

Toward End-to-End, Full-Stack 6G Terahertz Networks

Michele Polese, Josep Miquel Jornet, Tommaso Melodia, and Michele Zorzi

This article provides an overview of the issues that need to be overcome to introduce the terahertz spectrum in mobile networks, from a MAC, network, and transport layer perspective, with consideration on the performance of end-to-end data flows on terahertz connections.

ABSTRACT

Recent evolutions in semiconductors have brought the terahertz band into the spotlight as an enabler for terabit-per-second communications in 6G networks. Most of the research so far, however, has focused on understanding the physics of terahertz devices, circuitry, and propagation, and on studying physical layer solutions. However, integrating this technology in complex mobile networks requires proper design of the full communication stack, to address link- and system-level challenges related to network setup, management, coordination, energy efficiency, and end-to-end connectivity. This article provides an overview of the issues that need to be overcome to introduce the terahertz spectrum in mobile networks, from a MAC, network, and transport layer perspective, with consideration on the performance of end-to-end data flows on terahertz connections.

INTRODUCTION

Wireless communications will fundamentally shape the innovations of the digital society toward 2030. The unprecedented growth of mobile traffic has called for the integration of new portions of the spectrum in fifth generation (5G) mobile networks, that is, in the millimeter wave (mmWave) band. The Third Generation Partnership Project (3GPP) is already considering an extension to 71 GHz for 3GPP NR, as higher carrier frequencies come with larger bandwidth. For this reason, the terahertz bands are considered as a possible enabler of ultra-high data rates in sixth generation (6G) networks [1, 2]. Indeed, the spectrum from 100 GHz to 10 THz features wide chunks of untapped bandwidth for communications and sensing. Notably, the IEEE has developed a physical layer that spans 50 GHz of bandwidth, between 275 and 325 GHz.

Terahertz frequencies, however, bring to the extreme the communications and networking challenges of the lower mmWave band. In particular, the harsh propagation environment features high path loss, inversely proportional to the square of the wavelength, and thus to the size of a single antenna element, and, in addition, high molecular absorption in certain frequency bands. Moreover, terahertz signals do not penetrate common materials, and are thus subject to blockage. Finally, the manufacturing of terahertz devices has been a challenge for years, and only

very recent advances in electronics and photonics have enabled portable terahertz equipment [3].

Nonetheless, the promise of multi-gigabit-per-second capacity has sparked research efforts to overcome these challenges with novel, efficient, and high-performance physical layer techniques. Several studies have focused on increasing the communication range in macro scenarios [4], and on signal generation and modulation. Directional antennas are used to mitigate the increased path loss, as they can focus the power in narrow beams, which increase the link budget, and to enhance the security of wireless links, making eavesdropping more challenging. Furthermore, the small wavelength at terahertz allows many antenna elements to be packed in a small form factor (1024 in 1 mm² at 1 THz), thus enabling ultra-massive multiple-input multiple-output (UM-MIMO) techniques [3] and array-of-subarrays solutions [5]. Finally, reconfigurable electronic surfaces can act as smart reflectors and overcome blockage in non-line-of-sight (NLoS) [4].

These studies, however, mostly focused on modeling, designing, and optimizing the physical layer and the RF circuitry. Eventually, only a fraction of the hops between a client and a server will be on terahertz links. Therefore, it is necessary to study and understand the performance of this technology considering the integration in complex, end-to-end networks, where multiple nodes and layers of the protocol stack interact to deliver packets between two applications at the two endpoints of a connection. Moreover, the harsh propagation characteristics of the terahertz band, the limited coverage of a terahertz access point, the directionality, and the huge availability of bandwidth introduce new challenges and potential for the medium access control (MAC), network, and transport layers, and may call for radical re-design of traditional paradigms for user and control planes of wireless networks.

In this article, we discuss five key areas for the development of end-to-end terahertz networks, summarized in Fig. 1, by reviewing literature contributions and providing novel results based on full-stack, end-to-end ns-3 simulations [6].¹ To the best of our knowledge, this is the first contribution that provides a holistic perspective on networking challenges at terahertz frequencies, as prior work mostly focused on the lower portion of the spectrum [7] or on terahertz devices and the physical layer [4, 8]. Notably, we believe it will be crucial to study mechanisms to:

¹ The simulation scripts can be found at <https://github.com/mychele/toward-e2e-6g-terahertz-networks>

