Mitigating Interference in LTE Networks

With Sequans AIR™ - Active Interference Rejection

sequans.com

SEQUANS COMMUNICATIONS

Contents

Executive summary

As LTE (long term evolution) networks proliferate and network traffic increases, interference is becoming an issue for LTE operators. Because LTE spectrum is limited, most operators are deploying single frequency networks to maximize capacity; however, while single frequency networks increase spectral efficiency, they also increase the potential for interference. Network-based interference mitigation solutions are specified in future versions of the LTE standard, but these are not yet available to address the interference problems of today's LTE networks, and will not remove all interference. Terminal-based interference solutions, however, are available today and they offer operators a powerful weapon to combat interference. LTE chipmaker Sequans Communications has introduced a terminal-based interference solution called Sequans Active Interference Rejection – or Sequans AIR™. Sequans AIR provides key benefits to both end-users and operators: end users will experience higher throughput and better service continuity, and LTE operators will improve coverage and increase the capacity of their networks.

Key benefits

- Up to 3.5 x throughput increase at cell-edge
- Up to 2 x network capacity increase*
- Improved user experience in dense deployments

**Assumes all user terminals equipped with Sequans AIR*

1. Introduction

1.1. LTE market

LTE is a 4G wireless technology standardized by the 3GPP (3G Partnership Project) that is being deployed today by leading operators around the world to provide high-speed data and multimedia services. The LTE market is growing rapidly. According to telecom research firm IDATE, there were more than 10 million LTE subscribers and 50 deployed LTE networks at the end of 2011, growing to an estimated 118 million subscribers and 200 networks in 2013. All leading operators are moving to LTE.

The LTE standard is an evolving standard with several planned releases. Most operators began deployment with Release 8, and are now in the process of moving to Release 9. LTE-Advanced networks based on Release 10 are expected to begin deploying near the end of 2013.

Table 1 – LTE standard versions

Version	Date	Main features
Rel 8	2008 Q4	First LTE release. All-IP network (SAE).
Rel 9	2009 Q4	SAE enhancements, WiMAX and LTE/UMTS interoperability.
Rel 10	2011 Q1	LTE Advanced fulfilling IMT Advanced 4G requirements. Backwards compatible with release 8 (LTE).
Rel 11	2012 Q3	Advanced IP interconnection of services.

1.2. Inter-cell interference in LTE networks

Due to limited spectrum resources, most operators are deploying their LTE networks in a frequency reuse =1 configuration, which means that a single carrier frequency is reused in all cells of the network. This deployment scheme is also referred to as a single-frequency network, and it is different from schemes used in predecessor cellular networks, where predefined planning ensured limited inter-cell interference. Single frequency networks are the most efficient in terms of overall spectral efficiency, but by nature they are limited by inter-cell interference. See Figure 1.

> **Single frequency LTE networks are by nature limited by inter-cell interference.**

Figure 1 – Inter-cell interference

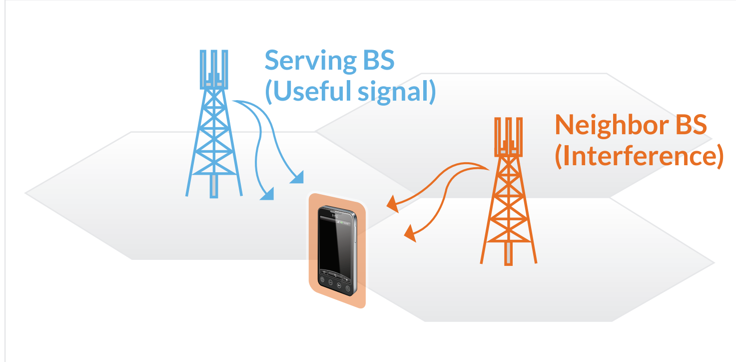

This is a well-known issue that has been the topic of many publications, such as Managing Interference in LTE Networks by Senza Fili Consulting [SENZA-FILI].

1.3. Impact of small cells

Data is replacing voice as the predominant application, and therefore overall capacity and continuity of service are becoming key concerns for operators as data hungry devices flood cellular networks.

As a means of increasing capacity, operators have been trending towards deploying small cells in a layered configuration, whereby a macro cell is deployed for coverage and several smaller overlapping pico or femto cells are deployed for capacity. The small cells can be turned off/on dynamically, depending on traffic, in order to save energy. See Figure 2.

