NSF WORKSHOP ON LOW-LATENCY WIRELESS RANDOM-ACCESS

Contents

1. Overview

The NSF workshop took place at MIT (Cambridge, MA) on November 2-3, 2017. It included 23 invited speakers from industry and academia and almost a hundred attendees. The workshop was hosted by the Laboratory for Information and Decision Systems and was made possible by a generous support of the CCF and CNS programs at NSF, and of the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-09-39370.

Below we present an article written by Greta Friar surveying the main points of the workshop, followed by detailed log of talks and discussion sessions.

2. Gearing up for the Internet of Things (article written by Greta Friar, MIT News, Nov. 22, 2017)

National Science Foundation workshop, hosted by LIDS, brought together academia and industry to prepare next-generation (5G) wireless for machine-to-machine communication.

Telecommunications is gearing up for explosive growth of the Internet of Things (IoT), the massive number of devicessmart watches, smart thermostats, traffic and energy monitorsthat will be given network connectivity so that they can communicate and exchange data.

The big question raised by the IoT is capacityhow can a telecommunications network with limited spectrum serve thousands or millions of devices at once? Existing networks cannot support the addition of exponentially more devices, nor the near-instantaneous connection speed necessary for machine-to-machine communication. The next generation network, 5G, must be designed to meet these requirements.

How to prepare global wireless networks for the IoT was the topic of discussion at the "NSF Workshop on Low-Latency Wireless Random-Access," hosted by MIT's Laboratory for Information and Decision Systems (LIDS) and sponsored by the National Science Foundation (NSF). The two-day event took place on campus November 2-3, during which time more than twenty speakers from MIT, other universities, and companies including Qualcomm, SigFox, and Huawei presented their work on how to solve the network challenges created by the IoT.

”Developing solutions for massive multi-access wireless communications is a fascinating, domain-spanning challenge for researchers with considerable practical applications,” said MIT associate professor and LIDS faculty member Yury Polyanskiy, who organized the workshop.

The IoT will transform industries such as agriculture, transportation, utilities, athletics and more by introducing smart devices and many networked monitors. The monitors will provide live environmental feedback, which the devices can use to optimize their functions. Determining 5G standards and protocols that will allow these transformative technological developments is a complex undertaking.

There are numerous hurdles to redesigning wireless infrastructure for the IoT. First, current networks are optimized for a relatively small set of users, mostly human, sending large amounts of data phone call, a videoin a continuous stream, with different connections centrally organized. To accommodate the IoT, networks will have to manage decentralized, intermittent transmission of many small data packets from many, many more users, mostly machines.

This vast increase in machines on the network will lead to more interference, which causes latency or loss of connectivity. Just like people have trouble getting a call through at a large event where there are too many other people on their phones, machines will have trouble sending their data if the network gets overloaded with devices needing connectivity. Devices that spend too much time searching for a working connection will also wear out their batteries, an issue for IoT monitors that are meant to be very low energy and low cost.

Another big challenge for the IoT is that it requires extremely low latency – or lag. As anyone who has used the internet knows, data doesn’t always transmit at the speed you want it to. And though a half-loaded video or delayed text may be a pain for humans, for machines even a few milliseconds of lag can have serious consequences. (A smart car interpreting traffic data, for instance, cannot afford any lag.) This is why the goal for 5G is an ambitious one millisecond latency between devices.

These issues were discussed at length during the workshop. Two sessions were dedicated to information theory, with academic speakers sharing their models for better massive machine-type communications. Sessions also covered topics including reliability and security.

“We don’t have one thing to solve, we have many things to solve,” said the first speaker of the event, Christophe Fourtet, co-founder and CSO of telecommunications company SigFox.

The workshop provided a somewhat rare opportunity for academia and industry to exchange ideas, as these communities are often siloed from each other.

"Academics played a major role in establishing early telecommunication standards, but this has become more rare recently. With 5G calling for radically new simple network-access methods, though, it's a great opportunity for academics to play a big role again," said Polyanskiy.

Presenter Swarun Kumar, assistant professor at Carnegie Mellon University (CMU) and an MIT alumnus, echoed Polyanskiy's sentiments. "The time from academic research to industry in this field is too slow," Kumar said. He made an appeal to the members of industry present at the workshop to reach out to him and other academics, noting that they could solve problems more efficiently together. Kumar presented a local network that he and colleagues at CMU had implemented around the campus as a proof of concept for their innovations to network infrastructure.