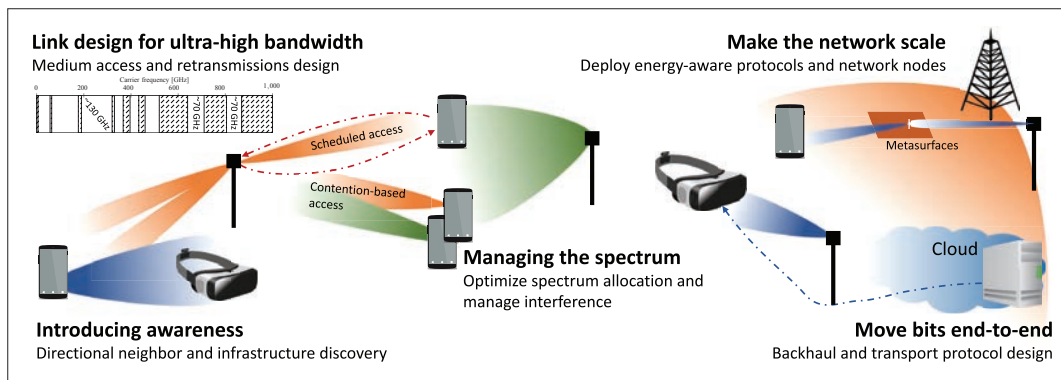


Figure 1. Main design challenges for end-to-end, full-stack terahertz networks.

- **Introduce awareness** of neighbors, fixed infrastructure, and channel usage, overcoming the deafness introduced by directional communications
- Account for the **ultra-high bandwidth in the link design**, by analyzing the complexity-efficiency trade-off of medium access schemes
- **Make the network scale** with a joint, energy-aware design of the network deployment, considering active and passive nodes
- **Manage the spectrum**, understanding the impact of interference and how the available bandwidth benefits wireless backhaul and multi-connectivity
- **Move bits end to end**, analyzing which kinds of transport protocols may provide the best performance

We selected these topics as they represent elements that are impacted by the characteristics of the terahertz channel and deployments, and can be actually optimized to provide a seamless end-to-end 6G experience on terahertz links. They are discussed in the remainder of this article: for each area, we first highlight the challenges introduced by terahertz scenarios, and then outline promising research direction.

INTRODUCING AWARENESS

Mobile nodes need to gather awareness of the surrounding environment. Even in a random access context, without any coordination, devices still need to be aware of incoming transmissions. Moreover, in a cellular or wireless LAN (WLAN) scenario, where mobile users exchange data with a fixed infrastructure (i.e., base stations and access points), each endpoint of the wireless link should have knowledge of the other: in cellular networks, users perform initial access to a base station, which is a gateway to the overall network, and schedules spectrum resources for the connected devices. Finally, they may need to sense the channel to understand if it is busy, and decide to transmit only if idle. Therefore, it is necessary for mobile transceivers to gain awareness of the fixed infrastructure, of their neighbors, and, if needed, of the channel occupancy.

Traditionally, cellular and WLAN networks at sub-6 GHz use signals broadcasted (quasi) omnidirectionally, which are not subject to deafness. A synchronization signal sent by Long Term Evolution (LTE) base stations, for example, can be received by all the users in their coverage area, simultaneously. In 5G networks with mmWave

communications, however, this paradigm has changed, as directionality prevents devices and base stations from transmitting and receiving omnidirectionally. Such systems, indeed, need beamforming to improve the link budget and extend the possible communication range [7]. To achieve the maximum gain, the endpoints need to align their transmit and receive beams, and this may introduce delays in the link setup, limit the awareness of the spectrum utilization, and impair the reliability of communications in highly mobile scenarios.

LINK BUDGET AND INITIAL ACCESS LATENCY COMPARISON

Figure 2a compares the signal-to-noise ratio (SNR) at different distances (5, 30, and 100 m) for mmWave and terahertz links. The SNR is given by the ratio between the received signal power and the noise power, without the beamforming gain, to analyze the impact of the higher carrier frequency on the propagation loss. The latter is computed using established models (i.e., the 3GPP model in an urban canyon scenario in the frequency range considered for 3GPP NR [9], and [8] for terahertz links in the 300–1000 GHz spectrum). The bandwidth is 400 MHz for mmWaves (i.e., the maximum bandwidth per carrier of 3GPP NR) and 50 GHz for terahertz, compliant with IEEE 802.15.3d.

For example, we observe that the SNR gap between a carrier at 30 GHz and one at 430 GHz is 37 dB. As mentioned earlier, this difference can be compensated for by increasing the number of antenna elements in each node, thanks to the smaller wavelength at terahertz when compared to mmWaves and sub-6 GHz networks. This translates into narrower beams, which improve the link budget but, at the same time, increase the deafness problem, limiting the awareness of mobile nodes.

The impact of highly directional antennas on the MAC and higher layers is shown in Fig. 2b, which reports the latency for the initial link establishment, with an exhaustive scan to identify the best transmit and receive beam, using the frame structure of 3GPP NR and the latency model from [7]. Notably, the base stations send 8, 16, 32, or 64 directional synchronization signals every T_{ss} seconds, which allow mobile nodes to assess the channel quality and decide on the best beam pair. For the mmWave links, we consider a setup with 16 antennas at the base stations and 4 at the user equipments (UEs). This yields a gain of up to 18 dB, which is enough to have a line-of-sight

To achieve the maximum gain, the endpoints need to align their transmit and receive beams, and this may introduce delays in the link setup, limit the awareness of the spectrum utilization, and impair the reliability of communications in highly mobile scenarios.

