



DISSERTATION

Address Assignment in Indoor Wireless Networks Using Deterministic Channel Simulation

Submitted at the Faculty of Electrical Engineering and Information Technology,
Vienna University of Technology in partial fulfilment of the requirements for the
degree of Doctor of Engineering Sciences (Dr. techn.)

under supervision of

Prof. Dr. Christoph Grimm
Dept. of Computer Science
University of Kaiserslautern

Prof. Dr. Markus Vincze
Automation and Control Institute
Vienna University of Technology

by

Dipl.-Ing. Edgar John Holleis
Matr. Nr. 9725754
Belvederegasse 39/2, 1040 Vienna

Abstract

A crucial step during commissioning of wireless sensor and automation networks is assigning high-level node addresses (e.g. floor / room / fixture) to nodes mounted at their respective location. This address assignment typically requires visiting every single node prior to, during or after mounting. For large-scale networks it also presents a considerable logistical effort. This thesis describes a new approach to automatically assign high-level addresses without visiting every node.

First, the wireless channel is simulated using a deterministic channel simulation in order to obtain node-to-node estimates of path loss. Next, the channel is measured by a pre-commissioning test procedure on the live network. In a third step, results from measurements and simulation are condensed into graphs and matched against each other. The resulting problem, identified as weighted graph matching, is solved heuristically.

Viability of the approach and its performance is demonstrated by means of (i) a publicly available test data set, (ii) measurements at the Vienna University of Technology using a custom built ZigBee network, and (iii) synthetically generated data. The quality of the resulting address assignment depends foremost on the geometrical layout of the network. This is demonstrated using formal and experimental results.

Common office style layouts result in network topologies that require unrealistically accurate measurement and simulation to solve flawlessly. Under realistic conditions, up to 15% mismatches are to be expected. Favourable layouts, e.g. when no two nodes of the same node type share the same room, can be solved flawlessly and reliably.

The thesis systematically analyses influencing factors, such as the network's topology, fading due to multi-path propagation, antenna radiation patterns, building materials and the impact of model accuracy. Moreover, the ZigBee network measurement procedure and the weighted graph matcher based on deterministic annealing and implemented using GPGPU (general purpose computation on graphics processing unit) are presented.

Deutsche Kurzfassung

Ein entscheidender Schritt bei der Inbetriebnahme von drahtlosen Sensor- und Aktuator-Netzwerken ist die Vergabe von semantischen Adressen (z. B. Stockwerk / Raum / Position) an die Netzwerkteilnehmer an ihrem Einbauort. Dies ist in der Regel ein händischer, pro Netzwerkteilnehmer ausgeführter Arbeitsschritt – vor, während oder nach der Montage. Für große Netzwerke stellt die Adressvergabe einen erheblichen logistischen Aufwand dar. Diese Arbeit beschreibt einen neuen Ansatz für die automatische Zuweisung der semantischen Adressen ohne die Notwendigkeit des manuellen Zugriffs auf jeden einzelnen Netzwerkteilnehmer.

Schritt Eins simuliert den drahtlosen Kanal unter Verwendung einer deterministischen Kanalsimulation mit dem Ziel die Kanaldämpfung der einzelnen Funkstrecken zu ermitteln. Im Schritt Zwei wird der Kanal mittels des installierten, aber noch nicht fertig in Betrieb genommenen Netzwerkes systematisch vermessen. Im dritten Schritt werden die Ergebnisse von Simulation und Messung zu gewichteten Graphen zusammengesetzt und miteinander verglichen. Das resultierende Problem ist das bekannte „Weighted Graph Matching Problem“, welches mittels Heuristik gelöst wird.

Die Effektivität des Ansatzes wird (i) anhand eines öffentlich verfügbaren Datensatzes, (ii) anhand eines eigenen ZigBee Netzwerkes und (iii) anhand synthetisch erzeugter Daten gezeigt. Die Qualität der resultierenden Adressvergabe hängt vor allem von der geometrischen Anordnung des Netzwerks ab. Dies zeigt sich anhand formaler und experimenteller Ergebnisse.

Die typische Architektur von Bürogebäuden resultiert in Netzwerktopologien, welche für ihre einwandfreie Lösung unrealistische Maßstäbe an Messung und Simulation legen. Unter realistischen Bedingungen führt dies zu Fehlerraten bis zu 15%. Günstige Topologien, z. B. wenn nie mehr als ein Netzwerkteilnehmer des gleichen Typs den selben Raum teilen, können problemlos und zuverlässig gelöst werden.

Die Arbeit analysiert systematisch die relevanten Einflussfaktoren, wie die Netzwerktopologie, Schwund aufgrund von Mehrwegausbreitung, die Richtcharakteristik der beteiligten Antennen, Gebäudebaustoffe, sowie die Präzision der geometrischen Modelle. Darüber hinaus legt die Arbeit das Messverfahren des ZigBee Netzwerks, sowie den heuristischen Lösungsalgorithmus des Gewichteten-Graphen-Problems dar. Letzterer basiert auf „Deterministic Annealing“ und ist mittels GPGPU (Allzweckberechnung auf Grafikprozessoren) implementiert.

Acknowledgements

I'd like to thank my supervisor Christoph Grimm, who kept me on the right track. I'd like to thank my colleagues and superiors at Tridonic – Reinhold Juen, Walter Werner, Heinz Seyringer and Stefan Zudrell-Koch, who let me loose on the subject and let me see it through. I'd like to thank my colleagues of the SmartCoDe project for three years of painstaking yet ultimately successful research. I'd like to thank my colleagues at the ICT – Stefan Krywult, Christian Hambeck, Franz Lukasch and Thomas Hartmann, who endured the joys and pains of PHD school with me, who provided invaluable input, and shared many a lively discussion and good laugh. I'd like to thank my parents and family, who believed in me and maintained me through the early years of my studies. I'd like to finally thank my significant other – Martina Maida, who never stopped supporting me with her love, despite the many evenings of my absence, weekends of stress and skipped vacations.

Table of Contents

1 Motivation & Problem Statement.....	1
1.1 Outline of Automatic Address Assignment.....	3
1.2 Contributions.....	5
1.3 Thesis Overview.....	6
2 Introduction into the Field.....	7
2.1 Modern Building Automation Networks.....	7
2.2 Wireless Building Automation Systems.....	10
2.3 The Problem with Cable Replacement.....	13
3 State of the Art.....	15
3.1 Address Assignment in Wired Building Automation Systems.....	16
3.2 Address Assignment in Wireless Building Automation Systems.....	20
3.3 Related Work – Localisation Systems.....	23
4 A New Method for Automatic Address Assignment.....	25
4.1 Formalization.....	25
4.2 The Impact of Network Topology.....	29
4.2.1 Probability of Matching Error in the Three-Node Network.....	33
4.2.2 The Impact of the Indoor Environment on the Probability of Matching Error.....	37
4.2.3 Probability of Matching Error in Networks of Four and More Nodes.....	42
4.2.4 Consequences of the Formal Analysis of Network Topology.....	48
4.3 Model of the Indoor Channel.....	51
4.3.1 Fading.....	53
4.3.2 Ray Tracing.....	54
4.3.3 Antenna Radiation Patterns.....	58
4.3.4 Walls and Materials.....	61
4.4 Assembling the Measured Connectivity Graph.....	63
4.5 Pre-Processing the Graphs / Noise Cancellation.....	64
4.6 Weighted Graph Matching.....	68
4.6.1 Implementation and Performance of the WGM Algorithm.....	70
4.6.2 Clean-Up Heuristic and Hill-Climber.....	75
5 Experimental Evaluation of Automatic Address Assignment.....	77
5.1 University of Mannheim Data Set.....	78

5.1.1 Modelling of the Demonstration Site.....	78
5.1.2 Completing the Test Data Set.....	80
5.1.3 Results.....	81
5.2 Vienna University of Technology Data Set.....	82
5.2.1 Impact of the Topology.....	83
5.2.2 Impact of Antennas.....	85
5.2.3 Large Data Set.....	89
5.3 Synthetic Data Sets.....	93
5.4 Summary of Findings.....	94
6 Conclusion.....	97
6.1 Future Research.....	99
Bibliography.....	103
Appendix A: ZigBee Graph Assembly Cluster.....	109
Appendix B: ZigBee Network Used for Experiments.....	113

1 Motivation & Problem Statement

Wireless technology has seen significant advances over the last decade. The trend is likely to continue for some time as emerging technologies like the Internet of Things, cyber-physical systems and machine-to-machine communication will see widespread adoption. This will allow vehicles to become safer, buildings to become aware the inhabitants and increase their comfort and our workplaces to become smarter and more productive. It will undoubtedly also offer new and more pervasive ways of surveillance.

Many applications and the underlying networks will be mobile, ad-hoc and the communication opportunistic. Other applications will be stationary and carefully planned to perform a designated function, such as managing, monitoring or controlling some process. Especially the latter category depends on wireless nodes to be aware of their location, function or role: For a sensor it is not enough to just exist and measure, to be useful it is crucial to know *what* is being measured *where*.

Answering the *what* and *where* – or in other words – assigning the role, function and location is the responsibility of the commissioning procedure. It is also the topic of this thesis: The commissioning of large scale wireless sensor and automation networks and ways to improve it. More precisely this work is about *address assignment*, which is the step of the commissioning procedure where a high-level address based on role, function or location is assigned to the wireless node (a more precise definition follows). Address assignment today is a manual task performed on a per-node basis.

The one application where large scale wireless automation networks exist today, is building automation. To be clear, *wireless* building automation is still considered an emerging technology and the vast majority of the market uses wired systems. But as of 2013, the list of wireless reference projects on industry supplier's homepages gets longer and market research hints at a slow but steady uptake. Wireless building automation networks therefore serve as example and reference against which the problem of address assignment is analysed and a novel method of address assignment developed.

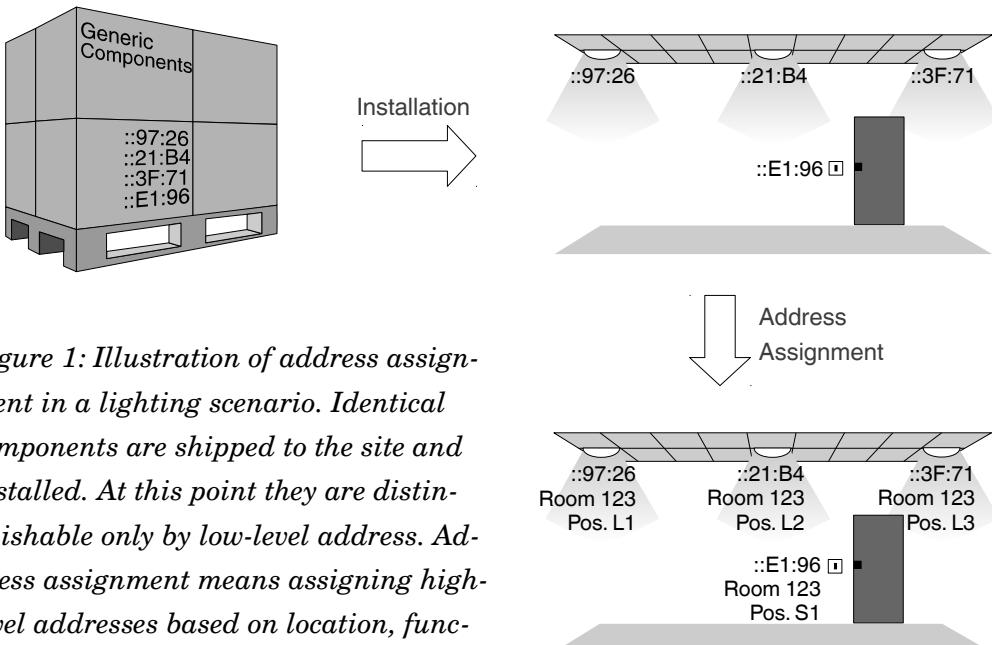


Figure 1: Illustration of address assignment in a lighting scenario. Identical components are shipped to the site and installed. At this point they are distinguishable only by low-level address. Address assignment means assigning high-level addresses based on location, function or role. The address assignment is the basis for computing node configuration, such as binding and group membership. The node configuration is then deployed using commissioning protocols.

Building automation is also both a conservative industry and a competitive industry. There is great reluctance to modify the established procedures and there is no disposition to raise unit costs by even a single cent unless it offers added value that the customer is willing to pay for. Address assignment procedures that do not increase unit costs are therefore preferable. Indeed, the radio's built-in capacity to measure signal strength can be used to infer location, and by extension the function and role of the node.

The topic of this thesis is automatic address assignment, i.e. an address assignment procedure that can automatically assign high-level addresses to wireless nodes without the need to visit or manually manipulate nodes. As a constraint, the procedure should fit into the established procedures of the building automation industry and should if possible rely only on built-in standard features of state-of-the-art wireless transceivers. This thesis sets out to answer the following questions: Can such a procedure work in the narrow setting of today's building automation? If not – what would it take to make it work? What are the relevant physical phenomena? What are favourable and unfavourable conditions?

Address assignment means attaching a high-level address which conveys locality (e.g. floor / room / fixture) to the nodes of the wireless network (figure 1). The goal is, in other words, to produce a mapping between addresses used by the media access control (MAC) layer, to logical addresses based on location or role. This is a non-trivial task especially in large-scale installations.

Some protocols are explicit in their distinction between low-level address and high-level address, e.g. KNX/RF. Other protocols, such as ZigBee, rely on low level addresses for communication, but support high-level roles with group membership or binding tables. High-level addresses become explicit again if dealt with in the context of a deployment plan or professional commissioning tool.

While this work deals with both location and wireless channel measurements, it is different from the field of sensor node localisation. Localisation is typically a continuous problem, address assignment is a combinatorial problem. The presented method of address assignment is range-free, i.e. it does not assume any empirical relation between signal attenuation and distance. It does not rely on fingerprinting the wireless channel and does not need anchor-nodes like indoor GPS. Instead, it uses a deterministic channel simulation to deal with the problems of multi-path propagation.

The novel approach is taken, because there is an ample body of literature showing that indoor localisation based on channel measurements alone is generally unable to provide the necessary precision to reliably and unambiguously identify mounting locations. Those localisation methods that do offer the necessary precision require an unacceptably high effort for producing the necessary channel maps or for installing the necessary anchor nodes.

This work is also not about auto-configuration protocols. The latter are the consumers of the address mapping produced by address assignment. A device without permanent storage can use an auto-configuration protocol to download its configuration from network management. But network management needs to have prior knowledge about the mapping between hardware-addresses and configuration records.

1.1 Outline of Automatic Address Assignment

The presented technique for address assignment works on a freshly installed network. The nodes are powered and an initial, maybe temporary wireless network has been formed. No further parametrization is needed, the only perquisite is that nodes are reachable by hardware address.

Automatic address assignment consists of the following three-step procedure:

1. *Channel estimation*: Using the deployment plan of the network and the floor plan of the building as inputs, a three-dimensional deterministic channel simulation produces the *simulated connectivity graph*. It is attributed with the nodes' logical addresses.

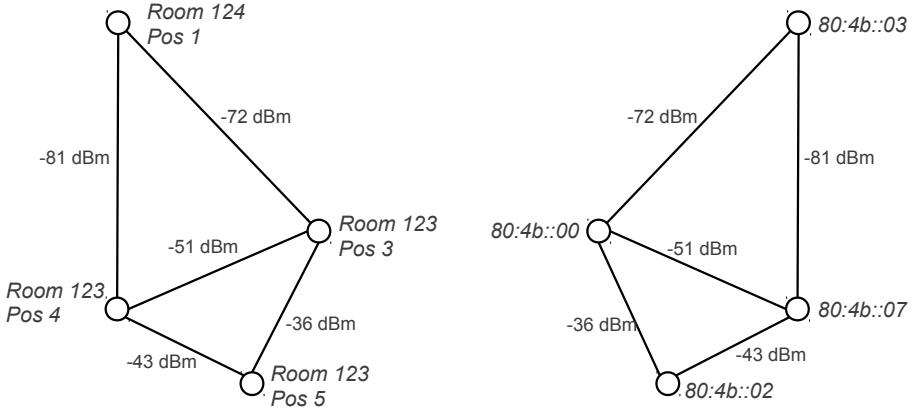


Figure 2: The connectivity graphs for illustration purposes. Real ones are dramatically bigger and weights have added noise. Left: Simulated connectivity graph attributed by logical addresses based on device location. Right: Measured connectivity graph attributed by MAC addresses.

2. *Channel measurement:* All participating nodes cooperate to perform measurements of the wireless channel, such as RSSI measurement. The measurements are transmitted to the management workstation to assemble the *measured connectivity graph*. It is attributed with the participating nodes' hardware addresses.
3. *Matching algorithm:* Simulated and measured connectivity graphs are isomorphic. They are thus matched against each other, producing the sought after mapping between hardware addresses and logical addresses.

Measured and simulated connectivity graphs are depicted in figure 2. For the channel estimation, ray-tracing is used. Channel estimation assumes the prior availability of detailed computer readable plans of the building, as well as a deployment plan of the network. This requirement is easily satisfiable for large, professional installations, because the necessary plans are provided by the planning agency as part of their contractual obligations. The plans need to be available before the responsible contractor can begin installation of the automation network. This is in line with the stated goal of devising an address assignment technique that naturally fits into established work flows.

The deviation between simulation and measurement is modelled as noise. Whether a particular data set, consisting of measured and simulated connectivity graph, results in a high quality address mapping depends on the amount of noise, and on the geometrical configuration of the network. As further analysed in section 4.2, some configurations can tolerate higher amounts of noise than others. Unfortunately, common office-style configurations are particularly sensitive to noise and therefore do not lead to error-free address assignments.

The most important contributing factors to noise are:

- Disturbances in the near field of the antennas resulting in effective radiation patterns that differ from the idealized radiation patterns used in the simulation
- Fading due to multi-path propagation which is only imperfectly compensated by sampling the channel at different frequencies
- Presence of materials that cannot be modelled by simple empirical approaches, such as anisotropic or frequency-selective attenuation, non-specular reflection and polarization

The crucial question is: What is an acceptable error rate such that automatic address assignment leads to cost savings? The answer depends on the costs of manual address assignment, the costs of automatic address assignment and the costs of mistakes of either method. The cost of developing automatic address assignment into a product also has to be factored in. Those numbers of course are either not available or confidential. A qualified estimate for the acceptable error rate is therefore difficult. Based on a back-of-the envelope calculation the author believes the break-even point to be between 1% and 5% error rate. Swapped pairs of nodes thereby count as one matching error.

1.2 Contributions

The central contribution is the method of automatic address assignment itself. There are automatic methods for wired networks (cf. section 3.1), but nothing of its kind for wireless networks, neither in academic nor patent literature. This includes the problem formalization, (section 4.1), the graph assembly protocol (4.4), the pre-processing (4.5) and weighted graph matching (4.6). The procedure, together with experimental data by means of the Mannheim data set (5.1), is published in ISRN Sensor networks [HOL13], and patented [HOL11].

As of yet unpublished is the formal analysis of the separability of different network topologies in the indoor channel (section 4.2). It is valuable also for the field of sensor node localization based on multilateration, because it gives insight into optimal placing strategies of anchor nodes.

The parallel implementation of the weighted graph matching algorithm (section 4.6), based on Nvidia's CUDA technology, may be of value for the fields of computer vision, especially object recognition and scene reconstruction. If the presented algorithm is able to deal with graphs of thousands of vertices in the time-frame of minutes, it might be fast enough to work on smaller graphs interactively.

The author's work on wireless network security is not incorporated into this dissertation. The unifying theme between address assignment and the security work is commissioning of wireless sensor and automation networks. However, the goals and methods are somewhat conflicting and a solution that unifies both – automatic secure address assignment and commissioning – is the subject of future work. Relevant publications are an introductory text to wireless network security [HOL13a], a commissioning procedure secured against man-in-the-middle attacks [HOL11a] and a diploma thesis created under the author's supervision [KRI13].

1.3 Thesis Overview

Section 2 gives an introduction into the organisational, commercial and technological environment of building automation in general. Building automation is the reference application for automatic address assignment. Constraints and assumptions of the latter sections are stem from the needs and logic of the industry. Section 2.2 goes on to explain the challenges of wireless from a traditional building automation perspective.

Section 3 analyses the state of the art of address assignment in wired and wireless building automation systems. Section 3.3 looks at related technology, namely sensor network localisation.

Section 4 is the core of the dissertation. It features the formalization of the problem (section 4.1), as well as the systematic discussion of influencing factors: The impact of network topology (4.2), the impact of the indoor channel (4.3), including fading, simulation effects, antennas, walls and materials. Further topics are the channel measurement procedure (4.4), data pre-processing (4.5) and finally matching (4.6).

Section 5 has an experimental evaluation of automatic address assignment. It is demonstrated by several different data sets. The Mannheim data set is a publicly available dataset of indoor channel measurements [KIN08]. It is used to demonstrate the approach as a whole, as well as the noise model and preprocessing procedure (5.1). Several other datasets are recorded at the Vienna University of Technology for the very purpose of analysing and demonstrating the different aspects of automatic address assignment (5.2): Topology, antennas, as well as a large dataset of 95 network nodes.

Section 6, finally, revisits the theoretical and practical considerations and recapitulates which goals can be achieved, which are demonstrated to be unachievable and which are left for future work.

2 Introduction into the Field

This section draws on the publications of the field-bus groups at Vienna University of Technology (key publications are [KAS05] and [SAU11]), as well as the authors experience while working at Tridonic GmbH & Co KG, a leading supplier of electronic lighting components.

2.1 Modern Building Automation Networks

Building automation is a mature industry with long standing, established procedures and division of work. The services found in modern buildings are typically safety alarms (e.g. fire), HVAC (heating, ventilation, air-conditioning), lighting (including emergency lighting), access control and intrusion alarms. Often, these services are procured separately and supplied by different vendors. Since different services have different communication requirements, separate standards have evolved for each of them. The standards converged and today's prevalent standards (KNX, BACnet, LonWorks) support all services.

Never the less, lighting and metering continue to rely on domain specific technologies, instead of adopting one of the big three. Lighting widely uses DALI and for metering applications M-Bus is the technology of choice in Europe. Even for services that do use standard protocols, vendors often rely on proprietary solutions and offer standards compliance via gateways that present a standards-compliant view to the rest of the system.

Increasingly, the formerly separate services become more and more interconnected, e.g. presence detectors can be used for lighting control and intrusion detection. They also become interconnected with modern IT systems, for remote management, web-based configuration, reporting and statistics. The increased system integration also concerns cabling and overall architecture. Modern building automation systems usually comprise a backbone network based on Ethernet. On the backbone, the automation protocols are encapsulated in TCP/IP frames. On each floor, one or several gateways connect the backbone to the service specific field buses. The gateways often double as local unit controllers.

In the semi-professional and enthusiast niches of the market, there are fully wireless automation systems, where base station, sensors and actors all communicate wirelessly. Examples would be Z-Wave, FS20, and others. At the core of the building automation market, commercial and functional buildings, wireless building automation networks compete with the field buses, to wirelessly connect sensors and actors to unit controllers or gateways. The Ethernet backbone and central network components continue to be wired.

Building automation systems generally follow a hierarchical architecture with minimal dependencies between zones of control. The architecture is geared towards cost-efficiency at the field-level, minimal organizational interdependency between contractors and graceful service degradation in the event of failure of some part of the system. It is characteristic for building automation systems to decentralize control and also to store configuration information in the periphery of the network. This lowers bandwidth requirements and enables disconnected operation.

The architecture of building automation networks is driven on the one hand by technical and legal requirements related to control latency, safety and cost-effectiveness, on the other hand by organizational requirements, i.e. the inner logic of the industry. Building automation networks are generally installed during construction or retrofitting of large, functional buildings. Stake holders include the owner, architect, planner, prime contractor, and several sub-contractors, down to the technicians that do the manual work. Construction sites are subjects to very evolved logistics, involving resource management, material, curing time, weather, and countless interdependencies. At every one time, various parts of the building or larger campus are typically in different stages completion and so are different parts of the building automation system. Some services like lighting can be commissioned floor by floor or room by room, other services like HVAC can be brought into service only after completion of the building at large. For sensors and actors mounted behind panels or above suspended ceilings, all necessary manual interventions have to be completed before said panels and ceilings are installed. This may require delaying final mounting until the building automation system is fully operational, or it may be solved by splitting the commissioning procedure into several smaller steps.

In order to fulfil contractual obligations, contractors need to be able to show that they did the contracted work correctly, on time, and they may need to produce documentation to prove these facts in the event of audits or disputes. In particular electricians need to be able to show that wires were installed according to plan, devices connected, powered and able to respond over the network. This entails that devices are generally able to perform some part of their function, such as responding to basic commands, even prior to commissioning, which is often done by a different contractor, at a later time. Technological solutions that want to gain

acceptance by the market should therefore support the characteristic division of labour. As a consequence, dependencies of different building services upon common infrastructure and upon each other should be kept to the technically necessary. Each additional dependency between the systems increases the complexity and cost of scheduling, contracting, logistics and introduces further risks of stalling the project. This is the main reason why building automation is generally considered a conservative market.

While control is generally decentralized, there are options with respect to management of the building automation system. One option is to have management panels at strategic locations of the building. These are embedded computer systems with touch screen and graphical user interface, capable of visualizing the system status (e.g. alarm status) and capable of performing common management, maintenance and commissioning tasks. For simple, non-sophisticated applications it is a cost-effective solution. Interconnected building automation systems usually come with a dedicated management workstation or management server based on mainstream computer hardware and operation system. This allows side-by-side installation of the management applications of the different system-vendors, software level integration, facility management data bases and interconnection of the building automation system with IT infrastructure.

The mainstream building automation systems (KNX, BACnet, LonWorks) all come with standardized commissioning protocols. These are standardized in terms of message formats. The precise procedures for using these low-level protocol building blocks in order efficiently commission large networks are however proprietary to the system vendors, in form of integrated software packages, comprising templates, import / export wizards, rule-based systems, etc. They may also use proprietary extensions to the standard commissioning protocols.

Commissioning of wired building automation networks generally consist of the following steps:

- Discovery & address assignment
- Assigning roles and generating device configuration
- Committing device configuration

Discovery means ascertaining that all devices that should have been connected to the network actually are. It is generally more useful to do after address assignment, because only if the locations of the devices are known is it possible to meaningfully react on any defects discovered.

Assigning roles is often already done earlier, during the planning stage. Together with the address assignment, final device configuration can be computed and committed to the networked devices via the commissioning protocols.

In addition to that, wireless networks require the additional step of *network forming*. Because the wireless channel is open and not bounded by the limits of the building, the devices need to be told which network to join. Network forming is closely related to the security architecture of the wireless network.

2.2 Wireless Building Automation Systems

Going from wired communication to wireless communication entails a considerable progression in complexity and overhead, not dissimilar with the step from point-to-point communication to bus communication, cf. figure 3. Without getting into too much detail, this section explains the physical and technological reasons for the comparatively higher complexity and overhead of wireless technologies.

Bus communication inherits all the problems from point-to-point (synchronisation, basic signalling) and adds the problems of bus arbitration, collisions and interconnection of bus segments. Every mainstream communication technology, be it telecommunication, computer networking or automation, has developed characteristic solutions to these problems, optimized to the respective domain.

Wireless communication has all the problems of bus communication and adds the following:

- Dramatically more noise and interference, dramatically increased dynamic range
- Open channel, not bounded by the length of the bus, cannot be physically protected
- Time-variance of the channel typically requires dynamic, self-healing routing protocols
- No efficient repeaters, broadcasting is complicated and expensive

Radio frequency front-ends are designed to cope with the formidable dynamic signal range found in wireless communications. Requirements with regards to cost effectiveness and en-

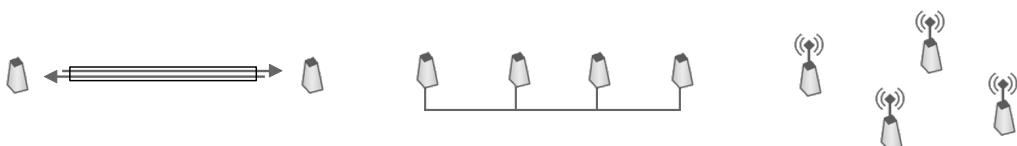


Figure 3: Technological progress from serial point-to-point communication, to field buses, to wireless. Challenges of serial communication are: framing and signalling; additional challenges of field buses are: arbitration, collision detection / avoidance; additional challenges of wireless are: the open, dynamic channel, fading, and inefficient repeating and broadcasting.

ergy efficiency can only be met with specialised mixed-signal integrated circuits. This means very high up-front costs of wireless technologies compared to wired, that can only be recouped with very high volume and economies of scale.

The increased noise and interference problematic is met with increased sophistication at the physical layer with regards to signal modulation, coding and decoding. Strategies include echo cancellation, channel equalization, spread spectrum modulation, or more recently MIMO (multiple input, multiple output) and UWB (ultra wide bandwidth). In contrast to mobile communication and wireless networking, the standards used in wireless building automation typically forego the most advanced strategies for the benefit of reduced complexity and improved energy efficiency. This does not mean that low-power wireless would rely on outdated technologies, it means that low-power wireless operates at a different point of equilibrium with respect to complexity, bandwidth, range, and energy requirements.

The open nature of the channel requires wireless networks to employ encryption in order to safeguard against unauthorized access. Encryption requires keys which complicate commissioning procedures. In wired automation networks, plugging a sensor or actuator into a field bus is an implicit statement about ownership. The wired node is justified to trust the network management of the field bus and start providing data and acting on commands. This is not the case in wireless networks. As soon as a node is powered, it is able to communicate, but it shouldn't generally trust any management node encountered, because that may just as well be under control of an attacker trying to subvert the wireless network.

The channel of wireless networks changes over time. Small changes in the configuration of the channel can make the difference between a reliable link between two nodes and one beset with frequent packet loss. The phenomena underlying this behaviour – free space loss, attenuation of the environment, reflection, diffraction, polarization and fading – are discussed in section 4.3. Indeed, this may be the core challenge of wireless networking: The fact that there is no way that a particular packet, sent between two nodes at a particular time and frequency, can be guaranteed to be correctly received at the intended destination.

Strategies such as CSMA/CD (carrier sense multiple access with collision detection) that can reliably detect collisions on wired protocols like Ethernet, have significant limitations in wireless networking. Channel availability can only be ascertained at the sender, not at the receiving end where the senders packet may collide with the packet of a third node that the sender is unaware of. This is called the hidden station problem and makes collision detection unreliable. The probability of packet loss is typically orders of magnitude higher in wireless networks than in wired networks.

Fading (cf. section 4.3.1) is destructive interference of a wireless signal with reflected versions of itself. It is the most incalculable effect of the indoor wireless channel and a direct

consequence of multi-path propagation. Fading is the reason why even nodes mounted in vicinity of each other, that ought to otherwise have good coverage, can experience occasional or permanent loss of connectivity. Highly dynamic channels, such as encountered in mobile communications, can use time statistics to combat fast fading (fading with high temporal variability). This is impractical in building automation, where the channel can be considered static over short time frames. Other strategies include antenna diversity, frequency hopping, ultra wide band and mesh routing. The common theme is to use the channel in multiple ways that are differently affected by fading. The destructive interference depends on a precise difference between path lengths. A slight variation of locality or frequency is typically able to avoid a deep fade. Mesh routing, where possible routes through the network are dynamically discovered by a distributed routing protocol, acts as an extreme form of antenna diversity.

The wide dynamic range between transmission power and reception power is the reason why wireless nodes cannot simultaneously send and receive. As a consequence, there are no efficient repeaters in wireless. Packets have to be buffered and can be forwarded only once the channel is free. As a packet is forwarded from hop to hop, each retransmission is not only received by the next, but also by the previous repeater in the chain. Nodes therefore need to keep some record such as not to re-enact the transaction based on an expired packet.

Related is the issue of broadcasting: Because of the unstable topology of wireless networks, there is essentially no way of ascertaining that a broadcast reaches all interested nodes, depicted in figure 4. There are a number of proposed algorithms that can efficiently reach coverage with high probability, but the standard approach taken by protocols is to let every routing node repeat the broadcast once. Compared to wired networks, broadcasts occupy the

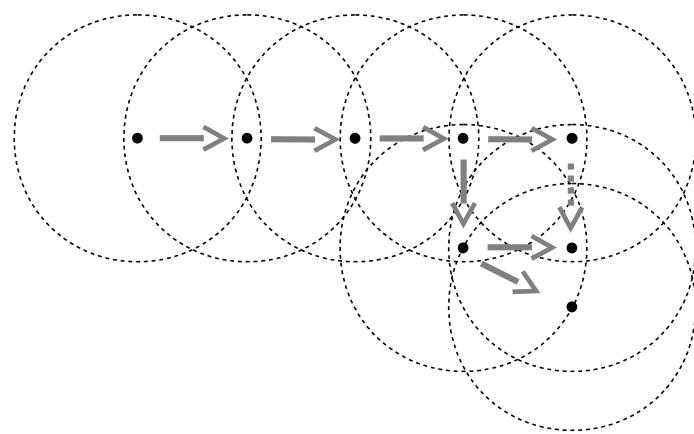


Figure 4: In wireless there is no inherent separation of network segments. Broadcasts are typically implemented by letting each routing node repeat the broadcast once, which causes the channel to be occupied for elongated periods.

channel for a long time and cause a comparatively huge overhead. Unfortunately, typical route-finding protocols rely heavily on broadcasting.

The open and volatile nature of the wireless channel is therefore met with increased sophistication at the media access and network layers of wireless protocols. This is necessary to achieve a level of service comparable to wired networks. Things that are relatively easy in wired networks, such as broadcasting, cause drastically more overhead in wireless networks.

As a result of all of the above, usage of channel bandwidth is an order of magnitude worse for wireless protocols than for wired protocols. Complexity of wireless protocols in terms of silicon area, lines of code and memory requirements is an order of magnitude higher for wireless protocols when compared to wired protocols of similar bandwidth and performance. In addition to that, section 2.3 describes common pitfalls that wire-replacement-projects commonly face: In short – usually the system architecture has to be re-engineered much more extensively than expected.

All this still does not mean that wireless is necessarily unreliable. But it means that in the indoor channel simple wireless links – planned, installed and tested – have a lower probability of performing flawlessly over the lifetime of the installation, than comparable wired links. In order to achieve comparable reliability, additional measures have to be taken across the different layers that result in increased complexity and costs.

Replacing wires with wireless therefore only makes economic sense if the increased complexity and associated costs can be offset by cost reductions. The cost reductions include not installing wires, but may also be based on efficiency gains, such as automatic address assignment.

2.3 The Problem with Cable Replacement

The costs or inconvenience of installing wires is typically the prime motivation for moving to wireless. When first exposed to the technology, many engineers fall into the trap of trying to design cable replacement systems. The resulting solutions often do not live up to the expectations. Cable replacement refers to a wireless point-to-point link with a simple MAC protocol acting as a drop-in replacement for a field bus segment.

The first issue of cable replacement systems is how the physics impact the reproducibility and predictability of system performance. Wired automation standards come with clear design rules with respect to maximum cable lengths, maximum number of sensors and actuators on a segment, maximum curvature radius of cables, maximum voltage drop on a segment, etc. These design rules are the result of careful analysis on the part of the designers of

standards. They are incorporated into automation systems at the planning stage by software tools and experienced planers. On-site installers install the system according to plan and test it. Testing in this sense means a simple functional test. The rationality is that if the design rules are respected and the systems passes a function test, then it will continue to work as planned over the design lifetime. There is a certain amount of degradation within the wire infrastructure, mainly due to corrosion of contact points, but those are accounted for in the safety margins of the design rules.

Design rules for wireless products lack the clarity of the wired counterparts. The specified parameters typically are related to the link budget, equivalent free space range, or the gain of possibly bundled antennas. Instead of robust design rules, planners need to work with estimated indoor ranges that are no more than rules of thumb. The burden of assuring performance is hence shifted to on-site installers. But even if additional figures like received signal strength are collected during the function test of the system, those figures merely have temporary significance. Wireless indoor channels do change over time, e.g. due to changes in utilization and furniture. Sometimes wireless indoor channels exploit propagation paths outside the structure and become subject to changing weather.

