LUND UNIVERSITY

Utilizing Massive MIMO for the Tactile Internet: Advantages and Trade-offs

Tärneberg, William; Karaca, Mehmet; Robertsson, Anders; Tufvesson, Fredrik; Kihl, Maria

Utilizing Massive MIMO for the Tactile Internet: Advantages and Trade-offs

William Tärneberg[1], Mehmet Karaca[1], Anders Robertsson[2], Fredrik Tufvesson[1], Maria Kihl[1]

[1]Department of Electrical and Information Technology, Lund University, Sweden

[2]Department of Automatic Control, Lund University, Sweden

Abstract—Controlling robots in real-time over a wireless interface present fundamental challenges for forthcoming fifth generation wireless networks. Mission critical real-time applications such as telesurgery over the tactile Internet require a communication link that is both ultra-reliable and low-latency, and that simultaneously serving multiple devices and applications. Wireless performance requirements for these applications surpass the capabilities of current wireless cellular standards. The prevailing ambitions for the fifth generation wireless specifications go beyond higher throughput and embrace the wireless performance demands of mission critical real-time applications in robotics and the Internet of Things. To accommodate these demands, changes have to be made across all layers of the wireless infrastructure. The fifth generation wireless standards are far from finalized but massive Multiple-Input Multiple-Output has surfaced as a strong radio access technology candidate and has great potential to cope with all these stringent requirements. In this paper, we investigate how Ultra-Reliable and Low-Latency Communication with massive MIMO can be achieved for bilateral teleoperation, an integral part of the tactile Internet. We conclude through simulation what the performance bounds are for massive MIMO and thus how to configure such a system for near deterministic latency and what the inherit trade-offs are.

Index Terms—Massive MIMO, URLLC, 5G, Robotics, Latency, Reliability, Tactile Internet

I. INTRODUCTION

The recent uptake in factory automation and robotization is just the beginning of a wider adaptation of Internet of Things (IoT), machine learning, big data, and cloud technologies to automate a large set of the professions that have come to characterise the twentieth century. This change is commonly referred to as the fourth industrial revolution [1].

This technological revolution is still in its infancy. There are great technological challenges yet to be addressed for the revolution to encompass cognitive and motorically intense professions. The fourth industrial revolution does not only imply that most tasks and professions will be automated at the rate of which technology matures. In fact, the ambition is not to indiscriminately eliminate human capital. Technology will instead be used to make use of the human cognitive advantage wherever it might be needed. This notion includes for example precise teleoperation, such as telesurgery. One can trivially imagine a future where machine decisions seamlessly inter-operate and complement physical human actions. Haptic feedback is a key enabling technology in this pursuit [2]. It has consequently been argued that the Internet as we know it today will shift from content delivery to labour delivery [3]. This paradigm shift is enabled by the tactile Internet. A tactile
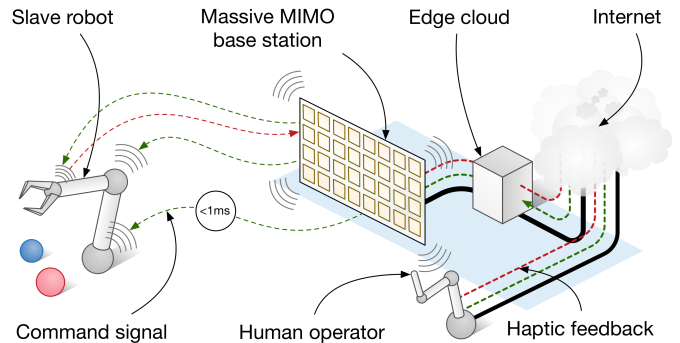

Fig. 1: The tactile Internet and massive MIMO.

Internet can include of amongst other things, a large number of connected tactile surfaces and robotic limbs, accessed remotely at high precision, see Figure 1.

