CSI sharing strategies for transmitter cooperation in wireless networks

Paul de Kerret, David Gesbert

CSI Sharing Strategies for Transmitter Cooperation in Wireless Networks

Paul de Kerret and David Gesbert
Mobile Communications Department, Eurecom
Campus SophiaTech, 450 Route des Chappes, 06410 Biot, France
{dekerret,gesbert}@eurecom.fr

Abstract

Multiple-antenna "based" transmitter (TX) cooperation has been established as a promising tool towards avoiding, aligning, or shaping the interference resulting from aggressive spectral reuse. The price paid in the form of feedback and exchanging channel state information (CSI) between cooperating devices in most existing methods is often underestimated however. In reality, feedback and information overhead threatens the practicality and scalability of TX cooperation approaches in dense networks. Hereby we addresses a "Who needs to know what?" problem, when it comes to CSI at cooperating transmitters. A comprehensive answer to this question remains beyond our reach and the scope of this paper. Nevertheless, recent results in this area suggest that CSI overhead can be contained for even large networks provided the allocation of feedback to TXs is made non-uniform and to properly depend on the network's topology. This paper provides a few hints toward solving the problem.

I. Introduction

Wireless communication has become essential to our lives in many ways, through a variety of services as well of devices ranging from pocket phones to laptops, tablets, sensors and controllers. The advent of multimedia dominated traffic poses extra-ordinary constraints on data rates, latency and above all spectral efficiency. In order to deal with the expected saturation of available resources in currently used bands, new wireless systems are designed based on $i)$ greater densification of infrastructure equipments (small cells), and $ii)$ a very aggressive spatial frequency reuse, which in turn results in severe *interference* conditions for cell-edge terminals.

The role played by multiple antenna combining in mitigating interference by means of zero forcing (ZF) (or related criteria) beamforming is well established. Over the last few years, the combination of multiple-antenna approaches together with the concept of cooperation among interfering wireless devices was explored, showing strong promise (see [1] and references therein). In particular, TX-based cooperation allows for avoidance of the interference before it even takes place (e.g. multi-cell MIMO or "Joint Processing CoMP"), or helps to shape it in a way which makes it easier for the receivers (RX) to suppress it (e.g. alignment). TX cooperation methods can be categorized depending on whether the data messages intended at the users must be known at several TXs simultaneously or not. For systems not allowing such an exchange (e.g. due to privacy regulations or low backhaul capabilities), interference alignment (IA) has been shown to be instrumental [2]. In contrast, when user data message exchange is made possible by a specific backhaul routing architecture, multi-cell, a.k.a. "network" MIMO, methods offer the best theoretical benefits [3].

A distinct advantage of TX cooperation over conventional approaches relying on egoistic interference rejection, lies in the reduced number of antennas needed at each RX to ZF residual interference. This gain is further amplified when user data messages exchange among TXs is made possible. For instance, in the case of three interfering two-antenna TXs, relying on RX based interference rejection alone requires three antennas to ZF the interference at each RX, while just two are needed when coordination is enabled via IA [4]. Further, if the three user messages are exchanged among the TXs, thus enabling network MIMO precoding, then just one antenna per TX and RX is sufficient to preserve interference-free transmission.

However, the benefits of multiple antenna transmit cooperation go at the expense of requiring channel state information (CSI) at the TXs. Indeed, whether one considers cooperation with or without user's data sharing, the TXs should in principle acquire the complete CSI pertaining to every TX and RX pair in the network. This is also the case for distributed schemes (e.g. [5]) where the computation of precoders typically relies on iterative techniques where each iteration involve the acquisition of local feedback. Yet, as local feedback is updated over the iterations, this approach implicitly allows each TX to collect information about the precoders and channels of other TXs, hence amounting to an iterative global CSI acquisition at all TXs.

At first glance, CSI feedback and sharing requirements grow unbounded with the network size. Since over the air feedback and backhaul exchange links are always rate and latency limited,

this means the practical application of TX cooperation in dense large networks is difficult.

In this paper, we challenge the common view that interfering TXs engaging in a cooperative scheme *can* or *should* share global (network-wide) CSI. Instead, we formulate the problem of a suitable CSI *dissemination* (or allocation) policy across transmitting devices while maintaining performance close to the full CSIT sharing scenario. We report a couple of findings revealing how the need for CSIT sharing can be alleviated by exploiting specific antenna configurations or decay property of signal strength versus distance, hence making TX cooperation distributed and scalable. We use interference alignment and network MIMO respectively as our driving scenarios.

More specifically, for the cooperation scenario without user data message sharing where alignment of interference is sought, we show how perfect alignment is possible in certain antenna topologies without knowledge of all the channel elements at some TXs. For the network MIMO scenario, this is not the case and we illustrate instead how power decay versus distance can be exploited to substantially reduce the CSI sharing requirements while fulfilling optimal asymptotic rate performance conditions. A common trait behind the findings is that different cooperating TXs can (and often must) live with their own individual partial version of the global CSIT. Hence, CSIT representation quality is bound to be non-uniform across TXs. Consequently, we discuss briefly the problem of multiple-antenna precoding with *TX-dependent* CSI.

II. Brief Notions in the Multiple-Antenna Transmitter Cooperation

We consider fast fading multiple-antenna wireless networks where the transmission can be mathematically represented by writing $\mathbf{y}_i$, the received signal at RX i, as

$$\mathbf{y}_i = \mathbf{H}_{ii}\mathbf{x}_i + \sum_{j \neq i} \mathbf{H}_{ij}\mathbf{x}_j \tag{1}$$

where $\mathbf{x}_i$ is the signal emitted by TX i and $\mathbf{H}_{ij}$ is a matrix containing the channel elements between TX j and RX i. The transmitted symbols $\mathbf{x} = [\mathbf{x}_1, \ldots, \mathbf{x}_K]^{\mathrm{T}}$ are then obtained from the user's data symbols $\boldsymbol{s} = [s_1, \ldots, s_K]^{\mathrm{T}}$ by multiplication with a precoder $\mathbf{T}$, i.e., $\mathbf{x} = \mathbf{T}\boldsymbol{s}$. If the user's data symbols are not shared between the TXs, the precoder $\mathbf{T}$ is restricted to a particular block-diagonal structure, while it can otherwise take any form. The received filter $\boldsymbol{g}_i^{\mathrm{H}}$ is then applied to the received signal $\boldsymbol{y}_i$ to obtain an estimate of the transmitted data symbol.

