King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Why to Decouple the Uplink and Downlink in Cellular Networks and How To Do It

Federico Boccardi[1], Jeffrey Andrews[2], Hisham Elshaer[3,4], Mischa Dohler[3]
Stefan Parkvall[5], Petar Popovski[6], Sarabjot Singh[7]

[1]Ofcom, London, UK
[2]The University of Texas at Austin, USA
[3]King's College London, UK
[4]Vodafone, Newbury, UK
[5]Ericsson, Stockholm, Sweden
[6]Aalborg University, Aalborg, Denmark
[7]Nokia, Berkeley, USA

5 November 2015

Abstract

Ever since the inception of mobile telephony, the downlink and uplink of cellular networks have been *coupled*, i.e. mobile terminals have been constrained to associate with the same base station (BS) in both the downlink and uplink directions. New trends in network densification and mobile data usage increase the drawbacks of this constraint, and suggest that it should be revisited. In this paper we identify and explain five key arguments in favor of Downlink/Uplink Decoupling (DUDe) based on a blend of theoretical, experimental and architectural insights. We then overview the changes needed in current (LTE-A) mobile systems to enable this decoupling, and then look ahead to fifth generation (5G) cellular standards. We demonstrate that decoupling can lead to significant gains in network throughput, outage and power consumption at a much lower cost compared to other solutions that provide comparable or lower gains.

1   **Introduction and Background**

From the first to the fourth generation (4G) of mobile networks, the downlink (DL) and uplink (UL) of a given communication session have been coupled: the mobile user equipment (UE) must associate with the same BS in both the DL and UL. Historically, this was a nearly optimal approach, since the strongest BS-UE connection was the same in both directions. However, this conventional approach has recently [1] come under scrutiny given the possible gains that can be achieved by decoupling the association in the context of a dense heterogeneous cellular network, wherein different BSs can have highly variable transmit powers and deployment topologies.

The arguments in favor of the coupled *status quo* are several. From a pure network design perspective, the logical, transport and physical channels are easier to design and operate; this pertains particularly to the synchronization of the acknowledgements, the call admission and handover procedures, DL/UL radio resource management, and power control, among others. Decoupling both links also requires strong synchronization and data connectivity (e.g. via fiber) between the BSs. From a deployment and topology perspective, until just a few years ago cellular systems have been designed and deployed under the assumption of a homogenous network with macro cells all transmitting with about the same power. From a traffic point of view, the load in both directions has been approximately the same in voice-centric 2G and early 3G systems. Moreover, 3.5G (e.g. HSPA) and 4G systems are dominated by downlink traffic, justifying the use of DL-centric association procedures rather than UL or decoupled ones.

The emergence of Heterogeneous Networks (HetNets) [2], where small cells at higher carrier frequencies and/or smaller transmit powers are deployed within the coverage area of macro cells, calls for revisiting the coupled association approach. Range extension has been included in 4G to add a bias in the cell association to offload more traffic from macro to small cells. Data and control plane separation has been introduced in [3]: the control information is sent by high-power nodes at lower frequencies, whereas the payload data is conveyed by low-power nodes at possibly higher frequencies. However, both range extension and data/control plane separation are based on a coupled DL/UL association, where DL and UL are associated to the same BS.

The motivation for Downlink/Uplink Decoupling (DUDe) emerges from a holistic view on the two-way (DL/UL) traffic and the association procedure of a UE, rather than adopting a coupled association *a priori* and then optimizing separately DL and UL transmissions. Since a coupled association is a particular sub-case of a decoupled one, a well-designed association policy based on Downlink/Uplink Decoupling (DUDe) can in principle outperform a coupled association. But by how much? And at what cost?

More specifically, the main questions this article attempts to answer are:

(i)   What recent trends in cellular network deployment and applications make the gains from DUDe more relevant now than in the past?

(ii)   What are the key benefits of a decoupled association in terms of throughput gain, reliability, and power conservation? What are the challenges? How can these gains be realized in current (e.g. LTE-A) and future 5G cellular networks?

(iii)   How disruptive will these changes be to the network architecture? Are the gains large enough to be worth the trouble?

We note that research developments on DUDe are recent and limited to a few contributions. The the interest in decoupling the downlink and uplink has been indicated in [4], [5] and [6], and has been further explored in a few subsequent contributions. In particular, [7], [8] and [9] studied the throughput and SINR gains from a theoretical perspective, while [1] assessed DUDe via detailed industry-standard simulations.

We begin the discussion in this paper with the five key arguments in favor of DUDe, and provide evidence for the corresponding gains from very recent theoretical analysis and simulation-based experiments. Then, we move on to discuss what changes will need to be made in the current and future cellular standards, and explain why, in our view, such changes are quite manageable. DUDe opens up many new interesting research questions as well, which we identify throughout the article.

2   Five Reasons to Decouple the Downlink and Uplink

We now articulate the five principle arguments in favor of DUDe. Our arguments are supported by a combination of recent theoretical and system-level simulation results by the present authors and others. In particular, the theoretical results are mostly sourced from the recent work [7], in which we perform a comprehensive SINR and rate analysis with DUDe in a multi-tier cellular network with spatially random UEs and BS. The UEs employ fractional UL power control and small-cell biasing is used to achieve cell range expansion: both very similar to LTE. The results are mathematical and thus transparent, albeit in some cases based on idealized models to allow tractability. We refer to this approach as the *analytical model*.