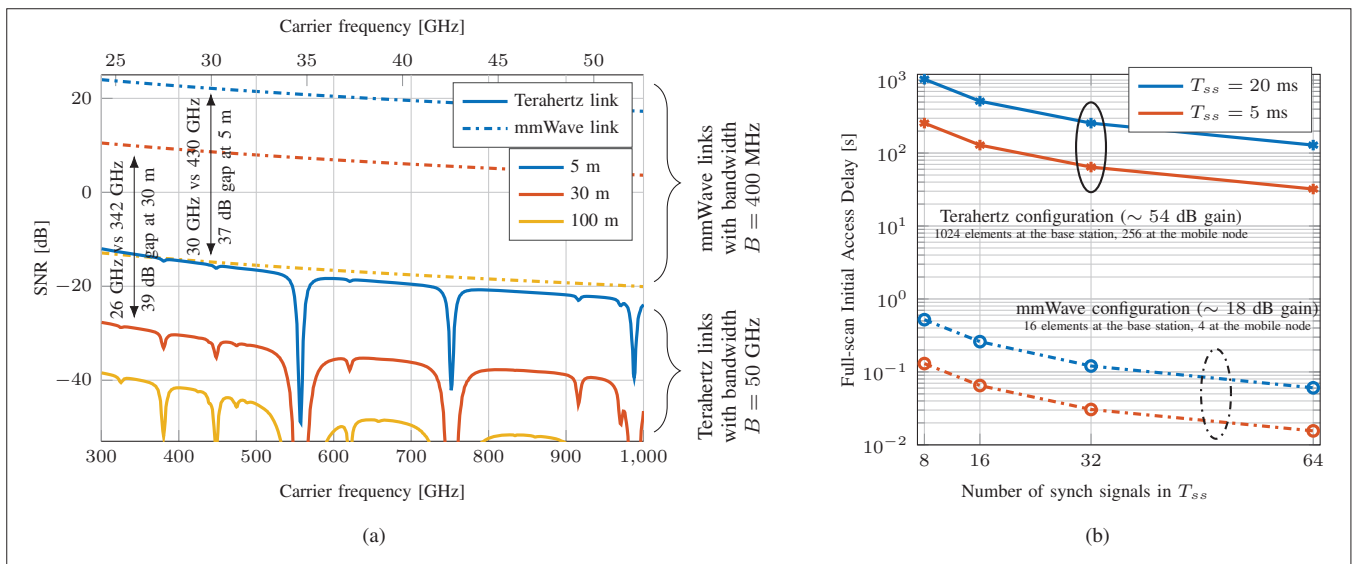


Figure 2. Link budget and impact of directionality at the MAC layer: a) link budget for mmWaves and terahertz links, without beamforming gain: $P_{tx} = 0.5$ W, noise figure $F = 10$ dB; b) link establishment delay computed with the 3GPP NR parameters and the model from [7], using antenna arrays that yield a comparable SNR at the receiver.

(LoS) SNR on the order of 0 dB at 100 m and 52.6 GHz. For terahertz, conversely, the configuration of Fig. 2b is designed to match the overall link budget of the mmWave setup, with 1024 elements at the base stations and 256 at the mobile endpoint, and a gain of up to 54 dB. While the link budget is the same, the tighter beam requires more directional synchronization signals for the same coverage area. Therefore, *an exhaustive search at the transmit and receive side on the 3GPP NR frame structure is slow but still feasible at mmWaves (with a delay smaller than 1 s in the worst configuration), but is not a viable option for terahertz, which would require more than 32 s to complete the scan in the best case.*

BEAM OPERATIONS FOR TERAHERTZ

Consequently, beam management protocols for the terahertz MAC layer have to carefully consider the trade-off between improving the link budget and the need for awareness of spectrum and infrastructure. Research on beam management solutions for mmWaves has already proposed several alternatives to a basic exhaustive search, which could be partially adapted to the terahertz domain [4]. Additionally, the characteristics of the terahertz spectrum could be exploited to improve directional operations [10]. Promising research directions will involve a combination of:

- A redesign of the frame structure with respect to that of NR considered in Fig. 2b, and of synchronization and reference signals, to exploit the larger bandwidth available at terahertz. Shorter pilots could result in more opportunities to transmit synchronization signals without impacting the control overhead, and multiple user-specific tracking signals could be multiplexed in frequency, if the link budget for the involved devices is high enough.
- Advanced antenna architectures could be exploited to transmit and receive directional signals from multiple transceivers at the same time (e.g., with digital beamforming), or to

infer angle of departure and arrival of terahertz signals (e.g., with leaky-wave antennas [11]). Further research efforts are required to design and realize digital architectures and antennas for the terahertz spectrum, but promising results have also been obtained with multi-beam solutions based on plasmonic nano-antennas [3], or by considering the information gathered not only through the main beam, but also with sidelobes.