Figure 2 - LTE small cell deployment

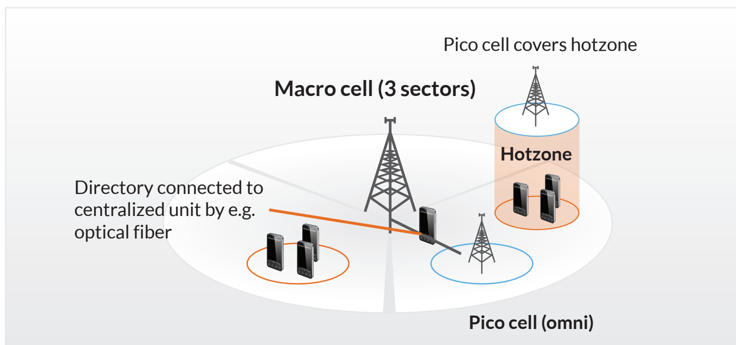

In these small cell deployment scenarios, interference can become greatly exacerbated.

1.4. Network-based interference management

The 3GPP is currently evaluating interference solutions. Future versions of the LTE standard (releases 10 and 11) will incorporate network-based interference management techniques. These techniques will help to manage interference, but will not be available before LTE-Advanced networks are deployed. Furthermore, these will only limit interference, not suppress it. Examples of such techniques are eICIC (enhanced inter-cell interference coordination) and CoMP (coordinated multipoint transmission and reception). These are well described in publications such as [SENZA-FILI].

1.5. Terminal-based Interference management

Due to the lack of availability and limitations of network-based interference solutions, terminal-based solutions that can be deployed now are gaining interest. These solutions must be able to operate on current LTE networks (Release 8 and 9), and be ready to operate on future LTE-Advanced networks (Release 10 and 11).

Terminals with embedded interference mitigation technology benefit users by providing:

- Overall higher throughput (especially at the cell edge)
- More stable performance across various locations in the cell
- Better continuity of service when moving across a network

Terminal-based interference solutions benefit operators by providing:

- Higher total capacity
- Improved coverage

These benefits will be described in detail in section 3.

2. Receiver design

2.1. Interference mitigation in LTE

There are numerous publications on the topic of interference cancelation techniques. For instance, [IRC-GSM] deals with interference mitigation for 2G systems, and the 3GPP has defined and standardized UE receiver classes related to interference cancelation capability for 3G systems (WCDMA) [TR25.963] for several receiver types as defined in table 2.

Table 2 - 3GPP reference receivers

3GPP Name		Reference receiver
Type 0	=	RAKE
Type 1	=	Diversity receiver (RAKE)
Type 2	=	Equalizer
Type 2i	=	Equalizer with interference awareness
Type 3	=	Diversity equalizer
Type 3i	=	Diversity equalizer with interference awareness
Type M	=	Multiple input multiple output (MIMO)

Similarly, the 3GPP has recently begun to investigate interference-aware receivers for Release 11, as described in [TR36.829].

While the theoretical aspects and algorithm principles of interference cancelation are well understood, numerous challenges lie in the implementation of these techniques in an LTE terminal:

- First, the LTE waveform is based on OFDMA modulation, which is not the same as 2G/3G modulation. New techniques have had to be developed because channel estimation and receiver design for the multi-carrier OFDMA modulation of LTE is very different from that used for the single-carrier and WCDMA modulation of 2G/3G. Practical implementation of interference mitigation theory in LTE requires intimate knowledge of OFDMA architecture.
- Second, LTE throughput is significantly higher than 2G/3G throughput, but the overall budget for power consumption is constrained in order to meet the requirements of battery-powered devices. Therefore, the implementation of interference mitigation techniques must be designed to require minimal hardware resources and consume minimal power. This goal can be achieved only by accounting for interference mitigation in the modem architecture from the initial design.
- Finally, the LTE standard has defined several transmission modes (see Table 3), from which specific interference mitigation techniques must be derived. Specific terminal feedback information transmitted to the eNodeB (for dynamic throughput optimization) must be taken into account in the design of interference mitigation techniques.

Table 3 – LTE transmit modes

TM	Description
1	Single transmit antenna
2	Transmit diversity
3	Open loop spatial multiplexing with cyclic delay diversity (CDD)
4	Closed loop spatial multiplexing
5	Multi-user MIMO
6	Closed loop spatial multiplexing using a single transmission layer
7	Beamforming
8	Dual-layer beamforming

2.2. Interference mitigation techniques

There are two possible approaches to implementing interference mitigation in a terminal receiver:

- Nonlinear: In this approach, the interfering signal is estimated and then subtracted from the received signal, possibly in an iterative manner. This requires explicit modeling of the interfering signal. Such an approach provides excellent performance, but is very sensitive to errors in the estimation of the interfering signal.