Here, we briefly discuss the leading techniques for MIMO based cooperation with or without user data message exchange. We point out commonly made assumptions in terms of CSIT sharing and feedback design.

A. Interference Alignment for Interference Channels

When the user's data symbols are not shared between the TXs, the setting is referred to as a *Interference Channel (IC)* in the communication theoretic literature. In MIMO ICs, a method called *Interference Alignment (IA)* has been recently developed and shown to achieve the maximal number of degrees-of-freedom (DoF), or pre-log factor, in many cases [2], [4]. As a consequence, IA has attracted a lot of interest in the community. In this work, we will take the DoF as our key performance metric, such that we focus on IA schemes.

IA is said to be *feasible* if the antenna configuration (i.e. the distribution of antenna elements at the TXs and the RXs) yields enough optimization variables to allow for the interference-free transmission of all user's data symbols, which means fulfilling [4]

$$\forall i, \forall j \neq i, \quad \boldsymbol{g}_i^{\mathrm{H}} \mathbf{H}_{ij} \boldsymbol{t}_j = 0. \tag{2}$$

Intuitively, IA consists in letting the TXs coordinate among themselves to beamform their signals such that the interferences received at each of the RXs are confined in a subspace of reduced dimensions, which can then be suppressed by linear filtering at the RXs with a smaller number of antennas.

B. Precoding in the Network MIMO

When the user's data symbols are shared between the TXs, the TXs form a *distributed antenna array* and a joint precoder can be applied at the transmit side [3]. Consequently, this setting becomes similar to the single TX multi-user MIMO downlink channel and the interference between the TXs can be completely canceled, e.g., by applying a global ZF precoder $\mathbf{T} \propto \mathbf{H}^{-1}$.

These two scenarios are schematically represented in Figure 1.

C. Limited Feedback Versus Limited Sharing

The limited feedback capabilities have been recognized as a major obstacle for the practical use of the precoding schemes described above. Consequently, a large literature has focused on

this problem and both efficient feedback schemes and robust transmission schemes have been derived, for Network MIMO [6], [7] and IA [8], [9].

Yet, all these works assume that the imperfect channel estimates obtained via limited feedback are *perfectly* shared between all the transmit antennas. This is a meaningful assumption when the TXs are colocated but less realistic otherwise, as we shall now see.

CSIT Sharing Issues: One obstacle to the sharing of global CSIT follows from the fact that the amount of CSI which has to be exchanged increases very quickly with the number of TXs. In fact, each TX needs to obtain the CSI relative to the full multi-user channel, which consists of $(NK)^2$ scalars in a K-user setting with N antennas at each node.

In addition, acquiring the CSI at a particular TX can be realized either by a direct broadcast of the CSI to all the listening TXs or by an over-the-air feedback to the *home* base station alone, followed by an exchange of the local CSIs over the backhaul, as it is currently advocated by 3GPP LTE-A standards [10]. See the illustrations for such scenarios in Figure 2. Note that exchange over the backhaul can involve further quantization loss and may lead to a different CSI-aging at each TX, due to protocol latency. Either case, the channel estimates available at the various TXs will not be exactly the same. This leads to a form of CSI discrepancy which is inherent to the cooperation among non-colocated TXs.

In order to capture multiple-antenna precoding scenarios whereby different TXs obtain an imperfect *and* imperfectly shared estimate of the overall multi-user channel, we denote by $\mathbf{H}^{(j)}$ the network-wide channel matrix estimate available at TX j. Consequently, the precoding schemes in Figure 1 have to be modified to take into account that each TX will compute its precoder based on its own channel estimate. Thus, TX j transmits $\mathbf{x}_j = \mathbf{T}^{(j)}\boldsymbol{s}$ based on the knowledge of $\mathbf{H}^{(j)}$ only.

The fundamental questions which arise are: (i) how complete and accurate should the estimate $\mathbf{H}^{(j)}$ be for each j while operating under reasonable CSI overhead constraints? and (ii) how should precoders be designed given the likely discrepancies between various channel estimates? Although these questions prove to be difficult and to a large extent remain open, we shed some light on the problem for two key scenarios in the following sections.

III. Aligning Interference with Incomplete CSIT

Let us first consider an IC, i.e., without user's data sharing. Feasibility studies for IA are typically carried out under the assumption of full CSIT. Yet, one can show that IA feasibility and the CSIT model are in fact tightly coupled notions. Assume for instance that all the RXs were given a generous number of antennas equaling or exceeding the number of TXs, it is well known that the interference could be suppressed at the RXs alone and no precoding, and hence no CSIT, is necessary. This example suggests the existence of a trade-off between the number of antennas and the CSI sharing requirements. Thus, it is possible to design IA algorithms using less CSIT that conventionally thought, without performance degradations by exploiting the availability of extra-antennas at a subset of devices. More specifically, the problem of finding the minimal CSIT allocation which preserves IA feasibility can be formulated. The minimality refers to the size of a CSIT allocation, defined as the total number of scalars sent through the multi-user feedback channel.

We differentiate between antenna configurations where IA is feasible and the number of antennas at the TXs and the RXs provide just enough optimization variables to satisfy alignment conditions, denoted as *tightly-feasible*, and the ones where extra antennas are available, denoted as *super-feasible*. Furthermore, we call a CSIT allocation *strictly incomplete* if at least one TX does not have the complete multi-user CSI. With such concepts in place, the following lesson can be drawn.

A. Tightly-feasible ICs

A strictly incomplete CSIT allocation implies that some TXs compute their precoders in order to fulfill IA inside a smaller IC formed by a subset of RXs and a subset of TXs. Most of the time, this creates additional constraints for the optimization of the other precoders which makes IA unfeasible. Yet, it can be shown that IA feasibility can be preserved under the following condition [11].