The significance of a passed function test as evidence of reliability is therefore much lower for wireless than for wired networks. In the long run, it is not uncommon for network segments with wireless cable replacement technology to experience transient outages, as well as an order of magnitude higher rates of packet loss compared with wired installations. The protocols, unless specifically designed for wireless, are not generally capable of dealing with the unreliable medium. The result is poor system performance. Put together, the modus operandi of planning and testing fails to produce reliable systems under the cable replacement paradigm.

Instead, systems need to be able to accommodate changes in the wireless channel, typically by building in enough redundancy so as to not rely on any one particular link. Possible solutions are various forms of mesh routing, but can also be frequency hopping à la Bluetooth which manages to avoid deep fades at particular frequencies. Requiring and supporting a certain level of redundancy also increases the ability to plan indoor networks, since indoor range estimates are always subject to statistics, as is shown for the data presented in section 5.2.3.

In the long run, vendors of wireless automation networks may need to adopt wireless network planning and simulation tools, such as those developed for mobile communication and wireless networking. These tools need to be integrated with building automation management systems in order to provide comprehensive support for wireless automation.

3 State of the Art

Automatic address assignment in wireless networks, such as the method investigated within the context of this work, is a multidisciplinary exercise. Relevant fields are telecommunications because of the wireless aspects, computer technology because of the field-bus and automation aspects, as well as computer science for the ray-tracing and graph-matching aspects. This section concentrates on the core aspect of commissioning and address assignment of building automation networks and reviews relevant works and literature. The other aspects receive a more terse treatment within the context of the relevant sections.

Because research and development in building automation happens primarily in industry and only to a lesser degree in academia, academic literature on the subject is scarce. It is therefore complemented by patent literature where applicable.

Commissioning works fundamentally different for professional installations and for enthusiast installations. The former typically use system engineering software to plan and execute the installation. Commissioning with software support is often called system level commissioning (S-mode in ZigBee parlance), commissioning semi-professional enthusiast systems typically uses E-mode “easy” commissioning. The latter does not require special software and hardware and follows procedures such as simultaneously pressing buttons to bind sensors and actors. Slightly more evolved procedures exist for grouping and scene definition.

Both, easy mode and system level commissioning, need to solve the address assignment problem. Easy mode does it implicitly: By simultaneously pressing buttons, sensor and actors learn each others' addresses, the roles are pre-determined and the actual location of the components is not relevant. System level commissioning is usually planned ahead of time, within the software tool. The virtual components within the software tool receive high-level addresses (IDs, names or labels). Installed components, on the other hand, receive low-level addresses (MAC addresses, physical or hardware addresses, serial numbers, hardware IDs). Address assignment then consists of creating an explicit mapping between low-level and high-level addresses, which is then saved into the database of the engineering software. The focus of the following analysis is on system level commissioning.

3.1 Address Assignment in Wired Building Automation Systems

An important distinction concerns the nature of the low-level addresses, often called MAC-addresses. The purpose of media access control (MAC) is to “ensure that the transmission medium is accessed and used in a fair way [HAL96]”. The role of the MAC address is to unambiguously identify data terminal equipment (DTE) connected to the medium. There are two trains of thought with respect to MAC addressing schemes:

- The first approach is to use only as many bits as necessary for MAC addressing. E.g. BACnet MS/TP uses 8 bits, DALI uses 6 bits, KNX physical addresses use 16 bits.
- The alternative approach is to use globally unique addresses as MAC address, such as those assigned by the IEEE Registration Authority on behalf of the ISO. 48 bit MAC addresses are used by amongst other as LonTalk serial number and by the Ethernet MAC. 64 bit MAC addresses are used amongst others by Firewire.
- IEEE 802.15.4 (the MAC layer underneath e.g. ZigBee) specifies both. Compliant networks can either use 64 bit globally unique “*long*” MAC addresses or 16 bit “*short*” MAC addresses.

In all cases, address assignment means associating locality to these low-level addresses. In the case of short MAC addresses there is the additional task of assuring uniqueness of the assigned MAC addresses. Assuring uniqueness is simpler than solving both problems. There are many solutions to automatically solve the unique address assignment problem:

- In DALI there is a procedure which lets devices on the bus chose random 24 bit temporary addresses and then find devices within the larger address space by exploiting the collision detection mechanism and binary search. Once the devices are isolated, 8 bit “physical” addresses are assigned [EN 62386-102].
- Grinberg convinces the US patent office that a virtually identical procedure applied to BACnet MS/TP satisfies the standards for novelty and non-obviousness [GRI11].
- Bohlmann et al. discloses a scheme for BACnet MS/TP, where “nomads” (devices without permanent MAC address assignment) choose random addresses from ranges that the “anchor” announces to be available. The anchor then searches the ranges and assigns permanent MAC addresses. In the event of a collision, the colliding nodes repeat the procedure until all nomads are found and reassigned [BOL08].
- ZigBee specifies two alternative schemes of assigning 16 bit short addresses. Under the cluster tree routing regime, joining nodes receive short addresses from the range administered by the parent handling the join. In ZigBee PRO the joining nodes get as-

signed a random address. Address collisions are detected and resolved using the 64 bit long addresses [ZIG08].

Ensuring uniqueness of low level addresses is the simple problem. Mapping the low level addresses to the location of the relevant devices is the really relevant step. This is generally a manual task.

In the lighting industry, address assignment is typically done after installation. First, the otherwise indistinguishable lighting components are installed and connected to the bus. During commissioning, the commissioning technician uses the DALI procedure described above to generate the low level addresses. Then, for each low level address, an IDENTIFY command is sent to the address which generally causes the associated lamp to blink. The commissioning technician then searches the blinking lamp by inspecting all possible fixtures connected to the bus, thereby mapping address to location. The commissioning technician repeats the step until all addresses are mapped.

The alternative procedure also generates low level addresses first. Then the commissioning technician visits each luminaire and briefly disconnects the lamp. The electronic ballast reacts by reporting the disconnection event to the controller node, thereby revealing its address.

Because the procedures described above are time-consuming and cumbersome, vendors have developed various proprietary extensions:

- Some DALI devices, especially emergency lights, can communicate their randomly assigned “physical” address by blinking a green and red signalling-LED conveying the raw binary data. The installer mimics the blinking pattern on a small commissioning tool (figure 5) featuring a red and green button and a display which shows the decoded numeric bus-address [WAL08].
- Often, unique serial numbers are printed on the device's housing. If the same serial number is queryable on the bus interface, then installers can read the serial number



Figure 5: Tridonic EM PRO Addressing Tool

off the housing at the time of installation, record it together with the location and derive the address mapping from the records. If the printed unique serial number is machine readable (e.g. a bar code), then the method can be semi-automated by a commissioning tool [HUB03].

- There are solutions where the devices can reveal their addresses using a particular blinking pattern [NOC10], or by invisibly modulating the light source [SHA06], [FER12].

The mainstream building automation systems also require the location of individual components to be known. In contrast to LonWorks and BACnet, KNX low level addresses already convey a measure of locality. The 16 bits are structured into area (4 b), line (4 b) and node (8 b), in KNX parlance separated by dots, as in “1.2.34”. The area / line / node tuple conveys topological location. This is mapped to physical location by the database of the network management system, a software called ETS (Engineering Tool Software), currently in version 4 and distributed by the KNX Association.

Address assignment in KNX is typically solved prior to or concurrently with installation. Before being connected to the bus, the components are labelled and programmed with their physical address [MEY13]. This may be done by the electrician on site, or by a contractor prior to delivery to the site. Once connected and powered, the parametrization of the components is done using ETS via the bus interface.

LonWorks uses globally unique, pre-programmed 48 bit IDs as low-level addresses. The network is configured and commissioned by a network management system such as Echelon LonMaker which needs to know the IDs and locations of the components. According to the LonMaker manual [LON06], address assignment can be done prior to mounting by entering or bar-code scanning the IDs and assigning them to the virtual representation of the components. It can be done during mounting using a hand-held tool to program unique “location strings” into the devices which are later extracted by network management. Address assignment is also supported after installation through network “discovery”: Components can reveal their ID to network management by physically pressing a “service pin”. Components can also be sent a “wink” command to blink or otherwise signal their “Wink LED”. LonWorks is therefore exceptionally flexible regarding available methods of address assignment.

Taking a more abstract perspective, independent of concrete building automation standards, the techniques described thus far can be assigned to one of three possibilities, also identified by Knauth et al.: Address assignment prior to, during or after mounting [KNA08].

- Address assignment prior to mounting: The mapping is produced off site, as soon as the hardware is available. Globally unique low level addresses are read off the com-

ponents, short low level addresses are programmed into components. A plan is devised that documents which of the individual components should be installed into exactly which position. The components may be labelled. The components and the plan is handed to the installer who is responsible for the plan to be followed without errors.

- Address assignment during mounting: The mapping is created on the fly, by the installer. He may be supported by a commissioning tool that can automatically read and record unique low level addresses from printed labels. Short low level addresses are programmed by the installer on the fly. The installer records the scanned or programmed addresses together with the locations, thereby creating the mapping.
- Address assignment after mounting: Otherwise undistinguishable components are mounted flexibly in any order. If necessary, short low level addresses are generated automatically. As part of the commissioning procedure, a technician visits every mounting location and records the low level address of the component thereby creating the mapping. Or the technician pushes a button or removes a lamp from a luminaire such that the component reveals its low level address to the network management, thereby creating the mapping. Or the network management sends some IDENTIFY command such that the technician can look for the blinking component, thereby creating the mapping. The important downside of address assignment after mounting is that components may already be covered by panels, or components may be mounted in inaccessible locations.

All of the above is executed by workers on a per devices basis, which is the major downside of all three approaches. It is a cost driver, a monotonous exercise, error-prone and requires extra training. The focus of this work, by contrast, is on fully automated address assignment. That are methods that do not require visiting all devices, scanning, labelling, or sorting of some sort. Manual steps are allowed, but only if they have network-wide impact.

Fully automated address assignment is state of the art in wired danger alarm systems. The difference between e.g. lighting and danger alarm system is that the latter require human readable location information to be readily available to emergency responders. All building automation services need to solve address assignment, but e.g. in lighting the details may stay confined to individual rooms or individual subsystems. Since manual address assignment is much more cumbersome if the bus spans the whole building, automatic address assignment is a higher priority for danger alarm systems. Relevant methods are:

- Alarm sensors are equipped with a relay switch that can be used to disconnect the tail of an otherwise linear bus from the central control unit. This is required by regulations to isolate bus faults. Following a broadcast command all alarm sensors open their relays. At that point, the control unit is connected only to the first sensor of the

daisy chain. The sensor is assigned an ID and then locked so that it doesn't react to any further ID assignment commands. The sensor is instructed to close its relay and the second sensor in the daisy chain comes online to be assigned [BUH82].

- If the previous method is combined with unique serial numbers and collision detection, the procedure can also deal with branches on the bus, at least so long as asymmetries of the topology can be used to distinguish between branches [UKA05].
- Berger et al. propose a solution where first a 24 bit address space is searched for unique IDs using collision detection and binary search. Once the IDs are known, each alarm sensor, one at a time, is instructed to emit a current pulse. The other devices record the IDs and direction (sign) of the corresponding current pulses. The information is transmitted back to the control unit which can reconstruct the topology [BER90].

3.2 Address Assignment in Wireless Building Automation Systems

Wireless automation systems are subject to an additional distinction: Whether nodes are mobile or not. If there are mobile nodes, it is necessary to distinguish between the role of the node in the network and the location of the node. For stationary systems, the node's location (as in mounting fixture) typically defines their role and vice versa. The situation in stationary wireless systems is therefore no different than that of wired systems. Wireless building automation networks are for the most part stationary, except for wireless remote controls. The location of mobile nodes, on the other hand, is the subject of localisation systems, discussed in section 3.3.

As explained in the introduction, the professional building automation market is conservative, and wireless technologies are a slowly evolving niche. Typically, wireless serves as an extension to a traditional wired building automation system. Fully wireless systems are more the domain of semi-professional and enthusiasts markets.

KNX-RF standardizes the KNX protocol over wireless media. It supports sending-only and bidirectional components. In contrast to standard KNX, KNX-RF uses globally unique 48 bit low level addresses. The details of interoperability are quite complicated:

Sending-only components use pre-programmed, invariant, globally unique 48 bit addresses (home and building electronic systems “HBES” serial numbers, [EN 50090-5-3]). Traditional KNX address fields are present but instead of their usual role are merely used to distinguish between endpoints of the sending-only component. 48 bit serial number plus 16 bit endpoint address together form the extended group address. Receiving (bi-directional) components

learn the extended group addresses by simultaneous button pressing (E-mode commissioning).

Bi-directional components, especially when part of a larger KNX installation, use an 48 bit RF domain address, instead of the serial number. The domain address is used to tell neighbouring installations apart, and typically equals the serial number of the primary bus coupler between KNX and KNX-RF. With a domain address present, the 16 bit of the traditional KNX address are used to index translation tables for KNX to KNX-RF interoperability. Bi-directional components can be bound to bus couplers, at which point they can be commissioned using the network management software, the same as wired KNX components. The bus coupler is responsible for address translation between the KNX addresses and extended KNX-RF addresses [EN 50090-5-3] [GRA06].

Address assignment is solved by KNX-RF in one of two ways: Simultaneous button pressing (E-mode) causes the sending component to transmit its address to the receiving component. S-mode is also possible, but may require vendor specific extensions¹: If mounting location and serial number of KNX-RF components is known and entered into the network management database, ETS (engineering tool system) can compute the address translation tables, program the bus-couplers and from there access the KNX-RF domain to program and commission KNX-RF components.

ZigBee is an industry driven, semi-open standard. Industry-driven means that it was created bottom-up, with support from the semi-conductor industry, and designed to fit into as many applications as possible. As a downside of the formidable flexibility, a common critique goes, ZigBee cannot really excel at any of the many applications, because too many compromises had to be made. It is semi-open, because the standard is freely downloadable, but protected by copyright. Selling ZigBee products requires joining the ZigBee Alliance and paying membership fees. ZigBee is subdivided into a family of standards, with domain specific profiles for the different applications. ZigBee is fairly common in academic research of sensor networks. The experiments of section 5 were conducted with a ZigBee network.

There are several ZigBee profiles with relevance to building automation²:

- ZigBee Home Automation – Scope: Enthusiast and semi-professionals, currently little market uptake, although major manufacturers (Schneider Electric, LG, Philips) have certified pilot products; wide range of devices specified: lighting, shading, HVAC, intruder alarm, locking; E-mode commissioning [ZHA10]

¹ Example: Hager Tebis TL131A media coupler, proprietary ETS plugin

² <http://www.zigbee.org/Standards/Overview.aspx>

- ZigBee Building Automation – Scope: Professional market, apparently little market uptake, probably defunct; in addition to ZigBee Home Automation supports BACnet tunnelling and S-mode commissioning [ZBA11]
- ZigBee Light Link – Scope: End users & enthusiasts, first products available by a major manufacturer (Philips, Apple, Osram); basically ZigBee Home Automation with co-ordinator-less commissioning and extended support for colour lighting and colour temperature lighting [ZLL12]
- ZigBee Smart Energy – Scope: Utility companies deploying advanced metering infrastructure, major deployments under way; currently in version 2 with supports an alternative IP-based network stack in addition to ZigBee infrastructure [ZSE08]

Outside the standardized profiles, there are incompatible proprietary ZigBee solutions. E.g. LeGrand has a proprietary home automation product line using ZigBee networking.

As described in subsection 3.1, ZigBee has both short 16 bit low-level addresses and globally unique 64 bit low-level addresses. For normal communication, short addresses are used. During network forming and for the purposes of commissioning, conflict resolution and cryptography, long addresses are used to unambiguously identify components [ZIG08].

ZigBee explicitly supports easy (E-mode), system-level (S-mode) and automatic (A-mode) commissioning and hence address assignment. A-mode typically means that the devices are sold pre-commissioned, and no further human interaction is necessary, hence the name automatic. E-mode means the usual simultaneous button pressing, in the case of ZigBee Light Link combined with rudimentary proximity sensing.

S-mode commissioning is used in ZigBee Building Automation and ZigBee Smart Energy. The former specifies use of a commissioning tool, as well as vendor specific out-of-band methods. The latter goes into details on possible scenarios, e.g. multiple utility meters sharing a common metering infrastructure: The electric meter is connected to a wide area network via power line communication and is at the same time the ZigBee Smart Energy coordinator and network trust centre (“Energy Services Portal”). Water meter or gas meter are ZigBee Smart Energy clients. Commissioning is solved by the utility companies communicating with each other, the e-meter hence learning the client-meter’s ID and trusting the client-meter’s certificate (“Certificate-Based Key Establishment”). In-premises displays for ZigBee Smart Energy can also join the network thanks to an IT-process and communication with the utility company. Address assignment and security is hence solved at the system level – either by a B2B process between utility companies, or by the consumer entering the in-premise display’s ID into a web-form or call centre system [ZSE08].

Enocean is a semi-open wireless building automation system optimized for extremely low power, energy-harvesting applications. Semi-open means that specifications are freely downloadable, but protected by copyright and patents. Implementing the standard requires joining the Ennociant Alliance and paying license fees. Address assignment and commissioning (“teach-in”) is done by simultaneous button pressing. Communication is mostly uni-directional, there is support for repeaters and bi-directional remote management. Enocean has a market in commercial-grade functional buildings and there are gateways to KNX, LonWorks, BACnet and DALI that expose the Enocean wireless domain to mainstream building automation³.

Other wireless building automation protocols, such as those surveyed in [REI07] and [GRA11], generally target the home automation market and therefore come without support for large-scale commissioning methods. Prevalent commissioning methods are push-button, as well as commissioning tools, sometimes integrated into controllers, with various forms of user interfaces.

Jost et al. patent a wireless commissioning tool that uses indoor localisation and knowledge of the deployment plan to solve the address assignment problem. The localisation method itself remains unspecified, but the claim-section enumerate localisation based on path-loss, time of flight and triangulation. If the localisation relies on the same building automation network that is undergoing commissioning, then there is a hen-and-egg problem, because localisation cannot work prior to address assignment. Apparently a subset of the network nodes are first commissioned manually in order to bring the localisation function online [JOS07].

3.3 Related Work – Localisation Systems

Node localisation is a related problem: It assumes no prior knowledge of mounting locations, or at least, does not make use of it. The field of node localisation is of substantial breadth. Based on the taxonomies of [SAV04], [HIGH01], [STO07] and [HLIU07] the presented work uses:

- Active localisation – all system components emit signals. The cooperation of localisation target and infrastructure is needed to compute locations. The algorithm is centralized.
- Sensing via RF RSS – measurements of received signal strength (RSS) or alternatively time of flight (TOF). In contrast to many systems based on RSS measurements, no empirical relationship between measured data and distance is assumed. The sys-

³ <http://www.enocean.com/en/white-papers/>

tem needs no dedicated hardware other than what is available in industry standard sensor nodes.

- Multi-hop localisation based on radio connectivity, multilateration, without anchors – physical locations are calculated based on the devices position in the networks connectivity graph. In contrast to other works, the usage of anchors – special nodes with preconfigured location – is optional.
- Locations are expressed symbolically (e.g. floor / room / fixture) and relative to the particular structure or building.

For indoor localisation based on signal attenuation, there are several approaches for overcoming the challenges presented by multi-path propagation: Fingerprinting, relies on building a detailed channel map of the building [PAP09] [BAH00] [KIN06]. This relies on first installing a number of anchors that broadcast regular beacons and then taking systematic measurements based on a regular, dense grid. During the operation phase of the system, locations are calculated by measuring the beacons and comparing measurements to the map. [KIN06] observes distance errors of less than 3 m at the 90th percentile using a dense network of WiFi access points as beacons. This approach is not well applicable to address assignment, since building the required channel map requires a higher effort than manual address assignment.

Alternatively, statistics can be used to combat the problems of multi-path propagation by using a very dense network of anchors, up to several per room. While multi-path propagation breaks the inverse square-law relationship between signal attenuation and distance for any particular measurement, it still applies on average. Zanca et al. compares several localisation algorithms [ZAN08] [PAT01] [CLIU07]. In the case of 5 beacons per room the average localisation error is no better than 2 m. This may be adequate for address assignment, the overhead of installing several anchors per room, however, is not.

Anchor selection algorithms, finally, combat the effects of multi-path propagation by sorting anchors into those having line-of-sight to the localisation target and those that have not [WAN12] [SIN07]. The high number of required anchors, again, makes the approach impractical for address assignment.

4 A New Method for Automatic Address Assignment

This section goes into the details of automatic address assignment. The problem is formalized (section 4.1). The major influencing factors, namely the network topology (section 4.2) and the indoor channel and its simulation (section 4.3) are studied. An efficient method for channel measurement is described (section 4.4). Preprocessing (section 4.5) and matching simulation against measurement (section 4.6) conclude the treatment.

4.1 Formalization

The informal description of the three-step process of address assignment is given in section 1.1. This section builds on the previous definitions of the *measured* and *simulated connectivity graph*. Both are weighted graphs, henceforth represented by their adjacency matrices \mathbf{G}_M and \mathbf{G}_S . The weights of the graphs, i.e. the matrix entries of \mathbf{G}_M and \mathbf{G}_S are realizations of random variables because the underlying exercise of simulation and measurement of wireless networks introduces intentional and unintentional deviations that are best modelled stochastically. Section 4.2 takes a closer look on the statistics.

Underlying the random nature of \mathbf{G}_M and \mathbf{G}_S is an ideal network connectivity graph \mathbf{G}_I . The relationship is given by:

\mathbf{G}_I Ideal connectivity graph	$\mathbf{G}_M = \mathbf{P}_M^\top \mathbf{G}_I \mathbf{P}_M + \mathbf{N}_M$ Measurement	$\mathbf{G}_M = \mathbf{P}^\top \mathbf{G}_S \mathbf{P} + \mathbf{N}$ Relation between \mathbf{G}_M and \mathbf{G}_S <i>(subject to the substitutions)</i>
	$\mathbf{G}_S = \mathbf{P}_S^\top \mathbf{G}_I \mathbf{P}_S + \mathbf{N}_S$ Simulation	$\mathbf{P} = \mathbf{P}_S^\top \mathbf{P}_M, \quad \mathbf{N} = \mathbf{N}_M - \mathbf{P}^\top \mathbf{N}_S \mathbf{P}$

$\mathbf{G}_I, \mathbf{G}_M, \mathbf{G}_S$ Adjacency matrices of the ideal, measured and simulated connectivity graph; symmetric, square; represent negative attenuation loss of the channel in dB (e.g. -75 dB)

$\mathbf{P}_M, \mathbf{P}_S, \mathbf{P}$	Permutation matrices with the sought after mapping (address assignment); permutation matrices are square binary matrices with exactly one “1” in each column and row
$\mathbf{N}_M, \mathbf{N}_S, \mathbf{N}$	Deviation between measurement and simulation in dB; symmetric, square matrices; the analysis of section 4.2 treats the elements of \mathbf{N} as random variables

Because graph vertices do not exhibit any natural order, and because measured and simulated connectivity graph are assembled from different data sets, they are permuted with respect to the ideal connectivity graph by random permutation matrices \mathbf{P}_M and \mathbf{P}_S . Inaccuracies of measurement and simulation are captured by the noise terms \mathbf{N}_M and \mathbf{N}_S . Unlike noise in the sense of stochastic processes and signal processing, they lack a linear, time-like component. The term *noise* is used never the less, because the \mathbf{N} -terms are a realization of random processes and serve the purpose conceptualizing the influence of randomness on the system.

The rightmost part of equation 1 describes the relation between measured and simulated connectivity graph as function of the mutual permutation \mathbf{P} and the combined noise \mathbf{N} . This is done by solving the simulation equation for \mathbf{G}_I and substituting the latter in the measurement equation. The resulting representation is advantageous, because \mathbf{G}_I and its dependent quantities are not directly observable, while the realizations of \mathbf{P} and \mathbf{N} are.

Equation 1 is a physical equation in the sense that it states the relationship between realizations of measured and simulated connectivity graph, permuted by a particular permutation and affected by a particular noise realization. In other words, the noise \mathbf{N} is the deviation between measurement and simulation given the correct permutation \mathbf{P} . In order to study the impact of permutations other than \mathbf{P} , the following function is defined:

$$\mathbf{D}(\mathbf{Q}) = \mathbf{G}_M - \mathbf{Q}^\top \mathbf{G}_S \mathbf{Q} \quad (2)$$

Function $\mathbf{D}(\mathbf{Q})$ is structurally equivalent to equation 1, but the independent variable \mathbf{Q} takes the place of \mathbf{P} , and the dependent variable \mathbf{D} (for deviation) takes the place of \mathbf{N} . If the permutation \mathbf{Q} is set to the correct deviation \mathbf{P} , then \mathbf{D} is equal to \mathbf{N} .

Since measured and simulated connectivity graph are attributed by respectively the low level addresses and the high level addresses, finding the correct address assignment is therefore equivalent to recovering the permutation \mathbf{P} . This can be achieved if measurement and simulation are sufficiently accurate and therefore \mathbf{N} zero-mean and of relatively small variance. Under these circumstances, \mathbf{P} can be recovered by minimizing the deviation $\mathbf{D}(\mathbf{Q})$:

$$\hat{\mathbf{P}} = \arg \min_{\mathbf{Q}} \|\mathbf{D}(\mathbf{Q})\|_F^2 = \arg \min_{\mathbf{Q}} \|\mathbf{G}_M - \mathbf{Q}^\top \mathbf{G}_S \mathbf{Q}\|_F^2 \quad (3)$$

P Permutation matrix recovered by the optimization

Q Independent variable of the objective function

Froebius Norm: $\|\mathbf{X}\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |x_{ij}|^2}$

Equation 3 is a least squared error optimization. The likelihood of finding the correct permutation, i.e. $\hat{\mathbf{P}} = \mathbf{P}$ and $\mathbf{D}(\hat{\mathbf{P}}) = \mathbf{N}$, depends on the following factors:

- The topology of the network and its fitness to provide adequate separability (subsection 4.2)
- The quality of measurement and simulation, i.e. the shape of \mathbf{N} (subsection 4.3)
- The ability of the matching algorithm to find the global minimum (subsection 4.6)

The noise \mathbf{N} can be further subdivided into a global offset, vertex-noise and edge-noise:

$$\mathbf{G}_M = \mathbf{P}^\top \mathbf{G}_S \mathbf{P} + off \mathbf{J}_{SP} + \mathbf{N}_V \mathbf{J}_{SP} \mathbf{J}_{SP} \mathbf{N}_V + \mathbf{N}_E \quad (4)$$

J_{SP} Logical matrix indicating the non-zero elements of $\mathbf{P}^\top \mathbf{G}_S \mathbf{P}$

$$\mathbf{J}_{SP} = \text{NZ}(\mathbf{P}^\top \mathbf{G}_S \mathbf{P}), \text{ with } \text{NZ}(\mathbf{M}): \text{nz}(\mathbf{M})_{ij} = \begin{cases} 1 & \text{if } m_{ij} \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

off J_{SP} Global offset between weights in measured and simulated connectivity graph;
off is scalar, in dB

N_V J_{SP} J_{SP} N_V Vertex noise, systematic offset affecting edges connected to a particular vertex;
 \mathbf{N}_V is diagonal, in dB

N_E Edge noise, any deviation between graph weights not better captured by global offset or vertex noise, in dB

All matrices are square matrices $m \times m$, with m being the number of vertices i.e. the number of nodes to be matched. Given the high dynamic range of physical quantities in wireless technology, the logarithmic notation seems adequate, also given typical measurement accuracy. As secondary effect, the logarithmic notation linearises the inverse square law.

The global offset off is a convenience device to keep vertex and edge noise zero-mean. While in theory it is possible to reach the same goal by eliminating all systematic sources of error from measurement and simulation, it also means a lot of effort. For practical reasons, the channel simulation does not cover all physical phenomena but only the most important. Conveniently, the global offset is relatively simple to estimate (cf. section 4.5).

Vertex noise captures all deviations between measurement and simulation affecting a specific node. The physical phenomena underlying those deviations are related to the node itself, and its immediate vicinity: most importantly the antenna and its misalignment or mistuning due to inappropriate placement or defects. Section 4.5 discloses a method of cancelling the effect of vertex noise.

Vertex noise and edge noise as defined above are expected to be Gaussian. This is due to the central limit theorem in the logarithmic domain: The de-logarithmized equivalents of the adjacency matrices $\exp(\mathbf{G}_M)$ and $\exp(\mathbf{G}_S)$ are the result of a multiplicative product of uncertain quantities (material coefficients, physical dimensions, reflections). $\exp(\mathbf{G}_M)$ and $\exp(\mathbf{G}_S)$ are expected to adhere to the log-normal distribution and hence \mathbf{G}_M and \mathbf{G}_S , as well as \mathbf{N}_V and \mathbf{N}_E , Gaussian. This is backed by experimental results of section 5. Section 4.2.4 has a more detailed analysis of the assumptions underlying the noise model and in how far they are justified.

As stated before, measured and simulated connectivity graphs are attributed by low-level and high-level addresses. In addition, other facts about the network nodes can be modelled as vertex attributes and used in the least squared error approximation. Especially node types, such as “lighting control”, “lighting fixture”, “temperature sensor”, etc. are known in the simulation, and are also gathered by the measurement protocol. These attributes can serve as powerful node discriminators and therefore improve reliability. Attributes are modelled as:

$$\mathbf{T}_{kM} = \mathbf{P}^T \mathbf{T}_{kS} \mathbf{P} \quad (5)$$

\mathbf{T}_{kM} and \mathbf{T}_{kS} are diagonal matrices holding one of the $k = 1..na$ vertex attributes common to measured and simulated connectivity graphs. They are accounted for in the least squared error optimization by modifying equation 3:

$$\hat{\mathbf{P}} = \arg \min_{\mathbf{Q}} \| \mathbf{G}_M - \mathbf{Q}^T \mathbf{G}_S \mathbf{Q} \|_F^2 + \sum_{k=1}^{na} v_k \| \mathbf{T}_{kM} - \mathbf{Q}^T \mathbf{T}_{kS} \mathbf{Q} \|_F^2 \quad (6)$$

The constants v_k serve to balance the optimization goals of correctly matching the attributes and correctly matching the graph weights.

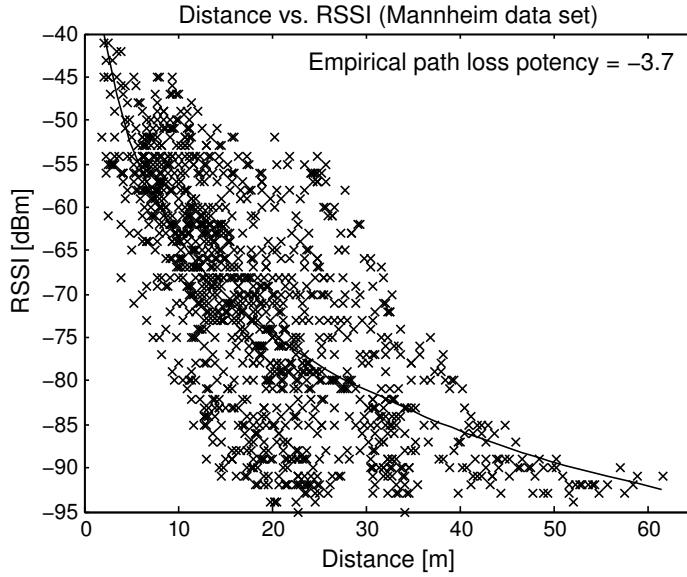


Figure 6: Distance vs. RSSI of the test data set. The solid line is an empirical inverse-square-law model fitted to the data, with an empirical path loss potency of -3.7.

Inferring distance directly from the values in \mathbf{G}_M and \mathbf{G}_S is essentially futile, even in the line-of-sight (LOS) case. Ignoring multi path, antenna gains and insertion losses, the RSSI reported as the result of a 0 dBm transmission at distance d in the 2.4 GHz band is given by the link budget equation: $RSSI = -20 \log_{10}(d) - 40.23$. While it is possible to distinguish between transmissions 1 m and 2 m away ($\Delta RSSI = 6$ dB), it is not possible to distinguish between 10 m and 11 m ($\Delta RSSI = 0.8$ dB), which is below the reporting accuracy of typical transceivers. Taking into account effects such as non-line-of-sight (NLOS) or wave-guiding along corridors, we arrive at figure 6: Distance vs. RSSI as found in the test data of King et al., as extensively used in section 5.1 [KIN08].

Sections 4.3 and 4.4 treat respectively channel simulation and channel measurement, section 4.5 preprocessing and vertex noise cancellation, section 4.6 the weighted graph-matching algorithm.

4.2 The Impact of Network Topology

The geometrical layout of the network within the building, or network topology, has an impact on the ability of the least squared error approximation to recover the correct permutation. This section deals with the theory of what topologies are suitable to be commissioned with automatic address assignment, what is the maximum amount of noise acceptable for a

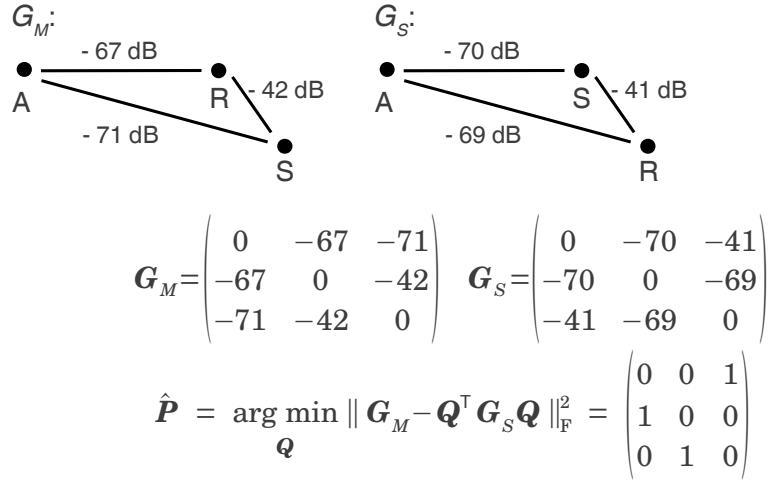


Figure 7: Example network of three nodes (A, R, S), moderate noise. The optimization has a unique solution.

particular topology and how to lay out networks so that automatic address assignment works as intended.

To illustrate the problem, figure 8 shows the most simple network that cannot be handled by the presented approach. Networks of three nodes, like the one depicted in figure 7, have a unique solution, provided there are no symmetries and the noise variance is small enough. With the numbers chosen in the example, $g_{S1,2} < g_{S2,3}$, 1 dB apart. If, due to noise, $g_{S1,2} > g_{S2,3}$, then the optimization would find a different $\hat{\mathbf{P}}$.

The topic of this section is to study the relationship between

- Network topology
- Noise variance
- Probability of matching error

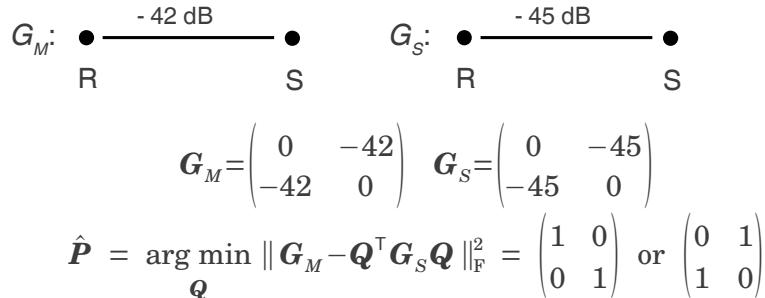


Figure 8: Example network of two nodes (R, S), measured and simulated connectivity graph. The optimization is inconclusive, because both alternatives have equal costs.

