Technical Report
Wireless Control and Optimisation

Milosz Hulboj & Vlad Lapadatescu
{mhulboj,vlad}@cern.ch

CERN — HPN openlab project
CH-1211, Genèva 23, Switzerland
http://cern.ch

October 5, 2010

1 Introduction

Nowadays, wireless networks are becoming ubiquitous, and the density of wireless networks is increasing. Compared to wired networks, wireless installations pose *many* more challenges due to the nature of the radio-wave propagation and limited spectrum available.

CERN provides a particularly challenging environment for wireless networks. The *demand* for wireless access is growing and so are the user *needs*. There is a wide range of different scenarios, *e.g.* conference rooms, auditoria, long hallways, offices but also warehouse-like buildings and assembly halls. Answering the expectations of such a large and diverse user community is a real challenge.

A new openlab team under the codename `WIND` (Wireless Infrastructure Network Deployment) was founded to carry out a research activity and provide new algorithms, guidelines and solutions that will support the deployment and operation of the Wi-Fi infrastructure at CERN.

This technical report attempts to provide a comprehensive overview of the recent advances in the modelling, simulation, deployment, monitoring and operation of the wireless networks. The presented selection is by no means exhaustive, we deliberately focused on the subjects that are relevant to control and optimisation of corporate wireless networks, while disregarding on purpose much of the research related to *e.g. ad hoc* routing and *mesh* networks.

2 Prologue

2.1 *Root* of all evil

Virtually *all* the challenges with the wireless networks are directly or indirectly related with the radio wave propagation (RF propagation). One has far less control over this process compared to working with the wired media.

Wireless networks can only operate with a limited power within a designated part of the electromagnetic spectrum. There are other applications that might be using the same frequencies and interfere with the network operation (*e.g.* microwave ovens, Bluetooth users or other networks).

Most of the wireless networks operate in an *indoor* environment. In such cases it is difficult to describe the RF propagation. The main problems comes from the great *variability* in the building layout, materials and even furniture. There is no such thing as a *"standard building"*. What is more, even seemingly *minutiæ* changes can greatly affect the propagation parameters (*e.g.* closing of a fire door, people gathering for a meeting, closing of the window shades).

But this is not the only challenge originating from the nature of wireless propagation. Interferences define the spatial boundaries for spectrum reuse, and have direct impact on the assignment network capacity and end-user experience. Unfortunately packet delivery under interference is poorly understood for real networks. There are three independent variables that can help *mitigate* the interference: access point placement, power management and channel management [1].

Some of the challenges stem from the 802.11 standard *itself*. Carrier sense relies on channel measurements at the sender to infer the probability of reception at the receiver. In some cases the correlation between channel conditions at the sender and those at the receiver is very *weak* [2, 3]. Initial version of standard did not include any form of closed-loop power control nor any sophisticated monitoring features.

2.2 Fathom the *unfathomable*

Modelling and simulating wireless networks is a *formidable* task. As it was mentioned, even apparently simple factors can have a significant impact (*e.g.* the inner structure of wall and floor panels [4]). Still, it is rather unlikely that anyone will try to model each single brick type as described in [5].

Invalid data entered into a simulation will result in a nonsensical output data according to the `GIGO`[1] principle. Because of that in most of the cases simulations are only used to *aid* the coverage planning. With some pessimistic model assumptions they can deliver results that are *sufficient* for many deployments. Section 3 describes the problems of modelling and simulations in more detail.

2.3 Running the show

Even if one carried out the detailed measurements of the given site one would get only a static image at the point of measurement. Site surveys do not have unlimited accuracy. There is *no guarantee* that the propagation conditions will remain anything close to what has been observed. This means that one *cannot* just deploy a wireless network and forget about it

Monitoring and troubleshooting are two crucial tasks for running a healthy wireless network. Regrettably monitoring was not a major concern at the time the standards were being shaped. Commonplace access points provide limited monitoring facilities and there is no direct way to see how the network is being perceived by the clients. New standards (802.11h, 802.11k and the 802.11v proposal) introduce some mechanisms to alleviate the situation, but they do not address the problem of legacy clients. One could imagine deploying special probes dedicated to monitoring that would fill in the missing information [6, 7, 8].

The troubleshooting of a wireless network is equally problematic. Without detailed information it is a daunting task depending on a gut feeling — much the same like a dowser who does the random search by following a divining rod. To put it briefly – the realm is complex, the fractional information is available *post factum* and the methodology is not mature enough.

It should be then of no surprise that these two tasks are being neglected in many deployments. Section 4 describes some of the approaches to wireless network operation.

2.4 Bow to the inevitable

Not so long ago, in the 1997, we saw the *dawn* of the wireless local area networks. At that time they were slow, expensive and the technology was still in its infancy. Not long after, faster standards appeared and the *critical mass* for wireless popularisation had been reached. The user demands are growing, however they remain blissfully *ignorant* of the intricacies of wireless networks. The famous panacea of *"adding of more access points"* is not a solution, in some cases it can even degrade the performance.

And what is happening now? We have left 802.11b in the dust, 802.11g is commonplace and 802.11n is quickly gaining grounds. More and more appliances are using wireless and the number of different applications is growing. Companies are beginning to treat wireless networks as an important asset. The attempts to move from *"Best Effort"* to *"Mission Critical"* have already started. An approach to effectively handle large and dense installations is needed.

2.5 The nuts and bolts

This report attempts to shed some light on the recent research and developments in the area of wireless networks deployment, control and monitoring. It is important to note that our project would like to build upon the existing findings when possible, as it does not make much sense to *"re-invent the wheel"*.

[1] Garbage In, Garbage Out – a pun on the phrase First-In, First-Out.

3 Propagation models

As already mentioned, most (if not all) of the wireless network problems come from the issues inherent to wireless propagation. It comes as no surprise that RF planning and modelling is a first step in deploying a wireless network. After all, this step gives a *"confidence-boost"* to those deploying the network.

One has to note the clear difference between two types of simulation scenarios: indoor and outdoor. Unlike mobile telephony, in case of 802.11 networks, one is primarily concerned with the *indoor* aspect of the propagation process. This poses a completely new set of challenges compared to the outdoor environment. The complexity increases *many-fold*: almost each building is different, the environment is highly variable and the NLOS[2] propagation is dominant.

The models used in outdoor simulations are static in nature at reasonable time scales. This allows for accurate simulation results. However such an approach will *not* work in an indoor case. There are two different approaches to indoor simulations.

3.1 Modelling and simulation philosophies

In the *"blueprint"* based methods it is necessary to have the detailed plans of the building layout, furniture, electrical properties of the elements that are going to be part of a simulation (*e.g.* walls, windows). The risk here is twofold. First, the results of the simulations are not robust with respect to model specification. Minutiæ changes (*i.e.* displacement of fraction of wavelength) can seriously alter the results. Secondly, indoor environment is anything but static. The building geometry changes all the time, with mundane actions like opening the doors, moving the furniture or even simple fact of people moving around. The *"blueprint"* approaches tend to be computationally complex and it is not possible to conduct the simulation in real-time to respond to changes in the environment. Such methods are usually being used in the deployment phase.

The other methods accept up-front that it is difficult to get good accuracy in real-time. They use simplifying assumptions, work with more generic concepts and do *not* need detailed modelling. Such methods tend to be far less accurate than the *"blueprint"* ones, but can provide the information in real-time. These methods are better adapted to the dynamic indoor scenarios and are being employed by the monitoring and control systems for wireless networks.

3.2 Peculiarities of indoor propagation

The simulation domain also becomes more complex indoors. The wireless access point is not usually in line of sight (LOS) of the receivers and thus simple path loss[3] formulae are no longer sufficient. With many more potential sources of interference indoors, designing a reliable wireless network is not a trivial task.

Reflection is one of the most basic physical effects. It occurs when wave impinges on the obstacle whose dimensions are considerably *larger* than the wavelength. Reflected components can carry a considerable part of the wave's energy into locations not within line-of-sight. They contribute to the multipath effect.

Diffraction represents the bending of a wave around an obstacle whose dimensions are considerably *larger* than the wavelength. Similarly to reflections it accounts for directing part of the wave's energy into NLOS areas.

Absorption is a measure of how much of the wave's energy has been lost when passing though a medium. The amount of energy lost is dependent on the material of which an electromagnetic wave passes and also on the wave's frequency.

[2]Non-Line-Of-Sight or near-line-of-sight is a term used to describe radio transmission across a path that is partially obstructed. In such cases phenomena like diffraction, refraction, scattering and multipath become central.

[3]Path loss refers to the radio propagation losses caused by the attenuation of the radio wave as it moves through free space, as it is absorbed or diffracted by obstacles and also includes losses caused by other phenomena.

Scattering occurs when the wave hits obstacles whose dimensions are *comparable* to the wavelength. They will each reflect part of the incoming wave into a great number of directions, each carrying some of the wave's energy with it. As scattering deals with obstacles comparable in size to the wavelength, it tends to be *difficult* to model and simulate (even a displacement of a fraction of wavelength can change the scattering pattern).

Wave guiding occurs when the radio wave travels along a long corridor, being reflected back and forth along the walls, following the corridor's path and even taking turns, distributing an important amount of energy to remote locations along that hallway.

Multipath effect is a result of one or more of the mentioned phenomena. It happens when coherent signals arrive at the receiver at different times due to difference in path length. This causes the fading effect – if the signals are in phase, they would amplify the resulting signal, otherwise they will attenuate the resulting signal. 802.11n standard takes advantage of multipath propagation by employing multiple antennas.

Attempts to model and simulate indoor propagation is in itself a Sisyphean endeavour. And, withal, a significant number of researchers stand in defiance trying to find a *robust* solution. This section presents[4] four major methods that are prevalent in the scientific literature to this date.

Although old, [9], [10] present a general overview of the propagation models available. Newer overviews can also be found in [11] or [12]. The following section will present four major approaches that were found in the scientific literature to date.

3.3 Empirical approaches

Empirical models are based on statistical propagation models. Although built around databases containing exhaustive measurements, they still lack the accuracy one would demand in an indoor scenario. This is because a lot of physical events cannot be taken into account by this type of model.

There are several models that have been created, ranging from very simple algorithms, to more complex ones. [13] offers a comparison between different types of models. Some of the models are described in more detail in [14][5], and in [15].

3.3.1 One slope model – `1SM`

This is one of the simplest models available. It is very fast but *highly* inaccurate. This model only takes into consideration the distance between the receiver and the transmitter. Path loss is determined by the following formula:

$$L_{1SM}(d) = L_0 + 10n \log d$$

where $L_0(dB)$ is a reference loss value for the distance of 1m, n is the power decay factor (path loss exponent defining slope) and d is the distance in metres.

3.3.2 Multi wall models – `MWM`

This is a family of semi-empirical models that provide much better accuracy than `1SM`. These types of models incorporate knowledge of the building into the path loss calculations. Path loss is determined by the following formula:

$$L_{MWM}(d) = L_{FSL}(d) + \sum_{i=1}^{n} k_{wi} L_{wi} + \sum_{j=1}^{m} k_{fj} L_{fj}$$

[4] This report will not present the numerical results from any of the works. There is no common framework for conducting the simulations and evaluating their accuracy. Thus it is impossible to compare the results in an objective manner.