Lesson 1. *In a tightly-feasible IC, there exists a strictly incomplete CSIT allocation preserving IA feasibility if there exists a tightly-feasible sub-IC strictly included in the full IC.*

Exploiting this result, a CSIT allocation algorithm is derived in [11] along with an algorithm which achieves IA based on this incomplete CSIT allocation. In a few words, it consists in

giving to each TX the CSI relative to the smallest tightly-feasible sub-IC to which it belongs. The precoders are then designed to align interference inside this smaller sub-IC, thence requiring only a part of the total CSIT, while IA feasibility is preserved. We will see in the simulations results presented in the following that the reduction in the CSIT size is significant. In fact, the reduction of the CSIT allocation feeds on the heterogeneity of the antenna configuration such that the more heterogeneous the antenna configuration is, the larger is the saving brought by using the minimal CSIT allocation. This is particularly appealing in regards to the future networks where mobile units and base-stations from different generations with different number of antennas are likely to co-exist.

B. Super-feasible ICs

In super-feasible settings, the additional antennas can be used to reduce the size of the minimal CSIT allocation. Yet, how to exploit optimally these additional antennas to reduce the feedback size is a very intricate problem. Still, a low complexity heuristic CSIT allocation can be derived [11]. The main idea behind the algorithm is to let some TXs or RXs ZF less interference dimensions such that small tightly-feasible settings are formed inside the original setting.

The effective CSIT reduction is illustrated in Figure 3 for a 3-user IC. The results are averaged over 1000 random distributions of the antennas across the TXs and the RXs. If 12 antennas are distributed between the TXs and the RXs, the setting is tightly-feasible and the previous CSIT allocation policy for tightly-feasible settings is used. With more than 12 antennas, the algorithm exploits every additional antenna to reduce the size of the CSIT allocation.

When the setting is tightly-feasible, the reduction in feedback size requires neither a DoF reduction nor any additional antenna and comes in fact “for free”: It simply results from exploiting the heterogeneity in the antenna numbers at the TXs and the RXs.

IV. A CSIT Allocation Policy for Network MIMO

When the user’s data symbols are jointly precoded at the TXs, complete CSIT allocation, in the sense defined above, is needed in all the practically relevant scenarios. So in this case, a different notion of reduced CSIT sharing must be advocated. The essential ingredient of this approach is the classical intuition that a TX should have a more accurate estimate for channels creating the strongest interference, i.e. originating from devices in the close neighborhood. This

means that the fact that interference decays with pathloss can be exploited in principle to reduce the CSIT sharing requirements. This concept was recently introduced in [12]. A mathematical tool known as generalized degree-of-freedom comes handy to capture the effect of path loss on the multiplexing gain of cooperating networks with partially shared CSIT. Additionally, a simplified model referred to as Wyner model is used in this context to aid analytical tractability and provide first insights into this problem.

A. Generalized Degrees of Freedom

TX cooperation methods are often evaluated through the prism of DoF performance. Unfortunately the DoF is essentially pathloss-independent, such that a DoF analysis fails to properly capture the behavior of a large (extended) network MIMO. An extension of the notion of DoF, introduced in [13] as the *generalized DoF*, offers a much better grip over the problem as it can better take pathloss models into account. Upon defining the *interference level* γ as $\gamma \triangleq \log(\text{INR})/\log(\text{SNR})$ with SNR denoting the signal-to-noise ratio and INR the interference-to-noise ratio, it is possible to define the *generalized DoF* as the DoF obtained when the SNR and the INR tend both to infinity for *a given interference level* γ.

For ease of exposition, we consider scenarios where all the TXs and RXs have a single antennas. The CSI is distributed, meaning that each TX has its own channel estimate based on which it computes its transmit coefficient without further exchange of information with the other TXs. We denote the estimate at TX j by $\mathbf{H}^{(j)}$ and its i-th row, which corresponds to the channel from all TXs to RX i, by $\boldsymbol{h}_i^{(j)}$. We consider a digital quantization with a number of bits quantization $\boldsymbol{h}_i^{(j)}$ equal to $B_i^{(j)}$. Therefore, TX j computes its own version of the precoding matrix $\mathbf{T}^{(j)}$ based on its own estimate $\mathbf{H}^{(j)}$. It then transmits $\mathrm{x}_j = \boldsymbol{e}_j^{\mathrm{T}}\mathbf{T}^{(j)}\boldsymbol{s}$.

B. The One Dimensional Wyner Model

In the simple 1D *Wyner model* [14], the TXs are regularly placed along a line and solely the direct neighboring TXs emit non-zero interference. The channel is thus represented by a tridiagonal matrix. Furthermore, we assume that the interference from the direct neighbors are attenuated with a coefficient $\mu = P^{\gamma-1}$, according to the generalized DoF model. The transmission in the Wyner model is schematically presented in Figure 4.

Our objective is to evaluate how small $B_i^{(j)}$ can be while guaranteeing the same DoF as a system with perfect CSIT. Obviously, sharing the most accurate CSIT to all the TXs is a possible solution, yet, the size of the CSI required *at each TX* grows then unbounded with the number of users K, making this solution both inefficient and unpractical. In contrast, a much more efficient CSI sharing policy achieving the maximal generalized DoF, denoted as *distance-based*, is summarized below [12].

C. Distance-Based CSIT Allocation

We are interested in a CSIT allocation strategy, referred below as "distance-based", whereby each TX receives a number of CSI scalars which remains *bounded* as the number of users K increases.

The distance-based CSIT allocation is obtained by setting for all i, j,

$$B_i^{(j)\,\mathrm{Dist}} = \lceil ([1+(\gamma-1)|i-j|]^+ + 2[\gamma + (\gamma-1)|i-j|]^+) \log_2(P) \rceil \tag{3}$$

where $[\bullet]^+$ is equal to zero if the argument is negative and to the identity function otherwise, and $\lceil \bullet \rceil$ is the ceiling operator. It can be shown that the CSIT allocation $\{B_i^{(j)\,\mathrm{Dist}}\}_{i,j}$ allows to achieve the maximal generalized DoF, i.e., those achieved in a system with perfect feedback [12]. The proof is based on the off-diagonal exponential decay of the inverse of the tridiagonal channel matrix.

Setting $\gamma = 1$ (no significant pathloss attenuation) in the previous equation, a conventional CSIT allocation is obtained where all channels are described with the same number of bits at all TXs (uniform allocation). For $\gamma < 1$ however, the number of bits allocated decreases with the distance $|i-j|$ between the considered TX and the index of the quantized channel, until no bit at all is used for quantizing the channel if the distance $|i-j|$ between the RX and the TX is larger than $\lceil 1/(1-\gamma) \rceil$. Crucially, this solution allocates to each TX a total number of bits which no longer grows with K as only the CSI relative to a *neighborhood* is shared at each TX.