The latter has been solved for structural (unweighted) graphs by [LUO01], but not for weighted graphs such as \mathbf{G}_M and \mathbf{G}_S . In formal terms: Sought after is $P\{\varepsilon\} = P\{\hat{\mathbf{P}} \neq \mathbf{P}\}$ – the probability that the recovered permutation is not the actual permutation. These matching errors occur, if the optimization (equation 3) recovers the wrong permutation, matrix. Formally:

$$P\{\varepsilon\} = P\{\|\mathbf{G}_M - \hat{\mathbf{P}}^\top \mathbf{G}_S \hat{\mathbf{P}}\|_F^2 < \|\mathbf{G}_M - \mathbf{P}^\top \mathbf{G}_S \mathbf{P}\|_F^2\}$$

Using equation 1 and the substitution $\hat{\mathbf{P}} = \mathbf{P} \mathbf{P}_\Delta$, where \mathbf{P}_Δ is itself a permutation matrix capturing the potential matching error, the problem is reformulated as:

$$P\{\varepsilon\} = P\{\|\mathbf{G}_M - \mathbf{P}_\Delta^\top \mathbf{G}_M \mathbf{P}_\Delta + \mathbf{P}_\Delta^\top \mathbf{N} \mathbf{P}_\Delta\|_F^2 < \|\mathbf{N}\|_F^2\} \quad (7)$$

The probability of a particular matching error, the mismatch given by \mathbf{P}_Δ , is thus expressed as function of \mathbf{G}_M and \mathbf{N} . A dual formulation in terms of \mathbf{G}_S and a modified noise term can be found with appropriate substitutions. For the following analysis, \mathbf{G}_M is treated as deterministic and \mathbf{N} as matrix of random variables.

When applying the Froebius Norm, all entries of $(\mathbf{G}_M - \mathbf{P}_\Delta^\top \mathbf{G}_M \mathbf{P}_\Delta)$ that are not permuted by \mathbf{P}_Δ cancel each other. The same applies to entries of \mathbf{N} that are not permuted by \mathbf{P}_Δ , because they equally appear on both sides of the inequality.

To simplify matters, \mathbf{P}_Δ is henceforth restricted to a transposition matrix, which is a permutation matrix that swaps exactly two vertices. This entails a moderate loss of generality, because cases can be constructed where a matching error involves the false attribution of three or more vertices whereas any pairwise permutation does not make the inequality true. The simplification, on the other hand, allows to produce the visualizations below and therefore facilitates a thorough analysis of the problem.

With \mathbf{P}_Δ restricted to a transposition matrix (swapping the vertices numbered r and s) and using the symmetry of \mathbf{G}_M and \mathbf{N} , the problem is restated in element-wise notation:

$$P\{\varepsilon\} = P\left\{\sum_{i \notin \{r, s\}} [(g_M)_{ir} - (g_M)_{is} + n_{is}]^2 + \sum_{i \notin \{r, s\}} [(g_M)_{is} - (g_M)_{ir} + n_{ir}]^2 < \sum_{i \notin \{r, s\}} n_{ir}^2 + \sum_{i \notin \{r, s\}} n_{is}^2\right\} \quad (8)$$

The notation exploits the fact that most summands in equation 7 cancel when applying the subtraction and Froebius norm. Equation 8 includes only those summands affected by the transposition \mathbf{P}_Δ , and is also exploiting symmetry of \mathbf{G}_M and \mathbf{N} . The situation is exemplified below for $r = 1$ and $s = 3$. Elements that cancel are crossed out:

$$\| \mathbf{M} - \mathbf{P}_{\Delta 13}^T \mathbf{M} \mathbf{P}_{\Delta 13} \|_F = \left\| \begin{pmatrix} \theta & m_{12} & \overline{m_{13}} & m_{14} & m_{15} \\ m_{21} & \theta & m_{23} & \overline{m_{24}} & \overline{m_{25}} \\ \overline{m_{31}} & m_{32} & \theta & m_{34} & m_{35} \\ m_{41} & \overline{m_{42}} & m_{43} & \theta & \overline{m_{45}} \\ m_{51} & \overline{m_{52}} & m_{53} & \overline{m_{54}} & \theta \end{pmatrix} - \begin{pmatrix} \theta & m_{32} & \overline{m_{31}} & m_{34} & m_{35} \\ m_{23} & \theta & m_{21} & \overline{m_{24}} & \overline{m_{25}} \\ \overline{m_{13}} & m_{12} & \theta & m_{14} & m_{15} \\ m_{43} & \overline{m_{42}} & m_{41} & \theta & \overline{m_{45}} \\ m_{53} & \overline{m_{52}} & m_{51} & \overline{m_{54}} & \theta \end{pmatrix} \right\|_F$$

Equation 8 can be rewritten as:

$$P\{\varepsilon\} = P\left\{ \sum_{i \notin \{r, s\}} \Delta(g_M)_i^2 \left[1 - \frac{\Delta n_i}{\Delta(g_M)_i} \right] < 0 \right\} \quad (9)$$

using: $\Delta(g_M)_i = (g_M)_{ir} - (g_M)_{is}$, $\Delta n_i = n_{ir} - n_{is}$

The equation has an intuitive explanation: \mathbf{G}_M and \mathbf{G}_S are adjacency matrices of graphs. The $\Delta(g_M)_i$ denote separability of vertices within the one graph \mathbf{G}_M by comparing the edge-weights connecting the vertices under test R and S to a third witness vertex I . This weight-difference is then put in relation to the noise-level Δn_i which is the expected noise-difference of corresponding edges between the graphs. The relationship between the variables is depicted in figure 9. The role of the sum is to aggregate separability of vertices R and S relative all other vertices of \mathbf{G}_M and the expected deviation of those edge-weights relative \mathbf{G}_S .

In order for the inequality to become true, i.e. the vertices R and S to become swapped in the least squared error optimization, the value of the sum must become negative which happens

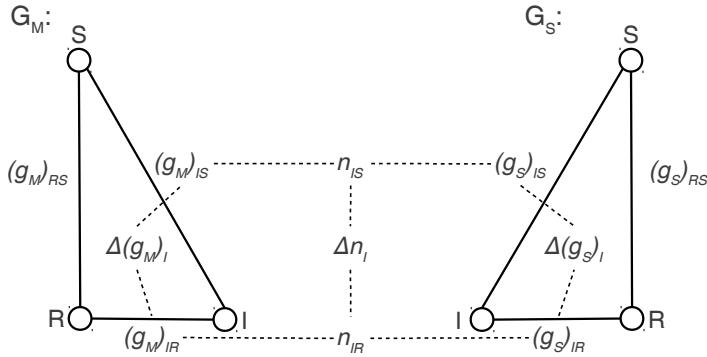


Figure 9: Illustration of the variables used in equation 9. $(g_M)_{XX}$ and $(g_S)_{XX}$ are edge-weights from the adjacency matrices \mathbf{G}_M and \mathbf{G}_S . $\Delta(g_M)_I$ is the difference between edge-weights connecting the witness node I to nodes under test R and S . n_{XX} are deviations of edge-weights between the graphs and Δn_I is the difference between deviations concerning the witness node I relative the nodes under test R and S .

when a qualified majority of the summands is negative. Each summand compares the weight difference of a pair of edges (RI , SI) to the noise difference of the same pair of edges. The summand is negative if the absolute value of the noise-difference is bigger than the absolute value of the weight difference, also considering the sign.

Equation 9 is a very important result, because it allows to study concrete topologies for their susceptibility to noise, and systematically calculate error probabilities. The remainder of the section presents the consequences of equation 9 as a series of examples and visualisations.

The graphs \mathbf{G}_M and \mathbf{G}_S represent channel loss between wireless nodes. In order to visualize those topologies, channel loss models have to be mapped to Cartesian space for visualization. Instead of using the ray-tracer, as documented in section 4.3, this section uses the inverse square law to define the equivalent free-space distance \mathbf{G}_{FSD} . The elements of \mathbf{G}_M and \mathbf{G}_{FSD} are related as:

$$(\mathbf{g}_M)_{ij} = \begin{cases} 0 & \text{if } i=j \\ k + s \log_{10}((\mathbf{g}_{FSD})_{ij}) & \text{otherwise} \end{cases} \quad (10)$$

The constant k is -40.25 in the 2.45 GHz band, but its value is not needed for error probabilities or visualizations. The constant s means the effective inverse square and is set to -2 for the remainder of the section.

It is usually not possible to map the equivalent free-space distances of a typical indoor network to Cartesian space. But as stated before, even if the inverse-square-law is a bad approximation for any particular link in the indoor wireless network, it still applies on average. Also, the inverse square law is only used in this section to generate visualisations of equation 9. The inverse square law was not used in the derivation.

4.2.1 Probability of Matching Error in the Three-Node Network

Under consideration is the following network of three points:

$$\vec{R} = \begin{pmatrix} -d/2 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{S} = \begin{pmatrix} d/2 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{A} = \begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad \text{subject to} \quad \begin{cases} |\vec{A} - \vec{R}| > d \\ |\vec{A} - \vec{S}| > d \end{cases} \quad (11)$$

d is the distance between R and S and simultaneously the minimal distance between any two nodes in the network. The first task is to express the probability of mismatching nodes R and S in relation to the noise variance and the position of an arbitrary witness node A . The latter should be placed away from R and S , at a distance greater or equal d , because node A should

act as a witness and not as likely target of a matching error itself. The corresponding free-space equivalent distance graph is given by

$$\mathbf{G}_{FSD} = \begin{pmatrix} 0 & d & \sqrt{(x+d/2)^2+y^2+z^2} \\ d & 0 & \sqrt{(x-d/2)^2+y^2+z^2} \\ \sqrt{(x+d/2)^2+y^2+z^2} & \sqrt{(x-d/2)^2+y^2+z^2} & 0 \end{pmatrix},$$

where nodes R , S and A are represented by the respectively the first, second and third column. Applying equation 10 and equation 9 leads to the to the probability of swapping R and S :

$$P\{\varepsilon\} = P\left\{1 - \frac{\Delta n_A}{\Delta(g_M)_A} < 0\right\} \quad (12)$$

using: $\Delta(g_M)_A = -10\log_{10}\left(\frac{(x+d/2)^2+y^2+z^2}{(x-d/2)^2+y^2+z^2}\right)$, $\Delta n_A = n_{AR} - n_{AS}$

If the non-diagonal elements of matrix \mathbf{N} are Gaussian random variables, independent from each other, with variance σ_N^2 , then the variance of Δn_A is $2\sigma_N^2$. The assumption of statistical independence is of course problematic. On the other hand, potential interdependence of the random variables comes from hundreds of individual factors such as imprecise or simplified simulation parameters. Like the assumption of Gaussianity, which stems from the central limit theorem, it is a practical necessity in the context of this section. After all, the current purpose is to study the matching error probability using formal methods – a goal which justifies certain simplifications.

Further simplifying the equation requires discriminating on basis of the sign of $\Delta(g_M)_A$:

$$P\{\varepsilon\} = \begin{cases} P\{\Delta n_A < \Delta(g_M)_A\} & \mid \Delta(g_M)_A < 0 \\ P\{\Delta n_A > \Delta(g_M)_A\} & \mid \Delta(g_M)_A > 0 \\ 0.5 & \mid \Delta(g_M)_A = 0 \end{cases} \quad (13)$$

The probability of matching error thus boils down to a comparison between separability $\Delta(g_M)_A$ of the nodes R and S relative the witness node A , and the noise which separates corresponding edges of the graphs \mathbf{G}_M and \mathbf{G}_S . The regions of zero separability are defined to have a probability of matching error of 50%, regardless which assignment for R and S is presumed.

The situation is in fact analogous to digital signal detection. The probability of matching error is analogous to bit error rate, separability is analogous to the received signal energy per

bit and noise is analogous to the spectral noise density. The separability to noise ratio (SNR) could be defined as $\text{SNR} = \Delta(g_M)_A / \Delta n_A$. At this point, however, the analogy ends. In digital signal detection, the bit error rate is dependent not only on SNR, but also on the modulation. There is no analogy for the concept of modulation in address assignment.

Figure 10 depicts the shape of $\Delta(g_M)_A$ (“signal”) as a function of the independent witness point A (x/y/z). Several facts can be gleaned from the plot:

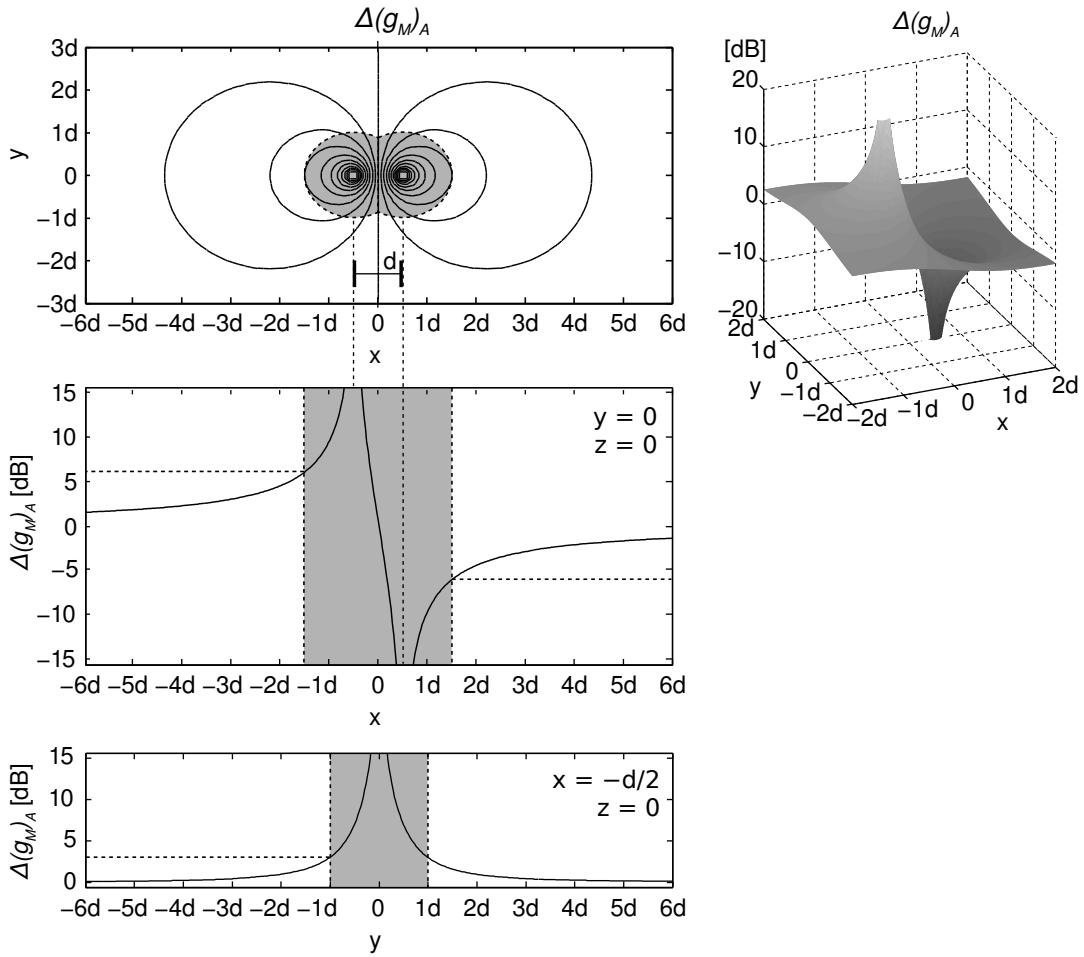


Figure 10: Separability of the nodes R ($-1/2d/0/0$) and S ($1/2d/0/0$) relative a third node A (x/y/z), the independent parameter: Top right – surface detail cut through x-y-plane. Top left – surfaces of equal separability are spheroids lined up along the x-axis with R and S located at the common focal points. Separability at the y-z-plane (a degenerate spheroid) is zero. Middle left – separability along the x-axis. Bottom left – separability along the y-axis at $x = -d/2$. The grey area is problematic for placement of A, because it violates the definition of d (cf. equation 11). The maximum attainable separability between R and S due to a single witness node is 6 dB along the x-axis and 3 dB along the y-axis.

- The grey area of figure 10 is the space forbidden by the problem definition (equation 11), where node A stops being a witness towards separability of R and S and instead becomes a likely target of a matching error itself
- The contribution to separability is maximal, if the point A is located on the straight RS , just outside the forbidden space, at the distance d to either side
- Zero separability is at any point where the equivalent free space distance of \overline{AR} and \overline{AS} are the same
- Points located away from the straight RS have significantly lower contribution to separability

The concept of forbidden space is only useful, if node A is otherwise indistinguishable from nodes R and S . If there is another way of distinguishing node A , for example using node attributes or making node A an anchor, then it can be placed arbitrarily close to R and S without the danger of mismatch. In that case, node A 's contribution to separability is maximal if it is placed very close to either R or S . This result also makes intuitive sense: If node A is placed close to node R and the two are separable by other means, than R and S become separable as well, because node R is the one with A next to it and node S is the other node.

From equation 13 it is a small step to probability of matching error. It is a function of noise variance σ_N^2 and separability $\Delta(g_M)_A$ and can be expressed as:

$$P\{\varepsilon\} = Q\left(\frac{|\Delta(g_M)_A|}{\sqrt{2}\sigma_N}\right) = Q\left(\frac{-10 \log_{10}\left(\frac{(x+d/2)^2+y^2+z^2}{(x-d/2)^2+y^2+z^2}\right)}{\sqrt{2}\sigma_N}\right) \quad (14)$$

where $Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty \exp\left(-\frac{u^2}{2}\right) du$

Figure 11 is a plot of $P\{\varepsilon\}$. Starting at point S a path along the straight RS and a path perpendicular to it are shown. The curves show that nodes aligned with R and S are much more valuable as witnesses for preventing matching errors, than nodes located in perpendicular direction. The minimum error probability for the three node network can also be read from figure 12. At a realistic noise level of $\sigma_N = 5$ dB, and witness node A located in the optimal position, then the probability of wrongly matching R and S is 20%.

This is not a very promising result, but in reality networks have more than three nodes and reliability can be greatly enhanced by using more than one witness node. However, the result

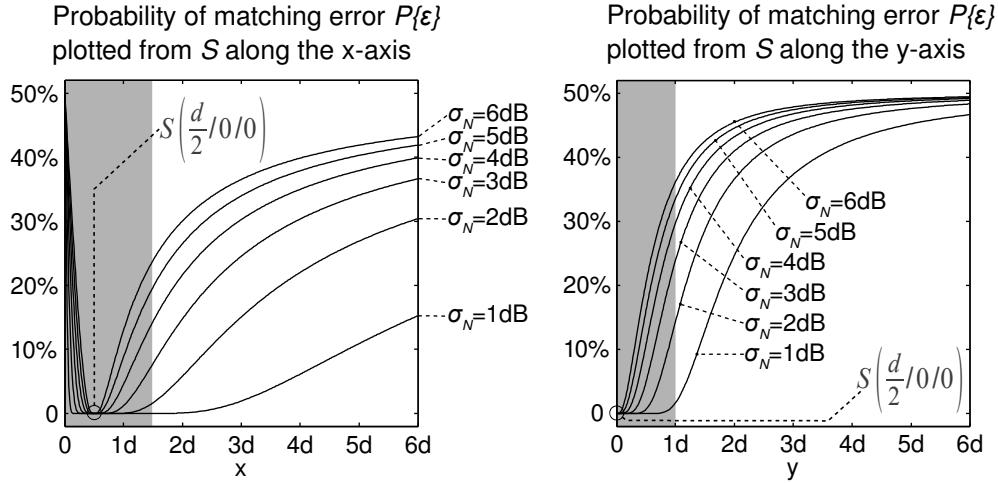


Figure 11: Probability of swapping nodes R ($-\frac{1}{2}d/0/0$) and S ($\frac{1}{2}d/0/0$) as a function of the witness node A . Left – node A aligned with the straight RS . Right – node A located perpendicular to straight RS . The grey area is problematic for placement of A , because it violates the definition of d (cf. equation 11).

is relevant for remote areas of bigger networks. If there happens to be a constellation where two nodes can only be told apart by a single witness node, and if all three nodes are indistinguishable by node-type, then the lower bound for the matching error probability is said 20% at the noise level of $\sigma_N = 5$ dB.

4.2.2 The Impact of the Indoor Environment on the Probability of Matching Error

The impact of the indoor environment is studied using the same three node network as above: Nodes under test R and S , witness node A . How the indoor environment influences this network depends on two factors:

- The indoor environment can have equal impact on the edge-weights $(g_M)_{AR}$ and $(g_M)_{AS}$, or it can impact each differently. Only the latter option changes $\Delta(g_M)_A$.
- While simulation inaccuracies are typically modelled as noise, this subsection takes a closer look on some specific problems and how they deterministically impact Δn_A . Also in this case, only environmental factors that impact the constituent edge weights separately become apparent in Δn_A .

Multi-path propagation is perhaps the most problematic environmental factor in the indoor channel. However, dealing with that is deferred to section 4.3.2, which details the simulation techniques used by ray-optical ray-tracing. This subsection instead looks at environmental impacts of the indoor channel on individual propagation paths.

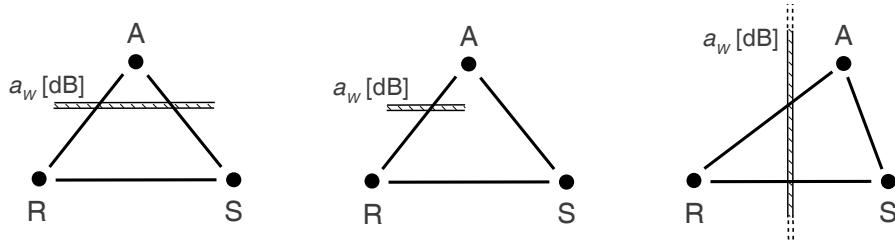


Figure 12: Attenuation due to walls in the three node network. Left – both links between witness node A and the nodes under test R and S are equally impacted by the environment. Middle – only one link is attenuated. Right – The wall bisects the space between R and S. The witness node A can either be on the same side as R or as S. Note that the independent variable of the analysis is the position of the witness node A (x/y/z). The precise geometry of the wall is irrelevant, only its impact on the links AR and AS is considered.

The most pronounced effect is attenuation due to walls and floors. Figure 12 shows the most interesting configurations: (i) The wall impacts both transmission links equally, (ii) the wall impacts one link but not the other, this includes the common case (iii) where the wall is located between the nodes under test R and S .

The impact of the additional attenuation on the three node network is modelled by adding the sources of attenuation: free-space loss and additional losses a_w , generally $a_w < 0$. Modified $(g_M)'_{AR}$ and $(g_M)'_{AS}$ become:

$$(g_M)'_{AR} = (g_M)_{AR} + (a_w)_{AR} = k - 20 \log_{10}(\sqrt{(x+d/2)^2 + y^2 + z^2}) + (a_w)_{AR}$$

$$(g_M)'_{AS} = (g_M)_{AS} + (a_w)_{AS} = k - 20 \log_{10}(\sqrt{(x-d/2)^2 + y^2 + z^2}) + (a_w)_{AS}$$

Fading, reflections, gains and losses due to different phenomena are equally covered by this model, since the sign and the causes of the $(a_w)_{XX}$ are irrelevant. As with the previous section, the independent variable of the analysis is the location of the witness node A (x/y/z). The geometry and location of the wall are not parametrized. The only assumption concerning the wall is that it attenuates radio signals travelling between witness node A and nodes under test R and S according to $(a_w)_{AR}$ and $(a_w)_{AS}$.

The impact on the noise term depends on whether the additional losses are correctly modelled in the simulation. If they are, then the noise term remains unchanged. If however additional losses are omitted in the simulation, then the elements of \mathbf{N} can be modelled as the sum of a deterministic part, and a random variable:

$$n'_{AR} = N_{AR} + (a_w)_{AR} \quad n'_{AS} = N_{AS} + (a_w)_{AS}$$

In the first case in figure 12, the additional losses equally impact transmission links between witness node A and nodes under test R and S , i.e. $(a_w)_{AR} = (a_w)_{AS}$.

$$\Delta(g_M)'_A = (g_M)'_{AR} - (g_M)'_{AS} = -10 \log_{10} \left(\frac{(x+d/2)^2 + y^2 + z^2}{(x-d/2)^2 + y^2 + z^2} \right) = \Delta(g_M)_A$$

$\Delta(g_M)'_A$ is thus the same as in equation 12, which means the presence of additional losses, such as walls, that have equal impact on the transmission links has no consequences on separability.

If the additional losses are equally accounted for the graph G_s , then there is no impact N , no impact on Δn_A , and therefore no impact on the probability of matching error $P\{\varepsilon\}$. The same is true even if the simulation G_s neglects the impact of the wall, because the deterministic part of the noise term cancels:

$$\Delta n'_A = n'_{AR} - n'_{AS} = N_{AR} - N_{AS} = \Delta n_A$$

$\Delta n'_A$ is the same as in equation 12. Additional losses, such as walls, that equally attenuate both transmission links between witness node A and nodes under test R and S therefore have no impact on the probability of matching error of the idealized three node network, even if the additional losses are inadequately represented by the simulation.

The situation is different, if the additional losses affects only one of the two links AR and AS , or if they affects the two links differently, i.e. $(a_w)_{AR} \neq (a_w)_{AS}$. Using $\Delta a_w = (a_w)_{AR} - (a_w)_{AS}$ the separability between R and S becomes:

$$\Delta(g_M)'_A = -10 \log_{10} \left(\frac{(x+d/2)^2 + y^2 + z^2}{(x-d/2)^2 + y^2 + z^2} \right) + \Delta a_w = \Delta(g_M)_A + \Delta a_w$$

While the shape of $\Delta(g_M)'_A$ is essentially the same as $\Delta(g_M)_A$ (cf. figure 10 vs. figure 13), it is offset by Δa_w . The plane of zero separability is no longer the y-z-plane, but instead one of the spheroids around R (negative Δa_w) or S (positive Δa_w). The last point is more important, than it appears: It says that at large distances away from R and S , the contributions to separability of the environmental influences dominate the contributions of free-space losses.

In case the additional losses are well represented in the simulation (original Δn_A), the probability of matching error computes to:

$$P\{\varepsilon\} = Q\left(\frac{|\Delta(g_M)'_A + \Delta a_w|}{\sqrt{2}\sigma_N}\right) = Q\left(\frac{\left|-10\log_{10}\left(\frac{(x+d/2)^2+y^2+z^2}{(x-d/2)^2+y^2+z^2}\right) + \Delta a_w\right|}{\sqrt{2}\sigma_N}\right)$$

Since the surface of zero separability is now a spheroid, not the x-y-plane, there is a second peak in the probability of matching error, cf. figure 13. Compared to figure 12, the probability is higher in the vicinity of the second peak. Aside from the second peak, the probability of matching error is lower because the attenuation of the environment increases separability. If $|\Delta a_w|$ is bigger than 6 dB (the maximum attainable separability due to free-space loss), then spheroid of zero separability withdraws fully into the forbidden zone and error-probability is lowered overall.

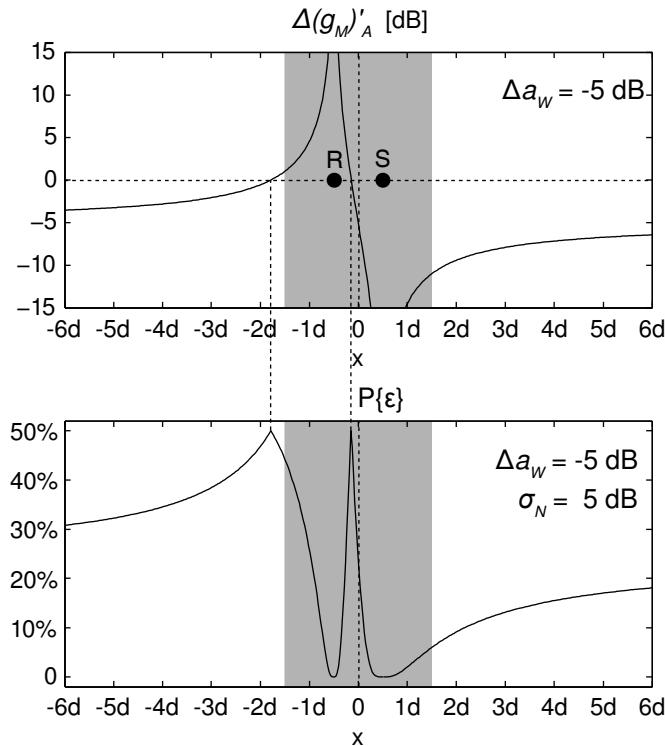


Figure 13: Separability and probability of matching error of the three node network, where link AR experiences environmental losses of -5 dB, e.g. a typical wall blocking line-of-sight. Top graph – separability intersects the abscissa twice; the plane of zero separability is a spheroid around R (cf. figure 10). With growing distance separability converges towards the contribution by environmental factors. Bottom – the left peak in the error probability is the point where attenuation due to free space loss and environment cancel. In the transitional regions separability is low and therefore error probability high. Outside the transitional regions, error probability is lower than the $\sigma_N = 5$ dB curve in figure 12.

A special case is a wall passing between R and S and completely bisecting the space. The witness node A can either share a side with R (Δa_w positive) or with S (Δa_w negative). Both cases make the spheroid of zero separability back off towards the other half-space and for typical parametrizations (wall close to the y-z-plane, $\Delta a_w \cong +/- 5$ dB) will make the areas of zero separability disappear completely. This is what happens at the positive semi-axis of figure 13, but applied to both negative and positive sides. A wall passing between R and S therefore globally lowers the probability of matching error, irrespective of the position of the witness node A . Intuitively, nodes separated by walls are easy to discern with low probability of mismatch.

The last point is important for the viability of the whole approach. The probability of matching error in the three-node network is markedly lower if it spans multiple rooms than if nodes R and S are located within the same room, all other parameters being equal. In short: Walls help.

The situation is different if the wall is not well represented in the simulation \mathbf{G}_S . The wall then impacts $\Delta n'_A$:

$$\Delta n'_A = N_{AR} - N_{AS} + \Delta a_w = \Delta n_A + \Delta a_w$$

Substituting the modified $\Delta(g_M)'_A$ and $\Delta n'_A$ in equation 13, leads to the following modified version:

$$P\{\varepsilon\} = \begin{cases} P\left\{\Delta n'_A < \Delta(g_M)'_A \right\} & \Delta(g_M)'_A < 0 \\ P\left\{\Delta n'_A > \Delta(g_M)'_A \right\} & \Delta(g_M)'_A > 0 \end{cases}$$

$$P\{\varepsilon\} = \begin{cases} P\left\{\Delta n_A < \Delta(g_M)_A \right\} & \Delta(g_M)_A < -\Delta a_w \\ P\left\{\Delta n_A > \Delta(g_M)_A \right\} & \Delta(g_M)_A > -\Delta a_w \end{cases}$$

$$P\{\varepsilon\} = Q\left(\frac{\text{sign}(\Delta(g_M)_A + \Delta a_w) \Delta(g_M)_A}{\sqrt{2}\sigma_N}\right) \quad \text{with} \quad \Delta(g_M)_A = -10 \log_{10} \left(\frac{(x+d/2)^2 + y^2 + z^2}{(x-d/2)^2 + y^2 + z^2} \right)$$

The expression for the probability of matching error (figure 14) is now discontinuous and includes regions with $P\{\varepsilon\} > 0.5$. This is because the presently studied case violates the premise of the least squared error approximation: The premise that the noise is zero mean and the random variables independent. Because the additional losses Δa_w are not represented in \mathbf{G}_S , but instead in the no longer zero mean noise $\Delta n'_A$, the implicit maximum likelihood discriminator (equation 3) becomes suboptimal and in the boundary region becomes more likely to choose the wrong assignment for R and S , than it is to choose the right assignment. The dis-

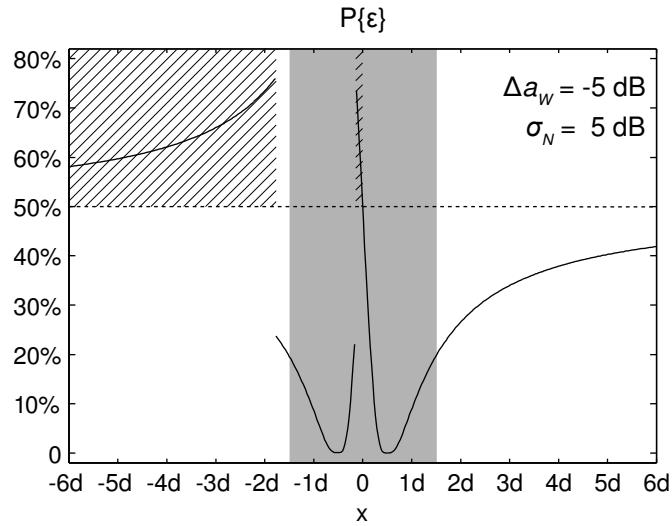


Figure 14: Probability of matching error of the three node network, where link AR experiences environmental losses of -5 dB that are not accounted for in the simulation \mathbf{G}_S . As a consequence, the maximum likelihood discriminator is more likely than not to chose the wrong assignment for R and S (hatched areas, error probability larger than 50%).

continuities occur at the now visible decision boundaries. Decision boundaries also do exist in equation 14, but are located at the plane of zero separability, where they do not cause discontinuities.

The subsection shows that in addition to free space loss, the environment is an additional source of separability. Depending on the geometry and the materials, the two sources can reinforce each other, or can cancel each other. The former case is advantageous and leads to low probability of matching error, the latter case makes matching unreliable. Whereas the deviation between simulation and measurement is generally modelled as noise, this subsection also looks at how low quality simulation impacts matching performance of the three node network: From the perspective of the witness node, if some unaccounted environmental losses make the closer node under test seem further away than the other node under test, then a matching error is a near certainty, except if the noise (some other unaccounted environmental loss) cancels the effect by chance.

4.2.3 Probability of Matching Error in Networks of Four and More Nodes

This subsection continues the discussion of equation 9 (repeated below), albeit for networks of more than three nodes. $P\{\epsilon\}$ is the probability that a particular pair of nodes under test, dubbed R and S , gets mismatched by the least square error optimization. In the case of net-

works of more than three nodes, the sum and the as yet neglected term $\Delta(g_M)_i^2$ have to be taken into consideration:

$$P\{\varepsilon\} = P\left\{ \sum_{i \notin [r, s]} \Delta(g_M)_i^2 \left[1 - \frac{\Delta n_i}{\Delta g_{Mi}} \right] < 0 \right\} \quad (9)$$

(repeated)

$$\text{using: } \Delta(g_M)_i = (g_M)_{ir} - (g_M)_{is}, \quad \Delta n_i = n_{ir} - n_{is}$$

The sum subdivides the problem by dissecting the whole network into a series of three node networks like the ones discussed in previous subsections by iterating over all possible witness nodes A_1, A_2, \dots . The overall probability of matching error is the weighted sum of the results of the three node network tests. The weights are given by $\Delta(g_M)_i^2$, which is the square of $\Delta(g_M)_A$ already discussed in figure 10. Since $\Delta(g_M)_A$ drops fairly quickly towards zero with increasing distance from the nodes under test, it follows that the probability of mismatching a pair of nodes in large networks is dominated by witness nodes geometrically adjacent to the nodes under test. Exceptions to that happen in cases where separability due to environmental losses and separability due to free space loss happen cancel each other in adjacent nodes.