Enabling the tactile Internet and haptic feedback for the fourth industrial revolution will require a communication intensive distributed system where a large set of resources are shared and decisions are made in distributed fashion in a network with many wireless links. It is a well know fact that the stability of tactile control loops are particularly sensitive to jitter. The tactile Internet will therefore require near-deterministic single-digit millisecond latencies. Additionally, at the rate at which tactile feedback loops operate, there is very little margin for error [4]. Traditional mechanisms that provide reliability such as Hybrid Automatic Repeat Query (HARQ) and Automatic Repeat Query (ARQ) found in Long Term Evolution (LTE) might not be feasible at such a low over-the-air latency. Unlike audio and video content, tactile feedback can currently not be scalably compressed in a lossy fashion. Tactile feedback, contrary to audio visual content, can therefore not trivially be adapted to prevailing wireless throughput capacity. Furthermore, the surfaces and robotic limbs that are part of tomorrow's tactile Internet do not operate in just one tactile dimension. They each rely on multiple multi-modal data inputs and outputs across multiple devices [5], requiring each device to simultaneously communicate at the same level of reliably and latency. Providing wireless Ultra-Reliable and Low-Latency Communication (URLLC) to the devices that constitute these services has been proven to be non-trivial.

Existing mobile specifications, as they are deployed, are unable to provide URLLC cost effectively at scale [6]. They also

lack the ability to deliver reliable low latency communication to multiple users simultaneously. Consequently, the proposed Fifth Generation Wireless Specifications (5G) are focusing on addressing the challenges of scale, device heterogeneity, and mission criticallity. In other words, the aim is to provide control communication for closing the global tactile control loop and not just deliver e.g. audio visual content.

Massive MIMO is a Multi-User Multiple Inputs Multiple Outputs (MU-MIMO) Radio Access Technology (RAT) where N User Equipments (UEs) share the same time-frequeny resource block. This allows Massive MIMO to simultaneously serve N at similar reliability and latency level, vital for realizing the tactile Internet. Additionally, Massive MIMO offers a 20-fold increase in spectrum efficiency over the current LTE MU-MIMO specifications [7]. Besides spectrum efficiency, the technology also offers reliable communication through the so called channel hardening effect and as a consequence of that, low latency communication. A low Bit Error Rate (BER) is essential for achieving ultra low over-the-air latency, beyond LTE. At very low over-the-air latencies, HARQ in LTE will introduce long tailed latencies and unwanted overhead in a system of thousands of devices. Moreover, the simultaneous support of multiple devices and applications with a desired reliability is one of the most attractive advantages of Massive MIMO, which is also essential for dense networks requiring URLLC. With the above properties, Massive MIMO is expected to be able to deliver URLLC for 5G networks powering the tactile Internet.

Work is beginning to emerge in the literature on how to generally realize URLLC and the tactile Internet using 5G. [4] taxomonizes the tactile Internet design challenges facing 5G. The work primarily addresses system design challenges and presents a reasonable foundation of performance requirements and limitations. There are also proposed system architectures for achieving URLLC [8]. The edge cloud will arguably also play an important role and undergo significant changes in realizing the services and the infrastructure of the tactile Internet [9]. Although the entire wireless system infrastructure needs to operate at a low latency, the over-the-air latency in the physical layer is fundamental in achieving URLLC. Consequently, there are for example proposed physical layer specifications for URLLC wireless networks [10]. However, none of these works have uniformly looked at the physical layer for closing the tactile feedback loop with the gains Massive MIMO can deliver. To the best of our knowledge, this has not previously been addressed.

In this work we investigate how to realize URLLC communication for the tactile Internet using Massive MIMO. We adopt the requirements of bilateral teleoperation, an application of haptic feedback, as a baseline, and investigate how its reliability and latency requirements can be fulfilled with Massive MIMO. Furthermore, we contribute with a performance analysis of Massive MIMO under URLLC-conditions which is used to formulate upper performance bounds for such a system. Additionally, the analysis allows us to constructively discuss trade-offs and specifications for an URLLC Massive MIMO system. The results in the analysis are contrasted with what is attainable with the current LTE specifications.

II. BILATERAL TELEOPERATION

In this section we take a closer look at the communication requirements for bilateral teleoperation which is an application of tactile feedback and one of the most promising applications of the tactile Internet. Generally speaking, closing force-feedback control loops for mechanical manipulators with interaction in stiff, as opposed to elastic, environments is challenging and becomes notoriously difficult from a stability point-of-view when uncertain delays are introduced in the loop. Also, in scenarios with less stiff interaction forces, transmission delays and communication jitter deteriorate performance and robustness. The concept of haptic teleoperation with bilateral force-velocity reflection between a "master" (human operator) and a so-called "slave" (slave robot) provides a mean of transparency and experienced interaction of contact forces and end-effector motions for the operator, see Figure 1.