The different CSIT allocations are compared in Figure 5 for $K = 15$ users and $\gamma = 0.5$. The conventional CSIT allocation consists in providing the best quality to all the TXs, while the other strategies have the same size as the distance-based CSIT allocation, but the feedback bits are respectively shared uniformly and according to a conventional clustering of size 3. It can be seen from (3) that the ratio between the size of the distance-based CSIT allocation and

the conventional CSIT allocation is independent of the SNR. Here, the distance-based CSIT allocation represents only 10% of the size of the conventional CSIT allocation.

Hence, the distance-based CSIT allocation achieves the maximal number of generalized DoF with only a small share of the total CSIT, and outperforms the other schemes of comparison. Additionally, the user's data sharing can also be reduced to a neighborhood without loss of performance. Consequently, this scheme can be seen as an alternative to clustering in which the hard boundaries of the clusters are replaced by a smooth decrease of the level of cooperation.

V. Precoding in the Network MIMO Channel with Distributed CSIT

As it becomes clear from the previous section, an efficient CSI dissemination policy naturally leads to a significant reduction of the CSI sharing requirements. As a consequence, the CSIT is represented non-uniformly across the TXs. Since non-uniform sharing is the best strategy in order to maximize performance under a given feedback overhead constraint, it is a natural consequence that some user's channels will be coarsely described at certain TXs and more accurately at others. Interestingly, the problem of designing precoders that can accommodate such a peculiar CSIT scenario is by and large open. In particular, new robust precoding schemes should be developed, as conventional precoders are designed under the assumption that the same imperfect CSIT is shared *perfectly* among the TXs.

Let us consider the model of distributed CSI described in Section IV where TX j receives its own channel estimate $\mathbf{H}^{(j)}$ with the i-th row, denoted by $\boldsymbol{h}_i^{(j)}$, obtained using $B_i^{(j)}$ quantizing bits. Assume a network where each TX has roughly the same average pathloss to each RX. The DoF which can be achieved with limited feedback is studied in [7] for the single TX MIMO downlink channel. We can extend this to the setting of non-uniform CSI so as to gain insight into the design of efficient precoders. In this case, *CSI scaling* coefficients $\alpha_i^{(j)}$ are introduced and defined by the limit of $B_i^{(j)}/((K-1)\log_2(P))$ when P goes to infinity.

ZF is widely used and well known to achieve the maximal DoF in the MIMO downlink channel with perfect CSIT [7]. One may wonder how conventional ZF performs in the presence of CSI discrepancies brought by imperfect sharing. The answer is strikingly pessimistic: The sum DoF achieved can be shown to be equal to just $K \min_{i,j} \alpha_i^{(j)}$ [15]. Intuitively, this can be restated as follows.

Lesson 2. *In the network MIMO with distributed CSI, the worst channel estimate across the TXs and the users limits the DoF achieved by each user using ZF precoding.*

This is in strong contrast to the single TX case studied in [7] where the quality of the feedback of user i relative to $\boldsymbol{h}_i$ solely impacts the DoF of user i. It shows clearly the disastrous impact of the CSI non-uniformity, since one inaccurate estimate at one TX degrades the performances of *all* the users. One may also wonder whether conventional-type robust precoders [6] can offer a better response. The answer is negative, unfortunately. Instead, a novel precoder design is needed that is tailored to the non-uniform CSIT sharing model.

Preliminary results to this end [15] suggest that it is possible to dramatically improve the DoF in certain scenarios. For instance, in the two-TX network, a scheme referred to as *Active-Passive (AP) ZF*, consisting in letting the TX with degraded CSIT arbitrarily fix its transmit coefficient while the other TX compensates to zero-out the interference, can be shown to recover the optimal DoF.

The average rate achieved with conventional ZF, AP ZF, and ZF with perfect CSIT are compared in Figure 6. In that case, the sum rate of conventional ZF saturates at high SNR while AP ZF is more robust and achieves a better DoF.

VI. Open Problems

New concepts for the CSIT sharing in wireless networks have been derived and their potential to reduce signaling overhead has been shown. We have presented some insights into a new problem which presents serious challenges, but also opportunities for the future. This leads to new intriguing open questions. Firstly, IA algorithms with incomplete CSIT are based on a DoF-preserving criterion only, i.e., on the performance at asymptotically high SNR. The impact of the incomplete CSIT on the performance at finite SNR should then be investigated to obtain practical solutions. Similarly, robust precoding schemes for the MIMO network with distributed CSI have so far considered DoF only as a metric. By and large, precoding over network MIMO with distributed CSI remains a challenging problem. Regarding the optimization of the CSIT allocation in a network MIMO channel, we believe that the distance-based CSIT allocation has the potential for impacting the design of cellular systems in practical settings. Yet, it has been examined for simplified channel models only and should be adapted to more realistic scenarios.

VII. Conclusion

The uniform allocation of CSI resources towards all TXs does not lead to an efficient use of the feedback and backhaul resources: The CSIT dissemination should be designed to allocate to each TX the right CSI, in terms of which elements to share and in terms of accuracy. For both network MIMO and IA, the CSIT sharing requirements have been significantly lowered by designing appropriate CSIT dissemination policies, thus making TX cooperation more practical and thereby paving the way for large performance improvement. Additionally, the problem of robust precoding schemes taking into account the inconsistency between the CSIs at the TXs has been tackled and shown to be a key element.

VIII. Acknowledgment

This work has been performed in the framework of the European research project HARP, which is partly funded by the European Union under its FP7 ICT Objective 1.1 - The Network of the Future.

References

[1] D. Gesbert, S. Hanly, H. Huang, S. Shamai (Shitz), O. Simeone, and W. Yu, "Multi-cell MIMO cooperative networks: a new look at interference," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1380–1408, Dec. 2010.

[2] V. R. Cadambe and S. A. Jafar, "Interference alignment and degrees of freedom of the K-user interference channel," *IEEE Trans. Inf. Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.

[3] M. K. Karakayali, G. J. Foschini, and R. A. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Communications*, vol. 13, no. 4, pp. 56–61, Aug. 2006.