The following paragraphs investigates a number of common network topologies. The first example is the double-sided linear network of $2L$ nodes (not counting nodes R and S , figure 15). This is the first practical network topology under consideration. It exists e.g. in long hallways, where adjacent rooms are not themselves equipped with nodes of the particular network. Reusing the general set-up of section 4.2.1 (equations 10, 11, 12), and witness points located along the x-axis at $A_{1,2,\dots} = \pm 1.5d, \pm 2.5d, \dots$ the probability of matching error amounts to:

$$P\{\varepsilon\} = P\left\{ \sum_{i=-L}^{-1} \left[-10 \log_{10} \frac{i^2}{(i-1)^2} \right]^2 \left[1 - \frac{\Delta n_i}{-10 \log_{10} \frac{i^2}{(i-1)^2}} \right] + \sum_{i=1}^L \left[-10 \log_{10} \frac{(i+1)^2}{i^2} \right]^2 \left[1 - \frac{\Delta n_i}{-10 \log_{10} \frac{(i+1)^2}{i^2}} \right] < 0 \right\}$$

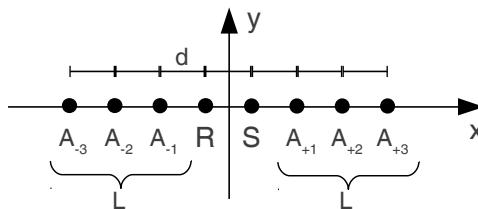


Figure 15: Topology of the double-sided linear network of $2L$ nodes (not counting R and S)

The equation can be rewritten to:

$$P\{\varepsilon\} = P\left\{ \sum_{i=1}^L 2 \left[20 \log_{10} \left(\frac{i+1}{i} \right) \right]^2 + \sum_{i=1}^L \Delta \Delta n_i \left[20 \log_{10} \left(\frac{i+1}{i} \right) \right] < 0 \right\}$$

with $\Delta \Delta n_i = \Delta n_{-i} + \Delta n_{+i}$

The right hand sum is a weighted sum of assumingly independent Gaussian random variables. With the variance of Δn_i being $2\sigma_N^2$ (equation 12), the variance of $\Delta \Delta n_i$ therefore being $4\sigma_N^2$, then the variance of the sum amounts to:

$$\sigma_\Sigma^2 = \sum_{i=1}^L \left[2\sigma_N 20 \log_{10} \left(\frac{i+1}{i} \right) \right]^2$$

The probability of matching error can now be expressed in closed form:

$$P\{\varepsilon\} = Q\left(\frac{20}{\sigma_N} \sqrt{\sum_{i=1}^L \left[\log_{10} \left(\frac{i+1}{i} \right) \right]^2} \right)$$

The first observation concerns the case $L \rightarrow \infty$: The resulting series converges.

$$\sum_{i=1}^{+\infty} \left[\log_{10} \left(\frac{i+1}{i} \right) \right]^2 = 0.184309\dots$$

This is unfortunate, because if instead it had diverged, the probability of matching error would have converged to zero for an infinite number of witness nodes. Instead, the extreme value leads to a lower bound for the probability of matching error in a linear network.

Figure 16 plots the probability of matching error as function of the standard deviation of the noise σ_N and the number $2L$ of nodes in the linear network, including the lower bound $L \rightarrow \infty$. At the lower bound, the error probability of 1% is reached at a noise level of $\sigma_N = 3.7$ dB.

The single-sided linear network of L nodes (not counting R and S) has a similar solution and is shown alongside the double-sided variant in figure 16. Its topology is relevant for the border areas of real networks. The probability of matching error is identical to the double-sided linear network, except for a factor of $1/\sqrt{2}$ of the argument of the Q function:

$$P\{\varepsilon\} = Q\left(\frac{20}{\sqrt{2} \sigma_N} \sqrt{\sum_{i=1}^L \left[\log_{10} \left(\frac{i+1}{i} \right) \right]^2} \right)$$

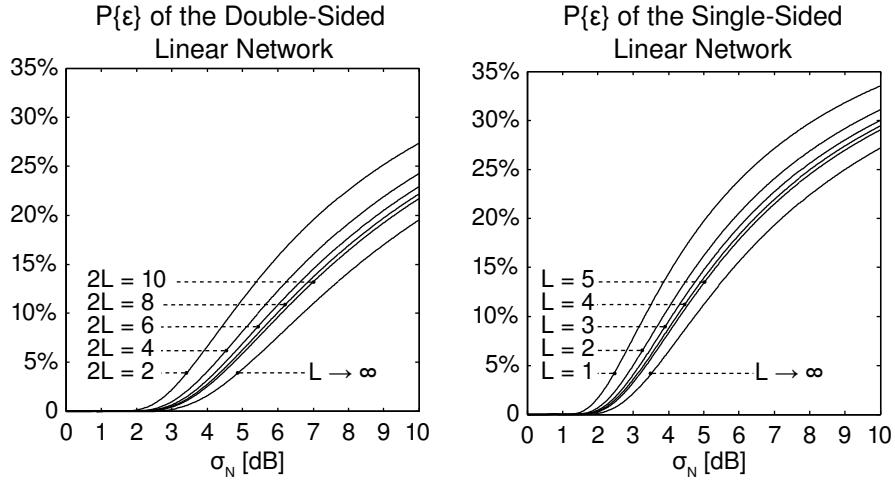


Figure 16: Probability of matching error of a double sided linear network of $2L$ nodes (left) and a single sided linear network of L nodes (right), in both cases not counting nodes under test R and S .

Of higher practical importance than the linear network are topologies of two rows, like the one depicted in figure 17. It can be found in open plan offices or side-wings of larger commercial buildings. A similar topology is also present in the Gusshaus dataset, studied in section 5.2. The interesting aspect of this topology is matching errors between the two rows, not along the main axis. Matching errors between the rows are more challenging, because the witness nodes are in an unfavourable position leading to a low overall separability. Using a similar ansatz as before, the probability of matching error for the one-sided and the two-sided version of the topology compute to:

$$\text{Double sided: } P\{\varepsilon\} = Q \left(\frac{\sqrt{2} \cdot 10}{\sigma_N} \sqrt{\sum_{i=1}^L \left[\log_{10} \left(\frac{i^2+1}{i^2} \right) \right]^2} \right)$$

$$\text{Single sided: } P\{\varepsilon\} = Q \left(\frac{10}{\sigma_N} \sqrt{\sum_{i=1}^L \left[\log_{10} \left(\frac{i^2+1}{i^2} \right) \right]^2} \right)$$

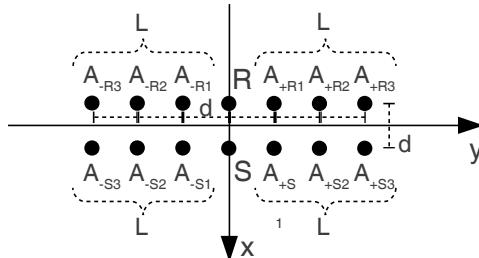


Figure 17: Topology of the two row linear network of $4L$ nodes (not counting R and S)

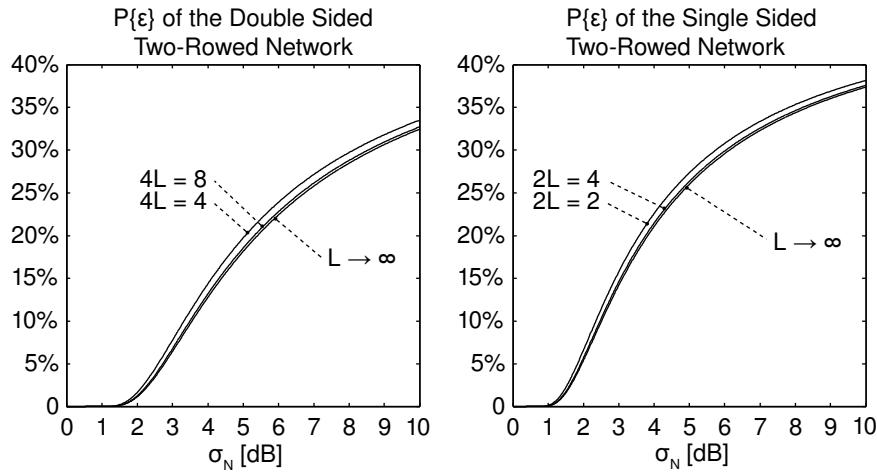


Figure 18: Probability of matching error of a double-sided two-rowed linear network of $4L$ nodes (left), and a single-sided two-row linear network of $2L$ nodes (right), in both cases not counting the nodes under test R and S .

Again, the series converges:

$$\sum_{i=1}^{\infty} \left[\log_{10} \left(\frac{i^2+1}{i^2} \right) \right]^2 = 0.103452$$

Compared to the previous case, the series converges much faster towards its limit, which is a consequence of the unfavourable placement of the witness nodes. Figure 18 is the plot of the probability of matching error. This result is a serious problem for the approach to address assignment: The double-row topology is not uncommon in real buildings, especially side-wings. At the far end of such a formation, the single sided topology is realized (right hand side of the figure). The probability of mismatching the hindmost two nodes at a realistic noise level of 5 dB amounts to 26 %. This is too high for a practical application in address assignment. Consequences and ways to cope with this finding are discussed in section 4.2.4.

The third dimension offers a way to pack more witness nodes in proximity with the nodes under test. This offers moderate improvements: Building on the two-row network, the number of witness nodes is increased by a factor of two by adding additional wings in -z and +z-direction. This yields a factor of $\sqrt{2}$ in the argument of the Q-function:

$$P\{\varepsilon\} = Q \left(\frac{20}{\sigma_N} \sqrt{\sum_{i=1}^L \left[\log_{10} \left(\frac{i^2+1}{i^2} \right) \right]^2} \right)$$

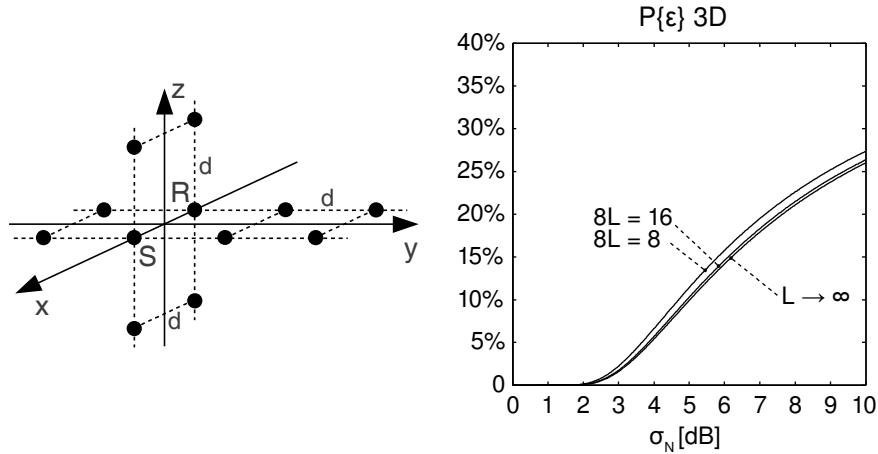


Figure 19: Two-row network extended into the third dimension. One L comprises the 8 adjacent nodes of the pair R, S .

Figure 19 shows the geometry and plot of $P\{\epsilon\}$. The probability of matching error at $\sigma_N = 5$ dB is now 11.4 %. With respect to the nodes under test R and S , the y and z -dimensions are interchangeable.

The y and z -dimensions are typically not interchangeable for real buildings, where floors and ceilings are made of reinforced concrete. Applying the results of subsections 4.2.2, those concrete ceilings have little impact on the separability of nodes within the same floor, because even if there are propagation paths impacted by ceilings, their impact likely cancels in the calculation of $\Delta(g_M)_i$ and Δn_i . The ceilings on the other hand make for excellent separability between nodes located on different floors. Since separability of nodes located on different floors is good, and the additional potential witness nodes improve separability of nodes on the same floor, it is therefore advantageous to match a multi-story buildings in one single step, rather than subdividing the network and match the smaller networks floor by floor. Since matching in a single step may exhaust available computing power, a compromise may be to match the building wing by wing.

The last issue of this subsection is the question whether matching errors lead to subsequent matching errors. Strictly speaking, there are no subsequent errors because there is no sequential process involved. While the matching algorithm may be iterative, the formal construct is not. The permutation matrices \mathbf{P} and $\hat{\mathbf{P}}$ as defined in equations 1 and 3 are the correct permutation and the global optimum of the least squared error optimization. There is no notion of committing to a particular erroneous matching for node R and as a consequence mismatching some other node. Within the formal framework each matching error stands alone, the notion of subsequent errors is nonsensical. The matching algorithm may produce subsequent errors, but that is the scope of section 4.6.

There is, however, the possibility that matching errors accumulate in a particular region of the network. It usually does not happen because of singular outliers in the elements of \mathbf{N} . Singular outliers may be part of many witness node relationships, but are typically dominant in only very few of them. Singular outliers tend to produce one, or at most two matching errors, most likely none at all.

Singular outliers of 15 dB can for example be the result of an uncharacteristically broadband fading effect. Fading is typically frequency selective with a coherency bandwidth smaller than the measurement bandwidth. But since a network of n nodes potentially has $n(n-1)/2$ links, improbable multi-path situations are bound to happen somewhere and produce singular outliers.

Just as probable, however, are problems with the simulation, such as incomplete models of anisotropic material behaviour. Those issues lead to locally increased σ_N and therefore to clustering of matching errors.

4.2.4 Consequences of the Formal Analysis of Network Topology

This subsection revisits the assumptions, presumptions and simplifications behind the analysis in order to evaluate the relevance of the findings with respect to real buildings and real networks. Some of these are indeed presumptions because they are known to be problematic and justifiable only as simplification of the real circumstances. But therein also lies value, because only through simplification it is possible to express error probability as closed form and produce the graphs of the previous subsections.

The most delicate simplifications concern the noise matrix \mathbf{N} and its random variable elements N_{XY} . Worthy of further discussion are (i) the availability and accessibility of statistical moments of the ensemble, as well as the assumption that the N_{XY} are (ii) Gaussian, (iii) zero mean and (iv) statistically independent. These presumptions are prerequisite for the simplified approach that the statistics of N_{XY} can be completely described by the standard deviations σ_N only. In fact, σ_N is in itself a simplification. More correctly, it is necessary to distinguish between ensemble statistics of the random variables N_{XY} and statistics of concrete realizations of \mathbf{N} . Indeed, each random variable N_{XY} is in fact be subject to its specific $(\sigma_N)_{XY}$.

Direct observation of the statistical ensemble of individual N_{XY} is difficult, real observations of individual N_{XY} are essentially limited to a single instance. Gathering additional observations would require installing and uninstalling the network several times with slight variations in position and orientation of the nodes, employed hardware, wall and furniture details. It would also require several simulation models using different ray-tracing engines and slight variations in modelling detail and parametrization.

This is relevant, because σ_N as used in previous subsections formally means $(\sigma_N)_{XY}$, the standard deviation of the statistical ensemble. Since measuring the latter is impractical, error probabilities are instead related to statistics of lower-case n_{XY} , the elements of a concrete realization of N , and its standard deviation $\hat{\sigma}_N$. It is assumed that the latter is an approximation for the ensemble. E.g. the variance of N_{12} , $(\sigma_N)_{12}^2 = E\{(N_{12} - \mu_{12})^2\}$ is approximated by $\hat{\sigma}_N^2$, the variance of the elements of a particular realization of N . This simplification is partly justified, because many contributors to randomness are ignorant of the concrete link. Examples:

- Variance of the ray-tracing simulation (a Monte-Carlo method)
- Hardware variations and installer errors
- Fading – details differ between line-of-sight (LOS) and non-line-of-sight (NLOS) and the rare case where multi-path propagation is suppressed by some unusual constellation; in practise the phenomenon is amenable to statistics with well-understood parametrization

Other influencing factors, however, are expected to differently impact the statistical ensemble versus the statistics of realizations of N : E.g. the parametrization of the simulation: A bad estimate for the attenuation of a particular wall type is capable of biasing the overall matrix N into one or the other direction, thereby exerting limited influence on the extracted statistical moments $\hat{\sigma}_N$, while the same effect is not manifest in the ensemble statistics $(\sigma_N)_{XY}$. The error probabilities of the previous subsections therefore do not adequately cover these classes of mistakes.

A related problem with the standard deviation $(\sigma_N)_{XY}$ of the ensemble N_{XY} is the fact that it is not constant between different statistical ensembles N_{XY} , e.g. $(\sigma_N)_{12} \neq (\sigma_N)_{23}$. Some positions of the network are simply easier to simulate with lower variance than others. Therefore, error probabilities may be overstated for some positions and understated for others. But given the experimental set-up, $\hat{\sigma}_N$ is the best possible estimate of ensemble variances available. Also, the knowledge that it is only an estimate does not invalidate the analysis of the topologies.

The presumption of Gaussianity of the N_{XY} remains unaffected by these considerations. Since the empirical data for N , presented in section 5, fits the Gaussian model reasonably well, it is justified to conjecture Gaussianity also about the statistical ensemble, based on the same argument around the central limit theorem.

The random variables N_{XY} are furthermore believed to be zero mean, because potential biases can be eliminated from the measurement and because ray-optical ray-tracing is an unbiased estimation technique [PHA10]. The mean of empirical N , however, is biased by certain factors, mainly because it is too cumbersome to identify and eliminate all sources. Instead, it is forced to zero by the pre-processing procedures described in section 4.5.

The most problematic presumption is statistical independence between random variables N_{XY} , elements of \mathbf{N} . For any real-world example, the elements of \mathbf{N} are definitely not statistically independent. Any simplification or inaccuracy of the ray-tracing simulation will very likely influence more than a single element of \mathbf{N} , and thereby be a source of correlation, mostly positive correlation. On the other hand, the goal of section 4.2 is to study topologies independently of the ray-tracer. Introducing a parametric model for the statistical dependencies between elements of \mathbf{N} into the calculations would fail to increase their usefulness, least of all because doing so would overwhelm the graphical depictions.

There is an analogy from radio communication: The AWGN channel (additive white Gaussian noise). Indeed the noise is never quite white and the signal path never quite linear. But replacing the AWGN model by something more sophisticated defies the purpose of having a simple, mathematically accessible model that is able to adequately explain a wide range of phenomena.

The omission of statistical dependencies therefore hides certain classes of matching errors from the analysis and therefore slightly underestimates the probability of matching error. On the other hand, the statistical dependencies will cancel in many situations due to the subtraction when computing Δn . Empirical observations suggest that correlation between elements of \mathbf{N} is small compared to other factors, especially after pre-processing.

A further problem with the studied topologies is the question whether those topologies actually exist. After all, real buildings have walls and walls impact separability of nodes, though often favourably. In fact, there are favourable topologies and unfavourable topologies and there are cases where unfavourable topologies are manifest and easy to spot in the floor plan, and there are cases where the unfavourable topologies arise out of subtle combinations of separability due to free space loss and due to walls.

In order to get matching error rates below 5%, a σ_N no worse than 4 dB is needed for the more favourable of the surveyed topologies (figure 16), and a σ_N no worse than 2 dB is needed for the more unfavourable topologies (figure 18). However, none of the surveyed topologies systematically considers the environment and the additional separability it provides. Literature, as well as empirical tests with the ray-tracer, suggest that for a typical indoor environments σ_N is in the order of 5 dB. This means that based on this analysis, and subject to the reservations above, automatic address assignment is not able to achieve the required minimum accuracy to achieve cost savings.

A positive outcome of the analysis is that unfavourable topologies can be automatically identified by directly applying equation 9 to the simulation model \mathbf{G}_S . Section 5.2.1 has the appropriate analysis for the data gathered at the Vienna University of Technology (VUT). Once the

hot spots for the error probability are known, the operator can be instructed to employ one of several counter measures:

- One or more nodes of the problematic constellation can be uncovered prior to graph matching by manually scanning the low-level address and entering it into the commissioning tool. The node henceforth serves as anchor.
- A helper node can be placed into the constellation at a suitable spot in order to enhance the separability. Figure 10 may serve as a guidance on exactly how and where helper nodes can be placed most efficiently.
- Extra care can be taken during channel modelling and simulation to ensure above quality simulation in the affected area of the network. Since the overall error-rate is dominated by those hotspots, it is justified to focus the modelling effort on geometry, materials and antenna configurations in the vicinity of the affected nodes. The ray-tracing simulation can be run with a higher number of traced rays for the affected links. If the noise standard deviation of the affected links can be pushed down to 2 dB, then the remaining probability of matching error is acceptable.
- Attributes, such as node types, can be used to dramatically improve separability. The same device could, for example, be manufactured in several versions of equal functionality, but easily differentiable for installers and queryable by node-type. Installers would then be directed to install e.g. devices with the blue marking next to windows and devices with the red marking on the opposite side.

Of course, these procedures are all manual steps, much like those the whole approach to address assignment is trying to prevent. This is a weak point of the approach, and it is shown using formal methods only. It can be worked around, but it cannot directly be fixed. A potential game-changer to counter the problem of difficult topologies would be a substantial improvement in deterministic channel modelling.

4.3 Model of the Indoor Channel

This section discusses the first step (channel estimation) of the three-step process outlined in section 1.1.

Since a detailed building and deployment plan are available to the channel simulation, and since section 4.2 identifies topologies that require high accuracy, channel simulation based on ray-tracing seems adequate. Section 4.3.2 discusses the advantages and disadvantages of competing approaches to channel simulation. For the prototype implementation, AWE Soft-

ware's Winprop channel modelling software is used. Unless otherwise noted, the 2.4 – 2.48 GHz ISM-band is primarily considered.

The channel is assumed to be static, i.e. static enough that time variations need not be considered and can not be taken advantage of. At commissioning-time of the building automation network, the building is assumed uninhabited and network components are mounted in their respective fixtures. Mobile nodes are not considered.

The two main challenges for static indoor channels are line-of-sight (LOS) vs. non-line-of-sight (NLOS), as well as fading due to multi-path propagation. Doppler fading is no concern for static channels. Polarization is dependent on antenna choices and wherever known, can in theory be adequately represented in the simulation. In practise, the radiation patterns depend not only on the antenna, but also on the fixture and geometry at the mounting location, cf. section 4.3.3. The channel simulation is therefore done with isotropic antenna models. The resulting simulation error contributes to the noise.

In the NLOS channel, good estimates for signal attenuation and reflection properties of the different materials are crucial. Signal attenuation varies between 1 dB for some types of dry walls, up to 60 dB for the thick walls of historical buildings (cf. section 4.3.4). Metal plating acts as near-perfect shielding up to the sensitivity limit of measurement equipment, although in practice gaps in the plating allows some of the signal through. Attenuation due to reflection on surfaces are in the order of 10 dB for walls and furniture and in the order of 0.1 dB for metal plated surfaces. Because the matching algorithm can tolerate some noise, empirical estimates based on approximate wall thickness and material class are sufficient (at least for favourable topologies). It is, however, essential to model large metal surfaces, such as radiators, metal suspended ceilings, metal doors, and metallized surface coating on windows and glass doors.

Easily overlooked metal surfaces are air- and cable conduits above suspended ceilings. In some cases, modelling of the latter surfaces is not necessary, because their overall contribution to the field configuration is limited enough – as long as they do not obstruct crucial propagation paths⁴. The Winprop software comes with demonstration projects that cover an adequate range of building materials. Additional own material measurements have been conducted and are documented in section 4.3.4.

⁴ The point becomes apparent if the radar equation is considered. A reflection on relatively small metal surfaces does not change the overall distribution of RF energy within the indoor channel.

4.3.1 Fading

Fading due to multi-path propagation needs to be considered in any meaningful discussion of indoor channels. It is a highly localized (centimetres) phenomenon, geometry and frequency dependent and causes local variations of the signal strength in the order of 30 dB [VAL97]. Figure 20 shows an example recorded with a network analyser between two monopole antennas, exhibiting multipath propagation. Fading is commonly modelled statistically, in particular by the Rayleigh distribution in the NLOS case and the Rician distribution in the LOS case. This is not sensible for the problem at hand, because there is too little variability in the static channel with respect to time and space, due to the fixed mounting locations. The statistical-inference approach is incompatible with deterministic channel simulation.

Still, the effects of fading are severe enough that they need to be accounted for. This can either be done in the channel simulation, or in the channel measurements. The Winprop software can do coherent simulations, although in the 2.4 GHz band, where the free-space wave length (λ_0) is merely 12 cm, the required precision with regards to entering dimensions into the simulation and the required precision with regards to placement of node-antennas makes the approach impractical.

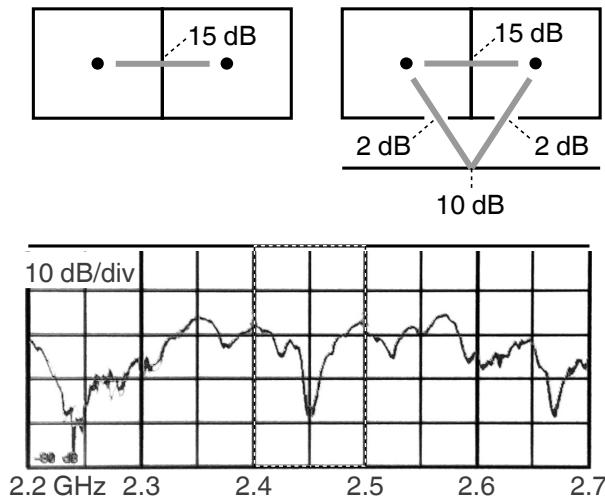


Figure 20: Top left – single dominant propagation path between two nodes in adjacent rooms, attenuated by 15 dB (transmission loss). Top right – multi-path propagation, multiple paths with similar transmission / reflection loss. This provokes interference of the received signal with itself. Bottom – received signal exhibiting frequency selective positive and negative phase interference (fading), recorded with a network analyser and two monopole antennas. At 2.45 GHz the received signal strength drops by approximately 20 dB. The curve bears a certain periodicity which is typical for this type of fading. The period is characteristic for the length difference between the different propagation paths.

The alternative approach is to measure the channel while avoiding fading effects. This is most easily done by measuring the channel at different frequencies. Wireless nodes compliant with IEEE 802.15.4 support 16 channels in the 2.4 GHz band, 5 MHz apart. If the coherence bandwidth (B_c) of the static indoor channel is smaller than the frequency range covered by IEEE 802.15.4, then the maximum observed RSSI over all channels is a good estimate for a non-fading channel. Jannsen observes this in 85 % of cases for the NLOS channel, where fading is a regular concern, and 55 % of cases for the LOS channel, where fading is a minor phenomenon [JAN96]. Even if this strategy cannot fully compensate for fading in all cases, the maximum of several measurements is a good estimate for the fading-free channel, because fading only causes negative spikes (selective cancellation), never positive spikes. Since there are no positive spikes, any value that is not a negative spike is already a sensible approximation. Of course there is constructive interference, but outside of parabolic reflectors it is hardly more than 3 dB. There is insufficient data to estimate the residual variance due to fading, but an educated guess would be: $\sigma_F < 3$ dB.

The described approach to fading compensation is not possible in sub-GHz radio, because the available frequency bands are too narrow to observe uncorrelated fading behaviour. Whether the larger wave-length makes coherent channel simulation practical, is subject of future research.

4.3.2 Ray Tracing

Ray-tracing and ray-casting are Monte-Carlo methods for predicting radiance distribution in cluttered environments [PHA10]. The ensemble of rays represent an unbiased sample approximating the field configuration. Clever sampling strategies, high quality geometry data and realistic material parameters lead to high-quality channel models. This includes effects such as wave-guiding along corridors and diffraction around edges. Landstorfer reports standard deviations between measurements and simulations σ_N between 4.2 dB and 11.9 dB [LAN99]. Valenzuela reports standard deviations σ_N better than 3.2 dB for LOS and better than 8.8 dB for NLOS [VAL98]. In both cases no specifics concerning the parametrization are given.

These numbers from literature are more than 10 years old at the time of writing, and advances in computational power, together with careful modelling, lead to standard deviations at the low end of the above range found during experiments (section 5). The simulation accuracy is therefore more than adequate for dealing with favourable topologies (cf. figure 16), but is produces matching errors with unfavourable topologies (cf. figure 18).

There exist more accurate simulation methods than ray-tracing, such as hybrid and full-wave approaches [ALI07] [CHAN06]. The problem is that increased simulation accuracy does not automatically translate to lower noise variance in relation to channel measurements, simply

because there are limits on how precisely it is possible to plan and execute the construction of a building and automation network, and how accurately the building materials can be modelled. Clearly, a typical plan prescribes quantity and type of network components and their location. The precise placement, including antenna alignment to within a fraction of lambda, which is relevant for a full-wave simulation, is often left to the discretion of the installer. The second problem concerns availability of integrated software packages that support full-wave simulation of indoor channels. Development of a custom packet is out of scope of this work. Therefore, a ray-tracing is favoured.

According to Matt Pharr's and Greg Humphreys' "*Physically Based Rendering*," the general approach is to simulate the distribution of electromagnetic energy emitted from various sources over the surfaces of the 3D model. At the core, ray-tracing is recursive Monte-Carlo integration. The incoming radiance at a particular point is given by the integral of incoming rays over the solid angle. The incoming rays are determined by the outgoing rays of nearby visible surfaces and the material of the pervaded space. The outgoing rays at the visible surfaces are computed by recursively calculating the incoming rays at the respective surface element and passing them through a surface material function, which again involves integration over the solid angle. Surface material can e.g. be modelled by bidirectional scattering and surface reflectance distribution functions. Integration can be done coherently (observing phase and polarization), or power-only, resulting in a model that observes or disregards fading [PHA10].

The first challenge behind ray-tracing lies in the sampling strategy: It is crucial to efficiently sample the integrand, thereby choosing to follow rays that are likely to carry relevant quantities of energy, not missing dominant propagation paths and all that without biasing the integral. Unbiased sampling and integration algorithms converge on the true value of the integral as the number of samples tends toward infinity. A ray-tracer using an unbiased sampler and integrator converges to the true radiance distribution if the number of followed rays tends towards infinity. The same ray-tracer using a finite amount of computation time will return results statistically centered around the correct result, but subject to a certain variance. Good ray-tracers minimize the variance with respect to the computation time. Good sampling strategies take the shape of the surface material function and the geometry into account.

The second challenge is to efficiently trace rays, which mainly comes down to collision testing between straight lines and polygons (or triangles). The most important methods involve acceleration data structures and space-time-trade-offs. Winprop's "intelligent ray-tracing" is all about pre-computing visibility relationships (geometric coupling) between surfaces and exploiting geometrical similarities between adjacent surface tiles [WOL99].

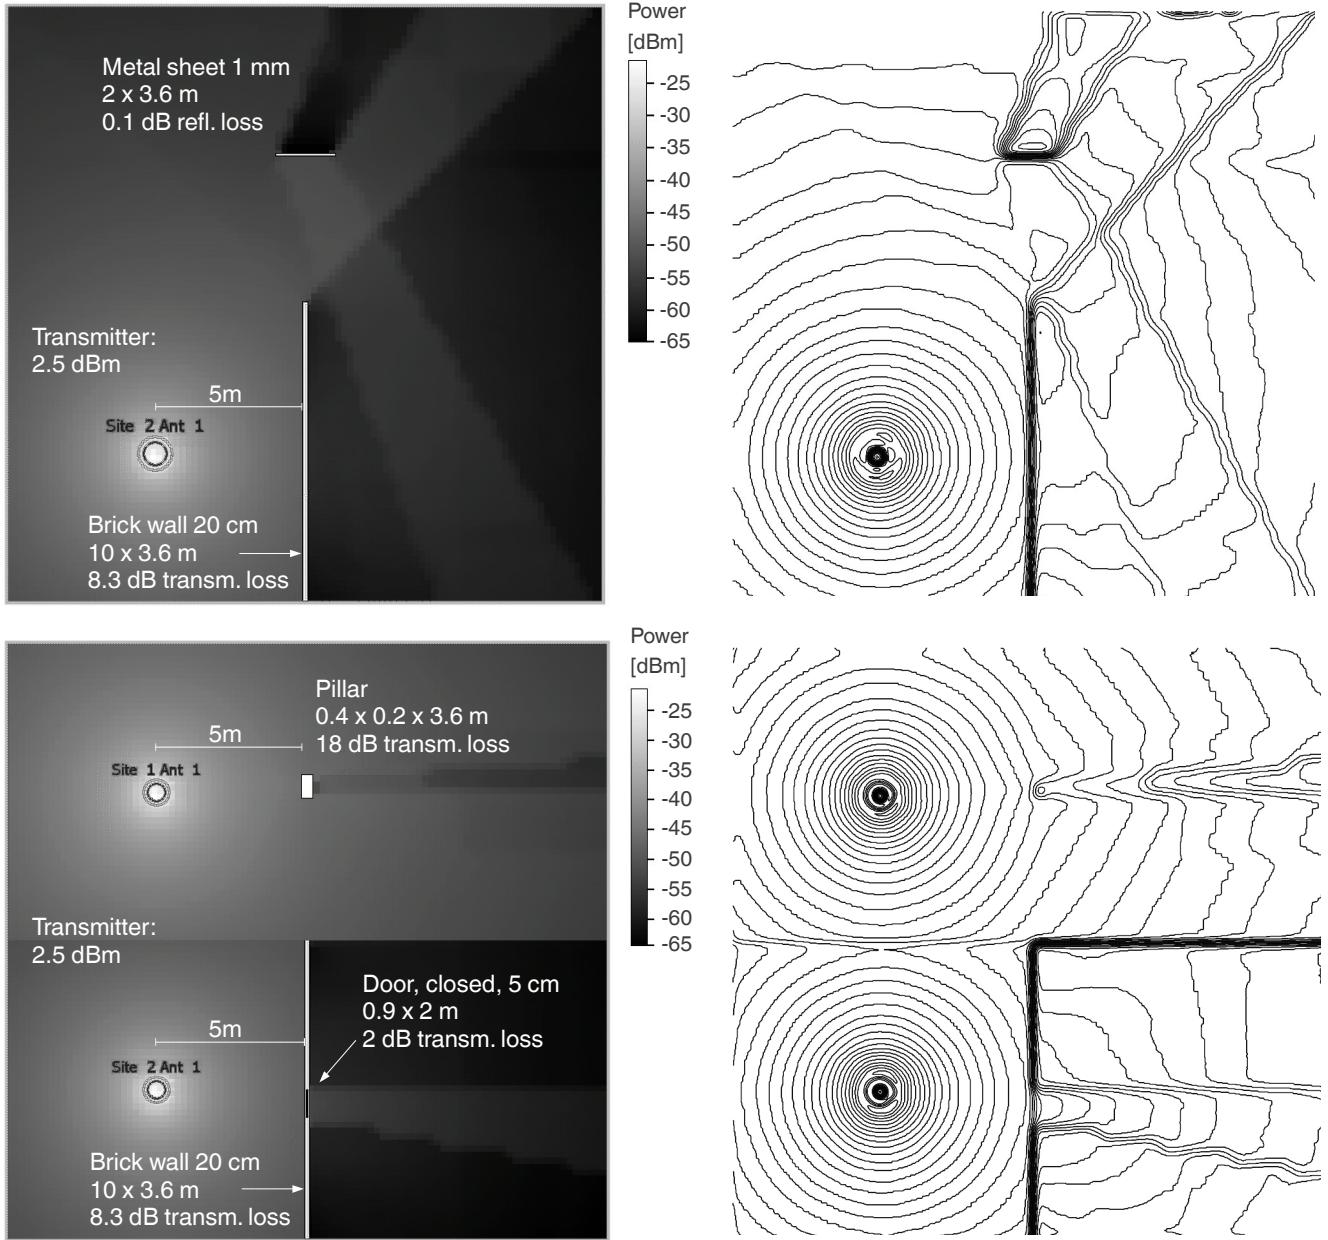


Figure 21: Demonstration of Winprop facilities. Left – annotated simulation output. Right – contour diagram, contour lines are 1 dB apart.

Top – transmission through brick wall, shadowing and specular reflection on metal surface. Note how the reflection cone does have better signal strength than the surroundings (by about 2 dB), but the reflecting surface is substantive. The impact of smaller reflective surfaces is therefore negligible.

Bottom – diffraction around a reinforced concrete pillar, and around a door and door post. Note how signal strength closely behind the pillar is hardly impacted. Note how contour lines behind the door are impacted even outside the signal cone.

Winprop settings: Isotropic antenna, empirical material models. Walls are 3 m high, concrete floor, no ceiling. Simulation height 1.5 m.

The recursion can be started at points of reception (classical ray-tracing), or at the energy sources (ray-casting or ray-launching). Practical ray-tracing uses a hybrid approach. Pursuing only ray-launching will result in many calculated rays that never reach a point of reception. The opposite will result in many calculated rays that never hit an energy source. The modern approach is to calculate a preliminary radiance distribution based on ray-launching, and then find paths between points of reception and radiated surfaces using ray-tracing. In computer graphics this approach is also called pre-computed light transport.

Winprop features a physical material model and an empirical material model. The former relies on physical material parameters like conductivity, dielectric permittivity and magnetic permeability. These are relatively complicated to measure or extract from building plans. More accessible is the empirical model which relies on transmission loss and reflection loss as material parametrization, which is comparatively easier to measure and the model therefore preferably used.