After each set of presentations, speakers returned to the front of the room for a panel Q&A, and in between sessions, speakers and attendees had a chance to mingle and talk over refreshments. The event organizers affirmed that these moments were some of the most valuable of the workshop: a chance for everyone present to open up dialogue and, perhaps, plant the seeds of collaborations that will build a better network.

3. Session 1: "IoT in Industry and Systems"

3.1. **Christophe Fourtet (SigFox): SDR Cognitive Networks The Future of Massive IoT.** Machine-Type Communications (MTC) and the Internet of Things (IoT) call for simple (low-cost and energy efficient devices), reliable and universal implementations. However, these criteria are in tension with one another, and might lead to a paradox. For instance, large cells, which are beneficial for minimizing capex/opex, collide with the requirement for massive capacity per km^2 while using low-power and ultra-low-cost (possibly imperfect) devices. Additional challenges are presented by the goal to maintain communication using only tiny pieces of free spectrum. To achieve this, allocation protocols for unpredictable spectrum are needed. Simplicity of the devices operated by the network is another high priority. However, this simplicity should not interfere with the high-reliability demand. SigFox offers a new approach to successfully take on these challenges. The proposed framework combines clever selection of waveforms, exploitation of modern computer techniques to boost the network's cognition capability, advanced signal processing methods, massive multi-input multi-output (MIMO) channels, real-time data processing, and more. Above all, cognitive networks is highlighted as being a top-priority. Building on cutting edge technology, a cognitive network can successfully handle random-access and high spectrum sharing rates, while reducing interference. The implementation relies on abundantly deploying AI throughout the network, which enables the collection of meta-data from global activity without the need for coding or transmitting. SigFox view cognitive networks as a pivotal component for the deployment of IoT.

3.2. **Mérouane Debbah (Huawei): Cellular IoT: Opportunities and Challenges.** The 4.5G network offers new services and experiences, potentially giving rise to new markets. Such a network proposes boosted performance both in terms of GBps and connectivity. Focusing on IoT, low-power wide area networks (LPWAs) are expected to constitute 70 percent of all cellular connections. To support this, narrow band (NB) IoT combined with enhanced MCT (eMTC) and existing 4G networks emerges as suitable solution. It offers a wide range of applications, including bicycle sharing, smart gas meter, pet tracking, etc. NB-IoT thus became the new benchmark for 5G-IoT. SingleRAN is an NB-IoT implementation that exploits eMTC to better support LPWA. This allows smoothly upgrading to 5G-IoT, while offering various exciting performance gains. For instance, NB-IoT achieves 20dB coverage gain over GSM/LTE, 10-year battery life verified in smart meter scenarios, 6 types of E2E security protection over three layers, and low latency. As many countries aim to introduce commercial-use of 5G-IoT around 2020, NB-IoT is key in bringing 5G-IoT to reality.

3.3. **Swarun Kumar (CMU): Enabling City-Scale Low-Power Wireless Networks.** Low-power wide area networking (LP-WAN) supports simple and cheap devices with a 10 year battery life. It allows such devices to communicate at few KBps with base stations several miles away. The main challenges in deploying LP-WAN are coping with interference and accounting for the communication range. Interfering collisions emerge from the sheer density of nodes and the simplicity of the current multiple access channel (MAC) protocols. The range of LP-WAN drops by a factor of 10 in urban areas due to excessive multipath, shadowing, etc. Past solution do not seem to account for these problems in a satisfactory manner. *Choir*, presented in [Eletreby, Zhang, Kumar, Yağan, SIGCOMM '17], is a LP-WAN implementation that overcomes the challenges. It utilizes hardware imperfections of low-cost devices to support decoding of large number of interfering transmissions at a simple, single-antenna LP-WAN base station. It also achieves range extension by exploiting spatial correlations of the data. The system is already successfully deployed in the Carnegie-Mellon University campus and shows great promise. In the future, the authors aim to extend it to a programmable LP-WAN that spreads throughout the entire city of Pittsburgh.