[4] C. Yetis, T. Gou, S. A. Jafar, and A. H. Kayran, "On feasibility of interference alignment in MIMO interference networks," *IEEE Trans. Signal Process.*, vol. 58, no. 9, pp. 4771–4782, Sep. 2010.

[5] K. Gomadam, V. R. Cadambe, and S. A. Jafar, "A distributed numerical approach to interference alignment and applications to wireless interference networks," *IEEE Trans. Inf. Theo.*, vol. 57, no. 6, pp. 3309–3322, Jun. 2011.

[6] M. B. Shenouda and T. N. Davidson, "Robust linear precoding for uncertain MISO broadcast channels," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2006.

[7] N. Jindal, "MIMO broadcast channels with finite-rate feedback," *IEEE Trans. Inf. Theory*, vol. 52, no. 11, pp. 5045–5060, Nov. 2006.

[8] H. Bolcskei and I. J. Thukral, "Interference alignment with limited feedback," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, 2009.

[9] O. E. Ayach and R. W. Heath, "Interference alignment with analog channel state feedback," *IEEE Trans. Wireless Commun.*, vol. 11, no. 2, pp. 626–636, Feb. 2012.

[10] S. Sesia, I. Toufik, and M. Baker, *LTE - The UMTS long term evolution: From theory to practice*, 2nd ed. Wiley, 2011.

[11] P. de Kerret and D. Gesbert, “Interference alignment with incomplete CSIT sharing,” submitted to IEEE Trans. Wireless Commun., Nov. 2012.

[12] ——, “CSI feedback allocation in multicell MIMO channels,” in *Proc. International Conference on Communications (ICC)*, 2012.

[13] R. Etkin, D. Tse, and H. Wang, “Gaussian interference channel capacity to within one bit,” *IEEE Trans. Inf. Theory*, vol. 54, no. 12, pp. 5534–5562, Dec. 2008.

[14] A. D. Wyner, “Shannon-theoretic approach to a Gaussian cellular multiple-access channel,” *IEEE Trans. Inf. Theory*, vol. 40, no. 6, pp. 1713–1727, Nov. 1994.

[15] P. de Kerret and D. Gesbert, “Degrees of freedom of the network MIMO channel with distributed CSI,” *IEEE Trans. Inf. Theory*, vol. 58, no. 11, pp. 6806–6824, Nov. 2012.

PLACE PHOTO HERE

Paul de Kerret (IEEE Student Member) was born in 1987 in Paris, France. In 2009, he graduated from Ecole Nationale Superieure des Telecommunications de Bretagne, France and obtained a diploma degree in electrical engineering from Munich University of Technology (TUM), Germany. He also earned a four year degree in mathematics at the Universite de Bretagne Occidentale, France in 2008. From january 2010 to september 2010, he has been a research assistant at the Institute for Theoretical Information Technology, RWTH Aachen University, Germany. Since october 2010, he is working toward a Ph.D. degree in the Mobile Communications Department at EURECOM, France.

PLACE PHOTO HERE

David Gesbert (IEEE Fellow) is Professor and Head of the Mobile Communications Department, EURECOM, France. He received his Ph.D in 1997 from Telecom Paris Tech. From 1997 to 1999 he has been with the Information Systems Laboratory, Stanford University. In 1999, he was a founding engineer of Iospan Wireless Inc, San Jose, Ca., a startup company pioneering MIMO-OFDM (now Intel). D. Gesbert has guest edited 6 special issues, published about 200 papers and several patents in the area wireless networks, among which three won international awards. He co-authored the book Space time wireless communications: From parameter estimation to MIMO systems, Cambridge Press, 2006.

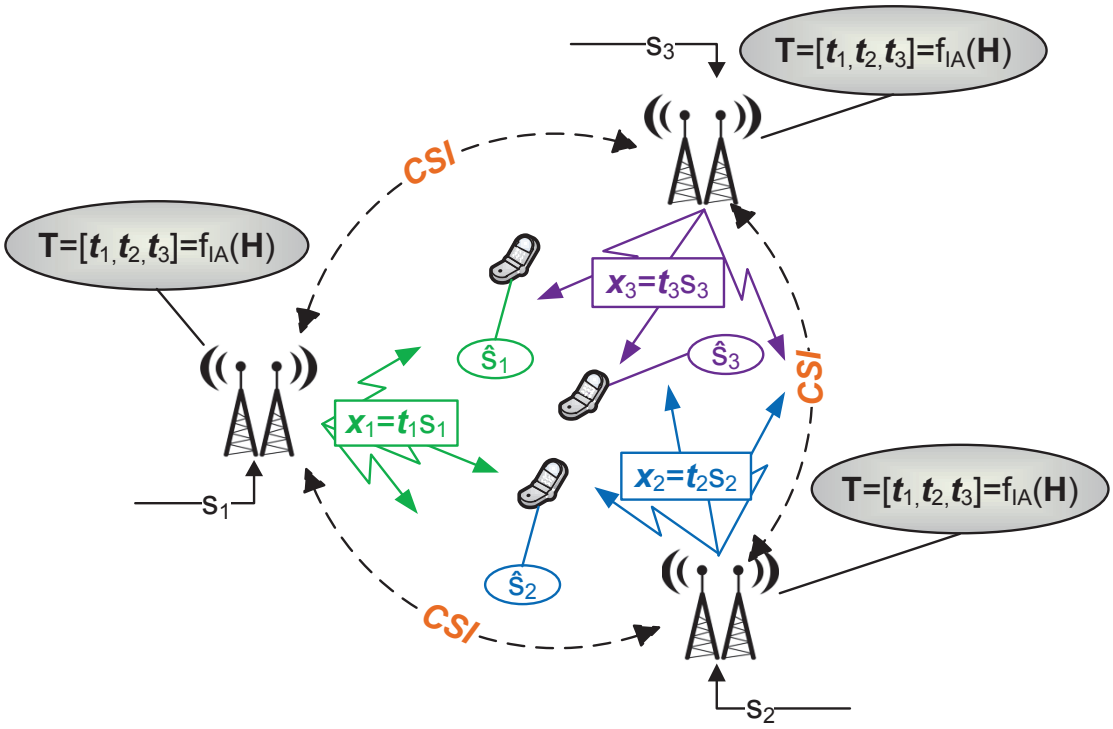

(a)