The phenomena modelled by ray-tracing are depicted in figure 21:

- Reflection – especially specular reflection; Diffuse reflections (scattering) have a comparatively small impact in the microwave bands as long as the materials are approximately homogeneous
- Transmission – optionally considering refraction; Both reflection and transmission may be modelled based on Fresnel equations and a physical material model or based on empirical material models
- Polarization – amendable to simulation, but in practise difficult to implement because it relies on the physical material model and elaborate antenna model

Ray-tracing is a ray-optical approximation of wave propagation phenomena and traditionally has difficulties in simulating diffraction. This is because a ray is equivalent to a locally planar wave front, a representation that is invalid in the proximity of conducting surfaces or objects coated with dielectric or magnetic materials. Different approximations for diffraction exist, the ones receiving the most attention being the geometrical theory of diffraction and its generalizations, as well as the uniform theory of diffraction [RAH12] [BOU93].

The disadvantage of these approximation is the reliance on physically exact material models. Winprop therefore alternatively supports a simpler empirical model of diffraction explained by Wölfle et al. [WOL00].

Assembling the simulated connectivity graph G_s from the simulation output is straight forward.

4.3.3 Antenna Radiation Patterns

Antennas are in theory easy to simulate with ray-tracing. In practise, acquiring accurate radiation models proves problematic. This is not due to the antenna per se, but inherent in node construction and mounting. The problems are:

- Missing ground planes – Monopole and whip antennas require antenna ground planes at their base to exhibit the typical dipole-like radiation pattern. Without it, the field lines hit the antenna feed and RF ground below the transceiver circuit.
- Given the overall design of low-power wireless nodes which is optimized for low costs, and only secondarily for RF performance, compromises are made with respect to the antenna, feeding and grounding. RF energy is radiated not only by the antenna, but also by support structures.

As a consequence, the effective radiation pattern is a function not only of the antenna, but of the whole assembly. And even though it is well known that conducting materials in the near field detune the antenna and further decrease RF performance, simplicity with respect to mounting and versatility outweigh. The effective radiation pattern therefore depends on the node and on the mounting location.

The demonstration node (cf. appendix B), as used in the measurement campaign at VUT, is based on the NXP/Jennic wireless ZigBee module JN5148-001-M00. It uses a printed circuit antenna based on a folded-”L” design (figure 22) [NIU00] [LEE00]. The site features a metal suspended ceiling. The measurement campaign includes positions next to light switches and positions next to luminaires, just below the metal suspended ceilings. Figure 23 compares the radiation pattern of the demonstration node mounted on a metal surface and without metal surface.

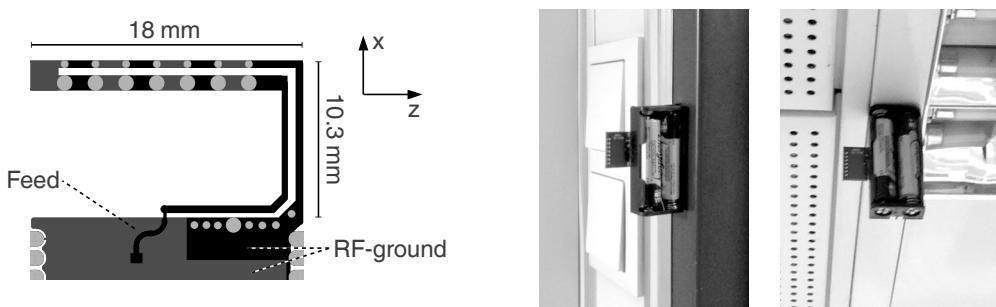


Figure 22: Left – NXP JN5148-001-M00 printed circuit antenna, bent folded “L” design. The pole meanders between top and bottom copper layer. Right – demonstration node mounted next to light switch and next to luminaire, below the metal suspended ceiling.

It is impractical to individually measure the radiation pattern of installed nodes as part of the address assignment procedure. It may be possible to estimate the radiation pattern based on simulating the characteristics of the node and precise geometry of the mounting fixture. But as explained in the previous subsection, reliance on overly accurate geometry data is also impractical.

In fact, figure 23 does not tell the whole story. The radiation patterns were recorded using non-professional equipment and represent only horizontal polarity. The radiation patterns therefore hide much of the actual complexity. Complete far-field radiation patterns are complex-valued functions of four dimensions: azimuth, elevation, frequency and polarization.

- Azimuth, elevation: The radial resolution of figure 23 is 15° , which is too coarse to show the real lobe structure. There may be narrow, deep fades between lobes. An accurate simulation would require very precise data not only on the position, but also on the angular alignment of the nodes and fixtures. A misalignment of 5° can mean the difference between hitting a lobe or the deep fade in between.

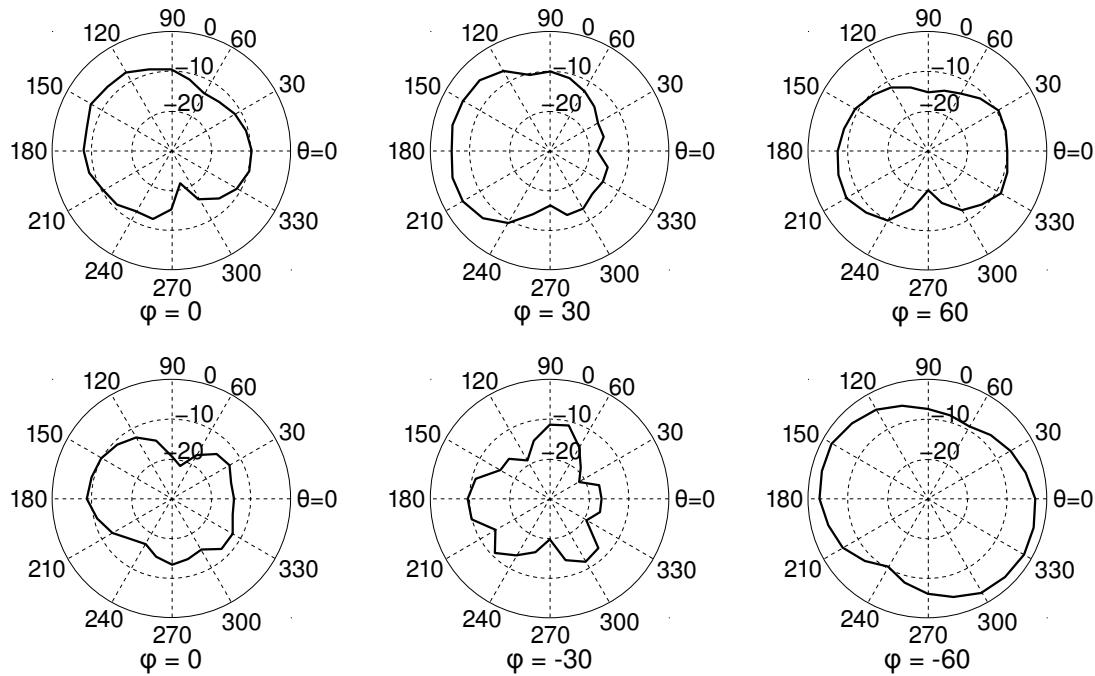


Figure 23: Radiation pattern of the demonstration node, 2.4 - 2.5 GHz, horizontal polarization, in dB, normalized to maximum gain. Top row – no conductive materials in the vicinity. The left pattern shows the XY-plane, the middle and right pattern radiate upwards at increasing elevations. Bottom row – the node is attached to a metal suspended ceiling. The left pattern shows the XY-plane, the middle and right pattern radiate downwards. At $\varphi = -30^\circ$ the incoming wave interferes destructively with the wave reflected off the ceiling. At $\varphi = -60^\circ$ the interference is constructive.

- Frequency: Every plotted value is in fact the maximum of 40 samples taken between 2.4 and 2.5 GHz. The radiation patterns at any particular single frequency look notably different. This mirrors the situation when collecting the measured connectivity graph, which happens over several channels of the 2.4 GHz ISM band because of fading-avoidance. In order to capture the frequency dependency of the radiation patterns, the ray-tracing simulation would have to be performed separately at every measured frequency.
- Polarization: As explained in section 4.3.2, simulation of polarization effects is avoided because it depends on a physical (rather than empirical) material model.

Other than precise measurements, the radiation patterns could in theory also be computed. But in order to produce quality far-field radiation patterns from near-field geometry, the situation around mounting fixtures would need to be specified with very high accuracy, in the order of $\lambda/10$, or 1.2 cm in the 2.4 GHz band. Additionally, highly accurate physical material models would be required. Unfortunately, if automatic address assignment is to fit into established procedures, it cannot require mounting and alignment precision in excess of what a typical installer would be able to produce. Without precise geometry and material models within the near field, any simulation of the effective radiation pattern of the node in its mounting location will produce poor results.

Experiments with various parametric radiation models based on the data of figure 23 failed to meaningfully lower variance between the measured and simulated graph. Since individual in-situ measurement of the effective radiation pattern, as well as its estimation, are impractical or ineffective, the consequence is to run the ray-tracer with ideal isotropic antennas. The resulting simulation error contributes to the noise variance.

Based on the raw data behind figure 23, the standard deviation between the isotropic model and the measured patterns is 4.9 dB for the radiation pattern without conductive materials and 5.9 dB for the radiation pattern below the metal suspended ceiling. The effective impact on simulated and measured connectivity graph is more nuanced, because of averaging over (i) frequency, (ii) polarization and (iii) sending and receiving node. If the vertical polarization of the radiation patterns (not measured) have a similar shape than the horizontal polarization, then the overall impact on the noise can be expected to be $\sigma_{N,RP} \approx 2.4$ dB. The estimate is based on averaging random points in the measured radiation patterns in the same way as the graph assembly procedure described in section 4.4 would.

4.3.4 Walls and Materials

As part of the series of experiments at VUT, in order to verify and complement the data supplied by Winprop, the radio characteristics of a number of different building materials is measured. Table 1 lists attenuation due to specular reflection and transmission of various building materials found at the site in the 2.4 GHz ISM band. Additionally, for the 11 cm interior dry wall and for the 9 cm brick wall, bi-directional transmission distribution and bi-directional scattering and surface reflectance distribution are provided (figure 24).

Table 1: Measured radio characteristics of selected building materials

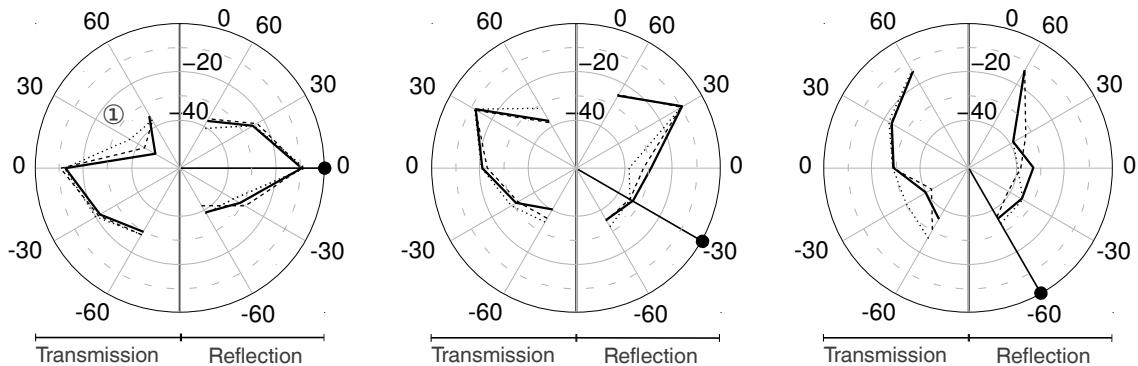
	Reflection	Transmission
Interior dry wall (10 cm)	-9 – -15 dB	-0.5 – -3.4 dB
Brick wall (9 cm)	-10 dB	-12 dB
Floor (reinforced concrete, 40 cm, screed, linoleum)	-9 dB	-20 dB
Metal suspended ceiling (transmission along joints)	0 dB	~ -36 dB
Interior blinds (metal)	0 dB (horiz. pol.) -2.5 – -4.7 dB (vert.)	
Glass window	-5.4 dB	
Glass door	-5 dB	-5 dB
Wooden door	-9 dB	-2.5 dB

The interior dry wall is made of rock wool between gypsum plasterboard, the latter reinforced with aluminium spacers. The panels (105 cm wide) are joined by aluminium braces. Because of the inner structure, reflection and transmission characteristics are given by their observed minimum and maximum. Between braces, the interior dry wall exhibits a very small transmission attenuation (-0.5 dB), with the antennas pointed at the braces, -3.4 dB are measured. Since it is impractical to model all joints between dry wall segments, mean values are used in the simulation. The resulting deviation between measurement and simulation contributes to noise.

The measurement set-up consists of two TP-Link TL-ANT2424B high-gain, directional, parabolic antennas (24 dBi) mounted on tripods and a Rhode&Schwarz ZVRE vector network analyser. The side lobe suppression of 20 dB ensures adequate suppression of secondary propagation paths. The measurement distance is 4 m, roughly 2 m at each side of the wall. A larger distance would be preferable, but especially the measurement of angular reflectivity and transmissibility require frontal and lateral clearance to both sides of a segment of wall not

Bidirectional transmission distribution, bidirectional scattering and surface reflectance distribution [dB]:

– Brick wall, plastered, 9 cm:



– Interior dry wall, rock wool between gypsum plasterboard, aluminium spacers & joints, 10 cm:

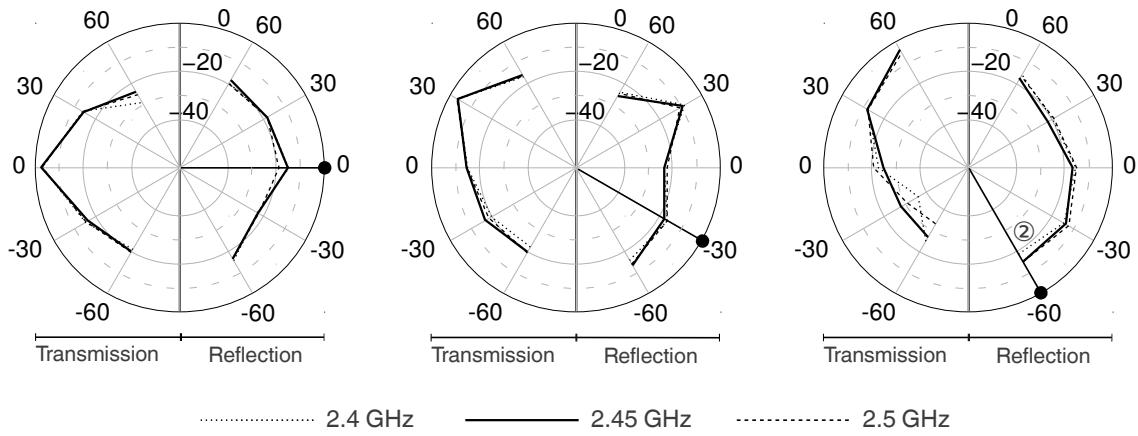


Figure 24: To read the figure: Left to right – the transmitting antenna, represented by •, is placed at different azimuth (0° , -30° , -60°) from the wall. The curves show attenuation of reflected and transmitted signal in dB, plotted as function of the azimuth of the receiving antenna. Transmitting and receiving antenna are pointed at a common point half-way inside the wall, at a distance of 2 m.

Reflection is dominated by specular reflection. Scattering is a minor phenomenon in case of the brick wall. The dry wall exhibits more scattering, most likely due to its inner structure containing aluminium joints and spacers, c.f. point ②. The angular resolution (30°) is too low to identify or rule out any possible anisotropic effects of the dry wall. Given above data, Winprop's empirical scattering model seems adequate.

Diverging of the three curves (at 2.4 GHz, 2.45 GHz and 2.5GHz) is evidence for fading, i.e. destructive interference involving an alternative propagation path. In the case of the anomaly at point ①, for example, the alternative path involves a nearby door and a side lobe of the sending antenna.

otherwise compromised by doors, furniture or electrical wiring. At 2 m distance the diameter of the first Fresnel zone at the wall under test is 0.7 m, albeit appropriately bigger at flat incidence and exit angles. Unless otherwise noted, horizontal polarization is measured.

4.4 Assembling the Measured Connectivity Graph

This section discusses the second step (channel measurement) of the three-step process outlined in section 1.1.

While the simulated connectivity graph can be created ahead of time, assembling the measured connectivity graph is an on-site activity. The wireless network is assumed non-commisioned, but the nodes are powered and an initial, maybe temporary, network has been formed by a procedure such as ZigBee network forming. Nodes are reachable via their low-level network or MAC-address. The wireless channel can be measured as a whole, or the problem can be subdivided rudimentarily, for example by only powering nodes on a particular floor or section of the building.

A procedure and protocol has been devised to gather the measured connectivity graph, c.f. appendix A. It is formulated in terms of ZigBee, but it is not directly dependent on particular ZigBee features. Support for many-to-one routing, however, is desirable for fast and efficient data collection. The procedure works as follows:

1. The network is put into graph assembly mode, temporarily suspending routing and other network-layer mechanisms.
2. Participating nodes visit a set of channels.
3. In each channel, participating nodes listen for incoming graph-assembly beacons. Receivers are in promiscuous mode, i.e. MAC-filtering is suspended, power-cycling as well. Received beacons are stored in a graph-assembly table (table 2), also recording the beacon's RSSI.
4. Every node also sends a small number of graph-assembly beacons at random points of time in each channel, containing mainly its low-level address and device type. Stay-time in each channel is in the order of 10 seconds – such that beacon collisions are sufficiently unprobable.
5. After the last visited channel, normal network operation is restored and the graph-assembly tables downloaded to the data sink.

Table 2: Structure of the Graph Assembly Table

<i>low-level address</i>	<i>node type descriptor</i>	<i>TX-power</i>	<i>RSSI channel 1</i>	...	<i>RSSI channel n</i>
:	:	:	:	:	:

Based on the downloaded graph tables (cf. table 2), assembly of the measured connectivity graph \mathbf{G}_M is straight forward:

- Each vertex is represented by one table; vertices are attributed by low level address and node type descriptor
- Each edge is represented by a pair of rows in the tables of the respective vertices. Edge weights are given by:
- $(g_M)_{AB} = (g_M)_{BA} = \max(\max(RSSI_{A \rightarrow B}) - TXpwr_A, \max(RSSI_{B \rightarrow A}) - TXpwr_B)$

$RSSI$ (in dBm) is the received signal strength indicator as reported by the transceiver when receiving the beacons. $TXpwr$ (in dBm) is the nominal transmit power of the node sending the beacon. Edge weights $(g_M)_{AB}$ (in dB) are estimates for the interference-corrected channel attenuation between nodes A and B. Antenna, matching and insertion losses at transmitter and receiver need not be accounted for, but are dealt with during pre-processing. A small number of missing graph assembly tables can be tolerated, because all relevant information is redundantly available in the assembly tables of neighbouring nodes.

An alternative method of calculating the edge weight is given by $(g'_M)_{AB} = \text{avg}(\dots)$. A large deviation between $(g_M)_{AB}$ and $(g'_M)_{AB}$ indicates the presence of an asymmetrical link. This can happen for a variety of causes, but if the majority of links of a particular node are asymmetric, it is a strong indication for a defect in the node's transmission or reception path. It does not matter for the purpose of address assignment, because the *max* function masks the problem. Problems concerning the antenna are usually symmetrical and cannot be detected by this mechanism.

Appendix B has details about the ZigBee network, the nodes and related procedures used during the experiments.

4.5 Pre-Processing the Graphs / Noise Cancellation

This section discusses the third step (matching algorithm) of the three-step process outlined in section 1.1. Specifically, the mechanics of pre-processing measured and simulated connectivity graphs are considered. Pre-processing involves the following steps:

- Optionally thinning out edges with very small weights (e.g. -95 dB and below)
- Estimation of the global offset
- Normalization, dynamic range compression
- Cancellation of vertex noise

There is a strong correlation between distance and path loss (cf. figure 6). While in the indoor NLOS environment it is not directly useful without the ray-tracer, the fact remains and it leaves its imprint on the data. One important consequence of this correlation is the fact that in the typical graph-assembly table the majority of neighbours have a low or very low RSSI. This is because the volume of a sphere increases with the cube of the radius, i.e. in a building with evenly distributed nodes more neighbours are near the limit of radio coverage, than there are neighbours which are reliable communication partners. This also means a disproportionately high number of edge weights in the measured connectivity graph are near the noise floor or the limit of receiver sensitivity – the weight distribution is heavy at the low end. Figure 28 in section 5.1.1 illustrates the point by means of Mannheim test data set.

The same applies in a slightly different way to the simulated connectivity graph: In the simulation, receivers are not limited with respect to sensitivity or noise. The low end of the weight distribution in measured and simulated connectivity graph therefore tend to look markedly dissimilar. There are more issues complicating things:

- Graph-assembly tables may be limited in size which results in edges missing in the measured but not the simulated connectivity graph, also predominately at the low end
- Edge noise also tends to be larger at the low end, because inaccuracies with respect to material parameters and geometry tend to accumulate in transmission paths that are subject to several transmissions and reflections

There is a trade-off with respect to choosing which edges to keep and which to discard. Discarding many edges at the low end reduces noise, thereby increases the accuracy of the match, and simultaneously lowers computation time of graph matching. The problem lies with choosing corresponding edges to discard in both graphs, especially considering that the graphs are permuted and the data is noisy. Eagerly discarding many edges, means potentially discarding non-corresponding edges in measured and simulated connectivity graph, and thereby amplifying a small amount of noise to a much bigger problem by having many un-matched edges.

The approach taken by the reference implementation is therefore to keep all edges. An alternative, but ultimately abandoned approach is to try to minimize the number of un-matched edges by forcing the number of edges in both graphs to be equal. This can be achieved by discarding edges from the bottom of the simulated weight distribution until the number becomes

equal. The strategy fails whenever there are outliers with respect to vertex noise. Vertices with high positive vertex noise end up almost without any edges which reduces the quality of the match.

Choosing the global offset is non-obvious: Aligning the weight distributions by its means is sub-optimal for later graph-matching, because the distributions are heavy at their low end which is also noisy. The median is better. Still better results are obtained by aligning at the 9th 10-quantile, i.e. choosing off such that the weight distributions of \mathbf{G}_M and $(\mathbf{G}_S + off \mathbf{J}_{SP})$ share a point where 90% of the weights of each distribution lie below and 10% above. This strategy relies mostly on the magnitude of LOS weights while at the same time being reasonably immune to outliers.

Dynamic range compression serves to emphasise the relative importance of edges weights with low noise at the top of the distribution over graph edges with high relative noise. It is combined with normalization of the graph edges between 0 and 1:

$$\begin{aligned} \mathbf{G}_M^{(C)} &= (c_1 \mathbf{G}_M + c_2 \mathbf{J}_M)^\alpha \\ \mathbf{G}_S^{(C)} &= (c_1 (\mathbf{G}_S + off \mathbf{J}_S) + c_2 \mathbf{J}_S)^\alpha \\ c_1 &= 1 / (\max(\max(\mathbf{G}_M), \max(\mathbf{G}_S + off \mathbf{J}_S)) - \min(\min(\mathbf{G}_M), \min(\mathbf{G}_S + off \mathbf{J}_S))) \\ c_2 &= -c_1 \min(\min(\mathbf{G}_M), \min(\mathbf{G}_S + off \mathbf{J}_S)) \end{aligned} \tag{15}$$

α is the parameter that governs the amount of compression applied. The reference implementation chooses $\alpha = 2.5$ which is empirically determined. $\mathbf{G}_M^{(C)}$ and $\mathbf{G}_S^{(C)}$ are the normalized and compressed adjacency matrices. \mathbf{J}_M and \mathbf{J}_S are logical matrices indicating the non-zero elements of \mathbf{G}_M and \mathbf{G}_S . c_1 and c_2 are derived from minimal and maximal edge weight in order to normalize the graphs.

Vertex noise elimination: The underlying idea is to mathematically transform the wireless network's connectivity graph into an analogue representation where vertex noise is baseless. In equation 4 each wireless node transmits a fixed amount of RF-energy (normalized to 0 dB) into the channel of which a small amount is received at its neighbours, represented by the edge-weights. Vertex noise is a systematic deviation between simulation and measurement where the energy balance between a node and its neighbours is biased. In the analogue representation, the energy balance is not normalized to the transmitted energy, but to the received energy: Find a multiplier for each vertex such that the sum of the received energy is 0 dB (in fact 1, because of normalization) in the transformed adjacency matrices $\mathbf{G}_M^{(S)}$ and $\mathbf{G}_S^{(S)}$:

$$\begin{aligned}\mathbf{G}_M^{(S)} &= \mathbf{D}_M \mathbf{G}_M^{(C)} \mathbf{D}_M \\ \mathbf{G}_S^{(S)} &= \mathbf{D}_S \mathbf{G}_S^{(C)} \mathbf{D}_S\end{aligned}\tag{16}$$

\mathbf{D}_M and \mathbf{D}_S are diagonal matrices with the vertex multipliers. $\mathbf{G}_M^{(S)}$ and $\mathbf{G}_S^{(S)}$ are doubly stochastic, symmetric matrices (columns and rows sum to 1). The elements are fractions between 0 and 1 and express how much a particular edge contributes to the hypothetical 0 dB received energy figure.

Finding the multipliers and stochastic matrices is achieved by applying the Sinkhorn-Knopp (SK) algorithm [SIN67] on the normalized matrices $\mathbf{G}_M^{(C)}$ and $\mathbf{G}_S^{(C)}$. This works as advertised as long as they are fully indecomposable. This seems to be the case as long as the underlying graphs are connected, i.e. the previous thinning out of edges has not lead to any unconnected subgraphs.

The transformed graphs $\mathbf{G}_M^{(S)}$ and $\mathbf{G}_S^{(S)}$ allow us to reformulate the original problem 4 as:

$$\mathbf{G}_M^{(S)} = \mathbf{P}^T \mathbf{G}_S^{(S)} \mathbf{P} + \mathbf{N}^{(S)}\tag{17}$$

Since $\mathbf{G}_M^{(S)}$ and $\mathbf{G}_S^{(S)}$ are doubly stochastic, $\mathbf{N}^{(S)}$ is zero-mean and therefore per definition free of vertex noise. Moreover, the original vertex noise does not otherwise significantly magnify $\mathbf{N}^{(S)}$. Temporarily setting $\alpha = 1$, $\mathbf{N}^{(S)}$ can be expressed as:

$$\begin{aligned}\mathbf{N}^{(S)} &= \mathbf{D}_M (c_1 \mathbf{P}^T \mathbf{G}_S^{(S)} \mathbf{P} + c_1 \text{off } \mathbf{J}_{SP} + c_2 \mathbf{J}_M) \mathbf{D}_M - \mathbf{D}_{SP} (c_1 \mathbf{G}_M^{(S)} + c_2 \mathbf{J}_{SP}) \mathbf{D}_{SP} + \\ &+ c_1 \mathbf{D}_M (\mathbf{N}_V \mathbf{J}_{SP} + \mathbf{J}_{SP} \mathbf{N}_V + \mathbf{N}_E) \mathbf{D}_M + c_1 \mathbf{D}_{SP} (\mathbf{N}_V \mathbf{J}_{SP} + \mathbf{J}_{SP} \mathbf{N}_V + \mathbf{N}_E) \mathbf{D}_{SP}\end{aligned}\tag{18}$$

The first line of equation 18 contains two terms that apply the vertex multipliers \mathbf{D}_M and $\mathbf{D}_{SP} = \mathbf{P} \mathbf{D}_S \mathbf{P}^T$ to the respective other adjacency matrix, but as defined it does so by observing the right permutational order. In the second line the original noise is down-scaled using the vertex multipliers. This equation can be further analysed in two important special cases.

Case 1 $\alpha = 1, \mathbf{P} = \mathbf{I}, \mathbf{N}_V = \mathbf{0}$ (edge noise only) – The inputs into the SK algorithm (equation 16) differ only by \mathbf{N}_E which is zero-mean (eq. 4). By the rules of the SK algorithm, this leads to identical vertex multipliers, i.e. $\mathbf{D}_S = \mathbf{D}_M$. The first line of equation 18 can therefore be rewritten to $\mathbf{P}^T \mathbf{G}_S^{(S)} \mathbf{P} - \mathbf{G}_M^{(S)}$, which is $-\mathbf{N}^{(S)}$. The whole equation simplifies to $\mathbf{N}^{(S)} = c_1 \mathbf{D} \mathbf{N}_E \mathbf{D}$: In the absence of vertex noise the preprocessed noise $\mathbf{N}^{(S)}$ is no more than correctly scaled edge noise.

Case 2: $\alpha = 1, \text{off} = 0, \mathbf{P} = \mathbf{I}, \mathbf{N}_E = \mathbf{0}$ (vertex noise only) – This case is studied by first introducing modified vertex noise \mathbf{N}_V' and modified edge noise \mathbf{N}_E' :

$$\mathbf{G}_M = \mathbf{N}_V' (\mathbf{P}^T \mathbf{G}_S \mathbf{P} + off \mathbf{J}_{SP}) \mathbf{N}_V' + \mathbf{N}_E'$$

The modified definition is largely equivalent to equation 4, the difference being that vertex specific deviations between \mathbf{G}_M and \mathbf{G}_S are represented by a multiplicative factor \mathbf{N}_V' , not a summand. From $\mathbf{N}_E' = 0$ and the other preconditions follow $\mathbf{G}_S^{(S)} = \mathbf{G}_M^{(S)}$, because the multiplicative \mathbf{N}_V' is drawn into the vertex multiplier \mathbf{D}_S by the SK algorithm. From there follows $\mathbf{N}^{(S)} = 0$. Under the preconditions of case 2, preprocessing perfectly eliminates modified vertex noise. Given the more useful original definition of vertex noise instead, but also given $\mathbf{N}_V < \mathbf{G}_M, \mathbf{G}_S$, it can be shown that there remains some residual $\mathbf{N}^{(S)}$ which is however zero mean and therefore no longer behaves like vertex noise.

In general of course, $\mathbf{P} \neq \mathbf{I}$, that is after all the point of address assignment. But the properties demonstrated by the special cases transfer well to the general case: Vertex noise is to a certain degree suppressed, to a certain degree mutated into edge noise and the latter transforms proportionally with the adjacency matrices.

The downside of vertex noise elimination is that the edge weight distribution of $\mathbf{G}_M^{(S)}$ and $\mathbf{G}_S^{(S)}$ compares unfavourably with the situation before vertex noise elimination – almost all edges very close to zero. Still, it is worth the effort in case of the demonstration data set of section 5. Whether this is true in all cases, remains to be subject of future research.

4.6 Weighted Graph Matching

This section discusses the third step (matching algorithm) of the three-step process outlined in section 1.1. The weighted graph matching problem (WGM) is well known in the field of combinatorial optimization. It is equivalent to the quadratic assignment problem, which is NP-complete. There exist no efficient algorithms for finding optimal solutions. Several heuristic approaches have been surveyed for suitability with respect to the address assignment procedure [CON04][ALM93][ZAV06][UME88][RAN96][ZAS09]. The *Graduated Assignment Algorithm* of Gold and Rangarajan [RAN96] was chosen, because it compares favourably along several dimensions important to address assignment:

- Performance – in the sense of low order computational complexity; Zaslavskiy et al. lists it amongst the best performing [ZAS09]
- Converges well on noisy data
- Deals with missing and spurious vertices and edges
- Takes advantage of optional vertex attributes
- Supports sparse problems

- Relatively simple and parallelizable algorithm, makes it suitable for implementation with GPGPU (General Purpose Computation on Graphics Processing Unit)

The graduated assignment approach is related to deterministic annealing. It relaxes or convexifies the problem by first allowing other values than “0” and “1” in the vertex assignments. Instead of a permutation matrix, the algorithm computes a doubly stochastic matrix, indicating the similarity of vertices given no particular assignment. As the control parameter increases, the problem becomes increasingly non-convex, the doubly stochastic matrix is pushed more and more towards a permutation matrix. The algorithm thereby avoids local minima, because it started off in the global minimum of the convex problem and tracks the evolution of this global minimum from the maximally convex version of the problem to the original non-convex version.

The algorithm is implemented using a combination of Matlab and CUDA. It performs well (minutes) for graphs up to 2000 vertices and 80000 edges, at which point the necessary data structures become too large to fit into GPU⁵ memory. Details of the implementation are presented subsection 4.6.1.

Weighted graph matching is least square error optimization. This leads to ambiguity in the interpretation of matching errors, i.e. mis-assigned nodes. An error can mean one of two things:

- The weighted graph matching algorithm converged to a sub-optimal solution
- The weighted graph matching algorithm converged to the optimal solution and noise was biased in a way that made the wrong solution optimal (c.f. section 4.2)

For all but for toy-sized problems (less than approx. 11 vertices), it is essentially impossible to compare the result with the real optimal solution. Instead, the cases can be distinguished by taking the correct solution (if known) and subjecting it to the hill climbing optimization (section 4.6.2). If the hill climb algorithm finds the same solution as the WGM algorithm, than the assignment error is a result of the noise not of the WGM algorithm. The converse is, however, not true.

The number of vertices in measured and simulated graph may be unequal. Failure to transmit the graph assembly table should not normally be the cause of missing vertices, because the existence of the underlying node is evident in the graph assembly tables of neighbouring nodes. Not working, not installed or spuriously installed nodes, however, result in graphs of unequal sizes. The weighted graph matching algorithm deals with that by introducing a suitable number of slacks. They should be introduced only in this stage, because zero-sum rows and columns cause problems during vertex-noise elimination.

⁵ Nvidia GeForce GTX 560 Ti, 384 CUDA cores, 2 GiB RAM

The matching algorithm also deals with anchors. Anchors are nodes for which the address assignment is known a-priori. Typically, this is the network's coordinator, nodes specially joined to the network for the sole purpose of commissioning, or any other node that the commissioning technician visits manually during the procedure. In terms of the matching algorithm, anchors are similar to vertex attributes (like device type), in that they amend the objective function and thereby bias or restrict the state space. Indeed, the matching algorithm is able to take advantage of anchors to guide the matching procedure toward high-quality node assignments.

