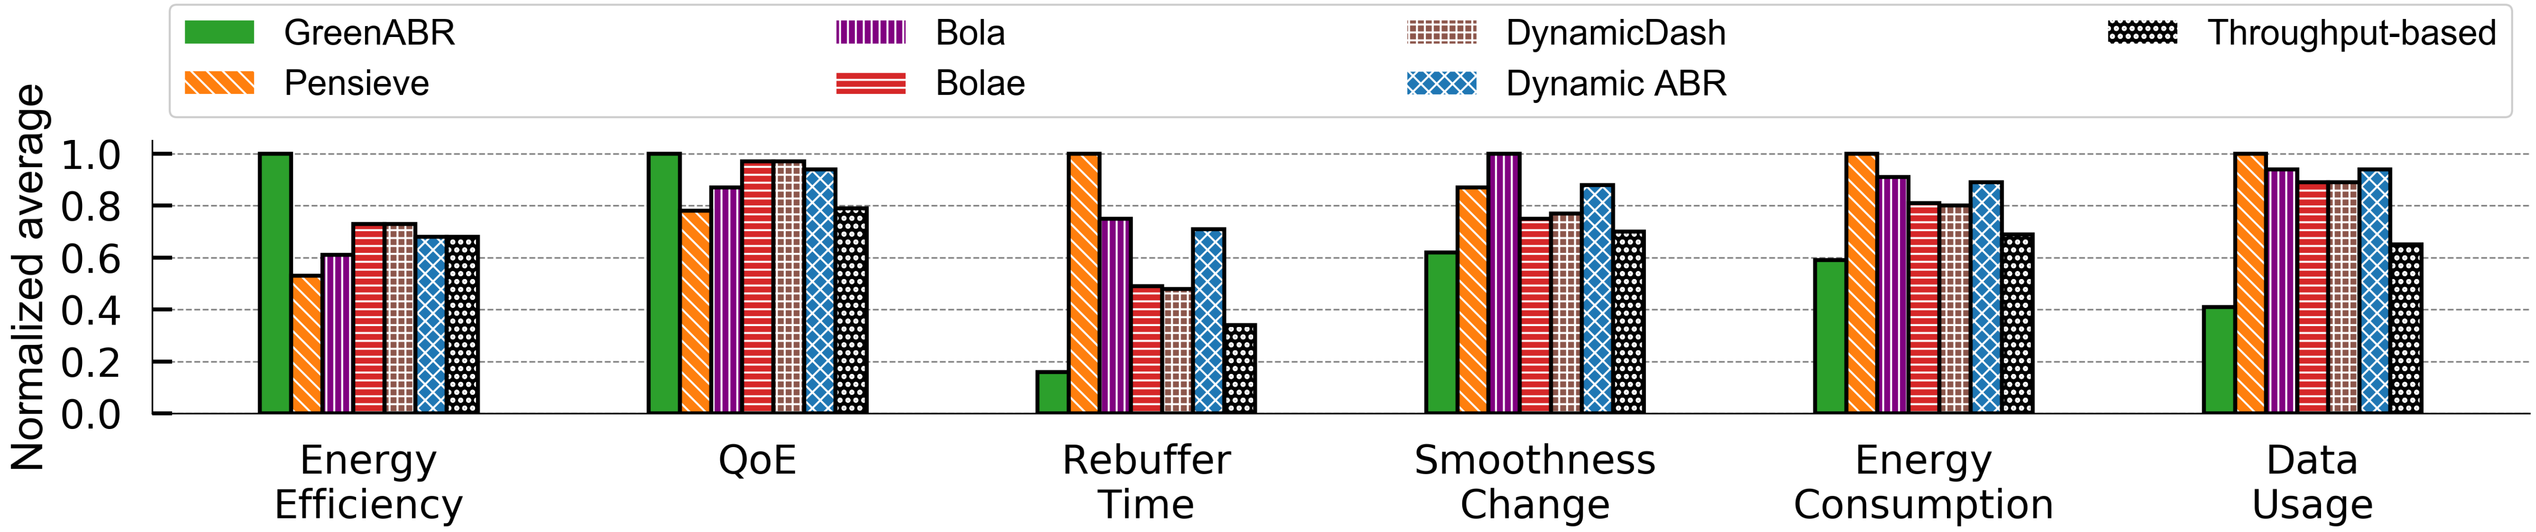

Resolution	Bitrate(kbps)
320 x 180	**150**
320 x 180	300
640 x 360	750
768 x 432	1200
1024 x 576	1850
1280 x 720	2850
1920 x 1080	4300
1920 x 1080	**6000**
1920 x 1080	**9000**
1920 x 1080	**12000**

- Average results of 3 videos for all traces
- Energy Efficiency (QoE/Energy) => 27-47% better
- QoE => 3-22% better
- Rebuffering => 53-84 % lower
- Smoothness => 16-31% lower
- Energy => 15-41% savings
- Data => 36-59% savings