Bilateral teleoperation has been studied extensively. The method of 'wave variables', introduced in [2], has successfully been extended and applied in remote-controlled mining, dental and medical surgery, and even for space applications with significant delays. In such systems, fundamental limitations will impose a trade-off between stability and quality in terms of experienced transparency of the bilateral teleoperation depending on the properties of the communication channel.

For remote control, a good complement to haptic feedback is streaming video from a remote site, which allows the operator on the master side to visually inspect the interaction. However, for a consistent user experience, it is vital that the different feedback channels, with possibly significantly varying amount of data, are synchronized without unnecessary delays and jitter.

In a typical system, each joint is controlled separately over wireless links. Robotic joints typically update at $250Hz$ [11]. A latency of at most $4ms$ is therefore desired, preferably $1 - 2ms$ to accommodate jitter. Furthermore, the jitter is fundamentally addressed with a high wireless link reliability, BER $< 1e-5$.

For a good survey of haptic bilateral teleoperation case, see [12]. In [13] the quality-latency trade-off for bilateral haptic teleoperation is investigated for different wireless standards. In [14], the effect of network quality on bilateral teleoperation was investigated. In that work, the performance metric was how accurate the slave system follows the command of the master system as well as how transparent the environment is to the operator. In that work, it was shown that packet loss (i.e., BER) affects the signal oscillations, while the latency in the network causes the steady-state tracking error increase. Based on this result, one can conclude that the reliability and latency issues must be addressed together. Next, we show that how massive MIMO can help us to achieve a desired reliability and latency with some expected trade-offs.

III. RELIABILITY

In this section, we evaluate how massive MIMO can be dimensioned to achieve the desired reliability level, BER, pre-

sented in Section II. The analysis was done through simulation using MATLAB's communications toolbox executed on Lund Unversity's cluster, LUNARC.

We begin by detailing the configuration of the simulated system. With the reliability challenges detailed in Section II coupled with the URLLC ambitions in [15], we are targeting a BER of $1e-5$. Because the UEs are relatively computationally underpowered and do not require a particularly high throughput, as discussed in [16], we use the Quadrature Phase Shift Keying (QPSK) modulation scheme.

In this work, we adopt the Independent and Identically Distributed random variables (i.i.d) Rayleigh fading channel model. Using i.i.d Rayleigh fading channels will form a reasonable upper performance bound, as a best case scenario. In reality, due to correlation between users, the i.i.d assumption will not hold and we can expect that the minimum requirements for URLLC haptic feedback will be more stringent. For fading channels, in order to improve the reliability of the channel it is common to use channel coding such as convolutional or turbo coding depending on the application. Nevertheless, using the coding incurs additional cost in terms of receiver complexity and decoding delay. Note that convolutional coding with rate 1/2 and constraint length 7 is used in this study since it yields a lower delay compared to Turbo coding and Low-Density Parity-Check (LDPC) codes. In terms of diversity combining, we investigate the use of both Maximum-Ration Combining (MR) and Zero-Forcing (ZF), which are both linear precoding schemes. Although MR is arguably not entirely beneficial in massive MIMO, we include MR as reference to a low-complexity mechanism. MR therefor acts as a lower performance bound and will effectively contract the channel properties with ZF.

As for our targeted haptic feedback system, in the scenario evaluated in [13], a 6-Degrees of Freedom (DoF) robot was considered. The authors of [4] propose segments of 48 data bits for a 3-DoF setup. We therefore adopt a packet size of 100 bits and that the traffic flow is near-constant when UE is operational. To detect errors, we also assume that Cyclic Redundancy Check (CRC) is applied. These parameters can also be considered to be true for other robotics systems.

A. The role of massive MIMO

Massive MIMO provides the means to significantly reduce the BER over existing LTE MU-MIMO specifications in a relatively straightforward manner [17]. Because of the focusing effect in massive MIMO [7], more UEs can be served with a lower BER at a lower Signal-to-Interference-plus-Noise Ratio (SNR) than in current deployed wireless specifications. The UEs of the tactile Internet operate over a wide range of power requirements. Some UEs are for example battery powered with a targeted lifespan expressed in years. Here massive MIMO offers an advantage over conventional techniques as the UEs can be made relatively simple as much of the complexity can be moved to the Radio Base Station (RBS). A predictable power consumption is an integral part of the reliability of a UE. A low transmission power typically results in a low SNR.