The simulation results and parameters follow largely from [1], and utilize an existing LTE HetNet deployment in conjunction with a high resolution 3D ray tracing channel model that takes into account clutter, terrain and building data. This ensures a highly realistic and accurate propagation model. The BS types and locations are based on a small cell test network in the London area and consist of five macro cells covering a one kilometer square area with a dense small cell deployment embedded in the square kilometer. The UE distribution is based on live traffic measurements, and the UEs use the same UL fractional power control as in the analytical model. We assume that the DL association is based on the DL Reference Signal Received Power (RSRP). We refer to this approach as the *simulation model*.

As we will see, these two distinct approaches to modeling and analyzing DUDe are quite unified in terms of the conclusions they offer. Table 1 contains the cellular network notation and simulation parameters. We also use the same parameters for numerical evaluation of mathematically derived results using the analytical model.

Table 1. Cellular network notation and parameter values.

Parameter	UEs	Macro cells	Small cells

Max Transmit Power	20 dBm	46 dBm	30 dBm (unless otherwise specified)
Antenna system (simulation)	1 Tx and 1 Rx Antenna gain = 0 dBi	2 Tx and 2 Rx Antenna gain = 17.8 dBi	2 Tx and 2 Rx Antenna gain = 4 dBi
Antenna system (analysis)	The analysis considers 1 Tx and 1 Rx isotropic antenna system for UEs, Macro cells and Small cells.		
Downlink Bias	N/A	0 dB	Varies from 0 to 8 dB
Spatial distribution (analytical)	Uniform	Poisson Point Process	Poisson Point Process
Spatial distribution (simulation)	Hotspot distribution based on realistic traffic measurements	Based on the Vodafone LTE network deployment.	Based on the Vodafone LTE small cell test network deployment.
Spatial Density	330 per km^2	5 per km^2	4 small cells per macro cell
Channel Model (analytical)	Rayleigh small scale fading Standard path loss with exponent = 3.5 Lognormal shadowing with standard deviation = 8 dB		
Channel Model (simulation)	3D ray-tracing propagation model		
Power control	Uplink fractional path loss compensation		
Operating frequency	2.6 GHz co-channel FDD deployment		
Bandwidth	20 MHz (100 frequency blocks)		
Scheduler	Equipartition of resources among UEs (analysis) Proportional fair (simulation)		

2.1 Increased uplink SNR and reduced transmit power

In a typical HetNet scenario the downlink coverage area of a macro cell is much larger than that of a low power BS; indeed, this is why they are often called "small cells". The coverage area disparity is primarily attributable to the differences in downlink transmit powers, but is also due to the BS heights and antenna gains. In contrast, in the uplink all transmitters have roughly the same maximum transmit power. Therefore, a device that is associated to a macro cell in the downlink might instead wish to be associated to a small cell in the uplink, to take advantage of the reduced path loss [8]. The positive effects are twofold. For UEs that are transmitting at the maximum power, a connection to a closer BS provides a higher SNR. Moreover, for a fixed target SNR, the reduced path loss alternatively allows transmit power reduction via power control.

In Fig. 1, we observe the decrease in transmit power via DUDe by comparing three cases via the simulation model. The first case is the baseline with a coupled downlink/uplink association, and no small cell bias. The need for bias arises in a HetNet scenario where, due to load balancing, the UEs are steered towards being associated to a small cell if their received power is lower (up to the bias value in dB) than the one received from the macro cell BS. The second case is still coupled, but the small cells have a 6 dB bias. We note that 6 dB has been shown in [10] to be a

reasonable value for the bias. The third case is for DUDe. DUDe yields 2.3 dB at 50% and 3dB at 95% CDF relative to the coupled association with 6 dB bias.

2.2   Improved Interference Conditions

DUDe also decreases the uplink interference, due to multiple complementary effects.

First, and as an obvious consequence of the transmit power reduction demonstrated in the previous section, the UL interference generated to other base stations is correspondingly reduced by about 2-3 dB. This is quite significant especially for the low SINR UEs in the uplink, since at low SINR in a dense network, decreasing the interference by 3dB implies an approximate doubling of data rate.

Second, DUDe provides the ability to independently select the association that minimizes interference at both the UE and the BS. Uplink interference in a given spectral band is an aggregation of many different UEs' transmissions in different cells, as received by a given BS, say BS0. The interference generated by each of these UEs depends on its location relative to its own desired BS, the amount of power control, its distance to BS0 and the uplink precoding weights. In contrast, the downlink interference at a given UE depends on the BSs' transmit power, the downlink beamforming weights and the distance to the different BSs. On top of this, the nearly independent scheduling and loading in each of the DL and UL causes further randomness in the interference. For all these reasons, average interference levels can be quite different in the downlink and uplink resources. Therefore, a decoupled association that allows the UE/network to seek out the best interference environment in the two links independently can be expected to substantially outperform a coupled association, which must "split the difference".

Third, DUDe will also prove a boon for Device-2-Device (D2D) communication which, as of 3GPP Rel. 12, will take place in the uplink bands. By lowering the UL transmit power and generating less interference, DUDe will create a more benign environment for D2D receivers and thus allow more D2D transmissions to take place.