Results - 10 representations HD

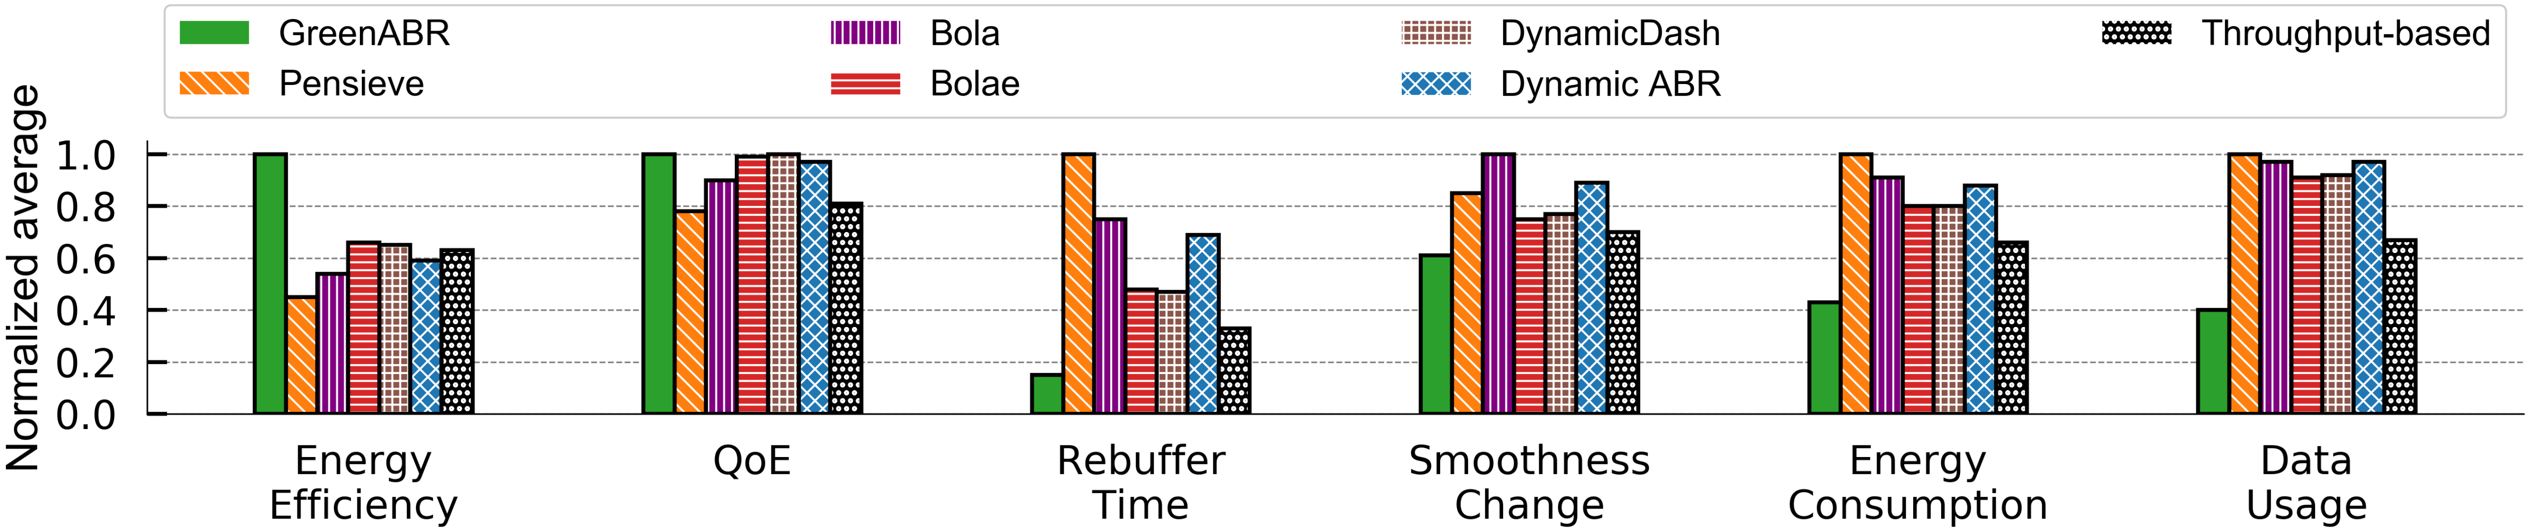

Resolution	Bitrate(kbps)
320 x 180	**150**
320 x 180	300
640 x 360	750
768 x 432	1200
1024 x 576	1850
1280 x 720	2850
1920 x 1080	4300
1920 x 1080	**6000**
2560 x 1440	**9000**
3840 x 2160	**12000**

- Average results of 3 videos for all traces
- Energy Efficiency (QoE/Energy) => 34-55% better
- QoE => 1-22% better
- Rebuffering => 55-85 % lower
- Smoothness => 13-39% lower
- Energy => 35-57% savings
- Data => 41-60% savings