Modulation	QPSK
Channel model	i.i.d Rayleigh fading
Precoding	MR, ZF
Channel coding	Convolutional code with rate 1/2, constraint length 7
SNR	$(-14, 10)$ dB
BER target	$1e-5$
Packet size	100 bits

TABLE I: Physical layer parameters.

To therefore sweep across SNR levels from as low as -14 dB to 10 dB. The targeted massive MIMO system's parameters are summarized in Table I.

B. Performance of massive MIMO

A fundamental differentiating design parameter in a massive MIMO system is the number of RBS antennas, M. A high M/N ratio yields a lower BER or allows the system to operate with a lower SNR. In our massive MIMO system, the UEs are assumed to operate with one antenna. For the sake of generality, we refer to the robotic joints and surfaces in this scenario as UEs.

We proceed by investigating the relationship between the number of antennas and the system's reliability by finding the minimum number of antennas M required to achieve a BER of $1e-5$ for a given N simultaneously served UEs at a certain SNR level. Here, SNR is defined as the input SNR where it is defined as the ratio of transmit and noise power. Since the N UEs in the system are low powered and we do not want to add additional delay as a results of computational complexity and since massive MIMO achieves an inherently low BER, we initially proceed without any channel coding. Using the scenario in [13] as our reference. The ABB robotic arms in that scenario have 3-6 Degrees Of Freedom (DOF) [11].

Figure 2 reveals the difference in performance between MR and ZF for i.i.d channels. The graphs can be read as either the minimum number of antennas required to achieve a certain BER or the degradation of BER as a function of the number of UEs and SNR. With either pre-coding schemes, there is no significant degradation in BER until SNR$= 0$. From this point, ZF's performance degrades at a relatively higher rate with than MR but still performs strictly better than MR. ZF outperforms MR on average a factor of 2 at high SNRs values and with a factor of almost 10 at low SNRs. ZF's gain over MR increases linearly with SNR. However, ZF's gain over MR increases quadratically with the number of UEs, N.

In Figure 3, channel coding is used to improve reliability and reveal its relative gain. Again, ZF outperforms MR on average of a factor of 0.85 across all configurations. Contrasting Figures 2 and 3 shows that adding channel coding to ZF provides on average a factor 0.5 improvement at low SNR levels and practically no improvement for high SNR level when it comes to the minimum number of required antennas. MR on the other hand, sees a more than four-fold increase in performance.