Figure 3 – Nonlinear interference cancelation

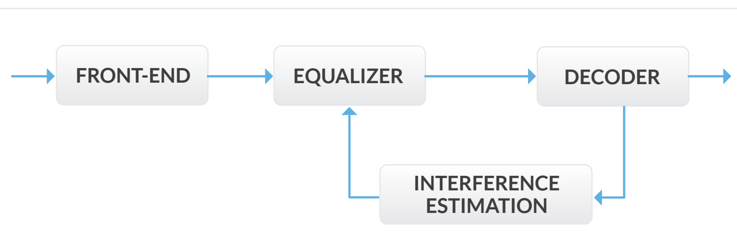

- Linear: In this approach, the receiver uses multiple antennas to perform spatial suppression of the interfering signal. Specifically, the receiver forms a receive antenna beam, with a spatial null in the direction of the interferer. This works best with a large number of receive antennas, but provided with proper spatial properties, this technique can handle a number of interferers. Such receivers are usually called IRC (interference rejection combiner) receivers.

Figure 4 - Receive beamforming

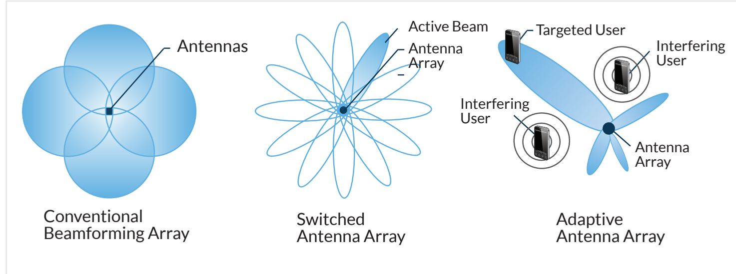

Note that for both of these approaches, channel estimation, whereby the channel and interference are accurately estimated, is a key step in the process. Another key step is to detect the presence or absence of interference. This eases overall processing in the UE, considering that interference is highly unpredictable and dependent on variable factors such as channel conditions, traffic from other terminals, and scheduling from the eNodeB.

2.3. Introducing Sequans AIR

Based on the specific requirements of interference mitigation in LTE, and considering the rapidly changing interference conditions in packet-switched networks, Sequans has designed Sequans AIR (active interference rejection), a compact LTE receiver that includes interference mitigation capability, suited to the various transmission modes of LTE. Sequans AIR adopts the linear approach to interference mitigation. It has been co-developed with technology partner ArrayComm, a pioneer in antenna processing and interference management techniques. Sequans has leveraged its own expertise with OFDMA and MIMO receivers, and the combined efforts of Sequans and ArrayComm have resulted in an innovative and powerful interference mitigation algorithm and an optimized implementation on silicon.

Sequans AIR mitigates interference not only from data channels, but also from control channels.

Sequans AIR has been designed to mitigate the interference not only from data channels but also from control channels. Even though control channels are designed to be more robust than data channels, they may also suffer from strong interference. If they do, the terminal may not be able to demodulate the control channel and may lose its connection to the network.

2.4. Support on Sequans products

Sequans AIR is designed for use on Sequans' latest LTE platforms:

- Andromeda (based on SQN3110 baseband IC) for handsets and tablets
- Mont-Blanc (based on SQN3120 baseband IC) for dongles, mobile hotspots, M2M applications, and other data-centric devices.

Sequans AIR is designed to work in any LTE network, regardless of eNodeB vendor, carrier frequency, channel bandwidth, or duplexing scheme (TDD/FDD).

Sequans will further enhance Sequans AIR on future generations of its products to support LTE Releases 10 and 11 and to cope with potential new interference scenarios.

3. Performance results

In this section, we present performance results of the Sequans AIR receiver, based on simulations at the link level and system level:

- The link performance data provides good information about the receiver performance in a range of interference conditions (from noise-limited to interference-limited). These results do not directly translate to show the benefits of Sequans AIR in a real system.
- The system level results were obtained through partnership with SIRADEL, a leading provider of advanced RF tools, using realistic geographical data. These results clearly demonstrate the benefits of Sequans AIR in real operational deployment conditions.

3.1. Link level performance

In order to evaluate the benefit of Sequans AIR in the receiver, the Sequans AIR algorithm was implemented in Sequans' LTE simulator, which is bit accurate and can represent the true performance of the chip. Only downlink is considered in these simulations. The first two simulations assume an ideal link adaptation where the best MCS (modulation and coding scheme) per SNR (signal to noise ratio) point is selected independently for a Sequans AIR receiver and a reference implementation for an industry-standard MRC (maximum ratio combining) receiver. The channel considered is the extended-vehicular-A channel in low mobility as defined by the 3GPP. The carrier frequency is 2.6 GHz and we assume cell planning such that the interfering cell(s) and the serving cell reference signals are non-overlapping. We assume up to three interfering cells in the link layer simulation.