3.4. **Piyush Gupta (Qualcomm): Massive Wide-Area Internet of Things: Some Design and Performance Aspects.** 5G cellular wireless networks will enhance performance while expanding the range of usage. Among other things, such networks offer enhanced mobile broadband, mission-critical services, massive IoT, etc. Massive IoT aims to connect anything and anywhere using efficient and low-cost communication. Such connectivity can bring to life ideas like smart cities, smart homes, agricultural and environmental sensing, asset tracking, etc. Two types of use case scenarios are to be considered. The 1st type is essentially static communication with light traffic (periodic or event-driven sporadic), e.g., meter sensing. The 2nd type has potentially high mobility and persistent traffic. Examples are wearable and asset tracking. Among the key issues that MTC/NB-IoT must account for are the lower power budget some devices are bound to and the potentially high pathloss between devices and base stations. Furthermore, in some cases only a small fraction of the devices in a certain proximity have a direct wide area network (WAN) connectivity. The solution proposed by Qualcomm uses relaying and mesh IoT. Namely, devices with high pathloss relays its signal through an another device with better pathloss. Even with one additional hop, the problem of high pathloss can often be significantly alleviated. Relaying can be preformed either by dedicated relay deployment or via mesh IoT. The first approach is well-studied but is bound high deployment cost. In the mesh IoT approach, where devices autonomously look for possible relays by utilizing the high density of IoT, many new design challenges are introduced. The main questions to be addressed are how to discover potential relay IoTs given that they are predominantly in a sleep state, and how to achieve reasonable end-to-end latency in spite of long sleep cycle. A possible design that accounts for these issues is proposed. IoT devices are in sleep state as often as possible to achieve the battery life goals. For IoT devices in coverage-challenged areas (e.g. building basement) to relay their data through better located devices, they should exploit the common awake period. Namely, the common awake period is used as the discovery period, where the transmitter and the receiver synchronize their awake cycles. To achieve this, devices send periodic broadcasts to enable discovery. Two modes for discovery, which are based on mobility, are considered: persistent mode for static devices and transient mode for mobile devices. The two modes are provisioned with separate resources and different modes of discovery can happen at different periodicity. Persistent mode is based on reservation, which allows utilizing the same resources from frame to frame since the devices are static. Transient mode is based on random selection of discovery resources in a way similar to long-term evolution directed (LTE-D).

3.5. **Discussion.** In the discussion session, the speakers shared their views on the future of IoT. Christophe Fourtet said that IoT is different from traditional communication in that it deals not with individuals, but with objects. He pointed out that there is no single technology that can cover IoT. He compared IoT with telephony, and concluded that IoT has pushed limits so low that new technological solutions are needed. Mérouane Debbah made three points. Firstly, IoT needs to learn from the space industry on satellite communication, because they both focus on long range low latency communication. Secondly, we need to think about how to make business out of IoT, which is currently unclear. Thirdly, there are some technologies that need to be built in, such as multiuser space-frequency block code. Swarun Kumar emphasized the importance of closing the gap between academia and industry. It would benefit both sides if academic ideas could rapidly be integrated in industrial implementations. Piyush Gupta agreed with the points made by other speakers. In addition, he mentioned that for IoT, different systems need to be brought together. A solution may span across different stacks. Afterwards, the audience asked several questions. A particular concern was regarding security issues related to IoT. All the panelists agreed that this is indeed a very important aspect that calls for further research.

4. Session 2: "Reliability in Multiple-Access Communication"

4.1. **Dariush Divsalar (NASA JPL): Low Rate LDPC Codes.** The space industry adopted low rate low-density parity-check (LDPC) codes as their standard for space communication, replacing previous convolution codes. To build such codes NASA JPL uses a proto matrix, where rows stand for check nodes and columns represent variable nodes. It can be described by a bipartite graph. The bipartite graph is then lifted by copying the original graph a few times and permuting the nodes. The optimization criteria is minimum distance with an iterative decoding threshold. By brute-force computations for various small cases it turns out that degree 1 variable nodes are important in constructing low rate codes. A construction of ultra-low rate protograph-based raptor-like (PBRL) LDPC is proposed. The PBRL code construction achieves rate compatibility and ease of coding, which constitutes a sufficiently good low rate codes for both short and long blocks.