As suggested by the results in Figures 2 and 3, with ZF, as long as M is sufficiently high (e.g, $M = 100$) we arguably stand to gain very little from channel coding. As seen in
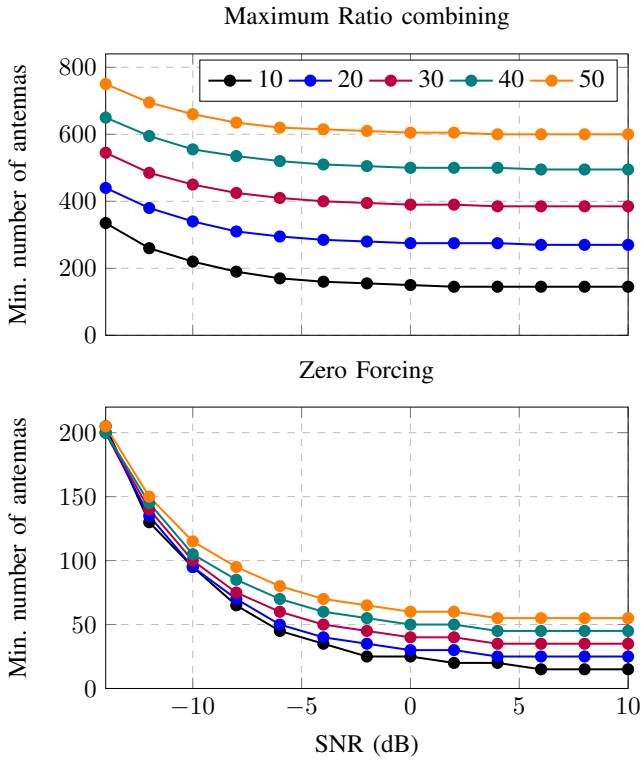
Fig. 2: Number of antennas required to simultaneously serve N UEs at a specific SNR with a BER of $1e-5$. Without channel coding.
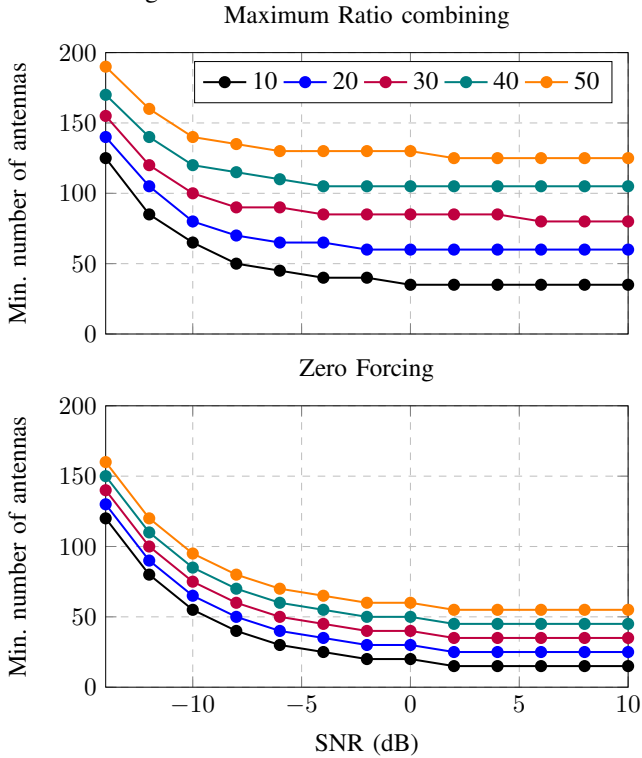

Fig. 3: Number of antennas required to simultaneously serve N UEs at a specific SNR with a BER of $1e-5$. With channel coding.

Figures 2 and 3 the BER is satisfied at SNR= 0, at which point the minimum number of antennas does not decrease with increased SNR. At this point, there is no gain in BER when using channel coding. It is also evident from Figures 2 and 3 that channel coding is not contributing when the number of UE high and SNR < 0. This is illustrated by the convergence of the minimum number of required antennas for each number of users as SNR is decreased. In fact, the the gain diminished on average quadratically with diminishing SNR. Channel coding gain is therefore only present for $K \leq 30$ and low SNR conditions. This can be attributed to the fact that the channel coding gain is very low when the intra-UE interference is high, e.g. when $K > 30$.

Furthermore, increasing the number of UEs N will require a factor 1 increase in the number of antennas on the RBS side. However, the number of simultaneously served UEs, N, does not only depend on M, and when the latency requirements are taken into account M and N should be carefully decided as we show next.

IV. LATENCY

In this section, we evaluate how ultra low latency can be achieved with multi-user massive MIMO. The starting point is a round-trip latency of $1ms$ specified in [4] and Section II.

The latency contributors can be decompose into three components: i-) processing/coding at the transmitter; ii-) over-the air transmission; iii-) the processing/decoding at the receiver. The first and third components are strongly related to the hardware, software, and coding scheme used at the transmitter and receiver UEs [17]. The second component is however strongly related to the frame structure used in the communication link. In the current LTE specifications, over-the air latency (i.e., transmission time interval (Transmission Time Interval (TTI)), is in part determined by to the symbol duration. This is already at $1ms$, which makes it impossible to achieve the desired low-latency. Therefore, TTI duration should be reduced. One solution is to reduce the symbol duration. More specifically, the current Orthogonal Frequency-Division Multiplexing (OFDM) symbol duration is too long. One way to do decrease it is to increase the sub-carrier spacing which will reduce OFDM symbol duration and consequently TTI, as according to the following relationship,

$$\text{TTI} = \tau \left(\frac{1}{\Delta f} + T_g + T_p \right), \qquad (1)$$

where τ, Δf, T_g and T_p are the number of OFDM symbols in one TTI, the sub-carrier spacing (in kHz), cyclic prefix duration (in μs) and some processing delay which may depend on the software and hardware on the UE, respectively. Furthermore, one OFDM duration (e.g., symbol duration) is expressed as,

$$T_s = \frac{1}{\Delta f} + T_g \qquad (2)$$

For example, in the LTE standard, $\Delta f = 15kHz$ and $T_g = 4.76\mu$s. Consequently, one OFDM duration in LTE is,