[5] Table 4.7.2 of [14] contains frequently used values for the coefficients used in the formulae presented in this section

where L_{FSL} [dB] is the free space loss (for 1SM $n = 2.0$), k_{wi} is a number of walls of i-th type between transmitter and receiver antennas, L_{wi} is attenuation factor for i-th wall type, N is a number of wall types, k_{fi} is a number of floors of a given type between transmitter and receiver and L_{fj} is the floor attenuation factor for the j-th floor type.

3.3.3 Linear attenuation model – LAM

Similarly to the 1SM model, it is a very simple model. It assumes that the path loss is linearly dependent on the distance:

$$L_{LAM}(d) = L_{FSL} + \alpha d$$

where L_{FSL} [dB] is the free space loss (for 1SM $n = 2.0$) and α is the attenuation coefficient

3.3.4 Remarks

The models presented here have the advantage of being computationally *light*. They are in complete opposition to the *"blueprint approach"* (*cf.* subsection 3.1).

This approach seems to fit well into the WIND project vision. The WIND project does not want to employ *much* computational resources to do the simulations, nor to depend on the detailed modelling of the building structure. Unfortunately this speed does *not* come for free. As mentioned in this section, the models do have a crippling feature – they lack accuracy. This may prove to be a hurdle in the context of the WIND project. Thus, other approaches have been investigated and are detailed below.

3.4 Ray-optical approaches

These approaches focus on the *particle-like* behaviour of electromagnetic waves (rather than the wave-like behaviour). They are based on optical geometry and geometrical theory of diffraction to compute the paths of photons in the simulation. It is obvious that such methods require information[6] about the environment (walls, doors, other obstacles). Ray-optical approach can deal to *some* extent with the NLOS propagation phenomena (reflections, waveguiding, diffraction, multipath). One can also easily account for the antenna radiation pattern.

These geometrical methods work best for the outdoor scenarios. One of the reasons behind this is the fact that the geometry of an indoor location is much more complex than the macro-scale outdoor environment. Ray-optical approaches do not scale well with the environment's complexity. Most of the researchers focus on reducing the computational complexity of these algorithms and improving their robustness with respect to modelling errors.

3.4.1 Ray-launching

This method relied on launching rays from a transmitter at a set of discrete angles in all directions. Their paths are traced until their power level drops beneath a certain threshold. One of the problematic features of this model is that spatial resolution decreases as the receiver is farther and farther away from the transmitter[7]. This model is seldom used because of the disadvantages it has.

3.4.2 Ray-tracing

This method looks for all the *valid* paths between the transmitter and receiver. Ray tracing is considered to yield high accuracy results, mainly because it implicitly takes into account, reflections, refractions, diffraction, etc. Regrettably there is a price to be paid. Normal ray-tracing is considered to be a power-hungry algorithm. There have been many attempts to shorten the computing time, while keeping the advantages brought by ray-tracing techniques.

[6]In most of the papers concerning this subject this information is being called a *database*.

[7]This approach can easy overlook small obstacles that happen to be located between the two rays.

Dominant path model (DPM) described in [16, 17, 18] focuses on the most *relevant* path during the propagation, thus shortening the computation time. The reasoning behind this idea is that the main part of the energy is delivered with along the main path. This approach has several advantages[8] – does not need the time consuming pre-processing, has relatively short[9] computation time and decreased dependence on the accuracy of the building databases. The waveguiding effects can now be taken into account[10].

Advanced database pre-processing is another attempt at reducing the complexity. It has been analysed in much detail in [19, 20, 21, 22] and also in [23]. It is being described in context of coverage analysis. During this process, one alters the position of the transmitter in order to see the corresponding changes of the coverage. The key observation is the fact that the building structure does *not* change and many of the visibility conditions and rays do *not* change as well. It means that some of the calculations could be done *a priori* thus reducing the footprint of the actual simulation.

There is one potential drawback of this approach. It seems to be suited towards static environments. Any change to the building database would require the repetition of the pre-processing step. Once that is done though, the simulation itself is fast. All the other characteristics (and flaws) follow from the ray-tracing approach.

Wave propagation using the photon path map Proposed by [24], the Photon Path Map is an extension of the Photon Map[11] algorithm. It facilitates the estimation[12] of the EM field strength in every point in space by tracing randomly generated photon paths.

Among the advantages mentioned in the paper, there are a few that stand out: the algorithms scales up *logarithmically* with regards to environment (scene) complexity, the accuracy of the propagation does not depend on the grid resolution and it can handle mixed scenarios (*indoor* and outdoor). However the approach presented in [24] models the transmission of only *one* sender. The interference of multiple senders is being detected by the network simulator.

Authors of [25] use this algorithm, in an attempt to visualise in real time the EM field patterns produced by moving transmitters.

Other Various different techniques also exist in scientific literature: [26] is another ray-tracing technique, optimised for 3D environments, but limited to a one floor representation of indoor scenes (multifloor extension seems straightforward provided that they have similar structure).

[27] and [28] extend the ray-tracing algorithm and introduce the concept of beams[13]. Authors state that it is well adapted to urban micro-cells and claim high accuracy and fast computation times (using NVIDIA CUDA platform). It appears that this algorithm could work well for indoor environment as well.

3.4.3 Remarks

These methods hold great potential because of their accuracy, simplicity and ability to simulate all the important physical phenomena (*e.g.* the delay spread). There are still some concerns with regard to the amount of required computing power.

---

[8]Among them authors [17] claim that DPM's accuracy at least as good as traditional ray-tracing

[9]When compared to other ray-tracing algorithms.

[10]Because with normal ray-tracing, the rays *"die"* off after several reflections, thus they do not end up contributing as much as they should to the propagation simulation

[11]Computer graphics algorithm that is being used to achieve the realistic lighting (Global Illumination) in 3D scenes. It employs a data structure for storing the incoming radiance on the surfaces. This data is then used to *estimate the radiance* of every pixel of the output image. The more photons there are per surface area, the more energy this surface has received.

[12]Authors call it a *discrete sampling* of the volumetric electromagnetic field by tracing stochastically generated photon paths.

[13]Beam is being defined as a continuum of rays.

3.5 Finite-Difference Time-Domain (FDTD) approaches

Unlike ray-tracing, these techniques focus on the *wave-like* aspect of the EM radiation. These approaches convert Maxwell's time dependent wave equations into difference equations, solving them for the electric field at one instant, then for the magnetic field at the next iteration. The main argument for the finite difference approach is that reflections and diffraction are *implicitly* taken into account (the simulation of waveguiding effects is natural – without performance penalties or dirty hacks).

These approaches, however, require a detailed description[14] of the environment. Worse still, the required discretisation grid is *directly* linked to the wavelength of the EM radiation one is struggling to simulate. In a nutshell the *shorter* the wavelength simulated, the *higher* resolution grid has to be used in order to produce correct results. This poses a real challenge for the wireless network simulations, as for 5GHz or even 2.4GHz the needed discretisation step becomes very[15] small making the simulations unreasonably complex. Most of the publications use some kind of *compromise*[16] to reduce the complexity at the expense of accuracy.

3.5.1 Finite-Difference Time-Domain

FDTD is a popular modelling technique that can be traced back to 1966[17]. One of the papers that applies this technique for wireless coverage simulations is [30]. Almost all the important physical effects are simulated at no additional overhead. The technique is simple to implement but demands large quantities of memory and it is very slow (computationally prohibitive in real-time). [31] tried to circumvent this problem by parallelising the algorithm and using GPUs[18] to take advantage of this. For big models, this resulted in large performance gains over the similar CPU based algorithm. CPUs still have the advantage of being able to manipulate larger model sizes.

3.5.2 ParFlow

ParFlow is a finite difference approach based on the Maxwell's wave equations. It is based on the concept of *partial flows*[19] and is equivalent to the TLM approach[20]. A more detailed description of the algorithm can be found *inter alia* in [32, 33, 34, 35, 36, 37, 38, 39]. As with FDTD, the computational load depends only on the discretisation level (number of cells). The low size[21] of each individual cell, required to simulate 802.11 radio propagation, makes this algorithm a very slow one indeed. A lower resolution can be used but it can produce artifacts and affects the simulation's accuracy.

3.5.3 MR-FDPF – Multi resolution Frequency Domain ParFlow

This approach proposes solving the Parflow equations in the *frequency domain*[22] using a multi-resolution formulation[23]. This approach was introduced by the author in [33]. A detailed description can be found in the aforementioned article. In the MR-FDPF approach the simulation computations are divided into *two* distinct steps: a long one-shot *preprocessing* step which does not depend on source characteristics, and a fast *propagation phase* which takes into account transmitter locations. It is well suited for complex, but *static* environments.

[14] Electric permittivity, magnetic permeability and electromagnetic conductivity

[15] On the order of a couple of centimeters, depending on the accuracy needed or the frequency used.

[16] Running the simulations at a much lower frequency than normal: 1GHz

[17] The basic FDTD algorithm was published in 1966 by Kane Yee in [29].

[18] Graphics Processing Unit

[19] In a 3D environment, a *vector field* in each of the discretisation points can be divided into 7 components: 6 flows bringing energy in a cardinal direction and additional stationary (inner) flow.

[20] Transmission Line Matrix method is a space and time discretising method for computation of electromagnetic fields.

[21] Suggested minimal spatial resolution is $\frac{\lambda}{6}$.

[22] Assuming that the time spreading is small, the frequency approach reduces the problem complexity to solving a linear system of equations.

[23] Introduces the concept of brick like elements, representing a group of pixels. The concept of bricks is also used in other scientific papers.

It is faster than most of the other FDTD methods presented in this report and appears to be a promising tool for the wireless network design phase. MR-FDPF is being continuously worked on and improved [34, 35, 40], extending its scope to 2.5D[24] [37] and then to full 3D [36]. It must be noted however that without good optimisations and simplifications the algorithm does not scale well[25] when moving from two dimensional to three dimensional environment.

3.5.4 Remarks

The approaches presented above produce some of the *most* accurate results that can be obtained through simulations. They *inherently* include important indoor events like reflections, refractions and diffraction without any additional computational penalty since they arise from the solutions of Maxwell's equations. On the other hand, they all *require* highly accurate blueprints of the environment that is being simulated, have long simulation times and are not well suited for dynamic environments. Due to these constraints the scope of applications of the FDTD models is, in principle, limited to the design phase of the wireless network. Furthermore the detailed information needed for setting up the simulations is seldom available.

3.6 Hybrid approaches

Shortcomings and limitations of all the models introduced thus far, such as lack of accuracy or high processing time, have given rise to so-called *hybrid* models. In general, *hybrid* methods are built by combining several approaches in an attempt to preserve as many of the advantages of the parent methods while inheriting as few drawbacks as possible. This is a daunting task, if one considers numerous underlying questions, *e.g.* how to correctly choose the models to be used in the framework, where to use each of the models as to maximise their efficiency and most importantly how to settle any incompatibility issues that may exist between the algorithms. Following subsection provides references to some examples of the *hybrid* approaches.