The Sequans AIR receiver yields 3.5 times more throughput than the default MRC receiver.

With respect to the interference profile, we assume that the interference consists of data transmitted in either the same transmit mode as the useful data from the serving eNodeB, or using a different transmit mode. In all cases, the downlink sub-frames are fully allocated, from both the serving and interfering sides. In the case of the PDCCH (physical downlink control channel), we assume for the sake of simulation that the serving and interfering eNodeBs consider the same aggregation level.

The next figures illustrate the performance of the AIR receiver compared to a reference MRC receiver. In Figure 5, we consider a single interferer, having a constant C/I (carrier to interferer) ratio of –3dB. This means that the power level of the interferer is twice the power level of the useful signal. The useful signal is using TM2 (the most robust way to transmit information within the various transmission modes), while the interferer is using TM1.

Figure 5 - Single-interferer link-level PDSCH performance

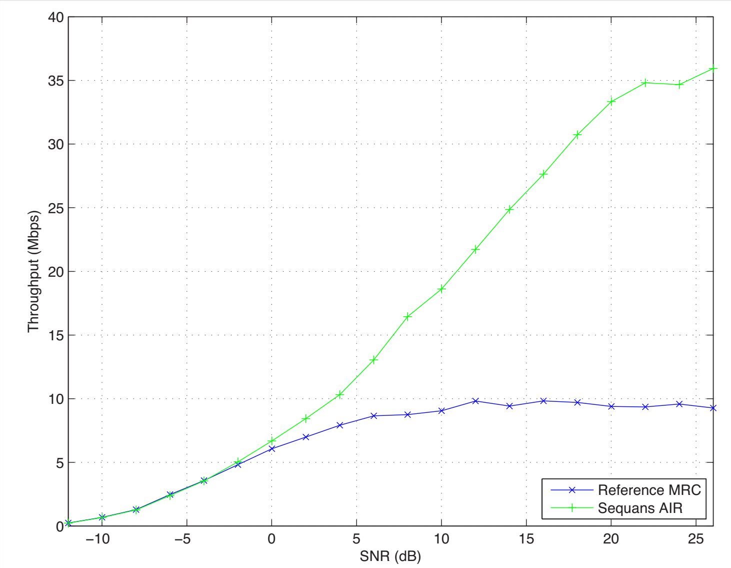

In this scenario, the MRC receiver has a throughput floor of about 10 Mb/s while the AIR receiver yields much higher throughput up to 35 Mb/s or about 350 percent of the reference MRC. In this scenario, even with a good SNR, the performance is interference-limited for the MRC and the AIR receiver therefore provides much higher throughput.

Figure 6 presents a very challenging scenario with three interferers. The first interferer has the same power as the serving cell while the second has power 3dB below it and the third, 6dB below. Even in this challenging scenario that requires rejecting interference from three interferers with only two UE antennas, the AIR receiver provides about 35 percent higher throughput than the reference MRC receiver.

Figure 6 - 3-interferer link-level PDSCH performance

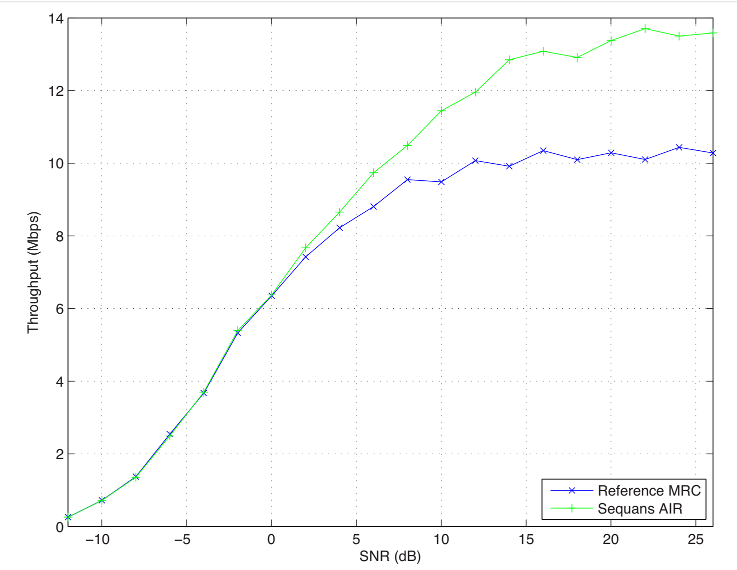

The two previous figures assume a constant level of interferer(s) power versus the serving cell(s) power, with a varying SNR. Another way to illustrate the link layer performance is to present the performance assuming a given SNR (depending typically on the distance from the receiver to the base station and its environment), and a varying C/I.