4.6.1 Implementation and Performance of the WGM Algorithm

The Graduated Assignment Algorithm of Gold and Rangarajan is a deterministic annealing approach, designed to minimize the following objective function (nomenclature adapted to match the rest of the document) [RAN96]:

$$\begin{aligned} \hat{\mathbf{P}} = \arg \min_{\mathbf{M}} E_{wg}(\mathbf{M}) &= \arg \min_{\mathbf{M}} \left[- \sum_{a=1}^{s_1} \sum_{i=1}^{s_2} \sum_{b=1}^{s_1} \sum_{j=1}^{s_2} m_{ai} m_{bj} c_{aibj}^{(1)} - \sum_{a=1}^{s_1} \sum_{i=1}^{s_2} m_{ai} c_{ai}^{(2)} \right] \\ \text{subject to } \forall a \ i \ m_{ai} \in \{0,1\} \ , \ \forall a \ \sum_{i=1}^{s_2} m_{ai} \leq 1 \ , \ \forall i \ \sum_{a=1}^{s_1} m_{ai} \leq 1 & \\ \text{where } c_{aibj}^{(1)} = \begin{cases} 0 & \text{if either } (g_1)_{ab} \text{ or } (g_2)_{ij} \text{ is NULL} \\ 1 - v^{(1)} |(g_1)_{ab} - (g_2)_{ij}| & \text{otherwise} \end{cases} \\ \text{and } c_{ai}^{(2)} = \sum_{k=1}^{na} 1 - v_k^{(2)} |(t_{k1})_a - (t_{k2})_i| & \end{aligned} \quad (19)$$

$\hat{\mathbf{P}}$	Sought after permutation matrix ($s_1 \times s_2$), optimization goal
\mathbf{M}	$s_1 \times s_2$ matrix, independent variable of the objective function $E_{wg}(\mathbf{M})$, constrained to be a permutation matrix
$\mathbf{G}_1, \mathbf{G}_2$	Adjacency matrices ($s_1 \times s_1, s_2 \times s_2$) of the two graphs to be matched against each other; in the context of automatic address assignment they are the adjacency matrices $\mathbf{G}_M^{(S)}$ and $\mathbf{G}_S^{(S)}$, measured and simulated connectivity graph after pre-processing
$\mathbf{T}_{k1}, \mathbf{T}_{k2}$	vectors (length s_1, s_2) of $k = 1..na$ vertex attributes; in the context of automatic address assignment they are the node types \mathbf{T}_{kM} and \mathbf{T}_{kS} defined in equation 6
s_1, s_2	Number of vertices of \mathbf{G}_1 and \mathbf{G}_2 respectively
$\mathbf{C}^{(1)}$	Four-dimensional matrix ($s_1 \times s_2 \times s_1 \times s_2$); the elements $c_{aibj}^{(1)}$ capture the distance between the edges of \mathbf{G}_1 and \mathbf{G}_2

$\mathbf{C}^{(2)}$	$s_1 \times s_2$ matrix; the elements $c_{ai}^{(2)}$ capture the distance between attribute vectors \mathbf{T}_{k1} and \mathbf{T}_{k2} of the vertex pair a, i
$v^{(1)}, v^{(2)}_k$	Normalization constants that make $\mathbf{C}^{(1)}$ and $\mathbf{C}^{(2)}$ zero mean

Equations 6 and 19 designate the same optimization problem. The latter ansatz defines $\mathbf{C}^{(1)}$ as kind of penalty function, specifying which kinds of edges in \mathbf{G}_1 and \mathbf{G}_2 are alike, and uses the permutation matrix \mathbf{M} as filter to gather exactly those $c_{aibj}^{(1)}$ that have big contributions towards the objective. The zero-meanness of $\mathbf{C}^{(1)}$ and $\mathbf{C}^{(2)}$ are not so important to the objective function, but are needed to ensure numerical stability of the graduated assignment algorithm. Without it, the exponentiation step below easily exceeds the valid domain of the floating point type *double*.

The algorithm as implemented in the technology demonstration consists of the following steps. For a thorough discussion, confer to the Gold and Rangarajan's publication [RAN96]:

1. Initialize β to β_0 , initialize the elements of \mathbf{M}^s to $(1 + \varepsilon)$
2. **Begin A:** Do until $\beta > \beta_f$
3. **Begin B:** Do until \mathbf{M} converges $\| \mathbf{M} - \mathbf{M}_{prev} \|_1 \leq \varepsilon_B$ (or maximum iterations i_b)
 4. Calculate the partial derivative \mathbf{Q} , the elements are given by:

$$q_{ai} = \frac{\partial E_{wg}(\mathbf{M}')}{\partial \mathbf{M}'} \Bigg|_{\mathbf{M}'=\mathbf{M}} = c_{ai}^{(2)} + \sum_{b=1}^{s_1} \sum_{j=1}^{s_2} m_{bj} c_{aibj}^{(1)}$$
 5. Calculate $\mathbf{M} = \exp_2((\beta/s) \cdot \mathbf{Q})$
 6. Incorporate known anchors into \mathbf{M}
 7. Add slack variables $\mathbf{M}^s = \mathbf{M}$
Subject \mathbf{M}^s to the Sinkhorn-Kopp algorithm (ε_{SK}, i_{SK}), c.f. section 4.5
Remove slack variables $\mathbf{M} = \mathbf{M}^s$
 8. Inject noise: $\mathbf{M} = \mathbf{M} \cdot \mathcal{U}(0.99, 1.01)$.. uniform distribution
9. **End B**
10. $\beta = \beta \cdot \beta_r$

End A

Variables of the algorithm:

\mathbf{M}	$s_1 \times s_2$ matrix, also holding the final result
--------------	--

\mathbf{M}^s	$s_1+1 \times s_2+1$ matrix, may occupy the same space as \mathbf{M} but adds one row and one column as slack variables during the SK-step
\mathbf{Q}	$s_1 \times s_2$ matrix, intermediary result
$\mathbf{C}^{(1)}, \mathbf{C}^{(2)}$	Defined by equation 19
$\beta, \beta_0, \beta_r, \beta_f$	The control parameter of the deterministic annealing process, its start value ($\beta_0 = 2$), the increase ($\beta_r = 1.05$) and the end value ($\beta_f = 250$)
$\varepsilon_B, \varepsilon_{SK}, i_B, i_{SK}$	Parameters of the convergence test (see below)
s	$s = \max(s_1, s_2)$

Convergence of step 3 is tested by taking the absolute accumulated deviation between \mathbf{M} and \mathbf{M}_{prev} of the previous pass (Manhattan norm). The threshold for convergence is $\varepsilon_B = 0.25$. If there is no convergence, the iteration is short circuited after $i_B = 15$ passes. The SK-algorithm of step 6 uses a convergence threshold of $\varepsilon_{SK} = 0.05$ and short circuits after $i_{SK} = 30$ passes.

The algorithm deviates from literature in three points:

- In the exponentiation (step 5) the control parameter β is used relative to the problem size s . The advantage is that the point of convergence of the algorithm becomes independent of the problem size, which otherwise wouldn't be the case.
- Step 6 is not found in the original description: Information about anchor nodes are incorporated by forcing the corresponding rows and columns of \mathbf{M} , corresponding to the respective vertices in \mathbf{G}_1 and \mathbf{G}_2 , to $\{0, 1\}$ as if it already were the finished permutation matrix. The alternative would be to model the anchors as vertex attributes and integrating them into $c^{(2)}_{ai}$. Instead, the information is exploited during the *softassign* step of the algorithm which is replaced by a definitive assignment for pairs of vertices where the mapping is available a-priori. The approach is especially usefully for noisy problems, where a small number of anchors can direct the algorithm towards convergence.
- Step 8 is not found in the original description: Noise injection is a common technique in deterministic annealing to support the algorithm to get out of unstable equilibria which sometimes manifest during phase transitions. Phase transitions are values of the control parameter where the state space undergoes substantial mutation on the way between the smooth problem and the rough problem [ROS10].

Two implementations have been created, one Matlab only, the other consisting of a mixture of Matlab and CUDA kernels (compute 2.0 architecture). Step 4, calculation of the partial derivative $\partial E_{wg} / \partial \mathbf{M}$, dominates the performance of the algorithm. The Matlab only implementation uses a pre-computed $\mathbf{C}^{(1)}$, which is expressed as two-dimensional cell array of two-dimen-

sional sparse matrices. This is practical until graph sizes of 120 vertices, where computation time reaches 19 min and the pre-computed data structure reaches 570 Mb (Intel Core i5-760, 2.8 GHz, 4 cores).

The largest potential performance gains lay in algorithmically optimizing the deterministic annealing process, especially the annealing rate β_r . Experiments with adaptive annealing rates did not deliver satisfactory results. More specifically, there are performance gains for datasets with low noise level, but performance suffers for datasets with typical noise levels and low separability as encountered in section 5. The reason is related to the behaviour during phase transitions:

Towards the end of the annealing process, as almost all vertices are already correctly matched, and as the algorithm goes through the final few phase transitions, the assignment of the few remaining vertices proves to be problematic. As the control parameter increases further, a decision is forced. During this stage, the non-adaptive algorithm often fails to converge to a fix-point (ε_B), but terminates the iteration due to reaching the maximum number of iterations (i_B). The adaptive version of the algorithm would at that point decrease the annealing rate and backtrack. This of course costs performance-wise. With the annealing rate deceased to the minimum, the adaptive algorithm finally advances through the phase transitions and terminates having spent almost all processing time on optimizing the assignment for a minority of nodes that in all likelihood cannot be correctly matched anyway, because they exhibit separability below the noise level.

Some effort went into identifying phase transitions that “are worth it” versus phase transitions like the ones described above, with mixed success. In the end, the adaptive version of the algorithm did exhibit performance gains for non-noisy datasets and comparable performance for noisy datasets, albeit at the price of increased complexity with regards to steering the annealing process and several empirical thresholds for recognizing those phase transitions that are worth the increased precision and processing time. Those thresholds might or might not work for datasets other than the ones encountered during development. An algorithm that performs slightly worse but exhibits greater generality is preferable to one that performs better in a subset of cases but may fail with new and unknown data. The approach was therefore abandoned and the presented algorithm uses a fixed annealing rate.

The chosen values for the parameters governing the annealing rate ($\varepsilon_B, \varepsilon_{SK}, i_B, i_{SK}$) are relatively conservative. Immediate, substantial performance gains, in the sense of running time, can be realized by lowering the i_x and increasing the ε_x . But of course, those values present trade-offs. The quality of the recovered vertex assignment will suffer at some point.

Performance improvements without quality compromises depends on optimizing the loops, especially the innermost loop (step 4), which is $O(s^4)$ for every single annealing step (the enclosing loops are bound by i_B and β_f). Pre-computation of $\mathbf{C}^{(1)}$ is impractical because of GPU memory limitations, but sparse matrix representation of \mathbf{G}_1 and \mathbf{G}_2 brings complexity of step 4 down to $O(s^2r^2)$, where r is the maximum number of neighbours of any graph vertex. For practical network connectivity graphs the neighbour count r is clearly much lower than the problem size s , or at least can be made so by pruning graph edges with very small weight (discussed in section 4.5). Practical graphs encountered during testing and created by the simulation have r roughly half of s , which lowers computation time by a factor of 4, all else being equal. Other than the problem size s , the maximum number of neighbours r is therefore a crucial performance figure.

The kernel for computing step 4 is organized around computing values q_{ai} in parallel. The indices a and i are therefore fixed in each thread⁶, whereas indices b and j are iterated over. More precisely, threads traverses sparse rows of \mathbf{G}_1 and \mathbf{G}_2 , which are stored in row-major “lists-of-lists” format, where individual entries contain the value and the column index. Indices b and j are hence not tightly iterated over, but loaded from the sparse matrix representation. This constitutes the remaining performance issue, because the necessary load from matrix \mathbf{M} , addressed via b and j is necessarily non-coalesced and hence serialized amongst the threads.

Non-coalesced access to matrix \mathbf{M} is avoided by the following strategy: The computation of q_{ai} is organized column-wise, such that index i is fixed across threads of the computation block. Furthermore, the loops b and j are organized such that j is the outer loop. Due to the fixed index i , when threads load index j from the sparse matrix \mathbf{G}_2 , they are guaranteed to find the same j and $(g_2)_{ij}$ over the whole computation block. The advantage of fixing j is that the load of m_{bj} now at least targets a single column of matrix \mathbf{M} , which opens up the opportunity of caching said column in shared memory before going into the inner loop over index b . Shared memory is on-chip and has support for non-coalesced access which removes the bottleneck. The overhead of caching the column of \mathbf{M} in the outer loop is small in relation to the additional throughput gained. The performance improvement of the optimized kernel is an additional 15 % (measured for $s = 500$, $r = 160$, computation time 240 s).

Other than the computation of \mathbf{Q} , the algorithm uses a number of simple reduction kernels, or in the case of the SK-algorithm (step 7), a reduction followed by scalar multiplication. Comparing the performance of the CUDA version to the Matlab only version, the same

⁶ “Thread”, “block” and “shared memory” are terms proprietary to Nvidia’s CUDA technology. The corresponding OpenCL terms are “work-item”, “work-group” and “local memory”.

120 node problem mentioned above takes 18 s to solve with the optimized CUDA version, which is 63 times faster than the Matlab only version.

4.6.2 Clean-Up Heuristic and Hill-Climber

On noisy, poorly separable data, the algorithm fails to unambiguously commit to a matching on a subset of nodes. More concretely – after the control parameter β reaches its end value β_f , the matrix \mathbf{M} has typically not converged to a permutation matrix. Increasing β_f is not an attractive option, because it elongates the computation time without improving the quality of the match, since the ambiguous nodes exhibit too little separability to be solved by least squares approximation. Instead, the matrix \mathbf{M} is coerced to a permutation matrix by a simple heuristic.

Analysing the result of the matching algorithm, matrix \mathbf{M} : From the SK-step (step 7) follows $m_{ai} \leq 1$. Pairs of vertices that correspond to each other with high certainty produce $m_{ai} \approx 1$. In order to avoid unmatched nodes or nodes with duplicate assignment, matches with high certainty are given priority. The heuristic works as follows:

1. **Begin A:** Iterate until $\mathbf{M}: m_{ai} \in \{0, 1\}$
2. Find maximum element of \mathbf{M} for which $m_{ai} < 1$,
set it to 1, set all elements of the same column and row to 0

End A

The approach may leave some nodes unmatched. For non-square problems, this is a desired result, since the algorithm thereby identifies nodes from the bigger graph having no correspondence in the smaller graph. If furthermore the correct nodes are identified as having no correspondence, they can be counted as correctly matched nodes. For square problems, any unmatched nodes may be assigned randomly which is no worse or better than leaving them unassigned and counting them as incorrectly matched by default.

The hill-climber is a useful tool for classifying test data into matching problems that potentially have a perfect solution with the least squared error approach and tell them apart from problems that have not. The hill-climber is of no use for doing the matching itself, because it inevitably runs into local minima after a few steps. It is also of no use for improving the result of the global weighted graph matcher, for if the latter has not found the global optimum, then a nearby local optimum encountered by the hill-climber is no better than any other nearby assignment. The remaining ambiguity is likely the result of poor separability, which is not anything that a hill-climb could fix.

Starting with known graphs \mathbf{G}_M , \mathbf{G}_S and the correct permutation \mathbf{P} , if the hill-climber finds a transposition that improves upon \mathbf{P} , then \mathbf{P} is not optimal and the quality of the simulation flawed, even though the variance may be low. This can be used as a starting point for further analysis, such as the study of simulation performance of the nodes identified by the hill-climber. The hill-climber works by directly computing equation 6 (the problem formalization), and thereby restricting the search to $\mathbf{Q} = \mathbf{P}\mathbf{Q}_\Delta$, where \mathbf{Q}_Δ is a transposition (a permutations of exactly two nodes).

5 Experimental Evaluation of Automatic Address Assignment

The theoretical analysis of automatic address assignment (section 4) are complemented by experimental results. The strengths and weaknesses are presented by means of a publicly available data set (section 5.1) and experiments conducted by the author at the Vienna University of Technology (section 5.2). Performance of the matching algorithm is demonstrated by means of synthetic data (section 5.3).



Figure 25: Floor plan of the Mannheim site as viewed in the Winprop simulation tool. "Site x" are WiFi access points, "Node x" are the selected points of measurement and together form the simulated connectivity graph.

5.1 University of Mannheim Data Set

The Mannheim data set leverages a publicly available data set originally collected by King et al. [KIN08]. It consists of 12 IEEE 802.11 “WiFi” access-points and 130 fingerprinting positions of RSSI measurements. The goal is to demonstrate the effectiveness of the approach under favourable conditions, also analysing the nature of vertex noise and edge noise. Out of the 12 access-points 2 are excluded, because one has no data associated and the other is redundant. Of the fingerprinting positions, 24 in regular 4.5 m grid are selected for the demonstration, leading to a total 34 vertices in the data set. Fading effects are countered using a slightly different strategy than the described in section 4.3.1: Instead of measuring at different frequencies, King et al. repeat each measurement several times with different orientations and slight displacements of the reception antenna.

As can be seen in figure 25, the topology is favourable to address assignment. The nodes along the corridor (fingerprinting positions) are aligned in a linear network. The bi-linear parts (upper right of the map) are relatively well separable because of the nodes (WiFi access points) along the periphery. The latter are separable because they are inside separate rooms.

5.1.1 Modelling of the Demonstration Site

The site is modelled with WinProp (cf. figure 25), using the floor plan in the publication as basis and filling in the details about the materials and mounting heights after a personal visit at the site. The walls throughout the surveyed corridors and offices are metal-plated, so are the doors. This makes the radio environment of the site somewhat unusual if not especially challenging. Wave propagation between rooms therefore happens primarily above the suspended ceiling. The high reflectivity of the environment requires a higher than usual number of samples (rays) to adequately approximate the RF-energy distribution.

Not included in the model are furniture, conduits above the suspended ceiling, columns of the building's reinforced frame, as well as details of the façade. Sending and reception antennas are modelled as dipoles. The state of the metal doors (fully open / half open / closed) at the

Table 3: The most important configuration parameters of the ray-tracing simulation with Winprop

Max transmissions / reflections / diffractions / scattering of each ray	4 / 4 / 2 / 1
Total max interactions (reflections + diffractions + scattering)	4
Uncorrelated simulation (not coherent, no phase)	
Empirical loss model (not based on Fresnel equation)	

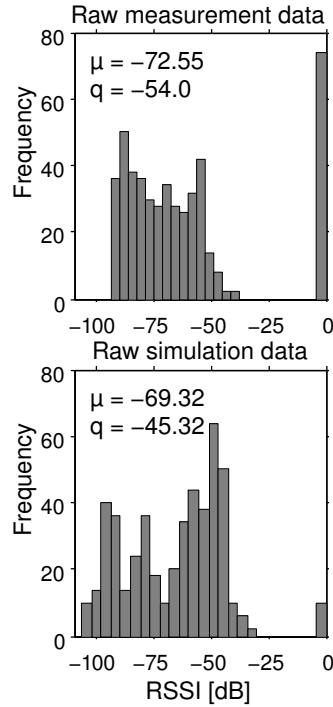


Figure 26: Raw measurement data vs. simulation data. Measurements are cut off at -95 dBm due to interference and receiver sensitivity. 'q' is the 9th 10-quantile used for estimating the global offset. Zeros stand for non-existent edges, they are not counted towards μ and q . The standard deviation after correction of the global offset, only counting edges that exist in both graphs, is $\sigma_{Mannheim} = 8.95 \text{ dB}$.

time of the measurements is unknown. In the simulation all doors are treated as closed. This may be one of the major contributors to noise in the test dataset.

Table 3 details the parametrization of the Winprop simulation. Figure 26 is a comparison between [KIN08]'s measurement and the simulation. The large deviation of the 9th 10-quantile (global offset *off*) can be attributed to unknowns with regards to the exact details of the RSSI measurement (antenna, cables, WiFi equipment). It is not relevant for the purpose of address assignment and is compensated in the pre-processing stage.

The deviation between simulation and measurement – here treated as vertex noise and edge noise – is depicted in figure 27. Relative edge noise [$\mathbf{N}_{ER} = \mathbf{N}_E \oslash (\frac{1}{2}\mathbf{G}_M + \frac{1}{2}\mathbf{P}^T\mathbf{G}_S\mathbf{P})$] fits a Gaussian model well with $\sigma = 0.08 \text{ dB/dB}$ (\oslash means element-wise division on non-zero elements). As explained in section 4, this is to be expected because noise is the consequence of a multiplicative product of uncertain quantities. Furthermore, edge noise is not constant over the range of edge-weights, with edges at the low end being more noisy than the rest. Extrac-

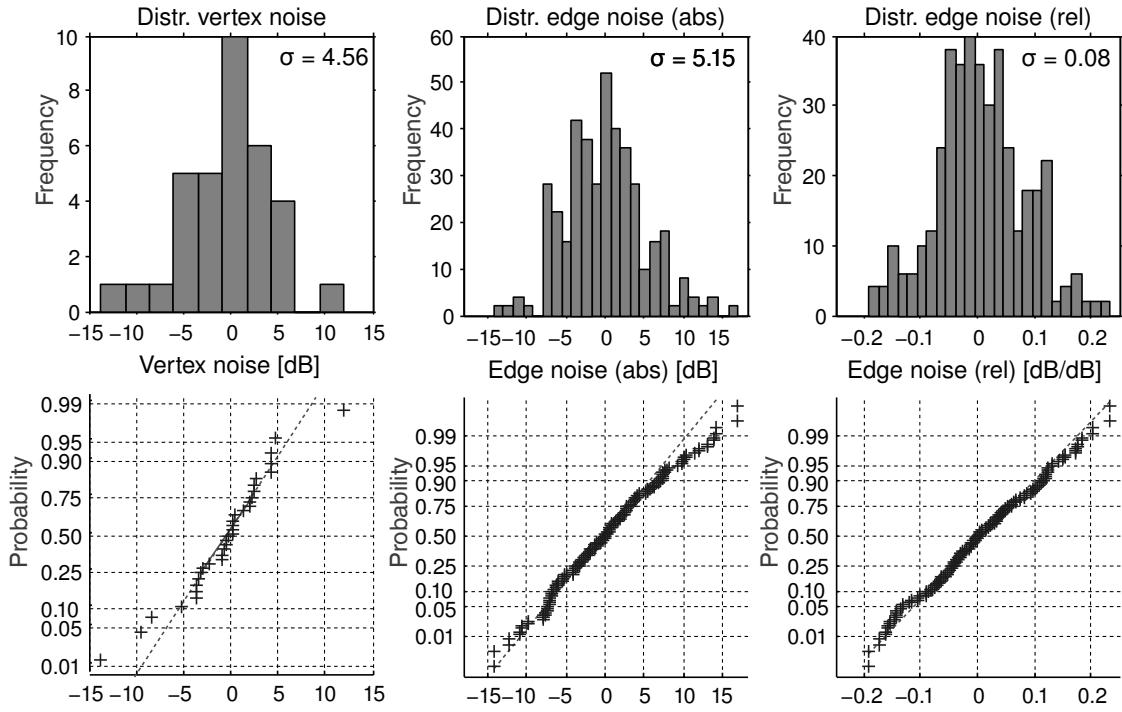


Figure 27: Measured vs. simulated data, noise distribution, no preprocessing except elimination of global offset. Outliers in vertex noise are due to modelling inaccuracies. As expected, edge noise fits the Gaussian model well, especially when expressed relative to edge weight.

tion of vertex noise from the test data set is done by directly solving equation 4 and optimizing for minimal edge noise⁷.

5.1.2 Completing the Test Data Set

The data set of [KIN08] was recorded for the purpose of fingerprinting the channel, rather than address assignment. Instead of wireless node to node measurement of RSSI, it contains measurements from access point to fingerprinting position. What is missing are RSSI measurements from access point to access point and from fingerprinting position to fingerprinting position.

The missing edges of the measured connectivity graph are estimated by taking values from the simulation and overlaying them with the noise profile extracted from comparing the available measurements with the corresponding data from the simulation, cf. figure 27. Vertex noise is applied unaltered, edge-noise is sampled from a random number generator. The measured data set thus consists of 203 real measured edges and 276 estimates with identical

⁷ This is possible with the test data set, because only in this case the permutation \mathbf{P} is known a priori.

noise properties. The resulting synthetic measurement-graph of 34 vertices and 479 edges (measurements and estimates) is a fair and realistic approximation of what could be expected by true node-to-node measurements of RSSI.

The experiment is set up as follows:

1. Analyse noise, fill in missing edges in \mathbf{G}_M , optionally enhance noise artificially
2. Apply random permutation \mathbf{P} in order to get simulated connectivity graph \mathbf{G}_S
3. Pre-processing, weighted graph matching to recover address assignment $\hat{\mathbf{P}}$
4. Compare \mathbf{P} and $\hat{\mathbf{P}}$, count matching errors

No anchors or additional vertex attributes are added to the data set.

5.1.3 Results

There are two sources of randomness in the test data set and graph matching procedure: The random Gaussian noise used to complement the measured connectivity graph, and the random permutation. The results of ten test runs are therefore averaged.

The presented method of address assignment is able to perfectly match (34 of 34 vertices) the test data set. Figure 28 details the matching performance if the noise is artificially inflated – vertex noise, edge noise and both.

Figure 28 suggest that the algorithm is stressed to its limit by the test data set. On the other hand, the noise being scaled is a logarithmic quantity. In the tails of the edge noise distribu-

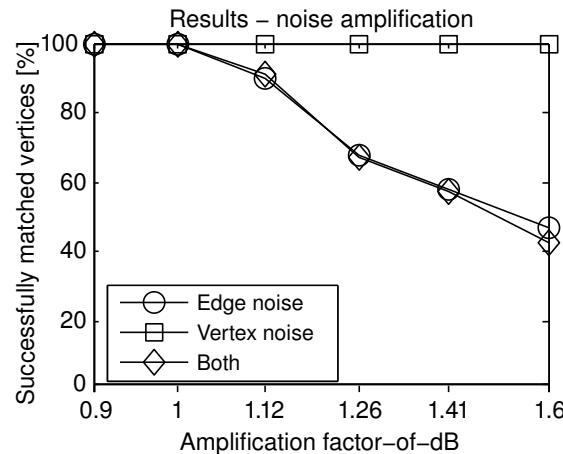


Figure 28: Artificial amplification of noise in the test data set. The factor in the abscissa scales a logarithmic quantity. Vertex noise is well suppressed, edge noise shouldn't be any bigger than $\sigma = 5.15$ dB. 5% more edge noise and sporadic matching errors manifest, 10% more edge noise and performance degrades to 90%.

tion (figure 28, tails at 15 dB, $\sigma = 5.15$ dB) the factor-of-dB of 1.26 means linearly 2.45 times more deviation between measurement and simulation.

5.2 Vienna University of Technology Data Set

The measurements at Vienna University of Technology were conducted in May and June 2013. They consist of several smaller and one large data set spanning rooms on the first, second and third floor of the building Gußhausstraße 27-29 of the Faculty of Electrical Engineering and Information Technology, figure 29. The measurements are presented in three parts: The first to show the impact of topology, the second to show the impact of antennas and the third is a large set of 95 node positions and 1839 unique links.

The building is a 70ies style high-rise office building, 5 floors, based on a reinforced concrete frame with large window fronts. The core structures (staircases and elevator shafts) and the reinforced concrete frame are load-bearing. The premises are used as offices, electronics laboratories, meeting room and lecture hall.

The site is modelled in greater detail than the Mannheim site, including the reinforced concrete frame, vent pipes, metallic white-boards in offices, black-boards of lecture halls, and even select pieces of laboratory and kitchen equipment with metal surfaces greater than 1 m²

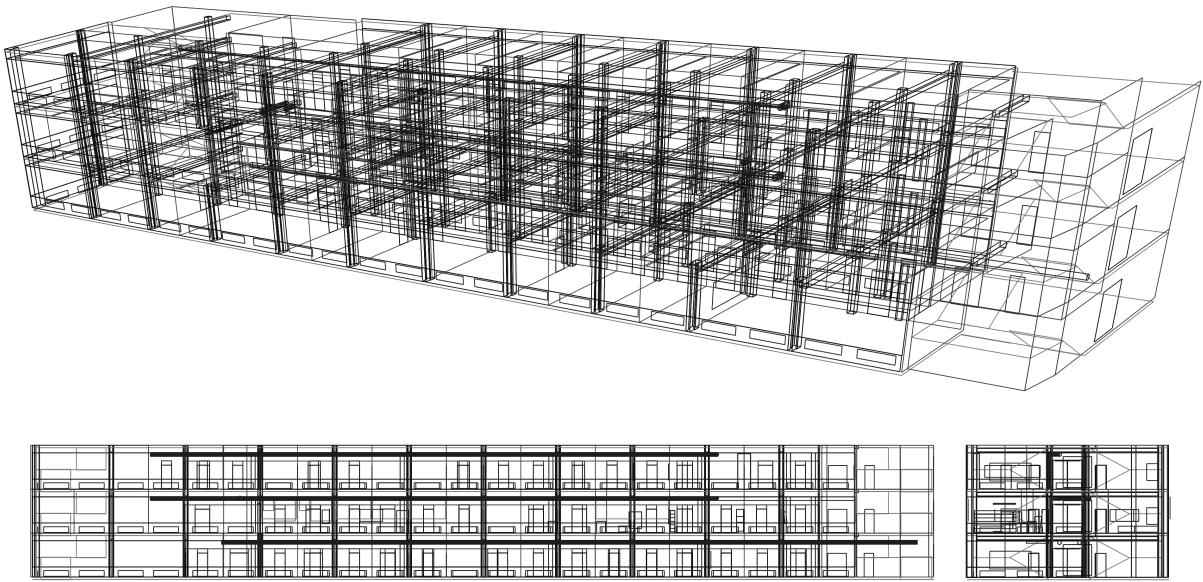


Figure 29: Building wing 'CA', levels 1 to 3, Faculty of Electrical Engineering and Information Technology at the Vienna University of Technology. Clearly visible is the reinforced concrete frame and the basic structure of the building. The model consists of 1574 polygons.

(figure 29). Custom measurements of select building materials, as presented in section 4.3.4, are also incorporated into the simulation.

The demonstration network is inspired by a lighting network, where demonstration nodes take the place of light switches and luminaires. Node type attributes are used to reproduce the distinction. The lighting network is chosen, because it characteristically features a higher density of nodes than the other sections of building automation.

5.2.1 Impact of the Topology

The topology is analysed independent of the weighted graph matching algorithm. Matching errors can occur as a result of noise only, even given a perfect matching algorithm that always finds the global optimum, see section 4.2. Here, a numerical comparison between two small topologies is presented.

The two topologies are (i) an office-like, rectangular topology which according to the theoretical analysis is difficult to handle and (ii) a favourable topology roughly shaped like a diamond. The rectangular dataset is in fact part of a larger dataset, presented in section 5.2.3, whereas the diamond shaped dataset was measured for the sole purpose. No anchors or node type attributes are used.

The analysis uses measured data \mathbf{G}_M only, simulation and noise are only evoked at the end of the section. The standard deviation of the graph edges are $\sigma_{E,Rect} = 8.8$ dB and $\sigma_{E,Diamond} = 5.4$ dB after fading compensation. The reason for the disparity is that the rectangular topology is more spread out, i.e. the difference between closest neighbour and farthest neighbour is on average bigger. Both graphs are normalized to $\sigma_{E,Norm} = 4$ dB standard deviation to facilitate a fair comparison. The concrete value of $\sigma_{E,Norm}$ is only relevant for generating the visualizations of separability. It is chosen to be 4 dB because a different data set used during development (not presented here) happens to exhibit an edge-weight standard deviation in the vicinity of 4 dB.

Separability $\Delta(g_M)_i = (g_M)_{ir} - (g_M)_{is}$ as defined in equation 9 is a property not of single graph vertices, but a property of a pair of vertices (R, S) relative a third vertex (I). Generalizing the idea, the separability of a pair of vertices relative the rest of the graph is given by

$$\delta_{RS} = \sqrt{\sum_{i \notin \{r, s\}} [(g_M)_{is} - (g_M)_{ir}]^2}$$

This is still difficult to visualize, because the δ_{RS} constitute a weighted graph of the same size as \mathbf{G}_M . Instead, each node is assigned the smallest separability that it shares with any other node:

$$\delta_R = \min_{S=1,\dots,n} (\delta_{RS})$$

Along the same lines, equation 9 can be used to directly estimate the probability of mismatching a pairs of nodes as a function of the connectivity graph and noise. Under the assumption of Gaussianity, linear independence and zero mean, it can be further developed to:

$$P\{\varepsilon_{RS}\} = Q\left(\frac{\delta_{RS}}{\sqrt{2}\sigma_N}\right) = Q\left(\frac{1}{\sqrt{2}\sigma_N} \sqrt{\sum_{i \notin \{r,s\}} ((g_M)_{is} - (g_M)_{ir})^2}\right) \quad (20)$$

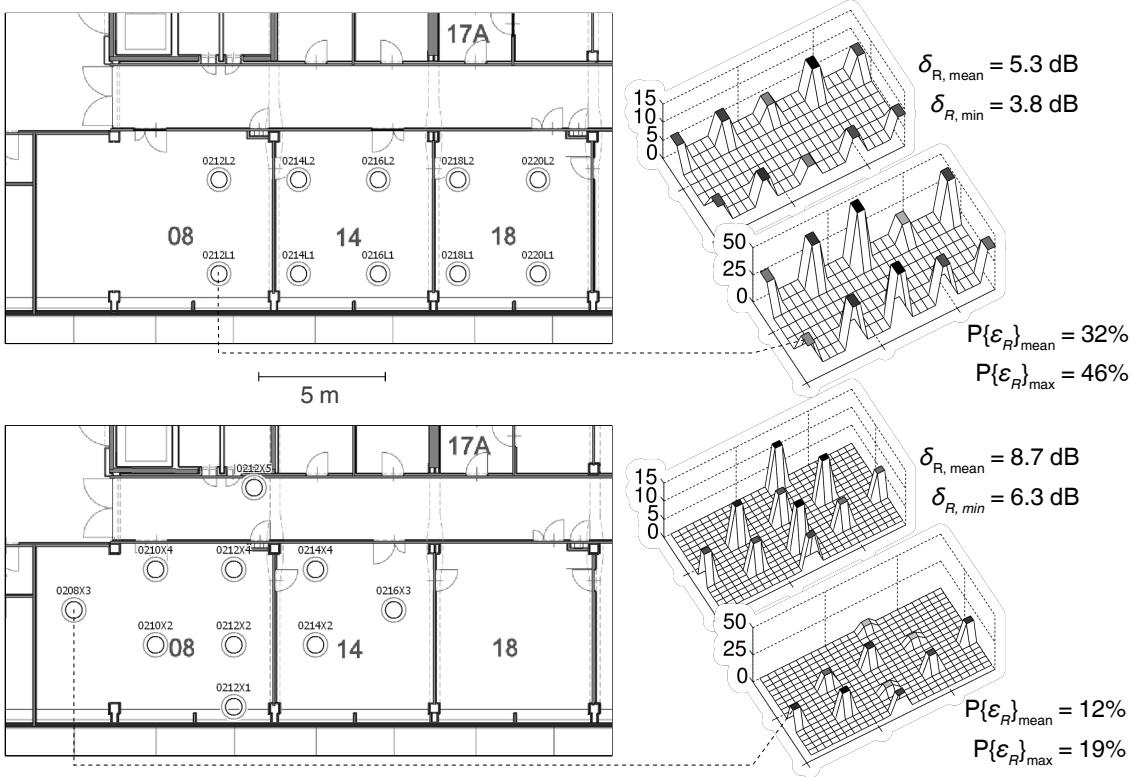


Figure 30: Separability and probability of mismatching a node, as a function of measured connectivity graph. Left – topologies on the floor plan: Rectangular topology (upper), diamond-shaped topology (lower). Right – Separability δ_R and probability of mismatching a node $P\{\varepsilon_R\}$ elaborated for the two topologies. The latter assumes a noise level of $\sigma_N = 4$ dB.

The diamond-shaped topology exhibits more separability and is therefore much less likely to experience mismatches. The highest separability of all, 11.5 dB, is displayed by node 0212X5 (diamond shaped topology, topmost node) located in the hallway and separated by a wall from all other nodes.

The result is again a fully connected, weighted graph where the vertices correspond to vertices of G_M and the edge-weights denote the probability of swapping the two vertices. From there it is straight forward to calculate the probability that a particular vertex is swapped with any other vertex:

$$P\{\varepsilon_R\} = 1 - \left(\prod_{S=1}^n (1 - P\{\varepsilon_{RS}\}) \right)$$

δ_R and $P\{\varepsilon_R\}$ are visualized for the rectangular and the diamond-shaped dataset in figure 30. The diamond-shaped topology has a significantly higher separability and therefore lower error probability.

As also discussed in section 4.2.4, the probabilities shouldn't be taken at face value, because the formal analysis relies on assumptions about the noise matrix that are not true in practise. Also, the noise level σ_N (a figure comparing simulated and measured connectivity graph) is chosen arbitrarily to match the normalized internal standard deviation of the graph edges $\sigma_{E,Norm}$. Given that the Q-function in equation 20 is highly non-linear, the actual probabilities in figure 30 are rough estimates at best. The separability, on the other hand, relies on solid measurement and the normalization constant $\sigma_{E,Norm}$. It is therefore comparable with every other measurement subject to the same normalization. The significance of the exercise is to serve as evidence of the correctness of equation 9 and the topology issue.