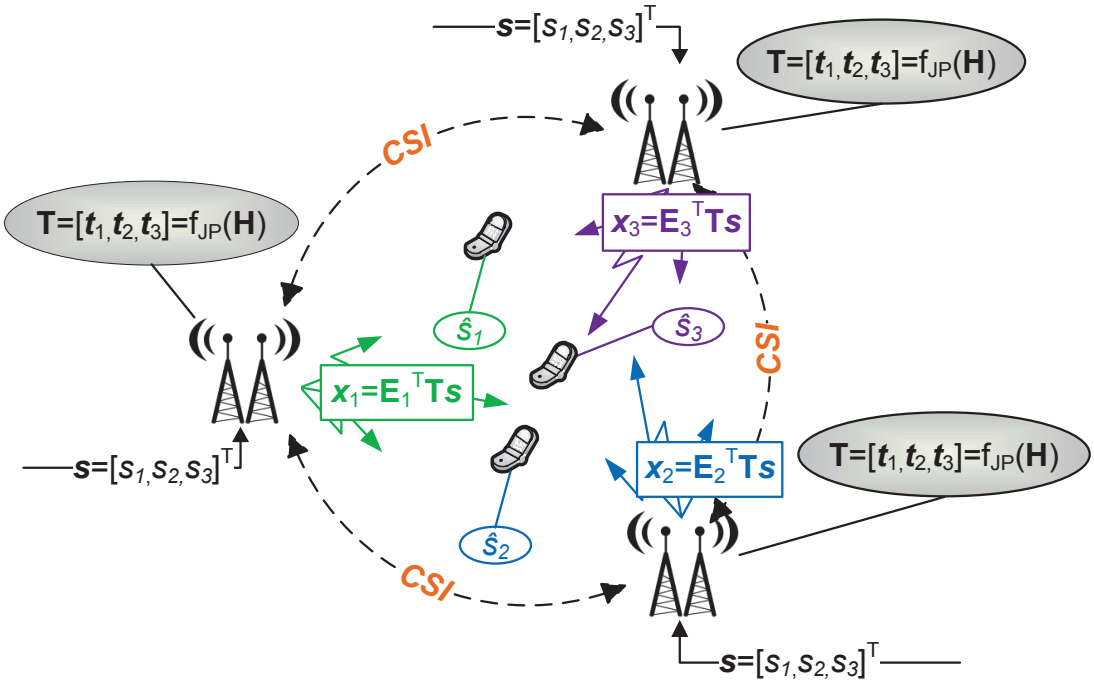

(b)

Fig. 1: The description of a distributed IA algorithm is done in Figure (a) while Figure (b) represents the distributed precoding in a Network MIMO with user's data sharing. The matrix $\mathbf{E}_i$ is a matrix which selects the rows of the multi-user precoder $\mathbf{T}$ corresponding to the antennas at TX i.

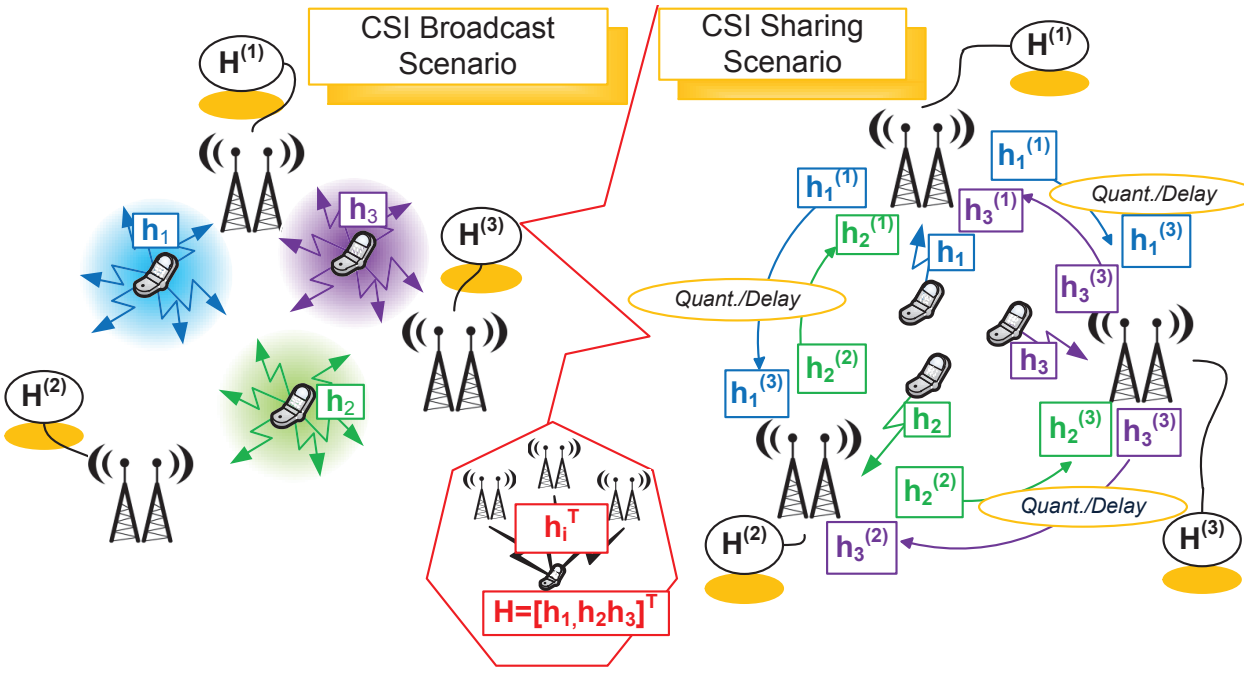

Fig. 2: Possible scenarios for the feedback of the CSI.

