

Poster: Bringing Temperature-Awareness to Millimeter-Wave Networks

Moh Sabbir Saadat; Sanjib Sur; Srihari Nelakuditi

Department of Computer Science and Engineering; University of South Carolina, Columbia, USA

msaadat@email.sc.edu; sur@cse.sc.edu; srihari@sc.edu

ABSTRACT

Millimeter-wave devices operate at very high frequency and ultra-wide bandwidth. They consume more energy, dissipate more power, and heat up faster. So, millimeter-wave (mmWave) would exacerbate the device overheating problem in the future. In this work, we first perform a thermal characterization of mmWave devices: it reveals that after only 10 s. of data transfer at 1.9 Gbps, the antenna temperature reaches 68°C; it reduces the link throughput by 21%, increases the standard deviation by 6×, and takes 130 s. to dissipate the heat completely. We then propose *Aquilo* to bring temperature-awareness in mmWave networks; *Aquilo* maintains relatively high throughput performance and cools down the devices substantially. Our testbed experiments in static conditions show that *Aquilo* reaches a median peak temperature just 1°C above the optimal with less than 10% throughput sacrifice only.

CCS CONCEPTS

• Networks → Network protocol design;

1 INTRODUCTION

Millimeter-wave (mmWave) is the core technology for applications demanding multi-Gbps throughput and ultra-low latency [1]: immersive VR/AR, tactile internet, autonomous vehicles safety, *etc.* But, the devices operate at very high frequency and bandwidth; so, compared to micro-wave, they consume more energy, dissipate more power, and subsequently heat up faster.

A temperature increase is disconcerting to the users, especially when devices are small, hand-held, body-worn, and near to the face and brain. Thus, investigating ways to mitigate thermal-inefficiencies in mmWave devices is of vital importance. While existing research works have extensively characterized channel, link, network, applications, and power consumption, to the best of our knowledge, *none has looked at the thermal characteristics of the mmWave devices*. To this end, we first characterize the thermal profile of a Commercial-Of-The-Shelf (COTS) 60 GHz mmWave smartphone [2]. Our measurements reveal that after only 10 s. of data transfer, at room temperature, with 1.9 Gbps bit-rate, the mmWave antenna temperature reaches up to 68°C; it reduces the average link throughput by more than 21%, increases the standard deviation of throughput by 6×, and takes about 130 s. to dissipate the heat completely. Driven by

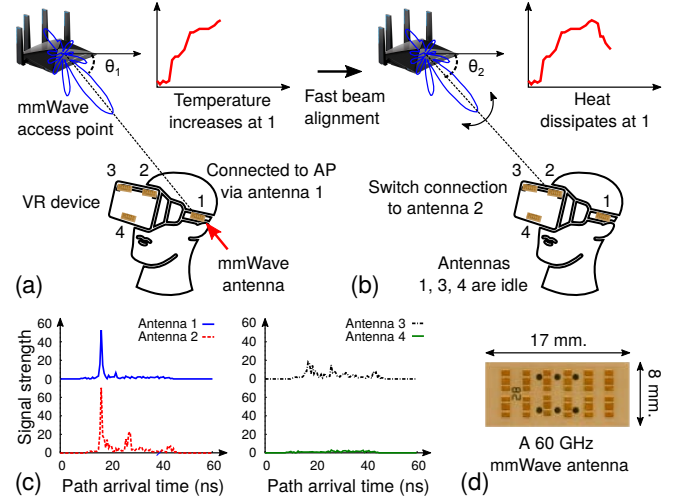


Figure 1: Core idea behind *Aquilo*: (a) VR device with 4 mmWave antennas connected to the access point (AP) via antenna 1; data transfer increases its temperature; (b) Scheduler switches to a relatively cooler antenna 2 with the best link to the AP; (c) Signal strength of the best beam from the 4 antennas; (d) Dimensions of a 60 GHz mmWave antenna.

the measurement insights, we also propose *Aquilo*¹ — a temperature-aware multi-antenna scheduler that cools down mmWave devices substantially. *Aquilo*'s key idea is intuitive: Before one antenna heats up excessively, its data stream may be switched or distributed to other redundant antennas, allowing it to dissipate the heat (see Figures 1[a–d]). We propose a smart, adaptive multi-antenna scheduling technique that exploits the near-past observation of the thermal profiles, and probe and switch scheme to maintaining relatively stable throughput performance while reducing the overall device temperature. We have validated *Aquilo*'s feasibility on an IEEE 802.11ad mmWave testbed (Figure 2[a]). Our testbed experiments demonstrate that, compared to a *throughput-only maximization* scheduling, under static conditions, *Aquilo* can effectively reduce the median peak temperature by 12°C. While this improvement comes from sacrificing 9.8% of throughput, *Aquilo* can still support at least 1.4 Gbps throughput at all times. In summary, we have two main contributions:

(1) **Thermal Characterization of 60 GHz mmWave Device:** To the best of our knowledge, we are the first to perform a study of the effects of mmWave device states and link performance on the device's temperature and the effect of temperature on the device's performance.

(2) **Temperature-Aware Multi-Antenna Scheduler:** Based on the insights from our thermal characterization, we propose, design, and

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MobiCom '20, September 21–25, 2020, London, United Kingdom

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7085-1/20/09...\$15.00

<https://doi.org/10.1145/3372224.3418159>

¹*Aquilo* was the Roman god of cold north wind and bringer of winter.

validate a temperature-aware multi-antenna scheduler and demonstrate its effectiveness in maintaining the link performance while reducing the temperature substantially.

2 THERMAL CHARACTERIZATION

Thermal Profile In Idle and Active States: We start by understanding the impact of the device’s idle and active states on the temperature. We first measure the temperature when the mmWave antenna is idle; it shows 48.63°C, about 24°C higher than the room temperature. This high idle temperature is due to the very high power consumption during idle listening [3] — the device needs to listen to the incoming mmWave packets and assess the clear channel condition continuously. Active states consume more power; thus, it increases the antenna temperature further. Nevertheless, we expect that, as soon as the data transfer stops, the antenna would quickly cool down and reach its idle temperature. But, heat acts just like stored energy in a capacitor — it takes a long time to dissipate, and more the data transfer, longer the heat dissipation time. Figures 2(b–e) plot two examples of antenna throughput and temperature for 5 s. and 10 s. of data transfers. They show that a 10 s. of transfer requires about 130 s. to dissipate the heat completely. An increase in the transfer duration also increases the dissipation time; our measurement in Figure 2(f) is in agreement. More importantly, *the average dissipation time can be up to 20× higher than the transfer duration.*

Performance Fluctuations with High Temperature: *High antenna temperature can adversely affect the throughput performance* due to increased thermal noise and leakage current. Figure 3(a) shows an example RSS and throughput profile; even under the stable channel condition, the higher temperature causes significant throughput fluctuation and average throughput reduction. Figure 3(b) further shows that the standard deviation of the throughput increases by almost 6×; it can reach up to 240 Mbps for a 14° rise in temperature. Thus, *higher antenna temperature not only degrades average performance but also increases its variations significantly.*

Effect of Periodic Idleness: A strawman approach to reducing the antenna temperature is to keep it idle in between the data transfers since *an idle period helps the cooling process.* To understand this effect, we run experiments with three data transfer modes: continuous for 270 s; periodic for 10 s. followed by idleness for 5 s. (66.7% duty cycle); and continuous for 180 s. (66.7% of 270 s.). Figures 3(c–d) show the throughput and temperature profiles. Under continuous transfer, the temperatures grow steadily, and it takes a long time to cool after the transfer stops. Introducing periodic idleness of 5 s. after every 10 s. of data transfer reduces the temperature substantially; the peak temperature is below 64°C, for 96% of the time. Unfortunately, for certain devices and applications, *e.g., VR/AR, periodic idleness is unaffordable, since they require stringent throughput and latency guarantees.*

3 AQUILO DESIGN

We propose *Aquilo* — a thermal profile based multi-antenna scheduler to maximize the throughput and minimize the temperature. *Aquilo* enables a mmWave AP and user device to select cooler antennas dynamically. Unfortunately, the objectives of minimizing temperature and maximizing performance, in a single antenna system, are perennially in conflict; this is because an antenna’s temperature rises steadily while transferring data continuously. Fortunately, upcoming

mmWave systems, like 5G NR smartphones and VR/AR devices, are being equipped with multiple antennas [4]; *Aquilo* leverages the presence of and coordination among these multiple antennas to reduce the overall system’s temperature.