In Figure 7, there are two interferers both using TM2 (transmit diversity), effectively making four interfering eNB antenna ports that must be rejected with only two UE antennas. Both interferers have the same power, but their power levels vary while each maintains the desired signal 15dB above the noise floor. The Sequans AIR receiver shows a gain of 2-3 dB above the reference MRC receiver, demonstrating that for a given performance the AIR receiver can handle nearly twice as much interference as the reference MRC receiver.

Figure 7 - Dual-interferer link-level PDSCH performance

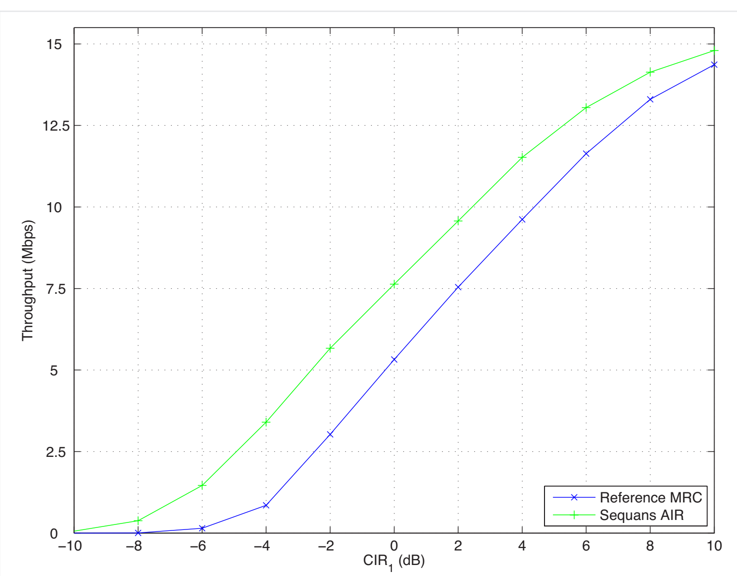

For a given performance, the Sequans AIR receiver can handle nearly twice as much interference.

The final performance curve in Figure 8 shows the effect of using the AIR receiver on the PDCCH control channel with an aggregation level of two, and a single interferer at a C/I = -3dB. In this scenario, the reference MRC fails to maintain the link since it cannot decode the PDCCH 35 percent of the time. The AIR receiver, on the other hand, operates well below a 1 percent error level, already at a low SNR, which means that it rejects the interference and receives the control information correctly.

This document presents only a few representative scenarios of the benefits of Sequans AIR, whereby the Sequans AIR receiver demonstrates a clear gain over the MRC receiver, even for one of the most challenging cases where the desired signal is interfered by many interferers and there are only two antennas at the UE. For both the traffic channel, PDSCH (physical downlink shared channel), and the control channel, PDCCH, large gains over a reference MRC receiver were observed, showing that a Sequans AIR mobile device would be able to decode the control channel and maintain connectivity, while the reference MRC mobile device would be disconnected.

Figure 8 - Single-interferer link-level PDSCH performance

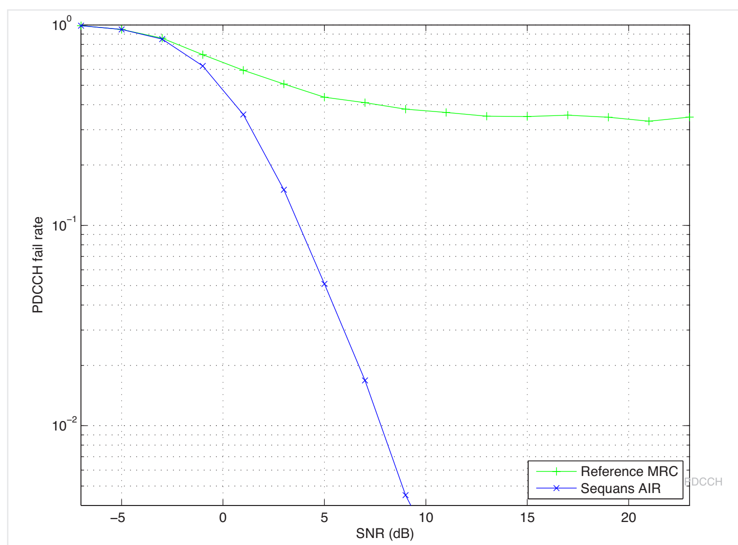

The Sequans AIR receiver demonstrates a clear gain over the MRC receiver, even in the most challenging situations.

Although link level performance can help to characterize the performance of a given receiver, it does not highlight the promise of having a network deployed with such receivers. The next section addresses this by looking at system level performance.