- Multi-stage beam management schemes, where beam configurations with different beam widths and, consequently, gains are used for different steps of either tracking or channel sensing procedures.
- Context-based schemes, which use additional information to gather awareness of the surrounding environment. Notably, at terahertz, the large bandwidth and the propagation characteristics make the medium particularly suitable for radio frequency sensing [2]. Therefore, the same radio interface could be used to transmit data and to map the surrounding environment, identifying, for example, sources of blockage.
- MAC protocols that rely on multi-connectivity (i.e., the availability of multiple radios in the mobile devices and base stations) to exploit different frequency bands for different tasks: sub-6 GHz links can provide a control overlay, and assist beam management operations at terahertz with reliable feedback links.

LINK DESIGN FOR ULTRA-HIGH BANDWIDTH

The unprecedented availability of bandwidth at terahertz, as shown in Fig. 1, and the highly directional transmissions also call for further research on medium access and retransmission mechanisms. Notably, the high physical layer rate may allow MAC protocols to still achieve high throughput, while accepting a lower bandwidth/spectral efficiency in favor of simplicity. Indeed, protocols that require a high level of coordination may be impractical for terahertz links.

For the medium access, the main design choice is between scheduled and contention-based. Each of these paradigms presents potentials and limitations that could be exploited for different use cases and scenarios. In particular:

- With scheduled access, the base stations of a cellular network allocate time and frequency resources to the connected users, avoiding collisions and, consequently, increasing spectrum efficiency. Additionally, the centralized control at the base stations enables prompt link adaptation and tracking. Indeed, the base stations can allocate specific tracking or reference signals to connected users, and coordinate with the infrastructure to understand the evolution of the channel dynamics for a certain user. However, it may be challenging to establish and maintain connectivity toward the fixed infrastructure, given the harsh propagation environment in the terahertz band, and the control plane needs to cope with directionality.
- Most research related to terahertz MAC layer design has proposed approaches based on contention-based medium access [6]. With respect to a scheduled MAC, a contention-based MAC does not need control plane connectivity toward the infrastructure. Moreover, the highly directional transmissions and their short duration (thanks to the ultra-high available data rate) limit the impact of collisions among different concurrent data exchanges. Nonetheless, the two endpoints of the communication still need to agree on the optimal beam pair; thus, a beam search step needs to be included at every channel access [10].

For retransmissions, unpredictable blockers and frequent beam or base station updates could disrupt the constant flow of acknowledgments between the receiver and the transmitter, thus impacting the design of retransmission policies. Additionally, given the high data rate, coding for forward error correction could make the transceiver design overly complex. Therefore, an efficient retransmission process remains an open issue. To this end, network coding techniques [12] have low-complexity implementations and simplify the retransmission process, as they do not require the retransmission of a specific packet, but rather of a random combination of the batch of packets that the receiver needs to decode.

MAKING THE NETWORK SCALE

mmWave base stations are being deployed as high-density small cells in 5G networks, since their coverage is limited [4]. Following this trend, and considering the higher path loss in the terahertz spectrum, the density of a cellular network operating at terahertz will likely increase even further. This will have an impact on the capital and operating expenditures, the energy consumption, and the complexity and scale of the backhaul infrastructure. In this section, we review the main challenges and research directions associated with the deployment and operation of a terahertz network at a massive scale.

Figure 3 illustrates the differences in deployment density that can be expected for mmWave

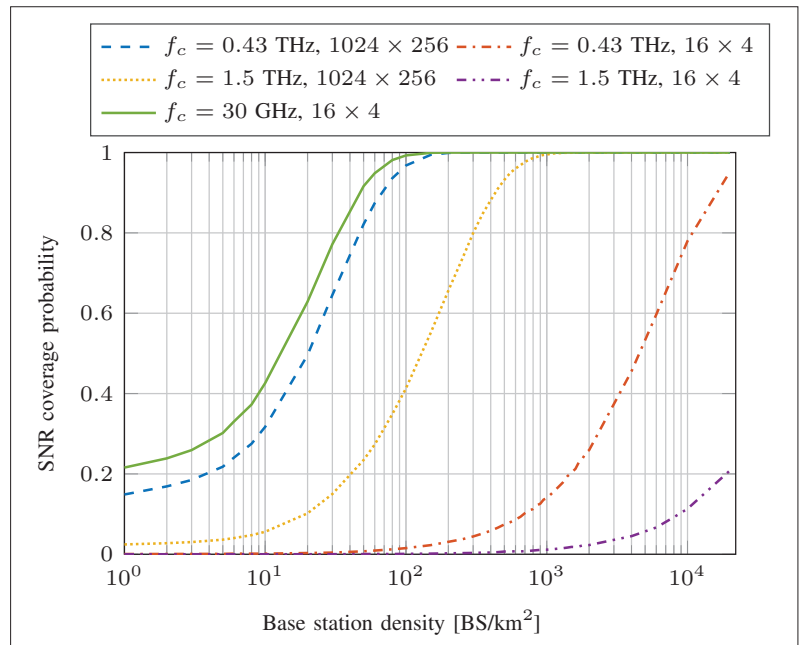


Figure 3. Probability of having an SNR higher than 0 dB for different deployment configurations. The mmWave channel is modeled according to [9] (i.e., through a combination of probabilistic LoS and path loss models for an urban scenario with microcells). For terahertz, the path loss of an LoS-only channel is modeled with [8]. The other link parameters are as in Fig. 2a.