$T_s = T_u + T_g = 71.4\mu s$ where $T_u = \frac{1}{\Delta f} = 66.7$ μs and there are 14 OFDM symbols in one TTI. Clearly, in order to reduce the TTI duration, Δf should be reduced since the other factors T_g and T_p are not controllable and usually depend on the channel characteristics and the type of the UE, respectively. As it can be seen from (1) that another alternative to reduce TTI is to use fewer OFDM symbols at each TTI (i.e., reduce τ) without needing to change Δf. One drawback of using a reduced number of OFDM symbols is that the resulting scheduling cost can increase.

In order for massive MIMO to operate efficiently, the RBS neededs to collect Channel State Information (CSI), which is achieved through the pilots symbols transmitted from each UEs to the RBS. In an OFDM based system, each pilot symbol correspond to a sub-carrier in one OFDM symbol. That is to say, some number of OFDM symbols should be dedicated for CSI. If β number of OFDM symbols out of τ symbols are used for pilots, then the number of UEs that can be simultaneously served by a massive MIMO RBS is,

$$K = \beta S = \beta \left(\frac{1}{\Delta f T_g} \right) \tag{3}$$

where $\beta < \tau$ and S is frequency smoothness as defined in [18]. In other words, S is the coherence bandwidth of the channel in terms of number of sub-carriers and over S sub-carriers the channel can be seen as constant and a reliable communication for the CSI transmission can be realized. Note that if β symbols are used for pilots then the actual data communication will be $(\tau - \beta)$ symbols, which are used by all the scheduled UEs at the same time. Here, the system efficiently is defined as $\xi = (1 - \beta/\tau)$. Clearly, with higher β values we can support more UEs. However, the amount of data that can be received or transmitted will be reduced.

Combining Equations (1) to (3) highlights an interesting trade-off. A higher Δf yields a lower TTI but also lowers the number of UEs K that can be simultaneously served, when $N \geq K$. Generally speaking, we have strict latency requirement, i.e., TTI_{thr} and proceeding to maximize K yields the following optimization problem,

$$\max \quad K \tag{4}$$
$$s.t. \quad \text{TTI} \leq \text{TTI}_{\text{thr}} \tag{5}$$

The solution is straightforward and given by,

$$\Delta f^* = \frac{1}{\frac{\text{TTI}_{thr}}{\tau} - T_g - T_p} \tag{6}$$

and

$$K^* = \beta \left(\frac{\text{TTI}_{thr}}{\tau T_g} - \frac{T_p}{T_g} - 1 \right) \tag{7}$$

A. System view

Figure 4 depicts the optimal sub-carrier spacing and the number of simultaneously supported UEs with varying T_g values by using Equations (6) and (7). As an example, when
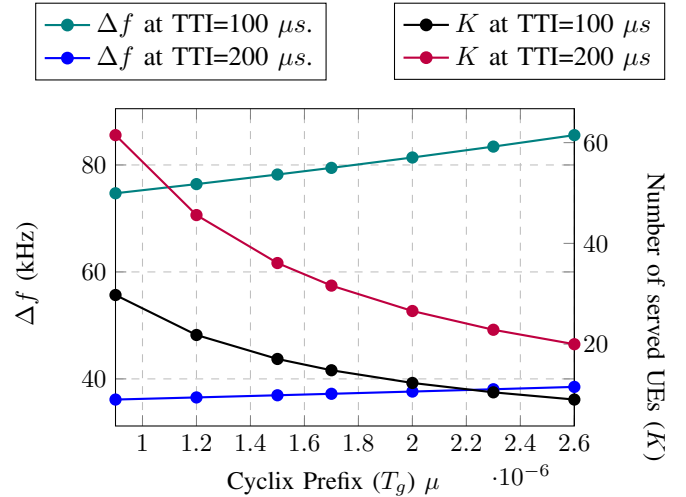