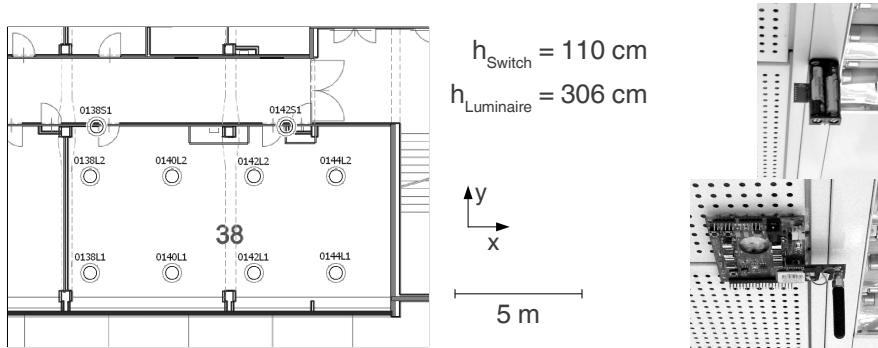
Concluding the analysis, the simulated connectivity graphs of the two topologies are then calculated using the ray-tracer and normalized to $\sigma_{E,Norm}$. The noise levels are found to be $\sigma_{N,Rect} = 2.7$ dB and $\sigma_{N,Diamond} = 2.6$ dB. The hill-climber, starting from the correct matching, finds that the correct is also the optimal matching in case of the diamond-shaped topology and finds the opposite in case of the rectangular topology. With and without edge-weight normalization, the matching procedure is able to solve the diamond-shaped topology flawlessly. The same is impossible for the rectangular topology, because the correct topology cannot be found using least squares approximation.

5.2.2 Impact of Antennas

The aim of this measurement is to assess the impact of different antennas on the measured connectivity graph and to show in how far this impact can be compensated by the inclusion of antenna models into the simulation. Figure 31 shows the different types of nodes and antennas used in the measurement. Since only ten units of the NXP / Jennic evaluation sensor boards are available, the measurement also uses ten positions (eight luminaire positions, two



Figure 31: Left – Minimal wireless node with PCB antenna, 30 units built for the purpose of the demonstration; Figures 22 and 23 show the PCB antenna in more detail, including the radiation pattern. Right – “Sensor Node” as distributed with the NXP / Jennic evaluation kit for the JN5148 wireless transceiver, 10 units available; off-the-shelf 2.4 GHz monopole antenna; note the insufficient ground plane around the base of the antenna. The axes indicate the local coordinate system of the antenna.



*Figure 32: Floor plan of room number CA0138, a small lecture hall / seminar room. Major reflectors are the blackboard at the far left side, and the metal suspended ceiling. 8 nodes occupy luminaire positions (*L1, *L2) and 2 more nodes occupy light switch position (*S1) next to the doors. Right – the nodes mounted in the luminaire positions.*

light switch positions). Figure 32 has the floor plan and close-ups of the node as mounted next to the luminaire.

There are four measured datasets, two with the NXP / Jennic sensor boards and monpole antennas (*M MP*), two with the custom built minimal nodes and printed antennas (*M PCB*). The first measurement of each node type is conducted with the nodes oriented towards the nominal direction, the second measurement with the nodes rotated 90° counter-clockwise around the vertical axis. The rotated measurement, when compared with the nominal measurement, serves to assess the impact of the non-isometry of the radiation pattern.

Table 4: Standard deviation of the path loss between corresponding edges in different datasets.

① The PCB antenna exhibits considerable anisotropy. ② Not so the monopole antenna. The remaining anisotropy can be explained by the missing antenna ground plane. ③ The isotropic antenna model approximates the PCB antenna best – better even than the empirical model created to approximate the PCB antenna. ④ The simulation approximates the rotated measurement dataset better than the nominal measurement. Most probably, because the predominant propagation paths hit extreme values of the antennas lobe structure more often in one dataset than the other.

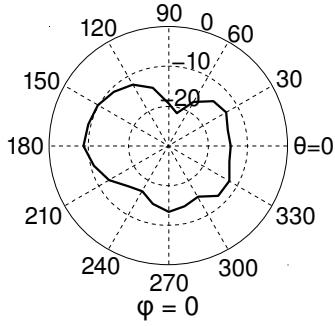
STD [dB]	M PCB	M PCB _{90°}	M MP	M MP _{90°}	S ISO	S PCB	S DP
M PCB		5.0	3.8	3.8	4.6	5.5	4.7
M PCB _{90°}	① 5.0		4.7	4.8	3.7	4.2	4.0
M MP	3.8	4.7		0.7	4.1	5.1	4.1
M MP _{90°}	3.8	4.8	② 0.7		4.2	5.1	4.1
S ISO	↑ 4.6	↔ ④ → 3.7	4.1	4.2		1.7	0.6
S PCB	③ → 5.5	4.2	5.1	5.1	1.7		1.8
S DP	↓ 4.7	4.0	4.1	4.1	0.6	1.8	

There are three simulated datasets, each with a different antenna model. *S ISO* uses no antenna model, *S MP* a simple monopole antenna model and *S PCB* uses the antenna model from figure 23.

The datasets represent path loss between nodes of the measurement run, or respectively simulation run. The datasets are compared by assembling the weighted connectivity graphs (e.g. $G_{M,PCB}$ and $G_{S,ISO}$) and computing the standard deviation of corresponding edge weights in dB. Measurements are fading corrected, but no other preprocessing is applied. All seven datasets are systematically compared against each other. The results are presented in table 4.

A downside of this approach is that the compared edge weights (representing path loss between wireless nodes) are neither statistically independent with respect to sampling of the antenna radiation pattern, nor are they well determined because of multi-path propagation. Almost all nodes have line-of-sight and the network topology is arranged in a regular grid. Most links will therefore hit the antenna radiation pattern in a way that cluster around certain favoured directions. The rotated measurements exhibit the same sampling bias, only that different parts of the radiation pattern are favoured.

Never the less, the results are consistent with theory. For example, the standard edge weight deviation between *M PCB* and *M PCB_{90°}* is 5 dB (table 4, ①). In *M PCB* the predominant clusters can be found in the radiation pattern (figure 23) at ($\phi=0^\circ / \theta=0^\circ$) \approx -15 dB and



Detail of figure 23 (repeated): Radiation pattern [dB] of the minimal demonstration node with PCB antenna, mounted under the metal suspended ceiling.

$(0^\circ/180^\circ) \approx -10$ dB. In $M_{PCB_{90^\circ}}$ they are at $(0^\circ/90^\circ) \approx -20$ dB and $(0^\circ/270^\circ) \approx -15$ dB. This adequately explains the standard edge weight deviation.

Dataset S_{PCB} uses the radiation pattern as antenna model for the ray-tracing simulation. According to table 4, ③ the approach fails to lower the deviation with respect to simulation without antenna model (S_{ISO}). There are several factors:

- The radiation pattern represents one particular node, slightly modified to contact the PCB antenna, mounted to a generic thin metal sheet as part of the measurement procedure. It may not be representative for other nodes mounted at different positions to the metal suspended ceiling.
- The nodes are assembled by hand and magnetically attached to the suspended ceiling. The reproducibility and precision is limited. Both the distance, as well as the mounting angles (φ, θ) with respect to ceiling and luminaire are subject to considerable variations.
- The radiation pattern is incomplete, as it fails to account for polarization and frequency dependent effects (cf. section 4.3.3).

Given these conditions, the discrepancy is no surprise to any person knowledgeable in radio frequency technology. As already stated in section 4.3.3, this is not a problem specific to the VUT dataset. Professional applications of automatic address assignment suffer the same limitations: It is simply impractical to plan, build or model the building, the fixtures and the network components with the necessary precision as to reproduce deterministic near-field antenna conditions. The conclusion is, again, to ignore the antenna problematic and accept it as source of variance.

Section 4.3.3 estimates the standard deviation due to ignoring antennas $\sigma_{N,RP}$ to be 2.4 dB. This number underestimates the antenna issue if the topology is the cause of clustering ef-

fects (table 4, ①). It overestimates the issue if the antenna is relatively isotropic to begin with (table 4, ②). Other numbers in table 4 are not directly comparable with $\sigma_{N,RP}$ because they include other sources of variance (e.g. the ray-tracer in case of ③ and ④).

The datasets and the measurement procedure are highly reproducible. Repeating the same measurement without manipulating the nodes between measurements yields a standard deviation of edge weights of 0.03 dB. The standard deviation between forward and backward measurement between two nodes amounts to 1.78 dB. The effect is mainly due to tolerances of receivers and transmitters, to a lesser degree due to non-reciprocity of the channel.

5.2.3 Large Data Set

The data set contains 95 node positions. Because only 30 nodes are available, the measurement is conducted in several steps. The 30 nodes are split into two groups of 15, and the node positions into seven realms (five realms of 15 positions and two realms of 10 positions). The five realms of 15 positions are ordered left to right in the second floor of the building wing, the two realms of 10 node positions are in the first and third floor, respectively, cf. figure 33. At every step, two adjacent realms are equipped with demonstration nodes and a partial connectivity graph is assembled. The recovered low-level addresses are prefixed with the realm identifier such that every physical node can represent several node positions in the combined connectivity graph.

The measured connectivity graph therefore contains 95 vertices and 1839 edges. Because only physically adjacent realms are measured, potential edges between none-adjacent realms are missing. This does not greatly impair the quality of the dataset because the great majority of these potentially missing edges would not actually result in a successful measurement, since path loss is already below the receiver sensitivity of the demonstration nodes.

The simulation is done in a single run and not realm-by-realm. Edges that lie in non-adjacent realms and therefore cannot exist in the measured connectivity graph, are removed from the simulated connectivity graph. The Winprop model is depicted in figure 29.

The premises contain several aspects that present a challenge to the simulation:

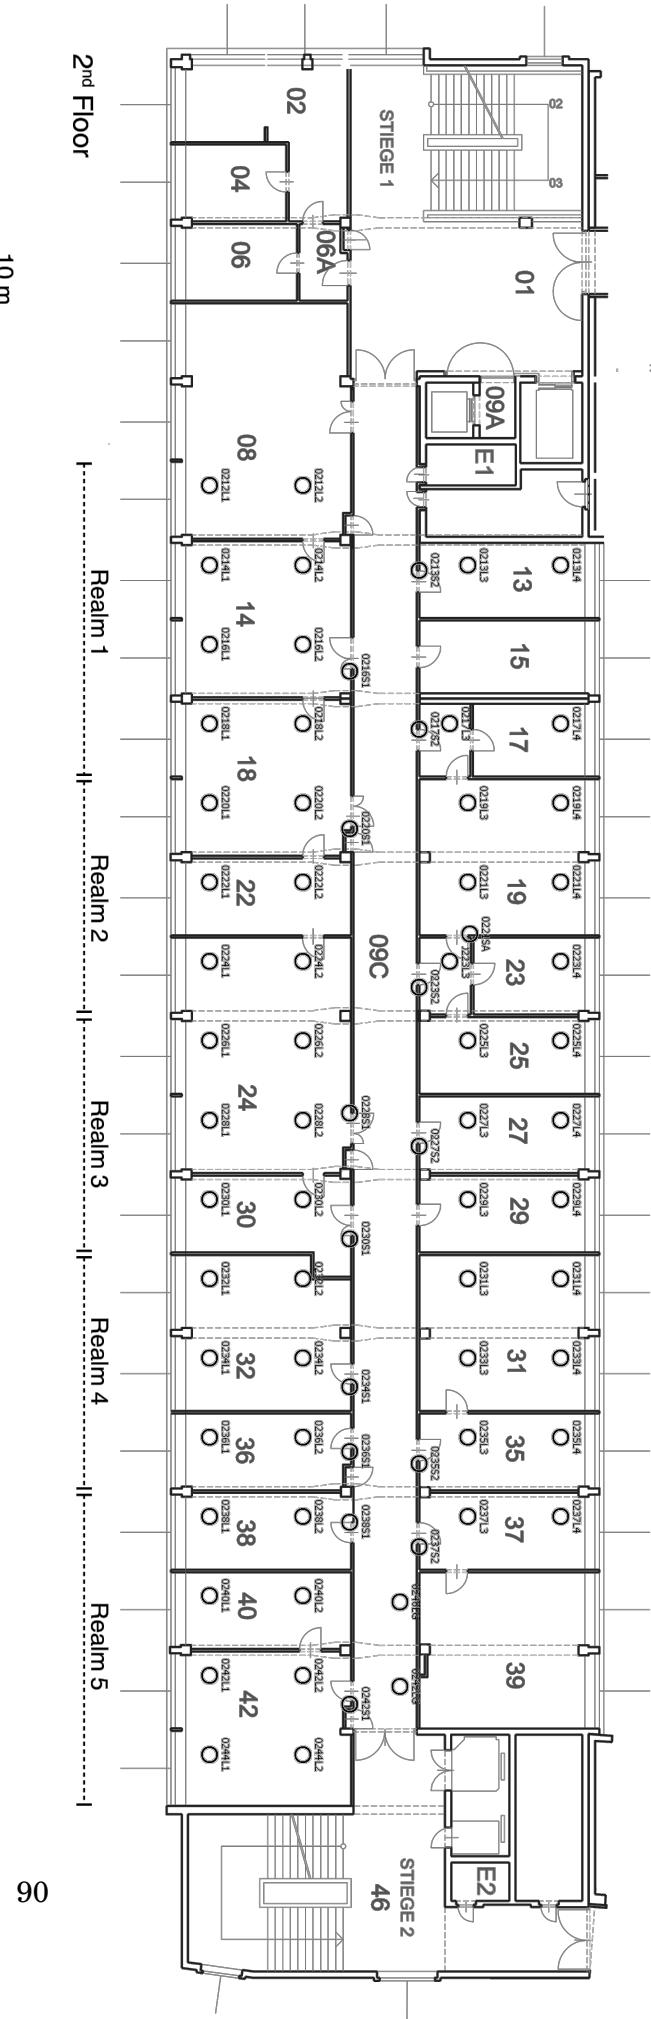
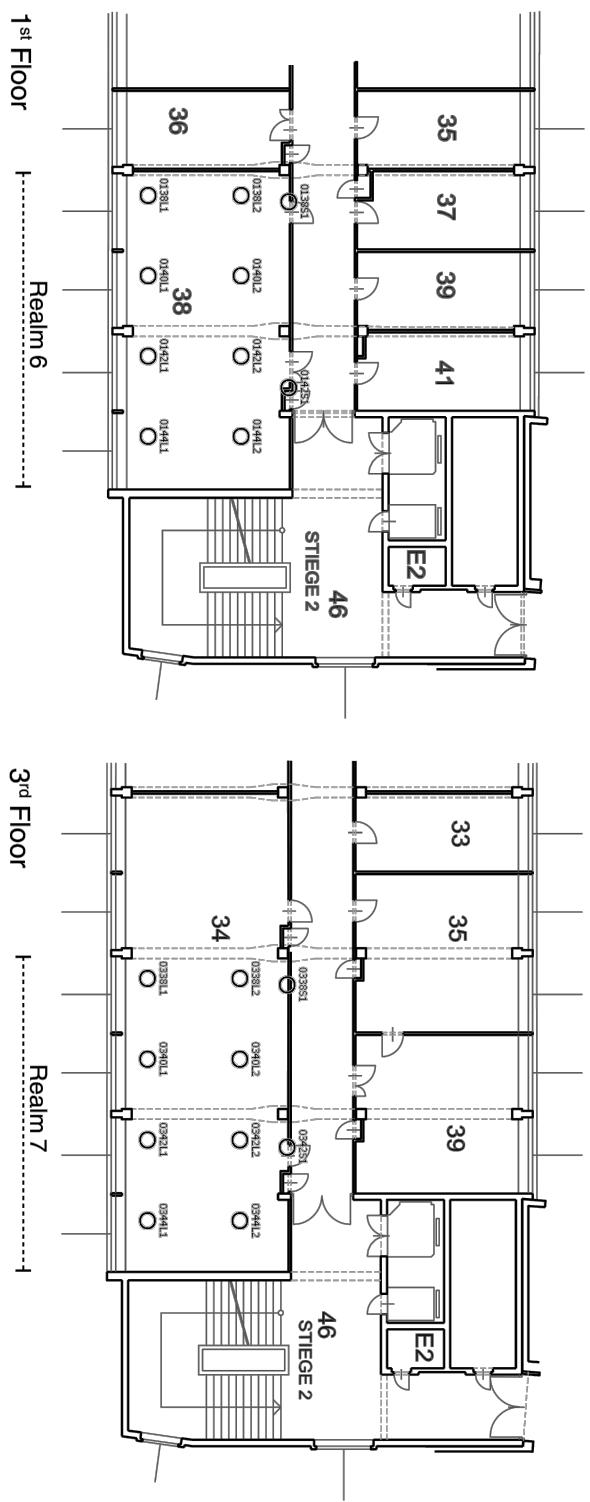
- The metal suspended ceiling prevents radio waves to propagate along the direct path between floors. Instead, the signal is diffracted around the edges of the metal suspended ceiling, inside the dry walls and window lining. Simulation quality therefore depends on the empirical diffraction model to be well calibrated.
- The ceiling height is 68 cm lower in the corridor than in the offices. Luminaire nodes mounted below the suspended ceiling of the offices are therefore above the metal suspended ceiling of the corridor. The space above the suspended ceiling of the corridor is

*Figure 33: Floor plan and deployment plan of the test network:
Node positions, location descriptors
(high level addresses), and realms.*

Mounting heights:

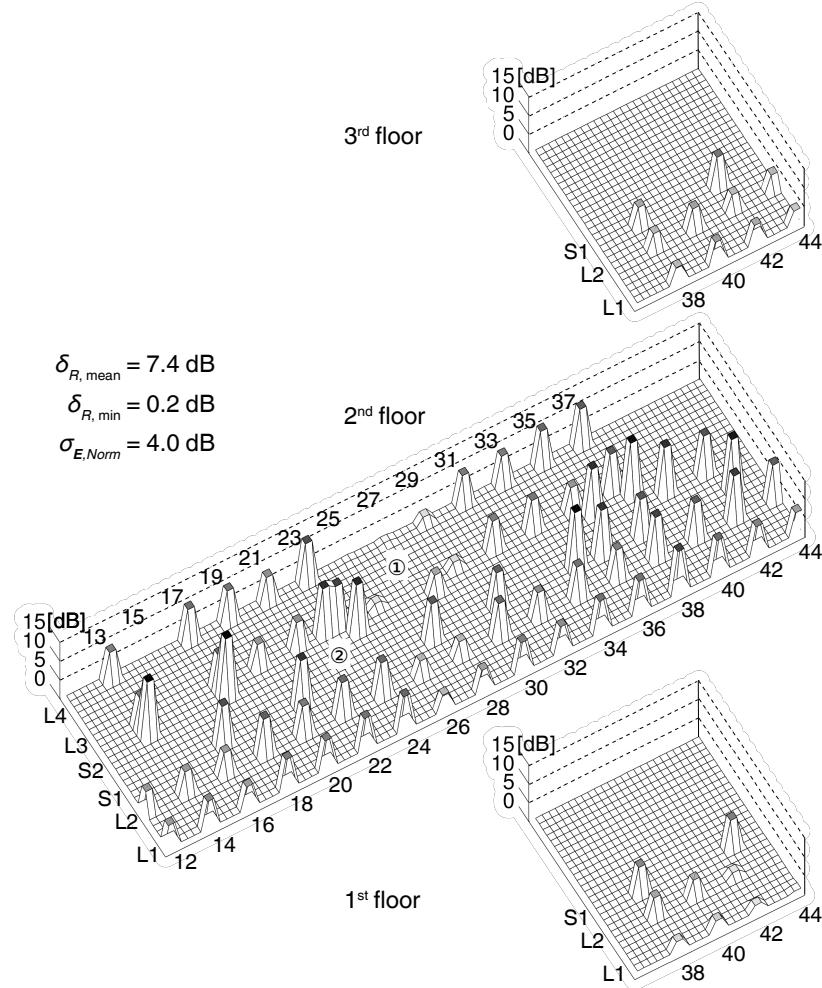
*L1, *L2, *L3, *L4: 306 cm

*LG: 238 cm *S*: 110 cm



therefore part of the dominant propagation path between nodes to both sides of the corridor. It is also occupied by metal air ducts and cable conduits, which are therefore included in the simulation.

- Offices are equipped with white boards and laboratories with tool cabinets, white goods and more sources of clutter.



*Figure 34: Analysis of the separability of the large dataset (standard deviation of edge weights again normalized to $\sigma_{E,\text{Norm}} = 4 \text{ dB}$). Notable: The light-switch nodes (*S1, *S2) are generally easier to separate than the luminaire nodes (*L1, *L2, *L3, *L4). The reason is that there are fewer light switch nodes, located towards the centre, surrounded by other nodes that provide separability. ① – The three rooms CA0225, -27 and -29 exhibit especially high symmetry of the measured edge weights and hence negligible separability. This is most likely due to chance. ② – Constellation of three nodes close together, but separable by node type. This results in high separability.*

Figure 34 contains the analysis of the node separability. To be comparable with the data in section 5.2.1, the edge distribution of the measured graph G_M is again normalized to $\sigma_{E,Norm} = 4$ dB during generation of the figure.

The separability of the light-switch nodes is comparable with the diamond-shaped topology (figure 30, bottom). What the two have in common is that many nodes are located at the centre of the topology and few at the boundary. For every two nodes that need separating, there is likely a nearby node that is in a good position to provide the separability.

The separability of the luminaire nodes, on the other hand, is comparable with the rectangular topology. The common property is that many nodes are located near the limits of the topology. The light-switch nodes cannot contribute much separability because they are vertically displaced ($\Delta h = 196$ cm) and therefore removed from the optimal axis along which they could provide separability (cf. figure 10). The luminaire nodes on the other side of the corridor are too far away, to the effect that the luminaire nodes approximate the unfavourable double row network of section 4.2.3. This is especially true for the somewhat isolated sub-networks on the 1st and 3rd floor.

Separability between the floors is good, because of the attenuation of the reinforced concrete ceiling and metal suspended ceiling. The nodes in one floor, however, fail to lend significant additional separability to the nodes below and above, because they are, again, removed from the axis along which separability would be needed. In addition to that, diffraction around the metal suspended ceiling is difficult to simulate and it is probable that some links between the floors do more harm than good.

After simulating the network with the ray-tracer, the standard deviation between measurement and simulation before preprocessing amounts to $\sigma_N = 6.74$ dB. This number includes all factors from measurement and simulation discussed so far: fading, antenna models, material models, component variances, ray-tracing variance.

The output of the matching procedure is analysed in figure 35. The most important aspect is that the algorithm succeeds in recognizing the global structure of the problem. This means that the solution to the graph matching problem is near global optimum. Secondly, there are no mismatches between floors, due to the good separability provided by the reinforced concrete ceiling.

There are no mismatches between light-switch nodes, also due to their better separability. The mismatches generally happen in areas with below-average separability, although not exclusively so. This is to be expected, because the analysis of separability is based solely on measured data with the simulation part replaced by a stochastic model.

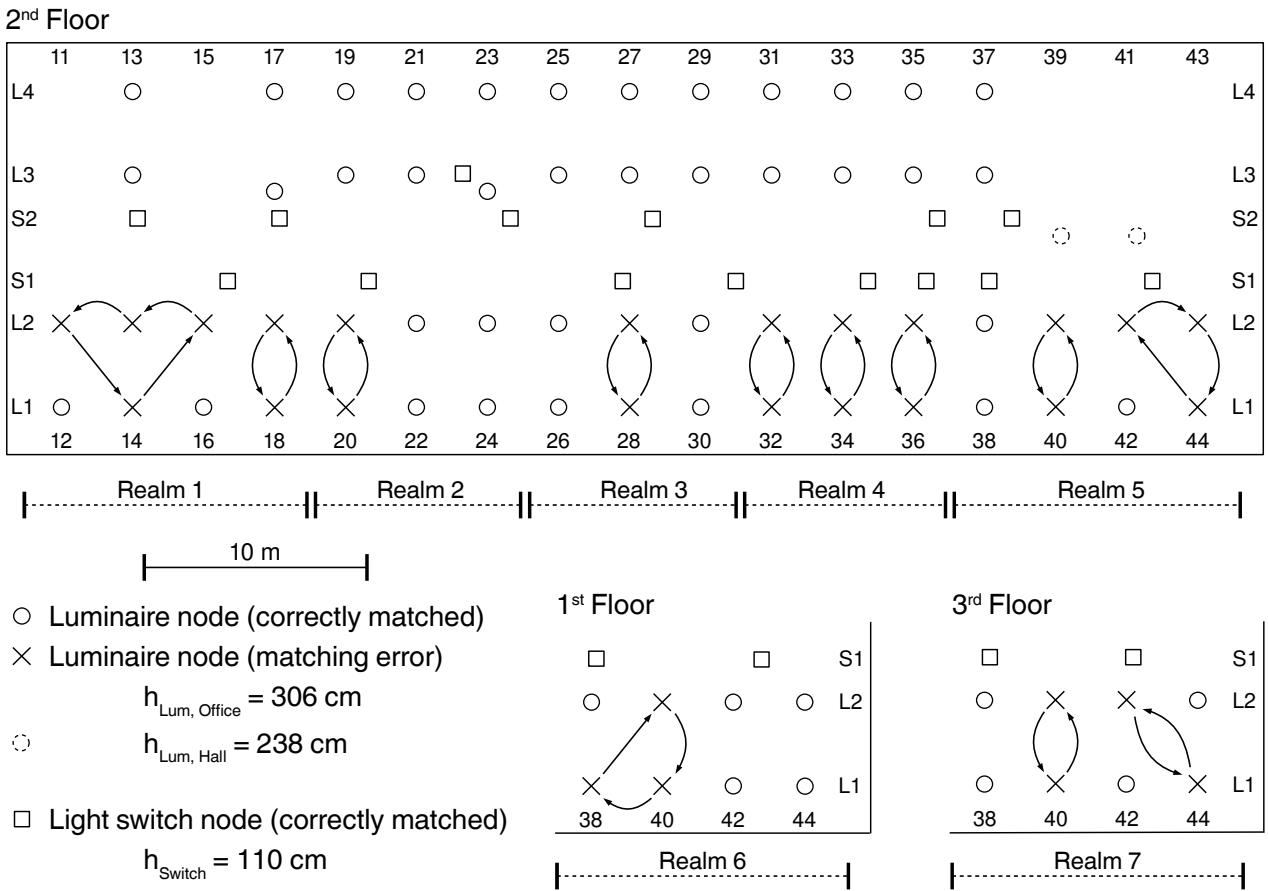


Figure 35: Analysis of the output of the matching procedure. 67 nodes are matched correctly, 28 nodes are mismatched. All mismatches concern neighbouring nodes which means the algorithm successfully recognizes the global structure. The problem however exhibits low separability of neighbouring nodes. There are no matching errors between the floors, because the ceiling ensures adequate separability. There are no mismatched light switch nodes, because they are completely surrounded which increases separability. The fact that all mismatches occur in the lower part (L1, L2) and none in the upper part (L3, L4) may be correlated with antenna anisotropy. It may also be the result of chance.

5.3 Synthetic Data Sets

A synthetic experiment is conducted with fully random graphs of 500 nodes, without node type attributes, without anchors. The graph G_M is constructed by choosing 500 uniformly distributed random points in a $50 \times 50 \times 100$ m rectangular area and calculating RSSI using the inverse square law with an empirical path loss potency of -3.5 and dipole antennas. G_S is derived by applying appropriate vertex and edge noise. Each data point is the average of 10 runs. Results are depicted in figure 36.

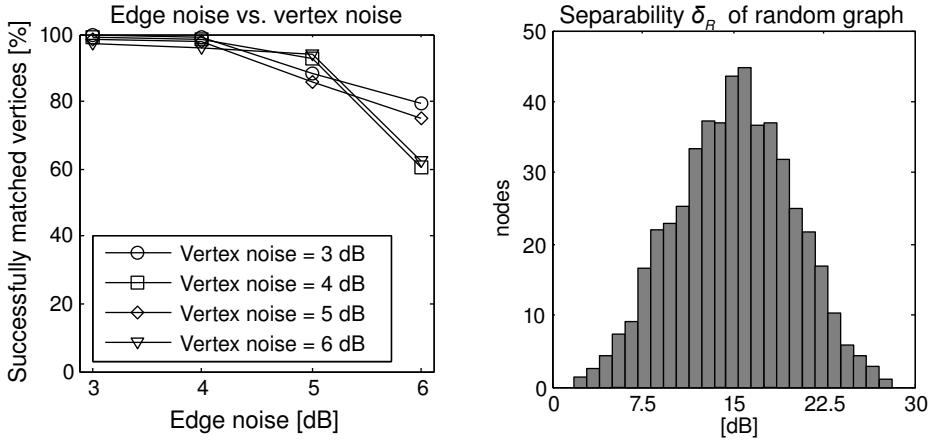


Figure 36: Left – Matching errors due to edge noise and vertex noise on random graphs of 500 vertices. Right – analysis of the separability of 500 vertex random graphs (σ_E normalized to 4 dB). The left tail, nodes that by chance have poor separability, are responsible for the matching errors.

The amount of matching errors at the different noise levels is consistent with the analysis of separability. The matching errors occur because the random topologies include a small number of poorly separable node-pairs which inevitably thwart the least squared error optimization approach. The sharp bend downward at high vertex noise and high edge noise levels cannot fully be explained by the separability histogram. The error rate of 40 % is indicative of the pre-processing and weighted graph matching being responsible for some of the mismatches.

Further observations about the synthetic experiment:

- The node pairs with poor separability are exclusively located near the boundaries of the virtual space. Separability is highest near the centre of the space. This is consistent with expectations, c.f. section 4.2.
- Run time of the matching algorithm for 500 node graphs: 240 s.

5.4 Summary of Findings

Table 5 lists the most important findings and influencing factors of sections 4 and 5. They are presented as standard deviation σ_x between measurement and simulation. The underlying assumption is that measurement and simulation are either free of biases, or biases can be eliminated by means of the global offset. The remaining deviation can be modelled statistically. If the statistics were Gaussian, which in some instances they are, then σ_x parametrizes

the distribution. The distributions are statistically independent, hence the individual σ_x can be accumulated by $\sigma_{\Sigma} = \sqrt{\sigma_{X1}^2 + \sigma_{X2}^2 + \dots}$.

Table 5: Summary of findings

Section	Impacting Factor		min	avg	max
4.3.1	Fading, residual after fading compensation	$\sigma_{N,F}$ [dB]			3
<i>Educated guess based on observation and results of Janssen [JAN96].</i>					
4.3.2	Ray-tracing, the method (not the models)	$\sigma_{N,RT}$ [dB]	1		4.2
<i>Estimation based on comparing forward and backward direction of the channel attenuation between simulated nodes. This figure only covers the variance of ray-tracing as Monte Carlo method, not shortcomings of the models. The actual value depends on the settings of the ray-tracer. The low value is observed with the “SRT”-setting, a slow but accurate method. The high value is observed with the “IRT”-setting where actual node locations are replaced with nearby grid positions.</i>					
4.3.3	Antenna radiation pattern	$\sigma_{N,RP}$ [dB]	0.7	2.4	5.0
<i>Low and high value are measured (section 5.2.2), avg value is estimated. The high value only takes effect if a grid-like topology fails to exert an averaging effect on the radiation pattern. The low value applies for antennas which are symmetric in directions favoured by the topology, e.g. dipoles.</i>					
5.2.2	Transceiver tolerances	$\sigma_{N,TP}$ [dB]		1.8	
<i>Measurement, based on comparing forward and backward direction of the channel attenuation between demonstration nodes.</i>					
4.2	Margin imposed by topology (5 % error rate)	σ_N [dB]	2	5	
<i>Estimation based on formal analysis. The low value applies to unfavourable topologies, the avg value to other grid-like topologies. Favourable topologies that take node type attributes into account can tolerate even more noise.</i>					

6 Conclusion

The novel method of automatic address assignment outlined in section 1.1 is based on sound theoretical foundations. Its successful applicability in wireless building automation depends on the scenario. Optimally, automatic address assignment should work

- without restricting network topologies,
- without precluding the choice of materials,
- without requiring higher than usual precision with respect to planning and execution.

These conditions would allow automatic address assignment to naturally fit into established industry procedures. Unfortunately, both theory (section 4) and practise (section 5) are in agreement that under those conditions the method fails to deliver address assignments at the required level of quality (error rate no worse than 5 %) – at least in the 2.4 GHz ISM band.

However, automatic address assignment does deliver high quality assignments in a number of scenarios that may make its application practical for niche applications:

- If the floor plan and network deployment plan were accurate to about $\lambda/10$ (1.2 cm), antennas and mounting fixtures perfectly specified and highest quality material models available, then the ray-tracer could produce coherent channel simulations accurate enough that even unfavourable topologies could be matched.
- Unfavourable topologies can be avoided by systematic use of node type attributes. If no two adjacent nodes share the same node type, then high quality address assignments can be delivered. This can be realized for example by artificially diversifying the node types and instructing deployment plans and installers accordingly.
- Unfavourable topologies can be mitigated if there is enough environmental separability. If interior walls are sufficiently thick and no two nodes of the same node type share a single room, then high quality address assignments can be achieved. These conditions are met in some HVAC applications, emergency lighting applications and danger alarm systems.

Conclusion

No definitive answer can be given for sub-GHz radio. While the topology issue stays quite the same, channel estimation, antennas and fading avoidance are different. The prevailing opinion is that indoor channel estimation is easier in sub-GHz radio because attenuation and shadowing are reduced. As a result, multi-path propagation and fading are also more nuanced, because comparatively more energy is transmitted along the direct path. On the other hand, antennas in sub-GHz are electrically short, typically badly defined whip antennas and therefore at least as difficult to accurately simulate as 2.4 GHz antennas. So without building a prototype network and trying it out, the question whether automatic address assignment works better or worse for sub-GHz must remain open.

To summarize the practical experiments: The results from the Mannheim data are contradictory. The data set has a particularly favourable topology – a fact that can be attributed to the original purpose of the dataset (localization by means of fingerprinting). It is also more challenging than the average site, because of the following aspects:

- The site is difficult to simulate because of metal-plated walls
- The situation during which the measurement was performed by [KIN08] was more challenging and less controlled than what can be expected during typical application of the address assignment procedure (furniture, unknown state of metal plated doors)
- Additional vertex attributes (especially device-type) are not present in the data set but would be in typical applications

The automatic address assignment procedure is able to solve the problem perfectly and reliably. This suggests that for favourable topologies there is margin for less than perfect simulation performance.

The VUT datasets, on the other hand, cannot be matched perfectly and this despite the lower noise variance. This again shows the central role of the network topology and the resulting presence or absence of separability of the nodes. Additional findings are:

- The matching errors happen as predicted by theory in areas of low separability.
- Walls help separability. Although there is one mismatch involving nodes left and right of a dry wall (attenuation < 1 dB), there are no mismatches between floors.
- The antenna, and per extension the node assembly, fixture and near-field, have a significant impact on the ability to accurately simulate the installed network. Easy antennas alone, however, are not enough to solve unfavourable topologies. Only extremely good antenna models and coherent simulation have the potential of doing so. This has not been tried experimentally though.

Of all analysed influencing factors and physical phenomena, the topology issue is the most problematic, because it defines the upper bound of acceptable noise variance. Fading, an-

tenna radiance patterns and simulation quality of the ray-tracer are all contributing to overall variance, but none stands out as singular reason. All are non-trivial to fix and together are responsible for making certain topologies impossible to solve reliably. Performance of the matching algorithm, on the other hand, proves to be unproblematic.

Potential game changers that would allow automatic address assignment to succeed even under unfavourable conditions:

- Any advance in deterministic channel estimation that would lower the standard deviation of the noise below 2 dB
- Inclusion of a high precision time-of-flight engine into low power radios – this may be the case for future chipsets with UWB capability (cf. section 6.1)

Indeed, automatic address assignment offers new possibilities besides cost advantages: Easy address assignment is an enabler for the Internet of Things, because it is an enabler for ubiquitous positioning: The address assignment procedure provides as its by-product a dense network of anchors (wireless nodes with known id and known location) and a detailed channel map. This is precisely what indoor positioning systems as surveyed in section 3.3 need in order to provide reliable and high-quality service.