3.2. System level performance

A typical urban LTE deployment has been simulated, using central Paris as an example. The zone of interest covers a 1 km^2 zone, with macro base stations deployed. Base stations outside the zone of interest are included to generate an accurate interference pattern inside the zone.

Figure 9 - Paris VI simulation zone

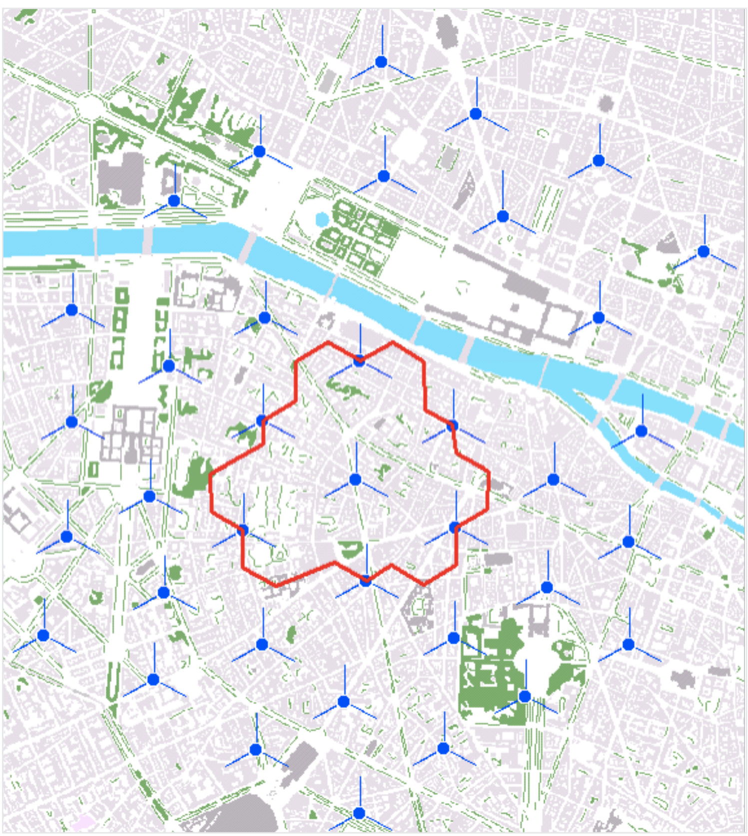

For the simulation, a typical hexagonal configuration was used, as described in Table 4.

Table 4 – System-level parameters

Modeling	Real environment (high-resolution geo map data.) Propagation model: Ray-based Volcano model.
Study area	0.87 km² corresponding to 12 cells.
System	LTE FDD - 2 x 10 MHz. Central frequency: 2.6 GHz. Transmission mode 3 (SU-MIMO) with 2 layers.
Macro-cell layout	Three sectors per site. Hexagonal site deployment: three rings around the central site, i.e. 37 sites corresponding to 114 cells. Average inter-site distance (ISD): 500 m. Average antenna height: 32 m above ground. Maximum transmit power per antenna: 46 dBm. Antenna: directional, 14 dBi gain. Antenna electric down-tilt: 6°. Number of antennas per sector: 2.
User	UE antenna heights: 1.5 m for outdoor UEs; 1.5 m, 13.5 m and 25.5 m above ground for indoor UEs. UE antenna: omni-directional, 0 dBi. Number of antennas: 2. UE noise figure: 9 dB.
Traffic	Downlink traffic load: -0% (no interference); -50%; -100%.

Three scenarios were simulated:

- One scenario without interference (meaning the neighboring cells have no traffic on downlink data channel).
- Two scenarios with, respectively, 50 percent and 100 percent downlink traffic loads.
  This traffic load represents the average portion of signal resources allocated to the cell users. The MAC layer abstraction does not consider any network-based interference mitigation technique, thus the set of resources allocated by each cell is viewed as random and independent.

We assume that the network is deployed at 2.6 GHz, using 10 MHz of bandwidth, and operating in TM3. In this simulation, we made the conservative assumption that TM3 was restricted to 1-layer transmission. Other simulation parameters used are given in Table 4.

To derive the throughput values, both outdoors and indoors, the system simulation [SIR] takes into account the SNR, CQI (channel quality indicator) and throughput values based on the properties of the receiver – the default MRC receiver or the Sequans AIR receiver (reusing link-level simulation results). The peak user data rate maps represent the net data rate experienced by a single user in the cell benefiting from the entire bandwidth (50 resource blocks).

For indoor performance evaluation, a specific outdoor-to-indoor channel model was considered using high-resolution 3D map data.

Figure 10 presents the interference-free case for the outdoor users. In this case, both the Sequans AIR and the MRC receivers display optimal performance. It is quite interesting to see that with this typical macro deployment, usually dimensioned to offer capacity, the network is not noise-limited outdoors (the maximum throughput is obtained almost everywhere).