Finally, in addition to reducing the amount of average interference, DUDe also allows a reduction of the uplink SINR variance, as shown in Fig. 2 (obtained via the simulation model) which translates into more efficient and effective UL schedulers and performance gains [11]. Specifically, with respect to a coupled association with a 6 dB bias (see Section 2.1) the decoupling yields a reduction of 1dB on average which is about 25% at 50% CDF.

2.3   Improved Uplink Data Rate

Unsurprisingly, increasing the desired received power and decreasing the interference leads to higher SINR, and hence a higher spectral efficiency and data rate. However, there are additional factors which can complicate the effect of DUDe on the uplink rate.

For example, consider an LTE HetNet with small cell range expansion and biasing. On average the optimal downlink bias is in the neighborhood of 5-10 dB as noted before, although with blanking or interference avoidance, up to 18-20 dB may be used in certain scenarios [2, 10]. DL biasing leads to a better association in both directions even with coupled association, since by expanding the DL small cell coverage region, more UEs associate with the nearby small cells in the UL as well, which is also the main point of DUDe.

Nevertheless, we still observe very substantial rate gains for DUDe even when compared with biased coupled associations. Detailed breakdowns of these rate gains in various configurations are given in [7] and [1], with our findings summarized in Table 2. Here, picocells have transmit power of 30 dBm while femtocells 20 dBm. The gains result mainly from the improved channel quality and also from the biasing as discussed above, which gives cell-edge (5th percentile) and median (50th percentile) UEs access to more resources which results in a UL higher rate. It is quite encouraging that two very different models and approaches to evaluating the rate gains both result in the conclusion that gains in the range of 100-200% are within reach, although the gains do erode somewhat with biasing since the baseline improves. Finally, we note that a recent paper based on optimization theory with a yet different model, also finds significant gains from DUDe [9].

Table 2: Summary of predicted uplink rate gains averaged over all UEs in the network, as a result of DUDe. Picocells have transmit power of 30 dBm while femtocells 20 dBm. We note that DUDe outperforms the baseline also when downlink biasing is used.

	DUDe vs picocells (Bias = 0)		DUDe vs picocells (Bias = 6 dB)		DUDe vs femtocells (Bias = 0)		DUDe vs femtocells (Bias = 8 dB)	
	Analysis	Simula-tion	Analysis	Simula-tion	Analysis	Simula-tion	Analysis	Simula-tion
5th percentile (cell-edge) rate gain	115%	90%	50%	30%	270%	260%	140%	95%
50th percentile (Median) rate gain	95%	150%	30%	60%	260%	230%	120%	180%

2.4 Different load balancing in the uplink and the downlink

The load that a given BS has in the UL may be different from the load that the same BS may have in the DL. This implies that is not optimal to have the same set of UEs connected to the same BS in both uplink and the downlink, such that at least some of the UEs should use decoupled access.

Additionally, DUDe allows pushing more UEs to under-utilized small cells in the UL only since it is not limited by interference as is the case in the DL. In Fig. 3 we show that this results in a

better distribution of the UEs among macro and small cells which, in turn, allows for a more efficient resource utilization and higher UL rates. We note that DUDe outperforms the baseline for both unbiased and biased association.

2.5 Low deployment costs with RAN centralization

Implementing a decoupled cell association in a real network requires excellent connectivity and modest cooperation between different base stations. As we will discuss in the subsequent section, the main requirement DUDe imposes is a low latency connection between the downlink and the uplink base stations, to allow fast exchange of control messages, like hybrid-HARQ messages. We emphasize that unlike the most sophisticated forms of Cooperative Multi-Point (CoMP), like joint processing, where a high throughput backhaul connection between BSs is required to allow rapid data exchange, DUDe does not impose a tight requirement on the backhaul capacity. Put another way, DUDe allows gains similar to uplink joint processing (about 100% edge and average throughput gain, as just seen), but with lower deployment costs. Compared to using MIMO or new spectrum to increase the throughput, the cost comparison is even more favorable to DUDe.

The ongoing trend towards using partial or full Radio Access Network (RAN) centralization in deployments where a high-speed backhaul is available, will be an enabler for downlink and uplink decoupling, as signaling will be routed to a central processing unit with low-latency connections. In particular, partial centralization refers to those local deployments (e.g. indoor) where the transmission points serving the same local area are all connected to the same baseband processing central unit. Full centralization, often referred as Cloud-RAN, extends this approach to larger areas, where a large number of RF units are connected to the same baseband processing central unit.

Given this already ongoing trend towards more centralized RAN architectures, which are underpinned by low-latency connectivity between BSs, the incremental cost of DUDe appears negligible in such scenarios.

3 DUDe in LTE-A: Enabling Architectures

DUDe can, depending on the deployment scenario and backhaul properties, already be supported by the existing LTE/LTE-A specifications. Illustrated in Fig. 4, three specific embodiments are discussed below.

3.1 Centralized Processing

As mentioned before, in a deployment scenario with multiple radio units with a different cell-ID connected to a centralized node (like in the case of a Centralized Radio Access Network C-RAN), DUDe is possible in LTE-A without additional standardization support (see Fig. 4a). The BS used for downlink transmission to a specific UE is selected using conventional means, typically based on downlink signal strength measurements. Uplink transmissions are received by one, or if macro diversity is desirable, multiple radio units as the specifications do not mandate the reception node. Uplink decoding could either be performed at the radio unit (or at the set of radio units) or sampled analog data could be forwarded to the centralized unit via a Common Public Radio Interface (CPRI) interface for further processing.