3.6.1 Brick tracing

Brick tracing technique represents a combination of two powerful algorithms: full-wave analysis of walls (*cf.* 3.5.1) and ray-tracing (*cf.* 3.4) to account for the multiple interactions between the wall elements. The walls are described in term of periodically arranged discrete units called *bricks*. Their electromagnetic response is calculated via FDTD, while the iterative field/current calculation algorithm[26] is used to compute the interaction between walls. This technique is presented in [41] and later extended to 3D environments [5, 42]. Authors claim the computational time is significantly reduced when compared to a normal full-wave simulation, but these claims should be taken with caution as this is still a slow algorithm. The heavy reliance on individual bricks is also noteworthy, as generating models of a building, by accounting for individual bricks is by no means a trivial task.

3.6.2 Others

[43] describes another algorithm that combines *ray-tracing* and FDTD methods. Unlike *brick tracing* though, *ray-tracing* is used in the analysis of wide areas, while FDTD focuses on areas close to complex discontinuities where the former algorithm does not produce sufficiently accurate results. Each zone to be analysed by FDTD is enclosed in a *"virtual"* box. This method still retains many of *ray tracing*'s vices though. Other hybrid methods worth mentioning are [44], or [45] which try to combine *ray-tracing* techniques with statistical modelling.

[24]In a multi-storey building each floor is being represented by interconnected 2D MR-FDPF simulations.

[25]Algorithm complexity in a 3D environment for an $N \times N \times N$ grid is $O(N^6)$.

[26]A generalised version of the *ray tracing* algorithm.

Propagation model family	Strong points	Shortcomings
Empirical models	Simple, high-speed (real-time), minimal *a priori* information required, model independence.	Low accuracy, inability to model most of the propagation phenomena
Ray-optical models	Good accuracy, Models the NLOS phenomena, accounts for antenna radiation pattern, similar algorithms are being used in computer graphics and scale well on the GPU platforms.	Requires *blueprint* environment knowledge, does not scale well with environment complexity, computational complexity.
FDTD models	High accuracy, all propagation effects are implicitly taken into account.	Requires detailed knowledge about environment (*blueprints*, material information, *etc.*), high computational complexity, does not scale well (dense discretisation grid needed).
Hybrid models	Trade-off between accuracy and complexity.	Additional complexity introduced by the hybrid framework, difficult to choose underlying algorithms and models.

Table 1: Summary of different modelling approaches.

3.6.3 Remarks

These has been considerable effort in building hybrid models that would provide a good compromise between the speed of computations and the accuracy of results. Unfortunately every approach still replicates many of the disadvantages that it was trying to mitigate in the first place. There is much ongoing research in this area and it is likely that there are better algorithms *in the offing*.

3.7 Conclusions

The field of propagation modelling continues to be a witness to a thriving research activity. There are tools on the market that help with the planning and deployment process. These tools operate in the *"blueprint"* mode, which forces the user to input the detailed information about the building layout, furniture and other elements that are to be part of the simulation Users are reluctant to spend much time on planning, they use the planning tools to gain some confidence as to their network design.

Because of that it is important to provide some information about propagation characteristics of the already existing wireless network, incorporating the real-life measurement information if possible. Nonetheless, *not many* of the available methods are either mature enough or efficient enough to be used for real-time wireless network *control*. Table 1 summarises briefly the described models.

4 Performance metrics

Proper monitoring is an *essential* element of any network. Monitoring provides both historical and instantaneous information about the network condition. This information facilitates diagnostics and troubleshooting. In addition to that, such knowledge is invaluable when it comes to planning network modifications. To complete the picture, one should remember that in order to do *any control* of the network[27] it is necessary to measure its state! How else could one tell that the actions have any effect and the *network performance* is *"increasing"*.

The WIND project is especially interested in *control* and *optimisation* of the network operation. In order to optimise the operation of a given system it is necessary to define some form of *cost function* that describes the goodness (or badness) of the current state. There is no single universal variable that would capture all the aspects of the network status[28].

The problem is that measuring the parameters contributing to the whole picture of the network is not as straightforward as one would expect. There is a multitude of metrics, yet it is far from trivial[29] to decide which carry *particularly* useful information. In consequence, before using a metric one has to *understand* what is really being captured, otherwise one might be in for a nasty surprise.

To make matters worse, useful metrics are being hidden in the physical layer of 802.11, forcing researchers to modify the APs' firmware in order to get ahold of the needed data. Of course, there is always the possibility of using probing and monitoring to estimate the metrics, but such approaches always come with the increased system complexity.

It is *desirable* that the metric employed be exposed through *standardised* and well known means like SNMP[30] or sFlow[31]. This would reduce the complexity of any monitoring software and insure that tools are not bound to a specific hardware vendor. Both SNMP and sFlow are widely used in the real world installations, yet the latter offers some unique features that make it a preferred choice for some of the monitoring activities. Unlike SNMP it operates in the *publisher*[32]/*subscriber*[33] model. The sFlow's *raison d'être* is twofold: to perform the random *packet sampling* of the traffic traversing the network device and to *collect* the vital device information[34]. This approach has a clear advantage over other monitoring methods since it scales well to complex and fast networks and its impact on the network load is insignificant. The information provided by sFlow is very detailed and facilitates nearly real-time access to the status of the network. Its success in the industry as well as in projects like CINBAD[35], make it a promising tool in the context of the WIND project.

Regrettably, monitoring was *not* a major concern when the 802.11 standards were being drawn. There was not much need for performance monitoring. The networks were slow, the number of the devices low and the conservative scheduling mechanism of PHY/MAC layer ensured smooth operation. Only some basic metrics that were needed for the operation of the PHY/MAC layer of the network were created (for example RSS[36]). Over time the need for additional metrics had become evident and new amendments were created. Unfortunately all the equipment produced up to that point does not support the new extensions. One can do nothing about this legacy equipment. Subsection 4.1 will treat about the new features that serve the monitoring purpose.

[27]By control one should understand tuning *certain* variables in order to achieve *better performance* of the system.

[28]*Performance* can be for example defined as coverage, throughput, maximum number of users, minimum roaming time, etc.

[29]RSSI (*cf.* section 4.2.1) is a good example of a metric that can be misleading and that has difficulties in conveying an accurate picture of the wireless environment.

[30]Simple Network Management Protocol used in network monitoring and management.

[31]sFlow is a standard for monitoring computer networks.

[32]Network device is running a special *agent* that is responsible for collecting the data and forwarding it to the subscribers.

[33]In most of the cases the Network Monitoring System will subscribe to receive the sFlow information.

[34]For example the interface SNMP counters.

[35]Joint CERN-HPN research project in the area of network monitoring and anomaly detection.

[36]Received Signal Strength Indicator is a measurement of the power present in a received radio signal. **Note**: 802.11 RSSI is acquired during the preamble stage of receiving an 802.11 frame (provided the reception was successful).

4.1 802.11 monitoring related amendments

4.1.1 802.11e – MAC Quality of Service Enhancements

Approved by the IEEE in 2005, this standard [46] attempts to alleviate some of the common issues that come with using wireless networks in *bandwidth* and *latency sensitive applications* like VoIP or video streaming. The original 802.11 standard implements two types of communication mechanisms for wireless stations: DCF[37] and PCF[38]. Given the nature of their design, mainly because they are *unable* to differentiate between different types of traffic, both of these methods perform inadequately in resource sensitive applications. This is the area where the 802.11e standard strives to improve. It introduces two new communication modes: EDCF and HCF.

EDCF or Enhanced DCF mechanism introduces the concept of *traffic categories/priorities* (eight categories in total). Stations with a low priority traffic will have to wait longer than higher priority ones in order to gain access to the medium.

EDCF Hybrid Coordination Function is an extension of PCF, also relying on a controller that pools stations and distributes specific start times and maximum durations for each transmission. More information on the communication mechanisms can be found in [47].

802.11e introduces several new information elements, but the most interesting is the QBSS Load[39] data. It is present in the Beacon frames, Probe Responses and Management frames. The QBSS Load contains information about the current station population and traffic levels in the QBSS. It contains different fields, among which Channel Utilization is worth mentioning. Channel Utilization is defined as *"the percentage of time, normalized to 255, the QAP (QoS AP) sensed the medium was busy, as indicated by either the physical or virtual carrier sense (CS) mechanism."* It is calculated by the following formula:

$$\frac{\text{channelbusytime}}{\text{dot11ChannelUtilizationBeaconIntervals} \times \text{dot11BeaconPeriod} \times 1024} \times 255$$

where *channel busy time* is the number of microseconds during which the CS mechanism has given the channel busy indication and *dot11ChannelUtilizationBeaconIntervals*[40] represents the number of consecutive beacon intervals during which the channel busy time is measured.

4.1.2 802.11h – Spectrum and Transmit Power Management Extensions

The introduction of this standard [48] was forced by the European regulations, which require a specific behaviour of the WLAN devices if a radar or satellite signal is being detected in 5GHz frequency band. 802.11h defines DFS[41] and TPC[42] mechanisms on top of 802.11 MAC and 802.11a/n PHY.

In order to successfully carry out these new duties, the AP has to *monitor* the status of the current channel as well as other channels. It can also request that the client stations carry out the measurements. There are three types of measurements that are needed for the DFS:

Basic measurement allows to check whether another BSS, a non-802.11 OFDM signal, an unidentified signal, or a radar signal is using the measured channel.

[37] Distributed Coordination Function relies on CSMA/CA for channel access. An optional RTS/CTS mechanism also exists.

[38] Point Coordination Function, only available in *infrastructure mode*, since it requires the AP to act as a traffic coordinator (point coordinator), deciding which stations can transmit at any given point in time.

[39] A BSS providing QoS is called a QBSS in 802.11e.

[40] Found in the IEEE 802.11 RRM MIB.

[41] Dynamic Frequency Selection mechanism for dynamically switching from one operational frequency to another in case of *radar* signal detection.

[42] Transmit Power Control mechanism that automatically reduces the used transmission output power when other networks are within range. This results in reduction of interference and increase of battery life.

Clear-Channel Assessment (CCA) measures the fractional duration over which the channel was busy during the measurement duration.

Received Power Indication (RPI) measures the histogram of the quantised received energy power levels as seen at the antenna connector. It is useful when assessing the general level of interference present on a channel.

In addition to that, 802.11h defines a TPC Report element that contains the information about the *transmit power* used to send the frame with the report itself and the link margin information[43].

Unfortunately this information is not necessarily exposed by the standard means (SNMP, sFlow, SOAP, etc.). On the client side it might be possible to access this information by tweaking the firmware. However the WIND project aims to use only the information provided by the *standardised* interface.

4.1.3 802.11k – Radio Resource Measurement of Wireless LANs

This is one of the first standardisation attempts that tackles the problem of the increasing density of wireless installations. IEEE 802.11k [49, 50, 51] specifies the set of measurements (PHY and MAC level) that can be performed by an AP and the client stations.

A wireless station (or an AP) can *request* that another station measure and report the information about[44]:

1. Channel Load[45] in which the measuring station reports the fractional duration over which either the physical or virtual carrier sense mechanism indicates that the medium was *busy*. It can help to asses the channel load.

2. Noise Histogram which provides the fractional time over which the energy detected was within a certain power range during a medium idle period (there are 8 different levels). More precisely, power histogram measures **non-IEEE 802.11** noise power. Sampling takes place when virtual Carrier sense determines an idle period and the station is not transmitting or receiving a frame. This report can be used to identify the expected value for the noise in a specific channel.