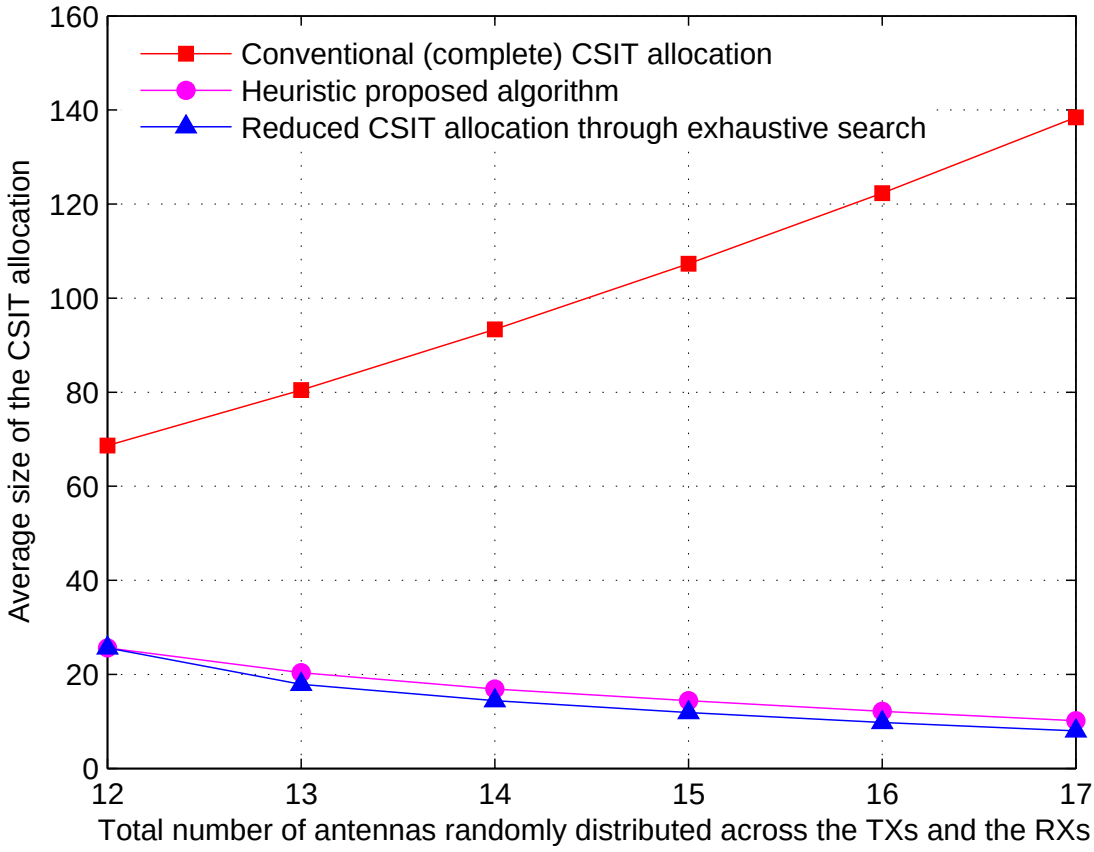

Fig. 3: Average CSIT allocation size in terms of the number of antennas randomly distributed across the TXs and the RXs for $K = 3$ users.

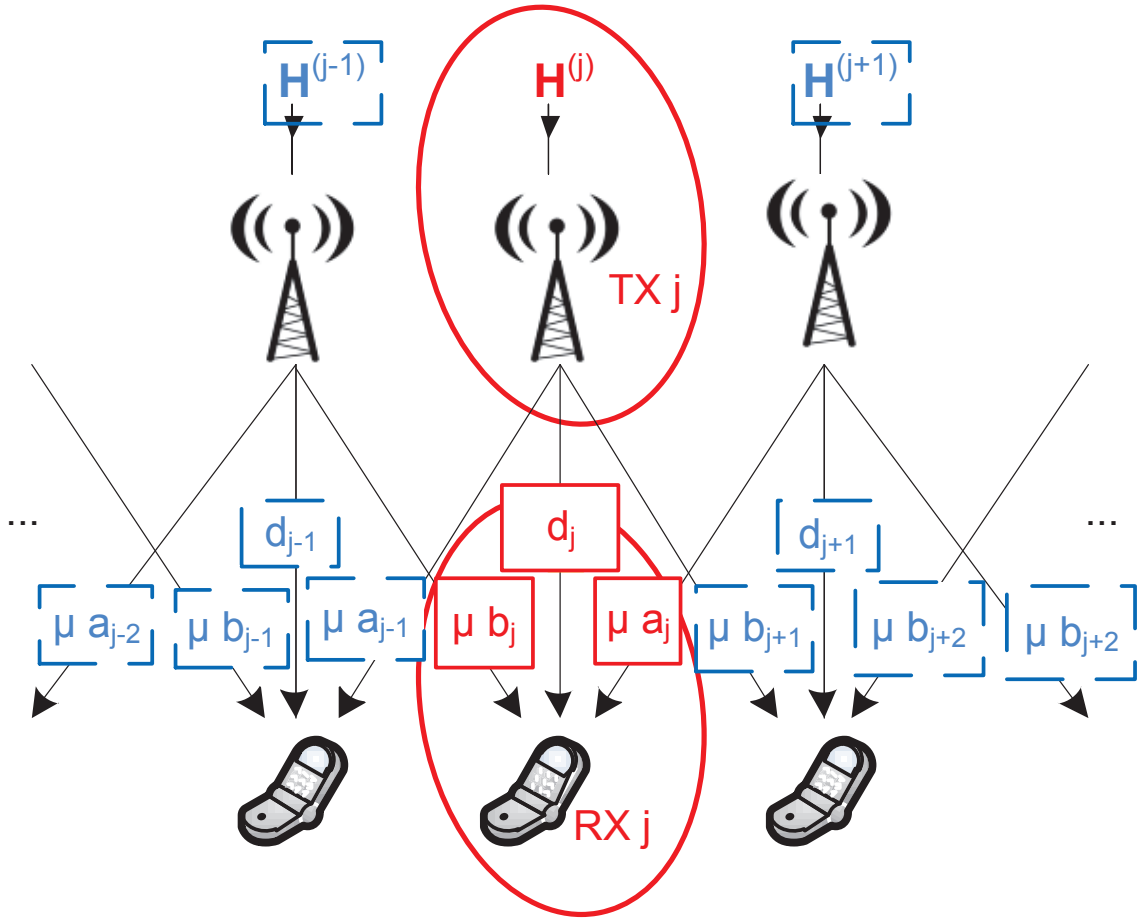

Fig. 4: Schematic representation for the Wyner model considered.