Uplink-related control signaling (including e.g. hybrid Automatic Repeat Request (ARQ) and power control commands) needs to be transmitted from the downlink node. In the same way, downlink-related control signaling from the terminal needs to be received by the uplink node and forwarded to the downlink node over the infrastructure.

3.2 Shared Cell-ID

An interesting extension of the approach described above is the so-called shared cell-ID approach [6] (see Fig. 4b), where radio units all belong to the *same* cell (i.e. have the same cell-ID). Here, Channel State Information (CSI) enhancements and quasi-co-location mechanisms introduced in Release 11 as part of the CoMP work are used to rapidly, independently and, from a terminal perspective, transparently switch transmission and reception points for a given terminal. This is a step away from the traditional cell-oriented paradigm towards viewing the antenna points as resources to be used in the best possible way to maximize performance. Furthermore, node association and mobility are handled via proprietary (non-standardized) solutions, transparent to the mobile terminal, and providing better mobility robustness in dense networks compared to methods relying on UE-centric measurements.

Although conceptually straightforward, both centralized processing and shared-ID approaches require a fairly low-latency backhaul to meet the timing requirements (e.g. to send hybrid-HARQ messages). In a practical LTE-A rollout, the deployment is thus limited to remote radio units connected to a centralized baseband processing node.

3.3 Dual Connectivity

While the two solutions described above require a very low-latency backhaul, usually achieved via connecting radio units to the same central unit, DUDe can also be implemented with a less ideal backhaul. *Dual Connectivity*, an extension first introduced in Release 12, allows for a terminal to be simultaneously connected to two cells and can be used for DUDe (Fig. 4c). We note that in Release 12, DUDe using dual connectivity is limited to inter-frequency deployments, i.e. to deployments where the two cells transmit over different frequency bands; nevertheless, later releases may add support for intra-frequency band deployments. The two cells operate separately, handling their own scheduling and control signaling (e.g. H-ARQ message) and thereby significantly relaxing the backhaul requirements compared to the centralized baseband approach and enabling the standardized X2 interface to be used for inter-BS communication.

This solution has advantages and disadvantages. On one hand, a low latency backhaul connection for the signaling is not needed. On the other hand, mobility must be handled using standardized mechanisms and the possibilities for proprietary optimization are limited.

4 DUDe in 5G and Beyond

The next few years will see intense research and development on 5G. The ITU is starting their work on requirements under "IMT-2020", and in 3GPP initial activities on 5G standardization are expected towards the end of 2015 with the overall goal of a large-scale trial around 2018 and commercial operation in 2020. Although any discussion of 5G is by definition speculative, there is an emerging consensus on the data rate requirements and likely key technical features of 5G, including extreme BS densification, massive MIMO, the introduction of millimeter wave bands, and possibly a "cell-less" architecture [5,12].

With this view of 5G, in this section we discuss whether 5G (and beyond) standards should include other features to *natively* support DUDe. In other words, we discuss whether a design that is optimized for DUDe from its inception, rather than amended *a posteriori*, could lead to even higher gains.

4.1 Major Architectural Changes?

An important question is whether a simple evolution of today's 3GPP architecture design discussed above would be able to efficiently support DUDe in emerging heterogeneous 5G deployments. In the previous section, we discussed how the LTE-A architecture already supports a DUDe implementation when different BSs are connected via fiber to the same radio unit. For the case of different base stations not connected to the same radio unit, we discussed how the support for DUDe in 4G is limited to different frequencies. Any future 5G releases in 3GPP should thus simply allow for a same-frequency dual connectivity, which – despite having implications on resource and interference management – is not considered to be a major upgrade.

A further tweak is needed to ensure proper encryption of all data and control channels, particularly when communication via the X2 interface is used between BSs. Whilst each eNB can support tens of IP Security (IPSec) tunnels, the management of security via IPsec is so cumbersome that operators tend to deploy only a few IPsec gateways (GWs) per country. Indeed, LTE has seen most IPsec GWs deployed close to the Serving GW (S-GW), which means that traffic logically going via the X2 is actually routed via the S-GW – this incurs a delay which renders DUDe inefficient. Whilst LTE-A enjoys some more IPsec GWs to be deployed closer to the mobile edge, future 5G designs ought to improve security mechanisms and implementations allowing the encryption of X2 traffic with lower latency.

Further, some integration work is needed with emerging paradigms which have proven useful for current coupled systems. First, as mentioned before, the integration of DUDe with the decoupled Control/Data-Plane and LAA will require some architecture modifications. Second, self-organizing networking (SON) paradigms will be instrumental in coordinating in a non-conflicting manner [13] the increased degrees of freedom in the system.

Given the above discussion, however, we conclude that a native support of DUDe does not require major design changes in 5G from an architectural perspective.