3. Neighbour Report which contains information about known neighbouring APs.

4. Beacon Report that contains a list of APs in range (on a given channel or channels), information about RCPI[46] levels, signal strength and SNR[47]. This measurement can be done using an active scan, passive scan or beacon table.

5. Frame Report which includes information about all the frames received in a certain time interval (count and average signal strength for each of the senders). This report can potentially give a picture of all the channel traffic.

6. Location Report that is used to exchange the location information (*e.g.* longitude, latitude and altitude) between the stations.

7. Link Measurement which indicates the instantaneous quality of a link. This measurement enables a pair of stations to compute link margins.

[43]Link margin is calculated as the ratio of the TPC Request frame signal strength to the minimum required by the station.

[44]This list is not exhaustive. Only the most interesting reports are being mentioned.

[45]Similar to the Clear Channel Assessment report used in 802.11h (see section 4.1.2).

[46]Received Channel Power Indicator is a replacement for the RSSI metric. Unlike the later one, RCPI measurement covers the entire received frame with defined absolute levels of accuracy and resolution.

[47]Signal to Noise Ratio is defined as the ratio between the incoming signal level and the interference level, calculated at the same time as RSSI.

8. `STA`[48] `Statistics` which basically correspond to set of different `MAC` counters standardised under `ieee802dot11`.

In addition to that some of the existing frames were extended to carry additional information (*e.g.* information about the `BSS` load in the beacons). The measured information should be exposed to network management via `SNMP`.

Up till recently the APs and clients were not able to share the channel information. 802.11k provides tools for building the global picture of the channel state. What is more, with the incoming reports from the clients, the AP will be able to track the *hidden nodes* that are located on the edge of their cells. The AP (or network management system) would be then able to simply direct those clients to an AP from which they would get a better service.

How can this new amendment help in network management? It defines a set of measurements and establishes a common framework for exchanging the information about the state of the network and environment between the devices. With 802.11k the APs and the network management system will have a more *complete* view of the wireless network. 802.11k might be instumental in improving and simplifying the wireless network operation and monitoring. Unfortunately it will take some time before enough 802.11k compliant devices proliferate the market. It might even be enough to do the firmware upgrade in case of the recent devices.

4.1.4 802.11v – management of the wireless networks

IEEE 802.11v [52] is a proposed amendment to the IEEE 802.11 standard to allow for the *configuration* of client devices while connected to wireless networks.

Currently a station will associate with the AP with the strongest signal. This approach can easily lead to a situation in which one of the APs is *overcrowded* while others remain *idle*. Signal strength alone is not enough as it does not tell anything about the real status of the network (channel load, number of devices associated to the AP, etc). 802.11v provides APs with a command to tell a station to associate with a different AP. It will be also possible to *force* an already associated station to move to a different AP. With these tools it will be possible to do the load-balancing without resorting to dirty hacks (see [7] for examples of pre-802.11v workarounds).

4.1.5 802.11T – Evaluating 802.11 wireless performance (cancelled draft)

The aim of this standardisation [53] attempt was to define a set of guidelines, methodologies and performance metrics that could be used by the hardware manufacturers to evaluate the performance of their products. All the metrics[49] are being defined in the context of certain use cases that are supposed to correspond to real-life situations[50]

However 802.11T is meant to be used mainly by the hardware manufacturers. All the tests have to be conducted in a set of specific and well-defined environments[51] in order to guarantee repeatability. The goal of 802.11T is to enable testing and comparison of 802.11 wireless devices based on a common and accepted set of performance metrics, measurement methodologies and test conditions. It cannot be applied directly to the real-life monitoring of an existing network.

4.2 Overview of metrics defined in various publications

Surprisingly there are not many publications analysing the performance metrics that could be directly adapted to describe *dense* wireless installations. Part of the research (*e.g.* [54, 55, 56, 57, 58, 59, 60]) stems from the exploration in the area of *ad hoc* routing in mesh networks. A metric reflecting the *quality* of the link is needed in order

[48] STAtion

[49] For example: latency, jitter, packet loss, throughput vs. path loss, fast BSS transition, receiver sensitivity, and access-point capacity and association performance.

[50] The three principal-use cases are data, latency sensitive and streaming media.

[51] Test setup is placed in a shielded chamber for isolation. RF cables connect the antenna ports of each device to other devices through programmable attenuators that emulate the path loss among the devices.

to select *the best route.* However typical conditions in which a mesh network operates and user requirements are *far* from those in dense indoor installations.

In addition to that some papers (*e.g.* [55]) analyse the metrics for *slow networks* (802.11, 802.11b). Some of these metrics are heavily speed dependent and might carry far less useful information or even worse their interpretation can be diametrically different. Most of the papers consider *stationary* 802.11 networks (yet the wireless medium is non-stationary and there is an increasing number of highly-mobile stations). Finally, some of the proposed solutions (*e.g.* [6, 8]) use dedicated *probes* to capture the needed information. For more information about the strengths and weaknesses of the probe based approach consult section 4.4.

Article [61] serves as a good introduction to performance indicators in wireless networks. It describes the metrics which could be used to assess the *link quality.* The authors have identified four metrics that are frequently used (*i.a.* in *ad hoc* routing research) to measure the quality.

Each of these metrics provides some insight into the *link quality.* However link quality alone is not sufficient to accurately describe the state of the network. Even worse, none of these metrics *alone* can give a robust link quality information.

4.2.1 Received Signal Strength Indication

RSSI, as already mentioned, is related to the strength of an arriving signal. It includes the energy from the incoming transmission as well as from other sources (e.g. noise, interfering transmissions). However, the RSSI is being measured *only* during the reception[52] of the frame preamble and header. This simple fact has far stretching consequences. First of all, as the interferences have high *temporal variability*, it is possible that the conditions change drastically between the preamble and data parts of the frame. What is more, the packet preamble is being transmitted at the *lowest rate*, which raises a question whether the metric is meaningful for higher transmission rates. Authors conducted simulations demonstrating that the RSSI is `invariant` with respect to interference power. Fortunately RSSI will soon[53] be replaced with RCPI which measures the received RF power over the data portion of the frame, which is a more useful piece of information as it does not have any of the RSSI defects.

Paper [55] analyses the relationship between the packet loss rate, transmitter-receiver distance and the SNR reported by the network interface. They confirm that the distance is not predictive of link reliability. Their experiments show that attempting to estimate the packet loss rate in real-time by observing recent transmissions can lead to large errors as the links exhibit *non-stationary* traits. Furthermore the values reported by the network interface can be misleading as they are being calculated only for *successfully received* packets. The authors conclude that in low-interference conditions, the SNR serves as a good indicator of link reliability, but as the external interference increases, the relationship is unpredictable.

[62] is yet another publication that builds up upon the RSSI as a performance metric despite its shortcomings. This paper introduces the models for describing the PHY behaviour of packet reception and carrier sense with interference in *static* networks. These models, after initial seeding, use the RSSI and *packet counts* measurements from a real network to capture its RF characteristics and then use this information to predict how the network would perform under different settings. The experiments attempt to characterise the stability of the wireless medium on different time scales. Losses tend to occur in *small bursts* but can be treated as independent for larger time intervals (*i.e.* tenths of seconds, minutes). Furthermore there is enough similarity between the measurements to make useful predictions over *moderate* time scales. It is thus necessary to measure the network at least for such time scales to get the stable predictions.

4.2.2 Packet Delivery Rate

PDR is simply the percentage of the transmitted packets that were successfully received. It provides a high-level overview of the link quality. This metric is tightly related with both the packet size and transmission rate. The shorter the packet is, the less likely it is to suffer from the interference. Similarly, as the transmission speed

[52]And only the *successful* reception is being considered.

[53]As soon as 802.11k becomes commonplace.

increases, the transmission time of each packet decreases[54]. Authors state that *"the PDR metric may not provide a unified estimate of the link quality"*. One could make up for the packet size and transmission rate dependency by simply deriving the PDR value for different packet sizes.

4.2.3 Bit Error Rate

BER is the ratio of the number of erroneous bits to the total number of received bits. It is not straightforward to estimate, as one has to exclude *outlier packets* (e.g. packets with corrupted length field, packets where the bits are shifted by some offset). There are two ways of measuring BER, both of which are not easy to implement on wide-scale. The first one requires comparing raw data bits received by the station with the ones that were sent. In order to get a decent estimate of the BER one could send a known pseudo-random sequence of bits. For that reason it is not being used in practical applications. The second approach involves reading the internal CRC mismatch counter in the station.

4.2.4 User mobility metrics

Link quality is not the only metric that might be important for accurate representation of the network state. There were also some attempts [63] to characterise and quantify the *user mobility*, thus attempting to uncover the behaviour patterns of the user population which is important for effective *load balancing*. For example typical laptop users tend to spend a large fraction of their time in a single location, albeit, according to [63], when they move away they do not reduce their data transfer rates (*ergo* their behaviour does not change). The authors provide a framework for quantifying the user mobility[55] and introduce the *persistence* that measures how long users stay associated with the same access point and *prevalence* that reflects how frequently users visit different locations. For more information about the user mobility and usage patterns consult section 5.

4.2.5 Channel activity based metrics

Available Bandwidth Publication [64] introduces the idea of *available bandwidth* which is being defined as a maximum rate that a new flow can send without impacting the rate achieved by the existing flows on the link. In the context of wireless networks, the link does *not* have a well-defined bandwidth due to the dynamic multirate adaptation. Similarly due to 802.11 MAC contention-based mechanism the packets may not follow the FIFO scheduling. Authors provide extensive information about the related work and introduce a new method called *ProbeGap* which is based on probing the link for *idle periods*. The Poisson-spaced series of probe packets containing the local timestamp information is being send over the link. If the link is free, then the experienced delay will be low. [65] uses throughput estimation for their WLAN frequency planning algortihm and notes that 802.11e *channel utilisation* parameter transmitted in beacons would provide much more precise information about the channel state.

`Channel free time` In the context of load balancing and monitoring, it is important to get a hold of as much information about the wireless environment as possible before taking any action on the network itself. Essential pieces of information are provided by metrics similar to the `Channel free time`. They play a key role in assessing the usage level of a wireless channel since they manage to capture several important factors: traffic generated between stations and their associated AP, background[56] traffic and interference from other sources. In addition to that `Channel free time` offers a good estimation about the load of a particular AP.

Heretofore the scientific community either needed to employ dirty hacks to surface this information, or estimate it through measurements and probing. Although available in some form or another in various 802.11 amendments, their slow adoption meant that devices providing this information were scarce, to say the least.

[54]One has to remember that the relation is not that simple, as high-rate transmissions use more aggressive modulations that are more susceptible to noise.

[55]The results are based on the data from a real network.

[56]Other wireless traffic on the same channel.

The `Channel free time` metric can me found in different variants, all portraying the *same* aspect of the wireless environment – percentage of the time for which the medium was sensed busy.

1. `Channel utilisation` (*cf.* 4.1.1) – This information is available neither via SNMP nor by `sFlow`. Only the beacon interval can be configured via SNMP (*dot11ChannelUtilizationBeaconInterval*).

2. `Clear Channel Assessment` (*cf.* 4.1.2) is part of the CSMA/CA mechanism. The novelty of this standard is that an AP can *request* its associated stations to perform a `Clear Channel Assessment`. This gives the AP a much broader view of the wireless environment since it no longer has to rely solely on the data from its immediate vicinity. This information is exposed neither via SNMP nor `sFlow`.