Figure 10 - Outdoor coverage without interference

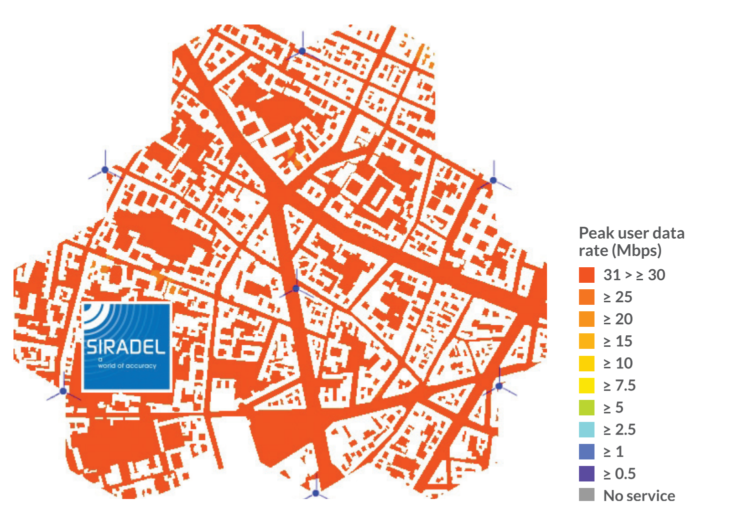

The performance for indoor users is quite different because of penetration losses. Figure 11 illustrates the performance of indoor users at ground floor. In this case, the maximum data rate is not achieved, as there is no use of small indoor cells.

Figure 11 - Indoor coverage without interference at ground floor

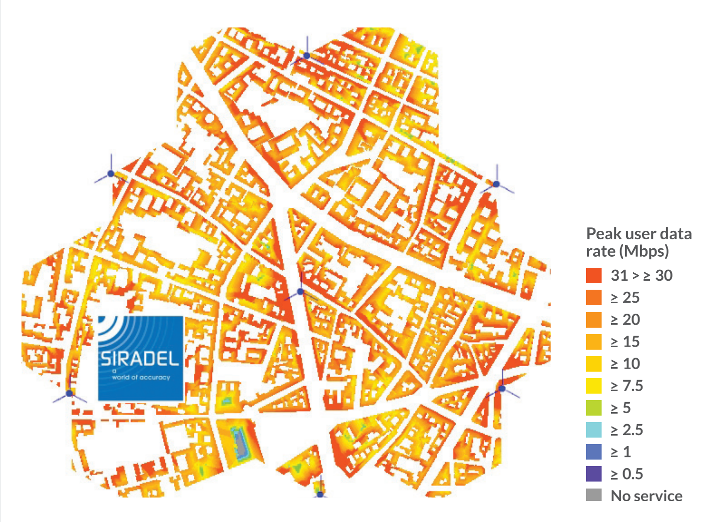

The two previous figures depict scenarios with no interference, and thus represent the performance boundary of a perfect interference cancelation receiver.

Now let us consider real-life scenarios with interference. Figure 12 shows the coverage map for the default MRC receiver and the Sequans AIR receiver.

Figure 12 - Outdoor coverage with 100% interference

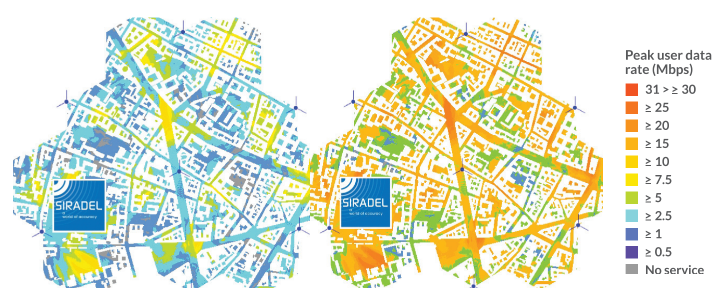

(a) - Default MRC receiver **(b) - Sequans AIR receiver**

When comparing the results of the full interference case with the ideal interference-free case, the following observations can be made:

- The network is interference-limited. Even at good SNR levels, the throughput drops considerably with the default receiver as compared to the Sequans AIR receiver.
- A standard receiver may not be able to connect even in a deployment that was designed for capacity (i.e. over-dimensioned with respect to coverage).
- A receiver able to mitigate interference can recover most of the degradation in a realistic deployment.

A similar result may be seen for indoor users as illustrated in Figure 13, although the number of areas with no service is even larger (reflecting the cumulative effects of interference and the loss of signal strength). However, in the deep indoor areas where the SNR is much affected by the indoor penetration losses, the performance of the default receiver and the Sequans AIR receiver become close. In this case communication is no longer interference-limited, but noise-limited.