4.2. **Slawomir Stanczak (Fraunhofer HHI): Harnessing Channel Collisions for Enabling Low-Latency Wireless Random-Access: From Theory to Practice.** Consider the problem of computation over a MAC, where each of the k nodes holds a value x_i and aims to compute the function value $f(x_1, \ldots, x_k)$. In typical solutions, transmissions are orthogonalized; i.e. communication and computations are separated. This can be highly inefficient. Alternatively, one can merge data transmission and function computation into a single step by exploiting the broadcast property of the wireless channel and the fact that every function has a nomographic representation $f(x_1, \ldots, x_k) = \psi\left(\sum \phi(x_i)\right)$. A practical application of such ideas is anomaly detection. Classification for anomaly detection is a common objective in condition monitoring, fault diagnosis and in general industrial environments. The proposed key idea exploits the interference for function computation. The particular problem of f-consensus is discussed, with a focus on the user activity detection scenario. In this setup with one receiver and $K = N + k$ users, only k of the total K users are active. Each user only knows whether he is active or not. The goal is to recover the exact set of active users at the receiver. This can be solved using combinatorial group testing, which is essentially Boolean searching. The solution is, at each step, to choose a random subset of the users, and perform a test to determine if there is at least one active user in the set. If so, no action is performed; otherwise, there are no active users in the group and the subset is removed from the set of potentially active users. The core questions are what is the smallest number of tests necessary to detect all active users, how this number scales with N and k, and how to design protocols for encoding the messages (disjunction coding). It is shown that one can achieve computationally efficient solutions even when communication is held over a noisy MAC.

4.3. **Petar Popovski (Aalborg University): Revising the Problem of Access Protocols in Ultra-Reliable Wireless Communications.** A new mode of communication that comes with 5G is massive and ultra-reliable wireless communication. In this mode, we typically have short packets from a large number of devices. For example, one has to deal with 10kbps packets from 10000 devices, rather than the more traditional 1Mbps packets from 100 devices. Even more so, the system is expected to be able to support 100kbps 99.999% of the time rather than 100Mbps 95% of the time. Current wireless systems are not able to handle these requirements. A protocol for short packets, referred to as 'reliable massive access' is proposed. It integrates authentication and random-access. Information is encoded into the overall access pattern when access re-attempts occur, thus improving connection establishment. The protocol improves on traditional methods in access reliability, access control overhead and access latency. In conclusion, ultra-reliable wireless has the potential to profoundly change systems and devices. Because of the large amount of short packets, it is essential to pay communication theoretic attention to the control information, while carefully designing every step of the communication protocol.

4.4. **Gianluigi Liva (German Aerospace Center): Advances in Random-Access Protocols for Satellite Communications.** Satellite communication has a large population of terminals, often scattered over wide areas and with long propagation delays, especially for geostationary satellites. The medium access problem over such systems, with particular emphasis on latency, is a great challenge. Random-access protocols are, therefore, essential for satellite communication. There are several modern random-access protocols for satellite communications, such as contention resolution diversity slotted ALOHA (CRDSA), coded slotted ALOHA (CSA) and frame-less ALOHA, enhanced spread spectrum ALOHA (E-SSA), time code division multiple access (TCDMA). The first two, which are adequate for the collision channels, are discussed and analyzed. It is argued that collision channel models are too simplistic to describe actual wireless channels. An approach to improve performance is to enrich the collision model starting from the base principles. Reworking the CSA protocols, the contention resolution ALOHA protocol was introduced. It is concluded that while modern random-access protocols attain near-optimum performance on the collision channel, they can be quite suboptimal in actual wireless environments. Nonetheless, adopting an information theoretic perspective on protocol design provides promising directions with potentially high efficiency gains.

4.5. **Discussion.** The organizer asked the speakers to talk about what is the novel aspect in the area of multiple-access communication. Divsalar said that the existing satellite communication technologies are now being used for a wider audience, such as 5G IoT. Popovski said that previously, random-access protocols were designed to operate in one of the two paradigms: (i) where everything is known; (ii) where nothing is known. New systems, however, call for combining these two paradigms into a single framework. Liva said that many of the timely topics discussed in the workshop track back to classic past works, and that there is still an abundance of open questions. Stanczak brought up several points. First, protocols that attain very low latency that meet manufacturing industry requirements are needed. Second, we would like to know more about the mathematics of machine learning so as to avoid treating systems as a blackbox. Furthermore, the mathematical models under consideration themselves should be reconsidered as there is currently a gap between these models and reality. Machine learning was proposed as a way to close the gap. Afterwards, the audience asked several questions.