3. `Channel Load` (*cf.* 4.1.3) extends the monitoring capabilities by allowing stations to perform `Channel Load` requests. This gives them the capability to perform intelligent decisions about the most effective way to utilise the wireless medium. This information is available in the IEEE 802.11 RRM MIB – *dot11ChannelLoadReport*.

4. `Channel Access Delay` (*cf.* [66]) is an *active* technique[57] and it measures the minimum time delay for a packet transmission in the network. Worth mentioning is the fact that CAD is dependent on the instantaneous network activities, thus it is subject to high variability over time. As such, the authors look at CAD *distributions* over a short period of time, in order to estimate the network congestion level. This in turn, produces results that vary depending on the number of CAD values available during an interval.

5. `Channel Busy Time` (*cf.* [66]) is a *passive* technique[58], and it measures the fraction of the time for which the medium was utilised during a certain measurement period. CBT was found to exhibit *strong* linear correlation to medium utilisation, making it a good candidate for revealing channel utilisation information.

6. `Free air time` (*cf.* [7]) is similar to the other metrics, as it represents the percentage of the time when the medium is *not* in use. This metric is *estimated*, by having the AP periodically broadcast a small packet at a fixed transmission on the highest priority driver queue. The amount of time between enqueueing the packet and successful dispatch serves (after subtracting the calibration data[59]) as the measure of the `Free air time`. This method gives a good approximation, but tends to underestimate the free air time to some degree.

4.3 *Ad hoc* network oriented metrics

As it was already pointed out, there is a significant number of papers examining the routing process in the *mesh* and *ad hoc* networks. These papers attempt to formulate best metrics for describing the link quality. Some of these methods (*e.g.* [67]) employ already described metrics (*viz.* RSSI, PDR, SINR) and analyse their behaviour (*e.g.* stationarity, transmission rate influence, interferences) in context of the *outdoor mesh* network.. The following sections provide a brief overview of the selected representative metrics.

4.3.1 Expected Transmission Count

Article [54] proposes a new metric for *multi-hop routing*. The ETX metric incorporates the effects of link loss ratios and the asymmetry in the loss ratios between the two directions of the links. The authors note that 802.11 ACK mechanism re-sends lost packets, making all but the worst links appear loss-free. Unfortunately these retransmissions reduce the path throughput and interfere with other traffic.

[57] CAD values are monitored for probes sent at regular intervals.

[58] The authors resort to modifications to the driver used by Atheros radios, in order to get the required information.

[59] On an idle channel the frame is sent immediately. This means the delay recorded is just the transmission time plus some additional overheads. None depends on channel conditions.

The ETX of a link is the predicted number of data transmissions required to send a packet over that link, including retransmissions[60]:

$$ETX = \frac{1}{d_f d_r}$$

Where d_f (forward delivery ratio) is the measured probability that a data packet successfully arrives and d_r (reverse delivery ratio) is the measured probability that an ACK packet is successfully received. Thus $d_f d_r$ is the expected probability that a transmission is successfully received and acknowledged. Delivery ratios are being measured using the dedicated probe packets. The authors make several simplifying assumptions about the 802.11 networks:

- Radios have a fixed transmit power.
- In real-life delivery ratios depend on both packet size and transmission speed
- 134-byte packets are being used to estimate link loss ratios. This results in the underestimation of the delivery ratio for ACK packets which are much smaller.

The simulations have been performed on an *ad hoc* network running 802.11b at 1Mbit/s transmission speed and were routing oriented. This metric is strongly rooted in the world of *ad hoc* networks and mesh routing. It is extended further in [58] in order to alleviate the fact that wireless channels experience variations at different time-scales. For example, some channels with low average packet loss ratio may have high variability, which is not taken into account by metrics depending solely on the mean loss ratio.

[59] introduces yet another metric for estimating link quality based on data traffic and compares it to other noteworthy approaches (*i.a.* ETR).

4.3.2 Efficient and Accurate link-quality monitoR

EAR is the link-quality measurement framework for *multi-hop* wireless *mesh* networks [57]. This approach combines three complementary measurement schemes: passive (utilising the real-traffic passing through a node), cooperative (by overhearing the traffic on the network) and active (by means of dedicated packets).

EAR focuses on *link cost* and *capacity* as link-quality parameters. *Link cost* is defined as the inverse of the *delivery ratio* (d) of MAC frames. This definition reflects the expected transmission count of each data frame. The cost ($\mathcal{C}$) of link $A \to B$ is defined as:

$$\mathcal{C} = \frac{1}{d_i} \text{ and } d_i = (1 - \alpha) \times d_{i-1} + \alpha \times \frac{N_s}{N_t}$$

Where d_i is a smoothed delivery ratio, α is a smoothing constant, N_s is the number of successful transmissions and N_t is the total number of transmissions and retransmissions during a measurement period of the ith cycle.

The authors are aware of the fact that the packet size greatly affects the delivery ratio. EAR monitors packets within a 100-byte range of three popular sizes used in the Internet: 60, 512 and 1448 bytes deriving the link cost for each size.

4.3.3 QUality ESTimation

Paper [68], written in collaboration with HP Labs, proposes a method for estimating wireless link quality in *mesh* networks without in-band signalling overhead. QUEST relies upon special *profiles* which describe the relation between the PDR and SNR that was established *experimentally*. The quality of a given link is being assesed by performing the *profile* lookup for any incoming packets.

[60]Assuming that each attempt to transmit a packet can be considered a Bernoulli trial.

4.4 Probes, monitors, sniffers

Currently the information provided by the 802.11 wireless networks that do not support 802.11k amendment is limited to the most basic metrics (*i.e.* RSSI). This has compelled some of the researchers (*e.g.* [6, 8]) to use special probes, monitors and sniffers to obtain the needed information.

Using dedicated hardware probes[61] to gather the data, certainly has its advantages. First of all the probe, unlike an AP, can do *virtually any* kind of monitoring (e.g. examine signal strength, noise level and data rate for individual packets[62], monitor error rates, retransmissions, etc) and run custom processing software. The main function of the AP is handling the wireless traffic leaving not much space for other activities (i.e. to do the monitoring of other channels AP has to temporarily suspend its operation). There are some APs available on the market that feature a dedicated radio interface for monitoring purposes allowing for continuous monitoring. However these APs need *more* power for the additional wireless interface and the current power provided by the most widespread PoE[63] standard (802.3af) is not sufficient. Recently ratified 802.3at standard is able to provide more power, but there are still not many devices supporting it.

What is more, probes can be installed *without* impacting the existing infrastructure (i.e. they are *transparent* to the all other network devices) and moved around freely. Passive monitoring does not need any form of interaction with the existing wireless networks.

Unfortunately using the probes is not as easy at it would seem (*cf.* [6, 8]). First of all a probe is an additional piece of hardware that has to be purchased and it is clear that one needs a certain number of such devices to effectively monitor the wireless network. Next pre-operation challenge involves deploying the probes, finding the best location that has *needed infrastructure* (*viz.* mount, cabling, Ethernet socket). Finally the sniffers' data has to be send to a central location and merged. This unavoidably brings us to the problem of time synchronisation, data aggregation and analysis[64].

Hopefully 802.11k amendment will provide information of comparable quality without the need for superfluous devices.

4.5 Location detection in wireless networks

While indoor location detection by means of a wireless network may not be of key interest for the WIND project, articles addressing this topic have proved to be an invaluable source of information. Still it is not that unexpected. Some kind of propagation modelling is needed in order to deliver the approximate location information about the individual users. Furthermore, the input data for the modelling comes from metrics collected by the network devices and dedicated probes.

The principle behind location detection is simple. Probes (monitors) that are being deployed in a well known location, listen to the transmissions from the other devices and record the signal strengths which is one of the fundamental values for estimating the position of the device. Publication [8] describes a *self-configuring*[65] location estimation engine. Their approach relies heavily upon the information provided by the dedicated probes[66].

The article [69] gives an alternative to using probes to monitor the network, by using the APs themselves. The authors divide their investigation in three parts: evaluation of the effects of users' presence on the RSSI information gathered, evaluation of the statistical properties of RSSI sets and evaluation of the properties of multiple RSSI sets. The study on the effects of users' bodies on the wireless propagation is not only important for its degrading impact on the location information gathered, but also on wireless performance in general (*cf.* 5.2). Also noteworthy are the following findings:

[61] The cited papers carried out experiments with standard PCs equipped with USB wireless card and dedicated software as probes.

[62] A small subset of this information can be obtained using sFlow protocol from the access point.

[63] Power over Ethernet

[64] Albeit one will face similar challenge when faced with the 802.11k data.

[65] The system employed by the Microsoft Research team, does not need any initial or subsequent manual calibration.

[66] Authors used simple USB wireless dongles attached to the campus PCs. The PC runs a special software that does the driver level packet monitoring of all the frames. Each frame from the driver is being delivered to a set of filters that analyse the data, summarise it in application specific manner and submit it to the central monitor.

- Good communication signals between the AP and the client *may not* always result in a good positioning signal.
- A limited number[67] of APs can provide location information for a small number of users.
- Distribution of `RSSI` is not usually Gaussian.
- `RSSI` sets from different APs are mostly uncorrelated.
- Interference from other APs working on the same frequency does not have a great impact on the observed `RSSI` pattern.

More information on this subject can be found in [70, 71, 72]

Location information could also be used in optimising the operation of a wireless network. It can help to answer some of the following questions [8]:

- Do active clients typically associate with the nearest (distance) AP?
- What is the relationship between distance and loss rate?
- How is transmission-rate selection affected by the distance?
- Which locations exhibit the heaviest utilisation?
- Does the AP do a good job covering the area (are the sensitivity and power settings correct)?
- How many clients in a particular region were able to connect to the wireless network and for how long?
- Are there any regions with no RF coverage?
- Are there any hidden nodes?
- Is connection duration correlated with location?
- Are there specific regions where clients rapidly switch back and forth between APs?
- Are there any clients that violate the protocol rules (*e.g.* due to buggy firmware or for selfish reasons)?

Answers to many of this non-exhaustive list of questions are far from trivial and are likely to be of much use for the network operation and optimisation. However it seems that Wi-Fi alone might not provide sufficiently detailed information.

4.6 Conclusions

Quantifying the *quality* of the wireless network is anything but simple. It is not obvious at all what values should be monitored and what is their relation to the overall network *goodness.* What is worse, till recently, measurements were not regarded important or essential by the 802.11 standards. Only the recent proliferation of mobile devices highlighted the need for a complete and standardised set of performance metrics. Unfortunately, for now, we have to work with devices that provide only the legacy measurements. This means that it is necessary to rely mainly on the limited information from the access points (*e.g.* RSSI, client rates, channel time estimates).

The first and most important task is to define what is to be considered as a *good*, *healthy* and *efficient* wireless network. There is more to it than meets the eye. Sufficient coverage is merely a prerequisite, while factors like throughput, fairness, reliability, roaming speed are far more important.

Each of the metrics provides some kind of average, implicit estimate of the wireless network quality. How accurate are these metrics? What is their robustness with respect to noise, interference, transmission speed and characteristics of the given site? We will try to answer some of these questions during the course of the WIND project.