Figure 13 - Indoor coverage with 100% interference at ground floor

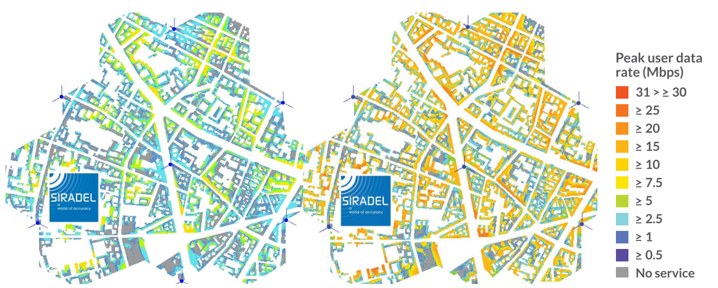

(a) - Default MRC receiver (b) - Sequans AIR receiver

Finally, in a scenario with a lower level of interference (neighboring cells have a traffic load of 50 percent), the relative gain of the Sequans AIR receiver compared to the default MRC receiver is lower than in the full interference scenario, as illustrated in Figure 14. Nonetheless, the Sequans AIR receiver is able to recover the interference-free performance, except in the few areas with a very low C/I.

Figure 14 - Outdoor coverage with 50% interference

(a) - Default MRC receiver (b) - Sequans AIR receiver

4. Conclusion

Interference is a key issue in LTE networks. Solutions implemented at the terminal side provide key benefits, both for end users and LTE operators. Sequans AIR is a solution that has been designed to fit on Sequans' chipset architecture, with proper hardware accelerators to enable full line-rate performance. Sequans AIR works in both TDD and FDD modes for all of the transmission modes defined in the LTE standard. The benefits of Sequans AIR have been proven, both at link-level and system-level.

5. Acknowledgements

Sequans wishes to thank its key technology partners, ArrayComm and Siradel.

5.1. ArrayComm

ArrayComm is a provider of physical layer solutions for wireless infrastructure and client device applications. ArrayComm is a world leader in multi-antenna signal processing, delivering commercial A-MAS™ software now that combines MIMO, beamforming, and interference cancelation to improve end user experience and radio network economics through gains in coverage, client data rates, and system capacity. The company's comprehensive and flexible PHY solutions include optimized DSP software and hardware accelerators that save development costs and time-to-market.
www.arraycomm.com

5.2. Siradel

SIRADEL (www.siradel.com) is a high-tech company (small-medium enterprise) created in 1994 and based in France, China (Hong-Kong) and Canada (Toronto). Siradel provides products and services for the ICT Industry in particular wireless telecommunications.

The portfolio of the company is composed of:

- 3D GIS data and RF measurements
- Advanced RF tools (Volcano, VolcanoLab)
- Management and technology consulting.

More than 50 people work at Siradel, serving more than 250 customers in about 50 countries. Siradel's solution brings more reliable and realistic assessments of wireless network and wireless equipment performance. The profile of its customers is diverse and includes wireless carriers, radio access equipment companies, manufacturers, regulatory bodies, utilities, and consultants.

6. Acronyms

3GPP	3rd Generation Partnership Project
CoMP	Coordinated multi-point transmission and reception
C/I	Carrier to interference ratio
CDD	Cyclic delay diversity
CQI	Channel quality indicator
DL	Downlink
eICIC	Enhanced inter-cell interference coordination
LTE	Long term evolution
MCS	Modulation and coding scheme
MRC	Maximum ratio combining
OFDMA	Orthogonal frequency division multiple access
PDCCH	Physical downlink control channel
PDSCH	Physical downlink shared channel
SNR	Signal to noise ratio
UE	User equipment (terminal)
WCDMA	Wideband code division multiple access

7. References

[IRC-GSM] J. Karlsson, J. Heinegkd, "Interference rejection combining for GSM", Proc. of 5th IEEE International Conference Universal Personal Communications, pp.433-437, Sep.1996

[SENZA-FILI] Senza Fili Consulting, "Managing Interference in LTE", April 2012

[TR25.963] 3GPP TR 25.963: Feasibility study on interference cancellation for UTRA FDD User Equipment (UE)

[TR36.829] 3GPP TR 36.829: Enhanced performance requirement for LTE User Equipment (UE) (Release 11).

[SIR] F. Letourneux, Y. Corre, E. Suteau, and Y. Lostanien, 3D performance analysis of a heterogeneous LTE network with urban femto-cells, COST IC1004 + iPLAN Joint Workshop on Small Cell Cooperative Communications, May 2012, Lyon, France.

SEQUANS COMMUNICATIONS