Fig. 4: K and Δf with various T_g

$\beta = 2$ and $\tau = 7$ (i.e., ξ=71%) and $\text{TTI}_{\text{thr}} = \{100, 200\}\mu s$. Less strict TTI requirements can create an opportunity to support more UEs simultaneously. For example, in the targeted scenario we require the TTI duration to be $100\mu s$, i.e., $\text{TTI}_{\text{thr}} = 100\mu s$ which is short enough to provide 1 ms end-to-end delay including encoding/decoding[1] and processing delays [4]. Thus, each OFDM symbol needs to be $14.28\mu s$, $T_g = 0.9\mu s$, and $T_u = 13.38\mu s$ as a consequence of Equation (6) $\Delta f^* = 75kHz$. When $\beta = 2$, by using Equation (7) the system can serve $K^* = 28$ UEs simultaneously.

Figure 4 can help us design the required frame structure by illustrating the primary system trade-offs. For example, after measuring the channel characteristics (e.g., the delay spread) and deciding the length of the cyclic prefix one can use the results in Figure 4 to decide the sub-carrier spacing depending on the latency requirements given in Section II. Then, the number of the supported UEs, K, can trivially be determined by using Equation (7). In order to support K UEs, the results in Section III can be utilized to determine how many antennas are needed at the base station. The design can also be realized in reverse, such that for a given number of RBS antennas, M, one can determine the number of simultaneously supported UEs, K, based on a given latency requirement by using Equation (7). If K is not supported with the given number of antennas, a lower K can be considered. However, a lower K will require an increased sub-carrier spacing to reduce the TTI further. The increased sub-carrier spacing comes at the expense of increased bandwidth. If there is a bandwidth shortage then there may not be sufficient number of pilot symbols, and consequently K will decrease.

B. Latency and reliability

In the real-time scenario presented in Section II, reliability and latency are inextricably linked. Irregardless of the over-the-air latency, a high BER will result in packet losses. With a BER of 0, the over-the-air latency would be deterministic.

---

[1]The encoding/decoding delay should be in the same order of one OFDM symbol duration in order to avoid any memory issue.

However, we are able to achieve BER of $1e-5$. Assuming that some degree of error detection mechanism is applied, retransmission mechanisms such as ARQ would contribute with jitter. Packet losses can either be remedied in the controller by compensating for the uncertainty permanently lost information introduces or through ARQ. However, when the TTI is reduced to $100\mu s$, the relative latency of retransmission increases dramatically. ARQ might therefore not be applicable in this scenario.

C. Precoding design

Lastly, we would like to discuss the impact of precoding design. This is an essential part of a massive MIMO system in terms of the end-to-end latency. It has been shown in [19] that in a typical massive MIMO system with 128 antennas at the base station and 8 UEs transmitting uplink data, the precoding delay due to the required matrix inversion and multiplications for ZF can be up to $150\mu s$. Since ZF has a complexity proportional to K^2M, the precoding latency will dramatically increase as K and M increase. This amount of latency is significant and a challenge for tactile Internet applications. One possible solution to reduce the precoding latency is to use more hardware resources, which however will increase the equipment cost.

V. CONCLUSIONS

In this paper, we have investigated the potential gains of utilizing Massive MIMO for realizing ultra-reliable, low-latency communication which is an essential part of the tactile Internet and thus applications that rely on haptic feedback. Although thsi paper specifically investigates the requirements for the tactile internet and haptic feedback, the results can generally be applied to any ultra reliable communication scenario in robotics, control, IoT, etc. We have addressed the minimum reliability and latency requirements for these type of applications and through systematic simulation studied and analyzed the performance of Massive MIMO given these requirements. The results reveal that depending on the precoding scheme used, the performance may vary but that ZF is highly preferable even without channel coding. Additionally, it arguably would be worthwhile to investigate the performance of Polar Codes in this scenario. Polar Codes [20] have been deemed beneficial for short packet transmission. The latency requirements can be achieved by modifying the frame structure but the trade-off between the latency and the number of simultaneously supportable devices must be taken into account in the design of the system. As a future work, we plan to implement a haptic teleoperation application through our Massive MIMO testbed (LuMaMi) at Lund University with the results founded in this paper as our design parameters.

ACKNOWLEDGEMENTS

REFERENCES

[1] K. Schwab, *The fourth industrial revolution.* Penguin UK, 2017.

[2] G. Niemeyer and J.-J. Slotine, "Stable adaptive teleoperation," *IEEE Journal of oceanic engineering*, vol. 16, no. 1, pp. 152–162, 1991.