and terahertz systems. We consider Monte Carlo simulations, in which macro base stations are randomly deployed outdoors according to a Poisson point process, and with the coverage probability computed as the probability of having an SNR above a threshold of 0 dB between a test user, at the center of the deployment, and at least one base station. For example, this can be a metric that is measured by the mobile devices during an initial access attempt [7]. The path loss and beamforming gain are modeled as in Fig. 2. We notice that with 16 antenna elements at the base station and 4 at the mobile device, the 30 GHz network has a coverage probability higher than 0.95 with 60 base stations/km², which corresponds to an average cell radius of 72 m. The same antenna configuration does not guarantee adequate performance at terahertz, since similar coverage probability requires $2 \cdot 10^4$ base stations/km² at 0.43 THz. As discussed earlier, antenna arrays with a larger number of antenna elements will be necessary when considering these frequency bands. With 1024 antenna elements for the base stations and 256 for the mobile devices, it is possible to reach a coverage probability of 0.95 with at least 100 base stations/km² at 0.43 THz, and 600 base stations/km² at 1.5 THz, with an average cell radius of 56 and 23 m, respectively.

ENERGY-AWARE NETWORK DESIGN

6G will further improve the energy efficiency of 5G networks [1], to balance the higher number of terahertz nodes to be powered up than at mmWaves or sub-6 GHz. Possible strategies toward the reduction of energy consumption at terahertz include:

- A lean control plane design that minimizes the control messages and always-on signals for control operations.

Interference management strategies should be re-designed to account for both active and passive users (e.g., to protect incumbents that use the spectrum for radio astronomy and Earth atmospheric science), with coordination loops across terahertz nodes that need to be fast enough to address a highly dynamic interference environment.

- Quick sleep cycles for base stations in dense clusters with low traffic. With the deployment density foreseen in Fig. 3, the number of base stations will likely approach and, in some areas, may even exceed that of active users in the network. Therefore, terahertz base stations should quickly turn off the main radio functionalities (to save energy) when there are no connected users. This will require coordination among neighboring base stations, which can notify each other of users relocating to switched-off base stations, so that they can swiftly resume operations, or wake-up radio systems, exploiting multi-connectivity with low-power radios.
- Energy-saving states for mobile devices without active transmissions, alternating stand-by with intervals in which the network could page the device. A specific challenge for terahertz is maintaining connectivity in stand-by with directional links, as the best beam pair could change if the mobile device in a low-energy mode moves through the network.
- Energy harvesting, with the infrastructure and the mobile devices exploiting modern harvesting circuitry to self-sustain during sleep cycles.

A CONTROL PLANE FOR REFLECTING ARRAYS AND METASURFACES

Even with large antenna arrays, uniform access coverage with terahertz links may be infeasible in terms of cost and energy consumption. Hotspots and indoor scenarios are more likely candidates for early terahertz deployments. For these, the traditional base station coverage can be enhanced with new network infrastructure elements, namely, reflecting arrays and metasurfaces [4]. These devices are based on phased arrays and nanomaterials, respectively, that can steer the terahertz wave impinging on them, thus reflecting the signal transmitted by a terahertz node toward a mobile user. Their deployment improves the link budget and the coverage in NLoS conditions, and reduces the density of full base stations. However, their integration with the fixed infrastructure requires the design of a dedicated control plane, with protocols and networking procedures to manage, among others, the handoff of users across different reflecting devices and the tracking in highly mobile scenarios. Moreover, the scale of terahertz networks will make manual configuration impractical, calling for intelligent procedures that automatically connect and jointly optimize the parameters of base stations and reflecting nodes in a plug-and-play fashion.

MANAGING THE SPECTRUM

The large untapped portions of spectrum at terahertz also create new opportunities for spectrum management. Indeed, besides allocating large bandwidths to the radio access, it is possible to enhance the network by deploying novel spectrum reuse schemes, in-band, high-capacity wireless backhaul, and multi-connectivity.

INTERFERENCE

First of all, it is important to characterize the impact of interference when considering the extremely directional links [13]. Indeed, as for mmWaves,

terahertz networks can be noise-limited thanks to beamformed transmissions and the limited coverage of each base station. However, at the same time, the high density that is needed to provide coverage (as discussed earlier) may introduce additional interference. Finally, the bursty transmissions complicate the tracking and prediction of possible interference sources. To this end, [14] proposes a reinforcement learning scheme to detect and mitigate intermittent interference sources, showing how adaptive, learning-based approaches can cope with challenging and dynamic terahertz scenarios. The behavior of interference at terahertz can affect the design of spectrum reuse and sharing schemes, which may be tuned as more or less aggressive according to the need to isolate from cross-cell interference. Additionally, interference management strategies should be re-designed to account for both active and passive users (e.g., to protect incumbents that use the spectrum for radio astronomy and Earth atmospheric science), with coordination loops across terahertz nodes that need to be fast enough to address a highly dynamic interference environment.