5. Session 3: "Information Theory I"

5.1. **Dongning Guo (Northwestern): Wireless Node Identification: Fundamental Limits and Practical Design.** Consider the problem of identifying the k out of l active users in a multiple access setting, when both the number of users and the blocklength of their codes simultaneously grow. For the Gaussian multiple access channel, it is shown that the user identification problem is equivalent to a compressed sensing problem. Fundamental limits are established for the user identification problem, giving the blocklength required to identify k out of l users. The capacity of the system is also characterized. In the considered regime, the capacity is not linear in the blocklength, unlike many familiar settings in information theory. Instead, an asymptotic bound on the number of bits per user for a fixed blocklength is derived. Although the identification problem reduces to a compressed sensing problem, the scheme solving the compressed sensing problem is impractical when the number of users is large. Therefore, an identification scheme based on second order Reed Muller codes is presented. This scheme is shown to be useful both with and without channel synchronization.

5.2. **Yury Polayanskiy (MIT): Energy-per-bit in Wireless Multiple-Access.** Consider the the setting where many low-power devices attempt to communicate with a single base station. As the number of users increases, each user must increase their power to support the same rate. How much more power does each user require as a function of the total number of users? To answer this question, discussions need to begin with how to properly define a "random-access code". A per-user error criterion is proposed as a new key ingredient. According to this criterion the error probability of a scheme is the fraction of incorrectly decoded messages to the total number of messages (as opposed to the classic definition where an error occurs when *any* of the multiple messages is decoded incorrectly). Additionally, in the proposed setup the receiver is only required to output the set of messages sent, agnostic to which user sent which message. The fundamental limit of minimum energy per bit as a function of user density (the ratio of number of users to blocklength) is studied. Non-matching achievability and converse bounds are provided. The achievability bound shows that far better schemes than ALOHA exist, though we do not yet know how to construct them.

5.3. **Or Ordentlich (Hebrew University): Low Complexity Schemes for the Random-Access Channel via Structured Codes.** Recently, Polyanskiy et al. proposed a new definition of random-access codes, giving a bound on the achievable minimum energy-per-bit of such schemes. However, a practical way to achieve this bound was not provided. This talk proposes a practical random-access scheme for communication over a Gaussian multiple access channel with a large number of users. The main idea is that good codes for a T user MAC channel make good random-access codes, as they can correct up to T collisions. The channel is divided into sub-blocks (as in ALOHA), and each user chooses to transmit their message in one sub-block, chosen uniformly at random. Because a T user MAC code is employed, if there are fewer than T users transmitting in the same sub-block, the messages can be resolved. A concatenated scheme is constructed using compute-and-forward as the inner code, and a BCH code for the outer code. It is shown that, in most cases, the energy-per-bit achieved by this scheme is far smaller than that achieved by ALOHA.

5.4. **Krishna Narayanan (Texas A&M): User-Independent Coding Schemes for Uncoordinated Unsourced Massive Multiple Access.** Recently, Polyanskiy and Ordentlich proposed the first practical random-access coding scheme for communication over the Gaussian MAC. This talk proposes a new scheme, which substantially improves on the previous one. The proposed paradigm is composed of four main ingredients. First, the transmission period is partitioned into sub-blocks, thereby instituting a slotted framework. Second, each message (data) is split into two parts, one of which chooses an interleaver for an LDPC type code. This part of the message is encoded by spreading codewords that are designed to be decoded by a compressed sensing type decoder. Third, another LDPC type code is used to encode the other part of the message. Decoding this part involves a joint message passing decoding algorithm designed for the T-user binary input real adder channel. Finally, each user repeats its codeword in multiple sub-blocks, with the transmission pattern being a deterministic function of message content and independent of the identity of the user. When decoded using a successive interference canceller type decoder, this coding scheme is shown to provide excellent performance.

5.5. **Discussion.** The discussion centered around the question: What pieces of this random-access coding theory are most useful practically? Guo emphasized the need for theory in the maximally uncoordinated setting, where transmitters do not possess knowledge of each other: each simply wakes up, sends its message, then goes back to sleep. All speakers in this session considered the Gaussian MAC, but the fading MAC is much more practical. Polyanskiy discussed how we intuitively understand the properties of good channel codes. For example, their output distributions should look like the capacity achieving output distribution, they have sufficient independence and symmetry properties, etc. However, for random-access codes, we have no intuition what a good scheme might look like. Hopefully, research in this direction will provide a better idea of the properties of good random-access schemes. Narayanan had a short discussion on the interest in software-defined implementations of random-access schemes, mentioning how FPGAs are now incredibly practical. Being cheap, requiring low-power, and being highly flexible, FPGAs serve as good candidates for such implementations. He stressed that schemes like the ones discussed during the session can be implemented using FPGA.