[67] Even two APs could be sufficient to distinguish a modest number of positions.

5 Network usage and mobility

Although we consider it to be an essential topic, playing a key role both predicting the performance of the wireless network and helping with optimization efforts during its life-cycle, it is surprisingly poorly documented. Understanding how the network is used, how peak load conditions are generated or what situations lead to *"choke"* points in terms of throughput and availability, provides vital elements in creating a healthy, stable and efficient network. User mobility also plays an important factor in designing a good wireless network. Sadly it is often forgotten, both in the scientific literature as in the network conception phase.

5.1 Usage and mobility

Valuable conclusions can be drawn from [73]. The authors try characterizing both mobility and network usage in a large corporate WLAN. Despite the fact that it deals with older wireless networks (802.11b), the patterns that were observed during the study can still apply for higher speed wireless networks (a, g, n).

The results presented in this article are compiled from a four week trace on a corporate network. More than 1300 unique[68] `MAC` addresses were seen active during this period.

At the end of the study, important conclusions are presented:

- The throughput of an AP does not depend on how many users are present, but rather on their type of behaviour[69] on the network.
- Users mobility *does not* influence their behaviour on the network.
- There is a low correlation between the time of day and network throughput.
- There is a low correlation between network usage and type of network.
- Personal differences among users affect how they use the network. High bandwidth consumers are only a small portion of the users[70].

5.2 User effects

[69] is another paper that deserves to be mentioned. While the main focus of the article is somewhat different than network usage and mobility, it carries out experiments on the effect of the user's presence on the performance of the wireless network. As expected there is a clear influence on the received signal strength, when any number of users are present. More importantly, this article also shows that even the user's body orientation has a non-negligible effect. In certain test conditions a difference in body orientations resulted in 9dB attenuation difference.

5.3 Conclusions

We see the need to introduce a system that can adapt in real time to changing conditions within a wireless network, as no simulation or site survey can mitigate or react to the events presented in this section.

[68] This number is an upper bound limit for the number of unique users, as some users could have owned more than one wireless device.
[69] The behaviour type refers to what the wireless network is used for: streaming, checking emails or webpages, *etc.*
[70] Pareto principle (also know as the 80/20 rule) – states that 80% of the effects come from the 20% of causes.

6 Operation and optimisation of wireless networks

One might have thought that after struggling with the deployment phase of the wireless network, the day-to-day operation would be plain and simple. Surprisingly, later, *the deployment issues seem like a picnic*. It is a real uphill struggle against the capricious wireless medium, limitations of the standard, hordes of relentless users, *etc.*

Wireless network are *complex systems*, made up of a large number of *interacting components*. One have a conglomerate of dissimilar users immersed in the highly variable environment that is difficult to measure. In most cases one has a heterogeneous wireless infrastructure that was either created by random purchasing decisions (which is bad) or conceived on purpose to avoid *vendor lock-in*. Providing redundancy in the wired network boiled down to providing additional network equipment and cabling, yet it is far more challenging due to the limited spectrum and the influence of interference. It should be obvious by now that providing sufficient coverage, acceptable performance and high-enough reliability is easier said than done.

Succinctly speaking, network operators do not have the tools to effectively monitor, diagnose and optimise their wireless networks. This problem has recently drawn attention of the researchers. In this section we will present several promising attempts to attack the problem.

6.1 About the recurring patterns

While looking at the literature tackling the problem of the dense wireless networks and at the performance metrics (*cf.* section 4) one can distinguish certain recurring themes. It comes as no surprise as there are not that many *degrees of freedom* when it comes to controlling and tweaking the operation of the wireless network. Among these are, *inter alia*, the frequency allocation across the access points, load-balancing of user associations across the access points[71] and the process of power control (or conversly – the receiver's sensitivity control).

6.2 MDG – Measurement-Driven Guidelines

[1] provides a good starting point for understanding how these control variables are correlated with the network performance. The authors wanted also to examine whether the above-mentioned areas of optimisation are interdependent or can be treated separately. In particular it turned out, that applying all the optimisations *independently* is not preferable as it can degrade the performance. The authors outline the conditions needed for efficient network tuning and formalise them within a Measurement Driven Guidelines framework. The optimisation algorithms comprising the MDG framework use the Gibbs sampling[72] for finding the optimum. The authors performed comprehensive evaluation of different scenarios in the testbed environment[73].

There is a clear distinction between the networks operating in the 5GHz band (with less bandwidth contention and multiple orthogonal channels available[74]) and the ones in 2.4GHz band (that is shared with plethora of other devices where there are only 3 orthogonal channels). The indicated paper states that for the 5GHz networks the optimisation of transmit power is *not needed* and the management process boils down to selecting the orthogonal (noise-free) channels and *load-balancing* across the access points.

Frequency Selection The goal of the optimisation is to allocate frequencies to the access point in such a way that the total amount of noise across the entire network is minimised and the amount of power sensed at each access point from its co-channel access points together is minimised:

$$\mathcal{F}_a = N_a + \sum_{b \neq a} s_{CH}(a,b)(P_b(a) + P_a(b))$$

[71] The default behaviour of associating with the strongest access point in the area can lead to pathological load distribution.

[72] Gibbs sampling generates a sequence of samples from the joint probability distribution of several random variables.

[73] Albeit, in our opinion, the number of clients used (30) was too small. Similarly using the *random* channel allocation as the reference baseline is far from being realistic, as the deployments tend to have *some* form of planned channel assignments. On the other hand the it is worth pointing out that the MDG performance was evaluated in the presence of external interferers.

[74] [1] recommends that in order to eliminate interferences, the neighbouring access points with mutual RSSI ≤ -40dBm should choose frequencies that are separated by at least 40MHz.

where N_a accounts for total thermal noise and interference from non-802.11 sources, $P_b(a)$ represents the power received at AP a from AP b[75] and s_{CH} represents the degree of orthogonality between the channels. According to the authors the channel selection process should depend on the AP-AP information, as the client-AP information exhibits great variability. We believe that this is an interesting approach and one could get even better results by employing better sources of information (*e.g.* 802.11 monitoring related amendments that provide all the necessary information).

User Association The default user association process is naïve and results in asymmetric load of the access points. The goal of the optimisation is to minimise the amount of time that a user needs to wait until the reception of a unit of information from its access point. This means that each client will associate with the access point providing the *minimum long-term delay*. A new metric is introduced, the Aggregated Transmission Delay that is calculated by measuring the time between queueing a packet at the MAC layer and receiving and ACK. Most likely information of similar quality could be obtained by using one of the channel utilisation metrics (*cf.* section 4.2.5). User association management should be applied in conjunction with frequency selection to maximise the overall network throughput.

The authors highlight that managing the user associations can be more beneficial for 5GHz. In 5GHz cells are smaller and the number of available channels is higher, thus it is more likely to find a suitable access point.

Power Control MDG uses an algorithm described in [74]. The goal is to tune the transmission power and CCA threshold[76] in order to balance the reduction in interference and the reduction in signal quality to the weakest client. The publication describes typical *scenarios* in which one can expect the algorithm to *work* and to *fail*. Intelligent frequency allocation can help reduce the number of aberrant cases, while blind application of the user association management can have the opposite effect.

MDG **operation** MDG is a decision framework that based on the measurement information decides which algorithms should be applied. Following information is needed:

1. Whether there are overlapping cells using the same channel. If yes then frequency selection algorithm is triggered.
2. The access point associations lest an overload should occur.
3. The signal quality in order to look for situations where power control could be applied.

The authors did not considered modifying the optimisation algorithms, instead they focused solely on the deciding what are the conditions for using each of the algorithms.

6.3 DenseAP

DenseAP [7] is a novel system for improving the performance of enterprise wireless networks using a *dense deployment* of the access points. This is contrary to the widespread deployment practises and inline with what most of the people thinks, albeit the underlying argumentation is different from the typical *"give us more access points"* mantra.

Using only few APs in the 5GHz range (where there are more orthogonal channels than in the 2.4GHz band) does not allow to fully utilise the available spectrum at each location. Similarly adding extra radios to the existing access points is not as effective as deploying a larger number of APs in different locations due to rapid signal fading in indoor environments. Having more access points would help mitigate the *rate anomaly*[77] problem.

[75] Information based on the RSSI.

[76] Channel is assumed to be idle if the channel power level is below the CCA threshold.

[77] Stations reduce data rates when signal strength is poor. As low-rate packets consume more airtime, the high-rate stations experience throughput degradation.

In order to increase the capacity, the access points must be assigned appropriate channels, the clients must make intelligent decisions about which AP to associate with and power (that translates to cell size) has to be tuned to limit the interference. The decisions about channel assignment and associations need to take into consideration the global view and state of the whole wireless network. The proposed solution *does not* require any modifications of the clients[78]. The central controller manages the association process and performs periodic load balancing by moving clients from overloaded APs to nearby APs with less load. The access points that have *no clients* are dedicated to *monitoring* all the channels.

Decision of the central controller are based on the following information from the access points:

1. List of associated clients – this information is readily available via `SNMP`.
2. Traffic pattern summaries describing for each client – it appears that similar information could be obtained via `sFlow`.
3. `RSSI` values of packet samples – *ut supra*,
4. *Current channel conditions* – captured in the `Available Channel` metric that belongs to the family of channel activity based metrics (*cf.* section 4.2.5 for information about obtaining similar information).
5. Information about the new clients requesting service from the network. Based on the state of the network the controller will decide with which access point the client could associate with.

In addition to that, the `DenseAP` approach allows for client location detection by means of algorithm described in [8].

The experimental testbed was deployed on a portion of the office floor and consisted of 24 access points[79] located in more or less every other office. The number of clients was also equal to 24 which makes the tests *somewhat* unrealistic. Most of the tests were conducted in the 5GHz band.

`Available Channel` **metric** Throughput that a client can expect to get with a given access point depends on several factors, such as quality of the channel, presence of other traffic/interference, rate adjustment algorithms in use and the `CCA` thresholds. The authors focus on the fact that throughput is mainly affected by:

1. The transmission rate used – if a client and an AP communicate at high transmission rate, then each frame will consume *less air time*, and the client will be able to send more data. Furthermore, other clients on the same channel will have *more opportunities to transmit.*
2. How busy the medium is (`free air time`) – if the channel has low utilisation then potentially *more clients* could use it and it is also acceptable to communicate at a low rate, because it will have little impact on other devices.

Given a channel (C), an AP (D) and a client (M), AC^{C}_{DM} is the product of `free air time` (*cf.* section 4.2.5) on C in the vicinity of D and M, and the `expected transmission rate`[80] between the D and M. The `free air time` is simply the percentage of time when the wireless medium is not in use. In principle the clients should associate with the least loaded access point that can provided the best data rate. This metric is also used in the channel assignment process.

Load balancing The access point is considered to be overloaded if it reports the `free air time` of less than 20%. The load balancing algorithm will look at the traffic patterns of the clients connected to the overloaded access point and try to find a less loaded one that would be capable of sustaining the client average transmission rate. The algorithm will move *at most* one client in each iteration and exclude this particular client from the next iteration in order to avoid *oscillations.*

[78]The association management and roaming are implemented using clever hacks on the access point side. Still the new standard amendments would surely simplify the process.