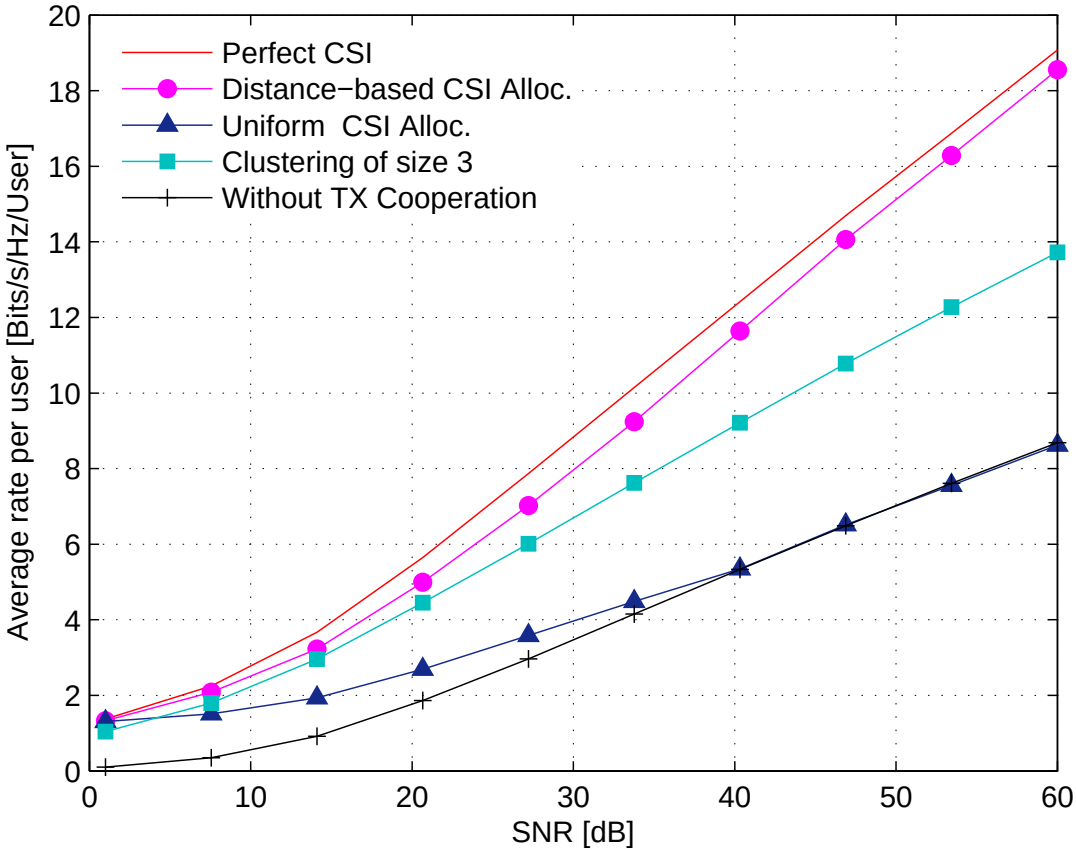

Fig. 5: Average rate per user in terms of the SNR. The distance-based CSIT allocation, the uniform allocation, and the clustering one have all a size equal to 10% of the size of the conventional CSIT allocation.

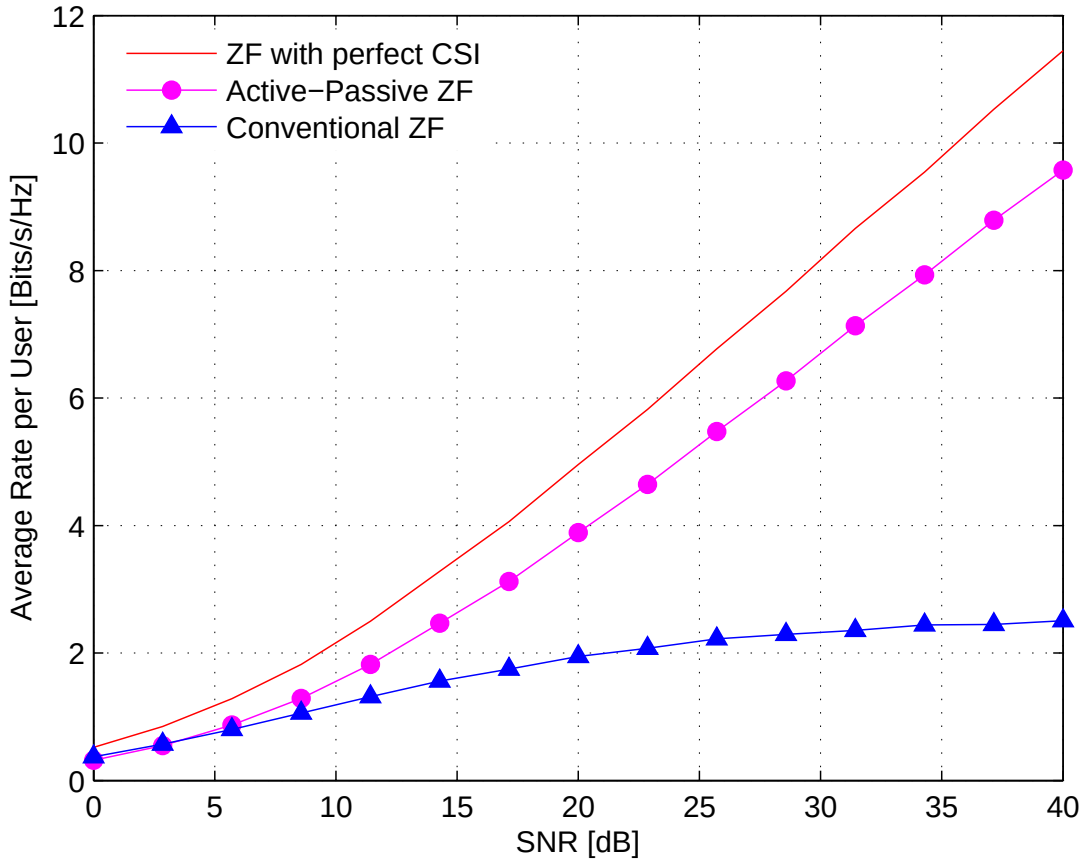

Fig. 6: Average rate per user in terms of the SNR for $\alpha_1^{(1)} = 1$, $\alpha_2^{(1)} = 0$, $\alpha_2^{(1)} = 0.5$, $\alpha_2^{(2)} = 0.7$.