4.2   DUDe and Hyper-densification

The importance of decoupled selection of the DL/UL access points may grow significantly in the coming years, as 5G will feature hyper-dense deployments in order to meet the high rate demands in crowded spots. One could argue that at extremely high densities of cells, DUDe will lead to lower gains since nearly all the devices will be associated to the nearest small cell in uplink and downlink. However, this will only be true if we assume that all the small cells will have the same power, traffic and deployment characteristics. This is an unrealistic assumption, since future cellular deployments will be characterized by a mixture of user deployed and operator deployed cells, with different power levels, using frequencies ranging from below 1 GHz to tens of GHz, providing services for very different types of traffic and natively supporting device-to-device communications. Downlink and uplink traffic flows will be routed via a mixture of licensed and unlicensed carriers, requiring different allocation criteria[1]. Therefore, we expect that DUDe gains in future deployments will be even higher with respect to the ones presented above, especially if we consider the generalized version in which a UE is associated with multiple points and selects the DL or UL direction dynamically, as a part of a scheduling and optimization process.

From a broader perspective, we believe decoupled access necessarily shifts the focus of algorithmic solutions and optimizations towards models that consider two-way traffic from each UE. This is part of a larger trend in wireless network optimization that encompasses full-duplex communication, two-way relaying, and dynamic TDD.

4.3   TDD, FDD, or a New Way of Duplexing?

DUDe can work with both FDD and TDD, with different implications from a system-level perspective and from a spectrum-related perspective.

TDD allows much more flexibility in trading downlink and uplink resources as compared to FDD. With decoupling, as we have seen, fewer uplink resources are needed to achieve the same uplink rate versus the coupled case, and those resources could be reassigned to the downlink via dynamic TDD, which is in line with the two-way network optimization discussed above.

---

[1] Recently 3GPP approved a work item on License Assisted Access (LAA), where licensed and unlicensed carriers are aggregated. LAA uses licensed spectrum for control-related transmissions while sending data over both licensed and license-exempt carriers.

Traditionally, another benefit of TDD is the possibility of estimating the downlink channel via uplink reference signals. This is particularly important for channels with large dimensionality, such as with massive MIMO. Unfortunately, when DUDe is used, downlink and uplink transmissions originate and terminate at different locations, respectively, breaking the channel reciprocity. Much of the existing spectrum is paired FDD spectrum, so for both of these reasons massive MIMO may need to be supported without channel reciprocity.

In the medium/long term, DUDe together with different emerging technology trends could require rethinking the traditional FDD/TDD dichotomy. DUDe, hyper-densification, the use of higher-frequencies and highly directional antenna arrays, could enable duplexing approaches over the spatial domain. For example, the same band could be used for two different devices, one receiving in downlink from a base station, and the other one transmitting in uplink to another base station. Effectively, assuming an effective downlink/uplink spatially coordinated scheduling mechanism, this could allow full-duplex like performance leveraging on the spatial domain [14]. In addition, once analog/digital interference-cancellation mechanisms become truly operational to support full temporal duplex, the DUDe concept is also beneficial since the generalized decoupling would allow the support of a downlink and not necessarily the same uplink user in the same spectral band.

4.4   DUDe with Millimeter Wave Frequencies

Above we discussed why DUDe could make channel estimation via channel reciprocity in TDD more difficult. This effect could be even more pronounced at millimeter wave (mmW) frequencies, where the large number of antenna elements used for beamforming would be enabled by channel reciprocity.

However, there are other factors that point to DUDe as an important enabler for mmW. For example, recent studies on electromagnetic field exposure [15] show that to be compliant with applicable exposure limits at frequencies above 6 GHz, the maximum transmit power in the uplink might have to be several dB below the power levels used for current cellular technologies. Since the transmit power has an important impact on uplink coverage (in particular for sounding over a non-precoded channel) we believe a pragmatic approach would be to allocate uplink over a lower frequency with a better link budget. That is, while in the rest of this paper we discussed associating a UE to a macro cell in the downlink and to a small cell in the uplink, for mmW the opposite strategy might prove fruitful: associating the UE to the mmW small cell in the downlink and to a sub 6 GHz macro cell in the uplink.

5   Conclusions

In traditional cellular networks, it is practically an axiom that the uplink (UL) connection is always associated with the same Base Station (BS) that has been selected for downlink (DL) reception. In this paper we revisit this axiom and introduce the features of Downlink/Uplink Decoupling (DUDe), a new architectural paradigm where downlink and uplink are not constrained to be associated to the same BS. This is becoming especially relevant in the wake of the densification expected in future cellular networks, where each terminal has multiple access

points in proximity. We have identified five key arguments that demonstrate the usefulness of DUDe, based on a blend of theoretical, experimental, and architectural insights. We have shown how DUDe can lead to significant gains in network throughput, outage and power consumption at a much lower cost compared to other solutions that provide comparable or lower gains. We have discussed the changes needed in the existing LTE-A systems in order to enable DUDe-based operation. We have then presented arguments why DUDe should natively be considered as a part of the future 5G systems. Interestingly, major changes to the radio access and core networking technologies are not needed. DUDe can be considered an innovative approach that affects the fundamentals of cellular networks and thus opens up many opportunities for research and design.

References

[1] H. Elshaer, F. Boccardi, M. Dohler and R. Irmer. "Downlink and Uplink Decoupling: a Disruptive Architectural Design for 5G Networks," *in Proc. of IEEE Global Communications Conference (GLOBECOM)*, pp. 1798 – 1803, Dec. 2014.

[2] A. Damnjanovic, J. Montojo, Y. Wei, T. Ji, T. Luo, M. Vajapeyam, T. Yoo, O. Song, and D. Malladi, "A Survey on 3GPP Heterogeneous Networks," *IEEE Wireless Commun.,* vol. 18, no. 3, pp. 10 – 21, Jun. 2011.