6. Session 4: "Low Latency Communication"

6.1. **Vincent Tan (National University of Singapore): An Exact Spectrum Formula for the Maximum Size of Finite Length Block Codes.** Inspired by the random coding proof of the Gilbert-Varshamov bound, this work gives an exact variational formula for the maximum size of a block code with minimum distance d. The formula is given in terms of a general distance function $\mu(x, x')$, where x and x' are arbitrary points in the space of codewords. Using the distance μ, one then defines the *distance spectrum* as $F_X(d) \triangleq \mathbb{P}\big(\mu(X, X') > d\big)$, where X and X' are independent and identically distributed according to P_X. The bound asserts that the maximum number of codewords at distance d is given by $\sup_{P_X} F_X(d)^{-1}$. A family of lower bounds for the maximum code size is derived from this bound, including the Gilbert-Varshamov bound. Additionally, the formula yields bounds on the second-order terms in the asymptotic expansion of the optimum finite length rate for block codes with a fixed normalized minimum distance.

6.2. **Albert Guillén i Fàbregas (ICREA and Universitat Pompeu Fabra): The Meta-Converse and Perfect Codes.** In a previous work of Fàbregas et al. the meta-converse bound proposed by Polyanskiy et al. was shown to be tight. The meta-converse gives a bound on the error probability in a Neyman-Pearson hypothesis test between two distribution: P_{XY} and $P_X \times Q_Y$, where P_X is a distribution on the input space, $P_{Y|X}$ is given by the channel, and Q_Y is an auxiliary distribution. In this work, it is shown that for symmetric channels, the error probability of the certain codes coincides with the error in a binary hypothesis test with certain parameters. Furthermore, the points in which they coincide are exactly the points where either a perfect or a quasi-perfect code exists. The main idea of the proof is to choose the auxiliary distribution Q_Y as the expurgated exponent, rather than the capacity achieving output distribution, which is used in most application of the meta-converse bound.

6.3. **Gonzalo Vazquez Vilar (Universidad Carlos III de Madrid): Computing the Meta-Converse Bound: A Saddlepoint Approximation of Hypothesis Testing.** Previously, Polyanskiy et al. proposed the meta-converse bound that gives an upper bound on the number messages in a channel code for a fixed blocklength and error probability, by viewing the problem as a binary hypothesis testing problem. Although useful theoretically, computing the bound can be quite difficult. This work gives a saddlepoint approximation to the error probability of binary hypothesis testing, which is then used to efficiently compute the meta-converse bound for moderate blocklengths. This gives a simple and elegant solution to an often quite difficult computational task. The approximation is compared with brute force computations of the bound, and is shown to be quite accurate in many cases.

6.4. **Tobias Koch (Universidad Carlos III de Madrid): Normal Approximations for Fading Channels.** This work characterizes the dispersion of the non-coherent, single-antenna, Rayleigh block fading channel in the high SNR regime. Unlike some other fading channel models where the dispersion is known, in the non-coherent setup, even the capacity is unknown. Consequently, there is little hope for getting a general expression for the dispersion. However, at high SNR, it is known that unitary space-time modulation gives an optimal input distribution, and that the capacity, in the limit as the SNR goes to infinity, can be determined. This work extends Marzetta and Hockwald's results by giving a closed form expression for the second order term in the asymptotic expansion of achievable rates. The normal approximation is then compared to non-asymptotic bounds, and shown to be faily accurate. This expression provides a new benchmark for evaluating codes over non-coherent fading channels. Furthermore, the result given insight into the tradeoff between using resources for channel estimation versus sending more information bits.

6.5. **Discussion.** The topic of discussion centered around the question "what future uses can finite-blocklength analysis provide in information theory?". The speakers agreed that, while progress in the area can be tough, with the increasing importance of low-latency communications, these results are as important as ever. Albert Fàbegras mentioned the new use of the expurgated exponent that arose in his work. Tobias Koch added that the wireless communication problem is now of incredible importance. Finite-blocklength results give the most honest analyses of achievable rates to date, which are relevant for many practical aspects of wireless communications. The following question also arose: finite-blocklength bounds are rather well understood for fixed blocklength codes, but how about convolutional codes? Specifically, how long does re-synchronization take, once one loses track of the state in a convolutional code? In this context, it is known that the closer the transmission rate is to capacity, the longer re-synchronization takes. Finite-blocklength methods were brought up as a possible way to get more explicit bounds.