WIRELESS BACKHAUL

The large available bandwidth can also be used for in-band backhaul for terahertz base stations. The high deployment density will make wired backhaul to each base station extremely expensive, thus calling for a fully wireless solution, such as that proposed with integrated access and backhaul (IAB) in 5G networks. Compared to IAB at mmWaves, terahertz benefits from the higher spectrum availability, which could improve the quality of service for end-to-end traffic flows over multiple wireless hops. Moreover, fixed relays simplify beam management, but network operators should carefully design the deployment and the topology of the wireless backhaul network to provide connectivity to all the base stations in the presence of blockage.

MULTI-CONNECTIVITY

Finally, 6G networks will rely on a combination of sub-6 GHz, mmWave and terahertz bands, and, possibly, optical wireless links [1, 2]. The network infrastructure and the mobile devices will therefore need to nimbly adapt and use the carrier that provides the best performance.

In Fig. 4 we compare the throughput of a terahertz link at 1.0345 THz, with 74 GHz of bandwidth (corresponding to the first available window in the spectrum above 1 THz [6]), and an mmWave link operating at 28 GHz with the maximum bandwidth allowed in 3GPP NR (i.e., 400 MHz), through full-stack simulations in ns-3 with the TeraSim [6] and mmWave modules [15]. The application is a constant bit rate source, with UDP at the transport layer. The wireless nodes are equipped with directional arrays: a rotating array for terahertz, as in the default *macro* scenario of TeraSim [6], and phased arrays with 64 and 16 antenna elements at the base station and the device [15].

This network configuration highlights two operating regimes, according to the distance between the base station and the mobile device. For short distances (i.e., 1 and 5 m), the terahertz link benefits from the larger bandwidth with respect to

the mmWave connection, and provides higher throughput. However, as the distance increases, the path loss for the terahertz link increases at a faster rate than for the mmWave one, which eventually becomes the better choice in terms of throughput at 10 and 20 m. Therefore, 6G terahertz devices should exploit multi-connectivity not only for the control plane and beam management, as mentioned earlier, but also for the user plane, forwarding data packets on the different available radio interfaces to provide diversity.

MOVING BITS END TO END

As discussed earlier, terahertz links in cellular or ad hoc networks will constitute only a fraction of the hops in an end-to-end connection, and will carry traffic generated by a wide range of different applications with various underlying transport protocols (e.g., TCP). The end-to-end performance is thus determined by the interaction between these applications and protocols and the resources at the physical layer, and a sub-optimal interplay may prevent full exploitation of the large bandwidth and data rates of terahertz connections. Prior research has shown that, at mmWaves, the highly intermittent channel and beamforming degrade the performance of traditional TCP congestion control schemes [15]. Therefore, the interplay with the transport layer should be considered when designing the protocol stack for terahertz links as well.

Figure 5 exemplifies the pitfalls of TCP for terahertz links by comparing the evolution of the congestion window for a single TCP flow on an mmWave link (28 GHz), with a scheduled MAC [7, 15], and a terahertz link (1.0345 THz) with two MAC layer configurations. The first, from [10], has beam management and contention, while the other is ideal (i.e., all the resources are always allocated to the same user with the best beam pair). We observe that the congestion window with the terahertz link and the realistic MAC is reduced multiple times, not because capacity is reached, but because of the inefficient interplay between the contention-based access and the TCP timers, which triggers timeouts and congestion recovery. Overall, this configuration performs much worse than the mmWave one, despite the larger bandwidth, with an average throughput of 66 Mb/s vs. 520 Mb/s. The ideal MAC configuration instead highlights another issue that TCP may suffer from on terahertz links: the sub-optimal use of the available physical layer resources due to the slow linear ramp-up of TCP in congestion avoidance.

Finally, current protocol stacks in mobile devices are not designed to handle data rates that can reach tens of gigabits per second. As the throughput increases, the CPU of the device becomes busier processing the received packets. Moreover, congestion and flow control decisions have to be made much more frequently. Therefore, the networking stack processing may quickly deplete the battery of mobile devices when operating at high throughput. This makes the case for additional simulation-based and experimental research on the design and performance of *simpler* network and transport protocols at terahertz (e.g., based on UDP, QUIC, or on future evolutions of TCP). Moreover, further analysis is needed to understand whether congestion control mechanisms