[3] Y. Kishiyama, A. Benjebbour, T. Nakamura, and H. Ishii "Future Steps of LTE-A: Evolution Towards Integration of Local Area and Wide Area Systems," *IEEE Wireless Commun.*, vol. 20, no. 1, pp. 12 – 18, Feb. 2013.

[4] J. G. Andrews, "Seven Ways that HetNets Are a Cellular Paradigm Shift," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 136-44, Mar. 2013.

[5] F. Boccardi, R. W. Heath, Jr., A. Lozano, T. L. Marzetta, and P. Popovski, "Five Disruptive Technology Directions for 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 74 – 80, Feb. 2014.

[6] S. Parkvall, E. Dahlman, G. Jöngren, S. Landström and L. Lindbom, "Heterogeneous Network Deployments in LTE," *Ericsson Review*, Feb. 2011.

[7] S. Singh, X. Zhang, and J. G. Andrews, "Joint Rate and SINR Coverage Analysis for Decoupled Uplink-Downlink Biased Cell Associations in HetNets", *IEEE Trans. Wireless Commun*., in press.

[8] K. Smiljkovikj, P. Popovski, and L. Gavrilovska, "Analysis of the Decoupled Access for Downlink and Uplink in Wireless Heterogeneous Networks", *IEEE Wireless Commun. Lett.*, vol. 4, no. 2, pp. 173 – 176, Jan. 2015.

[9] H. Boostanimehr and V. Bhargava, "Joint Downlink and Uplink Aware Cell Association in HetNets with QoS Provisioning", *IEEE Trans. Wireless Commun.*, in press.

[10] S. Singh and J. G. Andrews, "Joint Resource Partitioning and Offloading in Heterogeneous Cellular Networks", *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, pp. 888-901, Feb. 2014.

[11] N. Abu-Ali, A.-E.M. Taha, M. Salah and H. Hassanein, "Uplink Scheduling in LTE and LTE-Advanced: Tutorial, Survey and Evaluation Framework," *IEEE Communications Surveys & Tutorials,* vol.16, no.3, pp.1239-1265, Aug. 2014.

[12] J. G. Andrews, S. Buzzi, W. Choi, S. Hanly, A. Lozano, A. Soong, J. C. Zhang, "What will 5G be?", IEEE Journal on Sel. Areas in Communications, Vol. 32, no. 6, pp. 1065-82, June 2014.

[13] H.Y. Lateef, A. Imran, M.A. Imran, L. Giupponi, M. Dohler, "LTE-Advanced Self-Organising Network Conflicts and Coordination Algorithms," *IEEE Wireless Communications Magazine*, in press.

[14] H. Thomsen, P. Popovski, E. de Carvalho, N. K. Pratas, D. M. Kim, and F. Boccardi, "CoMPflex: CoMP for In-Band Wireless Full Duplex", submitted to *IEEE Wireless Commun. Lett.*, Jul. 2015, Available: http://arxiv.org/abs/1507.08574.

[15] D. Colombi, B. Thors, and C. Tornevik, "Implications of EMF Exposure Limits on Output Power Levels for 5G Devices Above 6 GHz,", *IEEE Antennas and Wireless Propag. Lett,* vol. 14, pp. 1247 – 1249, Feb. 2015.

Biographies

Federico Boccardi (senior member, IEEE) is a Principal Technology Analyst in Ofcom. He received the M.Sc and Ph.D. degrees in Telecommunication Engineering from the University of Padova, Italy, in 2002 and 2007 respectively, and the Postgraduate Diploma in Strategy and Innovation from the Oxford Saïd Business School in 2014. Before joining Ofcom, he was with Bell Labs (Alcatel-Lucent) from 2006 to 2013 and with Vodafone R&D in 2014. He holds more than 60 between issued patents and peer-reviewed international research papers. He is an Associate Editor for IEEE Transactions on Wireless Communications and IEEE Transactions on Cognitive Communications and Networking. His current interests fall at the intersection between technology and strategy.

Jeffrey G. Andrews (S'98, M'02, SM'06, F'13) received the B.S. in Engineering from Harvey Mudd College in 1995, and the M.S. and Ph.D. in Electrical Engineering from Stanford University. He is the Cullen Trust Endowed Professor (#1) of ECE at the University of Texas at Austin and Editor-in-Chief of the IEEE Transactions on Wireless Communication. He developed Code Division Multiple Access systems at Qualcomm from 1995-97, and has consulted for entities including Apple, Samsung, Verizon, AT&T, the WiMAX Forum, Intel, Microsoft, Clearwire, Sprint, and NASA. Dr. Andrews is an elected member of the Board of Governors of the IEEE Information Theory Society, has been co-author of eleven best paper award recipients, including the 2014 IEEE Stephen O. Rice Prize and the 2014 IEEE Leonard G. Abraham Award, and he recently received the 2015 ASEE/Hewlett Packard Terman Award.

Hisham Elshaer (S'13) received his B.Sc. from Ain-Shams University in Cairo, Egypt in 2008. He received his M.Sc. in wireless systems from the royal institute of technology (KTH) in Sweden in 2012. He was with Ericsson R&D in Stockholm as master thesis student in 2011. He worked as a consultant in Tieto Sweden AB from 2012-2013. Since July 2013 he has been with Vodafone group R&D in the UK as an early stage researcher under the EU Crossfire project

where he is working on developing new 5G technologies. He is also studying towards a PHD in King's College London. Hisham has won the best paper award at the IEEE Globecom 2014 conference for his work on Downlink and Uplink Decoupling.