7. Session 5: "Information Theory II"

7.1. **Victoria Kostina (CalTech): Random-Access Channels: Theory and Coding.** A random multiple access channel, with a potentially unbounded number of users, is discussed. This communication scenario gives insight into massive machine type communication. The setup differs from the traditional MAC in the sense that the subset of active users is random and unknown. Assuming a permutation invariant channel, achievability and converse bounds are derived via rateless coding and techniques from finite-blocklength information theory, respectively. It is first shown that an unbounded number of collisions can be resolved without apriori knowledge of the user-activity pattern and without any loss compared to the full channel state information (CSI) case. The second part of the talk focuses on coded random-access design. A two-layered coded random-access scheme for erasure channels that jointly accounts for collisions and erasures is discussed. The construction employs random LDPC codes and iterative decoding. The inner layer is designed to resolve collisions, while the outer layer deals with erasures. Simulation demonstrate that the proposed two-layered scheme outperforms separated schemes where collision resolution and erasure correction are dealt with independently.

7.2. **Uri Erez (Tel-Aviv University): Probability that the Symmetric Rate Capacity and the Sum-Capacity of a MAC Coincide (or nearly so): Non-Asymptotic Bounds.** The single-input-single-output Rayleigh slow fading MAC is considered. A protocol, where all users transmit at a rate just below the equal-rate capacity (per user) of the channel, is analyzed. The underlying assumption is that this rate is dictated to the users by the base station. Tight bounds on the distribution of the rate attained by the transmission scheme are established. In particular, these bounds characterize the probability that the dominant face of the MAC capacity region contains an equal-rate point, i.e., that the scheme strictly attains the sum-capacity of the channel. The analysis provides a non-asymptotic counterpart to the diversity-multiplexing trade-off of the MAC channel. Furthermore, a practical schemes based on integer-forcing and space-time precoding is discussed. It is shown that for the 2-user case, the outage probability of this scheme behaves close to the optimal. This near-optimality confirms the theoretic-based predictions.

7.3. **Ayfer Özgür (Stanford University): Sparse Combinatorial Group Testing Codes for Low-Energy Massive Multiple Access.** Random-access for IoT is studied. The proposed model involves a massive number of sporadically active devices, each observing a one-way communication channel to an access point. The communication is to be held under certain complexity and energy constraints imposed on the transmitters. Furthermore, node identity is assumed not to be mandatory, but re-transmission are forbidden. Specifically, assume there are d out of N ($1 << d << N$) active users at a time. The talk focuses on identifying the set of active users based on group testing. A naive approach is to employ $t = N$ orthogonal signatures. While being energy efficient, this solution is highly undesirable from spectral considerations. As an alternative approach, group testing is proposed as a method to boost spectral efficiency. The trade-off between spectral efficiency and energy efficiency is discusses. Tight upper and lower bounds on the spectral efficiency for given energy are provided. Thinking of the energy of a protocol as weight of its codewords, and of the spectral efficiency as the number of tests needed, the following is shown. For weight $w = ld + 1$ (where $l \in \mathbb{Z}$ and $ld + 1 \leq N^{1/(l+1)}$), the number of tests is $\Omega(d^{2/(l+1)} N^{1/(l+1)}) \leq t \leq O(dN^{1/(l+1)})$, with explicit constants for the upper bound.

7.4. **Discussion.** The discussion post the talk was very intriguing. A central question raised by the audience during the discussion is whether rateless coding is feasible to the industry. It was pointed out that rateless coding is used to account for non-ergodicity of the channels. Further, it was mentioned that there is a need to invest in feedback because of the half-duplicity of transmitters. Afterwards, the matter of communicating at subzero dB SNRs was brought up. It was widely argued that this is a highly challenging goal in practice because of issues like synchronization (of frequencies, times, etc.).