[79]In some experiments the number of access points was halved.

[80]Obtained from the `RSSI` of the probe frames by means of lookup table.

Power control The authors tested *some* adaptive power management schemes in their testbed and decided that power control at the access point is undesirable and the best policy was to simply use the fixed (optimum) power level (which confirms the observations from [74]).

Conclusions This paper also subscribes to the point of view that the centralised association management is crucial for achieving an efficient network. Numerous experiments carried out by the researchers show the benefits of the `DenseAP` system (*e.g.* graceful degradation of performance with increasing the number of clients, load balancing to achieve higher transmission speed, intelligent channel selection, *etc.*).

6.4 Load Balancing via 802.11k mechanisms

Both of the approaches mentioned so far stress the importance of the user association management and load balancing. [50, 51] gives the information how the new 802.11k (*cf.* section 4.1.3) standard can help in this vital tasks.

6.5 Conclusions

Efficient operation of wireless network is a complex task. We have already seen that it is difficult to quantify the state of the network (*cf.* section 4). It seems to us that having a centralised view is a prerequisite to even considering controlling a wireless network. The other necessary condition is having a dense deployment in order to guarantee sufficient performance and facilitate effective load balancing.

In many corporate environments deploying more access points should not pose significant problems. Of course, the deployment cost is important, but it is absolutely necessary to stress that a dense deployment can provide better performance and redundancy. The fact is that the proliferation of 802.11 clients is increasing. More and more applications run over wireless because the speed and quality of service are considered to be *"good enough"*. Even the critical applications start to notice the existence of the wireless users…There is a noticeable change in the way of thinking of many corporations – "we can now ran applications over wireless *safely*, gaining the mobility factor".

We think that the initial costs of the dense deployment would amortise themselves in no time while providing more efficient and robust wireless network.

7 Things to come

7.1 Beamforming

Wireless technologies have come a long way since the initial adoption of the 802.11 standard. Not only did the throughput consistently increase, but with the advent of 802.11n even crippling effects [81] were turned around and exploited to the benefit of wireless clients through the use of `MIMO`[82] technologies. And things are continuing to improve. One such technology is slowly maturing, bringing with it the promise of enhanced control over the wireless propagation and with it, higher throughput, superior reliability and resistance to interference. This technology is called *beamforming* (*viz. spatial filtering*) and stems from the domain of *array processing*.

Array processing is concerned with the analysis and extraction of the information received by an *array of sensors*. Such array is comprised of multiple[83] spatially distributed elements (*e.g.* antennas, hydrophones, microphones, geophones) that spatially sample the wavefront. Receiver of the infomation is in most cases concerned either with the content of the incoming signal (*i.e.* communication applications) or the location of the signal source (*i.e.* azimuth, elevation, distance). *Array processing* techniques have long been in use in military (*e.g* radar and sonar) and scientific (*e.g.* astronomy, seismology) applications. However their widespread adoption is limited due to the computational complexity of the algorithms involved.

In many cases in wireless applications one is interested with signals coming from a particular direction. *Beamforming* allows for the dynamic control of the radiation pattern of the *antenna array* by appropriate weighting and phase shifting of each elements' signal. One could direct the beam towards the clients locations, reducing the amount of energy output which could help mitigate the interference and allow for more efficient communication.

This is a giant leap forward in wireless propagation control and would prove to be an important asset in the context of the `WIND` project.

7.2 Beyond 802.11n

As we already pointed out, the introduction of 802.11n brought quantifiable performance improvements to wireless networks. However, this has not stopped pioneering work being done in this area, and already new IEEE task groups have been created to develop the next generation wireless standards.

802.11ac TG Responsible for the new standard called 802.11ac Very High Throughput[84], it seeks to achieve phenomenal speeds by today's standards: maximum 500Mbps single link throughput and maximum multi station throughput of at least 1Gbps (*cf.* [75]). This standard will be backwards compatible, coexisting with *"legacy"* 802.11 devices in the 5GHz unlicensed band. This standard is still some time away, as the current time projections (*cf.* [76]) put it for approval in late 2012.

802.11ad TG Responsible for the new standard called 802.11ad Very High Throughput[85], this is a sister project to 802.11ac, having the possibility of achieving much higher speeds[86], albeit with a much shorter range. With a projected date of approval similar to 802.11ac, it is still too early to tell how the finished specifications will materialize. Even so, by examining some released drafts and proposals (*cf.* [77]), we can see that companies like Sony, picture a standard complementing 802.11n/ac in the 2.4GHz/5GHz band, allowing for very fast, short range transfers. Sony even envisions a system where radios working in the 2.4GHz/5GHz bands would be used to detect stations out of reach for 60GHz transmitters and then using *adaptive beamforming* to extend and direct the *802.11ad* beam and possibly connect to those stations.

[81] Multipath (*cf.* section 3.1)
[82] Multiple In Multiple Out - Multiple transceivers are used to improve throughput and link range.
[83] At least two elements connected to the signal processing unit are needed to form an array.
[84] Working at frequencies smaller than 6GHz
[85] Working at frequencies of 60GHz
[86] Up to 7Gbps

8 Conclusions

This technical report attempts to summarise a selection of recent scientific papers in the area of wireless networking. The presented selection is by no means exhaustive, we deliberately focused on the subjects that are relevant to control and optimisation of corporate wireless networks, while disregarding on purpose much of the research related to *e.g.* *ad hoc* routing and *mesh* networks.

One obvious conclusion about is that almost all 802.11 network challenges stem from the wireless propagation phenomena. Unlike the wired medium, the wireless medium has a fixed amount of bandwidth which cannot be increased. One can of course deploy additional access points operating on orthogonal frequencies, but the number of such frequencies is very limited as well. The current tendency is to limit the cell size and create dense deployments by using many access points operating with low power and low sensitivity. This however, poses a challenge due to very limited monitoring and management facilities provided by the 802.11 standards. Effective load balancing and seamless roaming is not yet fully possible due to limited support of the newer standards (both on the client and access point side). Similarly, obtaining the monitoring information from the clients is currently a no-go.

We noticed yet another, even more, fundamental challenge. What is the answer to the following question: *"What is the characteristic of a good and healthy wireless network?"*. What are the client requirements and what can be guaranteed by the wireless installation? So far there is no single and universal answer to this question. Some of the papers tend to focus on the details (*cf.*section 6, *e.g.* reducing co-channel interference, load-balancing the users across different access points). Surely these details have impact on the overall performance of the network. But the initial question remains without a satisfactory answer. For example, stability of the connection is a must. The clients should not get disconnected, thus all disruptive operations (*e.g.* channel change, layer 3 roaming) should be limited to a minimum. Surely this is also a question of policy. In an organisation with many mobile `VoIP` users the requirements will be different from a library open-space. However it is important to identify the common features of a *good* network in order to have some initial bounds for assessing the performance.

One way of attacking the problem is to consider specific scenarios. A set of conditions determines a scenario (*e.g.* auditoria, long hallways, office space), and then according to the scenario and measurements, a solution is chosen. That solution is bound to a specific scenario, because it might be improving the network operation in some cases but not in others. But then another question arises: *"How to distinguish the scenarios?"*, and the answer is not obvious. Real life situations are not black and white – an auditorium can be surrounded by office space, office space can be located in an old warehouse, *etc.* It would be far better not to rely too much on the concept of scenarios...

WIND project should focus on defining what constitutes a good, healthy and reliable network. We have to look at the measurable parameters and see how they correlate with the end-user experience. Only after we establish a reasonably generic way of measuring the *"goodness"* of the network, we will be able to carry forward with tuning the network. Simultaneously we will look at different control variables for tuning the wireless networks. What do the access points support in terms of association management, sensitivity tuning, *etc.*. Do we see any situations in which these actions could improve the operation of the network? Then, combining the input variables and the control variables we could test different approaches to wireless network optimisation.

References

[1] Ioannis Broustis and Michalis Faloutsos. Mdg: measurement-driven guidelines for 802.11 wlan design. In *In Mobicom*, pages 254–265. ACM, 2007.

[2] Kyle Jamieson, Bret Hull, Allen Miu, and Hari Balakrishnan. Understanding the real-world performance of carrier sense. In *E-WIND '05: Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, pages 52–57, New York, NY, USA, 2005. ACM.

[3] Micah Z. Brodsky and Robert T. Morris. In defense of wireless carrier sense. In *SIGCOMM '09: Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 147–158, New York, NY, USA, 2009. ACM.

[4] P. Chandra, D.M. Dobkin, A. Bensky, R. Olexa, D. Lide, and F. Dowla. *Wireless Networking: Newnes Know It All.* Newnes Newton, MA, USA, 2007.

[5] M. Thiel and K. Sarabandi. An accurate 3d hybrid model for electromagnetic wave propagation in indoor wireless channels. In *IEEE Antennas and Propagation Society International Symposium, 2008. AP-S 2008*, pages 1–4, 2008.

[6] Jihwang Yeo, Moustafa Youssef, and Ashok K. Agrawala. A framework for wireless LAN monitoring and its applications, 2004.

[7] Rohan Murty, Jitendra Padhye, Ranveer Chandra, Alec Wolman, and Brian Zill. Designing high performance enterprise wi-fi networks. In *NSDI'08: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pages 73–88, Berkeley, CA, USA, 2008. USENIX Association.

[8] Ranveer Chandra, Jitendra Padhye, Alec Wolman, and Brian Zill. A location-based management system for enterprise wireless LANs. In *NSDI*. USENIX, 2007.

[9] Aleksandar Neskovic, Natasa Neskovic, and George Paunovic. Modern approaches in modeling of mobile radio systems propagation environment. *IEEE Communications Surveys and Tutorials*, 3(3), 2000.

[10] F. M. Landstorfer. Wave propagation models for the planning of mobile communication networks. March 15 1999.

[11] P. Almers, E. Bonek, A. Burr, N. Czink, M. Debbah, V. Degli-esposti, H. Hofstetter, P. Kyösti, D. Laurenson, G. Matz, A. F. Molisch, C. Oestges, and H. Özcelik. Survey of channel and radio propagation models for wireless mimo systems. *EURASIP Journal on Wireless Communications and Networking*.

[12] M. Khosroshahy. Ieee 802.11 and propagation modeling: A survey and a practical design approach. 2007.

[13] A. Zaballos, G. Corral, A. Carné, and J.L. Pijoan. Modeling new indoor and outdoor propagation models for WLAN.

[14] E. Damosso and L. Correia. COST 231 (Digital mobile radio towards future generation systems). *Final report, European Comission, Bruxelles*, 1999.

[15] S. Zvanovec, P. Pechac, and M. Klepal. Wireless LAN networks design: site survey or propagation modeling? *Radioengineering*, 12(4):42–49, 2003.

[16] G. Wolfle. Dominant paths for the field strength prediction. March 15 1998.

[17] G. Wolfle, R. Wahl, P. Wertz, P. Wildbolz, and F. Landstorfer. Dominant path prediction model for indoor scenarios. 2005.

[18] R. Wahl and G. Wolfle. Combined Urban and Indoor Network Planning Using the Dominant Path Propagation Model. 2006.

[19] G. Wolfle. Accelerated ray optical propagation modeling for the planning of wireless communication networks. March 15 2000.