[3] G. P. Fettweis, "The tactile Internet: Applications and challenges," *IEEE Vehicular Technology Magazine*, vol. 9, no. 1, pp. 64–70, 2014.

[4] A. Aijaz, M. Dohler, A. H. Aghvami, V. Friderikos, and M. Frodigh, "Realizing the tactile Internet: Haptic communications over next generation 5G cellular networks," *arXiv preprint arXiv:1510.02826*, 2015.

[5] Y.-L. Park, B.-R. Chen, and R. J. Wood, "Design and fabrication of soft artificial skin using embedded micro channels and liquid conductors," *IEEE Sensors Journal*, vol. 12, no. 8, pp. 2711–2718, 2012.

[6] C. Hoymann, D. Astely, M. Stattin, G. Wikstrom, J.-F. Cheng, A. Hoglund, M. Frenne, R. Blasco, J. Huschke, and F. Gunnarsson, "LTE release 14 outlook," *IEEE Communications Magazine*, vol. 54, no. 6, pp. 44–49, 2016.

[7] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive mimo for next generation wireless systems," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 186–195, 2014.

[8] P. Schulz, M. Matthe, H. Klessig, M. Simsek, G. Fettweis, J. Ansari, S. A. Ashraf, B. Almeroth, J. Voigt, I. Riedel *et al.*, "Latency critical IoT applications in 5G: Perspective on the design of radio interface and network architecture," *IEEE Communications Magazine*, vol. 55, no. 2, pp. 70–78, 2017.

[9] K. Moskvitch, "Tactile internet: 5g and the cloud on steroids," *Engineering & Technology*, vol. 10, no. 4, pp. 48–53, 2015.

[10] T. Wirth, M. Mehlhose, J. Pilz, B. Holfeld, and D. Wieruch, "5G new radio and ultra low latency applications: A PHY implementation perspective," in *2016 50th Asilomar Conference on Signals, Systems and Computers*, Nov 2016, pp. 1409–1413.

[11] ABB, "IRB 140 product specifications," http://new.abb.com/products/robotics/industrial-robots/irb-140, accessed on 2017-03-31.

[12] P. F. Hokayem and M. W. Spong, "Bilateral teleoperation: An historical survey," *Automatica*, vol. 42, no. 12, pp. 2035–2057, 2006.

[13] V. Millnert, "Quality-latency trade-off in bilateral teleoperation," Department of Automatic Control, Lund University, Sweden, Master's Thesis ISRN LUTFD2/TFRT--5951--SE, Nov. 2014.

[14] Y. Maddahi, R. A. Rahman, W. keung Fung, and N. Sepehri, "Effect of network quality on performance of bilateral teleoperated hydraulic actuators: a comparative study," *Control and Intelligent Systems*, vol. 41, 2013.

[15] B. Holfeld, D. Wieruch, T. Wirth, L. Thiele, S. A. Ashraf, J. Huschke, I. Aktas, and J. Ansari, "Wireless communication for factory automation: an opportunity for LTE and 5G systems," *IEEE Communications Magazine*, vol. 54, no. 6, pp. 36–43, 2016.

[16] E. Björnson, L. Sanguinetti, J. Hoydis, and M. Debbah, "Designing multi-user MIMO for energy efficiency: When is massive MIMO the answer?" in *2014 IEEE Wireless Communications and Networking Conference (WCNC)*, April 2014, pp. 242–247.

[17] S. Malkowsky, J. Vieira, L. Liu, P. Harris, K. Nieman, N. Kundargi, I. Wong, F. Tufvesson, V. Öwall, and O. Edfors, "The world's first real-time testbed for massive mimo: Design, implementation, and validation," *arXiv preprint arXiv:1701.01161*, 2016.

[18] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Transactions on Wireless Communications*, vol. 9, no. 11, pp. 3590–3600, 2010.

[19] H. Prabhu, J. N. Rodrigues, L. Liu, and O. Edfors, "3.6 a 60pj/b 300mb/s 128× 8 massive mimo precoder-detector in 28nm fd-soi," in *Solid-State Circuits Conference (ISSCC), 2017 IEEE International.* IEEE, 2017, pp. 60–61.

[20] E. Arikan, "Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3051–3073, 2009.