Mischa Dohler (S'01–M'03–SM'07-F'14) is full Professor in Wireless Communications at King's College London, Head of the Centre for Telecommunications Research, co-founder and member of the Board of Directors of the smart city pioneer Worldsensing, Fellow and Distinguished Lecturer of the IEEE, and Editor-in-Chief of the Wiley Transactions on Emerging Telecommunications Technologies and the EAI Transactions on the Internet of Things. He is a frequent keynote, panel and tutorial speaker. He has pioneered several research fields, contributed to numerous wireless broadband, IoT/M2M and cyber security standards, holds a dozen patents, organized and chaired numerous conferences, has more than 200 publications, and authored several books. He acts as policy, technology and entrepreneurship adviser. He has talked at TEDx and had coverage TV & radio.

Stefan Parkvall (senior member, IEEE) is currently a principal researcher at Ericsson Research working with research on 5G and future radio access. He is one of the key persons in the development of HSPA, LTE and LTE-Advanced radio access and has been deeply involved in 3GPP standardization for many years. Dr Parkvall is a senior member of the IEEE, served as an IEEE Distinguished lecturer 2011-2012, and is co-author of the popular books "3G Evolution – HSPA and LTE for Mobile Broadband", "HSPA evolution – the Fundamentals for Mobile Broadband", and "4G – LTE/LTE-Advanced for Mobile Broadband". He has numerous patents in the area of mobile communication. In 2005, he received the Ericsson "Inventor of the Year" award, in 2009 the Swedish government's *Major Technical Award* for his contributions to the success of HSPA, and in 2014 he and colleagues at Ericsson was nominated for the European Inventor Award, the most prestigious inventor award in Europe, for their contributions to LTE. Dr Parkvall received the Ph.D. degree in electrical engineering from the Royal Institute of Technology in 1996. His previous positions include assistant professor in communication theory at the Royal Institute of Technology, Stockholm, Sweden, and a visiting researcher at University of California, San Diego, USA.

PETAR POPOVSKI received his Dipl.-Ing. (1997) and Magister Ing. (2000) in communication engi- neering from Sts. Cyril and Methodius University, Skopje, Macedonia, and his Ph.D. from Aalborg University, Denmark (2004). He is currently a professor at Aalborg University. He is an editor for IEEE Transactions on Communications and has served in the past as an editor for IEEE Communications Letters, the IEEE JSAC Cognitive Radio Series, and IEEE Transactions on Wireless Communications. He is a Steering Committee member for the IEEE Internet of Things Journal. His research interests are in communication theory, wireless communications, and networking.

Sarabjot Singh (S'09, M'15) is a Research Scientist with Intel Corp., Santa Clara, CA, USA. He received the B. Tech. in Electronics and Communication Engineering from Indian Institute of Technology Guwahati, India, in 2010 and the M.S.E and Ph.D. in Electrical Engineering from University of Texas at Austin in 2013 and 2014 respectively. He has held industrial positions at Nokia in Berkeley CA;   Alcatel-Lucent Bell Labs in Crawford Hill, NJ; and Qualcomm Inc. in San Diego, CA. Dr. Singh is interested in the modeling, design, and analysis of heterogeneous

wireless networks with focus on novel MAC and PHY technologies. He was the recipient of the President of India Gold Medal in 2010 and the ICC Best Paper Award in 2013.
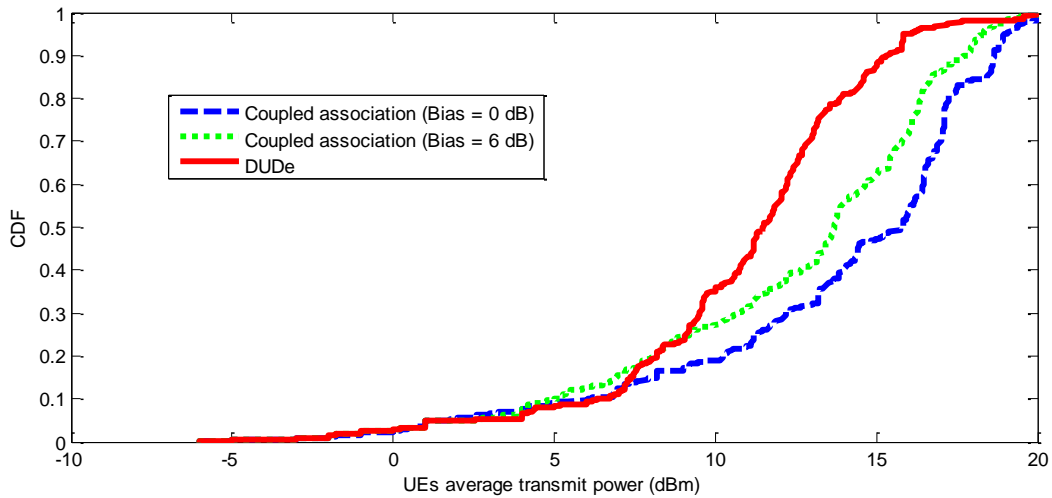