For industrial application of automatic address assignment, the different pieces of the solution have to be more closely integrated: Building automation planning tool, 3D modelling package, ray-tracer and last but not least the matching procedure. This has to be done in a way that integrates with established industry procedures, and with respect to roles and responsibilities: The architect or planner has to deliver 3D models of the structure, the planner of the automation network the deployment plan. Both have to be imported and refined in the ray-tracer. Using the network management station, the commissioning technician has to perform the channel measurement and do the graph matching. The resulting address assignment needs to be committed to the network by the network management station. Furthermore, the commissioning technician may decide on site to perform the procedure on the whole network, or on smaller parts. The technician may have last minute changes, or may define a small number of anchors, e.g. in parts of the building where the ray-tracing model is lacking for some reason. The proposed procedure supports all these scenarios, but the tools need to be created.

6.1 Future Research

Topics of further research include the study of applicability of the procedure to outdoor environments, for example in environmental monitoring or street lighting scenarios. More topics

Conclusion

include the application of the matching procedure to different or novel ranging techniques, especially ultra wide band (UWB) and time of flight (TOF) ranging.

The immediate advantage of UWB is that more spectrum is available for fading suppression and maybe also for compensation of the antenna radiation pattern by averaging over the available spectrum. Since the attenuation of walls helps separability, going to higher frequencies has the advantage of radio waves being more attenuated by typical building materials. Also, at higher frequencies electro-magnetic radiation behaves more like rays and less like waves which helps with ray-tracing. On the other hand, prediction of UWB channels is considered more challenging than standard channels.

TOF ranging is promising for one reason in particular: No inverse square law. As a consequence, the separability that a witness node provides to a pair of nodes under test does not decrease with the distance. A node 20 m away can be just as valuable a witness as the direct neighbour. This completely changes the dynamic of the topology issue, because many more nodes of the topology may contribute to separability.

There are downsides to TOF as well: The accuracy of TOF ranging depends critically on the chip rate of the modulation scheme. The chip rate of IEEE 802.15.4 in the 2.4 GHz band is 2 Mchips/s. The chip duration therefore amounts to 500 ns in which the wave-front travels roughly 150 m. That is also the base accuracy of TOF ranging in this technology. Using statistics over a large number of measurements, this base accuracy can be improved upon by roughly an order of magnitude, down to 15 m, which is what typical TOF engines claim to be capable of. This is most likely not enough for indoor address assignment, but it may be good enough for outdoor applications (e.g. street lighting, environmental monitoring, etc.).

Combining UWB and TOF is even more promising: Direct sequence UWB works with chip rates of 200 Mchips/s and more. This would give an order of magnitude improvement over the calculation above, or a base accuracy of 1.5 m without statistical tricks. TOF with UWB would provide sufficient accuracy for address assignment combined with an eased topology problematic. This combination is therefore very promising for general purpose indoor automatic address assignment.

There are still, however, challenges to overcome:

- With TOF ranging there remains the uncertainty whether a distance measurement is the result of measuring the direct path, or some reflection. This is especially true for UWB where the received signal is typically the result of a frequency selective multiple of direct and reflected paths. This undoubtedly negatively impacts the accuracy of distance measurements. Maybe ray-tracing can help in resolving the ambiguity.

- Gathering the measured connectivity graph is more time-consuming if TOF ranging is involved: Rather than simply measuring the received signal strength of a broadcast, TOF ranging requires every pair of two nodes to actively cooperate and therefore need exclusive access to the channel during the measurement. This requires more time and more coordination.

To conclude, TOF ranging is very promising, but presents additional challenges to overcome.

The final presented idea is to combine different channel measurement technologies in order to improve upon the results of this work. This includes combining several radio technologies, such as measuring attenuation at different frequencies, or combine different physical effects, such as measuring attenuation and time-of-flight. It could also mean combining radio frequency measurements with ultra-sound. The former works well between rooms, the latter may give better results within the same room. The disadvantage of the approach is that the dimensionality of the state-space is now greatly increased by the additional measurements. Combining different channels may therefore either increase reliability or overwhelm the matching algorithm. Never the less, it would be interesting to try.

Bibliography

- [ALI07]: A. Alighanbari, C. D. Sarris; Rigorous and Efficient Time-Domain Modeling of Electromagnetic Wave Propagation and Fading Statistics in Indoor Wireless Channels ; published in IEEE Transactions on Antennas and Propagation, IEEE Antennas and Propagation Society, 2007
- [ALM93]: H. A. Almohamad; A linear programming approach for the weighted graph matching problem; published in IEEE Transactions on Pattern Analysis and Machine Learning, IEEE Computer Society, 1993
- [BAH00]: P. Bahl, V.N. Padmanbhan; RADAR: An In-Building RF-based User Location and Tracking System; published in INFOCOM 2000 - Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies, 2000
- [BER90]: H. Berger, H. Politze, P. Ungemach; Method for determining the configuration of detectors of a danger alarm system and for determining the system configuration of suitable detectors; EPO EP0485878 B1, 1990
- [BOL08]: D. P. Bohlmann, J. R. Hall, B. D. Westphal, G. A. Zuercher; BACnet Protocol MS/TP Automatic MAC Addressing; USPTO US7987247 B2, 2008
- [BOU93]: D. P. Bouche, F. A. Molinet, R. Mittra; Asymptotic and Hybrid Techniques for Electromagnetic Scattering; published in Proceedings of the IEEE, IEEE, 1993
- [BUH82]: R. Buhler, J.. Muggli, A. Scheidweiler, E. Schibli; Method of transmitting measuring values in a monitoring system; EPO EP0093872 A1, 1982
- [CHAN06]: L. Chan-Ping, J. L. Volakis, K. Sertel, R. W. Kindt, A. Anastopoulos; Indoor Propagation Models Based on Rigorous Methods for Site-Specific Multipath Environments; published in IEEE Transactions on Antennas and Propagation, IEEE Antennas and Propagation Society, 2006
- [CLIU07]: C. Liu, T. Scott, K. Wu, D. Hoffman; Range-free sensor localisation with ring overlapping based on comparison of received signal strength indicator; published in International Journal of Sensor Networks, Inderscience Publishers, 2007
- [CON04]: D. Conte, P. Foggia, C. Sansone, M. Vento; Thirty Years of Graph Matching in Pattern Recognition; published in International Journal of Pattern Recognition and Artificial Intelligence, World Scientific, 2004
- [EN 50090-5-3]: Home and Building Electronic Systems - Part 5-3: Media and Media Dependent Layers - Radio Frequency; CENELEC, 2006

Bibliography

- [EN 62386-102]: Digital Addressable Lighting Interface - Part 102: General Requirements; IEC, 2009
- [FER12]: L. Fери, G. M. P. J. Linnartz, J. W. Rietman, W. C. T. Schenk, C. J. Talstra, H. Yang; Efficient address assignment in coded lighting systems; EPO EP 2417834 A1, 2012
- [GRA06]: W. Granzer, G. Neugschwandtner, F. Praus, W. Kastner; Wireless Communication in KNX/EIB; published in , 2006
- [GRA11]: K. Gravogl, J. Haase, C. Grimm; Choosing the best wireless protocol for typical applications; published in Architecture of Computing Systems - ARCS 201124th International Conference, Lake Como, Italy, February 24-25, 2011. Proceedings, 2022
- [GRI11]: P. Grindberg; Approach for Discovering Devices on a Common Bus Without Direct Communication; USPTO US8443110 B2, 2011
- [HAL96]: F. Halsall; Data Communications, Computer Networks, and Open Systems; ISBN ; Addison-Wesley; 1996
- [HIGH01]: J. Hightower, G. Borriello; Location Systems for Ubiquitous Computing; published in Computer, IEEE Computer Society, 2001
- [HLIU07]: H. Liu, H. Darabi, P. Banerjee, L. Jing; Survey of Wireless Indoor Positioning Techniques and Systems ; published in IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Systems, Man and Cybernetics Society, 2007
- [HOL11]: E. Holleis, S. Zudrell-Koch; Analysis and address allocation of wireless building networks; WIPO WO2013023955 A2, 2011
- [HOL11a]: E. Holleis; Managing device ownership and commissioning in public-key encrypted wireless networks; WIPO WO2013023968 A1, 2011
- [HOL13]: E. Holleis, C. Grimm; Address Assignment in Indoor Wireless Networks Using Deterministic Channel Simulation; published in ISRN Sensor Networks, Hindawi Publishing Corporation, 2013
- [HOL13a]: E. Holleis; Security Considerations ; published in C. Grimm, P. Neumann, S. Mahlknecht; Smart Embedded Appliances Networks; ISBN 978-1-4419-8794-5, Springer NY, 2013
- [HUB03]: A. Huber; Lighting system and method for its production; USPTO US20040232856 A1, 2003
- [JAN96]: G. J. M. Janssen; Wideband indoor channel measurements and BER analysis of frequency selective multipath channels at 2.4, 4.75, and 11.5 GHz; published in IEEE Transactions on Communications, IEEE Communications Society, 1996
- [JOS07]: F. Jost, D. Kottmann, O. Leuthold; Method for integrating network nodes; EPO EP1956455 B1, 2007

- [KAS05]: W. Kastner, G. Neugschwandtner, S. Soucek, H.M. Newmann; Communication Systems for Building Automation and Control; published in Proceedings of the IEEE, IEEE, 2005
- [KIN06]: T. King, S. Kopf, T. Haenselmann, C. Lubberger, W. Effelsberg; COMPASS: A probabilistic indoor positioning system based on 802.11 and digital compasses; published in WINTECH 06 - Proceedings of the 1st international workshop on wireless network testbeds, experimental evaluation & characterization, 2006
- [KIN08]: T. King, T. Haenselmann, W. Effelsberg; On-Demand Fingerprint Selection for 802.11-based Positioning Systems; published in WOWMOM '08 Proceedings of the 2008 International Symposium on a World of Wireless, Mobile and Multimedia Networks, 2008
- [KNA08]: S. Knauth, R. Kistler, D. Kaslin, A. Klapproth; SARBAU Towards Highly Self-Configuring IP-Fieldbus Based Building Automation Networks; published in AINA '08 Proceedings of the 22nd International Conference on Advanced Information Networking and Applications, 2008
- [KRI13]: C. Krieg; A Lightweight Kerberos Implementation for Resource-Constrained Systems in Wireless Sensor Networks (diploma thesis), 2013
- [LAN99]: F. M. Landstorfer; Wave Propagation Models for the Planning of Mobile Communication Networks; published in 29th European Microwave Conference, 1999
- [LEE00]: E. Lee, P. S. Hall, P. Gardner; Dual band folded monopole/loop antenna for terrestrial communication system; published in Electronics Letters, IET The Institution of Engineering and Technology, 2000
- [LON06]: Echelon Corporation; LonMaker User's Guide, Turbo Edition; Echelon Corporation, 2006
- [LUO01]: Bin Luo, Edwin R. Hancock; Structural Graph Matching Using the EM Algorithm and Singular Value Decomposition; published in IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Computer Society, 2001
- [MEY13]: Willi Meyer; KNX/EIB Engineering Tool Software ETS4; ISBN 3810103551; Hüthig & Pflaum Verlag, Heidelberg; 2013
- [NIU00]: F. Niu; Folded L-antennas; published in Antennas and Propagation Society International Symposium 2000, 2000
- [NOC10]: D. Nocker; Notbeleuchtungsanlage und ein Verfahren zur Initialisierung einer Notbeleuchtungsanlage; EPO EP2453719 A1, 2010
- [PAP09]: A. Papapostolou, H. Chaouchi; WIFE: Wireless Indoor Positioning Based on Fingerprint Evaluation; published in NETWORKING 2009 - Proceedings of the 8th International IFIP-TC 6 Networking Conference, Aachen, Germany, 2009
- [PAT01]: N. Patwari; Relative location in wireless networks; published in Proceedings of the 53rd Vehicular Technology Conference 2001, 2001

Bibliography

- [PHA10]: Matt Pharr, Greg Humphreys; Physically Based Rendering - From Theory to Implementation; ISBN 9780129750792; Elsevier; 2010
- [RAH12]: Y. Rahmat-Samii; GTD, UTD, UAT and STD: A Historical Revisit; published in IEEE-APS Topical Conference on Antennas and Propagation in Wireless Communication, 2012
- [RAN96]: S. Gold, A. Rangarajan; A graduated assignment algorithm for graph matching; published in IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Computer Society, 1996
- [REI07]: C. Reinisch, W. Kastner, G. Neugschwandtner, W. Granzer; Wireless Technologies in Home and Building Automation; published in Proceedings 5th IEEE International Conference on Industrial Informatics, 2007
- [ROS10]: F. Rossi, N. Villa-Vialaneix; Optimizing an Organized Modularity Measure for Topographic Graph Clustering: A Deterministic Annealing Approach; published in Neuocomputing, Elsevier, 2010
- [SAU11]: T. Sauter, S. Soucek, W. Kastner, D. Dietrich; The Evolution of Factory and Building Automation; published in IEEE Industrial Electronics Magazine, IEEE Industrial Electronics Society, 2011
- [SAV04]: A. Savvides, M. Srivastava, L. Girod, D. Estrin; Localization in Sensor Networks; published in ; Wireless Sensor Networks; ISBN , Springer US, 2004
- [SHA06]: R. Shane; System and Method for Light Source Identification; USPTO US20090026978 A1, 2006
- [SIN07]: K. Sinha; A Beacon Selection Algorithm for Bounded Error Location Estimation in Ad Hoc Networks ; published in ICCTA 07 - Proceedings of the International Conference on Computing: Theory and Applications, 2007
- [SIN67]: R. Sinkhorn, P. Knopp; Concerning Nonnegative Matrices and Doubly Stochastic Matrices; published in Pacific Journal of Mathematics, Mathematical Science Publishers, 1967
- [STO07]: R. Stoleru, J. A. Stankovic, S. H. Son; Robust node localization for wireless sensor networks; published in EmNets '07 Proceedings of the 4th workshop on Embedded networked sensors, 2007
- [UKA05]: U. Kästli, B. Lontka; Method for Determining the Configuration of a Danger Warning System, and Danger Warning System; EPO EP1703481 B1,
- [UME88]: S. Umeyama; An eigendecomposition approach to weighted graph matching problems; published in IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Computer Society, 1988
- [VAL97]: R. A. Valenzuela; Estimating local mean signal strength of indoor multipath propagation ; published in IEEE Transactions on Vehicular Technology, IEEE Vehicular Technology Society, 1997

- [VAL98]: R. A. Valenzuela; Indoor propagation prediction accuracy and speed versus number of reflections in image-based 3-D ray-tracing; published in VTC 98, 48th IEEE Vehicular Technology Conference, 1998
- [WAL08]: G. Wallhead, N. Aycliffe; Determination of the Bus Address of a Subscriber in an Illuminating Bus System; EPO EP2282612A2,
- [WAN12]: C. Wan, A. Mita, S. Xue; Non-Line-of-Sight Beacon Identification for Sensor Localization; published in International Journal of Distributed Sensor Networks, Hindawi, 2012
- [WOL00]: G. Wölfle, R. Hoppe, T. Binzer, F. M. Landstorfer; Radio Network Planning and Propagation Models for Urban and Indoor Wireless Communication Networks; published in Millennium Conference on Antennas & Propagation (AP2000), 2000
- [WOL99]: G. Wölfle, R. Hoppe, and F.M. Landstorfer; A Fast and Enhanced Ray Optical Propagation Model for Indoor and Urban Scenarios, Based on an Intelligent Preprocessing of the Database; published in 10th IEEE Internat. Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC) 1999, 1999
- [ZAN08]: G. Zanca, F. Zorzi, A. Zanella, M. Zorzi; Experimental comparison of RSSI-based localization algorithms for indoor wireless sensor networks; published in REALWSN 08 - Proceedings of the workshop on Real-world wireless sensor networks, 2008
- [ZAS09]: M. Zaslavskiy; A Path Following Algorithm for the Graph Matching Problem; published in IEEE Transactions on Pattern Analysis and Machine Learning, IEEE Computer Society, 2009
- [ZAV06]: M. M. Zavlanos; A Dynamical Systems Approach to Weighted Graph Matching; published in 45th IEEE Conference on Decision and Control, 2006
- [ZBA11]: ZigBee Building Automation Standard, 053516r12; ZigBee Alliance, 2011
- [ZHA10]: ZigBee Home Automation Public Application Profile, 053520r26; ZigBee Alliance, 2010
- [ZIG08]: ZigBee Specification 053474r17; ZigBee Alliance, 2008
- [ZLL12]: ZigBee Light Link Standard, 11-0037-10; ZigBee Alliance, 2012
- [ZOTA10]: ZigBee Over-the-Air Upgrading Cluster; ZigBee Alliance, 2010
- [ZSE08]: ZigBee Smart Energy Profile Specification, 075356r15; ZigBee Alliance, 2008

Appendix A: ZigBee Graph Assembly Cluster

Assembling the measured connectivity graph is an on-site activity. The wireless network is assumed non-commissioned, but the nodes are powered and an initial, maybe temporary, network has been formed by a procedure such as ZigBee network forming. Nodes are reachable via their low-level network or MAC-address. The wireless channel can be measured as a whole, or the problem can be subdivided rudimentarily, for example by only powering nodes on a particular floor or section of the building.

The procedure begins by putting the network into graph-assembly mode.

Command frame “**Enter Graph Assembly Mode Request**” (broadcast):
<countdown (seconds), channel list, stay time (seconds), n° of beacons >

countdown – To allow power-cycling devices to reliably receive the command, graph-assembly mode is only entered after a short countdown delay.

channel list – List of channels to be visited during graph-assembly. Should include the channels with lowest and highest frequency to ensure maximal spectral separation, plus one or two in the middle ranges.

stay time – Number of seconds the wireless nodes spend in each channel of *channel list*.

n° of beacons – Number of beacons to transmit in each channel of *channel list*.

In graph assembly mode, mechanisms such as mesh-routing and indirect delivery are suspended. All nodes, including low-power nodes that usually power-cycle their receivers, are required to keep their transmitters powered during graph assembly. After waiting countdown seconds, for each channel in *channel list*, all nodes go through the following procedure:

- Tune transceiver to next channel in *channel list*
- Choose *n° of beacons* random instances within the time window of *stay time* (millisecond resolution) to transmit graph assembly beacons; use clear channel assessment and random back-off to help avoiding beacon collisions

- Receive⁸ and store graph assembly beacons from nearby nodes, also recording the RSSI associated with the beacon

Command frame “**Graph Assembly Beacon**” (broadcast with radius = 1):
< low-level address, node type descriptor, TX-power (dBm) >

low-level address – Either a MAC address or other low-level addressing mechanism used by the underlying protocol.

node type descriptor – A tuple uniquely identifying the class, brand and type of node.

TX-power – The transmit power in dBm used to send the beacon. For each channel and node more than one beacon is generally received, but only the maximum RSSI value is retained in the table.

Table 6: Structure of the Graph Assembly Table

<i>low-level address</i>	<i>node type descriptor</i>	<i>TX-power</i>	<i>RSSI channel 1</i>	...	<i>RSSI channel n</i>
:	:	:	:	⋮	⋮

If the memory reserved for storing the table runs out and additional beacons are received, the table entry with the lowest maximum RSSI recorded is to be replaced by the new beacon. The size of each table entry depends on the underlying protocol and may vary between 8 bytes and 18 bytes (8 byte MAC addresses, 4 byte node type descriptor).

After visiting all channels, the nodes revert to normal behaviour (routing, power-cycling). They keep the saved beacons in storage until they are retrieved by graph retrieval commands.

Graph retrieval requests are broadcasted by the management station:

Command frame “**Graph Retrieval Request**” (broadcasted by management station):
< query jitter >

query jitter – Mechanism to randomly address some nodes, but not all [ZOTA10]. It is a number between 1 and 100. Receiving nodes draw a random number from the same range. If the random number is less than or equal to *query jitter*, they continue processing the message or discard it otherwise. *query jitter* should be chosen so that no more than ten nodes answer on average, so as not to overwhelm channel capacity near the management station.

⁸ End devices like ZigBee or IEEE 802.15.4 end devices may need to put their MAC-layer into promiscuous mode, meaning they need to accept packets from all nodes, not just their parent.

If the underlying wireless protocol supports such a mechanism, the node collecting the graph assembly tables, should first initiate a many-to-one route discovery [ZIG08]. While collecting the tables, the node temporarily acts as a data sink and runs the risk of locally overflowing routing tables. Many-to-one routing is a mechanism to alleviate the problem.

Nodes passing the *query jitter* test answer with Graph Retrieval Response:

Command frame “**Graph Retrieval Response**” (reply to management station)
 $\langle \text{originator}, n^{\circ} \text{ entries in packet}, \text{index first entry}, \text{total } n^{\circ} \text{ entries}, \text{list of table entries} \rangle$

originator – MAC address of the node having collected the graph assembly table. This allows nodes to send table fragments on behalf of other nodes, e.g. routers in behalf of their end-devices.

n° entries in packet, index first entry, total n° entries – Mechanism to transmit tables larger than fit in one packet. The numbers are respectively the number of transmitted table entries in the current packet, the index number of the first entry in the packet with respect to the whole table, and the total number of entries in the whole table.

If possible, Graph Retrieval Response packets are to be sent using source routing along the previously recorded many-to-one route. If possible, an acknowledged transmission method is to be used. After receiving acknowledgements for all table fragments, the graph assembly table can be deleted and the node can ignore further graph retrieval requests.

Assembling the measured connectivity graph G_M from the received graph assembly tables is straight forward:

- Each vertex is represented by one table; vertices are attributed by low level address and node type descriptor of the node originally assembling the table
- Each edge is represented by a pair of rows in the tables of the respective vertices. Edge weights are given by (RSSI and TXpwr in dBm, $g_{M,AB}$ in dB):

$$g_{M,AB} = g_{M,BA} = \max(\max(\text{RSSI}_{A \rightarrow B}) - \text{TXpwr}_A, \max(\text{RSSI}_{B \rightarrow A}) - \text{TXpwr}_B)$$

Edge weights are estimates for the interference-corrected channel attenuation between nodes A and B. Antenna, matching and insertion losses at transmitter and receiver are not accounted for, but dealt with during pre-processing. A small number of missing graph assembly tables can be tolerated, because all relevant information is redundantly available in the assembly tables of neighbouring nodes.

Appendix B: ZigBee Network Used for Experiments

Figure 37 shows the minimal ZigBee nodes used during the experiment at Vienna University of Technology, figures 38 and 39 show the programming tool used to flash the software.

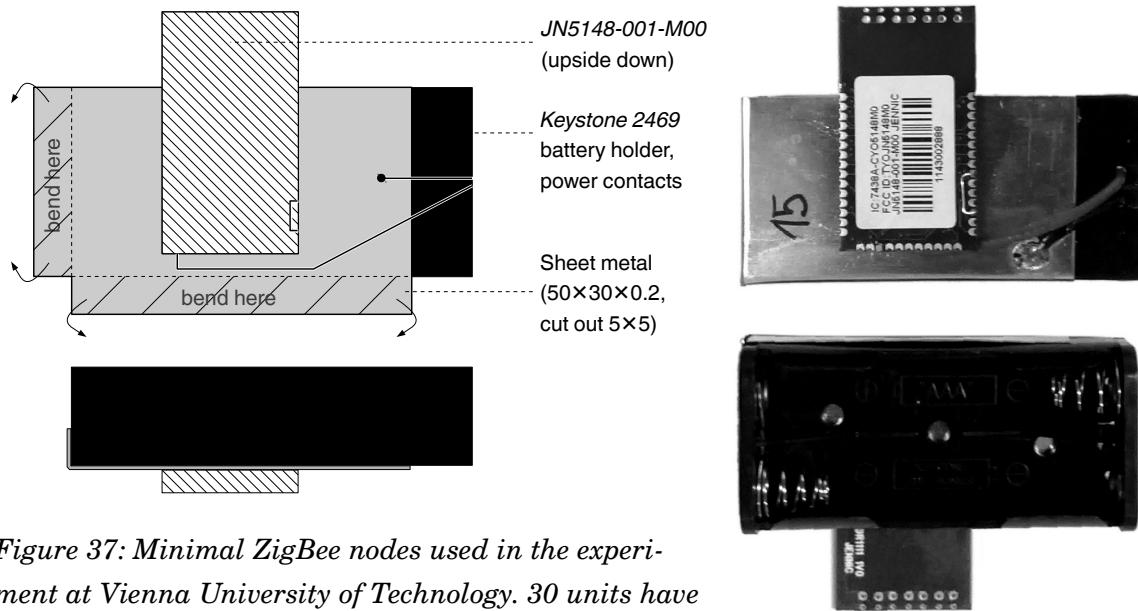
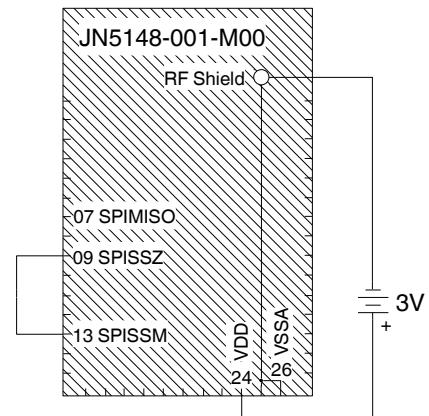


Figure 37: Minimal ZigBee nodes used in the experiment at Vienna University of Technology. 30 units have been produced and used in the experiment.

Top – construction drawing at full scale, all dimensions in mm. Top right – photos of the finished assembly. Bottom right – the schematic diagram of the design.

Construction: The RF-shield of the JN5148 module is soldered against the metal sheet using solder paste and a heat gun. The metal sheet itself is glued to the battery holder using double sided adhesive tape.

The node is powered by two “AAA” batteries and for lack of a power switch immediately boots into the software after inserting the batteries.



Appendix B: ZigBee Network Used for Experiments

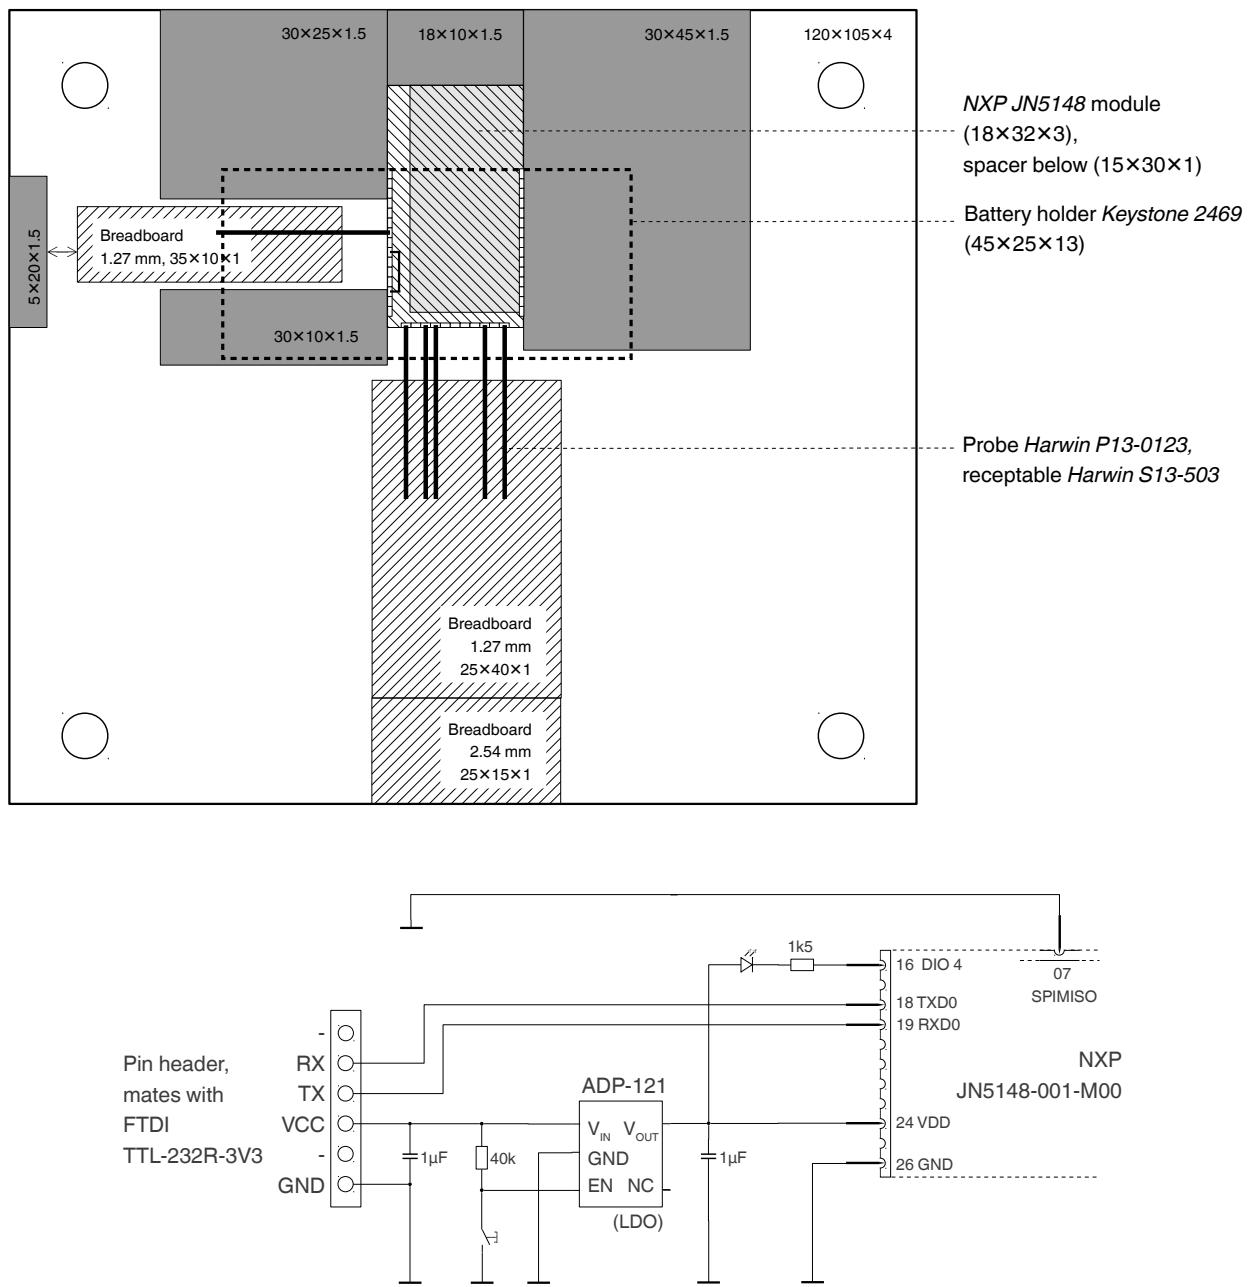


Figure 38: Top – construction drawing of the programming tool at full scale, all dimensions in mm. Bottom – schematic drawing of the programming tool.

The assembly is made of various pieces of acrylic glass that together with the spring-loaded needle probes forms a receptacle for the ZigBee node. Figure 39 shows the programming tool with and without inserted ZigBee node, and explains the programming procedure.

The programming tool provides the node with power, serial communication and a signalling LED. It is therefore important to remove the batteries from the node before inserting it into the programming tool.

The node software is derived from the SmartCoDe project. The majority of the code was written by the author for the purpose of the project and retrofitted for the address assignment experiment. It features a ZigBee coordinator and ZigBee router node, including joining, binding and life-cycle management, ZigBee Home Automation security, a logging framework, remote download of logfiles, over-the-air software upgrades, and a network management interface to the coordinator. In addition to that, it features an implementation of the ZigBee Graph Assembly Cluster specified in appendix A.

The graph assembly tables are downloaded to the ZigBee coordinator using the Graph Assembly Cluster and from there immediately forwarded to a PC over the management interface. An extension to the SmartCoDe management software *MTU Commander* is used to gather the individual tables, save them to files, and load them into the weighted graph matcher implemented in Matlab.

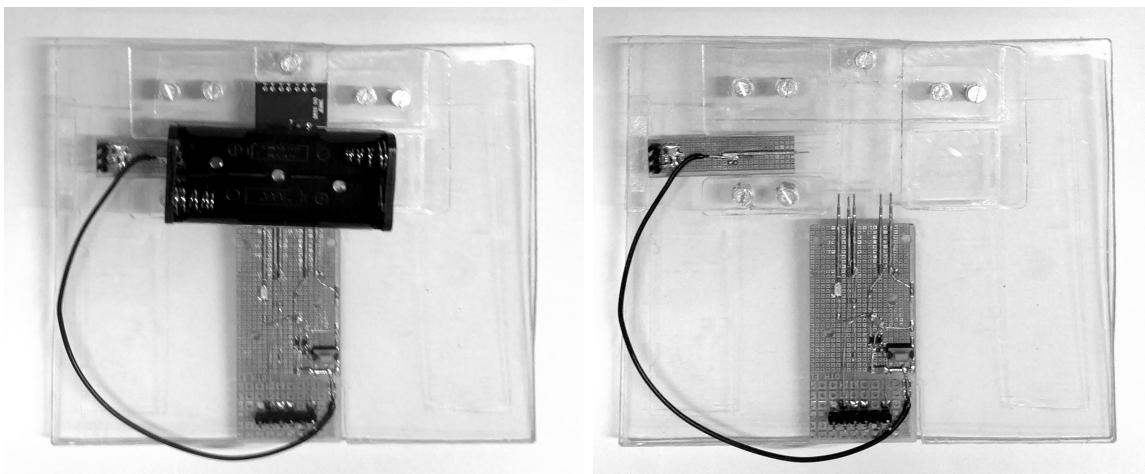


Figure 39: Programming tool with and without inserted node.

The programming procedure works as follows:

- *Connect the programming tool to the PC using the FTDI TTL-232R-3V3 cable (the USB to serial cable also used by the NXP JN5148 demonstration kit)*
- *Insert the node without batteries*
- *Keep the node in reset by holding down the push button*
- *Push left needle probe sideways to make contact with pin 7 of the ZigBee module (SPIMISO)*
- *Release the reset button*
- *Retract the left needle probe*
- *The node is now booted into the boot loader and ready to receive a program using the NXP Jennic flash programming software JN-SW-4007*

CURRICULUM VITAE



PERSONAL Edgar John Holleis, MSc.

Date of birth: 27th March 1979

Nationality: Austrian

CONTACT Belvederegasse 39/2

1040 Wien

+4369919133932

edgar@indoktrination.at

ACADEMIC **2009 – 2013** PHD study, Vienna University of Technology

CAREER **28.11.2008** Masters degree in Electrical Engineering (Telecommunication), Vienna University of Technology; Title of diploma thesis: "Inferring a Three-dimensional, Stable World View from Two-dimensional Sensor Data"

1997-2008 Electrical Engineering (Telecommunications), Vienna University of Technology; interrupted by
2003-2004 Civil Peace Service in Osijek, Croatia
2001-2002 Erasmus program, UPC-Barcelona TECH, Spain

PROFESSIONAL **Since 09/2013** Zühlke Engineering (Austria) GmBH: Embedded Software Engineer

CAREER **03/2010 – 12/2012** Tridonic GmBH & Co KG: EU-funded research project SmartCoDe (Smart-Grid & sensor networks)

08/2007 – 02/2010 Scientific assistant for Institute of Computer Technology, Vienna University of Technology. Research projects SENSE (AI & video surveillance), VOLARE (internet technology in air-traffic control), and ARS (AI & psycho-analysis).

Summer term 2006, 2007 University tutor (part-time) for Computer Languages Department: Lecture “Compiler Architecture”

PATENTS & PUBLICATIONS Edgar Holleis, Christoph Grimm; *Address Assignment in Indoor Wireless Networks Using Deterministic Channel Simulation*; ISRN Sensor Networks, Hindawi Publishing Corp., 2013

Edgar Holleis; *Managing Device Ownership and Commissioning in Public-Key Encrypted Wireless Networks*; European patent application WO2012EP65497 20120808 (pending)

Edgar Holleis, Stefan Zudrell-Koch; *Analysis and Address Allocation of Wireless Building Networks*; European patent application WO2012EP65417 20120807 (pending)

Edgar Holleis; *Security Considerations*; in Christoph Grimm, Peter Neumann, Stefan Mahlknecht (Eds.); Smart Embedded Appliances Networks (pp. 67-85), Springer NY. 2012

(+ co-author of two conference papers)