[20] G. Wolfle. Fast 3-D ray tracing for the planning of microcells by intelligent preprocessing of the data base. March 15 2000.

[21] G. Wolfle. Radio network planning with ray optical propagation models for urban, indoor, and hybrid scenarios. April 03 2001.

[22] R. Hoppe, P. Wertz, FM Landstorfer, and G. Wölfle. Advanced ray-optical wave propagation modelling for urban and indoor scenarios including wideband properties. *European Transactions on Telecommunications*, 14(1):61–69, 2003.

[23] M. Allegretti, M. Colaneri, R. Notarpietro, M. Gabella, and G. Perona. Simulation in urban environment of a 3d ray tracing propagation model based on building database preprocessing. *Proceedings of the XXVIIIth General Assembly of International Union of Radio Science, unpaginated CD-Rom*, 2005.

[24] Arne Schmitz and Leif Kobbelt. Wave propagation using the photon path map. pages 158–161, 2006.

[25] Arne Schmitz and Leif Kobbelt. Real-time visualization of wave propagation. 102:71–79, 2006.

[26] GE Athanasiadou and AR Nix. A novel 3-d indoor ray-tracing propagation model: The path generator and evaluation of narrow-band and wide-band predictions. *IEEE transactions on Vehicular Technology*, 49, 2000.

[27] A. Schmitz, T. Rick, T. Karolski, L. Kobbelt, and T. Kuhlen. Beam tracing for multipath propagation in urban environments. *EuCAP 2009*, 2009.

[28] A. Schmitz, T. Rick, T. Karolski, T. Kuhlen, and L. Kobbelt. Simulation of radio wave propagation by beam tracing.

[29] Kane Yee. Numerical solution of inital boundary value problems involving maxwell's equations in isotropic media. *Antennas and Propagation, IEEE Transactions on*, 14(3):302–307, May 1966.

[30] L. Talbi and GY Delisle. Finite difference time domain characterization of indoor radio propagation. *Progress In Electromagnetics Research*, 12:251–275, 1996.

[31] Samuel Adams, Jason Payne, and Rajendra Boppana. Finite difference time domain (FDTD) simulations using graphics processors. pages 334–338, June 2007.

[32] F. Guidec, P. Kuonen, and P. Calégari. Parflow++: A c++ parallel application for wave propagation simulation. *SPEEDUP Journal*, 10(1/2):68–73, 1996.

[33] J.M. Gorce and S. Ubeda. Propagation simulation with the parflow method: fast computationusing a multiresolution scheme. 3, 2001.

[34] J.M. Gorce, K. Runser, G. De La Roche, and I. Lyon. Fdtd based efficient 2d simulations of indoor propagation for wireless lan. In *IMACS, World Congress Scientific Computation, Applied Mathematics and Simulation, Paris, France*, 2005.

[35] Jean-Marie Gorce, Katia Jaffrès-Runser, and Guillaume De La Roche. The adaptive multi-resolution frequency-domain parflow (MR-FDPF) method for indoor radio wave propagation simulation. part I : Theory and algorithms, November 2005.

[36] G. De La Roche and J.M. Gorce. A 3d formulation of mr-fdpf for simulating indoor radio propagation. In *Antennas and Propagation, 2006. EuCAP 2006. First European Conference on*, pages 1–6, 2006.

[37] Guillaume de la Roche, Xavier Gallon, Jean-Marie Gorce, and Guillaume Villemaud. 2.5D extensions of the frequency domain parflow algorithm for simulating 802.11b/g radio coverage in multifloored buildings. pages 1–5, 2006.

[38] G. Roche, K. Jaffres-Runser, and J.M. Gorce. On predicting in-building wifi coverage with a fast discrete approach. *International Journal of Mobile Network Design and Innovation*, 2(1):3–12, 2007.

[39] G. de la Roche, J.M. Gorcey, and J. Zhang. Optimized implementation of the 3d mr-fdpf method for indoor radio propagation predictions. pages 2241–2245, 2009.

[40] J. Gorce, K. Jaffres-Runser, and G. De La Roche. Deterministic approach for fast simulations of indoor radio wave propagation. *IEEE transactions on Antennas and Propagation*, 55(3):938, 2007.

[41] M. Thiel and K. Sarabandi. A hybrid method for indoor wave propagation modeling. *IEEE Transactions on Antennas and Propagation*, 56(8 Part 2):2703–2709, 2008.

[42] M. Thiel and K. Sarabandi. 3d-wave propagation analysis of indoor wireless channels utilizing hybrid methods. *IEEE Transactions on Antennas and Propagation*, 57:1539–1546, 2009.

[43] W.Y.C. SK and S. Safavi-Naeini. A hybrid technique based on combining ray tracing and fdtd methods for site-specific modeling of indoor radio wave propagation. *IEEE Transactions on Antennas and Propagation*, 48(5), 2000.

[44] R.A. Valenzuela. Ray tracing prediction of indoor radio propagation. 1:140–144, 1994.

[45] M. Hassan-Ali and K. Pahlavan. A new statistical model for site-specific indoor radio propagation prediction based on geometric optics and geometric probability. *IEEE transactions on Wireless Communications*, 1(1):112–124, 2002.

[46] IEEE. *802.11e; Medium Access Control (MAC) Quality of Service Enhancements*. IEEE.

[47] S. Mangold, S. Choi, P. May, O. Klein, G. Hiertz, and L. Stibor. IEEE 802.11 e Wireless LAN for Quality of Service. 18:32–39, 2002.

[48] IEEE. *802.11h; Spectrum and Transmit Power Management Extensions*. IEEE.

[49] IEEE. *802.11k ; Radio Resource Measurement of Wireless LANs*. IEEE.

[50] Eduard Garcia Villegas, Rafael Vidal Ferré, and Josep Paradells Aspas. Load balancing in WLANs through IEEE 802.11k mechanisms. In Paolo Bellavista, Chi-Ming Chen, Antonio Corradi, and Mahmoud Daneshmand, editors, *ISCC*, pages 844–850. IEEE Computer Society, 2006.

[51] Eduard Garcia Villegas, Rafael Vidal, and Josep Paradells Aspas. Cooperative load balancing in IEEE 802.11 networks with cell breathing. In *ISCC*, pages 1133–1140. IEEE, 2008.

[52] IEEE. *802.11v* Proposal*; Wireless Network Management*. IEEE.

[53] IEEE. *802.11T* Cancelled draft*; Recommended Practice for the Evaluation of 802.11 Wireless Performance*. IEEE.

[54] Douglas S. J. De, Douglas S. J. Couto, Daniel Aguayo, Robert Morris, and John Bicket. A high-throughput path metric for multi-hop wireless routing, 2003.

[55] Michael R. Souryal, Luke Klein-berndt, Leonard E. Miller, and Nader Moayeri. Link assessment in an indoor 802.11 network. In *In Proc. of IEEE WCNC 2006, Las Vegas*, 2006.

[56] D. Valerio, L. De Cicco, S. Mascolo, F. Vacirca, and T. Ziegler. Optimization of ieee 802.11 parameters for wide area coverage. *Proceedings of MEDHOCNET*, 2006.

[57] Kyu han Kim and Kang G. Shin. On accurate measurement of link quality in multi-hop wireless mesh networks, 2006.

[58] Can Emre Koksal and Hari Balakrishnan. Quality-aware routing metrics for time-varying wireless mesh networks. *IEEE Journal on Selected Areas in Communications*, 24(11):1984–1994, 2006.

[59] Hongwei Zhang, Anish Arora, and Prasun Sinha. Learn on the fly: Data-driven link estimation and routing in sensor network backbones. In *INFOCOM*. IEEE, 2006.

[60] Hairong Zhou, Chi-Hsiang Yeh, and Hussein T. Mouftah. A reliable low-overhead MAC protocol for multi-channel wireless mesh networks. In *GLOBECOM*, pages 1370–1374. IEEE, 2007.

[61] Aggelos Vlavianos, Lap Kong Law, Ioannis Broustis, Srikanth V. Krishnamurthy, and Michalis Faloutsos. Assessing link quality in IEEE 802.11 wireless networks: Which is the right metric? In *PIMRC*, pages 1–6. IEEE, 2008.

[62] Charles Reis, Ratul Mahajan, Maya Rodrig, David Wetherall, and John Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In Luigi Rizzo, Thomas E. Anderson, and Nick McKeown, editors, *SIGCOMM*, pages 51–62. ACM, 2006.

[63] HyungJune Lee, Alberto Cerpa, and Philip Levis. Improving wireless simulation through noise modeling. In Tarek F. Abdelzaher, Leonidas J. Guibas, and Matt Welsh, editors, *IPSN*, pages 21–30. ACM, 2007.

[64] Karthik Lakshminarayanan, Venkata N. Padmanabhan, and Jitendra Padhye. Bandwidth estimation in broadband access networks. Technical Report MSR-TR-2004-44, Microsoft Research (MSR), May 2004.

[65] Timo Vanhatupa, Marko Hännikäinen, and Timo D. Hämäläinen. Evaluation of throughput estimation models and algorithms for wlan frequency planning. *Computer Networks*, 51(11):3110–3124, 2007.

[66] P.A.K. Acharya, A. Sharma, E.M. Belding, K.C. Almeroth, and K. Papagiannaki. Rate Adaptation in Congested Wireless Networks through Real-Time Measurements. *IEEE Transactions on Mobile Computing*, 2010.

[67] Daniel Aguayo, John C. Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. Link-level measurements from an 802.11b mesh network. In Raj Yavatkar, Ellen W. Zegura, and Jennifer Rexford, editors, *SIGCOMM*, pages 121–132. ACM, 2004.

[68] L. Verma, S. Kim, S. Choi, and S.J. Lee. Reliable, low overhead link quality estimation for 802.11 wireless mesh networks. *Proc. IEEE WiMesh*, 2008.

[69] K. Kaemarungsi and P. Krishnamurthy. Properties of indoor received signal strength for wlan location fingerprinting. pages 14–23.

[70] C.L. Wang and Y.S. Chiou. An adaptive positioning scheme based on radio propagation modeling for indoor WLANs. In *IEEE 63rd Vehicular Technology Conference, 2006. VTC 2006-Spring*, volume 6, 2006.

[71] Thomas King, Thomas Haenselmann, and Wolfgang Effelsberg. Deployment, calibration, and measurement factors for position errors in 802.11-based indoor positioning systems. In Jeffrey Hightower, Bernt Schiele, and Thomas Strang, editors, *LoCA*, volume 4718 of *Lecture Notes in Computer Science*, pages 17–34. Springer, 2007.

[72] M. Brunato and R. Battiti. Statistical learning theory for location fingerprinting in wireless lans. *Computer Networks*, 47(6):825–845, 2005.

[73] Magdalena Balazinska and Paul Castro. Characterizing mobility and network usage in a corporate wireless local-area network. 2003.

[74] Vivek Mhatre, Konstantina Papagiannaki, and François Baccelli. Interference mitigation through power control in high density 802.11 wlans. In *INFOCOM*, pages 535–543. IEEE, 2007.

[75] IEEE. 802.11ac: On the feasibility of 1gbps for various mac/phy architectures. 2008.

[76] IEEE. 802.11ac: Timeline. 2008.

[77] IEEE. 802.11ad: New technique proposal. 2010.