Multi-Antenna Coordination: At a high level, before one antenna heats up excessively, its data stream may be switched to the other antennas, allowing it to dissipate heat. However, our measurements show that the thermal behavior of mmWave antenna changes over time and space; moreover, the mmWave connectivity is also highly variable and unpredictable. So, *it is non-trivial to select an antenna that can not only sustain the connectivity but also reduce the system’s temperature.*

Online Thermal Profile Estimation: To overcome the challenges from variable thermal profile, *Aquilo* leverages online measurements from near-past time by recording the antennas’ temperature when they are active or idle. From our measurements in Figures 2(d–e), we can model the thermal profile of a mmWave antenna using two exponentials: *exponential gain*, $e^{\alpha \cdot t}$, in the active state; and *exponential fall*, $e^{-\beta \cdot t}$, in the idle state. Furthermore, we can estimate the α and β parameters from the near-past temperature observations.

Look-ahead Schedule, Antenna Probe and Switch: *Aquilo* leverages the estimated thermal profiles to find a list of antenna schedules. Under each schedule, the antennas will go through a temperature transformation, and thus, they would potentially reach different peak temperatures. Therefore, we can pre-compute the peak temperature attained by each schedule and select the one with the lowest peak. However, mmWave connectivity is highly variable and unpredictable; thus, the selected antennas may not have a strong enough link to sustain an application’s performance requirement. *Aquilo* proposes an adaptive scheme to overcome this challenge: Under each schedule from a temperature-based sorted list, *Aquilo* will probe and switch to only those antennas that can sustain connectivity. Since the probing takes a relatively small amount of time with state-of-the-art beam alignment protocols, *Aquilo*’s latency overhead will be very low.

Example Schedule and Performance from *Aquilo*: Figures 4(a–b) show example peak temperature under *Aquilo*, and contrast the results with a throughput optimization scheme. The throughput optimization achieves better throughput; it’s about 120 Mbps higher than *Aquilo*. However, it suffers from a very high system’s temperature; the peak is 71.85°C! *Aquilo* never crosses beyond 55.20°C; so, it achieves 16.65°C temperature reduction sacrificing only 6% throughput.

4 EVALUATION

We evaluated the preliminary effectiveness of *Aquilo* against other scheduling schemes: (a) *Best Case Temperature* serves as the upper-bound of temperature; (b) *Random Scheduling* selects antennas randomly but ensures they meet the performance requirement; (c) *Random with Non-Adjacency Criteria (NAC)* selects and assigns random antennas into nonadjacent timeslots; and (d) *Throughput Optimization* maximizes the throughput only. Figures 4(c–d) show the CDF of peak antenna temperatures and system throughput for 100 s. operation across 200 static scenarios. While the median of best-case peak temperatures is 53.64°C, the throughput optimization shows the worst-case temperature performance with its median beyond 67°C. A simple heuristic of NAC effectively improves the median peak

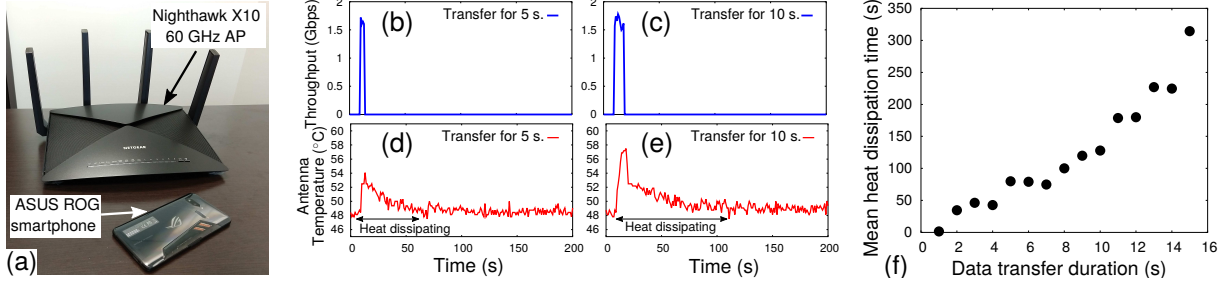


Figure 2: (a) Millimeter-wave AP and smartphone. (b–c) Data transfers for 5 s. and 10 s. (d–e) Antenna temperatures over time and heat dissipation durations. (f) Mean heat dissipation time.