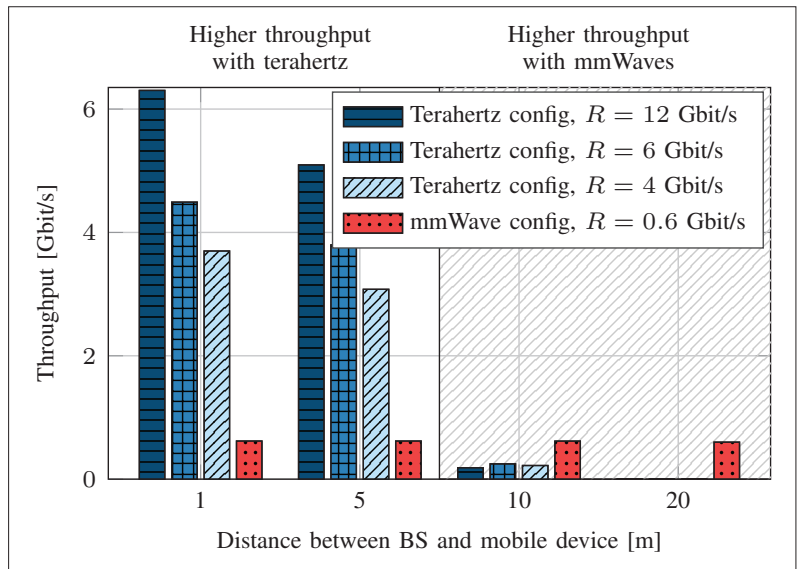


Figure 4. UDP throughput with different source rates R and distances between the base stations and mobile devices.

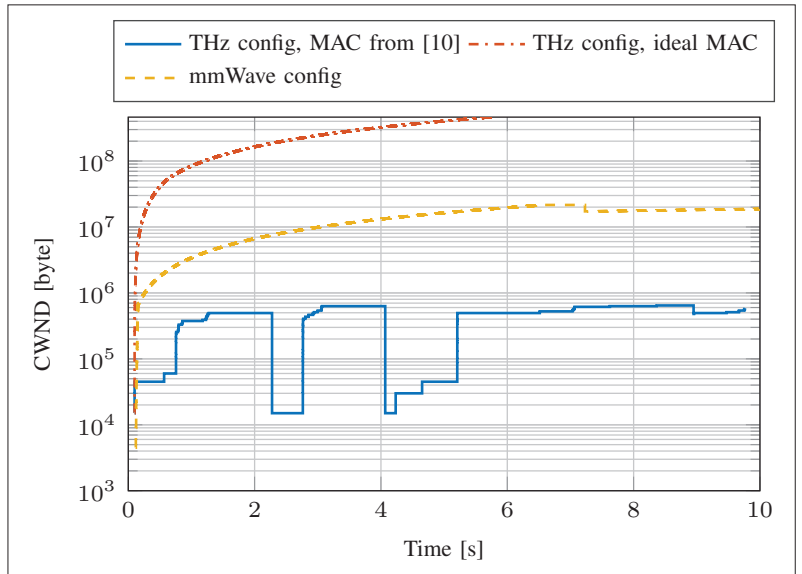


Figure 5. Evolution of the TCP CUBIC congestion window. The base station and mobile device are at a distance of 1 m, with the same configuration as in Fig. 4.

are useful for terahertz links, considering the data rates at stake.

CONCLUSIONS AND FUTURE WORK

The terahertz physical layer design still presents many open research challenges. This article focuses on key issues for the higher layers of the protocol stack, also discussing deployment and energy-related challenges.

We have identified relevant networking problems, providing insights and preliminary results that can drive future research on 6G terahertz networks at the MAC layer and above. Five key research questions emerged from our analysis:

- How can *beam management* and medium access schemes be designed to exploit the characteristics of the *terahertz spectrum*, for example, by combining sensing, large antenna arrays, and communications?

These research questions, and the insights we have provided in this article, could be used as a starting point to further progress the full-stack, end-to-end analysis and design of terahertz networks, to fully profit from the unprecedented amount of bandwidth available in this portion of the spectrum.

- Do terahertz networks need a *reliable control plane* and a scheduled MAC, or does contention-based access provide a better trade-off between *complexity* and *performance*? How can the control plane be extended to metasurfaces?
- Which are the most effective strategies to deploy a *high-density, energy-efficient* network?
- Which policies can be developed to manage the terahertz spectrum, considering dynamic *interference* sources, the large available bandwidth, and the possibility of using *multi-connectivity*?
- How can *transport protocol designs* and *implementations* evolve to satisfy the requirements of ultra-high bandwidth, highly variable links?

These research questions, and the insights we have provided in this article, could be used as a starting point to further progress the full-stack, end-to-end analysis and design of terahertz networks, to fully profit from the unprecedented amount of bandwidth available in this portion of the spectrum.

ACKNOWLEDGMENTS

The work of M. Polese and T. Melodia was supported in part by the U.S. Office of Naval Research under Grant N00014-20-1-2132. The work of J. M. Jornet was supported in part by the U.S. National Science Foundation, Grant CNS-2011411.