Figure 1. CDF of the UE's UL transmit power via the simulation model. Cell edge users (right side of figure) require higher transmit powers and thus achieve larger power reduction from DUDe.
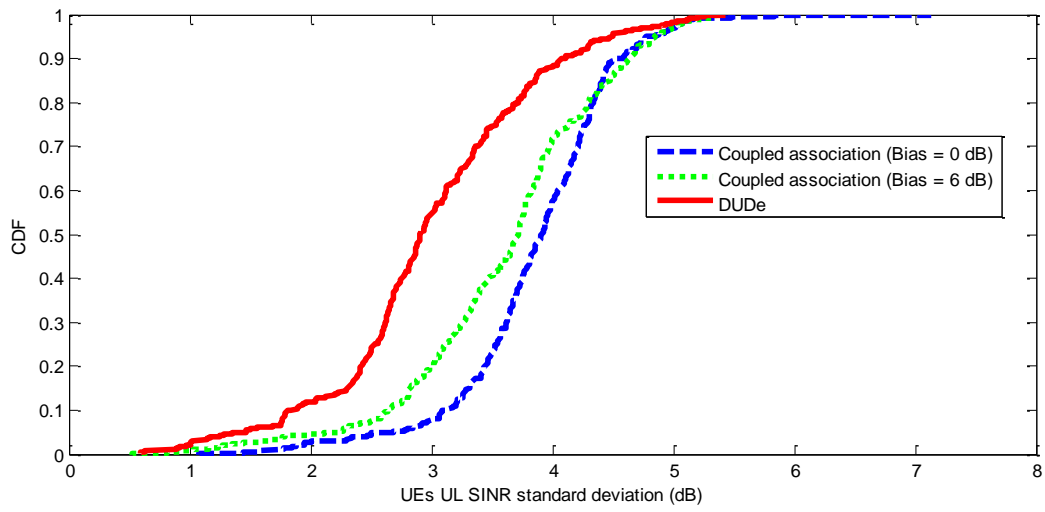
Figure 2. CDF of the UEs SINR standard deviation over time. DUDe reduces the variations and improves performance.
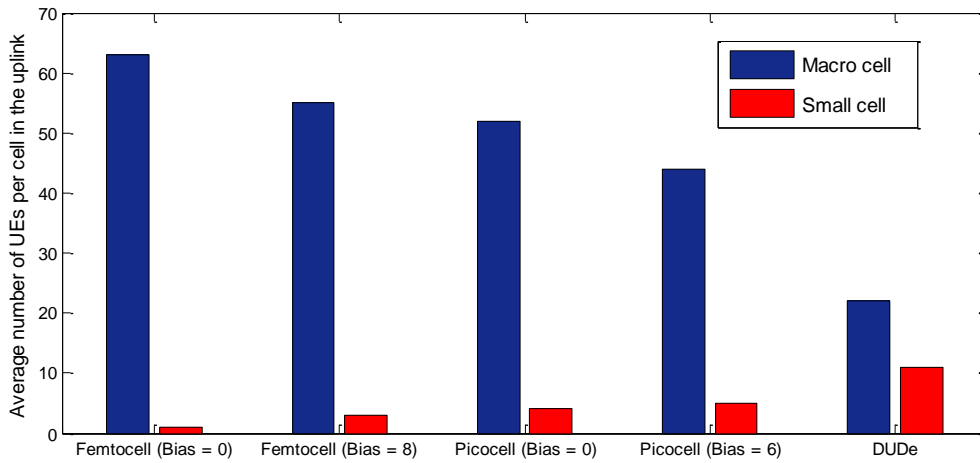
Figure 3. Average number of UEs per cell in the UL for the femtocell, picocell and DUDe cases. We note that the total number of UEs in the system is kept constant across the different cases (cf. Table 1).
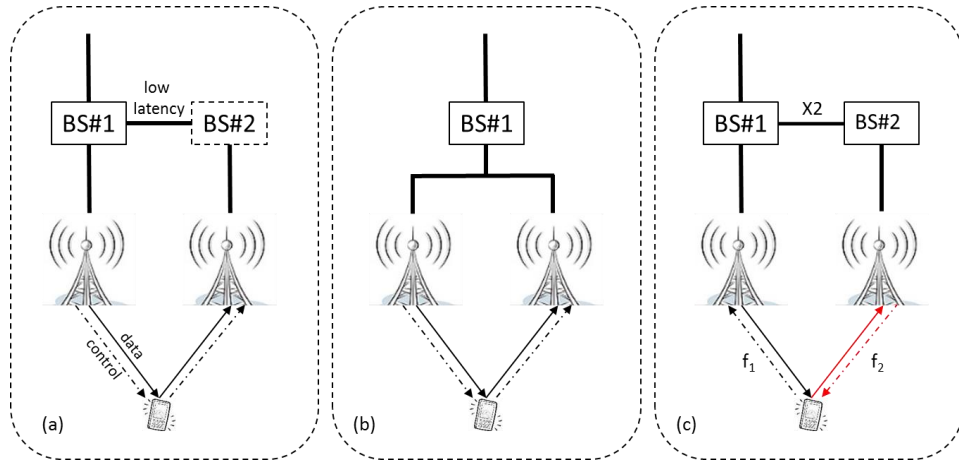
Figure 4. The three discussed embodiments of DUDe are: (a) centralized processing unit; (b) shared cell-ID; and (c) the dual connectivity option.