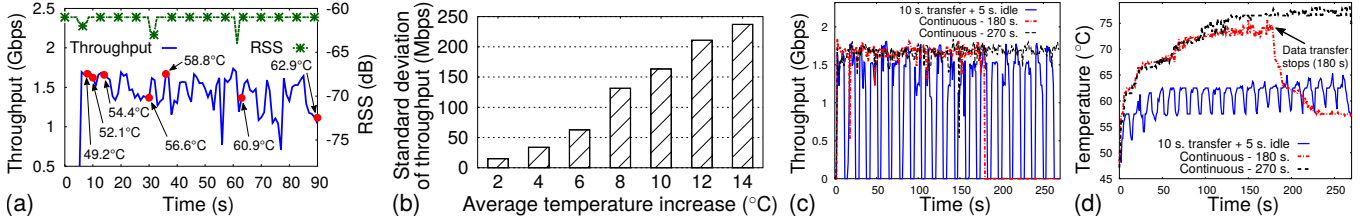


Figure 3: Link performance is affected by higher antenna temperature: (a) RSS is stable, but throughput fluctuates over time; every time index is marked for 2°C increase; (b) For every 2°C increase, we see an increase in the standard deviation of throughput. (c) Throughput of the link; and (d) Temperature profiles of the antenna under three data transfer modes.

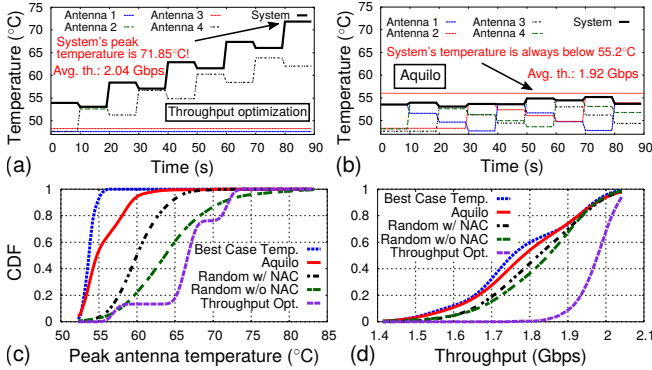


Figure 4: (a–b) Example temperature changes for a near-static 90 s. of data transfer under Aquilo and throughput optimization. Under throughput optimization, peak temperature can reach 71.85°C ; but, under Aquilo, the system’s temperature never crosses beyond 55.20°C . (c–d) Empirical CDF results from different schemes across various static conditions: peak antenna temperature; and system’s throughput.

temperature by more than 3.7°C from the random selection. *Aquilo* outperforms all the random schedulings and throughput optimization schemes in terms of temperature; its median peak temperature is just about 1°C above the best case. Compared to the throughput optimization, *Aquilo* suffers from around 200 Mbps throughput loss ($\sim 9.8\%$); but, in more than half of the cases, *Aquilo* reduces the peak temperature by approximately 12°C . In the future, we will evaluate *Aquilo* under different environmental dynamism, link throughput

requirements, system parameters, mobility patterns, and end-user applications.

5 CONCLUSION

Our work presents the first-of-a-kind study on mmWave thermal characterization; it reveals new challenges and opportunities to keep IEEE 802.11ad and 5G NR devices cool. We also propose *Aquilo*, a multi-antenna scheduler to bring temperature-awareness at mmWave, and show its preliminary effectiveness on COTS devices. In the future, we will investigate performance fluctuations and characterizing thermal performance under different applications and use cases, both indoor and outdoor; we will also design, implement, and evaluate a real-time *Aquilo*. Overall, we believe, this work helps allay concerns in some quarters about the health effects of 5G and accelerates its deployment broadly.

6 ACKNOWLEDGEMENT

We sincerely thank the anonymous reviewers. The authors were partially supported by the NSF under Grant CNS-1910853.

REFERENCES

- [1] 3GPP: A Global Initiative, “The Mobile Broadband Standard: Release 16,” 2019. [Online]. Available: <http://www.3gpp.org/release-16>
- [2] AsusTek Computer Inc., “ASUS ROG Phone,” 2018. [Online]. Available: <https://www.asus.com/us/Phone/ROG-Phone/>
- [3] Swetank Kumar Saha, et. al, “A Detailed Look into Power Consumption of Commodity 60 GHz Devices,” in *IEEE WoWMoM*, 2017.
- [4] Qualcomm Inc., “Qualcomm Announces 5G NR mmWave Prototype to Accelerate Mobile Deployments for Smartphones,” 2017. [Online]. Available: www.shorturl.at/hiqvy