REFERENCES

- [1] M. Giordani *et al.*, "Toward 6G Networks: Use Cases and Technologies," *IEEE Commun. Mag.*, vol. 58, no. 3, Mar. 2020, pp. 55–61.
- [2] T. S. Rappaport *et al.*, "Wireless Communications and Applications Above 100 GHz: Opportunities and Challenges for 6G and Beyond," *IEEE Access*, vol. 7, 2019, pp. 78,729–57.
- [3] I. F. Akyildiz and J. M. Jornet, "Realizing Ultra-Massive MIMO (1024×1024) communication in the (0.06–10) Terahertz Band," *Nano Commun. Networks*, vol. 8, June 2016, pp. 46–54.
- [4] I. F. Akyildiz, C. Han, and S. Nie, "Combating the Distance Problem in the Millimeter Wave and Terahertz Frequency Bands," *IEEE Commun. Mag.*, vol. 56, no. 6, June 2018, pp. 102–08.
- [5] L. Yan, C. Han, and J. Yuan, "A Dynamic Array-of-Subarrays Architecture and Hybrid Precoding Algorithms for Terahertz Wireless Communications," *IEEE JSAC*, vol. 38, no. 9, Sept. 2020.
- [6] Z. Hossain, Q. Xia, and J. M. Jornet, "TeraSim: An ns-3 Extension to Simulate Terahertz-Band Communication Networks," *Nano Commun. Networks*, vol. 17, Sept. 2018, pp. 36–44.
- [7] M. Giordani *et al.*, "A Tutorial on Beam Management for 3GPP NR at mmWave Frequencies," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 1, 1st qtr. 2019, pp. 173–96.

- [8] J. M. Jornet and I. F. Akyildiz, "Channel Modeling and Capacity Analysis for Electromagnetic Wireless Nanonetworks in the Terahertz Band," *IEEE Trans. Wireless Commun.*, vol. 10, no. 10, Oct. 2011, pp. 3211–21.
- [9] 3GPP, "TR 38.901, Study on Channel Model for Frequencies from 0.5 to 100 GHz," V16.1.0, 2020.
- [10] Q. Xia *et al.*, "A Link-layer Synchronization and Medium Access Control Protocol for Terahertz-band Communication Networks," *IEEE Trans. Mobile Comp.*, Early Access, 2019.
- [11] Y. Ghasempour *et al.*, "Single-Shot Link Discovery for Terahertz Wireless Networks," *Nature Commun.*, vol. 11, no. 1, Apr. 2020, pp. 1–6.
- [12] C. V. Phung *et al.*, "Improving THz Quality-of-Transmission with Systematic RLNC and Auxiliary Channels," *IEEE ICC Wksp.*, 2020.
- [13] V. Petrov *et al.*, "Interference and SINR in Millimeter Wave and Terahertz Communication Systems With Blocking and Directional Antennas," *IEEE Trans. Wireless Commun.*, vol. 16, no. 3, Mar. 2017, pp. 1791–1808.
- [14] R. Barazideh *et al.*, "Reinforcement Learning for Mitigating Intermittent Interference in Terahertz Communication Networks," *IEEE ICC Wksp.*, June 2020.
- [15] M. Mezzavilla *et al.*, "End-to-End Simulation of 5G mmWave Networks," *IEEE Commun. Surveys & Tutorials*, vol. 20, no. 3, 3rd qtr. 2018, pp. 2237–63.

BIOGRAPHIES

MICHELE POLESE [M'20] (m.polese@northeastern.edu) is a research scientist at Northeastern University, Boston, Massachusetts. He obtained his Ph.D. from the University of Padova, Italy, in 2020, where he also was a postdoctoral researcher and adjunct professor. He visited NYU, AT&T Labs, and Northeastern University. His research focuses on protocols and architectures for future wireless networks.

JOSEP MIQUEL JORNET [M'13] (jmjornet@northeastern.edu) received his Ph.D. degree in electrical and computer engineering (ECE) from the Georgia Institute of Technology in 2013. Between 2013 and 2019, he was with the Department of Electrical Engineering at the University at Buffalo. Since August 2019, he has been an associate professor in the Department of ECE at Northeastern University. His research interests are in terahertz-band communications and wireless nano-bio-communication networks. He has co-authored more than 120 peer-reviewed scientific publications and one book, and has been granted three U.S. patents. He is serving as the Principal Investigator on multiple grants from U.S. federal agencies, including the National Science Foundation, the Air Force Office of Scientific Research, and the Air Force Research Laboratory.

TOMMASO MELODIA [F'18] (melodia@northeastern.edu) received his Ph.D. degree in ECE from the Georgia Institute of Technology in 2007. He is the William Lincoln Smith Professor at Northeastern University, the director of the Institute for the Wireless Internet of Things, and the director of research for the PAWR Project Office. His research focuses on modeling, optimization, and experimental evaluation of wireless networked systems. He serves as Editor-in-Chief of *Computer Networks*.

MICHELE ZORZI [F'07] (zorzi@dei.unipd.it) is with the Information Engineering Department of the University of Padova, focusing on wireless communications research. He was Editor-in-Chief of *IEEE Wireless Communications* from 2003 to 2005, *IEEE Transactions on Communications* from 2008 to 2011, and *IEEE Transactions on Cognitive Communications and Networking* from 2014 to 2018. He served ComSoc as a Member-at-Large of the Board of Governors from 2009 to 2011, as Director of Education and Training from 2014 to 2015, and as Director of Journals from 2020 to 2021.