8. Session 7: "Massive Multiple-Access"

8.1. **Giuseppe Caire (Berlin Technical University): Fog Massive MIMO with on-the-fly Pilot Contamination Control.** Massive MIMO presents several attractive features for very-low-latency and high-reliability random-access communications. In particular, due to the large number of antennas, the wireless fading channel behaves almost deterministically, such that complicated adaptive rate schemes are not needed. Nevertheless, in a multi-cell dense deployment, frequent handovers, with per-cell pilot re-assignment, may still incur significant protocol overhead and latency. The talk presents a novel "Fog" massive MIMO architecture, where users seamlessly and implicitly associate to the most convenient multi-antenna Remote Radio Head (RRH) in a completely autonomous manner. Each user is associated with a unique uplink pilot sequence, and pilot contamination is mitigated by a novel "on-the-fly" mechanism. This scheme preserves the advantages of Cloud-RAN processing (in particular, no association between users and RRHs needs to be explicitly negotiated), without incurring in the latency of fully joint processing of the RRH signals at a common cloud center. Furthermore, the spectral efficiency of the resulting scheme is analyzed via stochastic geometry, using some recent results on unique coverage in Boolean models, which were developed specifically to analyze the proposed system. This Fog massive MIMO system is compared with a baseline massive MIMO cellular system and with the recently proposed cell-free architecture. The superiority of the proposed scheme is shown through analysis and simulation.

8.2. **Giuseppe Durisi (Chalmers University): Short Packets over Fading Channels: Pilot-Assisted or Non-Coherent Transmissions.** In most real-world wireless communication systems operating over fading channels, pilot symbols are multiplexed with the data symbols to facilitate channel estimation at the receiver. In this context, the talk addressed the following questions: how optimal is this strategy when packets are short, and shall one avoid learning the channel explicitly, and instead use non-coherent transmission schemes. By using finite-blocklength achievability and converse bounds, a tight normal approximation for the maximum coding rate at high SNR is obtained for non-coherent transmission schemes. RCUs bounds with unitary space-time modulation (USTM) codes are used to prove achievability results. A meta-converse theorem was the main tool for proving the converse. Using appropriate decoding metrics (ML-decoding) for non-coherent and pilot-assisted nearest neighbor (PAT-NN) decoding, numerical evaluations of achievability bounds show that non-coherent decoding is better in terms of spectral efficiency than regular PAT-NN. Constricting practical non-coherent schemes that achieve the theoretical bounds was posed as a goal for future research.

8.3. **Wei Yu (Univerity of Toronto): Massive Device Connectivity with Massive MIMO.** The talk considers a massive device communications scenario in which a large number of devices need to connect to a base-station in the uplink. User traffic is assumed to be sporadic so that, at any given coherence time, only a subset of users are active. For such a system, user activity detection and channel estimation are key issues. A two-phase framework is proposed for overcoming these issues. The first phase exploits compressed sensing techniques to identify the devices and their channels. Data transmission takes place in the second phase. Approximate message passing (AMP) is proposed for device identification. It is shown that the state evolution can be used to analytically characterize the missed-detection and false-alarm probabilities in AMP. This talk further considers the massive connectivity problem in the massive MIMO regime. It is analytically shown that massive MIMO can significantly enhance user activity detection, but the non-orthogonality of pilot sequences can nevertheless introduce significant channel estimation errors, hence limiting the overall rate. This effect is quantified and the optimal pilot length for massive uncoordinated device access is characterized.

8.4. **Wei Yang (Qualcomm): Privacy Amplification for Low-Latency Physical-Layer Security.** Physical-layer security is an alternative approach to cryptography for provably secure communication. It harnesses the intrinsic randomness of a noisy communication channel to encrypt secret information. The talk considers the highest attainable secret communication rate over a wiretap channel at a given blocklength. New achievability and converse bounds are derived, which are uniformly tighter than existing bounds. These new results give rise to the tightest available bounds on the second-order coding rate for discrete memoryless and Gaussian wiretap channels. The exact second order coding rate is established for semi-deterministic wiretap channels, which characterizes the optimal tradeoff between reliability and secrecy in the finite-blocklength regime. The underlying achievability scheme exploits the notion of privacy amplification. A stronger version of the classic privacy amplification result is provided, which opens the door for achieving secrecy under the strict semantic-security metric.

8.5. **Discussion.** This discussion focused on the question of why physical layer security is not being deployed in real-world systems. A remark from the audience emphasized that it has been implement in Bosch, Germany. Several issues of physical-layer security we brought up. The main concern was model mismatch and its potentially huge cost in terms of secret information leakage. Although a beautiful theory, it remained to be seen if can be implemented in real world applications on a large scale.